

# WeChat Neural Machine Translation Systems for WMT21

Xianfeng Zeng<sup>1\*</sup>, Yijin Liu<sup>12\*</sup>, Ernan Li<sup>1\*</sup>, Qiu Ran<sup>1\*</sup>, Fandong Meng<sup>1\*</sup>,  
Peng Li<sup>1</sup>, Jinan Xu<sup>2</sup>, and Jie Zhou<sup>1</sup>

<sup>1</sup> Pattern Recognition Center, WeChat AI, Tencent Inc, China

<sup>2</sup> Beijing Jiaotong University, Beijing, China

{xianfzeng,yijinliu,cardli,soulcaptran,fandongmeng,patrickpli,withtomzhou}@tencent.com  
jaxu@bjtu.edu.cn

## Abstract

This paper introduces WeChat AI’s participation in WMT 2021 shared news translation task on English→Chinese, English→Japanese, Japanese→English and English→German. Our systems are based on the Transformer (Vaswani et al., 2017) with several novel and effective variants. In our experiments, we employ data filtering, large-scale synthetic data generation (i.e., back-translation, knowledge distillation, forward-translation, iterative in-domain knowledge transfer), advanced finetuning approaches, and boosted Self-BLEU based model ensemble. Our **constrained** systems achieve 36.9, 46.9, 27.8 and 31.3 case-sensitive BLEU scores on English→Chinese, English→Japanese, Japanese→English and English→German, respectively. The BLEU scores of English→Chinese, English→Japanese and Japanese→English are the highest among all submissions, and that of English→German is the highest among all constrained submissions.

## 1 Introduction

We participate in the WMT 2021 shared news translation task in three language pairs and four language directions, English→Chinese, English↔Japanese, and English→German. In this year’s translation tasks, we mainly improve the final ensemble model’s performance by increasing the diversity of both the model architecture and the synthetic data, as well as optimizing the ensemble searching algorithm.

Diversity is a metric we are particularly interested in this year. To quantify the diversity among different models, we compute Self-BLEU (Zhu et al., 2018) from the translations of the models on the valid set. To be precise, we use the translation of one model as the hypothesis and the translations of other models as references to calculate an aver-

age BLEU score. A higher Self-BLEU means this model is less diverse.

For model architectures (Vaswani et al., 2017; Meng and Zhang, 2019; Yan et al., 2020), we exploit several novel Transformer variants to strengthen model performance and diversity. Besides the Pre-Norm Transformer, the Post-Norm Transformer is also used as one of our baselines this year. We adopt some novel initialization methods (Huang et al., 2020) to alleviate the gradient vanishing problem of the Post-Norm Transformer. We combine the Average Attention Transformer (AAN) (Zhang et al., 2018) and Multi-Head-Attention (Vaswani et al., 2017) to derive a series of effective and diverse model variants. Furthermore, Talking-Heads Attention (Shazeer et al., 2020) is introduced to the Transformer and shows a significant diversity from all the other variants.

For the synthetic data generation, we exploit the large-scale back-translation (Sennrich et al., 2016a) method to leverage the target-side monolingual data and the sequence-level knowledge distillation (Kim and Rush, 2016) to leverage the source-side of bilingual data. To use the source-side monolingual data, we explore forward-translation by ensemble models to get general domain synthetic data. We also use iterative in-domain knowledge transfer (Meng et al., 2020) to generate in-domain data. Furthermore, several data augmentation methods are applied to improve the model robustness, including different token-level noise and dynamic top-p sampling.

For training strategies, we mainly focus on scheduled sampling based on decoding steps (Liu et al., 2021b), the confidence-aware scheduled sampling (Mihaylova and Martins, 2019; Duckworth et al., 2019; Liu et al., 2021a), the target denoising (Meng et al., 2020) method and the Graduated Label Smoothing (Wang et al., 2020) for in-domain finetuning.

For model ensemble, we select high-potential

\* Equal contribution.

candidate models based on two indicators, namely model performance (BLEU scores on valid set) and model diversity (Self-BLEU scores among all other models). Furthermore, we propose a search algorithm based on the Self-BLEU scores between the candidate models with selected models. We observed that this novel method can achieve the same BLEU score as the brute force search while saving approximately 95% of search time.

This paper is structured as follows: Sec. 2 describes our novel model architectures. We present the details of our systems and training strategies in Sec. 3. Experimental settings and results are shown in Sec. 4. We conduct analytical experiments in Sec. 5. Finally, we conclude our work in Sec. 6.

## 2 Model Architectures

In this section, we describe the model architectures used in the four translation directions, including several different variants for the Transformer (Vaswani et al., 2017).

### 2.1 Model Configurations

Deeper and wider architectures are used this year since they show strong capacity as the number of parameters increases. In our experiments, we use multiple model configurations with 20/25-layer encoders for deeper models and the hidden size is set to 1024 for all models. Compared to our WMT20 models (Meng et al., 2020), we also increase the decoder depth from 6 to 8 and 10 as we find that gives a certain improvement, but deeper depths give limited performance gains. For the wider models, we adopt 8/12/15 encoder layers and 1024/2048 for hidden size. The filter sizes of models are set from 8192 to 15000. Note that all the above model configurations are applied to the following variant models.

### 2.2 Transformer with Different Layer-Norm

The Transformer (Vaswani et al., 2017) with Pre-Norm (Xiong et al., 2020) is a widely used architecture in machine translation. It is also our baseline model as its performance and training stability is better than the Post-Norm counterpart.

Recent studies (Liu et al., 2020; Huang et al., 2020) show that the unstable training problem of Post-Norm Transformer can be mitigated by modifying initialization of the network and the successfully converged Post-Norm models generally outperform Pre-Norm counterparts. We adopt these

initialization methods (Huang et al., 2020) to our training flows to stabilize the training of deep Post-Norm Transformer. Our experiments have shown that the Post-Norm model has a good diversity compared to the Pre-Norm Model and slightly outperform the Pre-Norm Model. We will further analyze the model diversity of different variants in Sec. 5.1.

### 2.3 Average Attention Transformer

We also use Average Attention Transformer (AAN) (Zhang et al., 2018) as we used last year to introduce more model diversity. In the Average Attention Transformer, a fast and straightforward average attention is utilized to replace the self-attention module in the decoder with almost no performance loss. The context representation  $g_i$  for each input embedding is as follows:

$$g_i = FFN\left(\frac{1}{i} \sum_{k=1}^i y_k\right) \quad (1)$$

where  $y_k$  is the input embedding for step  $k$  and  $i$  is the current time step.  $FFN(\cdot)$  denotes the position-wise feed-forward network proposed by Vaswani et al. (2017).

In our preliminary experiments, we observe that the Self-BLEU (Zhu et al., 2018) scores between AAN and Transformer are lower than the scores between the Transformer with different configurations.

### 2.4 Weighted Attention Transformer

We further explore three weighting strategies to improve the modeling of history information from previous positions in AAN. Compared to the average weight across all positions, we try three methods including decreasing weights with position increasing, learnable weights and exponential weights. In our experiments, We observe exponential weights perform best among all these strategies. The exponential weights context representation  $g_i$  is calculated as follows:

$$c_i = (1 - \alpha)y_i + \alpha \cdot c_{i-1} \quad (2)$$

$$g_i = FFN(c_i) \quad (3)$$

where  $\alpha$  is a tuned parameter. In our previous experiments, we test different alpha, including 0.3, 0.5, and 0.7, on the valid set and we set the alpha to 0.7 in all subsequent experiments as it slightly outperform the others.

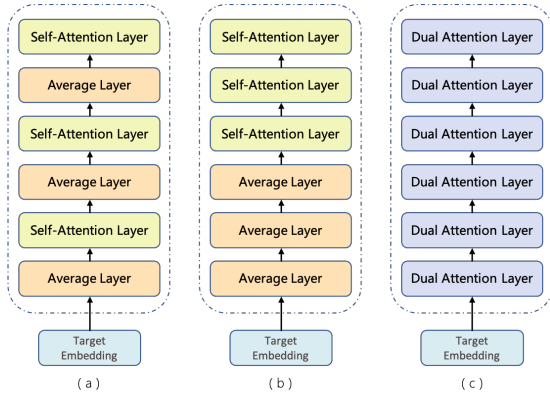


Figure 1: Mixed-AAN Transformers.

## 2.5 Mixed-AAN Transformers

Our preliminary experiments show that the decoder structure is strongly related to the model diversity in the Transformer. Therefore, we propose to stack different types of decoder layers to derive different Transformer variants. As shown in Figure 1, we mainly adopt three Mixed-AAN Transformer architectures: a) Alternately mixing the standard self-attention layer and the average attention layer, b) Continuously stacking several average attention layers on the bottom layers and then stacking self-attention layers for the rest layers. c) Stacking both the self-attention layer and average attention layer at each layer and using their average sum to form the final hidden states (named as ‘dual attention layer’).

In the experiments, Mixed-AAN not only performs better but also shows strong diversity compared to the vanilla Transformer. With four Mixed-AAN models, we reach a better ensemble result than the result with ten models which consist of deeper and wider standard Transformer. We will further analyze the effects of different architectures from performance, diversity, and model ensemble in Sec. 5.1

## 2.6 Talking-Heads Attention

In Multi-Head Attention, the different attention heads perform separate computations, which are then summed at the end. Talking-Heads Attention (Shazeer et al., 2020) is a new variation that inserts two additional learned linear projection weights,  $W_l$  and  $W_w$ , to transform the attention-logits and the attention scores respectively, moving information across attention heads. The calculation formula is as follows:

$$Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d_k}} W_l) W_w V \quad (4)$$

We adopt this method in both encoders and decoders to improve information interaction between attention heads. This approach shows the most remarkable diversity among all the above variants with only a slight performance loss.

## 3 System Overview

In this section, we describe our system used in the WMT 2021 news shared task. We depict the overview of our NMT system in Figure 2, which can be divided into four parts, namely data filtering, large-scale synthetic data generation, in-domain finetuning, and ensemble. The synthetic data generation part further includes the generation of general domain and in-domain data. Next, we proceed to illustrate these four parts.

### 3.1 Data Filtering

We filter the bilingual training corpus with the following rules for most language pairs:

- Normalize punctuation with Moses scripts except Japanese data.
- Filter out the sentences longer than 100 words or exceed 40 characters in a single word.
- Filter out the duplicated sentence pairs.
- The word ratio between the source and the target words must not exceed 1:4 or 4:1.
- Filter out the sentences where the fast-text result does not match the origin language.
- Filter out the sentences that have invalid Unicode characters.

Besides these rules, we filter out sentence pairs in which Chinese sentence has English characters in En-Zh parallel data. The monolingual corpus is also filtered with the n-gram language model trained by the bilingual training data for each language. All the above rules are applied to synthetic parallel data.

### 3.2 General Domain Synthetic Data Generation

In this section, we describe our techniques for constructing general domain synthetic data. The general domain synthetic data is generated via large-scale back-translation, forward-translation and knowledge distillation to enhance the models’ performance for all domains. Then, we exploit

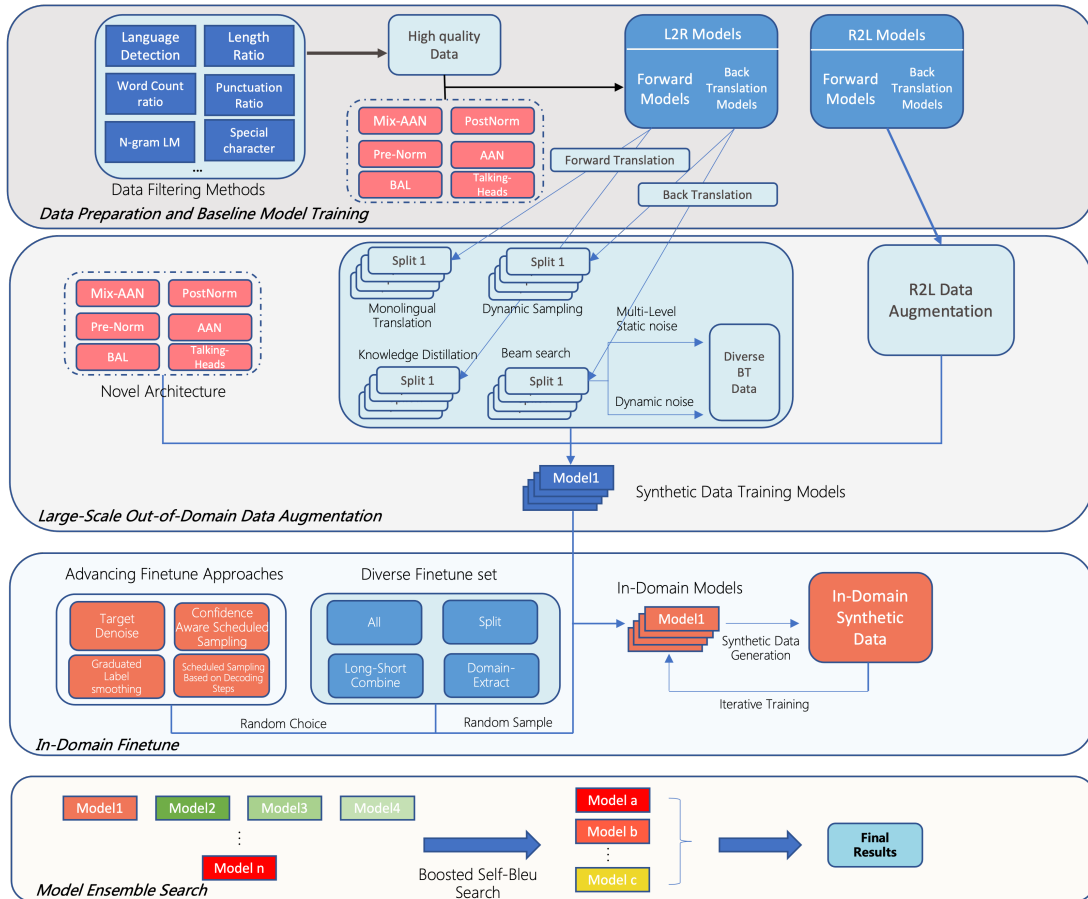


Figure 2: Architecture of our NMT system.

the iterative in-domain knowledge transfer (Meng et al., 2020) in Sec 3.3, which transfers in-domain knowledge to the vast source-side monolingual corpus, and builds our in-domain synthetic data. In the following sections, we elaborate the above techniques in detail.

### 3.2.1 Large-scale Back-Translation

Back-translation is the most commonly used data augmentation technique to incorporate the target side monolingual data into NMT (Hoang et al., 2018). Previous work (Edunov et al., 2018) has shown that different methods of generating pseudo corpus has a different influence on translation quality. Following these works, we attempt several generating strategies as follows:

- **Beam Search:** Generate target translation by beam search with beam size 5.
- **Sampling Top-K:** Select a word randomly from top-K (K is set to 10) words at each decoding step.

- **Dynamic Sampling Top-p:** Selected a word at each decoding step from the smallest set whose cumulative probability mass exceeds p and the p is dynamically changing from 0.9 to 0.95 during data generation.

Note that we also use Tagged Back-Translation (Caswell et al., 2019) in  $En \rightarrow De$  and Right-to-Left (R2L) back-translation in  $En \leftrightarrow Ja$ , as we achieve a better BLEU score after using these methods.

### 3.2.2 Knowledge Distillation

Knowledge Distillation (KD) has proven to be a powerful technique for NMT (Kim and Rush, 2016; Wang et al., 2021) to transfer knowledge from the teacher model to student models. In particular, we first use the teacher models to generate synthetic corpus in the forward direction (i.e.,  $En \rightarrow Zh$ ). Then, we use the generated corpus to train our student models.

Notably, Right-to-Left (R2L) knowledge distillation is a good complement to the Left-to-Right (L2R) way and can further improve model performance.

### 3.2.3 Forward-Translation

Using monolingual data from the source language to further enhance the performance and robustness of the model is also an effective approach. We use the ensemble model to generate high quality forward-translation data and obtain a stable improvement in En→Zh and En→De directions.

### 3.3 Iterative In-domain Knowledge Transfer

Since in-domain knowledge transfer (Meng et al., 2020) delivered a massive performance boost last year, we still use this technique in En↔Ja and En→De this year. It is not applied to En→Zh because no significant improvement is observed. We guess the reason is that the in-domain finetuning in the En→Zh direction does not bring a significant improvement compared to the other directions. And in-domain knowledge transfer is aiming at enhancing the effect of finetuning, so this does not have a noticeable effect in the English-Chinese direction.

We first use normal finetuning in Sec. 3.5 to equip our models with in-domain knowledge. Then, we ensemble these models to translate the source monolingual data into the target language. We use 4 models with different architectures and training data as our ensemble model. Next, we combine the source language sentences with the generated in-domain target language sentences as pseudo-parallel corpus. Afterwards, we retrain our models with both in-domain pseudo-parallel data and general domain synthetic data.

### 3.4 Data Augmentation

Once the pseudo-data is constructed, we further obtain diverse data by adding different noise. Compared to previous years' WMT competitions, we implement a multi-level static noise approach for our pseudo corpus:

- Token-level: Noise on every single subword after byte pair encoding.
- Word-level: Noise on every single word before byte pair encoding.
- Span-level: Noise on a continuous sequence of tokens before byte pair encoding.

The different granularities of noise make the data more diverse. The noise types are random replacement, random deletion and random permutation. We apply the three noise types in a parallel way for

each sentence. The probability of enabling each of the three operations is 0.2.

Furthermore, an on-the-fly noise approach is applied to the synthetic data. By using on-the-fly noise, the model is trained with different noises in every epoch rather than all the same along this training stage.

### 3.5 In-domain Finetuning

A domain mismatch exists between the obtained system trained with large-scale general domain data and the target test set. To alleviate this mismatch, we finetune these convergent models on small scale in-domain data, which is widely used for domain adaption (Luong and Manning, 2015; Li et al., 2019). We take the previous test sets as in-domain data and extract documents that are originally created in the source language for each translation direction (Sun et al., 2019). We also explore several advanced finetuning approaches to strengthen the effects of domain adaption and ease the exposure bias issue, which is more serious under domain shift.

**Target Denoising (Meng et al., 2020).** In the training stage, the model never sees its own errors. Thus the model trained with teacher-forcing is prone to accumulated errors in testing (Ranzato et al., 2016). To mitigate this training-generation discrepancy, we add noisy perturbations into decoder inputs when finetuning. Thus the model becomes more robust to prediction errors by target denoising. Specifically, the finetuning data generator chooses 30% of sentence pairs to add noise, and keeps the remaining 70% of sentence pairs unchanged. For a chosen pair, we keep the source sentence unchanged, and replace the  $i$ -th token of the target sentence with (1) a random token of the current target sentence 15% of the time (2) the unchanged  $i$ -th token 85% of the time.

**Graduated Label-smoothing (Wang et al., 2020).** Finetuning on a small scale in-domain data can easily lead to the over-fitting phenomenon which is harmful to the model ensemble. It generally appears as the model over confidently outputting similar words. To further preventing over-fitting of in-domain finetuning, we apply the Graduated Label-smoothing approach, which assigns a higher smoothing penalty for high-confidence predictions, during in-domain finetuning. Concretely, following the paper's setting, we set the smoothing penalty to 0.3 for tokens with confidence above 0.7,

zero for tokens with confidence below 0.3, and 0.1 for the remaining tokens.

**Confidence-Aware Scheduled Sampling.** Vanilla scheduled sampling (Zhang et al., 2019) simulates the inference scene by randomly replacing golden target input tokens with predicted ones during training. However, its critical schedule strategies are only based on training steps, ignoring the real-time model competence. To address this issue, we propose confidence-aware scheduled sampling (Liu et al., 2021a), which quantifies real-time model competence by the confidence of model predictions. At the  $t$ -th target token position, we calculate the model confidence  $conf(t)$  as follow:

$$conf(t) = P(y_t | \mathbf{y}_{<t}, \mathbf{X}, \theta) \quad (5)$$

Next, we design fine-grained schedule strategies based on the model competence. The fine-grained schedule strategy is conducted at all decoding steps simultaneously:

$$y_{t-1} = \begin{cases} y_{t-1} & \text{if } conf(t) \leq t_{golden} \\ \hat{y}_{t-1} & \text{else} \end{cases} \quad (6)$$

where  $t_{golden}$  is a threshold to measure whether  $conf(t)$  is high enough (e.g., 0.9) to sample the predicted token  $\hat{y}_{t-1}$ .

We further sample more noisy tokens at high-confidence token positions, which prevents scheduled sampling from degenerating into the teacher forcing mode.

$$y_{t-1} = \begin{cases} y_{t-1} & \text{if } conf(t) \leq t_{golden} \\ \hat{y}_{t-1} & \text{if } t_{golden} < conf(t) \leq t_{rand} \\ y_{rand} & \text{if } conf(t) > t_{rand} \end{cases} \quad (7)$$

where  $t_{rand}$  is a threshold to measure whether  $conf(t)$  is high enough (e.g., 0.95) to sample the random target token  $\hat{y}_{rand}$ .

**Scheduled Sampling Based on Decoding Steps.** We propose scheduled sampling methods based on decoding steps from the perspective of simulating the distribution of real translation errors (Liu et al., 2021b). Namely, we gradually increase the selection probability of predicted tokens with the growth of the index of decoded tokens. At the  $t$ -th decoding step, the probability of sampling golden tokens  $g(t)$  is calculated as follow:

---

### Algorithm 1 Boosted Self-BLEU based Ensemble

---

**Input:**

- List of candidate models  $M = \{m_0, \dots, m_n\}$
- Valid set BLEU for each model  $B = \{b_i, \dots, b_n\}$
- Average Self-BLEU for each model  $S = \{s_i, \dots, s_n\}$
- The number of models  $n$
- The number of ensemble models  $c$

**Output:** Model combinations  $C$

- 1: **for**  $i \leftarrow 1$  to  $n$  **do**
  - 2:    $score_i = (b_i - \min(B)) \cdot weight + (\max(S) - s_i)$
  - 3:    $weight = \frac{(\max(S) - \min(S))}{(\max(B) - \min(B))}$
  - 4: **end for**
  - 5: Add the highest score model to candidates list  $C = \{m_{top}\}$
  - 6: **while**  $|C| < c$  **do**
  - 7:    $index = \arg \min_i \frac{1}{|M-C|} \sum_{i \in M-C, j \in C} BLEU(i, j)$
  - 8:   Add  $m_{index}$  to candidate list  $C$
  - 9: **end while**
  - 10: **return**  $C$
- 

- Linear Decay:  $g(t) = \max(\epsilon, kt + b)$ , where  $\epsilon$  is the minimum value, and  $k < 0$  and  $b$  is respectively the slope and offset of the decay.
- Exponential Decay:  $g(t) = k^t$ , where  $k < 1$  is the radix to adjust the decay.
- Inverse Sigmoid Decay:  $g(t) = \frac{k}{k + e^{\frac{t}{k}}}$ , where  $e$  is the mathematical constant, and  $k \geq 1$  is a hyperparameter to adjust the decay.

Following our preliminary conclusions (Liu et al., 2021b), we choose the exponential decay and set  $k$  to 0.99 by default.

### 3.6 Boosted Self-BLEU based Ensemble (BSBE)

After we get numerous finetuned models, we need to search for the best combination for ensemble model. Ordinary random or greedy search is oversimplified to search for a good model combination and enumerate over all combinations of candidate models is inefficient. The Self-BLEU based pruning strategy (Meng et al., 2020) we proposed in last year’s competition achieve definite improvements over the ordinary ensemble.

However, diversity is not the only feature we need to consider but the performance in the valid

set is also an important metric. Therefore, we combine Self-BLEU and valid set BLEU together to derive a Boosted Self-BLEU-based Ensemble (BSBE) algorithm. Then, we apply a greedy search strategy in the top N ranked models to find the best ensemble models.

See algorithm 1 for the pseudo-code. The algorithm takes as input a list of  $n$  strong single models  $M$ , BLEU scores on valid set for each model  $B$ , average Self-BLEU scores for each model  $S$ , the number of models  $n$  and the number of ensemble models  $c$ . The algorithm return a list  $C$  consists of selected models. We calculate the weighted score for each model as line 2 in the pseudo-code. The weight calculated in line 3 is a factor to balance the scale of Self-BLEU and valid set BLEU. Then the list  $C$  initially contains the model  $m_{top}$  has a highest weighted score. Next, we iteratively re-compute the average Self-BLEU between the remaining models in  $|M - C|$  and selected models in  $C$ , based on which we select the model has minimum Self-BLEU score into  $C$ .

In our experiments, we save around 95% searching time by using this novel method to achieve the same BLEU score of the Brute Force search. We will further analyze the effect of Boosted Self-BLEU based Ensemble in section 5.2.

## 4 Experiments And Results

### 4.1 Settings

The implementation of our models is based on Fairseq<sup>1</sup> for En→Zh and EN→De, and OpenNMT<sup>2</sup> for En↔Ja. All the single models are carried out on 8 NVIDIA V100 GPUs, each of which has 32 GB memory. We use the Adam optimizer with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.998$ . The gradient accumulation is used due to the high GPU memory consumption. The batch size is set to 8192 tokens per GPU and we set the “update-freq” parameter in Fairseq to 2. The learning rate is set to 0.0005 for Fairseq and 2.0 for OpenNMT. We use warmup step = 4000. We calculate sacreBLEU<sup>3</sup> score for all experiments which is officially recommended.

### 4.2 Dataset

The statistics of all training data is shown in Table 1. For each language pair, the bilingual data is the combination of all parallel data released by

	En→Zh	En→De	En↔Ja
Bilingual Data	30.7M	74.8M	12.3M
Source Mono Data	200.5M	332.8M	210.8M
Target Mono Data	405.2M	237.9M	354.7M

Table 1: Statistics of all training data.

WMT21. For monolingual data, we select data from News Crawl, Common Crawl and Extended Common Crawl, it is then divided into several parts, each containing 50M sentences.

For general domain synthetic data, we use all the target monolingual data to generate back-translation data and a part of source monolingual data (about 80 to 100 million for different languages) to get forward translation data. For the in-domain pseudo-parallel data, we use the entire source monolingual data and bilingual data. All the test and valid data from previous years are used as in-domain data.

We use the methods described in Sec. 3.1 to filter bilingual and monolingual data.

### 4.3 Pre-processing and Post-processing

English and German sentences are segmented by Moses<sup>4</sup>, while Japanese use Mecab<sup>5</sup> for segmentation. We segment the Chinese sentences with an in-house word segmentation tool. We apply punctuation normalization in English, German and Chinese data. Truecasing is applied to English↔Japanese and English→German. We use byte pair encoding BPE (Sennrich et al., 2016b) with 32K operations for all the languages.

For the post-processing, we apply de-truecasing and de-tokenizing on the English and German translations with the scripts provided in Moses. For the Chinese translations, we transpose the punctuations to the Chinese format.

### 4.4 English→Chinese

The results of En→Zh on newstest2020 are shown in Table 2. For the En→Zh task, filtering out part of sentence pairs containing English characters in Chinese sentences shows a significant improvement in the valid set. After applying large-scale Back-Translation, we obtain +2.0 BLEU score on the baseline. We further gain +0.62 BLEU score after applying knowledge distillation and +0.24 BLEU from Forward-Translation. Surprisingly, we observe that adding more BT data from different

<sup>1</sup><https://github.com/pytorch/fairseq>

<sup>2</sup><https://github.com/OpenNMT/OpenNMT-py>

<sup>3</sup><https://github.com/mjpost/sacrebleu>

<sup>4</sup><http://www.statmt.org/moses/>

<sup>5</sup><https://github.com/taku910/mecab>

SYSTEM	En→Zh	En→Ja	Ja→En	En→De
Baseline	44.53	35.78	19.71	33.28
+ Back Translation	46.52	36.12	20.82	35.28
+ Knowledge Distillation	47.14	36.66	21.63	36.38
+ Forward Translation	47.38	–	–	36.78
+ Mix BT	48.17	37.22	22.11	–
+ <i>Finetune</i>	49.81	42.54	25.91	39.21
+ <i>Advanced Finetune</i>	50.20	–	–	39.56
+ 1st In-domain Knowledge Transfer	–	40.32	24.49	39.23
+ <i>Finetune</i>	–	43.66	26.24	–
+ <i>Advanced Finetune</i>	–	–	–	39.87
+ 2nd In-domain Knowledge Transfer	–	43.69	25.89	–
+ <i>Finetune</i>	–	44.23	26.27	–
+ <i>Advanced Finetune</i>	–	44.42	26.38	–
+ Normal Ensemble	50.57	45.11	28.01	40.42
+ BSBE	<b>50.94</b> *	<b>45.35</b> *	<b>28.24</b> *	40.59
+ Post-Process	–	–	–	<b>41.88</b> *

Table 2: Case-sensitive BLEU scores (%) on the four directions *newstest2020*, where ‘\*’ denotes the submitted system. Mix BT means we use multiple parts of Back Translation data with different generation strategies. The *Advanced Finetune* methods outperform the normal *Finetune* and we report the best results in single model. BSBE outperforms Normal Ensemble in all four directions.

shards with different generation strategy can further boost the model performance to 48.17. The finetuned model achieves a 49.81 BLEU score, which demonstrates that the domain of the training corpus is apart from the test set domain. The *advanced* finetuning further brings about 0.41 BLEU score gains compared to normal finetune. Our best single model achieves a 50.22 BLEU score.

In preliminary experiments, we select the best performing models as our ensemble combinations obtaining +0.4 BLEU score. On top of that, even after searching hundreds of models, no better results are obtained. With BSBE strategies in Sec. 3.6, a better model combination with less number of models are quickly searched, and we finally achieve 50.94 BLEU score. Our WMT2021 English→Chinese submission achieves a SacreBLEU score of 36.9, which is the highest among all submissions and chrF score of 0.337.

#### 4.5 English→Japanese

The results of En→Ja on *newstest2020* are shown in Table 2. For the En→Ja task, we filter out the sentence pairs containing Japanese characters in the English side and vice versa. The Back-Translation and Knowledge Distillation improve the baseline from 35.78 to 36.66. Adding more BT data further brings in 0.56 improvements. The improvement by finetuning is much larger than other directions, which is 5.32 BLEU. We speculate that

this is because there is less bilingual data for English and Japanese than for other languages, and the test results for Japanese are char level BLEU so this direction is more influenced by the in-domain finetuning. Two In-domain knowledge transfers improve BLEU score from 37.22 to 43.69. Normal finetune still provides 0.54 improvements after in-domain knowledge transfer. Then, we apply *advanced* finetuning methods to further get 0.19 BLEU improvements. Our final ensemble result outperforms baseline 9.57 BLEU.

#### 4.6 Japanese→English

The Ja→En task follows the same training procedure as En→Ja. From Table 2, we can observe that Back-Translation can provide 1.11 BLEU improvements from baseline. Knowledge Distillation and more BT data can improve the BLEU score from 20.82 to 22.11. The finetuning improvement is 3.8 which is slightly less than the En→Ja direction but still larger than En→Zh and En→De. We also apply two-turn in-domain knowledge transfer and further boost the BLEU score to 25.89. After normal finetuning, the BLEU score achieves 26.27. The *advanced* finetuning methods provide a slight improvement on Ja→En. After ensemble, we achieve 28.24 BLEU in *newstest2020*.



MODEL	EN-ZH	EN-JA	JA-EN	EN-DE
Transformer	49.92	44.27	26.12	39.76
Transformer with Post-Norm	49.97	-	-	-
Average Attention Transformer	49.91	44.38	26.31	39.62
Weighted Attention Transformer	49.99	-	-	39.74
Average First Transformer *	50.14	<b>44.42</b>	26.37	<b>39.87</b>
Average Bottom Transformer *	50.10	44.36	<b>26.38</b>	39.77
Dual Attention Transformer *	<b>50.20</b>	-	-	<b>39.87</b>
Talking-Heads Attention	49.89	-	-	39.70

Table 3: Case-sensitive BLEU scores (%) on the four translation directions *newstest2020* for different architecture. The model with ‘\*’ is the Mixed-AAN variants. The **bolded** scores correspond to the best single model scores in Table 2.

MODEL	Transformer	Post-Norm	AAN	Weighted	Avg-First	Self-First	Dual	TH
Transformer	100	78.12	76.02	75.08	74.47	74.02	73.51	72.63
Post-Norm	78.12	100	76.12	75.10	74.33	74.05	73.45	72.59
AAN	76.02	76.12	100	79.24	74.81	74.97	73.43	72.13
Weighted	75.08	75.10	79.24	100	74.72	74.93	73.55	72.21
Avg-First *	74.46	74.33	74.81	74.72	100	75.25	74.28	72.25
Avg-Bot *	74.02	74.05	74.97	74.93	75.25	100	74.21	72.33
Dual *	73.51	73.45	73.43	73.55	74.28	74.21	100	72.23
TH	<b>72.63</b>	<b>72.59</b>	<b>72.13</b>	<b>72.21</b>	<b>72.25</b>	<b>72.33</b>	<b>72.23</b>	100

Table 4: Self-BLEU scores (%) between different architectures. For simplicity, we refer to these models as Transformer (Pre-Norm Transformer), Post-Norm (Post-Norm Transformer), AAN (Average Attention Transformer), Weighted (Weighted Attention Transformer), Avg-First (Average First Transformer), Avg-Bot (Average Bottom Transformer), Dual (Dual Attention Transformer), TH (Talking-Heads Attention). The model with ‘\*’ is the Mixed-AAN variants.

#### 4.7 English→German

The results of En→De on *newstest2020* are shown in Table 2. After adding back-translation, we improve the BLEU score from 33.28 to 35.28. Knowledge Distillation further boosts the BLEU score to 36.58. The finetuning further brings in 2.63 improvements. After injecting the in-domain knowledge into the monolingual corpus, we get another 0.31 BLEU gain. We apply a post-processing procedure on En→De. Specifically, we normalize the English quotations to German ones in German hypotheses, which brings in 1.3 BLEU improvements.

### 5 Analysis

To verify the effectiveness of our approach, we conduct analytical experiments on model variants, finetune methods, and ensemble strategies in this section.

#### 5.1 Effects of Model Architecture

We conduct several experiments to validate the effectiveness of Transformer (Vaswani et al., 2017) variants we used and list results in Table 3. We also

investigate the diversity of different variants and the impacts on the model ensemble. The results is listed in Table 4 and Table 5. Here we take En→Zh models as examples to conduct the diversity and ensemble experiments. The results in other directions show similar trends.

**Performance.** As shown in Table 3, AAN performs slightly worse than other variants in En→Zh but Mixed-AAN variants outperform normal Transformer. Weighted Attention Transformer provides noticeable improvement compare to AAN and sometimes better than vanilla Transformer.

**Diversity.** The Self-BLEU scores in Table 4 demonstrate the difference between two models, more different models generally have lower scores. As we can see, AAN and all the variants with AAN have an absolutely lower Self-BLEU score with the Transformer. The Talking-Heads Attention has the minimum scores among all the variants.

**Ensemble.** In our preliminary experiments, we observe that more diverse models can significantly help the model ensemble. The results are listed

MODELS	newstest2020
Deeper & Wider Transformer	50.31
Weighted & Mixed-AAN	50.44
Ensemble with all models above	<b>50.62</b>

Table 5: Ensemble results with different architectures. The first row is the ensemble results with 10 deeper and wider models searched from dozens of ones. The second row is the ensemble results with only 4 Weighted Attention Transformer and Mixed-AAN models.

in Table 5. We get a more robust ensemble model with only four models using our novel variants than searching from dozens of Deeper and Wider Transformer models. Even these four models are trained with the same training data. After we combine the four models with Deeper and Wider Transformer, we can further get a significant improvement.

Take En→Zh as an example, our final submission consist of 1 Average First Transformer, 1 Average Bottom Transformer, 1 Dual Attention Transformer, 1 Weighted Attention Transformer and 1 Transformer with Post-Norm.

## 5.2 Effects of Boosted Self-BLEU based Ensemble

To verify the superiority of our Boosted Self-BLEU based Ensemble (BSBE) method, we randomly select 10 models with different architecture and training data. For our submitted system, we search from over 500 models. We use a greedy search algorithm (Deng et al., 2018) as our baseline. The greedy search greedily selects the best performance model into candidate ensemble models. If the selected model provides a positive improvement, we keep it in the candidates. Otherwise, it is added to a temporary model list and still has a weak chance to be reused in the future. One model from the temporary list can be reused once, after which it is withdrawn definitely. We compare the results of greedy search, BSBE and Brute Force and list the ensemble model BLEU and the number of searches in Table 6. Note that  $n$  is the number of models, which is 10 here. For BSBE, we need to get the translation result of every model to calculate the Self-BLEU. After that, we only need to perform the inference process once.

## 5.3 Effects of Advanced Finetuning

In this section, we describe our experiments on advanced finetuning in the four translation directions. As shown in Table 7, all the advanced fine-

ALGORITHM	BLEU	Number of Searches
Greedy	50.19	$2n$
Brute Force	50.44	$\sum_{i=1}^n C_n^i$
BSBE	<b>50.44</b>	$n + 1$

Table 6: Results of different search algorithm.  $n$  is the total number of models used for the search. The number of searches is number that the methods need to translate the valid set. Our BSBE achieves comparable BLEU score as Brute Force search and significantly reduces the searching time.

tuning methods outperform normal finetuning. For En→Zh, Schedule Sampling Based on Decoding Steps with Graduated Label Smoothing improves the model performance from 49.81 to 50.20. For En↔Ja, Target Denoising with Graduated Label Smoothing provides the highest BLEU gain, which are 0.19 and 0.11. For the En→De direction, Confidence-Aware Schedule Sampling with Graduated Label Smoothing performs the best, improving from 39.21 to 39.42. These findings are in line with the conclusion of Wang and Sennrich (2020) that links exposure bias with domain shift.

## 6 Conclusion

We investigate various novel Transformer based architectures to build robust systems. Our systems are also built on several popular data augmentation methods such as back-translation, knowledge distillation and iterative in-domain knowledge transfer. We enhance our system with advanced finetuning approaches, i.e., target denoising, graduated label smoothing and confidence-aware scheduled sampling. A boosted Self-BLEU based model ensemble is also employed which plays a key role in our systems. Our constrained systems achieve 36.9, 46.9, 27.8 and 31.3 case-sensitive BLEU scores on English→Chinese, English→Japanese, Japanese→English and English→German, respectively. The BLEU scores of English→Chinese, English→Japanese and Japanese→English are the highest among all submissions, and that of English→German is the highest among all constrained submissions.

## Acknowledgements

Yijin Liu and Jinan Xu have been supported by the National Key R&D Program of China (2020AAA0108001) and the National Nature Science Foundation of China (No. 61976015, 61976016, 61876198 and 61370130). The authors

FINETUNING APPROACH	EN-ZH	EN-JA	JA-EN	EN-DE
Normal	49.81	44.23	26.27	39.21
Graduated Label Smoothing	49.95	44.32	26.35	39.32
+ <i>Target Denoising</i>	50.09	<b>44.42</b>	<b>26.38</b>	39.34
+ <i>Confidence-Aware Schedule Sampling</i>	50.17	44.35	26.33	<b>39.42</b>
+ <i>Schedule Sampling Based on Decoding Steps</i>	<b>50.20</b>	44.36	26.33	39.40

Table 7: Case-sensitive BLEU scores (%) on the four translation directions *newstest2020* for different finetuning approaches. We report the highest score and bold the best result among different finetuning approaches.

would like to thank the anonymous reviewers for their valuable comments and suggestions to improve this paper.

## References

- Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. [Tagged back-translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 53–63, Florence, Italy. Association for Computational Linguistics.
- Yongchao Deng, Shanbo Cheng, Jun Lu, Kai Song, Jingang Wang, Shenglan Wu, Liang Yao, Guchun Zhang, Haibo Zhang, Pei Zhang, Changfeng Zhu, and Boxing Chen. 2018. [Alibaba’s neural machine translation systems for WMT18](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 368–376, Belgium, Brussels. Association for Computational Linguistics.
- Daniel Duckworth, Arvind Neelakantan, Ben Goodrich, Lukasz Kaiser, and Samy Bengio. 2019. [Parallel scheduled sampling](#). *arXiv preprint arXiv:1906.04331*.
- 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. Association for Computational Linguistics.
- 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, Melbourne, Australia. Association for Computational Linguistics.
- Xiao Shi Huang, Felipe Pérez, Jimmy Ba, and Maksims Volkovs. 2020. [Improving transformer optimization through better initialization](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4475–4483. PMLR.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- 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.
- Xiaodong Liu, Kevin Duh, Liyuan Liu, and Jianfeng Gao. 2020. [Very deep transformers for neural machine translation](#). *arXiv preprint arXiv:2008.07772*.
- Yijin Liu, Fandong Meng, Yufeng Chen, Jinan Xu, and Jie Zhou. 2021a. [Confidence-aware scheduled sampling for neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2327–2337.
- Yijin Liu, Fandong Meng, Yufeng Chen, Jinan Xu, and Jie Zhou. 2021b. [Scheduled sampling based on decoding steps for neural machine translation](#). In *Proceedings of EMNLP*.
- 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. Association for Computational Linguistics.
- Fandong Meng and Jinchao Zhang. 2019. [DTMT: A novel deep transition architecture for neural machine translation](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 224–231. AAAI Press.

- Tsvetomila Mihaylova and André F. T. Martins. 2019. [Scheduled sampling for transformers](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 351–356, Florence, Italy. Association for Computational Linguistics.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. [Sequence level training with recurrent neural networks](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- 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.
- Noam Shazeer, Zhenzhong Lan, Youlong Cheng, Nan Ding, and Le Hou. 2020. [Talking-heads attention](#). *arXiv preprint arXiv:2003.02436*.
- 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, Florence, Italy. Association for Computational Linguistics.
- 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 *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Chaojun Wang and Rico Sennrich. 2020. [On exposure bias, hallucination and domain shift in neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3544–3552, Online. Association for Computational Linguistics.
- 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.
- Shuo Wang, Zhaopeng Tu, Shuming Shi, and Yang Liu. 2020. [On the inference calibration of neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3070–3079, Online. Association for Computational Linguistics.
- Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. 2020. [On layer normalization in the transformer architecture](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 10524–10533. PMLR.
- Jianhao Yan, Fandong Meng, and Jie Zhou. 2020. [Multi-unit transformers for neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1047–1059, 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. Association for Computational Linguistics.
- Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. [Bridging the gap between training and inference for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4334–4343, Florence, Italy. Association for Computational Linguistics.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Txygen: A benchmarking platform for text generation models](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, pages 1097–1100. ACM.