

Unsupervised Pivot Translation for Distant Languages

Yichong Leng*

University of Science and Technology of China
lyc123go@mail.ustc.edu.cn

Xu Tan

Microsoft Research
xuta@microsoft.com

Tao Qin[†]

Microsoft Research
taoqin@microsoft.com

Xiang-Yang Li[†]

University of Science and Technology of China
xiangyangli@ustc.edu.cn

Tie-Yan Liu

Microsoft Research
tyliu@microsoft.com

Abstract

Unsupervised neural machine translation (NMT) has attracted a lot of attention recently. While state-of-the-art methods for unsupervised translation usually perform well between similar languages (e.g., English-German translation), they perform poorly between distant languages, because unsupervised alignment does not work well for distant languages. In this work, we introduce unsupervised pivot translation for distant languages, which translates a language to a distant language through multiple hops, and the unsupervised translation on each hop is relatively easier than the original direct translation. We propose a learning to route (LTR) method to choose the translation path between the source and target languages. LTR is trained on language pairs whose best translation path is available and is applied to the unseen language pairs for path selection. Experiments on 20 languages and 294 distant language pairs demonstrate the advantages of the unsupervised pivot translation for distant languages, as well as the effectiveness of the proposed LTR for path selection. Specifically, in the best case, LTR achieves an improvement of 5.58 BLEU points over the conventional direct unsupervised method.

1 Introduction

Unsupervised neural machine translation (NMT) (Artetxe et al., 2017b; Lample et al., 2017, 2018), which uses only monolingual sentences for translation, is of great importance

for zero-resource language pairs. Unsupervised translation relies on unsupervised cross-lingual word alignment or sentence alignment (Conneau et al., 2017; Artetxe et al., 2017a), where word embedding mapping (Artetxe et al., 2017b; Lample et al., 2017) and vocabulary sharing (Lample et al., 2018) are used for word alignment, and encoder/decoder weight sharing (Artetxe et al., 2017b; Lample et al., 2018) and adversarial training (Lample et al., 2017) are used for sentence alignment.

Unsupervised cross-lingual alignment works reasonably well for a pair of similar languages, such as English-German or Portuguese-Galician, since they have similar lexica and syntax and share the same alphabets and language branch. However, the alignment between a pair of distant languages, which are not in the same language branch¹, such as Danish-Galician is challenging. As a consequence, unsupervised translation between distant languages is usually of lower quality. For example, the unsupervised NMT model achieves 23.43 BLEU score on Portuguese-Galician translation, while just 6.56 on Danish-Galician translation according to our experiments. In this work, we focus on unsupervised translation of distant languages.

We observe that two distant languages can be linked through multiple intermediate hops where unsupervised translation of two languages on each

* The work was done when the first author was an intern at Microsoft Research Asia.

[†] Corresponding author

¹In this work, we use language branch to determine distant languages. We choose the taxonomy of language family provided by Ethnologue (Paul et al., 2009) (<https://www.ethnologue.com/browse/families>), which is one of the most authoritative and commonly accepted taxonomies. Distant languages can also be defined using other principles, which we leave to future work.

hop is easier than direct translation of the two distance languages, considering that the two languages on each intermediate hop are more similar, or the size of monolingual training data is larger. Therefore, we propose unsupervised pivot translation through multiple hops for distant languages, where each hop consists of unsupervised translation of a relatively easier language pair. For example, the distant language pair Danish→Galician can be translated by three easier hops: Danish→English, English→Spanish and Spanish→Galician. In this way, unsupervised pivot translation results in better accuracy (12.14 BLEU score) than direct unsupervised translation (6.56 BLEU score) from Danish to Galician in our experiments.

The challenge of unsupervised pivot translation is how to choose a good translation path. Given a distant language pair $X \rightarrow Y$, there exists a large amount of paths that can translate from X to Y^2 , and different paths may yield very different translation accuracies. Therefore, unsupervised pivot translation may result in lower accuracy than direct unsupervised translation if a poor path is chosen. How to choose a path with good translation accuracy is important to guarantee the performance of unsupervised pivot translation.

A straightforward method is to calculate the translation accuracies of all possible paths on a validation set and choose the path with the best accuracy. However, it is computationally unaffordable due to the large amount of possible paths. For example, suppose we consider at most 3 hops ($N = 3$) and 100 languages ($M = 100$), and assume each path takes an average of 20 minutes to translate all the sentences in the validation set using one NVIDIA P100 GPU to get the BLEU score. Then it will take nearly 1400000 GPU days to evaluate all candidate paths. Even if we just consider 20 languages ($M = 20$), it will still take 2200 GPU days. Therefore, an efficient method for path selection is needed. We propose a learning to route (LTR) method that adopts a path accuracy predictor (a multi-layer LSTM) to select a good path for a distant language pair. Given a translation path and the translation accuracy of each hop on the path, the path accuracy predictor can predict the overall translation accuracy following this

²Suppose we only consider translation paths with at most N hops. Given M candidate intermediate languages, there are $\frac{M!}{(M-N+1)!}$ possible paths.

path. Such a predictor is first trained on a training set of paths with known overall accuracy, and then used to predict the accuracy of a path unseen before.

We conduct experiments on a large dataset with 20 languages and a total of 294 distant language pairs to verify the effectiveness of our method. Our proposed LTR achieves an improvement of more than 5 BLEU points on some language pairs.

The contributions of this paper are as follows: (1) We introduce pivot translation into unsupervised NMT to improve the accuracy of distant languages. (2) We propose the learning to route (LTR) method to automatically select the good translation path. (3) Large scale experiments on more than 20 languages and 294 distant language pairs demonstrate the effectiveness of our method.

2 Related Work

In this section, we review the related work from three aspects: unsupervised neural machine translation, pivot translation, and path routing.

Unsupervised NMT As the foundation of unsupervised sentence translation, unsupervised word alignment has been investigated by (Conneau et al., 2017; Artetxe et al., 2017a), where linear embedding mapping and adversarial training are used to ensure the distribution-level matching, achieving considerable good accuracy or even surpasses the supervised counterpart for similar languages. Artetxe et al. (2017b); Lample et al. (2017) propose unsupervised NMT that leverages word translation for the initialization of the bilingual word embeddings. Yang et al. (2018) partially share the parameter of the encoder and decoder to enhance the semantic alignment between source and target language. Lample et al. (2018) further share the vocabulary of source and target languages and jointly learned the word embeddings to improve the quality of word alignment, and achieve large improvements on similar language pairs. Recently, inspired by the success of BERT (Devlin et al., 2018), Lample and Conneau (2019) leverage the BERT pre-training in the unsupervised NMT model and achieve state-of-the-art performance on some popular language pairs.

Previous works on unsupervised NMT can indeed achieve good accuracy on similar language pairs, especially on the closely related languages such as English and German that are in the same language branch. In this circumstance, they can

simply share the vocabulary and learn joint BPE for source and target languages, and share the encoder and decoder, which is extremely helpful for word embedding and latent representation alignment. However, they usually achieve poor accuracy on distant languages that are not in the same language branch or do not share same alphabets. In this paper, we propose pivot translation for distant languages, and leverage the basic unsupervised NMT model in (Lample et al., 2018) on similar languages as the building blocks for the unsupervised pivot translation.

Pivot Translation Pivot translation has long been studied in statistical machine translation to improve the accuracy of low/zero-resource translation (Wu and Wang, 2007; Utiyama and Isahara, 2007). Cheng et al. (2017); Chen et al. (2017) also adapt the pivot based method into neural machine translation. However, conventional pivot translation usually leverages a resource-rich language (mainly English) as the pivot to help the low/zero-resource translation, while our method only relies on the unsupervised method in each hop of the translation path. Due to the large amount of possible path choices, the accuracy drops quickly along the multiple translation hops in the unsupervised setting, unsupervised pivot translation may result in lower accuracy if the path is not carefully chosen. In this situation, path selection (path routing) will be critical to guarantee the performance of pivot translation.

Path Routing Routing is the process of selecting a path for traffic in a network, or between or across multiple networks. Generally speaking, routing can be performed in many types of networks, including circuit switching network (Girard, 1990), computer networks (e.g., Internet) (Huitema, 2000), transportation networks (Raff, 1983) and social networks (Liben-Nowell et al., 2005). In this paper, we consider the routing of the translation among multiple languages, where the translation accuracy is the criterion for the path selection. Usually, the translation accuracy of the multi-hop path is not simply the linear combination of the accuracy on each one-hop path, which makes it difficult to route for a good path.

3 Unsupervised Pivot Translation

Observing that unsupervised translation is usually hard for distant languages, we split the direct translation into multiple hops, where the unsupervised translations on each hop is relatively easier than the original direct unsupervised translation. Formally, for the translation from language X to Y , we denote the pivot translation as

$$X \rightarrow Z_1 \rightarrow \dots \rightarrow Z_n \rightarrow Y, \quad (1)$$

where Z_1, \dots, Z_n are the pivot languages and n is the number of pivot languages in the path. We set $n \in \{0, 1, 2\}$ and consider 3-hop path at most in this paper, considering the computation cost and accuracy drop in longer translation path. Note that when $n = 0$, it is the one-hop (direct) translation.

There exists a large amount of translation paths between X and Y and each path can result in different translation accuracies, or even lower than the direct unsupervised translation, due to the information loss along the multiple translation hops especially when unsupervised translation quality on one hop is low. Therefore, how to choose a good translation path is critical to ensure the accuracy of unsupervised pivot translation. In this section, we introduce the learning to route (LTR) method for the translation path selection.

3.1 Learning to Route

In this section, we first give the description of the problem formulation, and then introduce the training data, features and model used for LTR.

Problem Formulation We formulate the path selection as a translation accuracy prediction problem. The LTR model learns to predict the translation accuracy of each path from language X to Y given the translation accuracy of each hops in the path, and the path with the highest predicted translation accuracy among all the possible paths is chosen as the output.

Training Data We construct the training data for the LTR model in the following steps: (1) From M languages, we choose the distant language pairs whose source and target languages are not in the same language branch. We then choose a small part of the distant language pairs as the development/test set respectively for LTR, and regard the remaining part as the training set for LTR. (2) In order to get the accuracy of different translation paths for the distant language pairs,

as well as to obtain the input features for LTR, we train the unsupervised model for the translation between any languages and obtain the BLEU score of each pair. For M languages, there are total $M(M - 1)$ language pairs and BLEU scores, which requires $M(M - 1)/2$ unsupervised models since one model can handle both translation directions following the common practice in unsupervised NMT (Lample et al., 2018). (3) We then test the BLEU scores of the each possible translation path for the language pairs in the development and test sets, based on the models trained in previous steps. These BLEU scores are regarded as the ground-truth data to evaluate the performance of unsupervised pivot translation. (4) We just get the BLEU scores of a small part of the possible paths in the training set, which are used as the training data for LTR model³. We describe the features of the training data in the next paragraph.

Features We extract several features from the paths in training data for better model training. Without loss of generality, we take a 3-hop path $X \rightarrow Z_1 \rightarrow Z_2 \rightarrow Y$ as an example, and regard it as a token sequence consisting of languages and one-hops: X , $X \rightarrow Z_1$, Z_1 , $Z_1 \rightarrow Z_2$, Z_2 , $Z_2 \rightarrow Y$ and Y . We consider two kinds of features for the token sequence: (1) The token ID. There are a total of 7 tokens in the path shown above. Each token ID is converted into trainable embeddings. For a one-hop token like $Z_1 \rightarrow Z_2$, its embedding is simply the average of the two embeddings of Z_1 and Z_2 . (2) The BLEU score of each language and one-hop, where we get the BLEU score of each language by averaging the accuracy of the one-hop path from or to this language. For example, the BLEU score of the target language Z_1 in $X \rightarrow Z_1$ is calculated by averaging all the BLEU scores of the one-hop translation from other languages to Z_1 , while the BLEU score of the source language Z_1 in $Z_1 \rightarrow Z_2$ is calculated by averaging the BLEU scores of all the one-hop translation from Z_1 to other languages. We concatenate the above two features together in one vector for each language and one-hop token, and get a sequence of features for each path. The BLEU score of the path will be used as the label for the LTR model.

³As described in Footnote 2, we cannot afford to test the BLEU scores of all the possible paths, so we just test a small part of them for training.

Model We use a multi-layer LSTM model to predict the BLEU score of the translation path. The input of the LSTM model is the feature sequence described in the above paragraph. The last hidden of LSTM is multiplied with an one-dimensional vector to predict the BLEU score of the path.

3.2 Discussions

We make brief discussions on some possible baseline routing methods and compare them with our proposed LTR.

Random Routing: We randomly choose a path as the routing result.

Prior Pivoting: We set the pivot language for each language according to prior knowledge⁴. Denote P_X and P_Y as the pivot language for X and Y respectively. The path $X \rightarrow P_X \rightarrow P_Y \rightarrow Y$ will be chosen as the routing result by prior pivoting.

Hop Average: The average of the BLEU scores of each one-hop in the path is taken as the predicted BLEU score for this path. We select the path with the highest predicted BLEU score, as used in the LTR method.

Compared with these simple rule based routing methods described above, LTR chooses the path purely by learning on a part of the ground-truth paths. The feature we designed in LTR can capture the relationship between languages to determine the BLEU score and relative ranking of the paths. This data-driven learning based method (LTR) will be more accurate than the rule based methods. In the next section, we conduct experiments to verify effectiveness of our proposed LTR and compare with the baseline methods.

4 Experiments Design

Our experiments consist of two stages in general. In the first stage, we need to train the unsupervised NMT model between any two languages to get the BLEU scores of each one-hop path. We also get the BLEU scores for a part of multi-hop paths through pivoting, which are used as the training and evaluation data for the second stage. In the second stage, we train the LTR model based on the training data generated in the first stage. In this section, we give brief descriptions of the experiment settings for the unsupervised NMT model

⁴For the languages in each language branch, we choose the language with the largest amount of monolingual data in this branch as the pivot language. All languages in the same language branch share the same pivot language.

training (the first stage) and the LTR model training and path routing (the second stage).

4.1 Experiment Setting for Direct Unsupervised NMT

Datasets We conduct the experiments on 20 languages and a total of $20 \times 19 = 380$ language pairs, which have no bilingual sentence pairs but just monolingual sentences for each language. The languages involved in the experiments can be divided into 4 language branches by the taxonomy of language family: Balto-Slavic branch, Germanic branch, Italic branch and Uralic branch⁵. The language name and its ISO 639-1 code contained in each branch can be found in the supplementary materials (Section 1 and 2).

We collect the monolingual corpus from Wikipedia for each language. We download the language specific Wikipedia contents in XML format⁶, and use WikiExtractor⁷ to extract and clean the texts. We then use the sentence tokenizer from the NLTK toolkit⁸ to generate segmented sentences from Wikipedia documents.

To ensure we have the development and test set for the large amount of language pairs to evaluate the unsupervised translation accuracy in our experiments, we choose the languages that are covered by the common corpus of TED talks, which contains translations between more than 50 languages (Ye et al., 2018)⁹. In this circumstance, we can leverage the development and test set from TED talks for evaluation. Note that in the unsupervised setting, we just leverage monolingual sentences for unsupervised training and only use the bilingual data for developing and testing. In order to alleviate the domain mismatch problem that we train on monolingual data from Wikipedia but test on the evaluation data from TED talks, we also fine-tune the unsupervised models with the small size of monolingual data from TED talks¹⁰. The monolingual data from TED talks is merged with the monolingual data from Wikipedia in the

⁵The first three branches belong to Indo-European family while the last branch is actually a language family. We do not further split the 3 languages in Uralic family into different branches.

⁶For example, we download English Wikipedia contents from <https://dumps.wikimedia.org/enwiki>.

⁷<https://github.com/attardi/wikiextractor>

⁸<https://www.nltk.org/>

⁹<https://github.com/neulab/word-embeddings-for-nmt>

¹⁰https://github.com/ajinkyakulkarni14/TED-Multilingual-Parallel-Corpus/tree/master/Monolingual_data

fine-tuning process, which results in better performance for the unsupervised translation. The size of Wikipedia and TED talks monolingual data can be found in the supplementary materials (Section 3).

All the sentences in the bilingual and monolingual data are first tokenized with Moses tokenizer¹¹ and then segmented into subword symbols using Byte Pair Encoding (BPE) (Sennrich et al., 2016). When training the unsupervised model, we learn the BPE tokens with 60000 merge operations across the source and target languages for each language pair and jointly training the embedding using fastText¹², following the practice in Lample et al. (2018).

Model Configurations We use transformer model as the basic NMT model structure, considering it achieves state-of-the-art accuracy and becomes a popular choice for recent NMT research. We use 4-layer encoder and 4-layer decoder with model hidden size d_{model} and feed-forward hidden size d_{ff} being 512, 2048 following the default configurations in Lample et al. (2018). We use the same model configurations for all the language pairs.

Model Training and Inference We train the unsupervised model with 1 NVIDIA Tesla V100 GPU. One mini-batch contains roughly 4096 source tokens and 4096 target tokens, as used in Lample et al. (2018). We follow the default parameters of Adam optimizer (Kingma and Ba, 2014) and learning rate schedule in Vaswani et al. (2017). During inference, we decode with greedy search for all the languages. We evaluate the translation quality by tokenized case sensitive BLEU (Papineni et al., 2002) with multi-bleu.pl¹³.

4.2 Experiment Setting for Routing

Configurations for Routing We choose the distant language pairs from the 20 languages based on the taxonomy of language family: if two languages are not in the same language branch, then they are regarded as distant languages. We get 294 distant language pairs. As described in Section 3.1, we choose nearly 5% and 10% of the distant language pairs as the development and test set

¹¹<https://github.com/moses-smt/mosesdecoder/blob/master/scripts/tokenizer/tokenizer.perl>

¹²<https://github.com/facebookresearch/fastText>

¹³<https://github.com/moses-smt/mosesdecoder/blob/master/scripts/generic/multi-bleu.perl>

Source	Target	<i>DT</i>	<i>GT</i>	<i>GT</i> (Δ)	Pivot-1	Pivot-2	<i>LTR</i>	<i>LTR</i> (Δ)	Pivot-1	Pivot-2
Da	Gl	6.56	12.14	5.58	En	Es	12.14	5.58	En	Es
Bg	Sv	4.72	9.92	5.20	En	En	9.92	5.20	En	En
Gl	Sv	3.79	8.62	4.83	Es	En	8.62	4.83	Es	En
Sv	Gl	3.70	8.13	4.43	En	Es	8.13	4.43	En	Es
Be	It	2.11	6.40	4.29	Uk	En	5.24	3.13	En	En
Pt	Be	4.76	8.86	4.10	Ru	Ru	8.86	4.10	Ru	Ru
Gl	Da	7.45	11.33	3.88	Es	Es	11.33	3.88	Es	Es
Be	Pt	6.39	9.77	3.38	Ru	Ru	6.39	0.00	-	-
It	Be	2.24	5.19	2.95	Pt	Ru	4.64	2.40	Ru	Ru
Nl	Uk	4.69	7.23	2.54	De	De	7.12	2.53	Ru	Ru

Table 1: The BLEU scores of a part of the distant language pairs in the test set (Please refer to Section 1 and 4 in the supplementary materials for the corresponding full language name and full results). *DT*: direct unsupervised translation. *GT*: the ground-truth best path. *LTR*: the routing results of LTR. (Δ) is the BLEU gap between *GT* or *LTR* and *DT*. Pivot-1 and Pivot-2 are two pivot languages in the path, which will be the same language if the path is 2-hop and will both be empty if the path is 1-hop (direct translation).

Length	1-hop	2-hop	3-hop
Ratio (%)	7.1	53.6	39.3

Table 2: The length distribution of the best translation paths. The ratio is calculated based on all language pairs in the test set.

for routing. Note that if the language pair $X \rightarrow Y$ is in development (test) set, then the language pair $Y \rightarrow X$ will be also in development (test) set. We then enumerate all possible paths between any two language pairs in development and test set, and test the BLEU scores of the each possible path, which are regarded as the ground-truth data to evaluate the performance of the routing method. For the remaining 85% distant language pairs, we just test the BLEU score for 10% of all possible paths, and use these BLEU scores as the label for LTR model training.

We use 2-layer LSTM as described in Section 3.1. The dimension of input feature vector is 6, which includes the embedding of the token ID with size of 5, the BLEU score with size 1 (we normalize the BLEU score into 0-1). We change the depth and width of LSTM, but there is no significant gain in performance.

We use the mean square error as the training loss for the LTR model, and use Adam as the optimizer. The initial learning rate is 0.01. When applying the LTR model on unseen pairs, we predict the BLEU scores of all the possible paths (including 1-hop (direct translation), 2-hop and 3-hop translation path) between the source and target languages, and choose the path with the highest predicted BLEU score as the routing result. Note that when predicting the path with LTR in infer-

ence time, we do not include the pivot language which occurs less than 10 times in training set, which can improve that stability of the *LTR* prediction.

Methods for Comparison We conduct experimental comparisons on different methods described in Section 3 for path selection (routing), including *Random Routing (RR)*, *Prior Pivoting (PP)*, *Hop Average (HA)* and *Learning to Route (LTR)*. We also compare these routing methods with the direct unsupervised translation, denoted as *Direct Translation (DT)*. We list the BLEU score of the best multi-hop path (the ground truth) as a reference, which is denoted as *Ground Truth (GT)*.

5 Results

In this section, we introduce the performance of unsupervised pivot translation for distant languages. We first demonstrate the advantages of unsupervised pivot translation by comparing the best translation path (*GT*) with direction translation (*DT*), and then show the results of our proposed *LTR*. We also compare *LTR* with other routing methods (*RR*, *PP* and *HA*) to demonstrate its effectiveness.

5.1 The Advantage of Unsupervised Pivot Translation

In order to demonstrate the advantage of unsupervised pivot translation for distant languages, we first analyze the distribution of the length of the best translation paths (*GT*), as shown in Table 2. The direction translation (1-hop) only takes a ratio of 7.1%, which means that a majority (92.9%)




Figure 1: The CDF of the BLEU scores for the distant language pairs in the test set. The green curve represents the direct unsupervised translation (*DT*), and the black curve represents the best translation path (*GT*). The other three curves represent the three routing methods for comparison: blue for hop average (*HA*), cyan for prior pivoting (*PP*) and red for our proposed learning to route (*LTR*).

of the distant language pairs need multiple hops to improve the translation accuracy.

We further compare the BLEU score of the best path (*GT*, which is also the upper-bound of different routing methods) with the direct unsupervised translation, and show the results for a part of distant languages pairs in Table 1¹⁴. It can be seen that *GT* can largely outperform the direct translation *DT* with up to 5.58 BLEU points. We further plot the CDF of the BLEU scores on all the distant language pairs in the test set in Figure 1. It can be seen that the CDF curve of *GT* is always in the right part of *DT*, which means better accuracy and demonstrates the advantage of unsupervised pivot translation for distant languages.

5.2 Results of *LTR*

Accuracy of *LTR* Model As our *LTR* selects the good path by ranking according to the predicted BLEU score, we first report the accuracy of selecting the best path. *LTR* can achieve 57% in terms of top-1 accuracy and 86% in terms of top-5 accuracy. Although the top-1 accuracy is not so high, it is acceptable because there exists some other route path just a little bit lower than the best path. We show the routing results of our *LTR* for some language pairs in Table 1. Take the *Nl-Uk* language pair in Table 1 as an example. The routing result of *LTR* for this pair does not match with *GT*, which

¹⁴Due to space limitation, we leave the full results of the distant language pairs in the test set in the supplementary materials (Section 4).

Methods	<i>DT</i>	<i>RR</i>	<i>HA</i>	<i>PP</i>	<i>LTR</i>	<i>GT</i>
BLEU	6.06	3.40	6.92	7.12	8.33	8.70

Table 3: The performance of different routing methods. The BLEU score is averaged on all the distant language pairs in the test set. The compared methods include: *DT*: direct unsupervised translation, *RR*: random routing, *PP*: prior pivoting, *HA*: hop average, *LTR*: our proposed learning to route, and *GT*: the best translation path (the ground truth).

affects the top-1 accuracy. However, the BLEU gap between our selected path and the best path is as small as 0.09, which has little influence on the BLEU score of the selected path. Our further analysis in the next paragraph shows that the averaged BLEU score that *LTR* achieved in test set is close to that of *GT*.

BLEU Score of Selected Path We further report the BLEU score of the translation path selected by *LTR* as well as other routing methods in Table 3, where the BLEU score is averaged over all the distant language pairs in the test set. It can be seen that compared with direct unsupervised translation (*DT*) which achieves 6.06 averaged BLEU score¹⁵, our *LTR* can achieve 2.27 BLEU points improvement on average, and is just 0.37 points lower than the ground truth best path (*GT*). The small gap between the ground truth and *LTR* demonstrates that although *LTR* fails to select the best path in 43% of the distant pairs (just 57% in terms of top-1 accuracy), it indeed chooses the path which has a BLEU score slightly lower than the best path. Random routing (*RR*) even performs worse than *DT*, demonstrating the routing problem is non-trivial. *LTR* outperforms *PP* and *HA* by more than 1 BLEU point on average. We also show the CDF of the BLEU scores of different methods in Figure 1, which clearly shows that *LTR* can outperform the *PP* and *HA* routing methods, demonstrating the effectiveness of the proposed *LTR*.

5.3 Extension to Supervised Pivoting

In the previous experiments, we rely purely on unsupervised NMT for pivot translation, assuming that the translation on each hop cannot leverage any bilingual sentence pairs. However, there in-

¹⁵The averaged BLEU score seems not high, because the unsupervised translations between some hard languages in the test set obtain really low BLEU scores, which affects the average score.

Source	Target	<i>DT</i>	<i>GT</i> -unsup	<i>GT</i> -sup	Δ	Source	Target	<i>DT</i>	<i>GT</i> -unsup	<i>GT</i> -sup	Δ
Da	Gl	6.56	12.14	15.20	8.64	Pt	Be	4.76	8.86	13.03	8.27
Bg	Sv	4.72	9.92	9.92	5.20	Gl	Da	7.45	11.33	15.52	8.07
Gl	Sv	3.79	8.62	9.58	5.79	Be	Pt	6.39	9.77	14.50	8.11
Sv	Gl	3.70	8.13	9.38	5.68	It	Be	2.24	5.19	8.60	6.36
Be	It	2.11	6.40	9.26	7.15	Nl	Uk	4.69	7.23	8.07	3.38

Table 4: The BLEU scores of the same language pairs as shown in Table 1 (Please refer to Section 5 in the supplementary materials for the full results of the test set). *GT*-sup and *GT*-unsup represent the ground-truth best path with and without supervised pivoting. Δ is the BLEU gap between *GT*-sup and *DT*.

Methods	<i>DT</i>	<i>RR</i>	<i>HA</i>	<i>PP</i>	<i>LTR</i>	<i>GT</i>
BLEU	6.06	3.46	7.07	8.84	9.45	9.79

Table 5: The performance of different routing methods when enhanced with supervised pivoting. The BLEU score is averaged on all the distant language pairs in the test set. The compared methods include: *DT*: direct unsupervised translation, *RR*: random routing, *HA*: hop average, *PP*: prior pivoting, *LTR*: our proposed learning to route, and *GT*: the best translation path (the ground truth).

deed exist plenty of bilingual sentence pairs between some languages, especially among the popular languages of the world, such as the official languages of the United Nations and the European Union. If we can rely on some supervised hop in the translation path, the accuracy of the translation for distant languages would be greatly improved.

Take the translation from Danish to Galician as an example. The BLEU score of the direct unsupervised translation is 6.56, while the ground-truth best unsupervised path (Danish→English→Spanish→Galician) can achieve a BLEU score of 12.14, 5.58 points higher than direct unsupervised translation. For the translation on the intermediate hop, i.e., English→Spanish, we have a lot of bilingual data to train a strong supervised translation model. If we replace the unsupervised English→Spanish translation with the supervised counterpart, the BLEU score of the path (Danish→English→Spanish→Galician) can improve from 12.14 to 15.2, with 8.64 points gain over the direct unsupervised translation. Note that the gain is achieved without leveraging any bilingual sentence pairs between Danish and Galician.

Without loss of generality, we choose 6 popular languages (we select English, German, Spanish, French, Finish and Russian to cover each language branch we considered in this work) as the supervised pivot languages and replace the transla-

tions between these languages with the supervised counterparts. Note that we do not leverage any bilingual data related to the source language and target languages, and the supervised models are only used in the intermediate hop of a 3-hop path. For the bilingual sentence pairs between pivot languages, we choose the common corpus of TED talk which contains translations between multiple languages (Ye et al., 2018)¹⁶.

Table 4 shows the performance improvements on the language pairs (the same pairs as shown in Table 1). When enhanced with supervised pivoting, we can achieve more than 8 BLEU points gain over *DT* on 4 language pairs, without using any bilingual data between the source language or target language. We also compare our proposed learning to route method *LTR* with *RR*, *HA* and *PP*, as showed in Table 5. We conduct the experiments on the original development and test set, but removing the language pairs whose source and target languages belong to the supervised pivot languages we choose. It can be seen that *LTR* can still outperform *RR*, *HA* and *PP* and be close to *GT*, demonstrating the effectiveness of *LTR* in the supervised pivoting setting.

6 Conclusions and Future Work

In this paper, we have introduced unsupervised pivot translation for distant language pairs, and proposed the learning to route (*LTR*) method to automatically select a good translation path for a distant language pair. Experiments on 20 languages and totally 294 distant language pairs demonstrate that (1) unsupervised pivot translation achieves large improvements over direct unsupervised translation for distant languages; (2) our proposed *LTR* can select the translation path whose translation accuracy is close to the ground-

¹⁶This is the same dataset where we choose the development and test sets in Section 4.1. The data can be downloaded from <https://github.com/neulab/word-embeddings-for-nmt>.

truth best path; (3) if we leverage supervised translation instead of the unsupervised translation for some popular language pairs in the intermediate hop, we can further boost the performance of unsupervised pivot translation.

For further works, we will leverage more supervised translation hops to improve the performance of unsupervised translation for distant languages. We will extend our method to more distant languages.

References

- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017a. Learning bilingual word embeddings with (almost) no bilingual data. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 451–462.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2017b. Unsupervised neural machine translation. *arXiv preprint arXiv:1710.11041*.
- Yun Chen, Yang Liu, Yong Cheng, and Victor OK Li. 2017. A teacher-student framework for zero-resource neural machine translation. *arXiv preprint arXiv:1705.00753*.
- Yong Cheng, Qian Yang, Yang Liu, Maosong Sun, and Wei Xu. 2017. Joint training for pivot-based neural machine translation. In *Proceedings of IJCAI*.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Andre Girard. 1990. *Routing and dimensioning in circuit-switched networks*. Addison-Wesley Longman Publishing Co., Inc.
- Christian Huitema. 2000. *Routing in the Internet*. Prentice-Hall,.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *CoRR*, abs/1901.07291.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2017. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. Phrase-based & neural unsupervised machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 5039–5049.
- David Liben-Nowell, Jasmine Novak, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. 2005. Geographic routing in social networks. *Proceedings of the National Academy of Sciences*, 102(33):11623–11628.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA.*, pages 311–318.
- Lewis M Paul, Gary F Simons, Charles D Fennig, et al. 2009. Ethnologue: Languages of the world. *Dallas, TX: SIL International*. Available online at www.ethnologue.com/. Retrieved June, 19:2011.
- Samuel Raff. 1983. Routing and scheduling of vehicles and crews: The state of the art. *Computers & Operations Research*, 10(2):63–211.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- Masao Utiyama and Hitoshi Isahara. 2007. A comparison of pivot methods for phrase-based statistical machine translation. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 484–491.
- 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 *NIPS 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6000–6010.
- Hua Wu and Haifeng Wang. 2007. Pivot language approach for phrase-based statistical machine translation. *Machine Translation*, 21(3):165–181.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2018. Unsupervised neural machine translation with weight sharing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 46–55.
- Qi Ye, Sachan Devendra, Felix Matthieu, Padmanabhan Sarguna, and Neubig Graham. 2018. When and why are pre-trained word embeddings useful for neural machine translation. In *HLT-NAACL*.