

The Volctrans Machine Translation System for WMT20

Liwei Wu¹, Xiao Pan¹, Zehui Lin^{2*}, Yaoming Zhu^{3†}, Mingxuan Wang¹, Lei Li¹

¹ByteDance AI Lab, Beijing, China

{wuliwei.000, panxiao.94, wangmingxuan.89, lilei.lab}@bytedance.com

²Fudan University, Shanghai China

linzh18@fudan.edu.cn

³Shanghai Jiao Tong University, Shanghai China

ymZhu@apex.sjtu.edu.cn

Abstract

This paper describes our VolcTrans system on WMT20 shared news translation task. We participated in 14 translation directions. Our basic systems are based on Transformer (Vaswani et al., 2017), with several variants (wider or deeper Transformers, dynamic convolutions). The final system includes text pre-process, data selection, synthetic data generation, advanced model ensemble, and multilingual pre-training.

1 Introduction

We participated in the WMT2020 shared news translation task in 14 directions: English↔Chinese, English↔German, French↔German, English↔Polish, English↔Tamil, English↔Pashto, English↔Khmer, covering language pairs from high to low resources. In this year’s translation task, we mainly focus on exploiting self-supervised and unsupervised methods for NMT to make full use of the monolingual data (Lin et al., 2020; Yang et al., 2019).

We aim at building a general training framework which can be well applied to different translation directions. Our models are mainly based on the Transformer (Vaswani et al., 2017). Techniques used in the submitted systems include iterative back-translation, knowledge distillation. We also employed several tricks to improve in-domain BLEU scores, typically in-domain transfer learning. We also experimented with a multilingual pre-training technique which we proposed recently (Lin et al., 2020).

2 Baseline Models

We apply two different NMT skeletons for the shared news translation as our baseline systems.

We use the implementations in Fairseq (Ott et al., 2019). All models are trained with Adam optimizer (Kingma and Ba, 2014). We use the “inverse sqrt lr” scheduler with 4000 warm-up steps and set the max learning rate to $5e-4$. The betas are (0.9, 0.98). During training, the batches are made of similar length sequences, so we avoid extreme cases where most sequences in the batch are short and we are required to add lots of pad tokens to each of them because one sequence of the same batch is very long. We limit the batch size to 8192 tokens per GPU, to avoid running out of GPU memory. Meanwhile, to achieve a larger batch size to improve the performance (Ott et al., 2018), we set the parameter “update frequency” to 8, and train the model on 8 GPUs, resulting in an actual batch token size = $8192 \times 8 \times 8$. During training, we employ label smoothing of 0.1 and set dropout rate (Hinton et al., 2012) to 0.2.

2.1 Transformer

Following Sun et al. (2019); Wang et al. (2018), we use different architectures for Transformer (Vaswani et al., 2017) to increase the model diversity and potentially get a better ensemble model.

- Transformer 15e6d: According to Sun et al. (2019), a transformer with larger encoder layer number can learn better representation of source sentence and get better BLEU scores. We increase the number of encoder layers from 6 to 15 layers in the transformer big architecture which is the same as the Deeper Transformer in Sun et al. (2019).
- Transformer Mid 25e6d and Transformer Mid 50e6d: To get much better BLEU scores, we further increase the encoder layer number from 6 to 25 (Transformer Mid 25e6d) and 50 (Transformer Mid 50e6d) for the transformer

* Intern at ByteDance

† Intern at ByteDance

big architecture. However, the model is too large and can not be trained with GPU, so we decrease the feed forward size from 4096 to 3072 and the embedding size from 1024 to 768.

- Transformer 15000ffn. According to Sun et al. (2019), the performance of the Transformer model is largely dependent on the dimensions of feed forward network. We use the same architecture as Bigger Transformer in Sun et al. (2019) which increases the feed forward size from 4096 to 15000, the attention dropout from 0.1 to 0.3 and the relu dropout from 0.1 to 0.3. The number of encoder and decoder layers remains 6.
- Transformer 128hdim and Transformer 256hdim. Bhojanapalli et al. (2020) shows that a transformer model with larger attention dimensions can also get better BLEU score. We increase the head dimension from 64 to 128 (Transformer 128hdim) and 256 (Transformer 256hdim). The number of encoder and decoder layers remains 6.
- DLCL 25layers. Li et al. (2019) proposes a transformer variant call DLCL and shows that this architecture can make deep transformer get higher BLEU.

2.2 Dynamic Convolution

We also apply dynamic convolution (Wu et al., 2019) architectures.

- Dynamic Convolution 7e6d: The dynamic convolution model with 7 encoder layers and 6 decoder layers which is the same architecture proposed in Wu et al. (2019).
- Dynamic Convolution 25e6d: We increase the encoder layer number from 6 to 25. For layers above 7, we set the kernel size to 31.

3 Experiment Techniques

3.1 Parallel Data Up-sampling

According to the experiments, data diversity matters for the whole system. Apart from splitting the monolingual data into several disjoint parts, we sampled the parallel data so that each model has different deviations on the parallel data. We tested bagging sampling (sample with replacement) and up-sampling (sample with replacement under

the premise of using all data), experimental results show that when the amount of parallel data is inadequate with respect to the amount of model parameters (such as French \leftrightarrow German, English \leftrightarrow Polish, etc.), the bagging sampling method reduces the performance of the model; while when the amount of parallel data is abundant (such as English \leftrightarrow German), the bagging sampling method has no significant effects on the performance. On the contrary, the data up-sampling method never degrades the performance of the model.

mRASP: Multilingual Pre-training We employed a pre-training method mRASP, which pre-trains a universal multilingual neural machine translation model and fine-tune it on specific language directions. Basically, we pre-train a model using the provided parallel data on WMT2020 of English \leftrightarrow Khmer, English \leftrightarrow Inuktitut, French \leftrightarrow German, English \leftrightarrow Polish, English \leftrightarrow Pashto, English \leftrightarrow Tamil, on a shared vocabulary learned from the above parallel data plus provided monolingual data of all related languages. We learn a BPE sub-word vocabulary with 6000 merge operations. We up-sample the data from lower resource language data to balance data amount and only keep tokens that occur more than 10 times. Finally, we obtain a joint vocabulary of about 28000 tokens.

We fine-tuned the pre-trained model, for low-resource directions: Pashto \rightarrow English, English \leftrightarrow Khmer and English \leftrightarrow Tamil. The baseline model initialized by this method performs better than the randomly initialized baseline model by a large margin. We pre-trained three mRASP models using the same training data: Transformer big, Transformer 15000ffn and Dynamic Convolution. We report in Table 1 the best score in each setting and direction, and find that mRASP significantly outperforms the baseline.

3.2 Tag Back-Translation

Recently, back-translation (Edunov et al., 2018) is a standard method to improve the translation quality by leveraging the large scale monolingual data. Starting from WMT19, the source of the test set is the natural text and the of the test set is the translationese text. We find the tag back-translation (Caswell et al., 2019) method can achieve better BLEU compared with previous methods proposed in Edunov et al. (2018).

Testset	Ps→En	En→Km	Km→En	En→Ta	Ta→En
Random	10.2	39.3	12.7	7.4	14.0
w/ mRASP	13.8	42.8	14.4	9.2	17.9

Table 1: Comparison between randomly initialized baseline model and model initialized from mRASP model




Figure 1: Data Diversity Matters for Final System

To improve the data diversity among single models before model ensemble, we generated the back-translated data from different monolingual data using different baseline models, as illustrated in Figure 1. For high resource data (English, Polish, etc.), we divided monolingual data into several parts, each containing 10M sentences. However, for low resource data (Pashto), due to the lack of monolingual data, we use all monolingual data for all back-translation tasks.

3.3 Iterative Joint Training

Zhang et al. (2018) proposed an iterative joint training method for better usage of monolingual data from source side and target side. In each iteration, the S2T(source to target) model generates a S2T(target to source) synthetic data from the source side monolingual data and the T2S model generates a T2S synthetic data from the target side monolingual data. Then, the S2T and T2S model are trained with the new T2S and S2T synthetic data to improve the both models performance. In the next iteration, the S2T and T2S model can generate synthetic data with better quality and their performance can be improved further. We jointly trained the S2T and T2S model until they converge. Experiment results on English↔Polish shown in Table 2

3.4 Knowledge Distillation

Recently, knowledge distillation has been widely used to improve the performance of models (Sun et al., 2019; Li et al., 2019). In our knowledge distillation method, student model is trained to fit the output of teacher models. Concretely, we translate

Direction	En→Pl	Pl→En
Testset	news20 dev	news20 dev
Baseline	24.8	29.7
Iter 1	27.5	32.6
Iter 2	27.8	32.7
Iter 3	28.2	33.3

Table 2: Iterative Joint Training for English↔Polish

the source side monolingual data with an ensemble teacher and a right-to-left(R2L) (Liu et al., 2016) model teacher.

- **Ensemble Model.** We divided single models in the last joint training iteration into k groups (k=3 in our experiments, resulting in 3 models in each group) and ensemble models in one group to as the teacher model.
- **R2L Model.** We trained one R2L model for each ensemble group using the same data as anyone model in this group from the last iteration.

We then use pseudo parallel data from ensemble model as well as from R2L model to train the student model, without employing parallel data.

3.5 Advanced Tricks

Top-k Checkpoint Average Different from the conventional checkpoint average approach, which is to average continuous K checkpoints, we average K checkpoints which have the highest BLEU scores on the valid set, and find that this strategy usually leads to significant BLEU improvements over single checkpoints.

Random Ensemble We adopt a simple yet effective strategy in model ensemble. Rather than select the best checkpoint from each model (a.k.a. greedy search), we enlarge the search space: choose one checkpoint from top-k checkpoints from each model, and randomly select N combinations from the entire search space, see Figure 2.




Figure 2: Illustration of Random Ensemble

In domain Fine-tuning There exists a domain mismatch between the obtained system trained with provided parallel or monolingual training data and the target test set. In order to alleviate this mismatch, and to improve the translation performance in the domain of the target test set, we fine-tune the best single models with development sets for 1-2 epochs.

4 Settings and Results

For all news tasks, all ParaCrawl corpus is cleaned by the script proposed by Xia et al. (2019). We trained the baseline using the sampling method described in Section 3.1 with different architectures showed in Section 2. For the low resource language pair (English \leftrightarrow Pashto), as showed in Section 3.1, we pretrained three multilingual models with different model architectures (DLCL 25layers, Transformer 15000ffn and Dynamic Convolution 25e6d) on all parallel data available in WMT20 except the English \leftrightarrow Chinese to avoid a large dictionary¹ and fine-tuned the pre-trained models on the their own parallel data with different data sampling strategies to get 9 baseline models². Then we applied the tag back-translation, joint training, knowledge distillation and random ensemble methods as described in Section 3 to get the final translation system. All BLEU scores were reported with SacreBLEU(Post, 2018).

4.1 Chinese \rightarrow English

Final Submission We submitted our VolcTrans online system (unconstrained). The final submission achieves 36.6 BLEU. You can get access to VolcTrans online system on <http://translate.volcengine.cn/>.

¹Large dictionary leads to large parameter size in embedding

²We randomly combine sampling strategies and model architectures to get 9 baseline models for each direction. The performances of the 9 baselines is not the point, what we want is the model diversity among single models

Direction	En \rightarrow Zh
Testset	wmt19
Baseline	38.5
iterative BT	38.9
Ensemble KD	41.5
Ensemble System	42.0
BLEU on WMT20 testset submission	44.9

Table 3: Results of English \rightarrow Chinese by sacreBLEU

4.2 English \rightarrow Chinese

For English \rightarrow Chinese, we train English \leftrightarrow Chinese jointly. We use all parallel data available: News Commentary v15, Wiki Titles v2, UN Parallel Corpus V1.0, CCMT Corpus and WikiMatrix. After data filtering, XM parallel data remained. We use MosesTokenizer for English and Jieba for Chinese. After the pre-processing, separate BPE vocabulary is learned with 32000 merge operations for both English and Chinese on the parallel data. We sample parallel data of ratio 100%, 110% and 120% with replacement from all parallel data. Then we train 3 baselines with Transformer Mid 25e6d, Transformer Mid 50e6d and Dynamic Convolution 25e6d architectures respectively, resulting in 9 baseline models. We employ Newscrawl data as monolingual data for English. The total amount of monolingual data is 90M, containing all Newscrawl 2019 data and others sampled from Newscrawl 2014 to 2018. For Chinese, we employ Newscrawl data, CCMT data and LDC data. We split the Chinese into 3 parts, each contains 8M sentences. For iterative back translation stage and ensemble knowledge distillation stage, each model is combined with different English monolingual data part. Since there are only 3 parts of Chinese monolingual data, we use each part for 3 times at each stage. At the ensemble knowledge distillation stage, we also employ disjoint monolingual data as the distilling data. The detailed results of our system is reported in Table 3.

Final Submission We submitted the ensemble system of the 9 single models after ensemble knowledge distillation stage. The final submission achieves 44.9 BLEU.

4.3 English \leftrightarrow German

For English \leftrightarrow German, we train both directions jointly. We use all parallel data available: Eu-

roparl v10, ParaCrawl v5.1, Common Crawl corpus, News Commentary v15, Wiki Titles v2, Tilde Rapid corpus and WikiMatrix corpus. After data filtering, 28M parallel data remained. We use MosesTokenizer for both English and German. After the pre-processing, a joint BPE vocabulary is learned with 6000 merge operations on the parallel data. We sample parallel data of ratio 80%, 90% and 100% with replacement from all parallel data. Then we train 3 baselines with Transformer Mid 25e6d, Transformer Mid 50e6d and Dynamic Convolution 25e6d architectures respectively, resulting in 9 baseline models. We only employ Newscrawl data as monolingual data for both German and English. The total amount of monolingual data is 90M, containing all Newscrawl 2019 data and others sampled from Newscrawl 2014 to 2018. The 90M data was divided into 9 disjoint parts, each containing 10M sentences, to jointly train 9 systems separately. At the ensemble knowledge distillation stage, we also employ disjoint monolingual data as the distilling data. The detailed results of our system is reported in Table 4.

Direction	En→De		De→En	
Testset	news18	news19	news18	news19
Baseline	47.1	42.2	45.7	41.6
iterative BT	48.6	42.6	48.1	42.2
KD	49.7	44.3	48.4	43.3
Ensemble System	52.2	46.1	49.1	43.8
BLEU on WMT20 testset submission	38.2		43.5	

Table 4: Results of English↔German by sacreBLEU

Fine-tune In this step we use the development sets to handle the domain mismatch problem in WMT. For English→German direction, we fine-tune some of the best single models on news2018 for 1-2 epochs, and then get the final ensemble model from models with fine-tune and models without fine-tune.

Final Submission For either direction, the final submission is an ensemble system from single models with highest BLEU scores on development sets. For the final English→German submission, we replaced the English quote with the German quote. The final submissions on Test20 data achieve 38.2 BLEU for English→German direction and 43.5 BLEU for German→English direction.

4.4 French↔German

For French↔German, we train both directions jointly. The overall parallel data contains 13M sentences available including: Europarl v10, ParaCrawl v5.1, Common Crawl corpus, News Commentary v15, Wiki Titles v2 and WikiMatrix corpus. We train 9 baseline models, each with different architectures (Transformer 15e6d * 2, Transformer Mid 25e6d, Transformer Mid 50e6d, Transformer 15000ffn, Transformer 128hdim, Transformer 256hdim and Dynamic Convolution 25e6d) and each group of three models is combined with 3 different sampling strategies (no sample, up sample 120%, up sample 140%)³, resulting in 9 single models for each direction. We only employ Newscrawl data as monolingual data for both German and French. The monolingual data contains 90M sentences, including all Newscrawl 2019 data and others are sampled from Newscrawl 2014 to 2018. The data of 90M pairs is divided into 9 disjoint parts, each containing 10M sentences to jointly train 9 systems separately. The detailed experiment results are shown in Table 5.

Final Submission For either direction, the final submission is an ensemble system of all 9 models obtained after the knowledge distillation stage. For the final German→French submission, we replaced the English quote with the French quote. The final submissions on Test20 data achieve 35.7 BLEU for French→German direction and 35.3 BLEU for German→French direction.

Direction	Fr→De	De→Fr
Testset	news19	news19
Baseline	26.7	31.3
iterative BT	32.4	35.6
KD	32.7	36.8
Ensemble System	33.9	38.0
BLEU on WMT20 testset submission	35.7	35.3

Table 5: Results of French↔German by sacreBLEU

4.5 English↔Polish

Our English↔Polish systems are based on Europarl v10, ParaCrawl v5.1, Wiki Titles v2, Tilde Rapid corpus and WikiMatrix corpus. All data add

³There is little difference among the models with different sampling ratios, what we are concerned about is the data diversity caused by different sampling ratios.

Direction	En→Pl	Pl→En
Testset	news20 dev	news20 dev
Baseline	24.8	29.7
iterative BT	27.8	32.7
KD	28.2	33.3
Ensemble System	28.7	34.0
BLEU on WMT20 testset submission	26.1	34.4

Table 6: Results of English↔Polish by sacreBLEU

up to 8M sentences. We train 9 models with different architectures (Transformer 15e6d * 2, Transformer Mid 25e6d, Transformer Mid 50e6d, Transformer 15000ffn, Transformer 128hdim, Transformer 256hdim and Dynamic Convolution 25e6d) and each group of three models is combined with different sampling strategies (no sample, up sample 120%, up sample 140%) on both directions. We only used NewsCrawl as English monolingual data. The English monolingual data contains 90M sentences, including all NewsCrawl 2019 data and others sampled from NewsCrawl 2014 to 2018. We divide the data into 9 disjoint parts. Since Polish NewsCrawl corpus only contains 3M sentences, we additionally employ the Polish common crawl data. We sample 90M sentences from the Polish common crawl data which is then divided into 9 disjoint parts. Each part contains 10M common crawl sentences and 3M NewsCrawl sentences. Then we apply the joint training and knowledge distillation as described in Section 3. The detailed experiment results are shown in Table 6.

Final Submission Our final submission is an ensemble system consisting of all 9 models obtained after the ensemble knowledge distillation stage. The final submissions on Test20 data achieve 26.1 BLEU for English→Polish direction and 34.4 BLEU for Polish→English direction.

4.6 English↔Pashto

For English↔Pashto, we used all parallel data containing 13M sentences available as follows: ParaCrawl v5.1, Wiki Titles v2 and the Khmer and Pashto parallel data. For Pashto→English, we fine tune the three pre-trained models on all data with different sampling strategies. Each pre-trained model is fine tuned with three different sampling strategies and we get 9 models. For English→Pashto, we find that the models fine-

Direction	En→Ps	Ps→En
Testset	news20 dev	news20 dev
Baseline	8.4	10.2
+mRASP	-	13.8
iterative BT	9.6	16.4
Ensemble System	-	18.0
BLEU on WMT20 testset submission	10.6	20.0

Table 7: Results of English↔Pashto by sacreBLEU

tuned from the pre-trained models have lower BLEU score than the baseline model trained from scratch, so we use the 9 baseline models which are trained with different architectures and sampling strategies. The English monolingual data has 90M sentences containing all NewsCrawl 2019 data and others sampled from NewsCrawl 2014 to 2018. We divide the data into 9 groups. The detailed experiment results are shown in Table 7.

Final Submission For Pashto→English, our final submission is an ensemble model consisting of all 9 models obtained after the ensemble knowledge distillation stage. For English→Pashto, we find the ensemble model has lower BLEU score than the best single model, so we use the best single model as our final submission. The final submissions achieve 10.6 BLEU on wmt20 testset for English→Pashto direction and 20.0 BLEU for Pashto→English direction.

4.7 English↔Tamil

Direction	En→Ta	Ta→En
Testset	news20 dev	news20 dev
Baseline	7.4	14.0
+mRASP	9.2	17.9
iterative BT	11.8	23.8
BLEU on WMT20 testset submission	7.9	19.7

Table 8: Results of English↔Tamil by sacreBLEU

For English↔Tamil, we use all parallel data containing 533K sentences in total. We use all provided Tamil monolingual data. Following the procedure of English↔German, the English monolingual data contains 90M sentences, including all NewsCrawl 2019 data and others sampled from NewsCrawl 2014 to 2018. The 90M data was di-

vided into 9 disjoint parts, each containing 10M sentences. For Tamil, we don’t apply tokenizer and the raw text is directly split by BPE subword. We fine-tune the three pre-trained models on all parallel data with two different sampling strategies: no sample and up sample 120%, resulting in 6 models. We then conduct back translation for one iteration afterwards, each using one part of the English monolingual data for generating English→Tamil pseudo data, and all Tamil monolingual data for generating Tamil→English pseudo data. The detailed experiment results are shown in Table 8.

Final Submission For both English→Tamil and Tamil→English directions, our final submission is a single model. The final submissions achieve 7.9 BLEU for English→Tamil and 19.7 BLEU for Tamil→English.

4.8 English↔Khmer

Direction	En→Km	Km→En
Testset	news20 dev	news20 dev
Baseline	39.3	12.7
+mRASP	42.8	14.4
iterative BT	46.5	16.9
Ensemble System	-	17.8
BLEU on WMT20 testset submission	51.8	17.6

Table 9: Results of English↔Khmer by sacreBLEU

For English↔Khmer, we used all parallel data containing 4M sentences available as follows: ParaCrawl v5.1 and the Khmer and Pashto parallel data. For Khmer, we extract a dictionary from the Khmer and Pashto parallel data. The km data in this dataset is separated by a special token 200b. Loading this dictionary in the Jieba tokenizer, we get a Khmer tokenizer. We preprocess the Khmer data with our Khmer tokenizer followed by BPE subword. We fine tune the three pre-trained models on all data with different sampling strategies. Each pre-trained model is fine tuned with three different sampling strategies and we get 9 models. We use all provided Khmer monolingual data. The English monolingual data has 90M sentences containing all Newscrawl 2019 data and others sampled from Newscrawl 2014 to 2018. We divide the data into 9 groups. The detailed experiment results are shown in Table 9.

Final Submission For Khmer→English, our final submission is an ensemble model consisting of all 9 models after the iterative back-translation stage. For English→Khmer, we find the ensemble model has lower BLEU score than the best single model, so we use the best single model as our final submission. The final submissions achieve 51.8 BLEU on wmt20 testset for English→Khmer direction and 17.6 BLEU for Khmer→English direction.

5 Conclusion

This paper describes VolcTrans’s NMT systems for the WMT20 shared news translation task. For all directions, we almost adopted the same strategies, except for low-resource language pairs, we employed multilingual pre-training to boost the baseline models. We found that splitting the monolingual data into disjoint parts is an effective way to increase data diversity among single models, which is an important premise for building strong ensemble models. Our final systems achieved significant improvements, usually 3 to 5 BLEU scores, over baseline systems by integrating techniques such as tagged back-translation, iterative back-translation, random ensemble, knowledge distillation.

Acknowledgments

We thank Jun Cao, Zhuo Zhi, Runxin Xu for their support for filtering data, and Zherui Liu for supporting the computing resources. We would also like to thank the anonymous reviewers for their valuable comments.

References

- Srinadh Bhojanapalli, Chulhee Yun, Ankit Singh Rawat, Sashank J Reddi, and Sanjiv Kumar. 2020. Low-rank bottleneck in multi-head attention models. *arXiv preprint arXiv:2002.07028*.
- Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. Tagged back-translation. *arXiv preprint arXiv:1906.06442*.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. Understanding back-translation at scale. pages 489–500.
- Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*.

- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Bei Li, Yinqiao Li, Chen Xu, Ye Lin, Jiqiang Liu, Hui Liu, Ziyang Wang, Yuhao Zhang, Nuo Xu, Zeyang Wang, Kai Feng, Hexuan Chen, Tengbo Liu, Yanyang Li, Qiang Wang, Tong Xiao, and Jingbo Zhu. 2019. [The niutrans machine translation systems for WMT19](#). In *Proceedings of the Fourth Conference on Machine Translation, WMT 2019, Florence, Italy, August 1-2, 2019 - Volume 2: Shared Task Papers, Day 1*, pages 257–266. Association for Computational Linguistics.
- 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 *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2649–2663.
- Lemao Liu, Masao Utiyama, Andrew M. Finch, and Eiichiro Sumita. 2016. [Agreement on target-bidirectional neural machine translation](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 411–416. The 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. *arXiv preprint arXiv:1904.01038*.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. Scaling neural machine translation. *CoRR*, abs/1806.00187.
- 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, Belgium, Brussels. Association for Computational Linguistics.
- Meng Sun, Bojian Jiang, Hao Xiong, Zhongjun He, Hua Wu, and Haifeng Wang. 2019. Baidu neural machine translation systems for wmt19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 374–381.
- 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.
- Mingxuan Wang, Li Gong, Wenhuan Zhu, Jun Xie, and Chao Bian. 2018. Tencent neural machine translation systems for wmt18. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 522–527.
- Felix Wu, Angela Fan, Alexei Baevski, Yann N Dauphin, and Michael Auli. 2019. Pay less attention with lightweight and dynamic convolutions. *arXiv: Computation and Language*.
- Yingce Xia, Xu Tan, Fei Tian, Fei Gao, Weicong Chen, Yang Fan, Linyuan Gong, Yichong Leng, Renqian Luo, Yiren Wang, et al. 2019. Microsoft research asia’s systems for wmt19. *arXiv: Computation and Language*.
- Jiacheng Yang, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Yong Yu, Weinan Zhang, and Lei Li. 2019. Towards making the most of bert in neural machine translation. *arXiv preprint arXiv:1908.05672*.
- Zhirui Zhang, Shujie Liu, Mu Li, Ming Zhou, and Enhong Chen. 2018. Joint training for neural machine translation models with monolingual data. pages 555–562.