

Summer: WeChat Neural Machine Translation Systems for the WMT22 Biomedical Translation Task

Ernan Li, Fandong Meng and Jie Zhou

Pattern Recognition Center, WeChat AI, Tencent Inc, China
{cardli,fandongmeng,withtomzhou}@tencent.com

Abstract

This paper introduces WeChat’s participation in WMT 2022 shared biomedical translation task on Chinese→English. Our systems are based on the Transformer (Vaswani et al., 2017), and use several different Transformer structures to improve the quality of translation. In our experiments, we employ data filtering, data generation, several variants of Transformer, fine-tuning and model ensemble. Our Chinese→English system, named Summer, achieves the highest BLEU score among all submissions.

1 Introduction

This article describes the WeChat’s participation in WMT 2022 shared biomedical translation task on Chinese→English. We improve the translation quality of the system by increasing the diversity of model structure and data, fine-tuning the model with in-domain data, inserting tags at the beginning of each source sentence and selecting models with high diversity and good performance for ensemble.

For model architectures, our system adopt BIG and DEEP Transformer models which contain 10-layer and 20-layer encoders, 10240 and 4096 filter sizes, respectively, with TRANSFORMER-BIG setting (Vaswani et al., 2017). In order to increase the diversity of the model, we use structures such as Average Attention Transformer (AAN) (Zhang et al., 2018) and Mixed-AAN Transformer architecture (Zeng et al., 2021) in the decoder part.

For data generation, we use back-translation (Sennrich et al., 2016a), knowledge distillation (Kim and Rush, 2016), and forward-translation (Zeng et al., 2021) to improve data quality. And we use some data augmentation methods to improve the model robustness, such as adding synthetic noise and dynamic top-p sampling (Zeng et al., 2021). Furthermore, according to the different sources of the corpora,

we add tags at the beginning of the source sentence to perform domain adaptation.

For fine-tuning, we use in-domain bilingual corpus to fine-tune models from the general domain to the biomedical domain, and use target denoising (Meng et al., 2020) to improve the diversity of models and mitigate training-generation discrepancy.

For model ensemble, we use Self-BLEU (Zhu et al., 2018) to evaluate the similarity between models. We take the prediction of one model as the reference and use the prediction of the other model to calculate the BLEU score. The higher the Self-BLEU score, the lower the diversity of the models.

In the remainder of this paper, we start with presenting the data strategy in Section 2. Then we describe our system details in Section 3. Section 4 presents the experimental results. Finally, we conclude our work in Section 5.

2 Data

In this section, we introduce the details of bilingual and monolingual data used in this shared task.

2.1 Bilingual Corpus

Our baseline model is trained with out-of-domain (OOD) data from WMT 2022 shared task on general machine translation¹. Additionally, we use in-house data (depicted in Table 1 as OOD-IN-HOUSE) to improve performance of baseline model. With regard to in-domain data, firstly, we use the in-domain bilingual corpus provided by the WMT 2022 shared biomedical translation task² (depicted in Table 1 as IND-BIO). And we use the Champollion³ tool to align the sentences in the corpus. Then, we collect in-domain Chinese→English

¹<https://statmt.org/wmt22/translation-task.html>

²<https://github.com/biomedical-translation-corpora/corpora>

³<http://champollion.sourceforge.net/>

(depicted in Table 1 as IND-TAUS) sentence pairs from TAUS⁴.

2.2 Monolingual Corpus

The out-of-domain monolingual corpora are collected from WMT 2022 shared task on general machine translation and the in-house monolingual data. With regard to in-domain data, the English part of the bilingual corpus in other languages provided by the WMT 2022 shared biomedical translation task is used as in-domain monolingual data.

3 System overview

In this section, we introduce the details of our system used in the WMT 2022 shared biomedical translation task. Our system adopts data filtering, data generation, model architectures, fine-tuning and ensemble.

3.1 Data Filtering

For data filtering, we use the following rules for bilingual corpus:

- Normalize punctuation with Moses scripts on both English and Chinese.
- Filter out sentence pairs that are the same at the source and target.
- Filter out sentence pairs whose source sentence’s language recognition result is different from the original language.
- Filter out sentence pairs with a source-to-target length ratio greater than 1:3.
- Filter out the sentences longer than 150 words or exceed 40 characters in a single word.

Besides these rules, we use fast-align⁵ to filter out the sentence pairs with low alignment scores. We also filter out sentence pairs in which English sentences contain Chinese characters.

3.2 Data Generation

In this section, we introduce the approaches of data generation in our system, including back-translation, knowledge distillation, forward-translation, synthetic noise and tagging.

⁴<https://taus-corona-corpus.s3.amazonaws.com/en-zh.txt.gz>

⁵https://github.com/clab/fast_align

3.2.1 Back-Translation

Back-translation (Hoang et al., 2018) is the most commonly used data augmentation method in neural machine translation. Following the previous work (Edunov et al., 2018), we use following strategies to generate back translations to improve the diversity the training data:

- Beam search: We use beam search to generate the pseudo corpus with beam size setting to 4.
- Dynamic top-p sampling: Following the work (Zeng et al., 2021), at each decoding step, we select a word from the smallest set whose cumulative probability exceeds p , with p varying from 0.9 to 0.95 during the data generation process.

3.2.2 Knowledge Distillation

For knowledge distillation (Kim and Rush, 2016; Wang et al., 2021), we use the corpus generated from the teacher models to train the student models.

3.2.3 Forward-Translation

For forward-translation, we use an ensemble model to generate forward translations with the source-language monolingual corpus as input.

3.2.4 Synthetic Noise

For synthetic noise, we add different noises at the source side of the pseudo corpus to improve the diversity of the data and improve the robustness of the model:

- Randomly replace some source tokens with $\langle unk \rangle$.
- Randomly delete some tokens from the source sentence.
- Randomly swap the two tokens in the source sentence in the specify window.

3.2.5 Tagging

For tagging, inspired by (Johnson et al., 2017), we insert a tag at the beginning of each source sentence to denote its type: $\langle BT \rangle$ for the back-translation data, $\langle NOISE \rangle$ for the synthetic noise data, $\langle REAL \rangle$ for the ground-truth bilingual corpus and $\langle FT \rangle$ for the forward-translation data. Furthermore, we insert a tag at the second position of each sentence to denote its domain: $\langle BIO \rangle$ for the in-domain data, $\langle NEWS \rangle$ for the data from WMT22 general

LANGUAGE	OOD-NEWS	OOD-IN-HOUSE	IND-BIO	IND-TAUS
bilingual corpus	30.6M	90M	89K	0.4M
monolingual corpus	220M	50M	6.9M	–

Table 1: Data used for training the system, where *OOD-NEWS* is the out-of-domain data provided by WMT22 general translation task. *OOD-IN-HOUSE* is the out-of-domain data collected from in-house corpus. *IND-BIO* is the in-domain data provided by WMT22 shared biomedical translation task. And *IND-TAUS* is the in-domain data collected manually (not from MEDLINE, as depicted in 2.1). *M* denotes *million* and *K* denotes *thousand*.

translation task and $\langle INHOUSE \rangle$ for the data from our in-house corpus. At inference time, we always use the $\langle REAL \rangle$ and $\langle BIO \rangle$ tag.

3.3 Model Architectures

In this section, we introduce the model architectures used by our system, including Transformer (Big/Deep), Average Attention Transformer (AAN) and Mixed Average Attention Transformer (Mixed-AAN) (Zeng et al., 2021).

3.3.1 Transformer

Our baseline models are Big- and Deep-Transformer (Vaswani et al., 2017) models. In our experiments, we use multiple model configurations with 20-layer and 30-layer encoders for deep models and 10-layers encoders for big models, and use 6-layers decoders for all models. The hidden size is set to 1024 and the filtering size is set from 4096 to 10240.

3.3.2 Average Attention Transformer

To increase the diversity between models, we adopt Average Attention Transformer (Zhang et al., 2018), where the average attention is used to replace self-attention in the decoder. AAN summarizes the historical information of previous positions by means of cumulative average, which increases diversity with almost no harm to the quality of the model.

3.3.3 Mixed-AAN Transformers

Following the previous work (Zeng et al., 2021), we adopt the Mixed-AAN Transformers to further improve the diversity and quality of models. In this experiment, we only use two architectures of Mixed-AAN:

- Self-first: In the decoder part, we use self-attention as the first layer, and then use average attention and self-attention alternately.
- AAN-first: In the decoder part, we use average attention as the first layer, and then use self-attention and average attention alternately.

3.4 Fine-tuning

For fine-tuning, we mainly use the in-domain data provided by WMT22 shared biomedical translation task for domain adaption (Luong and Manning, 2015; Li et al., 2019). In order to prevent the model from overfitting, as well as to improve the diversity of the model after domain transfer, we adopt target denoising (Meng et al., 2020). We add synthetic noise at the decoder inputs during fine-tuning. Therefore, with target denoising, the model becomes more robust. The method of adding synthetic noise is described in Section 3.2.4.

3.5 Ensemble

After obtaining a variety of different models through the above methods, we need to find the best model combination to get the best result. In general, the better the model performance and the greater the diversity between models, the better the performance for the model ensemble. To measure diversity, we use Self-BLEU (Zhu et al., 2018) to evaluate the similarity between models. Overall, we select 6 models from 52 candidate models for ensemble. All the candidate models are generated by different combinations of data and different training strategies as described earlier.

4 Experiments

4.1 Settings

Our experiment is based on Fairseq⁶. The single models are carried out on 8 NVIDIA V100 / A100 GPUs. We adopt the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.998$. The batch-size is set to 4096 tokens, and the “update-freq” is set to 4, and the warmup step is set to 4000 and the learning rate is set to 0.0005.

4.2 Pre-processing and Post-processing

The Chinese sentences are segmented by a in-house segmentation tool and English sentences are seg-

⁶<https://github.com/pytorch/fairseq>

System	BLEU
Baseline	34.57
+ IND-TAUS	35.65
+ IND-BIO	40.96
+ OOD-IN-HOUSE	41.88
+ Back-Translation	42.8
+ Knowledge Distillation	43.12
+ Forward-Translation	43.32
+ Multi BT	44.11
+ Finetune	44.96
+ Target denoise finetune	45.1
Baseline_TAG	34.48
+ IND-TAUS	35.62
+ IND-BIO	41.07
+ OOD-IN-HOUSE	42.14
+ Back-Translation	43.91
+ Knowledge Distillation	44.14
+ Forward-Translation	44.39
+ Multi BT	45.23
+ Finetune	45.43
+ Target denoise finetune	45.54
+ Ensemble	46.91*

Table 2: Translation performance on WMT21 biomedical translation task testset. * is the system we submitted. Multi BT means the iterative back-translation (Hoang et al., 2018) which use with different part of data and different generation strategies.

mented by the tokenizer toolkit in Moses⁷. We normalize punctuation using Moses scripts on both English and Chinese. For handling uppercase and lowercase of the English letters, we add a special token at the beginning of a word to denote uppercase (`_UU_`) and title case (`_U_`). By this way to reduce the size of the word list and reduce the difficulty of model training. For instance, "We are together NOW." \rightarrow "`_U_` we are together `_UU_` now.". We use BPE (Sennrich et al., 2016b) with 32K operations for all the languages.

With the regard of post-processing, we use *detokenizer.perl* on the English translations provided in Moses.

4.3 Results

The experimental results of Chinese \rightarrow English on WMT21 OK-aligned biomedical test set are shown in Table 2.

Compared with the baseline model (Baseline_TAG), the in-domain bilingual data (+IND-BIO) provided by WMT22 shared biomedical

translation task brings a huge improvement, with 6.5 point increase in BLEU score. After adding the in-house out-of-domain corpus (+OOD-IN-HOUSE), we further gain +1.1 BLEU. We further obtain +1.8 BLEU by applying back-translation (+Back-Translat), and +0.23 BLEU by using knowledge distillation (+Knowledge Distillation), and +0.25 BLEU by using forward-translation (+ Forward-Transla). After using iterative back-translation (Hoang et al., 2018) (+Multi BT) described in Table 2, we further achieve improvement of +0.84 BLEU.

Additionally, we can find that the model with TAG was similar to the model without TAG in early stage experiments. As the number of data categories and data domains increases, the model with tags gradually demonstrates its advantages. Our best single model (+Target denoise fine) achieves 45.54 BLEU score, and we finally achieve 46.91 BLEU score by model ensemble (+Ensemble).

⁷<http://www.statmt.org/moses/>

5 Conclusion

We introduce WeChat’s participation in WMT 2022 shared biomedical translation task on Chinese→English. Our system is based on the Transformer (Vaswani et al., 2017), and uses several different Transformer structures such as Average Attention and Mixed-AAN to improve the performance. We use several data augmentation methods such as iterative back-translation, knowledge distillation, forward-translation and synthetic noise. We use tags to assist the model in domain learning and use in-domain fine-tuning with target denoising to domain transfer. Finally a Self-BLEU based ensemble method is used for model ensemble. Overall, our system achieves 46.91 BLEU score on WMT21 OK-aligned biomedical test set, and we achieve the highest BLEU score among all submissions.

References

- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium.
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. Iterative back-translation for neural machine translation. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 18–24.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351.
- 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.
- 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.
- Minh-Thang Luong and Christopher D Manning. 2015. Stanford neural machine translation systems for spoken language domains. In *Proceedings of the International Workshop on Spoken Language Translation*, pages 76–79.
- 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.
- 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.
- 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.
- 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.
- Fusheng Wang, Jianhao Yan, Fandong Meng, and Jie Zhou. 2021. [Selective knowledge distillation for 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 6456–6466, Online.
- Xianfeng Zeng, Yijin Liu, Ernan Li, Qiu Ran, Fandong Meng, Peng Li, Jinan Xu, and Jie Zhou. 2021. [WeChat neural machine translation systems for WMT21](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 243–254, Online. Association for Computational Linguistics.
- Biao Zhang, Deyi Xiong, and Jinsong Su. 2018. [Accelerating neural transformer via an average attention network](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1789–1798, Melbourne, Australia.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Tegygen: A benchmarking platform for text generation models. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1097–1100.