

Unifying the Convergences in Multilingual Neural Machine Translation

Yichong Huang Xiaocheng Feng Xinwei Geng Bing Qin

Harbin Institute of Technology, China

{ychuang, xcfeng, xwgeng, qinb}@ir.hit.edu.cn

Abstract

Although all-in-one-model multilingual neural machine translation (MNMT) has achieved remarkable progress, the *convergence inconsistency* in the joint training is ignored, *i.e.*, different language pairs reaching convergence in different epochs. This leads to the trained MNMT model over-fitting low-resource language translations while under-fitting high-resource ones. In this paper, we propose a novel training strategy named LSSD (**L**anguage-**S**pecific **S**elf-**D**istillation), which can alleviate the convergence inconsistency and help MNMT models achieve the best performance on each language pair simultaneously. Specifically, LSSD picks up *language-specific best checkpoints* for each language pair to teach the current model on the fly. Furthermore, we systematically explore three sample-level manipulations of knowledge transferring. Experimental results on three datasets show that LSSD obtains consistent improvements towards all language pairs and achieves the state-of-the-art¹.

1 Introduction

Neural machine translation (NMT) (Kalchbrenner and Blunsom, 2013; Sutskever et al., 2014) has witnessed enormous and significant progress, including network structures (Cho et al., 2014; Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017), attention mechanism (Luong et al., 2015; Liu et al., 2016; Shen et al., 2018) and decoding strategies (Xia et al., 2017; Geng et al., 2018; Zhou et al., 2019). While achieving promising performance (Wu et al., 2016; Hassan et al., 2018), widely-used *bilingual* NMT actually causes huge computational cost especially when tackling numerous language pairs, thereby facilitating the recent emergence of *multilingual* NMT (Ha et al., 2016; Firat et al., 2016; Johnson et al., 2017; Lu et al.,

¹Our code is publicly available at <https://github.com/OrangeInSouth/LSSD>.

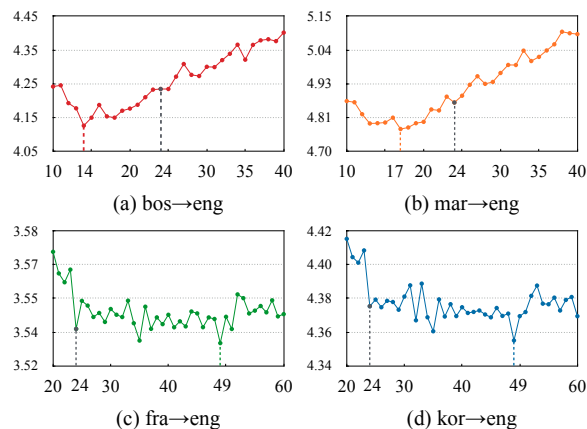


Figure 1: Loss curves of a multilingual NMT model trained to translate 4 languages into English. The upper two boxes illustrate low-resource directions. And the high-resource ones are placed in the bottom part. The x-axis and y-axis indicate training epochs and dev losses respectively. Dotted lines mark language-specific best checkpoints (colorful ones) as well as the overall best checkpoint (the black one).

2018; Aharoni et al., 2019). By directly translating multiple language pairs with one model, multilingual NMT (Tan et al., 2019; Zhang et al., 2020; Fan et al., 2021; Zhang et al., 2021a) quadratically accelerates deployment and effectively encourages transfer learning between similar languages, which greatly benefits low-resource directions (Arivazhagan et al., 2019) and successfully enables zero-shot translation (Gu et al., 2019).

Despite the remarkable success, multilingual NMT (Liu et al., 2020; Lin et al., 2021) clearly expresses a strong disagreement on the uniform convergence point across various translation corpora. The underlying reason is the imbalance and heterogeneity of available data in multilingual training (Wang et al., 2020a; Wu et al., 2021), which also explains why the disagreement between high-resource languages (HRLs) and low-resource languages (LRLs) is more pronounced, as illustrated in Figure 1. Particularly, a fairly common and impor-

tant observation about the learning of multilingual NMT is that HRLs typically encounter underfitting, while extremely severe overfitting generally arises in LRLs.

In this paper, we cast this convergence inconsistency as the performance deficit between the multilingual NMT model and its own language-specific best checkpoints, and aim to reduce this deficit. Towards tackling this problem, we propose a novel training strategy dubbed **L**anguage-**S**pecific **S**elf-**D**istillation (LSSD). At each training step, LSSD appoints recent best checkpoints oriented to each language pair as teacher models and treats the current training model as the student learning from each teacher model in a knowledge distillation manner. Differently depending on data selection and strength of performing knowledge distillation indeed, several practical strategies are employed to potentially provide a relatively fine-grained manipulation of knowledge transferring: 1) LSSD-WHOLE, which performs distillation on all data samples equally; 2) LSSD-SELECTIVE, which selects samples for distillation conditioned on whether the teacher performs better than the student; 3) LSSD-ADAPTIVE, which varies the distillation strength according to the sample-level performance ratio between the teacher and student.

Experimental results on TED talks² and WMT under *many-to-one* and *one-to-many* settings demonstrate that our method can obtain consistent and significant improvement through remedying the convergence inconsistency, ultimately achieving state-of-the-art translation performance.

2 Preliminaries

2.1 Neural Machine Translation

Bilingual neural machine translation (Bilingual NMT) translates a sentence x in source language into a sentence y in target language. Given a parallel corpus $D = \{(x, y) \in \mathcal{X} \times \mathcal{Y}\}$, the neural machine translation model is commonly trained with the Maximum Likelihood Estimation (MLE):

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{(x,y) \sim D} \sum_{i \leq |y|} \log P(y_i | x, y_{<i}; \theta), \quad (1)$$

where $P(\cdot | \cdot; \theta)$ is the conditional probability with model θ , which is usually implemented in an

²<https://www.ted.com/participate/translate>

encoder-decoder architecture (Bahdanau et al., 2015; Vaswani et al., 2017).

2.2 Multilingual Neural Machine Translation

Multilingual neural machine translation (multilingual NMT) translates multiple language pairs with one unified model. In this work, we follow Johnson et al. (2017) to train a multilingual NMT model jointly using training datasets of L language pairs $D^{train} = \{D_1^{train}, \dots, D_l^{train}, \dots, D_L^{train}\}$, where D_l^{train} is the dataset of language pair (S_l, T_l) . To encode and decode diverse languages to/from a shared semantic space, a large multilingual vocabulary \mathcal{V} is constructed. And a language tag is appended to the beginning of source sentences to specify the target language. Similar with bilingual NMT, the MNMT model is also trained with the same objective as Eq.1.

Model Selection Strategy The common practice saves a checkpoint at the end of each training epoch, and evaluates its performance on a set of development sets $D^{dev} = \{D_1^{dev}, \dots, D_l^{dev}, \dots, D_L^{dev}\}$. Finally the checkpoint with minimal average dev loss is selected as the *overall best checkpoint*. This average dev loss could be formalized as:

$$\mathcal{L}^{dev}(\theta, D^{dev}) = \frac{1}{L} \sum_{l=1}^L \mathcal{L}(D_l^{dev}; \theta). \quad (2)$$

In our work, we record language-specific dev losses \mathcal{L}_l^{dev} and save the *language-specific best checkpoint* towards each language pair l additionally.

2.3 Self-Distillation

Knowledge distillation is an effective model compression technology that distills knowledge from a high-capacity teacher model into the compact student model (Hinton et al., 2015). Self-distillation is an intriguing variation on knowledge distillation with the fundamental difference that self-distillation uses the same network for both the teacher and student model. (Yang et al., 2019; Zhang et al., 2019). Yang et al. (2019) use models in earlier epochs to guide the training of later epochs, which boosts the predictive accuracy in image classification by a large margin. In this paper, we extend this idea to let the student model learn multiple different targets in the same training epoch and the number of targets is constantly changing in different epochs.

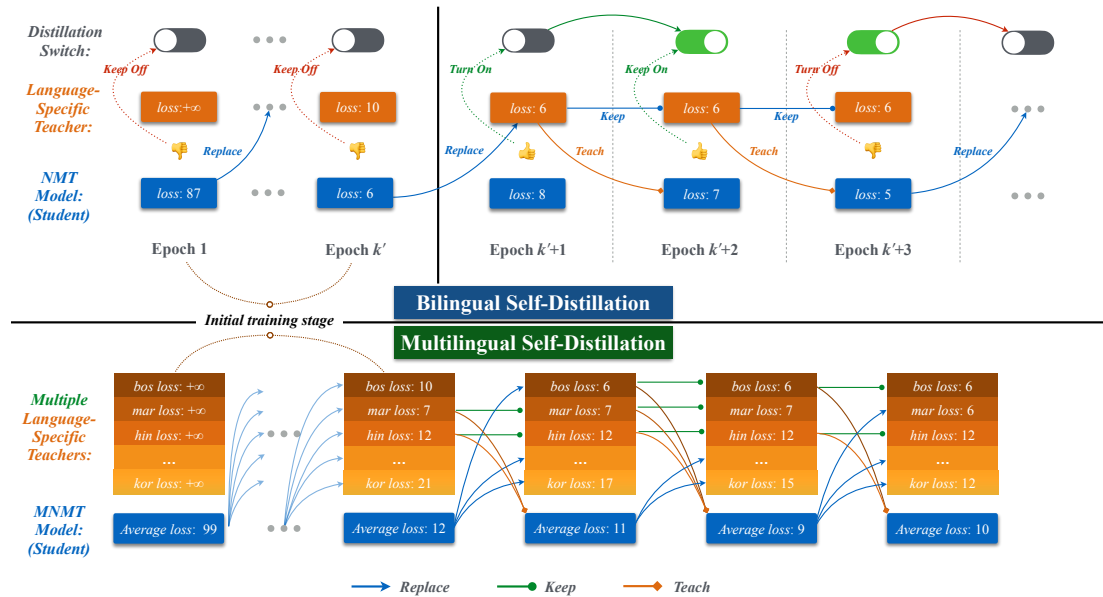


Figure 2: The upper part portrays the process of bilingual LSSD. The multilingual LSSD process is shown in the bottom part. The blue boxes indicates the student model (NMT Model or multilingual NMT model) and the inside value represents its dev loss in the current training epoch. The orange box indicates the teacher model and its loss means its dev loss in the corresponding training epoch. The distillation switch, which determines whether to perform distillation in the current epoch, is drawn as an ellipse with grey for off and green for on.

3 Language-Specific Self-Distillation

In this section, we first introduce the overall process of LSSD. Then, we provide a detailed formalization of LSSD, with special emphasis on the three fine-grained manipulations on knowledge transferring.

3.1 Overall

In this subsection, we present an overview of our training strategy for multilingual NMT, as illustrated in Figure 2. Specifically, we first take bilingual self-distillation as an example to show the distillation learning process. Then, we describe the multilingual self-distillation model and how it performs self-distillation towards multiple language pairs.

Bilingual LSSD Traditional bilingual neural machine translation builds a Sequence-to-Sequence model for training. In our bilingual LSSD, we introduce an additional teacher model and a distillation switch. The teacher model is used to guide the original machine translation model and the switch is designed to decide whether the teacher model is working or not in the current training epoch. We depict the bilingual LSSD process in the upper half of Figure 2.

Normally, at the beginning of training, the loss monotonically decreases, and we refer to this process as “initial training stage”. And k' denotes the

number of epochs the initial training stage lasts. It should be noted that the value of k' is not a hyperparameter but is up to the training process. During these k' epochs, we do not perform distillation (keep the distillation switch off) but replace iteratively the teacher with the better-performance student which has a lower loss. As we all know, the lower loss, the better. Therefore, if the teacher’s loss is lower than the student’s loss, we turn on the switch, as shown in the $k'+1$ epoch. In fact, we do not perform distillation due to the switch being closed at the beginning of this epoch. In the $k'+2$ epoch, since the switch is on, the teacher model distills the student model. And since the teacher’s loss is still lower than the student’s loss, the switch remains on. In the $k'+3$ epoch, the switch also shows turned on, the distillation learning is performed. But the teacher’s loss is higher than the student’s loss, we turn off the switch. At the same time, we replace the teacher model with the current student model to complete the teacher updating.

Multilingual LSSD Multilingual LSSD is a complex version of bilingual LSSD, which needs to maintain multiple language-specific teachers and conduct multi-objective distillation learning. We illustrate the multilingual LSSD process in the bottom half of Figure 2. The blue box also means the student model (the current training MNMT model)

Algorithm 1 Language-Specific Self-Distillation

Input: language pairs number L ; training sets $\{D_i^{train}\}_{i=1}^L$; dev sets $\{D_i^{dev}\}_{i=1}^L$; max training epochs K ; max training steps of one epoch T ; learning rate η ;

Initialize: initialize MNMT model θ ; for $l \in [1, L]$, set $\hat{\theta}_l = \emptyset$, $\Omega_l = off$, $\hat{\mathcal{L}}_l^{dev} = +\infty$.

```
1: for  $k \in [1, K]$  do                                ▷ For each training epoch
2:   for  $t \in [1, T]$  do                                ▷ Training stage
3:     Randomly sample a language pair  $l$ .
4:     Sample a mini-batch of sentence pairs  $B_t^l$ .
5:     if  $\Omega_l$  is on then
6:        $\mathcal{L} = \mathcal{L}_{NMT}(B_t^l; \theta) + \alpha \mathcal{L}_{LSSD}(B_t^l; \theta, \hat{\theta}_l)$ 
7:     else
8:        $\mathcal{L} = \mathcal{L}_{NMT}(B_t^l; \theta)$ 
9:     end if
10:    Update  $\theta$ :  $\theta = \theta - \eta \cdot \nabla_{\theta} \mathcal{L}$ 
11:  end for
12:  for  $l \in [1, L]$  do                                ▷ Validation Stage
13:     $\mathcal{L}_l^{dev} = \mathcal{L}_{NMT}(D_l^{dev}; \theta)$ 
14:    if  $\mathcal{L}_l^{dev} < \hat{\mathcal{L}}_l^{dev}$  then
15:       $\Omega_l = on$ ,  $\hat{\theta}_l = \theta$ ,  $\hat{\mathcal{L}}_l^{dev} = \mathcal{L}_l^{dev}$ 
16:    else
17:       $\Omega_l = on$ 
18:    end if
19:  end for
20: end for
```

and the value of loss represents the average loss of the student over multiple language-specific dev sets in the corresponding training epoch. The different kinds of orange boxes represent different language-specific teacher models, and each loss means different language dev loss in the current training epoch. For example, the darkest orange box `bos` refers to the teacher model in Bosnian-to-English translation.

At the beginning of training, multilingual LSSD also has an “initial training stage” for each language pair just like the bilingual LSSD, at which time the language-specific teacher doesn’t work. Note that the initial training stage of different language pairs may last different epochs. After the initial training stage, the language-specific teacher model begins working, and the student model can be guided by the teacher. The different kinds of orange lines mean the language-specific teacher is distilling the student model and the number of working teachers is determined by the language-specific switches, which are consistent with the distillation switch of bilingual LSSD. The blue line represents the language-specific teacher is replaced by the better-performance student from the last training epoch. In this paper, multilingual LSSD is the superposition of multiple bilingual LSSDs and they do not affect each other. We also summarize this process in Algorithm 1.

3.2 Formalization of LSSD

Formally, we denote the current model by θ and maintain a set of language-specific best checkpoints $\{\hat{\theta}_l\}_{l=1}^L$. In the validation stage, which is at the end of each epoch, we evaluate the performance of θ on each language pair. For each language pair l , if θ outperforms $\hat{\theta}_l$ in dev set D_l^{dev} , θ replaces $\hat{\theta}_l$. To control whether to perform teaching in the current epoch, we define a set of distillation switches $\{\Omega_l\}_{l=1}^L$. In the validation stage, if the current model θ exceeds the language-specific best checkpoint $\hat{\theta}_l$, we turn off the distillation switch, not performing teaching in the next epoch. Conversely, when the language-specific best checkpoint wins, we turn on the switch, performing teaching in the next epoch.

When the language-specific distillation switch is on, the parameters of the current model θ is updated by optimizing both \mathcal{L}_{NMT} and $\alpha \cdot \mathcal{L}_{LSSD}$. When the language-specific distillation switch is off, the model θ is updated only by \mathcal{L}_{NMT} , *i.e.*, $\alpha = 0$. The training loss is calculated as:

$$\mathcal{L} = \mathcal{L}_{NMT} + \alpha \mathcal{L}_{LSSD}, \quad (3)$$

where the α is the weight of distillation loss. And \mathcal{L}_{LSSD} is computed as the cross-entropy between the output distribution of $\hat{\theta}_l$ and θ , which is formalized as:

$$\mathcal{L}_{LSSD} = - \sum_{i \leq |y|} \sum_{w \in \mathcal{V}} P(w|x, y_{<i}; \hat{\theta}_l) \cdot \log P(w|x, y_{<i}; \theta). \quad (4)$$

3.3 Sample-level Manipulations for LSSD

To prevent the potential negative transferring on some translation samples where the teacher underperforms the student, we devise three sample-level manipulations for LSSD: LSSD-WHOLE, LSSD-SELECTIVE and LSSD-ADAPTIVE. All of them could be generalized as rescaling the distillation loss with a sample-level weight:

$$\mathcal{L}_{LSSD} = \mathcal{L}_{LSSD} \times \mathcal{G}, \quad (5)$$

where \mathcal{G} is the sample-level weight which is determined by the performance difference between teacher and student. And different operations correspond different implementation of \mathcal{G} .

Method	DIVERSE		RELATED		WMT	
	M2O	O2M	M2O	O2M	M2O	O2M
<i>Baselines</i>						
MULTILINGUAL	29.00	22.85	27.83	21.85	20.15	19.07
MULTI-DISTILL (Tan et al., 2019)	29.52	22.31	26.60	21.70	20.18	18.57
<i>Previous Works</i>						
MultiDDS-S (Wang et al., 2020a)*	27.00	18.24	25.52	17.32	–	–
MultiUAT (Wu et al., 2021)*	27.83	19.76	26.39	18.64	–	–
CCL-M (Zhang et al., 2021b)*	28.34	19.53	26.73	18.89	–	–
χ -IBR (Zhou et al., 2021)*	29.74	23.44	28.71	22.21	–	–
<i>Our Proposed Approaches</i>						
LSSD-Whole	30.57 [†]	23.55[†]	29.28 [†]	22.20 [†]	21.05 [†]	19.76[†]
LSSD-Selective	30.24 [†]	23.16	28.65 [†]	22.15 [†]	21.17[†]	19.32 [†]
LSSD-Adaptive	30.77[†]	23.39	29.40[†]	22.27[†]	20.96 [†]	19.48 [†]

Table 1: BLEU scores on TED-8-Diverse (DIVERSE) and TED-8-Related (RELATED) and WMT datasets. Bold indicates the highest BLEU value on each setting. “*” represents results taken from original papers. “M2O” means Many-to-One translation. “O2M” means One-to-Many translation. “†” means significantly better than MULTILINGUAL with t-test $p < 0.01$.

LSSD-WHOLE In this manner, we equally execute distillation with the same sample-level weight on all data samples. In practice, it is equivalent to setting $\mathcal{G} = 1$. LSSD-WHOLE can be viewed as the base version of LSSD.

LSSD-SELECTIVE To avert the negative impact of performing distillation on data samples where the teacher model errs on, we explore to only distilling samples where the teacher performs better than the student. And this is implemented as:

$$\mathcal{G}(x, y, \theta, \hat{\theta}_l) = \begin{cases} 0, & P(y|x; \hat{\theta}_l) < P(y|x; \theta) \\ 1, & P(y|x; \hat{\theta}_l) \geq P(y|x; \theta), \end{cases} \quad (6)$$

where $P(y|x; \theta)$ is the likelihood probability that model θ outputs with respect to the target sentence y given the input x . And we compute the sentence probability by averaging probabilities over all tokens.

LSSD-ADAPTIVE Since a one-size-fits-all rule prohibiting distillation on a subset of samples could limit flexibility, we design LSSD-ADAPTIVE changing the distillation weight according to the teacher-student performance ratio, which is implemented as:

$$\mathcal{G}(x, y, \theta, \hat{\theta}_l) = \min\left(\frac{P(y|x; \hat{\theta}_l)}{P(y|x; \theta)}, \sigma\right), \quad (7)$$

where σ is a hyperparameter to truncate the \mathcal{G} when higher than σ , value of which is set to 2 empiri-

cally³.

4 Experiments

4.1 Settings

Datasets We conduct experiments on three datasets: the widely-used TED-8-Diverse and TED-8-Related (Wang et al., 2020a), and a relative large-scale WMT dataset. The TED-8-Diverse contains 4 low-resource languages (bos, mar, hin, mkd) and 4 high-resource languages (ell, bul, fra, kor) to English. The TED-8-Related contains 4 low-resource languages (aze, bel, glg, slk) and 4 related high-resource language (tur, rus, por, ces) to English. Both of these two datasets have around 570K sentence pairs. We detail the data statistics and the interpretation of language codes in Appendix A.

For the WMT dataset, we consider 3 low-resource languages (et, ro, tr) and 3 high-resource languages (fr, de, zh) to English. Totally around 5M training sentences are sampled from the parallel corpus provided by WMT14, WMT16, WMT17, and WMT18. And we use the corresponding dev and test sets for validation and evaluation. The detailed data statistics are also placed in Appendix A. Compared to TED-8-Diverse and TED-8-related, the size of the WMT dataset is larger and distributed more unevenly over various languages.

³We search the optimal σ in $\{1.0, 1, 5, 2.0\}$

Dataset	Setting	Method	bos	mar	hin	mkd	ell	bul	fra	kor	Avg.
DIVERSE	M2O	MULTI-DISTILL LSSD	-0.76 +1.98	-0.39 +1.00	-0.86 +1.30	+0.58 +2.17	+1.65 +2.24	+1.23 +1.81	+1.30 +1.54	+1.36 +2.10	+0.52 +1.77
	O2M	MULTI-DISTILL LSSD	-3.44 +0.93	-0.87 +0.37	-1.26 +0.52	+0.41 +2.37	+0.49 +0.75	+0.35 +0.39	+0.06 +0.04	-0.04 +0.20	-0.54 +0.70
			aze	bel	glg	slk	tur	rus	por	ces	Avg.
RELATED	M2O	MULTI-DISTILL LSSD	-4.31 +1.53	-6.38 +2.35	-4.87 +1.70	+0.09 +1.47	+1.52 +1.18	+1.16 +1.04	+1.62 +1.43	+1.33 +1.82	-1.23 +1.57
	O2M	MULTI-DISTILL LSSD	-0.32 +0.15	-0.64 +0.12	+0.14 +0.50	+0.27 +0.32	-0.20 +0.68	-0.17 +0.20	+0.09 +0.54	-0.31 +0.32	-0.15 +0.35
			tr	ro	et	zh	de	fr	-	-	Avg.
WMT	M2O	MULTI-DISTILL LSSD	-0.84 +0.79	-1.40 +1.3	+0.10 +1.29	+0.78 +0.32	+0.71 +0.5	+0.82 +0.68	-	-	+0.03 +0.81
	O2M	MULTI-DISTILL LSSD	-2.42 +0.81	-2.26 +0.53	-0.69 +0.48	+0.58 +1.00	+0.87 +0.68	+0.92 +0.66	-	-	-0.50 +0.69

Table 2: BLEU score improvements of MULTI-DISTILL and our LSSD over the MULTILINGUAL baseline. For clarity, we take LSSD-ADAPTIVE and LSSD-WHOLE as the representatives of LSSD in M2O and O2M translation respectively. Bold indicates the best performance. Languages are ordered increasingly by data size from left to right.

For each dataset, we experiment in two multilingual translation scenarios: 1) MANY-TO-ONE (M2O): translating multiple languages to English in this work; 2) ONE-TO-MANY (O2M): translating English to different languages.

Hyperparameters We verify the effectiveness of LSSD on the Transformer (Vaswani et al., 2017) as implemented in fairseq (Ott et al., 2019) with 6 layers and 8 attention heads. And we use the same hyperparameters with the previous SOTA (Zhou et al., 2021) to obtain a strong baseline. The only difference with Zhou et al. (2021) is that we train all models for 300 epochs which is less than theirs⁴. For the TED-8-Diverse and TED-8-Related, we follow previous works Wang et al. (2020a); Zhang et al. (2021b) to preprocess both datasets using sentencepiece (Kudo and Richardson, 2018) with a vocabulary size of 8K for each language. For the WMT dataset, we preprocess data using sentencepiece with a vocabulary size of 64K for all languages. The complete set of hyperparameters can be found in Appendix B. All models are trained on 8 Tesla V100 GPUs. And the performance is evaluated with BLEU score using sacreBLEU (Papineni et al., 2002; Post, 2018). We set distillation weight α to 2.0 in M2O and 0.6 in O2M respectively (see section 5.4 for analysis on these choices).

Baselines We compare our LSSD with: 1) the standard-trained multilingual NMT model

⁴We also tried for more epochs, while no performance gains happened

(i.e., MULTILINGUAL) (Johnson et al., 2017); 2) MULTI-DISTILL (Tan et al., 2019), which is also a distillation-based strategy that guides the multilingual model in each translation direction utilizing bilingual models. To re-implement MULTI-DISTILL, we first train bilingual models of each language pair on all three datasets⁵. Follow previous works (Tan et al., 2019; Zhang et al., 2021b), we train bilingual models using the same model configuration and hyper-parameters with multilingual models. For all baselines and our LSSD, the same model configuration and hyper-parameters are applied.

4.2 Main Results

Overall results We summarize main results into Table 1. As we can see, 1) on all three datasets, our LSSD significantly outperforms the baselines under M2O and O2M settings, demonstrating the effectiveness of our approach; 2) compared with previous works, LSSD achieves higher BLEU scores on the TED-8-Diverse and TED-8-Related datasets, which indicates the superiority of our method; 3) in the comparison among the three variants of LSSD, LSSD-ADAPTIVE excels in M2O and LSSD-WHOLE performs best in O2M overall. To better understand this phenomenon, we analyzed the teacher-student performance ratio (detailed in Equation 7) in M2O and O2M respectively, and discovered that this ratio varies more significantly in the challenging O2M translation than in the M2O

⁵we report performance of bilingual models in Appendix C

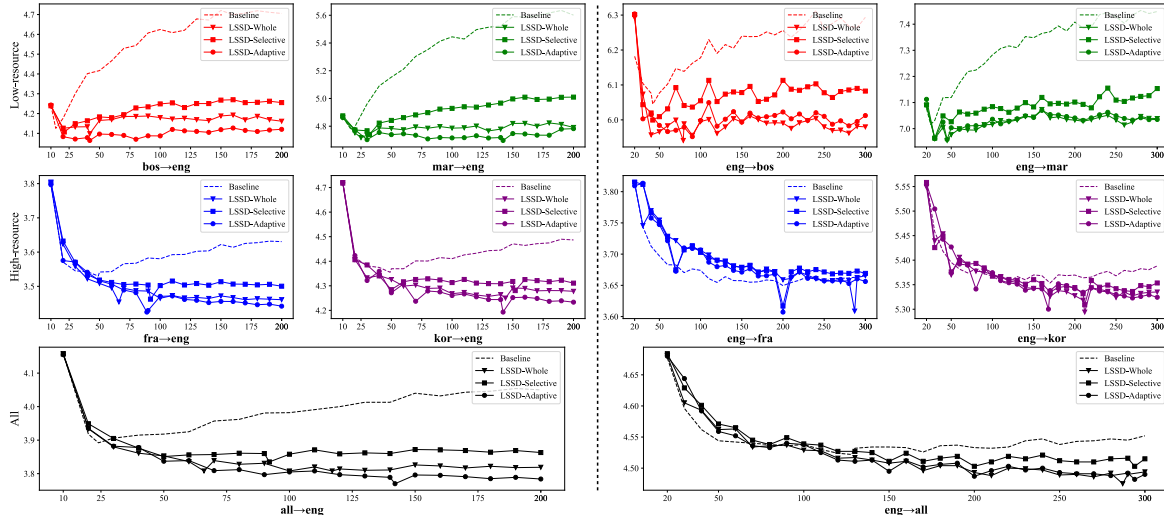


Figure 3: Loss curves of the MULTILINGUAL baseline and LSSD in M2O (left half part) and O2M (right half part) settings on the TED-8-Diverse dataset. The x-axis and y-axis indicate training epochs and dev losses respectively. Due to space limitation, we only display 4 language pairs (2 low-resource + 2 high-resource) in each setting.

Method	DIVERSE		RELATED		WMT	
	M2O	O2M	M2O	O2M	M2O	O2M
MULTILINGUAL	0.27	0.64	0.57	0.51	0.75	1.09
LSSD-WHOLE	0.21	0.15	0.32	0.30	0.25	0.32
LSSD-SELECTIVE	0.24	0.37	0.35	0.29	0.21	0.36
LSSD-ADAPTIVE	0.19	0.35	0.16	0.26	0.14	0.44

Table 3: Performance deficit of the MULTILINGUAL baseline and our LSSD. Bold indicates the lowest value under each setting. Compared to the baseline, LSSD reduces this deficit by 57% on average.

(the variance is 0.097 in O2M and is 0.039 in M2O), which may incur an unstable training for LSSD-Adaptive.

Results on each language Looking closer at results per languages for the MULTI-DISTILL and our LSSD, we calculate the difference between MULTILINGUAL and MULTI-DISTILL or LSSD separately in terms of BLEU, which is shown in Table 2. Firstly, on all datasets and settings, **LSSD consistently outperforms the MULTILINGUAL baseline across all language pairs**. The improvement is up to 2.37 (eng → mkd). Secondly, Multi-Distill struggles in low-resource directions across all three datasets, which is not unexpected considering that the training data is too scarce to train trustworthy bilingual teacher models. However, LSSD breaks this limitation, obtaining stable improvements on **both** low-resource and high-resource language pairs, by employing multilingual NMT models that own more balanced performance as teacher models.

5 Analysis

5.1 Convergence Inconsistency

In this work, we propose to formalize the convergence inconsistency as the loss gap between the overall best checkpoint and language-specific best checkpoints, which is referred to as “**Performance Deficit**”. Concretely, for each language pair l , we accumulate the dev loss differences between the overall best checkpoint θ and the corresponding language-specific best checkpoint $\hat{\theta}_l$. We give a formal definition as:

$$\text{DUB}(\theta, \{\hat{\theta}_l\}_{l=1}^L) = \sum_{l=1}^L (\mathcal{L}(\theta, D_l^{\text{dev}}) - \mathcal{L}(\hat{\theta}_l, D_l^{\text{dev}})). \quad (8)$$

Note that the performance deficit is a non-negative value because $\mathcal{L}(\theta, D_l^{\text{dev}})$ is always greater than or equal to $\mathcal{L}(\hat{\theta}_l, D_l^{\text{dev}})$. We list the performance deficit of the MULTILINGUAL baseline and our LSSD in Table 3. As observed, the performance deficit of LSSD is significantly lower than the baseline (decreased 57% on average), which proves the efficacy of LSSD in remedying convergence inconsistency.

5.2 Multiple Teachers vs. Single Teacher

To provide light on the necessity of using language-specific best checkpoints as teacher models, we conduct ablation study in Table 4. The STSD (Single Teacher Self-Distillation) uses the overall best

Method	DIVERSE		RELATED		WMT	
	M2O	O2M	M2O	O2M	M2O	O2M
LSSD	+1.57	+0.70	+1.45	+0.35	+0.90	+0.69
STSD	+0.89	+0.00	+0.86	-0.14	+0.18	-0.20

Table 4: BLEU score improvements of LSSD (Language-Specific Self-Distillation) and STSD (Single-Teacher Self-Distillation) over the MULTILINGUAL baseline on all three datasets.

checkpoint instead of language-specific best checkpoints to guide the training of multilingual NMT models. To make a fair comparison, we report the results of LSSD-WHOLE. As indicated, the gains from STSD over the baseline only account for 20% to 60% of LSSD’s gains in the M2O translation. Even more, STSD fails to enhance the O2M translation. These results show that the multilingual model gains a great deal from tailored language-specific teachers indeed.

5.3 Comparison of Loss Curves

To better comprehend how each approach affects the model training process, we analyze different methods from the perspective of loss curves, as illustrated in Figure 3. Firstly, comparing the loss curves of the MULTILINGUAL baseline with LSSD (solid lines vs. dotted lines), it is clear that the baseline suffers from serious over-fitting in low-resources (*e.g.*, *bos* \leftrightarrow *eng*). By letting the model *recall* previous checkpoints, LSSD mitigates the over-fitting, which delays the convergence to lengthen the training time for high-resource languages. Secondly, comparing the loss curves of baselines in M2O and O2M (dotted lines in left vs. right half section), it is observed that M2O suffers from more serious over-fitting than O2M. This explains why M2O benefits more from LSSD than O2M. Lastly, contrasting the three modes of LSSD, LSSD-ADAPTIVE performs better in M2O and achieves comparable with LSSD-WHOLE in O2M in terms of dev loss.

5.4 Effect of Distillation Weight α

As Equation 3 shows, LSSD trains multilingual NMT models with NMT loss and the α -weighted distillation loss jointly. We demonstrate the effect of different α on LSSD in Figure 4. As we can see, the optimal weight for O2M ($\alpha = 0.6$) is smaller than which for M2O ($\alpha = 2.0$). We conjecture this is due to the fact that O2M converges later than M2O (see Section 5.3), meaning that O2M

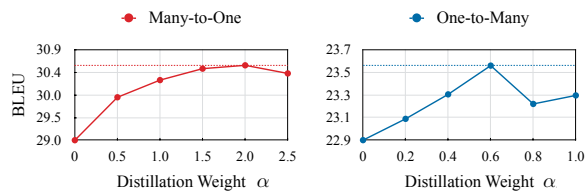


Figure 4: Effect of different distillation weights on LSSD. For clarity, we take LSSD-WHOLE as the representative of LSSD.

models might learn from immature teachers for more epochs. Consequently, a lower distillation strength is more suited to eliminating this danger.

6 Related Works

6.1 Advances in Multilingual NMT

Recently, multilingual NMT mainly focuses on: 1) designing effective parameters sharing strategy (Zhang et al., 2021a; Zhu et al., 2021; Xie et al., 2021; Lin et al., 2021); 2) obtaining language-agnostic representations (Zhu et al., 2020; Pan et al., 2021); 3) incorporating pre-training models (Siddhant et al., 2020; Wang et al., 2020b); 4) resolving the data imbalance among diverse languages (Wang et al., 2020a; Wu et al., 2021; Zhang et al., 2021b; Zhou et al., 2021). Different from them, LSSD is designed for alleviating the convergence inconsistency, which is ignored by existing works.

6.2 Knowledge Distillation in NMT

To the best of our knowledge, Kim and Rush (2016) first apply knowledge distillation in bilingual NMT and propose a sequence-level distillation. Wei et al. (2019) propose to avoid over-fitting by guiding the training process with best checkpoints.

In multilingual NMT, Tan et al. (2019) use knowledge distillation to close the gap between the multilingual NMT model and bilingual models. However, their work is based on the hypothesis that there are sufficient training data for each language pair to prepare a promising bilingual teacher. In fact, most languages face the resource-scarcity problem. In our work, LSSD’s teacher is the multilingual model, which has a better performance for low resource language translation via transfer learning. And we further develop three sample-level operations for LSSD via weighing teacher and student performance.

7 Conclusion

In this work, we propose a novel training strategy **Language-Specific Self-Distillation (LSSD)** to remedy the convergence inconsistency in multilingual neural machine translation. Moreover, we devise three sample-level manipulations for LSSD. Experimental results on three datasets demonstrate that LSSD achieves SOTA performance. Through analysis experiments, we also find that: 1) LSSD significantly mitigates the convergence inconsistency (decreased 57% on average), which is quantified by performance deficit; 2) both low-resource and high-resource languages benefit from LSSD.

8 Limitations

Same with other distillation-based works (Hinton et al., 2015; Tan et al., 2019), LSSD takes some extra training overhead. Taking the training in TED-8-Diverse M2O as an example, the BASELINE and LSSD-WHOLE spend 5.7 hours and 7.6 hours respectively. However, it is worth noting that our method doesn't affect the inference speed of the model.

Acknowledgements

Xiaocheng Feng is the corresponding author of this work. We thank the anonymous reviewers for their insightful comments. This work was supported by the National Key R&D Program of China via grant 2020AAA0106502, National Natural Science Foundation of China (NSFC) via grant 62276078 and the Major Key Project of PCL, PCL2022D01.

References

- Roei 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.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George F. Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. **Massively multilingual neural machine translation in the wild: Findings and challenges**. *CoRR*, abs/1907.05019.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. **Learning phrase representations using RNN encoder–decoder for statistical machine translation**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. **Beyond english-centric multilingual machine translation**. *Journal of Machine Learning Research*, 22(107):1–48.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. **Multi-way, multilingual neural machine translation with a shared attention mechanism**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 866–875, San Diego, California. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, and Yann Dauphin. 2017. **A convolutional encoder model for neural machine translation**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 123–135, Vancouver, Canada. Association for Computational Linguistics.
- Xinwei Geng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2018. **Adaptive multi-pass decoder for neural machine translation**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 523–532, Brussels, Belgium. Association for Computational Linguistics.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor O.K. Li. 2019. **Improved zero-shot neural machine translation via ignoring spurious correlations**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1258–1268, Florence, Italy. Association for Computational Linguistics.
- Thanh-Le Ha, Jan Niehues, and Alex Waibel. 2016. **Toward multilingual neural machine translation with universal encoder and decoder**. In *Proceedings of the 13th International Conference on Spoken Language Translation*, Seattle, Washington D.C. International Workshop on Spoken Language Translation.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, Shujie Liu, Tie-Yan Liu, Renqian Luo, Arul Menezes, Tao Qin, Frank Seide, Xu Tan, Fei Tian, Lijun Wu, Shuangzhi Wu, Yingce Xia, Dongdong Zhang, Zhirui Zhang, and Ming Zhou. 2018.

- Achieving human parity on automatic chinese to english news translation. *CoRR*, abs/1803.05567.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the knowledge in a neural network](#).
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Nal Kalchbrenner and Phil Blunsom. 2013. [Recurrent continuous translation models](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1700–1709, Seattle, Washington, USA. Association for Computational Linguistics.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Zehui Lin, Liwei Wu, Mingxuan Wang, and Lei Li. 2021. [Learning language specific sub-network for multilingual machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 293–305, Online. Association for Computational Linguistics.
- Lemao Liu, Masao Utiyama, Andrew Finch, and Ei-ichiro Sumita. 2016. [Neural machine translation with supervised attention](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3093–3102, Osaka, Japan. The COLING 2016 Organizing Committee.
- 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](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Yichao Lu, Phillip Keung, Faisal Ladhak, Vikas Bhardwaj, Shaonan Zhang, and Jason Sun. 2018. [A neural interlingua for multilingual machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 84–92, Brussels, Belgium. Association for Computational Linguistics.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. [Scaling neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.
- Xiao Pan, Mingxuan Wang, Liwei Wu, and Lei Li. 2021. [Contrastive learning for many-to-many multilingual neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 244–258, Online. Association for Computational Linguistics.
- 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*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, Shirui Pan, and Chengqi Zhang. 2018. [Disan: Directional self-attention network for rnn/cnn-free language understanding](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Aditya Siddhant, Ankur Bapna, Yuan Cao, Orhan Firat, Mia Chen, Sneha Kudugunta, Naveen Arivazhagan, and Yonghui Wu. 2020. [Leveraging monolingual data with self-supervision for multilingual neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2827–2835, Online. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014.

- Dropout: A simple way to prevent neural networks from overfitting. *JMLR*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NeurIPS*.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *CVPR*.
- Xu Tan, Yi Ren, Di He, Tao Qin, and Tie-Yan Liu. 2019. [Multilingual neural machine translation with knowledge distillation](#). In *International Conference on Learning Representations*.
- 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 *NeurIPS*.
- Xinyi Wang, Yulia Tsvetkov, and Graham Neubig. 2020a. [Balancing training for multilingual neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8526–8537, Online. Association for Computational Linguistics.
- Yiren Wang, ChengXiang Zhai, and Hany Hassan. 2020b. [Multi-task learning for multilingual neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1022–1034, Online. Association for Computational Linguistics.
- Hao-Ran Wei, Shujian Huang, Ran Wang, Xin-yu Dai, and Jiajun Chen. 2019. [Online distilling from checkpoints for 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 1932–1941, Minneapolis, Minnesota. Association for Computational Linguistics.
- Minghao Wu, Yitong Li, Meng Zhang, Liangyou Li, Gholamreza Haffari, and Qun Liu. 2021. [Uncertainty-aware balancing for multilingual and multi-domain neural machine translation training](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7291–7305, Online and Punta Cana, Dominican 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, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.
- Yingce Xia, Fei Tian, Lijun Wu, Jianxin Lin, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2017. [Deliberation networks: Sequence generation beyond one-pass decoding](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wanying Xie, Yang Feng, Shuhao Gu, and Dong Yu. 2021. [Importance-based neuron allocation for multilingual neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5725–5737, Online. Association for Computational Linguistics.
- C. Yang, L. Xie, C. Su, and A. L. Yuille. 2019. Snapshot distillation: Teacher-student optimization in one generation. In *CVPR*.
- Biao Zhang, Ankur Bapna, Rico Sennrich, and Orhan Firat. 2021a. [Share or not? learning to schedule language-specific capacity for multilingual translation](#). In *International Conference on Learning Representations*.
- Biao Zhang, Philip Williams, Ivan Titov, and Rico Sennrich. 2020. [Improving massively multilingual neural machine translation and zero-shot translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1628–1639, Online. Association for Computational Linguistics.
- Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. 2019. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *ICCV*.
- Mingliang Zhang, Fandong Meng, Yunhai Tong, and Jie Zhou. 2021b. [Competence-based curriculum learning for multilingual machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2481–2493, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Chunting Zhou, Daniel Levy, Xian Li, Marjan Ghazvininejad, and Graham Neubig. 2021. [Distributionally robust multilingual machine translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5664–5674, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Long Zhou, Jiajun Zhang, and Chengqing Zong. 2019. [Synchronous bidirectional neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 7:91–105.
- Changfeng Zhu, Heng Yu, Shanbo Cheng, and Weihua Luo. 2020. [Language-aware interlingua for multilingual neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1650–1655, Online. Association for Computational Linguistics.

Yaoming Zhu, Jiangtao Feng, Chengqi Zhao, Mingxuan Wang, and Lei Li. 2021. [Counter-interference adapter for multilingual machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2812–2823, Punta Cana, Dominican Republic. Association for Computational Linguistics.

A Data Statistics

Data statistics of the TED-8-Diverse and TED-8-Related are listed in Table 5. Data statistics of the WMT dataset is listed in Table 6.

DIVERSE		RELATED	
language	#num	language	#num
bos (Bosnian)	5,664	bel (Belarusian)	4,509
mar (Marathi)	9,840	aze (Azerbaijani)	5,946
hin (Hindi)	18,798	glg (Glacian)	10,017
mkd (Macedonian)	25,335	slk (Slovak)	61,470
ell (Greek)	134,327	cse (Czech)	103,093
bul (Bulgarian)	174,444	tur (Turkish)	182,470
fra (French)	192,304	por (Portuguese)	184,755
kor (Korean)	205,640	rus (Russian)	208,458

Table 5: Data statistics for the TED-8-Diverse dataset and the TED-8-Related dataset. ‘#num’ refers to the number of sentence pairs in the training set.

Language	Data Source	#num
tr (Turkish)	WMT17	5,000
ro (Romanian)	WMT16	10,000
et (Estonian)	WMT18	80,000
zh (Chinese)	WMT17	400,000
de (German)	WMT14	1,500,000
fr (French)	WMT14	3,000,000

Table 6: Data statistics for the WMT dataset. ‘#num’ refers to the number of sentence pairs in the training set.

B Hyperparameters

In this section, we list the details of hyperparameters we use for the experiments.

- We adopt the architecture with 6 layers and 8 attention heads.
- The embedding dimension is 512 and the FFN has a dimension of 2048.
- We use Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9, \beta_2 = 0.98$, and the same learning rate schedule as Vaswani et al. (2017).
- Batch size is set to 64K and half-precision training is adopted (Ott et al., 2018).
- For regularization, we use the dropout as 0.3 (Srivastava et al., 2014) and the label smoothing as 0.1 (Szegedy et al., 2016).
- For sampling strategy, we use temperature-based sampling (Arivazhagan et al., 2019) and

set $\tau = 1$ on the TED-8-Diverse and TED-8-Related. And we set $\tau = 5$ on the WMT dataset as the it is more imbalanced.

- For inference, we use beam search with beam size 5.

C Bilingual vs. Multilingual

We list the results of bilingual models and the multilingual model in Table 7, 8, 9.

	Method	bos	mar	hin	mkd	ell	bul	fra	kor	Avg.
M2O	BILINGUAL	5.64	3.35	9.90	19.66	37.00	38.85	40.67	19.25	21.79
	MULTILINGUAL	25.78	11.25	24.30	33.38	38.40	39.34	40.43	19.15	29.00
O2M	BILINGUAL	3.89	2.47	8.16	13.98	30.97	34.21	38.53	7.70	17.49
	MULTILINGUAL	17.04	4.96	15.99	25.34	33.27	36.31	40.81	9.08	22.85

Table 7: BLEU scores of bilingual models and the multilingual model on the TED-8-Diverse dataset. Languages are ordered increasingly by data size from left to right. Bold indicates the higher BLEU score. We can find that the MULTILINGUAL model consistently outperforms the BILINGUAL model on each language pair except for fra→eng and kor→eng.

	Method	aze	bel	glg	slk	tur	rus	por	ces	Avg.
M2O	BILINGUAL	2.16	1.8	10.1	23.62	26.51	24.97	44.53	25.77	19.93
	MULTILINGUAL	13.11	19.57	30.14	32.36	27.07	25.85	44.89	29.66	27.83
O2M	BILINGUAL	1.59	1.78	9.25	17.62	14.51	19.59	39.15	17.97	15.18
	MULTILINGUAL	7.40	13.09	25.71	25.66	16.93	20.94	41.68	23.35	21.85

Table 8: BLEU scores of bilingual models and the multilingual model on the TED-8-related dataset. Languages are ordered increasingly by data size from left to right. Bold indicates the higher BLEU score. We can find that the MULTILINGUAL model consistently outperforms the BILINGUAL model on all language pairs.

	Method	Tr	Ro	Et	Zh	De	Fr	Avg.
M2O	BILINGUAL	0.78	5.00	5.99	11.60	27.94	33.57	14.15
	MULTILINGUAL	9.89	23.12	16.92	12.72	26.68	31.57	20.15
O2M	BILINGUAL	0.34	3.34	4.76	19.38	24.11	36.07	14.67
	MULTILINGUAL	8.83	18.42	13.34	20.78	21.11	31.93	19.07

Table 9: BLEU scores of bilingual models and the multilingual model on the WMT dataset. Languages are ordered increasingly by data size from left to right. Bold indicates the better performance. We can see that the BILINGUAL model far outperforms the MULTILINGUAL on the high-source fr↔en and de↔en. However, when the training data is limited, the BILINGUAL model lags far behind the MULTILINGUAL.