

Improving Neural Machine Translation by Bidirectional Training

Liang Ding
The University of Sydney
ldin3097@sydney.edu.au

Di Wu
Peking University
inbath@163.com

Dacheng Tao
JD Explore Academy, JD.com
dacheng.tao@gmail.com

Abstract

We present a simple and effective pretraining strategy – bidirectional training (BiT) for neural machine translation. Specifically, we bidirectionally update the model parameters at the early stage and then tune the model normally. To achieve bidirectional updating, we simply reconstruct the training samples from “src→tgt” to “src+tgt→tgt+src” without any complicated model modifications. Notably, our approach does not increase any parameters or training steps, requiring the parallel data merely. Experimental results show that BiT pushes the SOTA neural machine translation performance across 15 translation tasks on 8 language pairs (data sizes range from 160K to 38M) significantly higher. Encouragingly, our proposed model can complement existing data manipulation strategies, i.e. back translation, data distillation and data diversification. Extensive analyses show that our approach functions as a novel bilingual code-switcher, obtaining better bilingual alignment.

1 Introduction

Recent years have seen a surge of interest in neural machine translation (NMT, Luong et al., 2015; Wu et al., 2016; Gehring et al., 2017; Vaswani et al., 2017) where it benefits from a massive amount of training data. But obtaining such large amounts of parallel data is not-trivial in most machine translation scenarios. For example, there are many low-resource language pairs (e.g. English-to-Tamil), which lack adequate parallel data for training.

Although many approaches about fully exploiting the parallel and monolingual data are proposed, e.g. back translation (Sennrich et al., 2016a), knowledge distillation (Kim and Rush, 2016) and data diversification (Nguyen et al., 2020), the prerequisite of these approaches is to build a well-performed baseline model based on the parallel data. However, Koehn and Knowles (2017); Lampl et al. (2018); Sennrich and Zhang (2019) em-

pirically reveal that NMT runs worse than their statistical or even unsupervised counterparts in low-resource conditions. Here naturally arise a question: *Can we find a strategy to consistently improve NMT performance given the parallel data merely?*

We decide to find a solution from *human learning behavior*. Pavlenko and Jarvis (2002); Dworin (2003); Chen et al. (2015) show that bidirectional language learning helps master bilingualism. In the context of machine translation, both the source→target and target→source language mappings may benefit bilingual modeling, which motivates many recent studies, e.g. dual learning (He et al., 2016) and symmetric training (Cohn et al., 2016; Liang et al., 2007). However, their approaches rely on external resources (e.g. word alignment or monolingual data) or complicated model modifications, which limit the applicability of the method to a broader range of languages and model structures. Accordingly, we turn to propose a simple data manipulation strategy and transfer the bidirectional relationship through *bidirectional training* (§2.2). The core idea is using a bidirectional system as an initialization for a unidirectional system. Specifically, to make the most of the parallel data, we first reconstruct the training samples from “ \vec{B} : source→target” to “ \overleftarrow{B} : source+target→target+source”, where the training data was doubled. Then we update the model parameters with \overleftarrow{B} in the early stage, and tune the model with normal “ \vec{B} source→target” direction.

We validated our approach on several benchmarks across different language families and data sizes, including IWSLT21 En↔De, WMT16 En↔Ro, WMT19 En↔Gu, IWSLT21 En↔Sw, WMT14 En↔De, WMT19 En↔De, WMT17 Zh↔En and WAT17 Ja↔En. Experimental results show that the proposed bidirectional training (BiT) consistently and significantly improves the translation performance over the strong Transformer (Vaswani et al., 2017). Also, we show that

BiT can complement existing data manipulation strategies, i.e. back translation, knowledge distillation and data diversification. Extensive analyses in §3.3 confirm that the performance improvement indeed comes from the better cross-lingual modeling and our method works like a novel code-switching method.

2 Bidirectional Training

2.1 Preliminary

Given a source sentence \mathbf{x} , an NMT model generates each target word \mathbf{y}_t conditioned on previously generated ones $\mathbf{y}_{<t}$. Accordingly, the probability of generating \mathbf{y} is computed as:

$$p(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^T p(\mathbf{y}_t|\mathbf{x}, \mathbf{y}_{<t}; \theta) \quad (1)$$

where T is the length of the target sequence and the parameters θ are trained to maximize the likelihood of a set of training examples according to $\mathcal{L}(\theta) = \arg \max_{\theta} \log p(\mathbf{y}|\mathbf{x}; \theta)$. Typically, we choose Transformer (Vaswani et al., 2017) as its SOTA performance. The training examples can be formally defined as follows:

$$\vec{\mathbf{B}} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N \quad (2)$$

where N is the total number of sentence pairs in the training data. Note that in standard MT training, the \mathbf{x} is fed into the encoder and $\mathbf{y}_{<t}$ into the decoder to finish the conditional estimation for \mathbf{y}_t , thus the utilization of $\vec{\mathbf{B}}$ is directional, i.e. $\mathbf{x}_i \rightarrow \mathbf{y}_i$.

2.2 Pretraining with Bidirectional Data

Motivation The motivation is when human learn foreign languages with translation examples, e.g. \mathbf{x}_i and \mathbf{y}_i . Both directions of this example, i.e. $\mathbf{x}_i \rightarrow \mathbf{y}_i$ and $\mathbf{y}_i \rightarrow \mathbf{x}_i$, may help human easily master the bilingual knowledge. Motivated by this, Levinboim et al. (2015); Liang et al. (2007) propose to modelling the invertibility between bilingual languages. Cohn et al. (2016) introduce extra bidirectional prior regularization to achieve symmetric training from the point view of training objective. He et al. (2018); Zheng et al. (2019); Ding et al. (2020a) enhance the coordination of bidirectional corpus with model level modifications. Different from above methods, we model both directions of a given training example by a simple data manipulation strategy.

Our Approach Many studies have shown that pretraining could transfer the knowledge and data distribution, hence improving the generalization (Hendrycks et al., 2019; Mathis et al., 2021). Here we want to transfer the bidirectional knowledge among the corpus. Specifically, we propose to first pretrain MT models on bidirectional corpus, which can be defined as follows:

$$\overleftrightarrow{\mathbf{B}} = \{(\mathbf{x}_i, \mathbf{y}_i) \cup (\mathbf{y}_i, \mathbf{x}_i)\}_{i=1}^N \quad (3)$$

such that the θ in Equation 1 can be updated by both directions. Then the bidirectional pretraining objective can be formulated as:

$$\begin{aligned} \overleftrightarrow{\mathcal{L}}(\theta) &= \overbrace{\arg \max_{\theta} \log p(\mathbf{y}|\mathbf{x}; \theta)}^{\text{Forward: } \overrightarrow{\mathcal{L}}_{\theta}} \quad (4) \\ &+ \underbrace{\arg \max_{\theta} \log p(\mathbf{x}|\mathbf{y}; \theta)}_{\text{Backward: } \overleftarrow{\mathcal{L}}_{\theta}} \quad (5) \end{aligned}$$

where the forward $\overrightarrow{\mathcal{L}}_{\theta}$ and backward $\overleftarrow{\mathcal{L}}_{\theta}$ are optimized iteratively.

From data perspective, we achieve the bidirectional updating as follows: 1) swapping the source and target sentences of a parallel corpus, and 2) appending the swapped data to the original. Then the training data was doubled to make better and full use of the costly bilingual corpus. The pretraining can acquire general knowledge from bidirectional data, which may help *better* and *faster* learning further tasks. Thus, we early stop BiT at 1/3 of the total training steps (we discuss its reasonability in §3.1). In order to ensure the proper training direction, we further train the pretrained model on required direction $\vec{\mathbf{B}}$ with the rest of 2/3 training steps. Considering the effectiveness of pretraining (Mathis et al., 2021) and clean finetuning (Wu et al., 2019b), we introduce a combined pipeline: $\overleftrightarrow{\mathbf{B}} \rightarrow \vec{\mathbf{B}}$ as our best training strategy. There are many possible ways to implement the general idea of bidirectional pretraining. The aim of this paper is not to explore the whole space but simply to show that one fairly straightforward implementation works well and the idea is reasonable.

3 Experiments

3.1 Setup

Data Main experiments in Table 1 are conducted on five translation datasets: IWSLT21 English \leftrightarrow German (Nguyen et al., 2020),

Data Source	IWSLT14		WMT16		IWSLT21		WMT14		WMT19		Δ
Size	160K		0.6M		2.4M		4.5M		38M		
Direction	En-De	De-En	En-Ro	Ro-En	En-Sw	Sw-En	En-De	De-En	En-De	De-En	Ave.
Transformer	29.2	35.1	33.9	34.1	28.8	48.5	28.6	32.1	39.9	40.1	–
+BiT	29.9 [†]	36.3 [‡]	35.2 [‡]	35.9 [‡]	29.9 [‡]	49.9 [‡]	29.7 [‡]	32.9 [†]	40.5 [†]	41.6 [‡]	+1.1

Table 1: Comparison with previous AT work on several widely-used benchmarks, including IWSLT14 En \leftrightarrow De, WMT16 En \leftrightarrow Ro, IWSLT21 En \leftrightarrow Sw, WMT14 En \leftrightarrow De and WMT19 En \leftrightarrow De. “[‡]/[†]” indicates significant difference ($p < 0.01/0.05$) from corresponding baselines, and this leaves as default symbol in Table 2-6.

WMT16 English \leftrightarrow Romania (Gu et al., 2018), IWSLT21 English \leftrightarrow Swahili¹, WMT14 English \leftrightarrow German (Vaswani et al., 2017) and WMT19 English \leftrightarrow German². The data sizes can be found in Table 1, ranging from 160K to 38M. Two distant language pairs in Table 2 are WMT17 Chinese \leftrightarrow English (Hassan et al., 2018) and WAT17 Japanese \rightarrow English (Morishita et al., 2017), containing 20M and 2M training examples, respectively. The monolingual data used for back translation in Table 3 is randomly sampled from publicly available News Crawl corpus³. We use same valid& test sets with previous works for fair comparison except IWSLT21 English \leftrightarrow Swahili, where we follow Ding et al. (2021d) to sample 5K/ 5K sentences from the training set as valid/ test sets. We preprocess all data via BPE (Sennrich et al., 2016b) with 32K merge operations. We use tokenized BLEU (Papineni et al., 2002) as the evaluation metric for all languages except English \rightarrow Chinese, where we use SacreBLEU⁴ (Post, 2018). The *sign-test* (Collins et al., 2005) is used for statistical significance test.

Model We validated our proposed BiT on Transformer (Vaswani et al., 2017)⁵. All language pairs are trained on Transformer-BIG except IWSLT14 En \leftrightarrow De and WMT16 En \leftrightarrow Ro (trained on Transformer-BASE) because of their extremely small data size. For fair comparison, we set beam size and length penalty as 5 and 1.0 for all language pairs. It is worth noting that our data-level approach neither modifies model structure nor adds extra training loss, thus it’s feasible to deploy on any frameworks, e.g. DynamicConv (Wu et al.,

2019a) and non-autoregressive MT (Gu et al., 2018; Ding et al., 2020b, 2021c), and training orders, e.g. curriculum learning (Liu et al., 2020a; Zhou et al., 2021; Zhan et al., 2021; Ding et al., 2021a). We will explore them in the future works.

Training For Transformer-BIG models, we empirically adopt large batch strategy (Edunov et al., 2018) (i.e. 458K tokens/batch) to optimize the performance. The learning rate warms up to 1×10^{-7} for 10K steps, and then decays for 30K (data volumes range from 2M to 10M) / 50K (data volumes large than 10M) steps with the cosine schedule; For Transformer-BASE models, we empirically adopt 65K tokens per batch for small data sizes, e.g. IWSLT14 En \rightarrow De and WMT16 En \rightarrow Ro. The learning rate warms up to 1×10^{-7} for 4K steps, and then decays for 26K steps. For regularization, we tune the dropout rate from [0.1, 0.2, 0.3] based on validation performance, and apply weight decay with 0.01 and label smoothing with $\epsilon = 0.1$. We use Adam optimizer (Kingma and Ba, 2015) to train the models. We evaluate the performance on an ensemble of last 10 checkpoints to avoid stochasticity.

Someone may doubt that BiT heavily depends on how to properly set the early-stop steps. To dispel the doubt, we investigate whether our approach is robust to different early-stop steps. In preliminary experiments, we tried several simple fixed early-stop steps according to the size of training data (e.g. training 40K En-De and early stop at 10K/ 15K/ 20K, respectively). We found that both strategies achieve similar performances. Thus, we decide to choose a simple and effective method (i.e. 1/3 of the total training steps) for better reproducibility.

3.2 Results

Results on Different Data Scales we experimented on 10 language directions, including IWSLT14 En \leftrightarrow De, WMT16 En \leftrightarrow Ro, IWSLT21

¹<https://iwslt.org/2021/low-resource>

²<http://www.statmt.org/wmt19/translation-task.html>

³<http://data.statmt.org/news-crawl/>

⁴BLEU+case.mixed+lang.en-zh+numrefs.1+smooth.exp+test.wmt17+tok.zh+version.1.5.1

⁵<https://github.com/pytorch/fairseq>

Data Source Size	WMT17		WAT17
	20M		2M
Direction	Zh-En	En-Zh	Ja-En
Transformer	23.7	33.2	28.1
+BiT	24.9 [‡]	33.9 [†]	28.8 [†]

Table 2: Performance on distant language pairs, including WMT17 Zh \leftrightarrow En and WAT17 Ja \rightarrow En. To perform BiT on languages in different alphabets, we share the sub-words dictionaries between languages.

En \leftrightarrow Sw, WMT14 En \leftrightarrow De and WMT19 En \leftrightarrow De. The smallest one merely contains 160K sentences, while the largest direction includes 38M sentence pairs. Table 1 reports the results, we show that BiT achieves significant improvements over strong baseline Transformer in 7 out of 10 directions under the significance test $p < 0.01$, and the rest of 3 directions also show promising performance under the significance test $p < 0.05$, demonstrating the effectiveness and universality of our proposed bidirectional pretraining strategy.

Notably, one advantage of BiT is it saved 1/3 of the training time for the reverse direction. For example, the pretrained BiT checkpoint for En \rightarrow De can be used to tune the reverse direction De \rightarrow En. This advantage shows that BiT could be an efficient training strategy for multiple training direction, e.g. multi-lingual MT tasks (Ha et al., 2016).

Results on Distant Language Pairs Inspired by Ding et al. (2021b), to dispel the doubt that BiT could merely be applied on languages within the same language family, e.g. English and German, we report the results of BiT on Zh \leftrightarrow En and Ja \rightarrow En language pairs, which belong to different language families (i.e. Indo-European, Sino-Tibetan and Japonic).

Table 2 lists the results, as seen, compared with baselines, our method significantly and incrementally improves the translation quality in all cases. In particular, BiT achieves averaged +0.9 BLEU improvement over the baselines, showing the effectiveness and universality of our method across language pairs.

Complementary to Related Work Recent studies start to combine pretraining and traditional data manipulation approaches for better model performance (Conneau and Lample, 2019; Liu et al., 2020b, 2021). To show the complementary between our proposed pretraining method BiT and

Model	BLEU
Transformer-BIG/+BiT	28.6/ 29.7 [‡]
+BT(Caswell et al., 2019)/+BiT	30.5/ 31.2 [†]
+KD(Kim and Rush, 2016)/+BiT	29.3/ 30.1 [†]
+DD(Nguyen et al., 2020)/+BiT	30.1/ 30.7 [†]

Table 3: Complementary to other works. “/+BiT” means combining BiT with corresponding works, and BLEU scores of BiT followed their counterparts with “/”. Experiments are conducted on WMT14 En-De.

Model	BLEU
Transformer-BIG	28.6
+mRASP (Lin et al., 2020)	29.3 [†]
+CSP (Yang et al., 2020)	29.4 [†]
+BiT (Ours)	29.7 [‡]

Table 4: Comparison with previous code-switch approaches on bilingual data, where we follow the best settings of “+mRASP” and “+CSP” as default without extra parameter tuning. For fair comparison, the pre-train/ finetune steps are identical with ours.

related data manipulation works, we list three representative data manipulation approaches for NMT: a) Tagged Back Translation (BT, Caswell et al. 2019) combines the synthetic data generated with target-side *monolingual data* and parallel data; b) Knowledge Distillation (KD, Kim and Rush 2016) trains the model with sequence-level distilled *parallel data*; c) data diversification (DD, Nguyen et al. 2020) diversifies the data by applying KD and BT on *parallel data*. As seen in Table 3, BiT can be applied on existing data manipulation approaches and yield further significant improvements.

3.3 Analysis

We conducted analyses to better understand BiT. Unless otherwise stated, all results are reported on the WMT14 En-De.

BiT works as a simple bilingual code-switcher Lin et al. (2020); Yang et al. (2020) employ the third-party tool to obtain the alignment information to perform code-switching pretraining, where partial of the source tokens is replaced with the aligned target ones. But training such alignment model is time-consuming and the alignment errors may be propagated. Actually, BiT can be viewed as a novel yet simple bilingual code-switcher, where the switch span is the whole sentence and both the source- and target-side sentences are replaced with

Model	AER	P	R
Transformer-BiG	27.1%	71.2%	74.7%
+BiT	24.3%	74.6%	76.9%

Table 5: The AER scores of alignments on En-De.

Model	En→Gu	Gu→En	Ave.Δ
Base	3.0	8.2	-
+BT	2.6	10.1	-
Base+BiT	4.2 [‡]	9.0 [‡]	+1.0
+BT	5.8 [‡]	12.4 [‡]	+2.8

Table 6: Results for En↔Gu on WMT2019 test sets. “Ave. Δ” shows the averaged improvements of “Base+BiT” v.s. “Base” and their corresponding “+BT” comparisons.

the probability 0.5. Take a sentence pair {“Bush held a talk with Sharon”→“布什与沙龙举行了会谈”} in English→Chinese dataset as an example, during pretraining phase, the reconstructed corpus contains {“Bush held a talk with Sharon”→“布什与沙龙举行了会谈”} and its reversed version “布什与沙龙举行了会谈”→“Bush held a talk with Sharon”, simultaneously. For the English→Chinese direction, the reversed sentence pair exactly belongs to the sentence-level switch with a probability of 0.5. For fair comparison, we implement Lin et al., 2020; Yang et al., 2020’s approaches in *bilingual data scenario*. Table 4 show the superiority of BiT, indicating BiT is a good alternative to code-switch in bilingual scenario.

BiT improves alignment quality Our proposed BiT intuitively encourages self-attention to learn bilingual agreement, thus has the potential to induce better attention matrices. We explore this hypothesis on the widely-used Gold Alignment dataset⁶ and follow Tang et al. (2019) to perform the alignment. The only difference being that we average the attention matrices across all heads from the penultimate layer (Garg et al., 2019). The alignment error rate (AER, Och and Ney 2003), precision (P) and recall (R) are evaluation metrics. Table 5 summarizes that BiT allows Transformer to learn better attention matrices, thereby improving alignment performance (24.3 vs. 27.1).

⁶<http://www-i6.informatik.rwth-aachen.de/goldAlignment>, the original dataset is German-English, we reverse it to English-German.

BiT works for extremely low-resource settings

Researches may doubt BiT may fail on extremely low-resource settings where *back-translation even does not work*. To dispel this concern, we conduct experiments on WMT19 English↔Gujurati⁷ in Table 6. Specifically, we follow Li et al. (2019) to collect and preprocess the parallel data to build the base model “Base” and our “Base+BiT” model. For a fair comparison, we sample the monolingual English and Gujurati sentences to ensure Parallel: Monolingual = 1:1 to generate the synthetic data. As seen, when directly applying back-translation (BT) on the En↔Gu Base model, there indeed shows a slight performance drop (-0.4 BLEU). However, our “BiT” significantly improves the initial Base model by averaged +1.0 BLEU, and making the BiT-equipped BT more effective compared to vanilla BT (+2.8 BLEU). These findings on extremely low-resource settings demonstrate that 1) our BiT consistently works well; and 2) BiT provides a better initial model, thus rejuvenating the effects of back-translation.

4 Conclusion and Future Works

In this study, we propose a pretraining strategy for NMT with parallel data merely. Experiments show that our approach significantly improves translation performance, and can complement existing data manipulation strategies. Extensive analyses reveal that our method can be viewed as a simple yet better bilingual code-switching approach, and improves bilingual alignment quality.

Encouragingly, with BiT, our system (Ding et al., 2021d) got the first place in terms of BLEU scores in IWSLT2021⁸ low-resource track. It will be interesting to integrate BiT into our previous systems (Ding and Tao, 2019; Wang et al., 2020) and validate its effectiveness on industrial level competitions, e.g. WMT⁹. It is also worthwhile to explore the effectiveness of our proposed bidirectional pretraining strategy on multilingual NMT task (Ha et al., 2016).

Acknowledgments

We would thank the anonymous reviewers and the area chair for their considerate proofreading and valuable comments.

⁷<http://www.statmt.org/wmt19/translation-task.html>

⁸<https://iwslt.org/2021/>

⁹<http://www.statmt.org/wmt21/>

References

- Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. Tagged back-translation. In *WMT*.
- Chao-Ying Chen, John Xuexin Zhang, Li Li, and Ruiming Wang. 2015. Bilingual memory representations in less fluent chinese–english bilinguals: An event-related potential study. *Psychological reports*.
- Trevor Cohn, Cong Duy Vu Hoang, Ekaterina Vymolova, Kaisheng Yao, Chris Dyer, and Gholamreza Haffari. 2016. Incorporating structural alignment biases into an attentional neural translation model. In *NAACL*.
- Michael Collins, Philipp Koehn, and Ivona Kučerová. 2005. Clause restructuring for statistical machine translation. In *ACL*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *NeurIPS*.
- Liang Ding and Dacheng Tao. 2019. The university of sydney’s machine translation system for wmt19. In *WMT*.
- Liang Ding, Longyue Wang, Xuebo Liu, Derek F. Wong, Dacheng Tao, and Zhaopeng Tu. 2021a. Progressive multi-granularity training for non-autoregressive translation. In *findings of ACL*.
- Liang Ding, Longyue Wang, Xuebo Liu, Derek F. Wong, Dacheng Tao, and Zhaopeng Tu. 2021b. Rejuvenating low-frequency words: Making the most of parallel data in non-autoregressive translation. In *ACL*.
- Liang Ding, Longyue Wang, Xuebo Liu, Derek F. Wong, Dacheng Tao, and Zhaopeng Tu. 2021c. Understanding and improving lexical choice in non-autoregressive translation. In *ICLR*.
- Liang Ding, Longyue Wang, and Dacheng Tao. 2020a. Self-attention with cross-lingual position representation. In *ACL*.
- Liang Ding, Longyue Wang, Di Wu, Dacheng Tao, and Zhaopeng Tu. 2020b. Context-aware cross-attention for non-autoregressive translation. In *COLING*.
- Liang Ding, Di Wu, and Dacheng Tao. 2021d. The usyd-jd speech translation system for iwslt2021. In *IWSLT*.
- Joel E Dworin. 2003. Insights into biliteracy development: Toward a bidirectional theory of bilingual pedagogy. *Journal of Hispanic Higher Education*, 2(2):171–186.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. In *EMNLP*.
- Sarthak Garg, Stephan Peitz, Udhyakumar Nallasamy, and Matthias Paulik. 2019. Jointly learning to align and translate with transformer models. In *EMNLP*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *ICML*.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2018. Non-autoregressive neural machine translation. In *ICLR*.
- Thanh-Le Ha, Jan Niehues, and Alex Waibel. 2016. Toward multilingual neural machine translation with universal encoder and decoder. In *IWSLT*.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, et al. 2018. Achieving human parity on automatic chinese to english news translation. In *arXiv*.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. Dual learning for machine translation. *NeurIPS*.
- Tianyu He, Xu Tan, Yingce Xia, Di He, Tao Qin, Zhibo Chen, and Tie-Yan Liu. 2018. Layer-wise coordination between encoder and decoder for neural machine translation. In *NeurIPS*.
- Dan Hendrycks, Kimin Lee, and Mantas Mazeika. 2019. Using pre-training can improve model robustness and uncertainty. In *ICML*.
- Yoon Kim and Alexander M. Rush. 2016. Sequence-level knowledge distillation. In *EMNLP*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *WNMT*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. Phrase-based & neural unsupervised machine translation. In *EMNLP*.
- Tomer Levinboim, Ashish Vaswani, and David Chiang. 2015. Model invertibility regularization: Sequence alignment with or without parallel data. In *NAACL*.
- Bei Li, Yinqiao Li, Chen Xu, Ye Lin, Jiqiang Liu, Hui Liu, Ziyang Wang, Yuhao Zhang, Nuo Xu, Zeyang Wang, et al. 2019. The niutrans machine translation systems for wmt19. In *WMT*.
- P. Liang, D. Klein, and Michael I. Jordan. 2007. Agreement-based learning. In *NeurIPS*.
- Zehui Lin, Xiao Pan, Mingxuan Wang, Xipeng Qiu, Jiangtao Feng, Hao Zhou, and Lei Li. 2020. Pre-training multilingual neural machine translation by leveraging alignment information. In *EMNLP*.
- Xuebo Liu, Houtim Lai, Derek F. Wong, and Lidia S. Chao. 2020a. Norm-based curriculum learning for neural machine translation. In *ACL*.

- Xuebo Liu, Longyue Wang, Derek F. Wong, Liang Ding, Lidia S. Chao, and Zhaopeng Tu. 2021. On the complementarity between pre-training and back-translation for neural machine translation. In *findings of EMNLP*.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020b. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *EMNLP*.
- Alexander Mathis, Thomas Biasi, Steffen Schneider, Mert Yuksekogul, Byron Rogers, Matthias Bethge, and Mackenzie W Mathis. 2021. Pretraining boosts out-of-domain robustness for pose estimation. In *WACV*.
- Makoto Morishita, Jun Suzuki, and Masaaki Nagata. 2017. Ntt neural machine translation systems at wat 2017. In *IJCNLP*.
- Xuan-Phi Nguyen, Shafiq Joty, Wu Kui, and Ai Ti Aw. 2020. Data diversification: A simple strategy for neural machine translation. In *NeurIPS*.
- Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational linguistics*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*.
- Aneta Pavlenko and Scott Jarvis. 2002. Bidirectional transfer. *Applied linguistics*.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *WMT*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Improving neural machine translation models with monolingual data. In *ACL*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. Neural machine translation of rare words with subword units. In *ACL*.
- Rico Sennrich and Biao Zhang. 2019. Revisiting low-resource neural machine translation: A case study. In *ACL*.
- Gongbo Tang, Rico Sennrich, and Joakim Nivre. 2019. Understanding neural machine translation by simplification: The case of encoder-free models. In *RANLP*.
- 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*.
- Longyue Wang, Zhaopeng Tu, Xing Wang, Li Ding, Liang Ding, and Shuming Shi. 2020. Tencent ai lab machine translation systems for wmt20 chat translation task. In *WMT*.
- Felix Wu, Angela Fan, Alexei Baevski, Yann N Dauphin, and Michael Auli. 2019a. Pay less attention with lightweight and dynamic convolutions. In *ICLR*.
- Lijun Wu, Yiren Wang, Yingce Xia, QIN Tao, Jianhuang Lai, and Tie-Yan Liu. 2019b. Exploiting monolingual data at scale for neural machine translation. In *EMNLP*.
- 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*.
- Zhen Yang, Bojie Hu, Ambyera Han, Shen Huang, and Qi Ju. 2020. Csp: Code-switching pre-training for neural machine translation. In *EMNLP*.
- Runzhe Zhan, Xuebo Liu, Derek F Wong, and Lidia S Chao. 2021. Meta-curriculum learning for domain adaptation in neural machine translation. In *AAAI*.
- Zaixiang Zheng, Hao Zhou, Shujian Huang, Lei Li, Xin-Yu Dai, and Jiajun Chen. 2019. Mirror-generative neural machine translation. In *ICLR*.
- Lei Zhou, Liang Ding, Kevin Duh, Ryohei Sasano, and Koichi Takeda. 2021. Self-guided curriculum learning for neural machine translation. In *IWSLT*.