

One Reference Is Not Enough: Diverse Distillation with Reference Selection for Non-Autoregressive Translation

Chenze Shao^{1,2}, Xuanfu Wu^{1,2}, Yang Feng^{1,2*}

¹ Key Laboratory of Intelligent Information Processing
Institute of Computing Technology, Chinese Academy of Sciences (ICT/CAS)

² University of Chinese Academy of Sciences
{shaochenze18z, wuxuanfu20s, fengyang}@ict.ac.cn

Abstract

Non-autoregressive neural machine translation (NAT) suffers from the multi-modality problem: the source sentence may have multiple correct translations, but the loss function is calculated only according to the reference sentence. Sequence-level knowledge distillation makes the target more deterministic by replacing the target with the output from an autoregressive model. However, the multi-modality problem in the distilled dataset is still nonnegligible. Furthermore, learning from a specific teacher limits the upper bound of the model capability, restricting the potential of NAT models. In this paper, we argue that one reference is not enough and propose diverse distillation with reference selection (DDRS) for NAT. Specifically, we first propose a method called SeedDiv for diverse machine translation, which enables us to generate a dataset containing multiple high-quality reference translations for each source sentence. During the training, we compare the NAT output with all references and select the one that best fits the NAT output to train the model. Experiments on widely-used machine translation benchmarks demonstrate the effectiveness of DDRS, which achieves 29.82 BLEU with only one decoding pass on WMT14 En-De, improving the state-of-the-art performance for NAT by over 1 BLEU.¹

1 Introduction

Non-autoregressive machine translation (Gu et al., 2018) has received increasing attention in the field of neural machine translation for the property of parallel decoding. Despite the significant speedup, NAT suffers from the performance degradation compared to autoregressive models (Bahdanau et al., 2015; Vaswani et al., 2017) due to the multi-modality problem: the source sentence may have

multiple correct translations, but the loss is calculated only according to the reference sentence. The multi-modality problem will cause the inaccuracy of the loss function since NAT has no prior knowledge about the reference sentence during the generation, where the teacher forcing algorithm (Williams and Zipser, 1989) makes autoregressive models less affected by feeding the golden context.

How to overcome the multi-modality problem has been a central focus in recent efforts for improving NAT models (Shao et al., 2019, 2020, 2021; Ran et al., 2020; Sun and Yang, 2020; Ghazvininejad et al., 2020; Du et al., 2021). A standard approach is to use sequence-level knowledge distillation (Kim and Rush, 2016), which attacks the multi-modality problem by replacing the target-side of the training set with the output from an autoregressive model. The distilled dataset is less complex and more deterministic (Zhou et al., 2020), which becomes a default configuration of NAT. However, the multi-modality problem in the distilled dataset is still nonnegligible (Zhou et al., 2020). Furthermore, the distillation requires NAT models to imitate the behavior of a specific autoregressive teacher, which limits the upper bound of the model capability and restricts the potential of developing stronger NAT models.

In this paper, we argue that one reference is not enough and propose diverse distillation with reference selection (DDRS) for NAT. Diverse distillation generates a dataset containing multiple reference translations for each source sentence, and reference selection finds the reference translation that best fits the model output for the training. As illustrated in Figure 1, diverse distillation provides candidate references “I must leave tomorrow” and “Tomorrow I must leave”, and reference selection selects the former which fits better with the model output. More importantly, NAT with DDRS does not imitate the behavior of a specific teacher but learns selectively from multiple references, which

*Corresponding author: Yang Feng

¹Source code: <https://github.com/ictnlp/DDRS-NAT>.




Figure 1: Illustration of diverse distillation and reference selection. Diverse distillation provides multiple references and reference selection selects the one that best fits the model output for the training.

improves the upper bound of the model capability and allows for developing stronger NAT models.

The object of diverse distillation is similar to the task of diverse machine translation, which aims to generate diverse translations with high translation quality (Li et al., 2016; Vijayakumar et al., 2018; Shen et al., 2019; Wu et al., 2020; Li et al., 2021). We propose a simple yet effective method called SeedDiv, which directly uses the randomness in model training controlled by random seeds to produce diverse reference translations without losing translation quality. For reference selection, we compare the model output with all references and select the one that best fits the model output, which can be efficiently conducted without extra neural computations. The model learns from all references indiscriminately in the beginning, and gradually focuses more on the selected reference that provides accurate training signals for the model. We also extend the reference selection approach to reinforcement learning, where we encourage the model to move towards the selected reference that gives the maximum reward to the model output.

We conduct experiments on widely-used machine translation benchmarks to demonstrate the effectiveness of our method. On the competitive task WMT14 En-De, DDRS achieves 27.60 BLEU with $14.7\times$ speedup and 28.33 BLEU with $5.0\times$ speedup, outperforming the autoregressive Transformer while maintaining considerable speedup. When using the larger version of Transformer, DDRS even achieves 29.82 BLEU with only one decoding pass, improving the state-of-the-art performance level for NAT by over 1 BLEU.

2 Background

2.1 Non-Autoregressive Translation

Gu et al. (2018) proposes non-autoregressive ma-

chine translation to reduce the translation latency through parallel decoding. The vanilla-NAT models the translation probability from the source sentence \mathbf{x} to the target sentence $\mathbf{y} = \{y_1, \dots, y_T\}$ as:

$$p(\mathbf{y}|\mathbf{x}, \theta) = \prod_{t=1}^T p_t(y_t|\mathbf{x}, \theta), \quad (1)$$

where θ is a set of model parameters and $p_t(y_t|\mathbf{x}, \theta)$ is the translation probability of word y_t in position t . The vanilla-NAT is trained to minimize the cross-entropy loss:

$$\mathcal{L}_{CE}(\theta) = - \sum_{t=1}^T \log(p_t(y_t|\mathbf{x}, \theta)). \quad (2)$$

The vanilla-NAT has to know the target length before constructing the decoder inputs. The target length T is set as the reference length during the training and obtained from a length predictor during the inference. The target length cannot be changed dynamically during the inference, so it often requires generating multiple candidates with different lengths and re-scoring them to produce the final translation (Gu et al., 2018).

The length issue can be overcome by connectionist temporal classification (CTC, Graves et al., 2006). CTC-based models usually generate a long alignment containing repetitions and blank tokens. The alignment will be post-processed by a collapsing function Γ^{-1} to recover a normal sentence, which first collapses consecutive repeated tokens and then removes all blank tokens. CTC is capable of efficiently finding all alignments \mathbf{a} which the reference sentence \mathbf{y} can be recovered from, and marginalizing the log-likelihood with dynamic programming:

$$\log p(\mathbf{y}|\mathbf{x}, \theta) = \log \sum_{\mathbf{a} \in \Gamma(\mathbf{y})} p(\mathbf{a}|\mathbf{x}, \theta). \quad (3)$$

Due to the superior performance and the flexibility of generating predictions with variable length, CTC is receiving increasing attention in non-autoregressive translation (Libovický and Helcl, 2018; Kasner et al., 2020; Saharia et al., 2020; Gu and Kong, 2020; Zheng et al., 2021).

2.2 Sequence-Level Knowledge Distillation

Sequence-level Knowledge Distillation (SeqKD, Kim and Rush, 2016) is a widely used knowledge distillation method in NMT, which trains the student model to mimic teacher’s actions at sequence-level. Given the student prediction p and the teacher prediction q , the distillation loss is:

$$\begin{aligned} \mathcal{L}_{SeqKD}(\theta) &= - \sum_{\mathbf{y}} q(\mathbf{y}|\mathbf{x}) \log p(\mathbf{y}|\mathbf{x}, \theta) \\ &\approx - \log p(\hat{\mathbf{y}}|\mathbf{x}, \theta), \end{aligned} \quad (4)$$

where θ are parameters of the student model and $\hat{\mathbf{y}}$ is the output from running beam search with the teacher model. The teacher output $\hat{\mathbf{y}}$ is used to approximate the teacher distribution otherwise the distillation loss will be intractable.

The procedure of sequence-level knowledge distillation is: (1) train a teacher model, (2) run beam search over the training set with this model, (3) train the student model with cross-entropy on the source sentence and teacher translation pairs. The distilled dataset is less complex and more deterministic (Zhou et al., 2020), which helps to alleviate the multi-modality problem and becomes a default configuration in NAT models.

2.3 Diverse Machine Translation

The task of diverse machine translation requires to generate diverse translations and meanwhile maintain high translation quality. Assume the reference sentence is \mathbf{y} and we have multiple translations $\{\mathbf{y}_1, \dots, \mathbf{y}_k\}$, the translation quality is measured by the average reference BLEU (rfb):

$$\text{rfb} = \frac{1}{k} \sum_{i=1}^k \text{BLEU}(\mathbf{y}, \mathbf{y}_i), \quad (5)$$

and the translation diversity is measured by the average pairwise BLEU (pwb):

$$\text{pwb} = \frac{1}{(k-1)k} \sum_{i=1}^k \sum_{j \neq i} \text{BLEU}(\mathbf{y}_i, \mathbf{y}_j). \quad (6)$$

Higher reference BLEU indicates better translation quality and lower pairwise BLEU indicates

better translation diversity. Generally speaking, there is a trade-off between quality and diversity. In existing methods, translation diversity has to be achieved at the cost of losing translation quality.

3 Approach

In this section, we first introduce the diverse distillation technique we use to generate multiple reference translations for each source sentence, and then apply reference selection to select the reference that best fits the model output for the training.

3.1 Diverse Distillation

The objective of diverse distillation is to obtain a dataset containing multiple high-quality references for each source sentence, which is similar to the task of diverse machine translation that aims to generate diverse translations with high translation quality. However, the translation diversity is achieved at a certain cost of translation quality in previous work, which is not desired in diverse distillation.

Using the randomness in model training, we propose a simple yet effective method called SeedDiv to achieve translation diversity without losing translation quality. Specifically, given the desired number of translations k , we directly set k different random seeds to train k translation models, where random seeds control the random factors during the model training such as parameter initialization, batch order, and dropout. During the decoding, each model translates the source sentence with beam search, which gives k different translations in total. Notably, SeedDiv does not sacrifice the translation quality to achieve diversity since random seeds do not affect the expected model performance.

We conduct the experiment on WMT14 En-De to evaluate the performance of SeedDiv. We use the base setting of Transformer and train the model for 150K steps. The detailed configuration is described in section 4.1. We also re-implement several existing methods with the same setting for comparison, including Beam Search, Diverse Beam Search (Vijayakumar et al., 2018), HardMoE (Shen et al., 2019), Head Sampling (Sun et al., 2020) and Concrete Dropout (Wu et al., 2020). We set the number of translations $k=3$, and set the number of heads to be sampled as 3 for head sampling. We also implement a weaker version of our method SeedDiv-ES, which early stops the training process with only $\frac{1}{k}$ of total training steps. We report the results of




Figure 2: Reference BLEU and pairwise BLEU scores of SeedDiv and other diverse translation methods on the test set of WMT14 En-De. We do not use compound split to keep consistency with previous work.

these methods in Figure 2.

It is surprising to see that SeedDiv achieves outstanding translation diversity besides the superior translation quality, which outperforms most methods on both translation quality and diversity. Only HardMoE has a better pairwise BLEU than SeedDiv, but its reference BLEU is much lower. The only concern is that SeedDiv requires a larger training cost to train multiple models, so we also use a weaker version SeedDiv-ES for comparison. Though the model performance is degraded due to the early stop, SeedDiv-ES still achieves a good trade-off between the translation quality and diversity, demonstrating the advantage of using the training randomness controlled random seeds to generate diverse translations. Therefore, we use SeedDiv as the technique for diverse distillation.

3.2 Reference Selection

3.2.1 Losses under Diverse Distillation

After diverse distillation, we obtain a dataset containing k reference sentences $\mathbf{y}_{1:k}$ for each source sentence \mathbf{x} . Traditional data augmentation algorithms for NMT (Sennrich et al., 2016a; Zhang and Zong, 2016; Zhou and Keung, 2020; Nguyen et al., 2020) generally calculate cross-entropy losses on all data and use their summation to train the model:

$$\mathcal{L}_{sum}(\theta) = -\frac{1}{k} \sum_{i=1}^k \log p(\mathbf{y}_i | \mathbf{x}, \theta). \quad (7)$$

However, this loss function is inaccurate for NAT due to the increase of data complexity. Sequence-level knowledge distillation works well on NAT by

reducing the complexity of target data (Zhou et al., 2020). In comparison, the target data generated by diverse distillation is relatively more complex. If NAT learns from the k references indiscriminately, it will not eventually converge to any one reference but generate a mixture of all references.

Using the multi-reference dataset, we propose to train NAT with reference selection to evaluate the model output with better accuracy. We compare the model output with all reference sentences and select the one with the maximum probability assigned by the model. We train the model with only the selected reference:

$$\mathcal{L}_{max}(\theta) = -\log \max_{1 \leq i \leq k} p(\mathbf{y}_i | \mathbf{x}, \theta). \quad (8)$$

In this way, we do not fit the model to all references but only encourage it to generate the nearest reference, which is an easier but more suitable objective for the model. Besides, when the ability of autoregressive teacher is limited, the NAT model can learn to ignore bad references in the data and select the clean reference for the training, which makes the capability of NAT not limited by a specific autoregressive teacher.

In addition to minimizing all losses $\mathcal{L}_{sum}(\theta)$ or the selected loss $\mathcal{L}_{max}(\theta)$, there is also an intermediate choice to assign different weights to reference sentences. We can optimize the log-likelihood of generating any reference sentence as follows:

$$\mathcal{L}_{mid}(\theta) = -\log \sum_{i=1}^k p(\mathbf{y}_i | \mathbf{x}, \theta). \quad (9)$$

The gradient of Equation 9 is equivalent to assigning weight $\frac{p(\mathbf{y}_i | \mathbf{x}, \theta)}{\sum_i p(\mathbf{y}_i | \mathbf{x}, \theta)}$ to the cross-entropy loss of each reference sentence \mathbf{y}_i . In this way, the model focuses more on suitable references but also assigns non-zero weights to other references.

We use a linear annealing schedule with two stages to train the NAT model. In the first stage, we begin with the summation $\mathcal{L}_{sum}(\theta)$ and linearly anneal the loss to $\mathcal{L}_{mid}(\theta)$. Similarly, we linearly switch to the selected loss $\mathcal{L}_{max}(\theta)$ in the second stage. We use t and T to denote the current time step and total training steps respectively, and use a constant λ to denote the length of the first stage. The loss function is:

$$\mathcal{L}(\theta) = \begin{cases} T_1 \mathcal{L}_{mid}(\theta) + (1 - T_1) \mathcal{L}_{sum}(\theta), & t \leq \lambda T \\ T_2 \mathcal{L}_{max}(\theta) + (1 - T_2) \mathcal{L}_{mid}(\theta), & t > \lambda T \end{cases} \quad (10)$$

Models	$\mathcal{L}_{max}(\theta)$		$\mathcal{L}_{sum}(\theta)$	
	For	Back	For	Back
AT	$k \times$	$1 \times$	$k \times$	$k \times$
vanilla-NAT	$k \times$	$1 \times$	$k \times$	$k \times$
CTC	$1 \times$	$1 \times$	$1 \times$	$1 \times$

Table 1: The calculation cost of $\mathcal{L}_{max}(\theta)$ and $\mathcal{L}_{sum}(\theta)$ for different models. ‘For’ and ‘Back’ indicate forward and backward propagations respectively.

where T_1 and T_2 are defined as:

$$T_1 = \frac{t}{\lambda T}, \quad T_2 = \frac{t - \lambda T}{T - \lambda T}. \quad (11)$$

In this way, the model learns from all references indiscriminately at the beginning, which serves as a pretraining stage that provides comprehensive knowledge to the model. As the training progresses, the model focuses more on the selected reference, which provides accurate training signals and gradually finetunes the model to the optimal state.

3.2.2 Efficient Calculation with CTC

To calculate the probability $p(\mathbf{y}|\mathbf{x}, \theta)$, the vanilla-NAT must set the decoder length to the length of \mathbf{y} . Therefore, calculating the probability of k reference sentences requires running the decoder for at most k times, which will greatly increase the training cost. Fortunately, for CTC-based NAT, the training cost is nearly the same since its decoder length is only determined by the source sentence. We only need to run the model once and calculate the probabilities of the k reference sentences with dynamic programming, which has a minor cost compared with forward and backward propagations. In Table 1, we show the calculation cost of $\mathcal{L}_{max}(\theta)$ and $\mathcal{L}_{sum}(\theta)$ for different models. We use CTC as the baseline model due to its superior performance and training efficiency.

3.2.3 Max-Reward Reinforcement Learning

Following Shao et al. (2019, 2021), we finetune the NAT model with the reinforcement learning objective (Williams, 1992; Ranzato et al., 2015):

$$\mathcal{L}_{rl}(\theta) = \mathbb{E}_{\mathbf{y}}[\log p(\mathbf{y}|\mathbf{x}, \theta) \cdot r(\mathbf{y})], \quad (12)$$

where $r(\mathbf{y})$ is the reward function and will be discussed later. The usual practice is to sample a sentence \mathbf{y} from the distribution $p(\mathbf{y}|\mathbf{x}, \theta)$ to estimate the above equation. For CTC based NAT, $p(\mathbf{y}|\mathbf{x}, \theta)$ cannot be directly sampled, so we sample from

the equivalent distribution $p(\mathbf{a}|\mathbf{x}, \theta)$ instead. We recover the target sentence by the collapsing function Γ^{-1} and calculate its probability with dynamic programming to estimate the following equation:

$$\mathcal{L}_{rl}(\theta) = \mathbb{E}_{\mathbf{a}}[\log p(\Gamma^{-1}(\mathbf{a})|\mathbf{x}, \theta) \cdot r(\Gamma^{-1}(\mathbf{a}))]. \quad (13)$$

The reward function is usually evaluation metrics for machine translation (e.g., BLEU, GLEU), which evaluate the prediction by comparing it with the reference sentence. We use $r(\mathbf{y}_1, \mathbf{y}_2)$ to denote the reward of prediction \mathbf{y}_1 when \mathbf{y}_2 is the reference. As we have k references $\mathbf{y}_{1:k}$, we define our reward function to be the maximum reward:

$$r(\mathbf{y}) = \max_{1 \leq i \leq k} r(\mathbf{y}, \mathbf{y}_i). \quad (14)$$

By optimizing the maximum reward, we encourage the model to move towards the selected reference, which is the closest to the model. Otherwise, rewards provided by other references may mislead the model to generate a mixture of all references.

4 Experiments

4.1 Experimental Settings

Datasets We conduct experiments on major benchmark datasets for NAT: WMT14 English \leftrightarrow German (En \leftrightarrow De, 4.5M sentence pairs) and WMT16 English \leftrightarrow Romanian (En \leftrightarrow Ro, 0.6M sentence pairs). We also evaluate our approach on a large-scale dataset WMT14 English \rightarrow French (En \rightarrow Fr, 23.7M sentence pairs) and a small-scale dataset IWSLT14 German \rightarrow English (De \rightarrow En, 160K sentence pairs). The datasets are tokenized into subword units using a joint BPE model (Sennrich et al., 2016b). We use BLEU (Papineni et al., 2002) to evaluate the translation quality.

Hyperparameters We use 3 teachers for diverse distillation and set the seed to i when training the i -th teacher. We set the first stage length λ to $2/3$. We use sentence-level BLEU as the reward. We adopt Transformer-base (Vaswani et al., 2017) as our autoregressive baseline as well as the teacher model. The NAT model shares the same architecture as Transformer-base. We uniformly copy encoder outputs to construct decoder inputs, where the length of decoder inputs is $3 \times$ as long as the source length. All models are optimized with Adam (Kingma and Ba, 2014) with $\beta = (0.9, 0.98)$ and $\epsilon = 10^{-8}$, and each batch contains approximately 32K source words. On WMT14 En \leftrightarrow De and WMT14 En \rightarrow Fr,

Models		Iterations	Speedup	WMT14		WMT16	
				EN-DE	DE-EN	EN-RO	RO-EN
AT	Transformer (Vaswani et al., 2017)	N	1.0×	27.51	31.52	34.39	33.76
	+ distillation ($k=3$)	N	1.0×	28.04	32.17	35.10	34.83
One-pass NAT	NAT-FT (Gu et al., 2018)	1	15.6×	17.69	21.47	27.29	29.06
	CTC (Libovický and Helcl, 2018)	1	–	16.56	18.64	19.54	24.67
	NAT-REG (Wang et al., 2019)	1	27.6×	20.65	24.77	–	–
	Bag-of-ngrams (Shao et al., 2020)	1	10.7×	20.90	24.61	28.31	29.29
	AXE (Ghazvininejad et al., 2020)	1	–	23.53	27.90	30.75	31.54
	SNAT (Liu et al., 2021)	1	22.6×	24.64	28.42	32.87	32.21
	GLAT (Qian et al., 2021)	1	15.3×	25.21	29.84	31.19	32.04
	Seq-NAT (Shao et al., 2021)	1	15.6×	25.54	29.91	31.69	31.78
	CNAT (Bao et al., 2021)	1	10.37×	25.56	29.36	–	–
	Imputer (Saharia et al., 2020)	1	–	25.80	28.40	32.30	31.70
	OAXE (Du et al., 2021)	1	–	26.10	30.20	32.40	33.30
	AligNART (Song et al., 2021)	1	13.4×	26.40	30.40	32.50	33.10
	REDER (Zheng et al., 2021)	1	15.5×	26.70	30.68	33.10	33.23
	CTC w/ DSLP&MT (Huang et al., 2021)	1	14.8×	27.02	31.61	34.17	34.60
	Fully-NAT (Gu and Kong, 2020)	1	16.8×	27.20	31.39	33.71	34.16
	REDER + beam20 + AT reranking	1	5.5×	27.36	31.10	33.60	34.03
Iterative NAT	iNAT (Lee et al., 2018)	10	2.0×	21.61	25.48	29.32	30.19
	CMLM (Ghazvininejad et al., 2019)	10	–	27.03	30.53	33.08	33.31
	RecoverSAT (Ran et al., 2020)	N/2	2.1×	27.11	31.67	32.92	33.19
	LevT (Gu et al., 2019)	2.05	4.0×	27.27	–	–	33.26
	DisCO (Kasai et al., 2020)	4.82	–	27.34	31.31	33.22	33.25
	JM-NAT (Guo et al., 2020b)	10	–	27.69	32.24	33.52	33.72
	RewriteNAT (Geng et al., 2021)	2.70	–	27.83	31.52	33.63	34.09
	Imputer (Saharia et al., 2020)	8	–	28.20	31.80	34.40	34.10
Our work	CTC	1	14.7×	26.09	29.50	33.55	32.98
	CTC + distillation ($k=3$)	1	14.7×	26.35	29.73	33.51	32.82
	DDRS w/o RL	1	14.7×	27.18	30.91	34.42	34.31
	DDRS	1	14.7×	27.60	31.48	34.60	34.65
	DDRS + beam20 + 4-gram LM	1	5.0×	28.33	32.43	35.42	35.81

Table 2: Performance comparison between our models and existing methods. The speedup is measured on WMT14 En-De test set. N denotes the length of translation. k means ensemble distillation (Freitag et al., 2017) from an ensemble of k AT models. – means not reported.

we train AT for 150K steps and train NAT for 300K steps with dropout 0.2. On WMT16 En \leftrightarrow Ro and IWSLT14 De \rightarrow En, we train AT for 18K steps and train NAT for 150K steps with dropout 0.3. We fine-tune NAT for 3K steps. The learning rate warms up to $5 \cdot 10^{-4}$ within 10K steps in pretraining and warms up to $2 \cdot 10^{-5}$ within 500 steps in RL fine-tuning, and then decays with the inverse square-root schedule. We average the last 5 checkpoints to obtain the final model. We use GeForce RTX 3090 GPU for the training and inference. We implement our models based on the open-source framework of fairseq (Ott et al., 2019).

Knowledge Distillation For baseline NAT models, we follow previous works on NAT to apply sequence-level knowledge distillation (Kim and Rush, 2016) to make the target more deterministic. Our method applies diverse distillation with $k = 3$ by default, that is, we use SeedDiv to generate 3 reference sentences for each source sentence.

Beam Search Decoding For autoregressive models, we use beam search with beam width 5 for the inference. For NAT, the most straightforward way is to generate the sequence with the highest probability at each position. Furthermore, CTC-based models also support beam search decoding optionally combined with n-gram language models (Kasner et al., 2020). Following Gu and Kong (2020), we use beam width 20 combined with a 4-gram language model to search the target sentence, which can be implemented efficiently in C++².

4.2 Main Results

We compare the performance of DDRS and existing methods in Table 2. Compared with the competitive CTC baseline, DDRS achieves a strong improvement of more than 1.5 BLEU on average, demonstrating the effectiveness of diverse distillation and reference selection. Compared with exist-

²<https://github.com/parlance/ctcdecode>

Models	BLEU	Speedup
Transformer-base	27.51	1.0×
Transformer-big (teacher)	28.64	0.9×
R2L-Transformer-big (teacher)	27.96	0.9×
DDRS-base	27.99	14.7×
+beam20&lm	28.95	5.0×
DDRS-big	28.84	14.1×
+beam20&lm	29.82	4.8×

Table 3: Performance of DDRS on the test set of WMT14 En-De with Transformer-big for distillation.

Models	En-Fr	De-En
Transformer	40.15	34.17
CTC	38.40	31.37
DDRS	39.91	33.12
DDRS +beam20&lm	40.59	34.74

Table 4: Performance of DDRS on the test sets of WMT14 En-Fr and IWSLT14 De-En.

ing methods, DDRS beats the state-of-the-art for one-pass NAT on all benchmarks and beats the autoregressive Transformer on most benchmarks with 14.7× speedup over it. The performance of DDRS is further boosted by beam search and 4-gram language model, which even outperforms all iterative NAT models with only one-pass decoding. Notably, on WMT16 En↔Ro, our method improves state-of-the-art performance levels for NAT by over 1 BLEU. Compared with autoregressive models, our method outperforms the Transformer with knowledge distillation, and meanwhile maintains 5.0× speedup over it.

We further explore the capability of DDRS with a larger model size and stronger teacher models. We use the big version of Transformer for distillation, and also add 3 right-to-left (R2L) teachers to enrich the references. We respectively use Transformer-base and Transformer-big as the NAT architecture and report the performance of DDRS in Table 3. Surprisingly, the performance of DDRS can be further greatly boosted by using a larger model size and stronger teachers. DDRS-big with beam search achieves 29.82 BLEU on WMT14 En-De, which is close to the state-of-the-art performance of autoregressive models on this competitive dataset and improves the state-of-the-art performance for NAT by over 1 BLEU with only one-pass decoding.

We also evaluate our approach on a large-scale dataset WMT14 En-Fr and a small-scale dataset IWSLT14 De-En. Table 4 shows that DDRS still

\mathcal{L}_{sum}	\mathcal{L}_{mid}	\mathcal{L}_{max}	λ	Reward	BLEU ₁	BLEU ₂
✓					24.61	25.97
	✓				25.23	26.90
		✓			25.31	26.88
	✓	✓	0		25.41	26.99
✓	✓	✓	1/3		25.48	27.13
✓	✓	✓	2/3		25.59	27.18
✓	✓		1		25.45	27.09
✓	✓	✓	2/3	<i>random</i>	25.63	27.26
✓	✓	✓	2/3	<i>average</i>	25.79	27.51
✓	✓	✓	2/3	<i>maximum</i>	25.92	27.60

Table 5: Ablation study on WMT14 En-De with different combinations of techniques. BLEU₁ is the BLEU score on validation set. BLEU₂ is the BLEU score on test set. The validation performance of CTC baseline is 24.57 BLEU. λ is the length of the first training stage. *random* means the reward of a random reference, *average* means the average reward, and *maximum* means the maximum reward among all references.

achieves considerable improvements over the CTC baseline and DDRS with beam search can outperform the autoregressive Transformer.

4.3 Ablation Study

In Table 5, we conduct an ablation study to analyze the effect of techniques used in DDRS. First, we separately use the loss functions defined in Equation 7, Equation 8 and Equation 9 to train the model. The summation loss $\mathcal{L}_{sum}(\theta)$ has a similar performance to the CTC baseline, showing that simply using multiple references is not helpful for NAT due to the increase of data complexity. The other two losses $\mathcal{L}_{mid}(\theta)$ and $\mathcal{L}_{max}(\theta)$ achieve considerable improvements to the CTC baseline, demonstrating the effectiveness of reference selection.

Then we use different λ to verify the effect of the annealing schedule. With the annealing schedule, the loss is a combination of the three losses but performs better than each of them. Though the summation loss $\mathcal{L}_{sum}(\theta)$ does not perform well when used separately, it can play the role of pre-training and improve the final performance. When λ is 2/3, the annealing schedule performs the best and improves $\mathcal{L}_{max}(\theta)$ by about 0.3 BLEU.

Finally, we verify the effect of the reward function during the fine-tuning. When choosing a random reference to calculate the reward, the fine-tuning barely brings improvement to the model. The average reward is better than the random reward, and the maximum reward provided by the selected reference performs the best.

Models	\mathcal{L}_{CE}	\mathcal{L}_{sum}	\mathcal{L}_{mid}	\mathcal{L}_{max}
AT	27.70	28.08	27.37	27.21
NAT	26.09	25.97	26.90	26.88

Table 6: The performance of AT and CTC-based NAT on the same diverse distillation dataset of WMT14 En-De with different loss functions. \mathcal{L}_{CE} is the cross-entropy loss with sequence-level distillation. \mathcal{L}_{sum} , \mathcal{L}_{mid} , and \mathcal{L}_{max} described in section 3.2.1 are losses for the diverse distillation dataset.

Methods	pwb ↓	rfb ↑	BLEU
HardMoe	53.57	24.77	24.51
Dropout	69.71	26.23	25.35
SeedDiv	59.87	26.99	25.92

Table 7: Pairwise BLEU (pwb) and reference BLEU (rfb) scores of diverse translation techniques and their DDRS performance on WMT14 En-De validation set. pwb and rfb scores are measured on WMT14 En-De test set without compound split.

4.4 DDRS on Autoregressive Transformer

Though DDRS is proposed to alleviate the multi-modality problem for NAT, it can also be applied to autoregressive models. In Table 6, we report the performance of the autoregressive Transformer when trained by the proposed DDRS losses. In contrast to NAT, AT prefers the summation loss \mathcal{L}_{sum} , and the other two losses based on reference selection even degrade the AT performance.

It is within our expectation that AT models do not benefit much from reference selection. NAT generates the whole sentence simultaneously without any prior knowledge about the reference sentence, so the reference may not fit the NAT output well, in which case DDRS is helpful by selecting an appropriate reference for the training. In comparison, AT models generally apply the teacher forcing algorithm (Williams and Zipser, 1989) for the training, which feeds the golden context to guide the generation of the reference sentence. With teacher forcing, AT models do not suffer much from the multi-modality problem and therefore do not need reference selection. Besides, as shown in Table 1, another disadvantage is that the training cost of DDRS is nearly k times as large, so we do not recommend applying DDRS on AT.

4.5 Effect of Diverse Distillation

In the diverse distillation part of DDRS, we apply SeedDiv to generate multiple references. There are also other diverse translation techniques that can

Models	WMT14		WMT16	
	EN-DE	DE-EN	EN-RO	RO-EN
CTC w/o RL	26.09	29.50	33.55	32.98
BLEU	26.48	30.02	33.64	33.31
METEOR	26.44	29.95	33.68	33.25
GLEU	26.58	29.96	33.59	33.34
BERTScore	26.51	30.20	33.69	33.42
BLEURT	26.66	30.05	33.71	33.35

Table 8: BLEU scores on WMT test sets when using different automatic metrics as reward to finetune CTC.

be used for diverse distillation. In this section, we evaluate the effect of diverse distillation techniques on the performance of DDRS. Besides SeedDiv, we also use HardMoe (Shen et al., 2019) and Concrete Dropout (Wu et al., 2020) to generate multiple references, and report their performance in Table 7. When applying other techniques for diverse distillation, the performance of DDRS significantly decreases. The performance degradation indicates the importance of high reference BLEU in diverse distillation, as the NAT student directly learns from the generated references.

4.6 Effect of Reward

There are many automatic metrics to evaluate the translation quality. To measure the effect of reward, we respectively use different automatic metrics as reward for RL, which include traditional metrics (BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), GLEU (Wu et al., 2016)) and pretraining-based metrics (BERTScore (Zhang* et al., 2020), BLEURT (Sellam et al., 2020)). We report the results in Table 8. Comparing the three traditional metrics, we can see that there is no significant difference in their performance. The two pretraining-based metrics only perform slightly better than traditional metrics. Considering the performance and computational cost, we use the traditional metric BLEU as the reward.

4.7 Number of References

In this section, we evaluate how the number of references affects the DDRS performance. We set the number of references k to different values and train the CTC model with reference selection. We report the performance of DDRS with different k in Table 9. The improvement brought by increasing k is considerable when k is small, but it soon becomes marginal. Therefore, it is reasonable to use a middle number of references like $k = 3$ to balance the distillation cost and performance.

k	1	2	3	5	8
BLEU	25.0	25.6	25.9	26.0	26.1

Table 9: Performance of DDRS with different number of references k on WMT14 En-De validation set.

Models	Distill	Train	Total	BLEU
CTC (64K)	5.5h	26.4h	31.9h	26.34
CTC (128K)	5.5h	52.5h	58.0h	26.59
DDRS (32K)	16.5h	19.3h	35.8h	27.60

Table 10: Performance and time cost of models with different batch sizes on WMT14 En-De. The time cost is measured on 8 GeForce RTX 3090 GPUs. The cost of Distill includes training teacher models and decoding source sentences.

4.8 Time Cost

The cost of preparing the training data is larger for DDRS since it requires training k teacher models and using each model to decode the training set. We argue that the cost is acceptable since the distillation cost is minor compared to the training cost of NAT, and we can reduce the training cost to make up for it. In Table 10, we report the performance and time cost of models with different batch sizes on the test set of WMT14 En-De. DDRS makes up for the larger distillation cost by using a smaller training batch, which has a similar cost to the CTC model of 64K batch and achieves superior performance compared to models of 128K batch.

5 Related Work

Gu et al. (2018) proposes non-autoregressive translation to reduce the translation latency, which suffers from the multi-modality problem. A line of work introduces latent variables to model the non-determinism in the translation process, where latent variables are based on fertilities (Gu et al., 2018), vector quantization (Kaiser et al., 2018; Roy et al., 2018; Bao et al., 2021) and variational inference (Ma et al., 2019; Shu et al., 2020). Another branch of work proposes training objectives that are less influenced by the multi-modality problem to train NAT models (Wang et al., 2019; Shao et al., 2019, 2020, 2021; Sun et al., 2019; Ghazvininejad et al., 2020; Shan et al., 2021; Du et al., 2021). Some researchers consider transferring the knowledge from autoregressive models to NAT (Li et al., 2019; Wei et al., 2019; Guo et al., 2020a; Zhou et al., 2020; Sun and Yang, 2020). Besides, some work propose iterative NAT models that refine the model outputs

with multi-pass iterative decoding (Lee et al., 2018; Gu et al., 2019; Ghazvininejad et al., 2019; Ran et al., 2020; Kasai et al., 2020). Our work is most related to CTC-based NAT models (Graves et al., 2006; Libovický and Helcl, 2018; Kasner et al., 2020; Saharia et al., 2020; Zheng et al., 2021; Gu and Kong, 2020), which apply the CTC loss to model latent alignments for NAT. In autoregressive models, translations different from the reference can be evaluated with reinforcement learning (Ranzato et al., 2015; Norouzi et al., 2016), probabilistic n-gram matching (Shao et al., 2018), or an evaluation module (Feng et al., 2020).

Our work is also related to the task of diverse machine translation. Li et al. (2016); Vijayakumar et al. (2018) adjust the beam search algorithm by introducing regularization terms to encourage generating diverse outputs. He et al. (2018); Shen et al. (2019) introduce latent variables with the mixture of experts method and use different latent variables to generate diverse translations. Sun et al. (2020) generates diverse translations by sampling different attention heads. Wu et al. (2020) train the translation model with concrete dropout and samples different models from a posterior distribution. Li et al. (2021) generate different translations for the input sentence by mixing it with different sentence pairs sampled from the training corpus. Nguyen et al. (2020) augment the training set by translating the source-side and target-side data with multiple translation models, but they do not evaluate the diversity of the augmented data.

6 Conclusion

In this paper, we propose diverse distillation with reference selection (DDRS) for NAT. Diverse distillation generates a dataset containing multiple references for each source sentence, and reference selection finds the best reference for the training. DDRS demonstrates its effectiveness on various benchmarks, setting new state-of-the-art performance levels for non-autoregressive translation.

7 Acknowledgement

We thank the anonymous reviewers for their insightful comments. This work was supported by National Key R&D Program of China (NO.2017YFE9132900).

References

- Dzmitry Bahdanau, KyungHyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. 3rd International Conference on Learning Representations, ICLR 2015 ; Conference date: 07-05-2015 Through 09-05-2015.
- Satanjeev Banerjee and Alon Lavie. 2005. **METEOR: An automatic metric for MT evaluation with improved correlation with human judgments**. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Yu Bao, Shujian Huang, Tong Xiao, Dongqi Wang, Xinyu Dai, and Jiajun Chen. 2021. **Non-autoregressive translation by learning target categorical codes**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5749–5759, Online. Association for Computational Linguistics.
- Cunxiao Du, Zhaopeng Tu, and Jing Jiang. 2021. Order-agnostic cross entropy for non-autoregressive machine translation. In *ICML*.
- Yang Feng, Wanying Xie, Shuhao Gu, Chenze Shao, Wen Zhang, Zhengxin Yang, and Dong Yu. 2020. **Modeling fluency and faithfulness for diverse neural machine translation**. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 59–66. AAAI Press.
- Markus Freitag, Yaser Al-Onaizan, and Baskaran Sankaran. 2017. Ensemble distillation for neural machine translation. *arXiv preprint arXiv:1702.01802*.
- Xinwei Geng, Xiaocheng Feng, and Bing Qin. 2021. **Learning to rewrite for non-autoregressive neural machine translation**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3297–3308, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Marjan Ghazvininejad, Vladimir Karpukhin, Luke Zettlemoyer, and Omer Levy. 2020. Aligned cross entropy for non-autoregressive machine translation. In *ICML*.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. **Mask-predict: Parallel decoding of conditional masked language models**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6112–6121.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. **Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks**. In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, page 369–376, New York, NY, USA. Association for Computing Machinery.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor O.K. Li, and Richard Socher. 2018. **Non-autoregressive neural machine translation**. In *International Conference on Learning Representations*.
- Jiatao Gu and Xiang Kong. 2020. **Fully non-autoregressive neural machine translation: Tricks of the trade**.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. **Levenshtein transformer**. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Junliang Guo, Xu Tan, Linli Xu, Tao Qin, Enhong Chen, and Tie-Yan Liu. 2020a. **Fine-tuning by curriculum learning for non-autoregressive neural machine translation**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:7839–7846.
- Junliang Guo, Linli Xu, and Enhong Chen. 2020b. **Jointly masked sequence-to-sequence model for non-autoregressive neural machine translation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 376–385, Online. Association for Computational Linguistics.
- Xuanli He, Gholamreza Haffari, and Mohammad Norouzi. 2018. **Sequence to sequence mixture model for diverse machine translation**. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 583–592, Brussels, Belgium. Association for Computational Linguistics.
- Chenyang Huang, Hao Zhou, Osmar R Zaiane, Lili Mou, and Lei Li. 2021. Non-autoregressive translation with layer-wise prediction and deep supervision. *ArXiv*, abs/2110.07515.
- Lukasz Kaiser, Samy Bengio, Aurko Roy, Ashish Vaswani, Niki Parmar, Jakob Uszkoreit, and Noam Shazeer. 2018. **Fast decoding in sequence models using discrete latent variables**. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2390–2399. PMLR.
- Jungo Kasai, James Cross, Marjan Ghazvininejad, and Jiatao Gu. 2020. Non-autoregressive machine translation with disentangled context transformer. In *ICML*.
- Zdeněk Kasner, Jindřich Libovický, and Jindřich Helcl. 2020. **Improving fluency of non-autoregressive machine translation**.

- 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.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. [Deterministic non-autoregressive neural sequence modeling by iterative refinement](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182, Brussels, Belgium. Association for Computational Linguistics.
- Jicheng Li, Pengzhi Gao, Xuanfu Wu, Yang Feng, Zhongjun He, Hua Wu, and Haifeng Wang. 2021. [Mixup decoding for diverse machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 312–320, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. [A simple, fast diverse decoding algorithm for neural generation](#).
- Zhuohan Li, Zi Lin, Di He, Fei Tian, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2019. [Hint-based training for non-autoregressive machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5708–5713, Hong Kong, China. Association for Computational Linguistics.
- Jindřich Libovický and Jindřich Helcl. 2018. [End-to-end non-autoregressive neural machine translation with connectionist temporal classification](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3016–3021, Brussels, Belgium. Association for Computational Linguistics.
- Ye Liu, Yao Wan, Jianguo Zhang, Wenting Zhao, and Philip Yu. 2021. [Enriching non-autoregressive transformer with syntactic and semantic structures for neural machine translation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1235–1244, Online. Association for Computational Linguistics.
- Xuezhe Ma, Chunting Zhou, Xian Li, Graham Neubig, and Eduard Hovy. 2019. [FlowSeq: Non-autoregressive conditional sequence generation with generative flow](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4282–4292, Hong Kong, China. Association for Computational Linguistics.
- Xuan-Phi Nguyen, Shafiq R. Joty, Kui Wu, and Ai Ti Aw. 2020. [Data diversification: A simple strategy for neural machine translation](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Mohammad Norouzi, Samy Bengio, Zhifeng Chen, Navdeep Jaitly, Mike Schuster, Yonghui Wu, and Dale Schuurmans. 2016. [Reward augmented maximum likelihood for neural structured prediction](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1723–1731.
- 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 NAACL-HLT 2019: Demonstrations*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Lihua Qian, Hao Zhou, Yu Bao, Mingxuan Wang, Lin Qiu, Weinan Zhang, Yong Yu, and Lei Li. 2021. [Glancing transformer for non-autoregressive 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 1993–2003, Online. Association for Computational Linguistics.
- Qiu Ran, Yankai Lin, Peng Li, and Jie Zhou. 2020. [Learning to recover from multi-modality errors for non-autoregressive neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3059–3069, Online. Association for Computational Linguistics.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2015. [Sequence level training with recurrent neural networks](#). *arXiv preprint arXiv:1511.06732*.
- Aurko Roy, Ashish Vaswani, Arvind Neelakantan, and Niki Parmar. 2018. [Theory and experiments on vector quantized autoencoders](#). *arXiv preprint arXiv:1805.11063*.
- Chitwan Saharia, William Chan, Saurabh Saxena, and Mohammad Norouzi. 2020. [Non-autoregressive machine translation with latent alignments](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1098–1108, Online. Association for Computational Linguistics.

- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. 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.
- 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.
- Yong Shan, Yang Feng, and Chenze Shao. 2021. Modeling coverage for non-autoregressive neural machine translation. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Chenze Shao, Xilin Chen, and Yang Feng. 2018. [Greedy search with probabilistic n-gram matching for neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4778–4784, Brussels, Belgium. Association for Computational Linguistics.
- Chenze Shao, Yang Feng, Jinchao Zhang, Fandong Meng, Xilin Chen, and Jie Zhou. 2019. [Retrieving sequential information for non-autoregressive neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3013–3024, Florence, Italy. Association for Computational Linguistics.
- Chenze Shao, Yang Feng, Jinchao Zhang, Fandong Meng, and Jie Zhou. 2021. [Sequence-Level Training for Non-Autoregressive Neural Machine Translation](#). *Computational Linguistics*, 47(4):891–925.
- Chenze Shao, Jinchao Zhang, Yang Feng, Fandong Meng, and Jie Zhou. 2020. [Minimizing the bag-of-ngrams difference for non-autoregressive neural machine translation](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 198–205. AAAI Press.
- Tianxiao Shen, Myle Ott, Michael Auli, and Marc’Aurelio Ranzato. 2019. [Mixture models for diverse machine translation: Tricks of the trade](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5719–5728. PMLR.
- Raphael Shu, Jason Lee, Hideki Nakayama, and Kyunghyun Cho. 2020. Latent-variable non-autoregressive neural machine translation with deterministic inference using a delta posterior. In *AAAI*.
- Jongyoon Song, Sungwon Kim, and Sungroh Yoon. 2021. [AlignNART: Non-autoregressive neural machine translation by jointly learning to estimate alignment and translate](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1–14, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zewei Sun, Shujian Huang, Hao-Ran Wei, Xinyu Dai, and Jiajun Chen. 2020. [Generating diverse translation by manipulating multi-head attention](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8976–8983. AAAI Press.
- Zhiqing Sun, Zhuohan Li, Haoqing Wang, Di He, Zi Lin, and Zhihong Deng. 2019. [Fast structured decoding for sequence models](#). In *Advances in Neural Information Processing Systems 32*, pages 3016–3026.
- Zhiqing Sun and Yiming Yang. 2020. [An EM approach to non-autoregressive conditional sequence generation](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9249–9258. PMLR.
- 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 *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pages 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Ashwin K. Vijayakumar, Michael Cogswell, Ramprasaath R. Selvaraju, Qing Sun, Stefan Lee, David J. Crandall, and Dhruv Batra. 2018. [Diverse beam search for improved description of complex scenes](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 7371–7379. AAAI Press.
- Yiren Wang, Fei Tian, Di He, Tao Qin, ChengXiang Zhai, and Tie-Yan Liu. 2019. [Non-autoregressive machine translation with auxiliary regularization](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, Honolulu, Hawaii, USA, 2019*, pages 5377–5384. AAAI Press.
- Bingzhen Wei, Mingxuan Wang, Hao Zhou, Junyang Lin, and Xu Sun. 2019. [Imitation learning for non-autoregressive neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1304–1312, Florence, Italy. Association for Computational Linguistics.

- Ronald J. Williams. 1992. [Simple statistical gradient-following algorithms for connectionist reinforcement learning](#). *Mach. Learn.*, 8(3–4):229–256.
- Ronald J Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2).
- Xuanfu Wu, Yang Feng, and Chenze Shao. 2020. [Generating diverse translation from model distribution with dropout](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1088–1097, Online. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Jiajun Zhang and Chengqing Zong. 2016. [Exploiting source-side monolingual data in neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1535–1545.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Zaixiang Zheng, Hao Zhou, Shujian Huang, Jiajun Chen, Jingjing Xu, and Lei Li. 2021. [Duplex sequence-to-sequence learning for reversible machine translation](#).
- Chunting Zhou, Jiatao Gu, and Graham Neubig. 2020. [Understanding knowledge distillation in non-autoregressive machine translation](#). In *International Conference on Learning Representations*.
- Jiawei Zhou and Phillip Keung. 2020. [Improving non-autoregressive neural machine translation with monolingual data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1893–1898, Online. Association for Computational Linguistics.