

The NiuTrans Machine Translation Systems for WMT21

Shuhan Zhou¹, Tao Zhou¹, Binghao Wei¹, Yingfeng Luo¹, Yongyu Mu¹,
Zefan Zhou¹, Chenglong Wang¹, Xuanjun Zhou¹, Chuanhao Lv¹, Yi Jing¹,
Laohu Wang¹, Jingnan Zhang¹, Canan Huang¹, Zhongxiang Yan¹,
Chi Hu¹, Bei Li¹, Tong Xiao^{1,2} and Jingbo Zhu^{1,2}

¹NLP Lab, School of Computer Science and Engineering, Northeastern University

²NiuTrans Research, Shenyang, China

zhoushuhan199710@163.com, zhoutao_neu@outlook.com

{xiaotong, zhujingbo}@mail.neu.edu.cn

Abstract

This paper describes NiuTrans neural machine translation systems of the WMT 2021 news translation tasks. We made submissions to 9 language directions, including English \leftrightarrow {Chinese, Japanese, Russian, Icelandic} and English \rightarrow Hausa tasks. Our primary systems are built on several effective variants of Transformer, e.g., Transformer-DLCL, ODE-Transformer. We also utilize back-translation, knowledge distillation, post-ensemble, and iterative fine-tuning techniques to enhance the model performance further.

1 Introduction

Our NiuTrans team participated in the WMT 2021 news translation shared tasks, including English \leftrightarrow Chinese (EN \leftrightarrow ZH), English \leftrightarrow Japanese (EN \leftrightarrow JA), English \leftrightarrow Russian (EN \leftrightarrow RU), English \leftrightarrow Icelandic (EN \leftrightarrow IS) and English \rightarrow Hausa (EN \rightarrow HA), nine submissions in total. All of our systems were built with constrained data sets. We adopt some effective models and useful methods, which have been witnessed the success in previous papers (Wang et al., 2018; Li et al., 2019; Zhang et al., 2020; Meng et al., 2020; Wu et al., 2020b; Chen et al., 2020; Yu et al., 2020; Wu et al., 2020a; Wei et al., 2020).

To enhance the performance of the single model, we choose pre-normalized Transformer-DLCL (Wang et al., 2019) and ODE-Transformer (Li et al., 2021a) as the backbone. All systems are built upon the relative position representation (Shaw et al., 2018) due to its strong performance when models are deep (Li et al., 2020). For the system combination, we adopt the post-ensemble (Kobayashi, 2018) to find the most similar hypothesis among several ensemble outputs, which could be regarded as a reranking technique without pre-training. Previous works have emphasized the importance of

diversity when building ensemble systems. Besides the architecture diversity, we also adopt iterative ensemble knowledge distillation leveraging the source-side monolingual data to enlarge the diversity. More details please refer to (Li et al., 2019).

Our data preparation pipeline consists of three-fold: (i) For the data filtering. We use a stricter cleaning process than last year (Zhang et al., 2020). Details will be discussed in Section 2.1. (ii) For the data augmentation, both iterative back-translation (Sennrich et al., 2016a) method, and iterative knowledge distillation (Freitag et al., 2017) method are employed to take the full advantage of monolingual data provided by the WMT organization. In the back-translation stage, we leverage target-side monolingual sentences to generate source-side pseudo sentences and use a nucleus sampling (Holtzman et al., 2019) decoding strategy to improve the generalization ability. Furthermore, we leverage in-domain source-side monolingual data by applying iterative knowledge distillation. (iii) For data selection, it’s hard to find massive in-domain data for low-resource languages to train a neural language model, so we use a statistical n-gram language model (XenC toolkit¹) instead.

Domain finetuning is quite essential to improve the translation system given a certain target domain. We use domain adaptation to migrate the models from the general domain to the news domain by iterative finetuning. After in-domain finetuning, we use multiple ensemble combinations by the post-ensemble method.

This paper is structured as follows: In Section 2, we introduce several effective techniques, including data preprocessing, deeper and wider Transformer models, iterative back-translation, itera-

¹<https://github.com/antho-rousseau/XenC>

Model	Depth	Hidden Size	Filter Size	RPR	Batch size	update freq
Transformer	6	512	2048	✗	4096	1
Transformer (Pre-Norm)	24	512	4096	✓	2048	4
Transformer-DLCL	25	512	4096	✓	2048	4
Transformer-DLCL	30	512	2048	✓	2048	4
Transformer-DLCL	30	512	4096	✓	2048	4
ODE Transformer	6	1024	4096	✓	2048	8
ODE Transformer	12	1024	4096	✓	2048	8

Table 1: The details of several model architectures we used.

tive knowledge distillation, fine-tuning and post-ensemble. In Section 3, we show the experiment settings and report the experimental results of the validation set (newstest2020). Finally, we draw the conclusion in Section 4.

2 System Overview

2.1 Data Preprocessing and Filtering

For word segmentation, we use different tools in six languages. English, Russian, Hausa and Icelandic sentences were segmented by Moses (Koehn et al., 2007), while Chinese and Japanese used NiuTrans (Xiao et al., 2012) and MeCab² separately. Then BPE (Sennrich et al., 2016b) with 32K operations is used for five languages sides independently, except for 36K operations in Russian.

The quality of the parallel training data is crucial to the performance of the models, so we use rigorous data filtering scheme as the suggestion in Zhang et al. (2020)’s work. For most language pairs, rules are as follows:

- Filter out sentences that contain long words over 40 characters or over 150 words.
- The word ratio between the source word and the target word must not exceed 1:3 or 3:1.
- Use Unicode to filter sentences with more than 10 other characters.
- Filter out the sentences which contain HTML tags or duplicated translations.
- In monolingual data, some sentences contain two or more sentences. We write a script to cut them into several sentences.

We use these rules to filter bilingual and monolingual data, detecting low-quality sentences with misalignment, translation errors, illegal characters, and missing translation.

²<https://github.com/taku910/mecab>

2.2 Model Architectures

As shown in previous work (Li et al., 2019; Zhang et al., 2020; Meng et al., 2020), deep Transformers bring significant improvements than the baseline on various machine translation benchmarks. In their work, the performance of the model was significantly improved by increasing the encoder depth. We keep the decoder depth unchanged as the brought benefit is marginal when the encoder is strong enough (Li et al., 2021b).

Hence, we train two deep models in our experiment: Transformer DLCL (Wang et al., 2019) and ODE Transformer (Li et al., 2021a) with a larger filter size. ODE Transformer is designed from the ordinary differential equations (ODE) perspective. Higher-order ODE solutions can gain fewer truncation errors, thus reducing the global error and improving the model performance. The details of several models we mainly experimented with are summarized in Table 1.

In addition, we incorporate relative position representation (RPR) into the self-attention mechanism on both the encoder and decoder sides. Preliminary experiments demonstrate that only relative key information is enough, and we set the relative window size to 8.

2.3 Large-scale Back-Translation

Back-translation (BT) is an effective data augmentation technique to boost the performance of NMT models, which use monolingual data to generate pseudo-training parallel data. Back-translation is divided into three stages:

- Using bilingual parallel data to train a target-to-source intermediate ensemble of models.
- Utilizing the ensemble of reverse direction models to translate the target monolingual corpus into the source corpus.




Figure 1: Process of the post-ensemble method. S_n denotes the sentence generated by the n -th ensemble of models. ALD denotes the average Levenshtein distance of a sentence with other sentences. For example, $ALD 1 = \frac{1}{2} * (Levinstein_distance(S1, S2) + Levinstein_distance(S1, S3))$. Finally, we select the sentence of the smallest ALD.

- Training models with the bilingual parallel corpus and the synthetic parallel corpus together.

Select in-domain monolingual data during back-translation can significantly alleviate domain adaptation problems (Zhang et al., 2020). Our in-domain data consist of the test sets released in recent years and the News Commentary high-quality monolingual data. Due to insufficient data in the domain, we used a statistical method to select in-domain data, the XenC toolkit. Furthermore, to avoid the high ranking of short sentences, we choose the in-domain source side sentences according to the distribution of sentence tokens number in the previous years’ test set.

For all tasks, we employ the beam search and Nucleus Sampling approaches to generate pseudo corpus and the scale of the pseudo corpus was about 1:1 to the real corpus.

2.4 Iterative Knowledge Distillation

Knowledge distillation (KD) has been proven to be a powerful technique to improve the performance of the student model by transferring knowledge from the teacher model (Li et al., 2019; Zhang et al., 2020). Here, we regard the ensemble models as a teacher model and single models as student models. Specifically, we first use the ensemble model to generate synthetic corpus in the forward direction. Then, we merge the synthetic parallel corpus with the bilingual parallel corpus to teach student models. And by searching for better model ensemble combinations, we can provide stronger teacher models for the next round of knowledge distillation. Our experiment found that the gap between the single model and the integrated model gradually narrowed as the iteration progressed. So for the nine tasks we participated in, two iterations

of knowledge distillation deliver the best performance.

2.5 Finetuning

Domain adaptation plays an important role in improving the performance of the models. A practical method of domain adaptation is to train models on large-scale out-domain corpus and then fine-tune the models with in-domain corpus (Luong and Manning, 2015). For all tasks, we mainly reuse an iterative fine-tuning process (Zhang et al., 2020) and use the development sets and the test sets of previous years as in-domain corpus.

It is worth noting that, in order to be consistent with the composition of the test set, we select parallel sentences pair from the previous development sets and test sets in which the source side is real and the target side is manually translated. Moreover, we found that iterative fine-tuning can better improve the translation quality of the names of news organizations in the news field.

2.6 Post-ensemble

Ensemble learning is a technique widely used in several WMT shared tasks, which improves performance by using multiple single models. In neural machine translation, a practical method of the model ensemble is to combine the probability distribution on the target vocabulary of different models in each step of sequence prediction. Here, we adopted their method, which uses a greedy-based strategy to find a better combination of models on the development set. However, enumerating all combinations of candidate models is an inefficient and cumbersome way.

In our ensemble experiments, we set the number of the ensemble to four and six. We observed that simply expanding the scale of the ensemble does not necessarily improve translation performance.

Besides, brute force search for all models is costly and unrealistic. As the number of models increases, the ensemble easily exceeds the computer capacity limit. Therefore, for all tasks, we finally search for four single models as an ensemble.

In addition, we use a simple but effective unsupervised ensemble method, post-ensemble, which uses a clustering method to select a majority-like output from multiple ensembles. As shown in the figure 1, we first choose several ensemble combinations composed of different models to obtain more diversity. Then we use these ensembles to generate multiple sentences, respectively. Next, we calculate the Levenshtein distance between each sentence, and finally, we select the sentence of the smallest average Levenshtein distance with other sentences.

For more detailed content, please refer to the original paper (Kobayashi, 2018). This technology can further improve the performance of the system based on ensemble learning.

3 Experiment

3.1 Experiment Settings

The implementation of our models is based on Fairseq (Ott et al., 2019). All models were trained on 8 RTX 2080Ti GPUs. We selected the pre-norm Transformer-base as the baseline for all tasks and enhanced our deep or wide models by enlarging the model depth and the hidden size, respectively. We used Adam optimizer (Kingma and Ba, 2014) with $\beta_1 = 0.9$, $\beta_2 = 0.997$ during training. As suggested in Ott et al. (2018) and Wang et al. (2019)’s work, models with larger capacities tend to perform much better within large batch size and learning rate. Due to the high GPU memory consumption, accumulated gradients every two steps where each batch contains 2048 tokens. Training for 15 epochs is sufficient for most tasks, and models have shown convergence in validation perplexity. The max learning rate and warmup step were set to 0.002 and 8000 for deep models, and 0.0016 and 16000 for deep and wide models, e.g., Transformer-DLCL, whose hidden dimension is 768. All the dropout probabilities were set to 0.1, including the residual dropout, attention dropout, and the ReLU dropout. We also used FP16 mix-precision training to accelerate further the training process with almost no loss in BLEU.

3.2 EN \leftrightarrow ZH

For EN \leftrightarrow ZH tasks, the training data consists of ParaCrawl, News Commentary v16, WikiMatrix, UN Parallel Corpus V1.0, and the CCMT Corpus. We regarded the newstest2019 as the valid set and the newstest2020 as the test set to tune the hyperparameters. After filtering the data, we sampled the top 12 and 20 million data according to the XenC score as the bilingual dataset. For the ZH \rightarrow EN task, we used 12 and 20 million data to train the baseline model, respectively, and found that the model trained by 12 million data is 1 and 1.2 BLEU point higher than the model trained by 20 million data in the valid and test set. We found that the data quality of the bottom 8 million is lower and also selected the 12 million data as our training data.

During the first-step back-translation, we sampled 8 million monolingual data from the combination of News crawl, News Commentary, News discussions, and News crawl. Then we used the baseline model to generate the hypotheses via the beam search strategy as the pseudo dataset. In the second-step back-translation, we utilized the same amount of pseudo data while using nucleus sampling, whose p is 0.9. For ZH \rightarrow EN and EN \rightarrow ZH, we got BLEU improvements of 1.8 and 2.9 in the first back-translation and further BLEU improvements of 0.5 and 0.8 in the second back-translation, respectively.

In addition, we implemented knowledge distillation twice to iteratively enhance the single model with the ensemble outputs. The main goal is to make the single student mimic the behavior of the ensemble models, thus obtaining stronger ensemble teachers in the next step. We used the test sets in previous years as in-domain data in EN \rightarrow ZH and EN \rightarrow ZH directions respectively, and we used the XenC tool to sample 3 million from the large scale monolingual data based on in-domain data. Then we used the best ensemble of models to construct pseudo data by decoding them and merge them to the original training data to continue training for each model. We got BLEU improvements of 1.1 and 0.6 in the first knowledge distillation and further BLEU improvements of 0.6 and 0.3 in the second knowledge distillation in ZH \rightarrow EN and EN \rightarrow ZH.

After knowledge distillations, we used the newstest2017-2019 to fine-tune our models for five epochs with the 0.0001 learning rate and got 2 and 0.4 BLEU improvements in EN \rightarrow ZH and EN \rightarrow ZH

System	EN→ZH	ZH→EN	EN→JA	JA→EN	EN→HA
Baseline	41.9	30.1	34.5	21.4	10.9
DLCL30-RPR	42.8	31.0	35.6	21.6	11.9
+Iteratively BT	46.5	33.3	38.4	21.7	16.5
+Iteratively KD	47.4	35.0	41.8	25.9	18.2
+Fine-tune	47.8	37.0	42.0	26.4	-
+Ensemble	48.8	37.2	42.7	27.4	18.5
+Post-ensemble & Post edit	49.0	37.5	43.6	27.4	-

Table 2: BLEU evaluation results on the WMT 2020 EN↔ZH, EN↔JA test sets and WMT 2021 EN→HA development sets.

directions, respectively. In the final stage, we add newstest2020 to the fine-tuning data. Finally, we searched for the best five combinations of 4 out of 12 models for post-ensemble to ensure the diversity of the models. Based on the ensemble method, post-ensemble further brought us +0.2 and +0.3 BLEU in ZH→EN and EN→ZH directions. Our main results showed in table 2, we find that iterative back-translation, iterative knowledge distillation, and iterative fine-tune are effective methods to get significant improvements.

3.3 EN↔JA

For EN↔JA tasks, we chose ParaCrawl v7.1, News Commentary v16, WikiMatrix, Japanese-English Subtitle Corpus, The Kyoto Free Translation Task Corpus, TED Talks total of six parallel data corpora about 17.5 million. For the ParaCrawl v7.1, we only selected 8.5 million data according to the score of sentences provided by the dataset. We chose all of News Crawl and News Commentary and 12 million data sampled from Common Crawl for the Japanese monolingual data. After merging corpora into training data, we found that there were many-to-one situations in both the target side and the source side. Therefore, we sorted sentences and calculated the Levenshtein ratio of two adjacent sentences to remove duplication sentences. We applied this method to all version data before training models and removed 10 percent of the total data. We randomly selected one out of many sentences in which Levenshtein ratios are greater than or equal to 0.9.

We also implemented tagged back-translation, which brought us +2.8 BLEU in the EN→JA task. In addition, beam search and nucleus sampling were used to generate two parts of translations to increase data diversity, and each part contains 12 million data. An interesting phenomenon is that

back-translation is useful for EN→JA task while knowledge distillation is helpful for JA→EN task. We suspect this is because the domain of Japanese monolingual more fits the field of the test set.

We also implemented knowledge distillation and fine-tuned iteratively. During the knowledge distillation phase, we used FDA³ and XenC to select monolingual data more like newstest2020 and generated pseudo data by using both post-ensemble and ensemble methods. During the fine-tuning phase, we used the WMT 2020 valid set and opposite direction test set. After performance stopped increasing at the second fine-tune, we utilized the best ensemble models to regenerate pseudo data by back-translation and knowledge distillation. Then, we retrained multiple deep models. Finally, we put all models together to greedy search for the best combination of 13 models. And this method brought us +0.7 BLEU in JA→EN task. Our main results are shown in table 2.

3.4 EN↔RU

For EN↔RU tasks, we used only two parallel datasets, including ParaCrawl v8 and News Commentary. After the data filter, about 12M sentence pairs were left to build our system. Additionally, we set the merge operations of BPE to 36K.

We also used iterative back-translation, iterative knowledge distillation, and fine-tuned to enhance the model. During the back-translation, English monolingual data is the same as the EN↔JA part, and Russian monolingual data sources consist of News Crawl and News Commentary. During the knowledge distillation, we used FDA to select 4 million sentence pairs from the monolingual dataset according to the newstest2020 and newstest2019. Then we merged them with the official development set to continue training our

³<https://github.com/bicici/FDA>

System	EN→RU	RU→EN	EH→IS	IS→EN
Baseline	22.0	35.6	20.9	28.4
ODE big6-RPR	22.7	36.8	22.4	30.5
+Iteratively BT	23.0	38.2	28.5	34.9
+Iteratively KD	23.3	38.9	30.7	36.0
+Fine-tune	24.4	39.4	-	-
+Ensemble	24.8	39.9	31.2	36.4

Table 3: BLEU evaluation results on the WMT 2020 EN↔RU test sets and WMT 2021 EN↔IS development sets.

Task	Submission	Task	Submission
EN→ZH	35.8	EN→RU	28.4
ZH→EN	31.9	RU→EN	41.8
EH→JA	46.2	EH→IS	30.6
JA→EN	27.2	IS→EN	39.2
EH→HA	19.7	-	-

Table 4: Our final submission results in nine tasks.

models for five epochs. After KD, we used the newstest2019 and newstest2018 to fine-tune our models for five epochs with the 0.0001 learning rate and got 1.1 and 0.5 BLEU improvements in EN→RU and RU→EN.

The detailed and full results can be described in Table 3. Iterative BT, KD, and fine-tune are still very effective and improved 2.8 and 4.3 compared with the base model in EN→RU and RU→EN tasks, respectively.

3.5 EN↔IS

The process of EN↔IS tasks is similar to EN→HA task but more complicated. Concretely, we used four parallel datasets, including ParaCrawl v7.1, Wiki Titles v3, WikiMatrix, and ParIce. After the data filtering, about 5.5 million sentence pairs were left to build the baseline system. The experimental results are listed in Table 3. We obtained significant improvements of 6.1 and 4.4 BLEU in EN→IS and IS→EN directions, respectively.

Then we implemented iterative KD two times and sampled 3 million in-domain source data according to WMT2021 development sets. Table 3 shows that it’s a very effective method to get 2.2 and 1.1 improvements. Furthermore, we fine-tuned models iteratively twice to transfer the knowledge into the target domain. Due to implementing two ensemble combinations to decode sentences, the model ensemble still gained 0.7 and 0.8 improvements.

3.6 EN→HA

In the EN→HA direction, we used ParaCrawl v8, Khamenei corpus, and English-Hausa Opus corpus three data sets, obtaining 1.43M parallel data after cleaning. We collected News crawl, Extended Common Crawl, and Common Crawl for the monolingual data, resulting in 5.7M Hausa monolingual data. Considering the insufficient scale of Hausa, we used all monolingual data in each round of back-translation. The implementation details of iterative knowledge distillation and back-translation are almost the same as the EN↔ZH tasks.

Table 2 summarized the results. We can observe that the wide and deep models were still effective in low-resource language pairs. Through the back-translation and knowledge distillation techniques, we gain 4.6 and 1.7 BLEU improvements, respectively.

3.7 Submission Results

The results we finally submitted are shown in table 4. We participated in nine tasks this year. On the whole, all of our systems performed competitively, especially in EH→IS and RU→EN directions.

Through all the experimental results, we found that different methods perform differently on nine tasks. Among them, iterative BT is effective for almost all tasks, except for the JA→EN task. Iterative KD performs better for EN↔ZH, EN↔JA and EH↔IS tasks, while fine-tune is more suitable for ZN→EN and EN→RU tasks.

4 Conclusion

This paper introduced our submissions on WMT21 nine tasks. Our main exploration is using a new effective architectures ODE Transformer and utilizing post-ensemble technology to enhance the system. And we experimented with iterative back-translation by different decoding strategies, iterative knowledge distillation, iterative fine-tuning, model ensembling and post-ensemble.

References

- Peng-Jen Chen, Ann Lee, Changhan Wang, Naman Goyal, Angela Fan, Mary Williamson, and Jiatao Gu. 2020. [Facebook AI’s WMT20 news translation task submission](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 113–125, Online. Association for Computational Linguistics.
- Markus Freitag, Yaser Al-Onaizan, and Baskaran Sankaran. 2017. [Ensemble distillation for neural machine translation](#). *CoRR*, abs/1702.01802.
- Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2019. [The curious case of neural text degeneration](#). *CoRR*, abs/1904.09751.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Hayato Kobayashi. 2018. [Frustratingly easy model ensemble for abstractive summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4165–4176, Brussels, Belgium. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Bei Li, Quan Du, Tao Zhou, Shuhan Zhou, Xin Zeng, Tong Xiao, and Jingbo Zhu. 2021a. [ODE transformer: An ordinary differential equation-inspired model for neural machine translation](#). *CoRR*, abs/2104.02308.
- 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 (Volume 2: Shared Task Papers, Day 1)*, pages 257–266, Florence, Italy. Association for Computational Linguistics.
- Bei Li, Ziyang Wang, Hui Liu, Quan Du, Tong Xiao, Chunliang Zhang, and Jingbo Zhu. 2021b. [Learning light-weight translation models from deep transformer](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 13217–13225. AAAI Press.
- Bei Li, Ziyang Wang, Hui Liu, Yufan Jiang, Quan Du, Tong Xiao, Huizhen Wang, and Jingbo Zhu. 2020. [Shallow-to-deep training for neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 995–1005, Online. Association for Computational Linguistics.
- Minh-Thang Luong and Christopher D. Manning. 2015. Stanford neural machine translation systems for spoken language domains.
- Fandong Meng, Jianhao Yan, Yijin Liu, Yuan Gao, Xianfeng Zeng, Qinsong Zeng, Peng Li, Ming Chen, Jie Zhou, Sifan Liu, and Hao Zhou. 2020. [WeChat neural machine translation systems for WMT20](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 239–247, Online. 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.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468, New Orleans, Louisiana. Association for Computational Linguistics.
- Qiang Wang, Bei Li, Jiqiang Liu, Bojian Jiang, Zheyang Zhang, Yinqiao Li, Ye Lin, Tong Xiao, and Jingbo Zhu. 2018. [The NiuTrans machine translation system for WMT18](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages

528–534, Belgium, Brussels. Association for Computational Linguistics.

Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. 2019. [Learning deep transformer models for machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1810–1822, Florence, Italy. Association for Computational Linguistics.

Daimeng Wei, Hengchao Shang, Zhanglin Wu, Zhengzhe Yu, Liangyou Li, Jiabin Guo, Minghan Wang, Hao Yang, Lizhi Lei, Ying Qin, and Shiliang Sun. 2020. [HW-TSC’s participation in the WMT 2020 news translation shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 293–299, Online. Association for Computational Linguistics.

Liwei Wu, Xiao Pan, Zehui Lin, Yaoming Zhu, Mingxuan Wang, and Lei Li. 2020a. [The volctrans machine translation system for WMT20](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 305–312, Online. Association for Computational Linguistics.

Shuangzhi Wu, Xing Wang, Longyue Wang, Fangxu Liu, Jun Xie, Zhaopeng Tu, Shuming Shi, and Mu Li. 2020b. [Tencent neural machine translation systems for the WMT20 news translation task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 313–319, Online. Association for Computational Linguistics.

Tong Xiao, Jingbo Zhu, Hao Zhang, and Qiang Li. 2012. [NiuTrans: An open source toolkit for phrase-based and syntax-based machine translation](#). In *Proceedings of the ACL 2012 System Demonstrations*, pages 19–24, Jeju Island, Korea. Association for Computational Linguistics.

Lei Yu, Laurent Sartran, Po-Sen Huang, Wojciech Stokowiec, Domenic Donato, Srivatsan Srinivasan, Alek Andreev, Wang Ling, Sona Mokra, Agustin Dal Lago, Yotam Doron, Susannah Young, Phil Blunsom, and Chris Dyer. 2020. [The DeepMind Chinese–English document translation system at WMT2020](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 326–337, Online. Association for Computational Linguistics.

Yuhao Zhang, Ziyang Wang, Runzhe Cao, Binghao Wei, Weiqiao Shan, Shuhan Zhou, Abudurexiti Reheman, Tao Zhou, Xin Zeng, Laohu Wang, Yongyu Mu, Jingnan Zhang, Xiaoqian Liu, Xuanjun Zhou, Yinqiao Li, bei Li, Tong Xiao, and Jingbo Zhu. 2020. [The NiuTrans machine translation systems for WMT20](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 338–345, Online. Association for Computational Linguistics.