

Imitation Learning for Non-Autoregressive Neural Machine Translation

Bingzhen Wei¹, Mingxuan Wang, Hao Zhou, Junyang Lin^{1,3}, Xu Sun^{1,2}

¹MOE Key Lab of Computational Linguistics, School of EECS, Peking University

²Deep Learning Lab, Beijing Institute of Big Data Research, Peking University

³School of Foreign Languages, Peking University

{weibz, linjunyang, xusun}@pku.edu.cn

xuanxuans27@gmail.com, haozhou0806@gmail.com

Abstract

Non-autoregressive translation models (NAT) have achieved impressive inference speedup. A potential issue of the existing NAT algorithms, however, is that the decoding is conducted in parallel, without directly considering previous context. In this paper, we propose an imitation learning framework for non-autoregressive machine translation, which still enjoys the fast translation speed but gives comparable translation performance compared to its auto-regressive counterpart. We conduct experiments on the IWSLT16, WMT14 and WMT16 datasets. Our proposed model achieves a significant speedup over the autoregressive models, while keeping the translation quality comparable to the autoregressive models. By sampling sentence length in parallel at inference time, we achieve the performance of 31.85 BLEU on WMT16 Ro→En and 30.68 BLEU on IWSLT16 En→De.

1 Introduction

Neural machine translation (NMT) with encoder-decoder architectures (Sutskever et al., 2014; Cho et al., 2014) achieve significantly improved performance compared with traditional statistical methods (Koehn et al., 2003; Koehn, 2010). Nevertheless, the autoregressive property of the NMT decoder has been a bottleneck of the translation speed. Specifically, the decoder, whether based on Recurrent Neural Network (RNN) (Hochreiter and Schmidhuber, 1997; Cho et al., 2014) or attention mechanism (Vaswani et al., 2017), sequentially generates words. The latter words are conditioned on previous words in a sentence. Such bottleneck disables parallel computation of decoder, which is serious for NMT, since the NMT decoding with a large vocabulary is extremely time-consuming.

Recently, a line of research work (Gu et al., 2017; Lee et al., 2018; Libovick and Helcl, 2018;

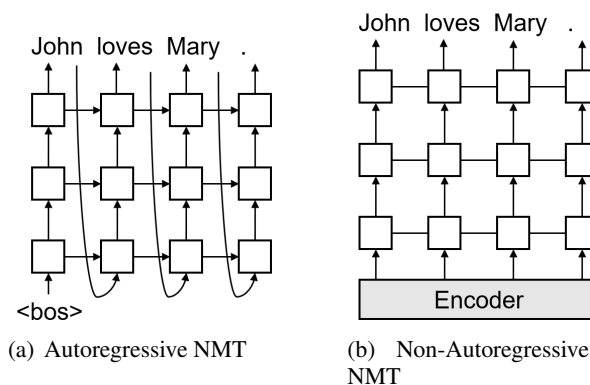


Figure 1: Neural architectures for Autoregressive NMT and Non-Autoregressive NMT.

Wang et al., 2018) propose to break the autoregressive bottleneck by introducing non-autoregressive neural machine translation (NAT). In NAT, the decoder generates all words simultaneously instead of sequentially. Intuitively, NAT abandon feeding previous predicted words into decoder state at the next time step, but directly copy source encoded representation (Gu et al., 2017; Lee et al., 2018; Guo et al., 2018; Wang et al., 2019) as inputs of the decoder. Thus, the generation of the NAT models does not condition on previous prediction. NAT enables parallel computation of decoder, giving significantly fast translation speed with moderate accuracy (always within 5 BLEU). Figure 1 shows the difference between autoregressive and non-autoregressive models.

However, we argue that current NAT approaches suffer from delayed supervisions (or rewards) and large search space in training. NAT decoder simultaneously generates all words of the translation, the search space of which is very large. For one time step, decoding states across layers (more than 16 layers) and time steps could be regarded as a 2-dimensional sequential decision process. Every decoding state has not only

to decide which part of target sentence it will focus on, but also to decide the correct target word of that part. All decisions are made by interactions with other decoding states. Delayed supervisions (correct target word) will be obtained by decoding states in the last layer, and intermediate decoding states will be updated by gradient propagation from the last layer. Therefore, the training of NAT is non-trivial and it may be hard for NAT to achieve a good model, which is the same case that reinforcement learning (Mnih et al., 2013, 2015) is hard to learn with large search space. The delayed supervision problem is not severe for autoregressive neural machine translation(AT) because it predicts words sequentially. Given the previous words, contents to be predicted at current step are relatively definite, thus the search space of AT is exponentially lower than NAT. We blame the delayed supervision and large search space for the performance gap between NAT and AT.

In this paper, we propose a novel imitation learning framework for non-autoregressive NMT (imitate-NAT). Imitation learning has been widely used to alleviate the problem of huge search space with delayed supervision in RL. It is straightforward to bring the imitation learning idea for boosting the performance of NAT. Specifically, we introduce a knowledgeable AT demonstrator to supervise each decoding state of NAT model. In such case, Specifically, We propose to employ a knowledgeable AT demonstrator to supervise every decoding state of NAT across different time steps and layers, which works pretty well practically. Since the AT demonstrator is only used in training, our proposed imitate-NAT enjoys the high speed of NAT without suffering from its relatively lower translation performance.

Experiments show that our proposed imitate-NAT is fast and accurate, which effectively closes the performance gap between AT and NAT on several standard benchmarks, while maintains the speed advantages of NAT (10 times faster). On all the benchmark datasets, our imitate-NAT with LPD achieves the best translation performance, which is even close to the results of the autoregressive model.

2 Background

In the following sections, we introduce the background about Autoregressive Neural Machine Translation and Non-Autoregressive Neural Ma-

chine Translation.

2.1 Autoregressive Neural Machine Translation

Sequence modeling in machine translation has largely focused on autoregressive modeling which generate a target sentence word by word from left to right, denoted by $p_\theta(Y|X)$, where $X = \{x_1 \cdots, x_T\}$ and $Y = \{y_1, \cdots, y_{T'}\}$ represent the source and target sentences as sequences of words respectively. θ is a set of parameters usually trained to minimize the negative loglikelihood:

$$\mathcal{L}_{AT} = - \sum_{i=1}^{T'} \log p(y_i | y_{<i}, X). \quad (1)$$

where T and T' is the length of the source and the target sequence respectively.

Deep neural network with autoregressive framework has achieved great success on machine translation, with different choices of architectures. The RNN-based NMT approach, or RNMT, was quickly established as the de-facto standard for NMT. Despite the recent success, the inherently sequential architecture prevents RNMTs from being parallelized during training and inference. Following RNMT, CNNs and self-attention based models have recently drawn research attention due to their ability to fully parallelize training to take advantage of modern fast computing devices. However, the autoregressive nature still creates a bottleneck at inference stage, since without ground truth, the prediction of each target token has to condition on previously predicted tokens.

2.2 Non-Autoregressive Neural Machine Translation

As a solution to the issue of slow decoding, Gu et al. (2017) recently proposed non-autoregressive model (NAT) to break the inference bottleneck by exposing all decoder inputs to the network simultaneously. NAT removes the autoregressive connection directly and factorizes the target distribution into a product of conditionally independent per-step distributions. The negative loglikelihood loss function for NAT model become is then defined as:

$$\mathcal{L}_{NAT} = - \sum_{i=1}^{T'} \log p(y_i | X). \quad (2)$$

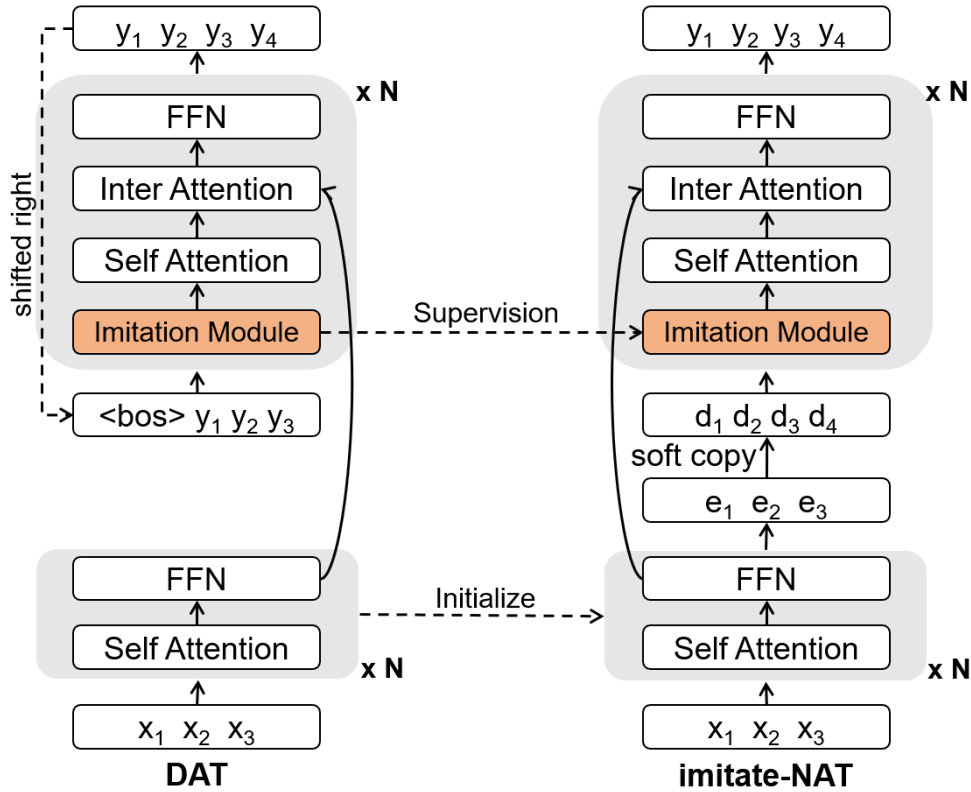


Figure 2: Illustration of the proposed model, where the black solid arrows represent differentiable connections and the dashed arrows are non-differentiable operations. Without loss of generality, this figure shows the case of $T=3$, $T'=4$. The left side of the figure is the DAT model and the right side is the imitate-NAT. The bottom is the encoder and the top is the decoder. The internal details of Imitation Module are shown in Figure 3.

The approach breaks the dependency among the target words across time, thus the target distributions can be computed in parallel at inference time.

In particular, the encoder stays unchanged from the original Transformer network. A latent fertility model is then used to copy the sequence of source embeddings as the input of the decoder. The decoder has the same architecture as the encoder plus the encoder attention. The best results were achieved by sampling fertilities from the model and then rescore the output sentences using an autoregressive model. The reported inference speed of this method is 2-15 times faster than a comparable autoregressive model, depending on the number of fertility samples.

This desirable property of exact and parallel decoding however comes at the expense of potential performance degradation. Since the conditional dependencies within the target sentence (y_t depends on $y_{<t}$) are removed from the decoder input, the decoder is not powerful enough to leverage the inherent sentence structure for prediction. Hence the decoder has to figure out such target-side information by itself just with the source-side informa-

tion during training, which leads to a larger modeling gap between the true model and the neural sequence model. Therefore, strong supervised signals could be introduced as the latent variable to help the model learn better internal dependencies within a sentence.

In AT models, the generation of the current token is conditioned on previously generated tokens, which provides strong target side context information. In contrast, NAT models generate tokens in parallel, thus the target-side dependency is indirect and weak. Consequently, the decoder of a NAT model has to handle the translation task conditioned on less and weaker information compared with its AT counterpart, thus leading to inferior accuracy.

3 Proposed Method: imitate-NAT

In this section, we propose an imitation learning framework (imitate-NAT) to close the performance gap between the NAT and AT.

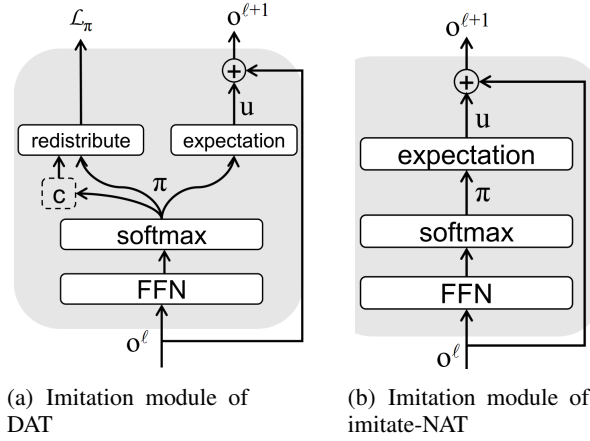


Figure 3: The imitation module of AT demonstrator and NAT learner.

3.1 Preliminary of imitate-NAT

We bring the intuition of imitation learning to non-autoregressive NMT and adapt it to our scenario. Specifically, the NAT model can be regarded as a *learner*, which will imitate a knowledgeable *demonstrator* at each decoding state across layers and time steps. However, obtaining an adequate demonstrator is non-trivial. We propose to employ an autoregressive NMT model as the demonstrator, which is expected to offer efficient supervision to each decoding state of the NAT model. Fortunately, the AT demonstrator is only used in training, which guarantees that our proposed imitate-NAT enjoys the high speed of NAT model without suffering from its relatively lower performance.

In following parts, we will describe the *AT demonstrator* and the *NAT learner* in our imitate-NAT framework, respectively.

3.2 AT Demonstrator

For the proposed AT, we apply a variant of the transformer model as the demonstrator, named DAT. The encoder stays unchanged from the original Transformer network. A crucial difference lies in that the decoder introduces the imitation module which emits actions at every time step. The action brings sequential information, thus can be used as the guidance signal during the NAT training process.

The input of each decoder layer $O^\ell = \{o_1^\ell, o_2^\ell, \dots, o_{T'}^\ell\}$ can be considered as the observation (or environment) of the IL framework, where ℓ donates the layer of the observation. Let $A^\ell = \{a_1^\ell, a_2^\ell, \dots, a_{T'}^\ell\} \in \mathcal{A}$ denotes an action sequence from the action space \mathcal{A} . The action space

\mathcal{A} is finite and its size n is a hyperparameter, representing the number of action categories. The distribution of the action of DAT can be then fed to the NAT model as the training signal. Let Π denotes a policy class, where each $\pi^\ell \in \Pi$ generates an action distribution sequence A^ℓ in response to a context sequence O^ℓ .

Predicting actions A^ℓ may depend on the contexts of previous layer O^ℓ and policies π^ℓ can thus be viewed as mapping states to actions. A roll-out of π given the context sequence O^ℓ to determine the action sequence A^ℓ , which is:

$$a_t = \arg \max(\pi^\ell(o_t^\ell)) \quad (3)$$

where

$$\pi^\ell(o_t^\ell) = \text{softmax}(\text{FFN}(o_t^\ell)). \quad (4)$$

The distribution $\pi^\ell(o_t^\ell)$ represents the probability of the decision depends on the current state or environment o_t^ℓ . The discrete operation $\arg \max(\cdot)$ suffers from the non-differentiable problem which makes it impossible to train the policy from an end to end framework.

Note that unlike the general reinforcement or imitation learning framework, we consider to compute the action state which as the expectation of the embedding of the action a_t :

$$u_t^\ell = \mathbb{E}_{a_t^\ell \sim \pi^\ell(o_t^\ell)} \delta(a_t^\ell), \quad (5)$$

where $\delta(a_t^\ell) \in \mathbb{R}^k$ returns the embedding of the action a_t^ℓ and k denotes the embedding dimension. The states of next layer are then based on the current output of the decoder state and the emitted action state:

$$o_t^{\ell+1} = \text{Transfer}(u_t^\ell + o_t^\ell), \quad (6)$$

where $\text{Transfer}(\cdot)$ denotes the vanilla transformer decoding function including a self-attention layer, an encoder-decoder attention layer and followed by a FFN layer (Vaswani et al., 2017).

3.2.1 Action Distribution Regularization

The supervised signal for the action distribution $\pi(o_t)$ is not direct in NAT, thus the action prediction can be viewed as an unsupervised clustering problem. One potential issue is the unbalanced distribution of action. Inspired by Xie et al. (2016), we introduce a regularization method to

increase the space utilization. Formally, an moving average c is applied to calculate the cumulative activation level for each action category:

$$c \leftarrow \alpha \cdot c + (1 - \alpha) \sum_{t=1}^{T'} \pi(o_t) / T' \quad (7)$$

We set α 0.9 in our experiments. Then $\pi'(o_i)$ can be re-normalized with the cumulative history c :

$$\pi'(o_t) = \frac{\pi(o_t)^2 / c}{\sum_j \pi(o_t)_j^2 / c_j} \quad (8)$$

The convex property of the quadratic function can adjust the distribution to achieve the purpose of clustering. The role of c is to redistribute the probability distribution of $\pi(o_t)$, which leads to a more balanced category assignment.

We define our objective as a KL divergence loss between $\pi(o_t)$ and the auxiliary distribution $\pi'(o_t)$ as follows:

$$\mathcal{L}_\pi = \sum_t \pi'(o_t) \log \frac{\pi'(o_t)}{\pi(o_t)} \quad (9)$$

3.3 NAT learner

3.3.1 Soft Copy

To facility the imitation learning process, our imitate-NAT is based on the AT demonstrator described in section 3.2. The only difference lies in that the initialization of the decoding inputs. Previous approaches apply a UniformCopy method to address the problem. More specifically, the decoder input at position t is the copy of the encoder embedding at position $\text{Round}(T't/T)$ (Gu et al., 2017; Lee et al., 2018). As the source and target sentences are often of different lengths, AT model need to predict the target length T' during inference stage. The length prediction problem can be viewed as a typical classification problem based on the output of the encoder. we follow Lee et al. (2018) to predict the length of the target sequence.

The proposed *Round* function is unstable and non-differentiable, which make the decoding task difficult. We therefore propose a differentiable and robust method named SoftCopy following the spirit of the attention mechanism (Hahn and Keller, 2016; Bengio, 2009). The weight $w_{i,j}$ depends on the distance relationship between the source position i and the target position j .

$$w_{i,j} = \text{softmax}(-|j - i|/\tau) \quad (10)$$

τ is a trainable parameters used to adjust the degree of focus when copying. Then the input of the target at position j can be computed as :

$$y_j = \sum_{i=0}^T w_{ij} x_i, \quad (11)$$

where x_i is usually the source embedding at position i . It is also worth mentioning that we take the top-most hidden states instead of the word embedding as x_i in order to cache the global context information.

3.3.2 Learning from AT Experts

The conditional independence assumption prevents NAT model from properly capturing the highly multimodal distribution of target translations. AT models takes already generated target tokens as inputs, thus can provide complementary extension information for NAT models. A straightforward idea to bridge the gap between NAT and AT is that NAT can actively learn the behavior of AT step by step.

The AT demonstrator generate action distribution $\pi_{AT}(O) \in \mathbb{R}^n$ as the posterior supervisor signal. We expect the supervision information can guide the generation process of NAT. The imitate-NAT exactly follows the same decoder structure with our AT demonstrator, and emits distribution $\pi_{NAT}(O) \in \mathbb{R}^n$ to learn from AT demonstrator step by step. More specifically, we try to minimize the cross entropy of the distributions between the two policies:

$$\mathcal{L}_{IL} = H(\pi_{AT}(o_t), \pi_{NAT}(o_t)) \quad (12)$$

$$= -\mathbb{E}_{\pi_{AT}(o_t)} \log \pi_{NAT}(o_t) \quad (13)$$

3.4 Training

In the training process, the action distribution regularization term described in 3.2.1 is combined with the commonly used cross-entropy loss in Eq. 1:

$$\mathcal{L}_{AT}^* = \mathcal{L}_{AT} + \lambda_1 \mathcal{L}_\pi \quad (14)$$

For NAT models, the imitation learning term are combined with the commonly used cross-entropy loss in Eq. 2:

$$\mathcal{L}_{NAT}^* = \mathcal{L}_{NAT} + \lambda_2 \mathcal{L}_{IL} \quad (15)$$

where λ_1 and λ_2 are hyper-parameters, which are set to 0.001 in our experiments.

| Models | WMT14 | | WMT16 | | IWSLT16 | |
|--|-------|-------|-------|-------|---------|---------|
| | En→De | De→En | En→Ro | Ro→En | En→De | Speedup |
| Transformer (Vaswani et al., 2017) | 27.41 | 31.29 | / | / | 30.90 | 1.00× |
| AT Demonstrator | 27.80 | 31.25 | 33.70 | 32.59 | 30.85 | 1.05× |
| NAT-FT(Gu et al., 2017) | 17.69 | 21.47 | 27.29 | 29.06 | 26.52 | 15.60× |
| NAT-FT(+NPD s=10) | 18.66 | 22.41 | 29.02 | 30.76 | 27.44 | 7.68× |
| NAT-FT(+NPD s=100) | 19.17 | 23.20 | 29.79 | 31.44 | 28.16 | 2.36× |
| NAT-IR($i_{dec} = 1$) | 13.91 | 16.77 | 24.45 | 25.73 | 22.20 | 8.90× |
| NAT-IR($i_{dec} = 10$) | 21.61 | 25.48 | 29.32 | 30.19 | 27.11 | 1.50× |
| LT | 19.80 | / | / | / | / | 5.78× |
| LT(rescoring 10) | 21.0 | / | / | / | / | / |
| LT(rescoring 100) | 22.5 | / | / | / | / | / |
| NAT without imitation | 19.69 | 22.71 | / | / | 25.34 | 18.6× |
| imitate-NAT | 22.44 | 25.67 | 28.61 | 28.90 | 28.41 | 18.6× |
| imitate-NAT (+LPD, $\Delta T = 3$) | 24.15 | 27.28 | 31.45 | 31.81 | 30.68 | 9.70× |

Table 1: The test set performances of AT and NAT models in BLEU score. NAT-FT, NAT-IR and LT denotes the competitor method in (Gu et al., 2017), (Lee et al., 2018) and (Kaiser et al., 2018) respectively. imitate-NAT is our proposed NAT with imitation learning.

4 Experiments

We evaluate our proposed model on machine translation tasks and provide the analysis. We present the experimental details in the following, including the introduction to the datasets as well as our experimental settings.

Datasets We evaluate the proposed method on three widely used public machine translation corpora: IWSLT16 En-De(196K pairs), WMT14 En-De(4.5M pairs) and WMT16 En-Ro(610K pairs). All the datasets are tokenized by Moses Koehn et al. (2007) and segmented into $32k$ -subword symbols with byte pair encoding Sennrich et al. (2016) to restrict the size of the vocabulary. For WMT14 En-De, we use newstest-2013 and newstest-2014 as development and test set respectively. For WMT16 En-Ro, we use newsdev-2016 and newstest-2016 as development and test sets respectively. For IWSLT16 En-De, we use test2013 as validation for ablation experiments.

Knowledge Distillation Datasets Sequence-level knowledge distillation is applied to alleviate multimodality in the training dataset, using the AT demonstrator as the teachers (Kim and Rush, 2016). We replace the reference target sentence of each pair of training example (X, Y) with a new target sentence Y^* , which is generated from the teacher model(AT demonstrator). Then we use the new dataset (X, Y^*) to train our NAT model. To avoid the redundancy of running fixed teacher

models repeatedly on the same data, we decode the entire training set once using each teacher to create a new training dataset for its respective student.

Model Settings We first train the AT demonstrator and then freeze its parameters during the training of imitate-NAT. In order to speed up the convergence of NAT training, we also initialize imitate-NAT with the corresponding parameters of the AT expert as they have similar architecture. For WMT14 En-De and WMT16 En-Ro, we use the hyperparameter settings of base Transformer model in Vaswani et al. (2017)($d_{model} = 512, d_{hidden} = 512, n_{layer} = 6$ and $n_{head} = 8$). As in Gu et al. (2017); Lee et al. (2018), we use the small model ($d_{model} = 278, d_{hidden} = 507, n_{layer} = 5$ and $n_{head} = 2$) for IWSLT16 En-De. For sequence-level distillation, we set beam size to be 4. For imitate-NAT, we set the number of action category to 512 and found imitate-NAT is robust to the setting in our preliminary experiments.

Length Parallel Decoding For inference, we follow the common practice of noisy parallel decoding (Gu et al., 2017), which generates a number of decoding candidates in parallel and selects the best translation via re-scoring using AT teacher. In our scenario, we first train a module to predict the target length as \hat{T} . However, due to the inherent uncertainty of the data itself, it is

hard to accurately predict the target length. A reasonable solution is to generate multiple translation candidates by predicting different target length $\in [\hat{T} - \Delta T, \hat{T} + \Delta T]$, which we called LPD (length parallel decoding). The model generates several outputs in parallel, then we use the pre-trained autoregressive model to identify the best overall translation.

5 Results and Analysis

Competitor We include three NAT works as our competitors, the NAT with fertility (NAT-FT) (Gu et al., 2017), the NAT with iterative refinement (NAT-IR) (Lee et al., 2018) and the NAT with discrete latent variables (Kaiser et al., 2018). For all our tasks, we obtain the baseline performance by either directly using the performance figures reported in the previous works if they are available or producing them by using the open source implementation of baseline algorithms on our datasets. The results are shown in Table 1.

1. imitate-NAT significantly improved the quality of the translation with a large margin. On all the benchmark datasets, our imitate-NAT with LPD achieves the best translation performance, which is even close to the results of the autoregressive model, e.g. 30.68 vs. 30.85 on IWSLT16 En→De tasks, and 31.81 vs. 32.59 on WMT16 Ro→En tasks. It is also worth mentioning that introducing the imitation module to AT demonstrator does not affect both the performance and the inference speed compared with the standard transformer model.

2. imitate-NAT Imitation learning plays an important role on bridging the gap between imitate-NAT and AT demonstrator Clearly, imitate-NAT leads to remarkable improvements over the competitor without imitation module (over almost 3 BLEU score on average). To make a fair comparison, the competitor follow exactly the same training steps with imitate-NAT, including the initialization, knowledge distillation, and Soft-Copy. The only difference comes from the imitation module.

3. imitate-NAT gets better latency. For NAT-FT, a big sample size(10 and 100) is required to get satisfied results, which seriously affects the inference speed of the model. Both NAT-FT and NAT-IR, the efficiency of models with refinement technique drop dramatically($15.6\times \rightarrow 2.36\times$ of

NAT-FT and $8.9\times \rightarrow 1.5\times$ of NAT-IR). Our imitate-NAT gets even better performance with faster speed. The speedup compared with AT model is $9.7\times$.

5.1 Ablation Study

To further study the effects brought by different techniques, we show in Table 2 the translation performance of different NAT model variants for the IWSLT16 En-De translation task.

Soft-Copy v.s. Uniform-Copy The experimental results show that Soft-Copy is better than Uniform-Copy. Since Uniform-Copy employs a hard copy mechanism and directly copies the source embeddings without considering the global information, which increases the learning burden of the decoder. Our model takes the output of encoder as input and proposes a differentiable copy mechanism which gets much better results(25.34 vs. 20.71, see in line 3 and 2).

Imitation Learning v.s. Non Imitation Learning The imitation learning method leads to an improvement of around 3 BLEU points(28.41 vs. 25.34, see line 6 and 3). NAT without IL degenerates into a normal NAT model. As discussed in section 1, current NAT approaches suffer from delayed supervisions (or rewards) and large search space in training. NAT decoder simultaneously generates all words of the translation, the search space of which is very large.

Length Parallel Decoding Compared with the greedy beam search, LPD technique improves the performance around 2 BLEU points(30.68 vs. 28.41, from line 7 and 6). The observation is in consist with our intuition that sampling from the length space can improve the performance.

Complementary with Knowledge Distillation In consist with previous work, NAT models achieved +4.2 BLEU score from sequence level knowledge distillation technique (see in row 1 and row 2). imitate-NAT without knowledge distillation obtained 23.56 BLEU score which is comparable to non-imitation NAT with knowledge distillation (see in row 3 and row 4). More importantly, we found that the imitation learning framework complemented with knowledge distillation perfectly. As shown in row 3 and 6, imitate-NAT substantially improves the performance of non-imitation NAT knowledge distillation up by +3.3 BLEU score.

| | Distill | UniformCopy | SoftCopy | LPD | Imitation Learning | BLEU |
|---|---------|-------------|----------|-----|--------------------|-------|
| 1 | | ✓ | | | w/o | 16.51 |
| 2 | ✓ | ✓ | | | w/o | 20.72 |
| 3 | ✓ | | | | w/o | 25.34 |
| 4 | | | ✓ | | w/ | 23.56 |
| 5 | | | ✓ | ✓ | w/ | 24.35 |
| 6 | ✓ | | ✓ | | w/ | 28.41 |
| 7 | ✓ | | ✓ | ✓ | w/ | 30.68 |

Table 2: Ablation study on the dev set of IWSLT16. *w/* indicates *with* and *w/o* indicates *without*. LPD indicates length parallel decoding.




Figure 4: Action category assignment distribution. Redistribute method leads to a more balanced distribution(blue), otherwise, it will be extremely unbalanced(red).

Action Distribution Study One common problem in unsupervised clustering is that the results are unbalanced. In this paper, we call that an action is selected or activated when its probability in $\pi(o_t)$ is maximum. Then the space usage can be calculated by counting the number of times each action is selected. We evaluate the space usage on the development set of IWSLT16, and the results are presented in Figure 4. We greatly alleviate the problem of space usage through the category redistribution technique(Eq.7, Eq.8). When building the model without category redistribution, most of the space is not utilized, and the clustering results are concentrated in a few spatial locations, and the category information cannot be dynamically and flexibly characterized. In contrast, category redistribution makes the category distribution more balanced and more in line with the inherent rules of the language, so the clustering results can effectively guide the learning of the NAT model.

6 Related Work

Gu et al. (2017) first developed a non-autoregressive NMT system which produces the outputs in parallel and the inference speed

is thus significantly boosted. However, it comes at the cost that the translation quality is largely sacrificed since the intrinsic dependency within the natural language sentence is abandoned. A bulk of work has been proposed to mitigate such performance degradation. Lee et al. (2018) proposed a method of iterative refinement based on latent variable model and denoising autoencoder. Libovick and Helcl (2018) take NAT as a connectionist temporal classification problem, which achieved better latency. Kaiser et al. (2018) use discrete latent variables that makes decoding much more parallelizable. They first auto encode the target sequence into a shorter sequence of discrete latent variables, which at inference time is generated autoregressively, and finally decode the output sequence from the shorter latent sequence in parallel. Guo et al. (2018) enhanced decoder input by introducing phrase table in SMT and embedding transformation. Wang et al. (2019) leverage the dual nature of translation tasks (e.g., English to German and German to English) and minimize a backward reconstruction error to ensure that the hidden states of the NAT decoder are able to recover the source side sentence.

Unlike the previous work to modify the NAT architecture or decoder inputs, we introduce an imitation learning framework to close the performance gap between NAT and AT. To the best of our knowledge, it is the first time that imitation learning was applied to such problems.

7 Conclusion

We propose an imitation learning framework for non-autoregressive neural machine translation to bridge the performance gap between NAT and AT. Specifically, We propose to employ a knowledgeable AT demonstrator to supervise every decoding state of NAT across different time steps and lay-

ers. As a result, imitate-NAT leads to remarkable improvements and largely closes the performance gap between NAT and AT on several benchmark datasets.

As a future work, we can try to improve the performance of the NMT by introducing more powerful demonstrator with different structure (e.g. right to left). Another direction is to apply the proposed imitation learning framework to similar scenarios such as simultaneous interpretation.

Acknowledgement

We thank the anonymous reviewers for their thoughtful comments. Xu Sun is the corresponding author of this paper.

References

- Yoshua Bengio. 2009. Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1):1–127.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP 2014*, pages 1724–1734.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor O. K. Li, and Richard Socher. 2017. [Non-Autoregressive Neural Machine Translation](#). *arXiv:1711.02281 [cs]*. ArXiv: 1711.02281.
- Junliang Guo, Xu Tan, Di He, Tao Qin, Linli Xu, and Tie-Yan Liu. 2018. Non-autoregressive neural machine translation with enhanced decoder input. *CoRR*, abs/1812.09664.
- Michael Hahn and Frank Keller. 2016. Modeling human reading with neural attention. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 85–95.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- ukasz Kaiser, Aurko Roy, Ashish Vaswani, Niki Parmar, Samy Bengio, Jakob Uszkoreit, and Noam Shazeer. 2018. [Fast Decoding in Sequence Models using Discrete Latent Variables](#). *arXiv:1803.03382 [cs]*. ArXiv: 1803.03382.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-Level Knowledge Distillation](#). *arXiv:1606.07947 [cs]*. ArXiv: 1606.07947.
- Philip Koehn. 2010. Statistical machine translation. Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *ACL*.
- Philipp Koehn, Franz Josef Och, and Daniel Marcu. 2003. Statistical phrase-based translation. In *HLT-NAACL*.
- Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. [Deterministic Non-Autoregressive Neural Sequence Modeling by Iterative Refinement](#). *arXiv:1802.06901 [cs, stat]*. ArXiv: 1802.06901.
- Jindich Libovick and Jindich Helcl. 2018. [End-to-End Non-Autoregressive Neural Machine Translation with Connectionist Temporal Classification](#). *arXiv:1811.04719 [cs]*. ArXiv: 1811.04719.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. 2013. Playing atari with deep reinforcement learning. *CoRR*, abs/1312.5602.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nature*, 518:529–533.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *ACL 2016*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NIPS, 2014*, pages 3104–3112.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *CoRR*, abs/1706.03762.
- Chunqi Wang, Ji Zhang, and Haiqing Chen. 2018. [Semi-Autoregressive Neural Machine Translation](#). *arXiv:1808.08583 [cs]*. ArXiv: 1808.08583.
- Yiren Wang, Fei Tian, Di He, Tao Qin, ChengXiang Zhai, and Tie-Yan Liu. 2019. [Non-Autoregressive Machine Translation with Auxiliary Regularization](#). *arXiv e-prints*, page arXiv:1902.10245.
- Junyuan Xie, Ross B. Girshick, and Ali Farhadi. 2016. Unsupervised deep embedding for clustering analysis. In *ICML*.