

Joint Training for Pivot-based Neural Machine Translation

Yong Cheng[#], Qian Yang[#], Yang Liu^{†*}, Maosong Sun[†] and Wei Xu[#]

[#]Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China

[†]State Key Laboratory of Intelligent Technology and Systems

Tsinghua National Laboratory for Information Science and Technology

Department of Computer Science and Technology, Tsinghua University, Beijing, China

chengyong3001@gmail.com, qian-yang12@mails.tsinghua.edu.cn

{liuyang2011, sms, weixu}@tsinghua.edu.cn

Abstract

While recent neural machine translation approaches have delivered state-of-the-art performance for resource-rich language pairs, they suffer from the data scarcity problem for resource-scarce language pairs. Although this problem can be alleviated by exploiting a pivot language to bridge the source and target languages, the source-to-pivot and pivot-to-target translation models are usually independently trained. In this work, we introduce a joint training algorithm for pivot-based neural machine translation. We propose three methods to connect the two models and enable them to interact with each other during training. Experiments on Europarl and WMT corpora show that joint training of source-to-pivot and pivot-to-target models leads to significant improvements over independent training across various languages.

1 Introduction

Recent several years have witnessed the rapid development of neural machine translation (NMT) [Sutskever *et al.*, 2014; Bahdanau *et al.*, 2015], which advocates the use of neural networks to directly model the translation process in an end-to-end way. Thanks to the capability of learning representations from training data, NMT systems have achieved significant improvements over conventional statistical machine translation (SMT) across a variety of language pairs [Junczys-Dowmunt *et al.*, 2016; Johnson *et al.*, 2016].

However, there still remains a major challenge for NMT: large-scale parallel corpora are usually non-existent for most language pairs. This is unfortunate because NMT is a data-hungry approach and requires a large amount of data to fully train translation models. Without sufficient training data, NMT tends to learn poor estimates on low-count events. Zoph *et al.* [2016] indicate that NMT obtains much worse translation quality than SMT when only small-scale parallel corpora are available.

As a result, improving neural machine translation on resource-scarce language pairs has attracted much attention [Firat *et al.*, 2016; Zoph *et al.*, 2016; Johnson *et al.*, 2016].

Most existing methods focus on leveraging parallel corpora of multiple resource-rich language pairs to improve NMT for resource-scarce language pairs. Firat *et al.* [2016] propose multi-way, multilingual neural machine translation to achieve direct source-to-target translation even without parallel data available. Zoph *et al.* [2016] present a transfer learning method that transfers the model parameters trained for resource-rich language pairs to initialize and constrain the translation model training of resource-scarce language pairs. Johnson *et al.* [2016] introduce a universal NMT model for all language pairs, which takes advantage of multilingual data to improve NMT for all languages involved.

Bridging source and target languages with a *pivot* language is another important direction, which has been intensively studied in conventional SMT [Cohn and Lapata, 2007; Wu and Wang, 2007; Utiyama and Isahara, 2007; Bertoldi *et al.*, 2008; Zahabi *et al.*, 2013; El Kholly *et al.*, 2013]. Pivot-based approaches assume that there exist source-pivot and pivot-target parallel corpora, which can be used to train source-to-pivot and pivot-to-target translation models, respectively. One of the most representative approaches, the triangulation approach [Wu and Wang, 2007], is to construct a source-to-target phrase table through combining source-to-pivot and pivot-to-target phrase tables. Another representative approach adopts a pivot-based translation strategy [Wu and Wang, 2007]. As a result, source-to-target translation can be divided into two steps: the source sentence is first translated into a pivot sentence using the source-to-pivot model, which is then translated to a target sentence using the pivot-to-target model. Pivot-based approaches have been widely used in SMT due to its simplicity, effectiveness, and minimum requirement of multilingual data. Recently, Johnson *et al.* [2016] adapt pivot-based approaches to NMT and show that their universal model without incremental training achieves much worse translation performance than pivot-based NMT.

However, pivot-based approaches often suffer from the error propagation problem: the errors made in the source-to-pivot translation will be propagated to the pivot-to-target translation. This can be partly attributed to the discrepancy between source-pivot and pivot-target parallel corpora since they are usually loosely-related or even unrelated. To aggravate the situation, source-to-pivot and pivot-to-target translation models are trained independently, which further enlarges the gap between source and target languages.

*Corresponding Author

In this work, we propose an approach to joint training for pivot-based neural machine translation. The basic idea is to connect the source-to-pivot and pivot-to-target NMT models and enable them to interact with each other during training. This can be done either by encouraging the sharing of word embeddings on the pivot language or by maximizing the likelihood of the cascaded model on a small source-target parallel corpus. Experiments on the Europarl and WMT corpora show that joint training of source-to-pivot and pivot-to-target models obtains significant improvements over independent training.

2 Background

Given a source language sentence \mathbf{x} and a target language sentence \mathbf{y} , we use $P(\mathbf{y}|\mathbf{x};\theta_{x\rightarrow y})$ to denote a standard attention-based neural machine translation model [Bahdanau *et al.*, 2015], where $\theta_{x\rightarrow y}$ is a set of model parameters.

Ideally, the source-to-target model can be trained on a source-target parallel corpus $D_{x,y} = \{(\mathbf{x}^{(s)}, \mathbf{y}^{(s)})\}_{s=1}^S$ using maximum likelihood estimation:

$$\hat{\theta}_{x\rightarrow y} = \operatorname{argmax}_{\theta_{x\rightarrow y}} \left\{ \mathcal{L}(\theta_{x\rightarrow y}) \right\} \quad (1)$$

where the log-likelihood is defined as:

$$\mathcal{L}(\theta_{x\rightarrow y}) = \sum_{s=1}^S \log P(\mathbf{y}^{(s)}|\mathbf{x}^{(s)}; \theta_{x\rightarrow y}) \quad (2)$$

Unfortunately, parallel corpora are usually not readily available for low-resource language pairs. Instead, one can assume that there exist a third language called *pivot* with source-pivot and pivot-target parallel corpora available. As a result, it is possible to bridge the source and target languages with the pivot [Cohn and Lapata, 2007; Wu and Wang, 2007; Bertoldi *et al.*, 2008; Zahabi *et al.*, 2013; El Kholy *et al.*, 2013]. Figure 1 illustrates the idea of pivot-based translation.

Let \mathbf{z} be a pivot language sentence. The source-to-target model can be decomposed into two sub-models by treating the pivot sentence as a latent variable:

$$\begin{aligned} & P(\mathbf{y}|\mathbf{x}; \theta_{x\rightarrow z}, \theta_{z\rightarrow y}) \\ &= \sum_{\mathbf{z}} P(\mathbf{z}|\mathbf{x}; \theta_{x\rightarrow z})P(\mathbf{y}|\mathbf{z}; \theta_{z\rightarrow y}) \end{aligned} \quad (3)$$

Let $D_{x,z} = \{(\mathbf{x}^{(m)}, \mathbf{z}^{(m)})\}_{m=1}^M$ be a source-pivot parallel corpus, and $D_{z,y} = \{(\mathbf{z}^{(n)}, \mathbf{y}^{(n)})\}_{n=1}^N$ be a pivot-target parallel corpus. The source-to-pivot and pivot-to-target models can be **independently** trained on the two parallel corpora, respectively:

$$\hat{\theta}_{x\rightarrow z} = \operatorname{argmax}_{\theta_{x\rightarrow z}} \left\{ \mathcal{L}(\theta_{x\rightarrow z}) \right\} \quad (4)$$

$$\hat{\theta}_{z\rightarrow y} = \operatorname{argmax}_{\theta_{z\rightarrow y}} \left\{ \mathcal{L}(\theta_{z\rightarrow y}) \right\} \quad (5)$$

where the log-likelihoods are defined as:

$$\mathcal{L}(\theta_{x\rightarrow z}) = \sum_{m=1}^M \log P(\mathbf{z}^{(m)}|\mathbf{x}^{(m)}; \theta_{x\rightarrow z}) \quad (6)$$

$$\mathcal{L}(\theta_{z\rightarrow y}) = \sum_{n=1}^N \log P(\mathbf{y}^{(n)}|\mathbf{z}^{(n)}; \theta_{z\rightarrow y}) \quad (7)$$

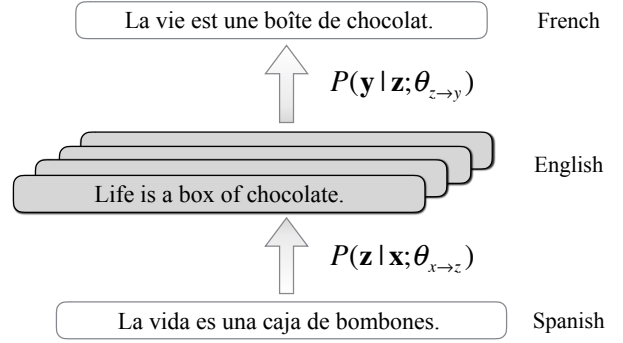


Figure 1: The illustration of translation on Spanish-French with English as the pivot language. The Spanish-English NMT model $P(\mathbf{z}|\mathbf{x};\theta_{x\rightarrow z})$ first transforms a Spanish sentence into latent English sentences, from which English-French NMT model $P(\mathbf{y}|\mathbf{z};\theta_{z\rightarrow y})$ attempts to generate a French sentence corresponding to the Spanish sentence.

Given an unseen source sentence to be translated \mathbf{x} , the decision rule is given by:

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y}} \left\{ \sum_{\mathbf{z}} P(\mathbf{z}|\mathbf{x}; \hat{\theta}_{x\rightarrow z})P(\mathbf{y}|\mathbf{z}; \hat{\theta}_{z\rightarrow y}) \right\} \quad (8)$$

Due to the exponential search space of the pivot language, the decoding process is usually approximated with two steps. The first step translates the source sentence \mathbf{x} into a pivot sentence:

$$\hat{\mathbf{z}} = \operatorname{argmax}_{\mathbf{z}} \left\{ P(\mathbf{z}|\mathbf{x}; \hat{\theta}_{x\rightarrow z}) \right\} \quad (9)$$

Then, the pivot sentence is translated to a target sentence:

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y}} \left\{ P(\mathbf{y}|\hat{\mathbf{z}}; \hat{\theta}_{z\rightarrow y}) \right\} \quad (10)$$

Although pivot-based approaches are widely used for addressing the data scarcity problem in machine translation, they suffer from cascaded translation errors: the mistakes made in the source-to-pivot translation as shown in Eq. (9) will be propagated to the pivot-to-target translation as shown in Eq. (10). This can be partly attributed to the **model discrepancy** problem: the source-to-pivot and pivot-to-target models are quite different in terms of vocabulary and parameter space because the source-pivot and pivot-target parallel corpora are usually loosely-related or even unrelated. To make things worse, the source-to-pivot model $P(\mathbf{z}|\mathbf{x}; \theta_{x\rightarrow z})$ and the pivot-to-target model $P(\mathbf{y}|\mathbf{z}; \theta_{z\rightarrow y})$ are trained on the two parallel corpora **independently**, which further increases the discrepancy between two models.

Therefore, it is important to reduce the discrepancy between source-to-pivot and pivot-to-target models to further improve pivot-based neural machine translation.

3 Joint Training for Pivot-based NMT

3.1 Training Objective

To alleviate the model discrepancy problem, we propose an approach to joint training for pivot-based neural machine translation. The basic idea is to connect source-to-pivot and pivot-to-target models and enable them to interact with each other during training. Our new training objective is given by:

$$\begin{aligned} \mathcal{J}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y}) \\ = \mathcal{L}(\boldsymbol{\theta}_{x \rightarrow z}) + \mathcal{L}(\boldsymbol{\theta}_{z \rightarrow y}) + \lambda \mathcal{R}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y}) \end{aligned} \quad (11)$$

Note that the training objective consists of three parts: the source-to-pivot likelihood $\mathcal{L}(\boldsymbol{\theta}_{x \rightarrow z})$, the pivot-to-target likelihood $\mathcal{L}(\boldsymbol{\theta}_{z \rightarrow y})$, and a connection term $\mathcal{R}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y})$. The hyper-parameter λ is used to balance the preference between likelihoods and the connection term.

We expect that the connection term associates the source-to-pivot model $\boldsymbol{\theta}_{x \rightarrow z}$ with the pivot-to-target model $\boldsymbol{\theta}_{z \rightarrow y}$ and enables the interaction between two models during training. In the following subsection, we will introduce the three connection terms used in our experiments.

3.2 Connection Terms

It is difficult to connect the source-to-pivot and pivot-to-target models during training because the source-to-pivot and pivot-to-target models are distantly-related by definition. More importantly, NMT lacks linguistically interpretable language structures such as phrases in SMT to achieve a direct connection at the parameter level [Wu and Wang, 2007].

Fortunately, both the source-to-pivot and pivot-to-target models include the word embeddings of the pivot language as parameters. It is possible to connect the two models via pivot word embeddings.

More formally, let $\mathcal{V}_{x \rightarrow z}^z$ be the pivot vocabulary of the source-to-pivot model and $\mathcal{V}_{z \rightarrow y}^z$ be the pivot vocabulary of the pivot-to-target model. We use w to denote a word in the pivot language and $\boldsymbol{\theta}_{x \rightarrow z}^w \in \mathbb{R}^d$ to denote the vector representation of w in the source-to-pivot model. $\boldsymbol{\theta}_{z \rightarrow y}^w \in \mathbb{R}^d$ is defined in a similar way.

Our first connection term encourages the two models to generate the same vector representations for pivot words in the intersection of two vocabularies:

$$\begin{aligned} \mathcal{R}_{\text{hard}}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y}) \\ = \prod_{w \in \mathcal{V}_{x \rightarrow z}^z \cap \mathcal{V}_{z \rightarrow y}^z} \delta(\boldsymbol{\theta}_{x \rightarrow z}^w, \boldsymbol{\theta}_{z \rightarrow y}^w) \end{aligned} \quad (12)$$

where $\delta(\boldsymbol{\theta}_{x \rightarrow z}^w, \boldsymbol{\theta}_{z \rightarrow y}^w) = 1$ if the two vectors $\boldsymbol{\theta}_{x \rightarrow z}^w$ and $\boldsymbol{\theta}_{z \rightarrow y}^w$ are identical. Otherwise, $\delta(\boldsymbol{\theta}_{x \rightarrow z}^w, \boldsymbol{\theta}_{z \rightarrow y}^w) = 0$. This constraint requires that $\boldsymbol{\theta}_{x \rightarrow z}^w$ and $\boldsymbol{\theta}_{z \rightarrow y}^w$ share the same embeddings of the intersected pivot words.

As word embeddings seem hardly to be exactly identical due to the divergence of natural languages, an alternative is to soften the above hard matching constraint by penalizing the Euclidean distance between two vectors:

$$\begin{aligned} \mathcal{R}_{\text{soft}}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y}) \\ = - \sum_{w \in \mathcal{V}_{x \rightarrow z}^z \cap \mathcal{V}_{z \rightarrow y}^z} \|\boldsymbol{\theta}_{x \rightarrow z}^w - \boldsymbol{\theta}_{z \rightarrow y}^w\|_2 \end{aligned} \quad (13)$$

The third connection term assumes that there is a small bridging source-target parallel corpus $\bar{D}_{x,y} = \{(\mathbf{x}^{(l)}, \mathbf{y}^{(l)})\}_{l=1}^L$ (*Bridging Corpus*) available. The connection term is defined as the log-likelihood of the bridging data:

$$\begin{aligned} \mathcal{R}_{\text{likelihood}}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y}) \\ = \sum_{l=1}^L \log P(\mathbf{y}^{(l)} | \mathbf{x}^{(l)}; \boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y}) \end{aligned} \quad (14)$$

$$= \sum_{l=1}^L \log \sum_{\mathbf{z}} P(\mathbf{z} | \mathbf{x}^{(l)}; \boldsymbol{\theta}_{x \rightarrow z}) P(\mathbf{y}^{(l)} | \mathbf{z}; \boldsymbol{\theta}_{z \rightarrow y}) \quad (15)$$

3.3 Training

In training, our goal is to find the optimal source-to-pivot and pivot-to-target model parameters that maximize the training objective:

$$\hat{\boldsymbol{\theta}}_{x \rightarrow z}, \hat{\boldsymbol{\theta}}_{z \rightarrow y} = \underset{\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y}}{\operatorname{argmax}} \left\{ \mathcal{J}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y}) \right\} \quad (16)$$

The partial derivative of $\mathcal{J}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y})$ with respect to the parameters $\boldsymbol{\theta}_{x \rightarrow z}$ of the source-to-pivot model can be calculated as:

$$\begin{aligned} \frac{\partial \mathcal{J}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y})}{\partial \boldsymbol{\theta}_{x \rightarrow z}} \\ = \sum_{m=1}^M \frac{\partial \log P(\mathbf{z}^{(m)} | \mathbf{x}^{(m)}; \boldsymbol{\theta}_{x \rightarrow z})}{\partial \boldsymbol{\theta}_{x \rightarrow z}} + \\ \lambda \frac{\partial \mathcal{R}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y})}{\partial \boldsymbol{\theta}_{x \rightarrow z}} \end{aligned} \quad (17)$$

The partial derivative with respect to the parameters $\boldsymbol{\theta}_{z \rightarrow y}$ can be calculated similarly.

The gradients of the first and second connection terms $\mathcal{R}_{\text{hard}}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y})$ and $\mathcal{R}_{\text{soft}}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y})$ with respect to model parameters are easy to calculate. However, calculating the gradients of the third connection term $\mathcal{R}_{\text{likelihood}}(\boldsymbol{\theta}_{x \rightarrow z}, \boldsymbol{\theta}_{z \rightarrow y})$ involves enumerating all possible pivot sentences in an exponential search space (see Eq. (15)).

To alleviate this problem, we follow standard practice to use a subset to approximate the full space [Shen *et al.*, 2016; Cheng *et al.*, 2016]. Two methods can be used to generate a subset: sampling k translations from the full space [Shen *et al.*, 2016] or generating a top- k list of candidate translations [Cheng *et al.*, 2016]. We find that using top- k lists leads to better results than sampling in our experiments.

We use standard mini-batched stochastic gradient descent algorithms to optimize model parameters. In each iteration, three mini-batches are constructed by randomly selecting sentence pairs from the source-pivot parallel corpus $D_{x,z}$, the pivot-target parallel corpus $D_{z,y}$, and the bridging source-target parallel corpus $D_{x,y}$ (only available for the third connection term), respectively. After separate gradient calculation in each mini-batch, the gradients are collected to update model parameters.

4 Experiments

4.1 Setup

We evaluated our approach on two translation tasks:

Corpus	Lang.		Source	Target
Europarl	es-en	# Sent.	850K	
		# Word	22.32M	21.44M
		Vocab.	118.81K	78.37K
	de-en	# Sent.	840K	
		# Word	20.88M	21.91M
		Vocab.	242.87K	80.44KM
	en-fr	# Sent.	900K	
		# Word	22.56M	25.00M
		Vocab.	80.08K	98.50K
WMT	es-en	# Sent.	6.78M	
		# Word	183.01M	166.28M
		Vocab.	0.98M	0.91M
	en-fr	# Sent.	9.29M	
		# Word	227.06M	258.95M
		Vocab.	0.23M	1.19M

Table 1: Characteristics of Spanish-English, German-English and English-French datasets on the Europarl and WMT corpora. “es” denotes Spanish, “en” denotes English, “de” denotes German, and “fr” denotes French.

1. Spanish-English-French: Spanish as the source language, English as the pivot language, and French as the target language,
2. German-English-French: German as the source language, English as the pivot language, and French as the target language.

Table 1 shows the statistics of the Europarl and WMT corpora used in our experiments. We use the `tokenize.perl` script for tokenization. For each language pair, we remove the empty lines and retain sentence pairs with no more than 50 words. To avoid the intersection of the source-pivot and pivot-target corpora, we split the overlapped pivot-language sentences of source-to-pivot and pivot-to-target corpora into two separate parts with equal size and merge them separately with the non-overlapping parts for each language pair.

The Europarl corpus consists of 850K Spanish-English sentence pairs with 22.32M Spanish words and 21.44M English words, 840K German-English sentence pairs with 20.88M German words and 21.91M English words, and 900K English-French sentence pairs with 22.56M English words and 25.00M French words. The WMT 2006 shared task datasets are used as the development and test sets. The evaluation metric is case-insensitive BLEU [Papineni *et al.*, 2002] as calculated by the `multi-bleu.perl` script.

The WMT corpus is composed of the Common Crawl, News Commentary, Europarl v7 and UN corpora. The Spanish-English parallel corpus consists of 6.78M sentence pairs with 183.01M Spanish words and 166.28M English words. The English-French parallel corpus comprises 9.29M sentence pairs with 227.06M English words and 258.95M French words. The *newstest2011* and *newstest2012* datasets serve as development and test sets. We use case-sensitive BLEU as the evaluation metric.

We use the attention-based neural machine translation system RNNSEARCH [Bahdanau *et al.*, 2015] in our experiments. For the Europarl corpus in Table 1, we set the vocab-

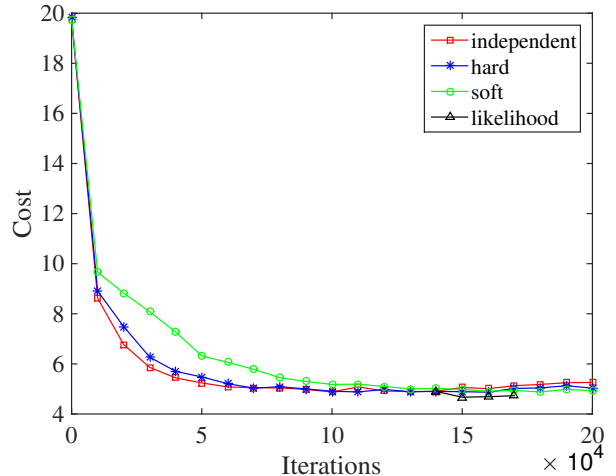


Figure 2: Learning curves of independent training and joint training on different connection terms.

ulary size of all the languages to 30K which covers over 99% of words for English, Spanish and French and over 97% for German. We follow Jean *et al.* [2015] to address rare words. For Spanish-English and English-French corpora from the WMT corpus, due to large vocabulary size, we adopt byte pair encoding [Sennrich *et al.*, 2016b] to split rare words into sub-words. We set the size of sub-words to 43K, 33K, and 43K for Spanish, English, and French, respectively. These sub-words cover 100% of the text.

We set the hyper-parameter λ for balancing between likelihood and the connection term to 1.0. The threshold of gradients is set to 0.1. The bridging source-target parallel corpus contains 100K sentence pairs that are sampled from the training corpus but do not overlap with the training corpus through removing them. We set k to 10 for calculating top- k lists. The parameters for the source-to-pivot and pivot-to-target translation models in the likelihood connection term are initialized by pre-trained model parameters.

4.2 Results on the Europarl Corpus

Table 2 shows the comparison results between our joint training on three connection terms and independent training on the Europarl Corpus. For the source-to-target translation task, we present source-to-pivot, pivot-to-target and source-to-target translation results compared with independent training. In Spanish-to-French translation task, soft connection achieves significant improvements in Spanish-to-French and Spanish-to-English directions although hard connection still performs comparably with independent training. In German-to-French translation task, soft and hard connections also achieve comparable performances with independent training.

In contrast, we find that likelihood connection dramatically improves translation performance on both Spanish-to-French and German-to-French corpora (up to +2.80 BLEU scores in Spanish-to-French and up to 2.23 BLEU scores in German-to-French). The significant improvements for source-to-pivot and pivot-to-target directions are also observed. This sug-

Training	Connection	Dataset	Spanish-French			German-French		
			es → en	en → fr	es → fr	de → en	en → fr	de → fr
indep.	-	Dev.	31.53	30.46	29.52	26.52	30.46	23.67
		Test	31.54	31.42	29.79	26.47	31.42	23.70
joint	hard	Dev.	31.81	30.18	29.11	26.48	30.47	23.87
		Test	31.55	31.13	29.93	26.58	31.35	23.88
	soft	Dev.	32.11**	30.41	30.24**	26.92	30.39	23.99
		Test	31.96*	31.40	30.57**	26.55	31.33	23.79
	likelihood	Dev.	33.35**	31.63**	32.45**	27.90**	31.49**	25.21**
		Test	33.54**	32.33**	32.59**	28.01**	32.34**	25.93**

Table 2: Comparison between independent and joint training on Spanish-French and German-French translation tasks using the Europarl corpus. English is treated as the pivot language. The BLEU scores are case-insensitive. “*”: significantly better than independent training ($p < 0.05$); “**”: significantly better than independent training ($p < 0.01$). We use the statistical significance test with paired bootstrap resampling [Koehn, 2004].

GROUNDTRUTH	source	uno no debe empezar a dudar en público del valor , tampoco del valor inmediato en el aspecto material , de esta ampliación .
	pivot	it makes little sense to start to doubt in public the value , including the direct value at a material level , of this enlargement .
	target	il ne faut pas commencer à douter en public de la valeur , ni de la valeur immédiate , de la portée matérielle de cet élargissement .
INDEP.	pivot	one should not begin <i>to doubt in</i> terms of <i>the value</i> of courage , or of the immediate effect on material , of <i>enlargement</i> . [BLEU: 13.33]
	target	<i>il ne</i> faudrait pas se tromper en termes de valeur de courage ou d ’ effet immédiat sur le matériel , l ’ <i>élargissement</i> . [BLEU: 8.69]
HARD	pivot	one must not <i>start to doubt in</i> the public , not the immediate value in the material , <i>this enlargement</i> . [BLEU: 19.02]
	target	<i>il ne faut pas</i> que l ’ on <i>commence à douter</i> , ni au public , ni à la <i>valeur immédiate</i> , à l ’ <i>élargissement</i> . [BLEU: 25.36]
SOFT	pivot	one cannot start thinking of the value of <i>the value</i> , and the immediate courage , <i>of this enlargement</i> . [BLEU: 21.57]
	target	on <i>ne peut pas commencer à</i> penser à la valeur <i>de la valeur</i> , au courage immédiat , <i>de cet élargissement</i> . [BLEU: 26.60]
LIKLIHHOD	pivot	one must not <i>start to</i> question the value of <i>the value</i> , either of the immediate value in the material aspect , <i>of this enlargement</i> . [BLEU: 24.60]
	target	<i>il ne faut pas commencer à</i> remettre en question la valeur <i>de la valeur</i> , <i>ni de la valeur immédiate</i> de l ’ aspect matériel , <i>de cet élargissement</i> . [BLEU: 56.40]

Table 3: Examples of pivot and target translations using the pivot-based translation strategy. We observe that our approaches generate better translations for both pivot and target sentences. We italicize *correct translation segments* which are no short than 2-grams.

Method	Dataset	Spanish-French (WMT)		
		es → en	en → fr	es → fr
indep.	Dev.	27.62	27.90	24.92
	Test	29.03	25.82	24.60
likelihood	Dev.	28.92**	28.52**	26.24**
	Test	30.43**	26.36**	25.95**

Table 4: Results on Spanish-French translation task from WMT corpus. English is treated as the pivot language. “**”: significantly better than independent training ($p < 0.01$).

gests that introducing source-to-target parallel corpus to maximize $P(y|x; \theta_{x \rightarrow z}, \theta_{z \rightarrow y})$ with z as latent variables makes the source-to-pivot and pivot-to-target translation models improved collaboratively.

Table 3 shows pivot and target translation examples of in-

Systems	<i>newstest2012</i>	<i>newstest2013</i>
Firat et al. [2016]	22.08	21.70
<i>this work</i>	25.95**	25.78**

Table 5: Comparison with Firat et al. [2016] on Spanish-French translation task from WMT corpus. “**”: significantly better than independent training ($p < 0.01$).

dependent training and our approaches. Apparently, our approaches improve the translation quality of both pivot sentences and target sentences.

According to Eq. (3), the cost of the source-to-target model can be decomposed into the costs of source-to-pivot and pivot-to-target models. As we have a small test trilingual corpus, (Spanish, English, French), we use the English sentence to approximate the latent variables in Eq. (3). Then

Corpus	Lang.	source-target	source-pivot-target
Europarl	es → fr	26.37	29.79
	de → fr	14.02	23.70
WMT	es → fr	11.75	24.60

Table 6: Translation performance on bridging corpora.

# Sent.	es → en	en → fr	es → fr
0	31.53	30.46	29.52
1K	32.64	30.29	30.23
10K	32.92	30.93	31.51
50K	33.29	31.57	32.40
100K	33.35	31.63	32.45

Table 7: Effect of the data size of source-to-target parallel corpora (*Bridge Corpora*) used in LIKELIHOOD.

we calculate the cost of Spanish-to-French on the trilingual corpus. Figure 2 shows the learning curves of the test cost of independent training and joint training on three connection terms. We can find that hard and soft connections learn slower than the independent training. Likelihood connection drives its cost lower after fine-tuning based on pre-trained parameters in just 10K iterations.

4.3 Results on the WMT Corpus

Likelihood connection obtains the best performance in our three proposed connection terms according to experiments on the Europarl corpus. To further verify its effectiveness, we evaluate all the methods on the WMT corpus, which is much larger than Europarl. As shown in Table 4, we find that likelihood connection still outperforms independent training significantly on Spanish-to-English, English-to-French and Spanish-to-French directions (up to +1.18 BLEU scores in Spanish-to-French).

We also compare our approach with Firat et al. [2016]. Although our parallel training corpus is much smaller than theirs, Table 5 shows that our approach achieves substantial improvements over them (up to +4.32 BLEU).

4.4 Effect of Bridging Corpora

As bridging corpora are used in likelihood connection term for “bridging” the source-to-pivot and pivot-to-target translation models, why do not we directly build NMT systems with these corpora?

We train source-to-target models using bridging corpora and show translation results in Table 6. We observe that performance is much worse than that in Table 2 and Table 4 using the pivot-based translation strategy. It indicates that NMT yields poor performance on low-resource languages and the pivot-based translation strategy remedies the drawback to alleviate data scarcity effectively.

We also investigate the effect of the data size of bridging corpora on the likelihood connection. Table 7 shows that using a small parallel corpus (1K sentence pairs) has made a measurable improvement. When more than 50K sentence pairs are added, the further improvements become modest.

This finding suggests that a small corpus suffices to enable the likelihood connection to reach the reasonable performance.

5 Related Work

Our work is inspired by two lines of research: (1) machine translation with pivot languages and (2) incorporating additional data resource for NMT.

5.1 Machine Translation with Pivot Languages

Machine translation suffers from the scarcity of parallel corpora. For low-resource language pairs, a pivot language is introduced to “bridge” source and target languages in statistical machine translation [Cohn and Lapata, 2007; Wu and Wang, 2007; Utiyama and Isahara, 2007; Zahabi *et al.*, 2013; El Kholy *et al.*, 2013].

In NMT, Firat et al. [2016] and Johnson et al [2016] propose multi-way, multilingual NMT models that enable zero-resource machine translation. They also need to apply pivot-based approaches into NMT to ameliorate the performance of zero-resource machine translation. Zoph et al. [2016] adopt transfer learning to fine-tune parameters of the low-resource language pairs using trained parameters on the high-resource language pairs. However, our approach aims to jointly train source-to-pivot and pivot-to-target NMT models, which can alleviate the error propagation of pivot-based approaches. We use connection terms to “bridge” these two models and make them benefit each other.

5.2 Incorporating Additional Data Resources for NMT

Due to the limit in quantity, quality and coverage for parallel corpora, additional data resources have raised attention recently. Gulcchre et al [2015] propose to incorporate target-side monolingual corpora as a language model for NMT. Sennrich, Haddow, and Birch [2016a] pair the target monolingual corpora with its automatic, NMT-generated translations, then merge them with parallel corpora for retraining source-to-target model. Zhang and Zong [2016] propose two approaches, self-training algorithm and multi-task learning framework, to incorporate source-side monolingual corpora. Cheng et al. [2016] introduce an autoencoder framework to reconstruct monolingual sentences using source-to-target and target-to-source NMT models. The proposed model can exploit both source and target monolingual corpora. In contrast to Cheng et al. [2016], the objective of our likelihood connection is to maximize the probability of target-language sentences through pivot sentences given source sentences. We use a small source-to-target parallel corpus to train source-to-pivot and pivot-to-target NMT models jointly.

6 Conclusion

We present joint training for pivot-based neural machine translation. Experiments on different language pairs confirm that our approach achieves significant improvements. It is appealing to combine source and pivot sentences for decoding target sentences [Firat *et al.*, 2016] or train a multi-source model directly [Zoph and Knight, 2016]. We also plan to study better connection terms for our joint training.

Acknowledgements

This research is supported by the 863 Program (2015AA015407), the National Natural Science Foundation of China (No. 61532001, No. 61522204, No. 61432013), Tsinghua Initiative Research Program grants 20151080475, MOE Online Education Research Center (Quantong Fund) grant 2017ZD203, and gift funds from Huawei and Ant Financial. We sincerely thank the reviewers for their valuable suggestions.

References

- [Bahdanau *et al.*, 2015] Dzmitry Bahdanau, KyungHyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR*, 2015.
- [Bertoldi *et al.*, 2008] Nicola Bertoldi, Madalina Barbaiani, Marcello Federico, and Roldano Cattoni. Phrase-based statistical machine translation with pivot languages. In *IWSLT*, 2008.
- [Cheng *et al.*, 2016] Yong Cheng, Wei Xu, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. Semi-supervised learning for neural machine translation. In *Proceedings of ACL*, 2016.
- [Cohn and Lapata, 2007] Trevor Cohn and Mirella Lapata. Machine translation by triangulation: Making effective use of multi-parallel corpora. In *Proceedings of ACL*, 2007.
- [El Kholy *et al.*, 2013] Ahmed El Kholy, Nizar Habash, Gregor Leusch, Evgeny Matusov, and Hassan Sawaf. Language independent connectivity strength features for phrase pivot statistical machine translation. In *Proceedings of ACL*, 2013.
- [Firat *et al.*, 2016] Orhan Firat, Baskaran Sankaran, Yaser Al-Onaizan, Fatos T Yarman Vural, and Kyunghyun Cho. Zero-resource translation with multi-lingual neural machine translation. In *Proceedings of EMNLP*, 2016.
- [Gulcehre *et al.*, 2015] Caglar Gulcehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, Loïc Barrault, Hui-Chi Lin, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. On using monolingual corpora in neural machine translation. arXiv:1503.03535 [cs.CL], 2015.
- [Jean *et al.*, 2015] Sebastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. On using very large target vocabulary for neural machine translation. In *Proceedings of ACL*, 2015.
- [Johnson *et al.*, 2016] Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *arXiv preprint arXiv:1611.04558*, 2016.
- [Junczys-Dowmunt *et al.*, 2016] Marcin Junczys-Dowmunt, Tomasz Dwojak, and Hieu Hoang. Is neural machine translation ready for deployment? a case study on 30 translation directions. arXiv:1610.01108v2, 2016.
- [Koehn, 2004] Philipp Koehn. Statistical significance tests for machine translation evaluation. In *Proceedings of EMNLP*, 2004.
- [Papineni *et al.*, 2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of ACL*, 2002.
- [Sennrich *et al.*, 2016a] Rico Sennrich, Barry Haddow, and Alexandra Birch. Improving neural machine translation models with monolingual data. In *Proceedings of ACL*, 2016.
- [Sennrich *et al.*, 2016b] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of ACL*, 2016.
- [Shen *et al.*, 2016] Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. Minimum risk training for neural machine translation. In *Proceedings of ACL*, 2016.
- [Sutskever *et al.*, 2014] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. In *Proceedings of NIPS*, 2014.
- [Utiyama and Isahara, 2007] Masao Utiyama and Hitoshi Isahara. A comparison of pivot methods for phrase-based statistical machine translation. In *HLT-NAACL*, 2007.
- [Wu and Wang, 2007] Hua Wu and Haifeng Wang. Pivot language approach for phrase-based statistical machine translation. *Machine Translation*, 2007.
- [Zahabi *et al.*, 2013] Samira Tofighi Zahabi, Somayeh Bakhshaei, and Shahram Khadivi. Using context vectors in improving a machine translation system with bridge language. In *Proceedings of ACL*, 2013.
- [Zhang and Zong, 2016] Jiajun Zhang and Chengqing Zong. Exploiting source-side monolingual data in neural machine translation. In *Proceedings of EMNLP*, 2016.
- [Zoph and Knight, 2016] Barret Zoph and Kevin Knight. Multi-source neural translation. In *Proceedings of NAACL*, 2016.
- [Zoph *et al.*, 2016] Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. Transfer learning for low-resource neural machine translation. In *Proceedings of EMNLP*, 2016.