

Reference Language based Unsupervised Neural Machine Translation

Zuchao Li^{1,2,3}, Hai Zhao^{1,2,3,*}, Rui Wang^{4,*}, Masao Utiyama⁴, and Eiichiro Sumita⁴

¹Department of Computer Science and Engineering, Shanghai Jiao Tong University (SJTU)

²Key Laboratory of Shanghai Education Commission for Intelligent Interaction and Cognitive Engineering, Shanghai Jiao Tong University, Shanghai, China

³MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, China

⁴National Institute of Information and Communications Technology (NICT), Kyoto, Japan

charlee@sjtu.edu.cn, zhaohai@cs.sjtu.edu.cn, {wangrui, mutiyama, eiichiro.sumita}@nict.go.jp

Abstract

Exploiting a common language as an auxiliary for better translation has a long tradition in machine translation and lets supervised learning-based machine translation enjoy the enhancement delivered by the well-used pivot language in the absence of a source language to target language parallel corpus. The rise of unsupervised neural machine translation (UNMT) almost completely relieves the parallel corpus curse, though UNMT is still subject to unsatisfactory performance due to the vagueness of the clues available for its core back-translation training. Further enriching the idea of pivot translation by extending the use of parallel corpora beyond the source-target paradigm, we propose a new reference language-based framework for UNMT, RUNMT, in which the reference language only shares a parallel corpus with the source, but this corpus still indicates a signal clear enough to help the reconstruction training of UNMT through a proposed reference agreement mechanism. Experimental results show that our methods improve the quality of UNMT over that of a strong baseline that uses only one auxiliary language, demonstrating the usefulness of the proposed reference language-based UNMT and establishing a good start for the community.

1 Introduction

Recently, the application of neural machine translation (NMT) (Sutskever et al., 2014;

* Corresponding authors. This paper was partially supported by National Key Research and Development Program of China (No. 2017YFB0304100), Key Projects of National Natural Science Foundation of China (U1836222 and 61733011), Huawei-SJTU Long Term AI Project, Cutting-edge Machine Reading Comprehension and Language Model. Rui Wang was partially supported by JSPS grant-in-aid for early-career scientists (19K20354): “Unsupervised Neural Machine Translation in Universal Scenarios” and NICT tenure-track researcher startup fund “Toward Intelligent Machine Translation”.

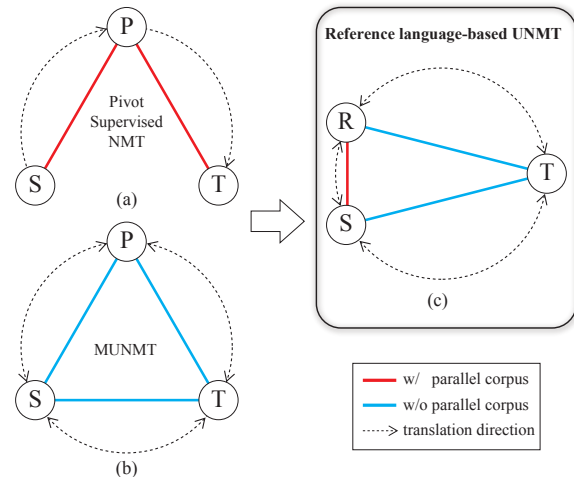


Figure 1: Schemas of (a) pivot supervised NMT, (b) MUNMT, (c) our proposed RUNMT, where \mathcal{S} stands for source language, \mathcal{T} for target language, \mathcal{P} for pivot language in pivot translation, and \mathcal{R} for the reference language in RUNMT.

Bahdanau et al., 2015) to standard benchmarks has achieved great success (Wu et al., 2016; Gehring et al., 2017; Vaswani et al., 2017) because of advances in deep learning and the availability of large-scale parallel corpora; however, the applicability of MT systems is limited because of their reliance on large parallel corpora for the majority of language pairs. In real-world situations, the majority of language pairs have very little parallel data, although large volumes of monolingual data are available for each language. UNMT removes the dependence on parallel corpora, relying only on monolingual corpora in each language (Reddi et al., 2018; Lample et al., 2018a,b; Conneau and Lample, 2019; Li et al., 2019b).

UNMT uses translation symmetry for dual learning in each language direction. Existing UNMT models are mainly built on the encoder-decoder schema. The essence of UNMT is to

learn unsupervised cross-lingual word alignment and/or sentence alignment. For unsupervised word alignment, the most popular methods are word embedding mapping (Conneau et al., 2017; Lample et al., 2018a; Sun et al., 2019), vocabulary sharing (Lample et al., 2018b), and language modeling (Conneau and Lample, 2019). Weight sharing can also be adopted in the encoder/decoder, adversarial training, and back-translation (BT) processes for unsupervised sentence alignment.

BT aims to train models using iteratively generated pseudo-parallel data, thus overcoming the lack of cross-language signals. Specifically, monolingual data in the source language is translated to the target language using a source-to-target translation model, and then the pseudo-parallel data (including both the generated and the original data) is used to train the target-to-source translation model, and vice versa.

Unfortunately, as the input sentences in the pseudo-parallel data are generated by unsupervised models, random errors and noise are inevitably introduced, resulting in low-quality parallel data for model training and bad translation performance. In addition, when vocabulary sharing UNMT models for two distant languages (that is, very little vocabulary overlap between the source language and target language) are trained with BT, the unsupervised model may generate the words in the source language instead of in the target language under source-to-target forward translation. As a result, although the reconstruction loss is small if the forward translation generation is very similar to the input, the model is not sufficiently optimized because the pseudo-parallel corpus contains very little cross-lingual sentence alignment information.

Multilingualism (Edwards, 2002; Clyne, 2017) is a powerful fact of communication across speech communities. In multilingualism, an important “lingua franca” (or common language) often serves as an aid to cross-group understanding, usually representing the language of a potent and prestigious society with a large number of users. For machine translation, the parallel corpora between languages and some lingua franca are usually more abundant. Thus, conventional Pivot Translation (PT) usually leverages a resource-rich language (mainly English) as the pivot to help the low/zero-resource translation (see Appendix A.1 for a detailed analysis). Although UNMT no longer requires parallel corpora, this feature

is still worth exploring and can be used to enhance current UNMT systems under low- or zero-resource scenarios. In addition, we can further use the transfer learning capabilities of the model to transfer the translation capabilities of languages and lingua francas to any two languages that need to be learned.

In this work, taking the merits of pivot language translation in both supervised NMT and UNMT as shown in Figure 1, we propose the reference language-based UNMT framework in which the reference language shares a parallel corpus with only the source language (using only the target language follows a similar pattern). In the framework, we use multilingualism and propose a reference agreement mechanism. Exploiting the accurate alignment clues between source and reference languages, we can more confidently enhance source-target UNMT by taking into account the translation agreement within the source, reference, and target languages. Specifically, this previously irrelevant parallel data plays a role in controlling the quality of the pseudo-sentence pairs through a cross-lingual equivalence (translation agreement). The proposed mechanism is orthogonal to the common multilingual transfer learning methods and different from the general pivot translation method. Empirical results on popular benchmarks and distant languages show that the reference agreement mechanism consistently improves the performance of UNMT systems. In addition, we explore the impact of multilingual information on the basis of our multilingual UNMT baseline and proposed method.

2 UNMT

UNMT is a recently proposed MT paradigm that attempts to achieve the co-growth of MT models in two directions while relying solely on monolingual data and for example, would benefit both English-to-French vs. French-to-English. It is a special kind of dual learning (He et al., 2016; Xia et al., 2017a,b; Li et al., 2019a) in both directions of language pairs. Currently, state-of-the-art UNMT models are based on a sequence-to-sequence encoder-decoder architecture using Transformer (Vaswani et al., 2017), similar to supervised NMT models.

For ease of expression, in the remainder of this paper, we denote the monolingual training data space of the source \mathcal{S} and target \mathcal{T} languages




Figure 2: Illustration of (a) Back-Translation (BT), (b) Reference Agreement Translation (RAT), (c) Reference Agreement Back-Translation (RABT), and (d) Cross-lingual Back-Translation (XBT) in UNMT. The green ellipses represents the name of the corresponding process, the arrow pointing to the ellipses represents the input for loss calculation, while pointing out indicates optimization target.

as ϕ_S and ϕ_T . The parallel training data space between languages S and T is represented as ϕ_{S-T} . The translation direction symmetry of the UNMT model training implies that the translation direction problem $S \rightarrow T$ is the same as $T \rightarrow S^1$.

In general, the NMT model with parameters $\theta_{S \rightarrow T}$ models the conditional probability $\mathbb{P}(\mathbf{t}|\mathbf{s})$ of the translated sequence \mathbf{t} . The model parameters $\theta_{S \rightarrow T}$ are trained to maximize the following likelihood on the parallel training data space:

$$\mathcal{L}(\theta_{S \rightarrow T}) = \mathbb{E}_{(\mathbf{s}, \mathbf{t}) \sim \phi_{S-T}} [-\log \mathbb{P}(\mathbf{t}|\mathbf{s}; \theta_{S \rightarrow T})]. \quad (1)$$

As there is a lack of cross-lingual sentence alignment information, the current UNMT models, despite their differences in training methods and structure, reach a consensus over the use of the parallel data that was iteratively generated by the BT method. Specifically, for a monolingual sentence of target language $\mathbf{t} \in \phi_T$, a source translation $\tilde{\mathbf{s}}$ is generated using the primal $T \rightarrow S$ translation model $\mathbb{P}(\cdot|\mathbf{t}, \theta_{T \rightarrow S})$, then $\tilde{\mathbf{s}}$ and \mathbf{t} form a pseudo-parallel pair $\langle \tilde{\mathbf{s}}, \mathbf{t} \rangle$ for $S \rightarrow T$ model training. Similarly, the generated pseudo-parallel pair $\langle \tilde{\mathbf{t}}, \mathbf{s} \rangle$ for a monolingual sentence \mathbf{s} in the source language is also used for training the $T \rightarrow S$ model.

The likelihood of the reconstructions $\mathbf{t} \rightarrow \tilde{\mathbf{s}} \rightarrow \mathbf{t}$ and $\mathbf{s} \rightarrow \tilde{\mathbf{t}} \rightarrow \mathbf{s}$ for the UNMT model is maximized

¹In UNMT, translation is bidirectional, so ‘‘source’’ and ‘‘target’’ languages only indicate translation direction for using model. Essentially, S and T are symmetrical and exchangeable.

over the BT process according to:

$$\mathcal{L}(\theta_{S \rightarrow T}) = \mathbb{E}_{\tilde{\mathbf{s}} \sim \mathbb{P}(\cdot|\mathbf{t}, \theta_{T \rightarrow S}), \mathbf{t} \sim \phi_T} [-\log \mathbb{P}(\mathbf{t}|\tilde{\mathbf{s}}; \theta_{S \rightarrow T})], \quad (2)$$

$$\mathcal{L}(\theta_{T \rightarrow S}) = \mathbb{E}_{\mathbf{s} \sim \phi_S, \tilde{\mathbf{t}} \sim \mathbb{P}(\cdot|\mathbf{s}, \theta_{S \rightarrow T})} [-\log \mathbb{P}(\mathbf{s}|\tilde{\mathbf{t}}; \theta_{T \rightarrow S})]. \quad (3)$$

Finally, the BT process is optimized by minimizing the following objective function:

$$\mathcal{L}_{\text{BT}}(S, T) = \mathcal{L}(\theta_{S \rightarrow T}) + \mathcal{L}(\theta_{T \rightarrow S}). \quad (4)$$

3 Reference Language based UNMT

In this section, we introduce the reference language-based UNMT framework and present our three kinds of reference agreement utilization approaches: reference agreement translation (RAT), reference agreement back-translation (RABT), and cross-lingual back-translation (XBT). These approaches are illustrated in Figure 2.

3.1 Framework and Reference Agreement

Figure 1(a) demonstrates the traditional pivot translation schema in supervised NMT, subfigure 1(b) shows the multilingual UNMT, and subfigure 1(c) is our proposed reference language-based UNMT framework. When applying pivot translation to UNMT, any language pair in UNMT can be directly trained without any parallel data, which allows translation in both directions due to the nature of UNMT. Thus, the traditional pivot schema ($S \rightarrow P \rightarrow T$) is not necessary

when applying pivot translation to UNMT; using a third language (usually a common language) is a more suitable practice for UNMT. In order to distinguish from the pivot language in traditional pivot translation, we define the language used to enhance the performance of translation $\mathcal{S} \rightarrow \mathcal{T}$ in UNMT as the reference language \mathcal{R} , regardless of whether the translation schema is $\mathcal{S} \rightarrow \mathcal{R} \rightarrow \mathcal{T}$ as the bridge or $\mathcal{S} \rightarrow \mathcal{T}$ directly.

In this paper, the reference agreement refers to the cross-lingual equivalence (i.e., translation agreement) provided by bilingual parallel sentence pairs between the reference language and the source or target language of the translation.

3.2 Reference Agreement Translation

In the absence of supervision signals, the quality of machine translation across languages cannot be effectively evaluated. That is, a suitable cross-lingual quality evaluation function $quality(s, \tilde{t})$ cannot be defined in cases where only the source and target generation are provided. As a result, the quality of synthetic pseudo-parallel pairs $\langle \tilde{s}, \tilde{t} \rangle$ and $\langle \tilde{t}, \tilde{s} \rangle$ in BT cannot be guaranteed, which limits the performance of UNMT.

RAT refers to the simultaneous translation of the parallel sentences of languages \mathcal{S} and \mathcal{R} into the target language \mathcal{T} . The two translations should be in agreement (i.e., the same). Therefore, this agreement in the translations from different sources can be used to collaboratively evaluate the generated quality, and it thus forms a new quality evaluation function $quality(s, r, \tilde{s}, \tilde{r})$.

Based on this premise, we propose a detailed implementation for the RAT approach, enabling reference agreement functions with BT during the UNMT training process and resulting in improved translation agreement, as shown in Figure 2(b). Specifically, RAT requires the two translation models to generate an agreed-upon translation by taking votes. We use this agreed-upon translation as the target and form pseudo-parallel data from the input of each language to train both of the models.

Specifically, for a parallel sentence pair $\langle s, r \rangle$, we would ideally have $\mathbb{P}(\cdot | s; \theta_{\mathcal{S} \rightarrow \mathcal{T}}) = \mathbb{P}(\cdot | r; \theta_{\mathcal{R} \rightarrow \mathcal{T}})$, as stated for RAT; however, as the two models $\theta_{\mathcal{S} \rightarrow \mathcal{T}}$ and $\theta_{\mathcal{R} \rightarrow \mathcal{T}}$ are trained on different data, the agreement may be corrupted. Therefore, we combine the two models to obtain the agreed-upon translation output \tilde{t}_a :

$$\tilde{t}_a \sim \mathbb{P}(\cdot | s, r; \theta_{\mathcal{S} \rightarrow \mathcal{T}}, \theta_{\mathcal{R} \rightarrow \mathcal{T}}), \quad (5)$$

where $\mathbb{P}(\cdot | s, r; \theta_{\mathcal{S} \rightarrow \mathcal{T}}, \theta_{\mathcal{R} \rightarrow \mathcal{T}})$ is

$$\prod_{i=1}^J \left[\frac{1}{2} (\mathbb{P}(\cdot | s, \tilde{t}_{<i}; \theta_{\mathcal{S} \rightarrow \mathcal{T}}) + \mathbb{P}(\cdot | r, \tilde{t}_{<i}; \theta_{\mathcal{R} \rightarrow \mathcal{T}})) \right], \quad (6)$$

where $\tilde{t}_{<i}$ stands for tokens that have been generated prior to the i -generation step. Finally, two synthetic sentence pairs $\langle s, \tilde{t}_a \rangle$ and $\langle r, \tilde{t}_a \rangle$ are used to train the models $\mathcal{S} \rightarrow \mathcal{T}$ and $\mathcal{R} \rightarrow \mathcal{T}$. Since the silver learning target is optimized, the smoothed cross-entropy loss \mathcal{L}_ϵ is used instead of the ordinary cross-entropy loss \mathcal{L} . The learning objective for RAT can be written as:

$$\mathcal{L}_{\text{RAT}}(\mathcal{S}, \mathcal{T}, \mathcal{R}) = \mathcal{L}_\epsilon(\theta_{\mathcal{S} \rightarrow \mathcal{T}}) + \mathcal{L}_\epsilon(\theta_{\mathcal{R} \rightarrow \mathcal{T}}), \quad (7)$$

where ϵ is the smoothing control value indicating the uncertainty of the target for the model.

3.3 Reference Agreement Back-translation

Motivated by the RAT approach, the input language sentences and agreed-upon translations form two synthetic parallel sentences. With these regularized pseudo-parallel sentences, we not only train the $\mathcal{S} \rightarrow \mathcal{T}$ and $\mathcal{R} \rightarrow \mathcal{T}$ forward-translation models (as the generation direction is the same as the training direction), but also train the BT models, i.e., $\mathcal{T} \rightarrow \mathcal{S}$ and $\mathcal{T} \rightarrow \mathcal{R}$. This gives the RABT training approach shown in Figure 2(c). The learning objective of RABT can be described as:

$$\mathcal{L}_{\text{RABT}}(\mathcal{S}, \mathcal{T}, \mathcal{R}) = \mathcal{L}(\theta_{\mathcal{T} \rightarrow \mathcal{S}}) + \mathcal{L}(\theta_{\mathcal{T} \rightarrow \mathcal{R}}). \quad (8)$$

3.4 Cross-lingual Back-translation

The traditional BT analyzed in Section 2 and illustrated in Figure 2(a) allows us to train a $\mathcal{T} \rightarrow \mathcal{S}$ model with the help of an $\mathcal{S} \rightarrow \mathcal{T}$ model, and vice versa; however, this mutually beneficial training is performed entirely within one language pair. Multilingual UNMT (MUNMT) (Sun et al., 2020) is a special case of UNMT that is capable of translating between multiple source and target languages. Although multiple language pairs are trained jointly in MUNMT, there is an obvious shortcoming for BT: translating between language pairs that do not occur together during training, i.e., lack of optimization across language pairs. Joint training across language pairs can be performed through forced high-order BT in UNMT, which takes the form $L_1 \rightarrow L_2 \rightarrow \dots \rightarrow L_{O+1} \rightarrow L_1$, where O is the translation order indicating the number of bridge languages in BT. This approach may fail because decoding through multiple noisy

channels ($L_i \rightarrow L_{i+1}$) accumulates latency and compounds errors, resulting in low-quality final pseudo-parallel data between L_{O+1} and L_1 .

Although this high-order BT can expose multiple language pairs for simultaneous training, it also introduces the problem of uncontrollable intermediate translation quality. Therefore, we propose XBT based on the reference agreement. This method allows BT to remain first order while training across language pairs. XBT is a new training approach for UNMT that translates language S to \mathcal{T} and then back-translates it to \mathcal{R} , or from \mathcal{R} to \mathcal{T} and then to S , based on the reference agreement provided by the bilingual parallel data $\phi_{S-\mathcal{R}}$ between languages S and \mathcal{R} . This training approach is illustrated in Figure 2(d). The objective function of XBT is:

$$\mathcal{L}_{\text{XBT}}(S, \mathcal{T}, \mathcal{R}) = \mathcal{L}(\theta_{\mathcal{T}_S \rightarrow \mathcal{R}}) + \mathcal{L}(\theta_{\mathcal{T}_R \rightarrow S}), \quad (9)$$

where \mathcal{T}_S and \mathcal{T}_R indicate language sentences translated from S and \mathcal{R} , respectively.

4 Experiments and Results

4.1 Datasets

We consider multilingual UNMT for four languages: English (*en*), French (*fr*), Romanian (*ro*), and Chinese (*zh*). To compare the impact of the relationship between the chosen reference language and the considered language pairs on the UNMT performance, we constructed two language scenarios: English–French–Romanian (*en-fr-ro*) and English–Chinese–Romanian (*en-zh-ro*), where English–Romanian (*en-ro*) is the main language pair considered. French and Chinese are used as the reference languages, providing the parallel corpora of English–French (*en-fr*) and English–Chinese (*en-zh*), respectively, to aid the UNMT of English–Romanian. English and Romanian belong to the Indo-European language family, but English belongs to the Germanic branch, whereas Romanian and French belong to the Romance branch. French is selected to evaluate the effect of the reference language being in the homologous family. Chinese belongs to the Sino-Tibetan language family, which is a distant language from Romanian and is selected to study a different language family reference language.

For English, French, and Romanian, we used the same monolingual sentences as those extracted from the WMT News Crawl datasets for the period 2007–2017 by [Conneau and Lample \(2019\)](#) for a

fair comparison and limited the maximum number of sentences in each language to 50 million(M), which results in 50M, 50M, and 14M sentences, respectively. For Chinese, we combined all of the sentences available in the WMT News Crawl datasets with the source sentences from the WMT’17 Chinese–English translation task, leading to 26M sentences. For the parallel data of *en-fr* and *en-zh* introduced by the two experimental settings, we only use those provided by MultiUN ([Ziemski et al., 2016](#)). Finally, the size of the resulting language pair parallel dataset is about 10M.

In both scenarios, we evaluated each language pair except for *en-fr* and *en-zh*, for which the relevant parallel data was used for reference agreement. Following previous studies, newstest 2016 was used to evaluate the *en-ro* language pair. For *fr-ro*, we sampled 5K sentence pairs from OPUS ([Tiedemann, 2012](#)) for evaluation, while for *zh-ro*, we use the religious and educational parallel data for out-of-domain evaluation and collected 2K news parallel sentences for in-domain evaluation. In detail, as data for *fr-ro*, we used GlobalVoices², OpenSubtitles ([Lison and Tiedemann, 2016](#)), and MultiParaCrawl³, whereas for *zh-ro*, Bible-uedin ([Christodouloupoulos and Steedman, 2015](#)), Tanzil, and the QCRI Educational Domain Corpus (QED) ([Abdelali et al., 2014](#)) were used. Because these parallel corpora between *zh-ro* are in religious and educational domains only, which are far away from the news domain of training data, we also collected a parallel corpus (2K in size) of *zh-ro* for in-domain evaluation.

The Moses scripts ([Koehn and Knowles, 2017](#)) were used for tokenization of *en*, *fr*, and *ro*, and the jieba toolkit⁴ was used for word segmentation on *zh*. In particular, following [Sennrich et al. \(2016\)](#), we removed diacritics from *ro*. For *zh*, to avoid confusion between Hong Kong Standard Traditional Chinese (*zh_hk*: QED), Taiwan Standard Traditional Chinese (*zh_tw*: Bible-uedin), and Simplified Chinese (*zh*: Tanzil and monolingual training data), we used `opencc`⁵ to convert *zh_hk* and *zh_tw* to simplified Chinese.

4.2 Baselines

Our baseline models follow XLM ([Conneau and Lample, 2019](#)), with the following refinements:

²<http://casmacat.eu/corpus/global-voices.html>

³<http://paracrawl.eu>

⁴<https://github.com/fxsjy/jieba>

⁵<https://github.com/BYVoid/OpenCC>

	<i>en-fr-ro</i>				<i>en-zh-ro</i>				#
	<i>en→ro</i>	<i>ro→en</i>	<i>fr→ro</i>	<i>ro→fr</i>	<i>en→ro</i>	<i>ro→en</i>	<i>ro→zh</i>	<i>zh→ro</i>	
PBSMT + NMT	25.13	23.90	n/a	n/a	25.13	23.90	n/a	n/a	1
XLM	33.30	31.80	n/a	n/a	33.30	31.80	n/a	n/a	2
MASS	<u>35.20</u>	<u>33.10</u>	n/a	n/a	<u>35.20</u>	<u>33.10</u>	n/a	n/a	3
UNMT	34.45	32.42	25.26	27.99	34.45	32.42	8.66 [2.31]	10.92 [3.56]	4
MUNMT	34.44	32.60	25.31	27.91	33.79	31.82	8.85 [2.63]	11.55 [3.87]	5
+ RAT	35.83	33.52	25.66	28.25	34.59	32.12	9.73 [3.02]	12.44 [3.95]	6
+ RABT	36.05	33.74	25.65	28.44	35.23	32.67	10.09 [3.30]	12.95 [4.00]	7
+ XBT	36.08	33.84	25.78	28.45	34.76	32.30	10.54 [3.32]	13.66 [4.03]	8
+ALL	<u>36.14</u>	<u>34.12</u>	25.60	<u>28.89</u>	<u>35.66</u>	<u>32.88</u>	<u>10.83</u> [3.44]	<u>13.75</u> [4.24]	9
MUNMT + RNMT	36.39	33.85	25.53	28.57	35.50	33.66	10.98 [3.64]	14.42 [4.39]	10
+ RAT	<u>36.65</u>	<u>34.07</u>	25.78	28.63	<u>36.26</u>	<u>34.18</u>	<u>11.26</u> [3.87]	<u>14.77</u> [4.78]	11
+ RABT	36.84	34.32	25.75	29.04	36.78	34.26	11.52 [3.90]	14.79 [5.01]	12
+ XBT	37.13	34.66	26.02	29.11	36.31	34.14	11.80 [4.03]	14.86 [4.98]	13
+ALL	37.27	34.85	26.50	29.45	37.01	34.55	11.92 [4.07]	15.02 [5.11]	14

Table 1: Comparison of the proposed methods with previous work (MultiBLEU). Overall best results are shown in bold (all our best results are better than the corresponding baselines at significance level $p < 0.01$ (Collins et al., 2005)). PBSMT + NMT: (Lample et al., 2018b), XLM: (Conneau and Lample, 2019), MASS: (Song et al., 2019). In the form $x[y]$, x and y respectively indicate results on in-domain and out-of-domain sets. Note, the BLEU used in $ro→zh$ is based on Chinese words segmented by the *jieba* toolkit.

	<i>en</i>	<i>fr</i>	<i>ro</i>	<i>zh</i>
<i>en-ro</i>	6.5 / 64.3	-	4.9 / 68.3	-
<i>fr-ro</i>	-	4.1 / 68.7	4.9 / 68.5	-
<i>zh-ro</i>	-	-	5.3 / 65.8	11.5 / 52.9
<i>en-fr-ro</i>	6.9 / 60.1	4.2 / 68.4	5.0 / 68.1	-
<i>en-zh-ro</i>	7.4 / 53.8	-	5.5 / 64.9	11.4 / 53.4

Table 2: Perplexity / Accuracy for masked language modeling in different languages joint pre-training.

UNMT Lample et al. (2018a); Aharoni et al. (2019); Song et al. (2019) have demonstrated the importance of pre-training, which is a key ingredient of UNMT. Conneau and Lample (2019) used masked language modeling (MLM) to pre-train the full model for the initialization step before applying a denoising autoencoder and BT training step. Therefore, we take the XLM architecture proposed by Conneau and Lample (2019) as our backbone baseline model.

MUNMT Our method studies the impact of adding a reference language to the existing UNMT language pair, which makes our model essentially multilingual. Therefore, MUNMT is the baseline for comparison. We adopt a multi-language joint vocabulary and training with a shared encoder and decoder for language model pre-training, denoising, and BT as the basis of our backbone, UNMT (XLM). Thus, with these settings, the MUNMT model can take advantage of multilingualism.

MUNMT + RNMT Furthermore, as we use a parallel corpus that exists between the reference language and the unsupervised translation language, for a fairer comparison, we consider adding a supervised neural machine translation between the source and reference language (RNMT) as an extra training step on the basis of MUNMT so that supervised and unsupervised training are performed jointly. This baseline is named MUNMT + RNMT.

In all our baselines, the byte pair encoding (BPE) code size is set to 60K, and the model hyperparameters are consistent with those of XLM. The smoothing value ϵ in RAT is set to 0.1.

4.3 Main Results and Analysis

This section examines the effectiveness of the proposed RUNMT framework⁶. The main results⁷ are presented in Table 1. Row #4 reports the replicated results of the XLM architecture (Conneau and Lample, 2019) based on the training of each language pair individually. Our UNMT basically reproduces XLM’s results, and it also

⁶Code available at <https://github.com/bcmi220/runmt>.

⁷Notably, concurrent works (Liu et al., 2020; Bai et al., 2020; Garcia et al., 2020) also explore the case of using auxiliary parallel data effects under the MUNMT setting, where all of these works share similarities in multilingualism motivation. Due to the inconsistency of the parallel corpora used, the results are not directly comparable, so we don’t include their results in the table.

makes some improvements over the original (probably because of differences in data sampling). Thus, our approach offers a strong baseline performance. Compared with the current state-of-the-art method MASS (Song et al., 2019), our baseline performance is slightly lower. This is because MASS adopts the new masked sequence to sequence the pre-training method, and the improvement of our method is orthogonal to the pre-training improvement.

For the MUNMT baseline, as shown in #5, the results are basically consistent with the UNMT results we replicated in #4, with some slight fluctuations, indicating the joint training of language pairs alone cannot make full use of multilingualism. Compared with MUNMT, MUNMT + RNMT (#10) is a very strong method for using an otherwise irrelevant corpus through a reference language.

As shown in Table 2, the performance (perplexity/accuracy) of joint pre-training on all languages is worse than that of pre-training on individual language pairs; however, for distant language pairs, adding a close reference language for joint pre-training will improve performance compared to pre-training on only the distant language pair. Therefore, in #5 and #10, the performance of *en-ro* in *en-fr-ro* and *en-zh-ro* is inconsistent in part due to pre-training. Similarly, comparing the performance of *en-ro* and *zh-ro* in UNMT and MUNMT, the performance of *zh-ro* in MUNMT is better than that in UNMT, indicating that transfer learning plays a role in joint training and the performance in *en-ro* worsens, indicating that joint training a close language pair with a distant language will result in a decline in its UNMT results.

The three specific approaches (RAT, RABT, and XBT) of the proposed RUNMT framework have achieved performance improvements over strong baselines, showing the effectiveness of our proposed approaches. Among them, RAT and RABT both use agreed-upon translations and their inputs to form pseudo-parallel data for training the model: RAT uses the noisy synthetic data as the target, while RABT uses the noisy synthetic data as the source. The results in #6 and #10 show that although RAT with a smoothing mechanism can improve the baselines' performance, the improved result is weaker than RABT in #7 and #12, which use the golden

sentences as the target. Comparing RABT and XBT, the gap in performance is relatively small. XBT has a greater average improvement (#8 and #13), indicating that agreement across language pairs is more effective in MUNMT than agreement within a language pair. In addition, combining the three approaches by optimizing them one by one in an update step, with the results shown in #9 and #14, further improved the performance, indicating that the agreement across language pairs and internal agreement within a language pair are complementary.

In Table 1, we also report the results of different domains within *zh-ro*, where the results in-domain are significantly higher than the results out-of-domain, indicating that the domain problem is also important for UNMT. Our approaches have also obtained consistent improvements over different domains, further verifying the effectiveness of the method.

4.4 Comparison with Pivot Translation

To alleviate the difficulty of lack of bilingual corpora, there are two solutions, the latest uses UNMT in an NMT framework, while the previous solution is pivot translation (usually in an SMT setting), in which the pivot language acts as a bridge creating a path from source to target languages, i.e. $\mathcal{S} \rightarrow \mathcal{P}$ and $\mathcal{P} \rightarrow \mathcal{T}$ across parallel corpora. Our proposed RUNMT is similar to pivot translation, as both seek help from a third language when there is a lack of parallel corpora between the source language and target language. The difference is that our RUNMT requires only one parallel corpus between source and reference languages, while pivot translation requires two: between source and pivot languages, and between pivot and target languages.

In order to make a fairer comparison between the proposed RUNMT framework and the pivot translation framework, we conducted the following experiments in *zh-ro* translation: choosing *en* as the reference language (in RUNMT) or the pivot language (in PT). The two frameworks are evaluated in two settings: one in which only one parallel corpus (*zh - en*) is provided as claimed in RUNMT, and the other in which two parallel corpora (*zh - en* and *en - ro*) are provided as required in PT. Since adding a parallel corpus in our proposed RUNMT framework requires only adding additional training techniques without modifying

	$zh \rightarrow ro$		$ro \rightarrow zh$	
	$zh - en \cdots ro$	$zh - en - ro$	$ro \cdots en - zh$	$ro - en - zh$
PT	14.93	19.78	6.95	13.90
RUNMT	15.02	20.10	11.92	14.56

Table 3: Comparison between RUNMT with traditional PT framework, where “ \rightarrow ” represents the direction of translation, “ $-$ ” represents supervised NMT with parallel corpus, and “ \cdots ” represents UNMT with only monolingual data.

the structure or training a new model, our RUNMT can also conveniently adapt to the setting of two parallel corpora. In order to adapt to the setting where only one parallel corpus is provided, the PT framework adopts the supervised NMT model that trains \mathcal{S} to \mathcal{P} ($zh \rightarrow en$ or $en \rightarrow zh$ in this experiment) and the UNMT model that trains $\mathcal{P}-\mathcal{T}$ ($en - ro$). For the $en - zh$ parallel corpus added in this setting, since MultiUN does not contain this pair, we use the training set provided by WMT’16.

The experimental results show that RUNMT is effective in not only the new case of only one parallel corpus provided, but also the traditional case of two parallel corpora provided, indicating that RUNMT generally makes better use of multilingualism. Additionally, it can be seen from the results that if the first pass of pivot translation is performed by a worse-performing model, error propagation will affect the overall performance, while direct translation in RUNMT will not be affected by this.

5 Ablation

Effects of Parallel Data Scale In order to analyze the influence of the scale of reference and source language parallel data on the performance of MUNMT and our proposed approaches, we compared the performance of $en \rightarrow ro$ on five different parallel corpus sizes: 1K, 10K, 100K, 1M, 10M together with UNMT baseline, and the results are shown in Figure 3.

It shows that although MUNMT + RNMT has been a very strong method for using an otherwise irrelevant corpus through a reference language compared to MUNMT, our proposed RUNMT framework can still improve on various parallel data scales, which verifies the generalization of our method. In addition, in the setting with low parallel data, RABT shows a better growth effect than XBT, and when the parallel resources reach a certain scale, XBT surpasses RABT, indicating




Figure 3: Performances of different parallel data sizes for MUNMT + RNMT $en \rightarrow ro$ with RUNMT.

that agreement training for cross-language pairs requires more parallel data than does agreement within a language pair. Furthermore, the effect of back-translation enhancement in all cases is better than that of forward-translation, which shows that the golden target is better than the silver target in UNMT. Finally, in low-resource settings, our methods have achieved a greater relative improvement, indicating that our methods mine the information of partially relevant parallel data to a greater extent for enhancing UNMT.

Analysis of Intermediate Translation Quality in BT To verify the problem of uncontrollable intermediate quality in the back-translation, we perform experiments on the distant language pair $zh-ro$ and report the results of translation direction $ro \rightarrow zh$. The reason for choosing $zh-ro$ is that Chinese and Romanian characters can be directly distinguished by using unicode encoding. We define BT-BLEU as the BLEU of $s \in \mathcal{S}$ and \tilde{s} generated in the $\mathcal{S} \rightarrow \mathcal{T} \rightarrow \mathcal{S}$ back-translation process, and we introduce this metric in the evaluation phase. We calculate the ratio of the generated Chinese token (subword) to the total number of generated tokens to reflect the intermediate quality of the back-translation from the side. The experimental results are shown in Table 4.

The results show that the growth trend of BLEU is consistent with the downward trend of the ratio of Chinese tokens in the Romanian translations, which has a notable correlation, indicating that this ratio can indeed reflect the training effect of the model to a certain extent.

	BLEU	BT-BLEU	RATIO
UNMT	10.92	31.47	2.17%
MUNMT	11.55	31.98	2.01%
MUNMT + RABT	12.95	32.52	1.69 %
MUNMT + XBT	13.66	33.16	1.62 %

Table 4: Intermediate Translation Quality in BT.

In addition, compared with MUNMT, our methods improve the quality of intermediate translations, bring BT-BLEU improvement, and reduce the proportion of Chinese tokens in Romanian translations, thus verifying the effectiveness of our methods.

6 Related Work

With the development of the deep neural network (He et al., 2018; Cai et al., 2018; Li et al., 2018; Zhou and Zhao, 2019; Zhang et al., 2019b,a; Zhou et al., 2019), UNMT (Artetxe et al., 2017; Lample et al., 2018a,b; Conneau and Lample, 2019; Song et al., 2019) has attracted widespread attention in academic research, as only large-scale monolingual corpora are required for training. The performance of UNMT has benefited from language model pre-training, denoising autoencoders, and BT techniques between similar languages such as English and French, but still lags behind that of supervised NMT for distant languages such as Chinese and English. Conneau and Lample (2019) extended the generative language model pre-training approach to multiple languages and showed that cross-lingual pre-training could be effective for MUNMT. Aside from the convenience of translation among multiple language pairs, including unseen language pairs, transfer learning should be considered when low-resource languages are trained together with rich-resource ones. As discussed by Arivazhagan et al. (2019), MUNMT usually performs worse than pivot-based supervised NMT; however, the pivot-based method easily experiences a computationally expensive quadratic growth in the number of source languages and suffers from the error propagation problem. Arivazhagan et al. (2019) addressed the zero-shot generalization problem that some translation directions have not been optimized well due to a lack of parallel data. Al-Shedivat and Parikh (2019) introduced a consistent agreement-based training method that encourages the model to produce equivalent translations of parallel sentences in

zero-shot translation, which share similarities with our RAT approach. However, in terms of a specific implementation, because of the differences between UNMT and NMT, we have provided three new UNMT methods, and have alleviated the problem of uncontrollable intermediate BT quality in UNMT. Arivazhagan et al. (2019) addressed the issue of transfer learning between language pairs with parallel data where there is a lack of parallel corpora in multilingual supervised NMT. As for the agreement in UNMT, (Sun et al., 2019) investigate the enhancement of unsupervised bilingual word embedding agreement in the UNMT training. Leng et al. (2019) propose a multi-hop UNMT that automatically selects a good translation path for a distant language pair during UNMT. Baijun et al. (2019) proposed a cross-lingual pre-training approach that makes use of the source-pivot data to pre-train the language model.

As for the multilingualism, Liu et al. (2020) proposes a multilingual denoising pre-training technique to improve machine translation tasks. Bai et al. (2020) and Garcia et al. (2020) both studied the agreement across language pairs. Their method is much the same as one of our proposed approaches, XBT, which relies on the supervision signals from a parallel corpus to build a bridge between language pairs in MUNMT. Compared with these two concurrent works, the other two settings of our proposed approaches, RAT and RABT, which use the internal agreement within language pair to improve the translations, can be used not only for MUNMT, but also for semi-supervised NMT to enhance the effect of the only two languages.

7 Conclusion

In this work, we capitalize on the supervised NMT and UNMT use of the pivot language in pivot translation. We propose the reference language-based UNMT framework, in which a reference agreement mechanism is introduced in several implementations to better leverage the reference agreement in parallel data brought by the reference language to reduce the uncontrollable intermediate quality problem in back-translation. The experimental results show that we achieved an improvement over our strong baseline, and our proposed RUNMT framework is compatible with and exceeds the traditional pivot translation framework.

References

- Ahmed Abdelali, Francisco Guzman, Hassan Sajjad, and Stephan Vogel. 2014. [The AMARA corpus: Building parallel language resources for the educational domain](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1856–1862, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Roe Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Maruan Al-Shedivat and Ankur Parikh. 2019. [Consistency by agreement in zero-shot neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1184–1197, Minneapolis, Minnesota. Association for Computational Linguistics.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Roe Aharoni, Melvin Johnson, and Wolfgang Macherey. 2019. The missing ingredient in zero-shot neural machine translation. *arXiv preprint arXiv:1903.07091*.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017. [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)*, pages 451–462, Vancouver, Canada. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*.
- Hongxiao Bai, Mingxuan Wang, Hai Zhao, and Lei Li. 2020. Unsupervised neural machine translation with indirect supervision. *arXiv preprint arXiv:2004.03137*.
- Ji Baijun, Zhang Zhirui, Duan Xiangyu, Zhang Min, Chen Boxing, and Luo Weihua. 2019. Cross-lingual pre-training based transfer for zero-shot neural machine translation. *arXiv preprint arXiv:1912.01214*.
- Jiaxun Cai, Shexia He, Zuchao Li, and Hai Zhao. 2018. [A full end-to-end semantic role labeler, syntactic-agnostic over syntactic-aware?](#) In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2753–2765, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Yong Cheng, Qian Yang, Yang Liu, Maosong Sun, and Wei Xu. 2017. Joint training for pivot-based neural machine translation. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 3974–3980. AAAI Press.
- Christos Christodouloupoulos and Mark Steedman. 2015. A massively parallel corpus: the bible in 100 languages. *Language resources and evaluation*, 49(2):375–395.
- Michael Clyne. 2017. Multilingualism. *The handbook of sociolinguistics*, pages 301–314.
- Michael Collins, Philipp Koehn, and Ivona Kučerová. 2005. [Clause restructuring for statistical machine translation](#). In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, ACL '05, pages 531–540, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems*, pages 7057–7067.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.
- John Edwards. 2002. *Multilingualism*. Routledge.
- Bent Fuglede and Flemming Topsoe. 2004. Jensen-shannon divergence and hilbert space embedding. In *International Symposium on Information Theory, 2004. ISIT 2004. Proceedings.*, page 31. IEEE.
- Xavier Garcia, Pierre Foret, Thibault Sellam, and Ankur P Parikh. 2020. A multilingual view of unsupervised machine translation. *arXiv preprint arXiv:2002.02955*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1243–1252. JMLR. org.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. Dual learning for machine translation. In *Advances in Neural Information Processing Systems*, pages 820–828.
- Shexia He, Zuchao Li, Hai Zhao, and Hongxiao Bai. 2018. [Syntax for semantic role labeling, to be, or not to be](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2061–2071, Melbourne, Australia. Association for Computational Linguistics.
- Yunsu Kim, Petre Petrov, Pavel Petrushkov, Shahram Khadivi, and Hermann Ney. 2019. [Pivot-based](#)

- transfer learning for neural machine translation between non-English languages. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 866–876, Hong Kong, China. Association for Computational Linguistics.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, et al. 2018a. Unsupervised machine translation using monolingual corpora only. In *International Conference on Learning Representations*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018b. [Phrase-based & neural unsupervised machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5039–5049, Brussels, Belgium. Association for Computational Linguistics.
- Yichong Leng, Xu Tan, Tao Qin, Xiang-Yang Li, and Tie-Yan Liu. 2019. [Unsupervised pivot translation for distant languages](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 175–183, Florence, Italy. Association for Computational Linguistics.
- Zuchao Li, Shexia He, Jiayun Cai, Zhuosheng Zhang, Hai Zhao, Gongshen Liu, Linlin Li, and Luo Si. 2018. [A unified syntax-aware framework for semantic role labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2401–2411, Brussels, Belgium. Association for Computational Linguistics.
- Zuchao Li, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, Zhuosheng Zhang, and Hai Zhao. 2019a. Explicit sentence compression for neural machine translation. *arXiv preprint arXiv:1912.11980*.
- Zuchao Li, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, Zhuosheng Zhang, and Hai Zhao. 2019b. Data-dependent gaussian prior objective for language generation. In *International Conference on Learning Representations*.
- Pierre Lison and Jörg Tiedemann. 2016. Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 923–929.
- Chao-Hong Liu, Catarina Cruz Silva, Longyue Wang, and Andy Way. 2018. Pivot machine translation using chinese as pivot language. In *China Workshop on Machine Translation*, pages 74–85. Springer.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *arXiv preprint arXiv:2001.08210*.
- Michael Paul, Hirofumi Yamamoto, Eiichiro Sumita, and Satoshi Nakamura. 2009. [On the importance of pivot language selection for statistical machine translation](#). In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*, pages 221–224, Boulder, Colorado. Association for Computational Linguistics.
- Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. 2018. Unsupervised neural machine translation. In *International Conference on Learning Representations*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Edinburgh neural machine translation systems for WMT 16](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 371–376, Berlin, Germany. Association for Computational Linguistics.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. Mass: Masked sequence to sequence pre-training for language generation. In *International Conference on Machine Learning*, pages 5926–5936.
- 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 *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1235–1245, Florence, Italy. Association for Computational Linguistics.
- Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2020. Knowledge distillation for multilingual unsupervised neural machine translation. *arXiv preprint arXiv:2004.10171*.
- I Sutskever, O Vinyals, and QV Le. 2014. Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems*.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- 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, Rochester, New York. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Hua Wu and Haifeng Wang. 2007. [Pivot language approach for phrase-based statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 856–863, Prague, Czech Republic. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Yingce Xia, Jiang Bian, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2017a. Dual inference for machine learning. In *International Joint Conferences on Artificial Intelligence*, pages 3112–3118.
- Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, and Tie-Yan Liu. 2017b. Dual supervised learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3789–3798. JMLR. org.
- Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2019a. Semantics-aware bert for language understanding. *arXiv preprint arXiv:1909.02209*.
- Zhuosheng Zhang, Hai Zhao, Kangwei Ling, Jiangtong Li, Zuchao Li, Shexia He, and Guohong Fu. 2019b. Effective subword segmentation for text comprehension. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(11):1664–1674.
- Junru Zhou, Zuchao Li, and Hai Zhao. 2019. Parsing all: Syntax and semantics, dependencies and spans. *arXiv preprint arXiv:1908.11522*.
- Junru Zhou and Hai Zhao. 2019. [Head-Driven Phrase Structure Grammar parsing on Penn Treebank](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2396–2408, Florence, Italy. Association for Computational Linguistics.
- Michał Ziemski, Marcin Junczys-Dowmunt, and Bruno Pouliquen. 2016. [The united nations parallel corpus v1.0](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3530–3534, Portorož, Slovenia. European Language Resources Association (ELRA).