

DePA: Improving Non-autoregressive Machine Translation with Dependency-Aware Decoder

Jiaao Zhan¹, Qian Chen, Boxing Chen, Wen Wang, Yu Bai¹, Yang Gao^{1*}

¹School of Computer Science and Technology,
Beijing Institute of Technology, Beijing, China

jiaao_zhan@163.com

{lukechan1231, chenboxing, wwang.969803}@gmail.com

{yubai, gyang}@bit.edu.cn

Abstract

Non-autoregressive machine translation (NAT) models have lower translation quality than autoregressive translation (AT) models because NAT decoders do not depend on previous target tokens in the decoder input. We propose a novel and general **Dependency-Aware Decoder (DePA)** to enhance target dependency modeling in the decoder of fully NAT models from two perspectives: decoder self-attention and decoder input. First, we propose an autoregressive forward-backward pre-training phase before NAT training, which enables the NAT decoder to gradually learn bidirectional target dependencies for the final NAT training. Second, we transform the decoder input from the source language representation space to the target language representation space through a novel attentive transformation process, which enables the decoder to better capture target dependencies. DePA can be applied to any fully NAT models. Extensive experiments show that DePA consistently improves highly competitive and state-of-the-art fully NAT models on widely used WMT and IWSLT benchmarks by up to **1.88 BLEU** gain, while maintaining the inference latency comparable to other fully NAT models.¹

1 Introduction

Autoregressive translation (AT) systems achieve state-of-the-art (SOTA) performance for neural machine translation (NMT) and Transformer (Vaswani et al., 2017) encoder-decoder is the prevalent architecture. In AT systems, each generation step depends on previously generated tokens, resulting in high inference latency when output is long. Non-autoregressive translation (NAT) models (Gu et al., 2018) significantly accelerate inference by generating all target tokens independently and simultaneously. However, this independence assumption

leads to degradation in accuracy compared to AT models, as NAT models cannot properly learn target dependencies. *Dependency* in prior works and our work takes its standard definition in NLP, i.e., syntactic relations between words in a sentence.

The mainstream NAT models fall into two categories: iterative NAT models and fully NAT models. Iterative NAT models (Gu et al., 2019; Ghazvininejad et al., 2019; Lee et al., 2018) improve translation accuracy by iteratively refining translations at the expense of slower decoding speed. In contrast, fully NAT models (Gu et al., 2018; Bao et al., 2022) have great latency advantage over AT models by making parallel predictions with a single decoding round, but they suffer from lower translation accuracy. In this paper, we aim at improving the translation accuracy of **fully NAT models** while preserving their latency advantage.

Previous research (Gu and Kong, 2021) argues that reducing dependencies is crucial for training a fully NAT model effectively, as it allows the model to more easily capture target dependencies. However, dependency reduction limits the performance upper bound of fully NAT models, since models may struggle to generate complex sentences. Previous studies show that multi-modality (Ran et al., 2020) is the main problem that NAT models suffer from (Huang et al., 2021; Bao et al., 2022), i.e., the target tokens may be generated based on different possible translations, often causing over-translation (token repetitions), under-translation (source words not translated), and wrong lexical choice for polysemous words. Table 1 Row3 shows all three multi-modality error types from the highly competitive fully NAT model GLAT (Qian et al., 2021) with modeling only forward dependency (F-NAT) in our experiments. We observe that **lack of complete dependency modeling could cause multi-modality errors**. For example, for the source text (in German) “*Woher komme ich?*” in the last column of Table 1, “*Woher*” means both “*where*” and “*how*”.

*Corresponding Author

¹We released our code at: https://github.com/zhanjiaao/NAT_DePA.

	Under-Translation	Over-Translation	Wrong Lexical Choice
Source	Wir haben es <u>weltweit</u> in 300 Gemeinden gemacht.	Einige leute wollten ihn einfach König nennen	Woher komme ich ? Wer bin ich ?
Target Reference	We 've done it in 300 communities around the world.	Some people just wanted to call him King .	Where am I from ? Who am I ?
F-NAT	We did it the world in 300 communities.	Some people just wanted to call <u>him him</u> king.	<u>How</u> do I come from ? Who am I ?
FB-NAT	We 've done it in 300 communities around the world.	Some people just wanted to call him king.	Where do I come from? Who am I ?

Table 1: Case studies of our proposed **FBD** approach on the highly competitive fully NAT model GLAT (Qian et al., 2021) for alleviating three types of multi-modality errors on the IWSLT16 DE-EN validation set. Repetitive tokens are in red. Source words that are not semantically translated are in bold and underlined. Wrong lexical choices (for polysemous words) and redundant words are in blue. **F-NAT** denotes only modeling forward dependencies while **FB-NAT** denotes modeling both forward and backward dependencies, the same as the models in Table 5. Case studies of our proposed **IT** approach are in Appendix.

The NAT model modeling only forward dependency (F-NAT) incorrectly translates “*woher*” into “*how*” and outputs “*How do I come from?*”; whereas the model modeling both forward and backward dependency (**FB-NAT**) translates it correctly into “*Where do I come from?*”. Therefore, instead of dependency reduction, we propose a novel and general Dependency-Aware Decoder (**DePA**), which enhances the learning capacity of fully NAT models and enables them to learn *complete* and *complex* forward and backward target dependencies in order to alleviate the multi-modality issue.

Firstly, we enhance the NAT decoder to learn complete target dependencies by exploring decoder self-attention. We believe that previous works (Guo et al., 2020a) incorporating only forward dependency modeled by AT models into NAT models are inadequate to address multi-modality. Therefore, we propose an effective forward-backward dependency modeling approach, denoted by **FBD**, as an auto-aggressive forward-backward pre-training phase before NAT training, using curriculum learning. The FBD approach implements *triangular attention masks* and takes different decoder inputs and targets *in a unified framework* to train the model to attend to previous or future tokens and learn both forward or backward dependencies.

Secondly, we enhance target dependency modeling within the NAT decoder from the perspective of the decoder input. Most prior NAT models (Gu et al., 2018; Wang et al., 2019; Wei et al., 2019) use a copy of the source text embedding as the decoder input, which is independent from the target representation space and hence makes target dependency modeling difficult. We transform the initial decoder input from the source language representation space to the target language representation space through a novel attentive transformation process, denoted by **IT**. Previous works on transforming the decoder input cannot guarantee that the decoder input is in the exact target representa-

tion space, resulting in differences from the true target-side distribution. Our proposed **IT** ensures that the decoder input is in the *exact* target representation space hence enables the model to better capture target dependencies.

Our contributions can be summarized as follows: (1) We propose a novel and general **Dependency-Aware Decoder (DePA)** for fully NAT models. For DePA, we propose a novel approach **FBD** for learning both forward and backward dependencies in NAT decoder, through which the target dependencies can be better modeled. **To the best of our knowledge, our work is the first to successfully model both forward and backward target-side dependencies explicitly for fully NAT models.** We also propose a novel decoder input transformation approach (**IT**). **IT** could ease target-side dependency modeling and enhance the effectiveness of FBD. DePA is **model-agnostic** and can be applied to any fully NAT models. (2) Extensive experiments on WMT and IWSLT benchmarks demonstrate that our DePA consistently improves the representative vanilla NAT model (Gu et al., 2018), the highly competitive fully NAT model GLAT (Qian et al., 2021) and the current SOTA of fully NAT models, CTC w/ DSLP & Mixed Training (denoted by **CTC-DSLP-MT**) (Huang et al., 2021) (**DSLP** denotes Deep Supervision and Layer-wise Prediction), by up to **+0.85 BLEU** on the SOTA CTC-DSLP-MT, **+1.88 BLEU** on GLAT, and **+2.2 BLEU** on vanilla NAT, while reserving inference latency as other fully NAT models, about **15 \times** speed-up over AT models. Experiments show that DePA achieves greater BLEU gains with less speed-up loss than DSLP when applied to various fully NAT models.

2 Related Work

Forward and Backward Dependencies Prior works explore bidirectional decoding to improve modeling of both forward and backward depen-

dependencies in phrase-based statistical MT (Finch and Sumita, 2009) and RNN-based MT (Zhang et al., 2018). For NAT, Guo et al. (2020a) and Wei et al. (2019) use forward auto-regressive models to guide NAT training. Liu et al. (2020) introduces an intermediate semi-autoregressive translation task to smooth the shift from AT training to NAT training. However, backward dependencies are rarely investigated in NAT.

Decoder Input of Fully NAT Models The decoder input of AT models consists of previously generated tokens. However, selecting appropriate decoder input for fully NAT models could be challenging. Most prior NAT models (Gu et al., 2018; Wang et al., 2019; Wei et al., 2019) use uniform copy (Gu et al., 2018) or soft copy (Wei et al., 2019) of the source text embedding as the decoder input, which is independent of the target representation space hence hinders target dependency modeling. Methods such as GLAT (Qian et al., 2021) and (Guo et al., 2020a,b) attempt to make the NAT decoder input similar to the target representation space by substituting certain positions in the decoder input with the corresponding target embedding. However, this creates a mismatch between training and inference. Guo et al. (2019) uses phrase-table lookup and linear mapping to make the decoder input closer to the target embedding, but this method still causes difference between the decoder input and the real target-side distribution.

Fully NAT Models To address multi-modality for fully NAT models, various approaches are proposed. Gu et al. (2018) uses knowledge distillation (KD) (Kim and Rush, 2016) to reduce dataset complexity. Libovický and Helcl (2018) and Saharia et al. (2020) use connectionist temporal classification (CTC) (Graves et al., 2006) for latent alignment. Sun et al. (2019) utilizes CRFs to model target positional contexts. Kaiser et al. (2018), Ma et al. (2019) and Shu et al. (2020) incorporate latent variables to guide generation, similar to VAEs (Kingma and Welling, 2013). Guo et al. (2020c) initializes NAT decoders with pretrained language models. Huang et al. (2021) proposes CTC with Deep Supervision and Layer-wise Prediction and Mixed Training (CTC-DSLMP-MT), setting new SOTA for fully NAT models on WMT benchmarks. DA-Transformer (Huang et al., 2022) represents hidden states in a directed acyclic graph to capture dependencies between tokens and gener-

ate multiple possible translations. In contrast, our DePA utilizes forward-backward pre-training and a novel attentive transformation of decoder input to enhance target dependency modeling. Under same settings and with KD, DA-Transformer performs only comparably to CTC-DSLMP-MT; however, performance of DA-Transformer benefits notably from *Transformer-big* for KD while CTC-DSLMP-MT uses *Transformer-base* for KD. DDRS w/ NMLA (Shao and Feng, 2022) benefits greatly from using diverse KD references while CTC-DSLMP-MT uses only a single KD reference. Hence, **CTC-DSLMP-MT is still the current SOTA for fully NAT models on WMT benchmarks.**

Non-autoregressive Models Besides fully NAT models, iterative NAT models are proposed such as iterative refinement of target sentences (Lee et al., 2018), masking and repredicting words with low probabilities (Ghazvininejad et al., 2019), edit-based methods to iteratively modify decoder output (Stern et al., 2019; Gu et al., 2019), and parallel refinement of every token (Kasai et al., 2020). Iterative NAT models improve translation accuracy at the cost of slower speed. Non-autoregressive models are practically important due to high efficiency. Other than MT, they are applied to various tasks such as image captioning (Gao et al., 2019), automatic speech recognition (Chen et al., 2019), and text-to-speech synthesis (Oord et al., 2018).

3 Methodology

3.1 Problem Formulation

NMT can be formulated as a sequence-to-sequence generation problem. Given a sequence $X = \{x_1, \dots, x_N\}$ in the source language, a sequence $Y = \{y_1, \dots, y_T\}$ in the target language is generated following the conditional probability $P(Y|X)$. NAT models are proposed to speed up generation by decoding all the target tokens in parallel, using conditional independent factorization as:

$$P_{NA}(Y|X) = P_L(T|x_{1:N}) \cdot \prod_{t=1}^T P(y_t|x_{1:N}; \theta) \quad (1)$$

where the target sequence length T is modeled by the conditional distribution P_L , and dependence on previous target tokens is removed. Compared to AT models, NAT models speed up inference significantly at the expense of translation quality, because the conditional independence assumption

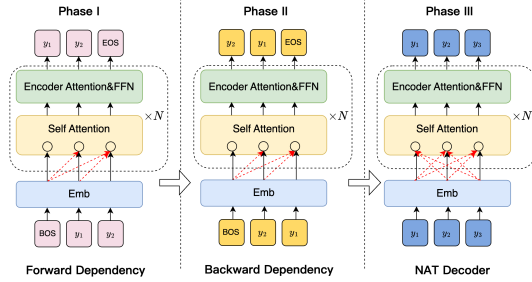


Figure 1: The proposed forward-backward dependency modeling (FBD) with triangular attention masks in a unified framework. The red dashed lines indicate the attention masks. We use different colors to highlight the difference of inputs and targets in each phase.

in Eq.1 enables parallel processing but lacks explicit modeling of dependency between target tokens. To enhance target dependency modeling, we propose two innovations as incorporating both forward and backward dependency modeling into the training process (Section 3.2) and transforming the decoder input into the target representation space (Section 3.3).

3.2 Target Dependency Modeling with Curriculum Learning (FBD)

Prior work (Guo et al., 2020a) utilizes forward dependency in AT models to initialize model parameters for NAT. However, as discussed in Section 1, for fully NAT models, only modeling forward dependency is inadequate for addressing the multimodality problem (Finch and Sumita, 2009; Zhang et al., 2018) (the Row for F-NAT in Table 1). Our innovations include incorporating both forward and backward dependency modeling into NAT models, via *triangular attention masks* in a *unified framework* through *curriculum learning* (Figure 1), and investigating efficacy of different curricula. In Figure 1, the **NAT decoder** phase denotes standard NAT training of any NAT decoder Dec . The **Forward Dependency** and **Backward Dependency** phases serve pre-training for NAT training, learning left-to-right and right-to-left dependencies to initialize NAT models with better dependencies. Forward Dependency and Backward Dependency training phases apply *the same upper triangle attention mask* on Dec . We use KD data from AT models for each phase but the inputs and the targets are different. The Forward Dependency training phase uses y_1 to predict y_2 and so on. The Backward Dependency training phase reverses the target sequence and uses y_2 to predict y_1 and so on. The NAT Train-

ing phase uses features of each word to predict the word itself. We make the following hypotheses: (1) Considering the nature of languages, learning forward dependency in Phase 1 is easier for the model for language generation. (2) Modeling backward dependency relies on learned forward dependency knowledge, hence it should be in the second phase. In fact, we observe the interesting finding that **the best curriculum remains forward-backward-forward-NAT (FBF-NAT) for both left-branching and right-branching languages**, proving our hypotheses. We speculate that NAT training may benefit from another forward dependency modeling in Phase 3 because the order of left-to-right is more consistent with characteristics of natural languages, hence adding the second forward dependency modeling after FB (i.e., FBF) smooths the transition to the final NAT training. Detailed discussions are in Section 4.3.

3.3 Decoder Input Transformation (IT) for Target Dependency Modeling

Given the initial decoder input z as a copy of source text embedding, we propose to directly select relevant representations from target embedding to form a new decoder input z' (Figure 2). z is used as the query and the selection is implemented as a learnable attention module. The learnable parameters bridge the gap between training and inference while the selection guarantees consistency between the decoder input matrix and the target representation space (i.e., the output embedding matrix of the decoder). This way, the decoder input is in the *exact* target-side embedding space and more conducive to modeling target dependencies for NAT models than previous approaches using source text embedding or transformed decoder input.

Decoder Input Transformation To transform z into the target representation space, we apply attention mechanism between z and the output embedding matrix $Emb \in \mathbb{R}^{d \times v}$, where d and v denote sizes of hidden states and the target vocabulary. Since NAT models usually have embedding matrix Emb including both source and target vocabularies, first, we conduct a filtering process to remove source vocabulary (mostly not used by the decoder) from the decoder output embedding matrix (the linear layer before decoder softmax). We build a dictionary that contains only target-side tokens in the training set. We then use this dictionary to filter Emb and obtain the new output embedding matrix




Figure 2: The proposed Decoder Input Transformation (IT) from z to z' , where $z \in \mathbb{R}^{T \times d}$ denotes the initial decoder input copied from the source text embedding x_{emb} , T and d denote the length of the target text y and the size of hidden states, respectively. $Emb \in \mathbb{R}^{d \times v}$ denotes the output embedding matrix of the decoder (the target representation space), where v denotes the size of the target vocabulary.

of the decoder $Emb' \in \mathbb{R}^{d \times v'}$, where v' denotes size of the filtered vocabulary. This filtering process guarantees that Emb' is strictly from the target representation space. The attention process starts with a linear transformation:

$$z^l = W_q \cdot z \quad (2)$$

Next, the dot-product attention is performed on z^l (as query) and Emb' (as key and value):

$$Sim = \text{softmax}(z^l \cdot Emb') \quad (3)$$

Sim represents similarity between each z_i^l and each embedding in the target vocabulary. Finally, we compute a weighted sum z' of target embedding based on their similarity values:

$$z' = Sim \cdot Emb'^T \quad (4)$$

Since z' is a linear combination of Emb' which is strictly in the target representation space, z' is also strictly in the target representation space, hence using z' as the decoder input provides a more solid basis for target dependency modeling.

Target-side Embedding Compression To reduce the computational cost of IT, we propose a target-side embedding compression approach to compress the large target embedding matrix. We process Emb' through a linear layer to obtain a new target embedding $Emb^* \in \mathbb{R}^{d \times v^*}$:

$$Emb^* = (W_c \cdot Emb'^T)^T \quad (5)$$

where $W_c \in \mathbb{R}^{v^* \times v'}$ is trainable and the size of compressed vocabulary v^* is set manually. The result Emb^* is still in the target representation space.

Since we can manually set v^* as a relatively small number (e.g., 1000, 2000), the computational cost of the attention mechanism can be greatly reduced. We hypothesize that target-side embedding compression may also alleviate over-fitting on small datasets and confirm this hypothesis in Section 4.3.

4 Experiments

4.1 Experimental Setup

Datasets We compare our methods with prior works on widely used MT benchmarks for evaluating NAT models: WMT14 EN \leftrightarrow DE (4.5M pairs), WMT16 EN \leftrightarrow RO (610K pairs). Also, we use IWSLT16 DE-EN (196K pairs), IWSLT14 DE-EN (153K pairs), and SP EN-JA² (50K pairs) for further analysis. For WMT16 EN \leftrightarrow RO and IWSLT16 DE-EN, we adopt the processed data from (Lee et al., 2018). For WMT14 EN \leftrightarrow DE, we apply the same preprocessing and learn subwords as Gu and Kong (2021). For IWSLT14 DE-EN, we follow preprocessing in (Guo et al., 2019). For SP EN-JA, we use sentencepiece³ to tokenize the text into subword units following Chousa et al. (2019). Following prior works, we share the source and target vocabulary and embeddings in each language pair in Emb , except EN-JA. Also following prior works (Gu et al., 2018; Qian et al., 2021), all NAT models in our experiments are trained on data generated from **pre-trained AT Transformer-base** with *sequence-level knowledge distillation (KD)* for all datasets except EN-JA.

²https://github.com/odashi/small_parallel_enja

³<https://github.com/google/sentencepiece>

Baselines and Training We implement the baseline models based on their released codebases. We implement the representative vanilla NAT (Gu et al., 2018; Qian et al., 2021; Huang et al., 2021)⁴, the highly competitive fully NAT model GLAT (Qian et al., 2021)⁵, and **current fully NAT SOTA CTC w/ DSLP & Mixed Training (CTC-DSLP-MT)** (Huang et al., 2021)⁶ and apply our methods to them. Following Qian et al. (2021), we use base-Transformer ($d_{model}=512$, $n_{head}=8$, $n_{layer}=6$) for WMT datasets and small-Transformer ($d_{model}=256$, $n_{head}=4$, $n_{layer}=5$) for IWSLT and SP EN-JA datasets. We use the same training setup for training the three models, Vanilla NAT, GLAT, and CTC-DSLP-MT as in their original papers cited above. We train models with batches of 64K tokens for WMT datasets, and 8K tokens for IWSLT and SP EN-JA datasets, using NVIDIA V100 GPUs. For GLAT, we use Adam optimizer (Kingma and Ba, 2015) with $\beta = (0.9, 0.999)$ and set dropout rate to 0.1. For Vanilla NAT and CTC-DSLP-MT, we use Adam optimizer (Kingma and Ba, 2015) with $\beta = (0.9, 0.98)$. For WMT datasets, the learning rate warms up to $5e-4$ in 4K steps and gradually decays according to inverse square root schedule (Vaswani et al., 2017). As for IWSLT and SP EN-JA datasets, we adopt linear annealing (from $3e-4$ to $1e-5$) as in Lee et al. (2018). We choose the model with the best performance on the validation set as the final model and evaluate the final model on the test sets. For experiments using our method FBD (Section 3.2), we use the **FBF-NAT configuration** (as in Section 4.3) and train the same number of steps at each phase (including NAT training phase), with 300K steps for each phase for WMT datasets and 100K steps for each phase for IWSLT datasets and SP EN-JA. IT by default is **without Target-side Embedding Compression** (Section 3.3).

Evaluation To evaluate the translation accuracy, we use **SacreBLEU** (Post, 2018) for all experiments and ChrF (Popovic, 2015) (also using the SacreBLEU tool) additionally for ablation study on IWSLT benchmark. To evaluate the inference latency, following Gu and Kong (2021), we measure the wall-clock time for translating the entire WMT14 EN-DE test set with `batch_size=1` on a

single NVIDIA V100 GPU, then compute the average time per sentence. We report **Speed-up** based on the inference latency of *Transformer-base AT (teacher)* and fully NAT models.

4.2 Main Results

Table 2 shows the main results on the WMT benchmarks. For EN \leftrightarrow RO, we report the mean of BLEU from 3 runs with different random seeds for Row 12-13, all with quite small standard deviations (≤ 0.16)⁷. We apply our proposed DePA, which includes IT and FBD, to vanilla NAT, GLAT, and the current fully NAT SOTA CTC-DSLP-MT, on WMT, IWSLT, and EN-JA benchmarks. We use the same hyperparameters and random seeds to fairly compare two models. **It is crucial to point out that accuracies of vanilla NAT, GLAT, and CTC-DSLP-MT models have plateaued out after 300K training steps on WMT datasets hence original papers of these three models set max training steps to 300K.** We verify this observation in our own experiments as we also see no gains on these models after 300K training steps on the WMT datasets. Hence, although our DePA trains $300K \times 4 = 1200K$ steps on WMT datasets due to **FBF** pre-training as in Section 4.3, **all comparisons between baselines w/ DePA and w/o DePA are fair comparisons.** Table 2 shows that DePA consistently improves the translation accuracy for both vanilla NAT and GLAT on each benchmark, achieving **mean=+1.37 and max=+1.88** BLEU gain on GLAT and **mean=+2.34 and max=+2.46** BLEU gain on vanilla NAT. DePA also improves the SOTA CTC-DSLP-MT by **mean=+0.42 and max=+0.49** BLEU gain on the WMT test sets (Table 2), **+0.85** BLEU gain on the IWSLT16 DE-EN validation set and **+1.43** BLEU gain on the EN-JA test set (Table 3). All gains from DePA on vanilla NAT, GLAT, and CTC-DSLP-MT are **statistically significant** ($p < 0.05$) based on a paired bootstrap resampling test conducted using 1K resampling trials and the SacreBLEU tool.

Table 2 also shows that on each benchmark, the average improvement from DePA on three models (vanilla NAT, GLAT, and CTC-DSLP-MT) is within **[0.90,1.56]** (Row15), always larger than the average improvement from w/DSLP on

⁴https://github.com/facebookresearch/fairseq/tree/main/examples/nonautoregressive_translation

⁵<https://github.com/FLC777/GLAT>

⁶<https://github.com/chenyangh/DSLP>

⁷WMT14 EN \leftrightarrow DE is much larger than WMT16 EN \leftrightarrow RO. Since standard deviations of BLEU from multiple runs with different random seeds on WMT14 EN \leftrightarrow DE are very small, ≤ 0.08 (Huang et al., 2022), following prior works, we report single-run BLEU on WMT14 EN \leftrightarrow DE to save energy.

Row#	Models	Speed-up \uparrow	WMT'14		WMT'16	
			EN-DE	DE-EN	EN-RO	RO-EN
1	Transformer- <i>base</i> (teacher)	1.0 \times	27.48	31.39	33.70	34.05
2	+ KD	2.5 \times	27.34	30.95	33.52	34.01
3	Vanilla NAT	15.6 \times	20.36	24.81	28.47	29.43
4	w/ DSLP*	14.8 \times	22.72	25.83	30.48	31.46
5	w/ DePA (Ours)	15.4 \times	23.15	26.59	30.78	31.89
6	GLAT	15.3 \times	25.21	29.84	31.19	32.04
7	w/ DSLP*	14.9 \times	25.69	29.90	32.36	33.06
8	w/ DePA (Ours)	15.1 \times	26.43	30.42	33.07	33.82
10	CTC*	15.5 \times	25.72	29.89	32.89	33.79
11	w/ DSLP*	14.8 \times	26.85	31.16	33.85	34.24
12	w/ DSLP & Mixed Training	14.8 \times	27.02	31.61	33.99	34.42
13	w/ DSLP & Mixed Training & w/ DePA (Ours)	14.7 \times	27.51	31.96	34.48	34.77
14	Average improvement from DSLP	-	1.32	0.78	1.38	1.17
15	Average improvement from DePA (Ours)	-	1.50	0.90	1.56	1.53

Table 2: BLEU and Speed-up from our **DePA** and existing methods on WMT benchmark test sets. **Speed-up** is measured on WMT14 EN-DE test set. BLEUs without rescoring are reported, with the best BLEU scores in bold for each group. * denotes the results are copied from previous work (Huang et al., 2021), other results are obtained by our implementation. Average improvements of DSLP are re-calculated using our results, which are slightly different from Table 1 in (Huang et al., 2021).

them, **[0.78,1.38]** (Row14). DePA brings consistent improvement over SOTA CTC-DSLP-MT on all benchmarks (Table 2 Row13-over-Row12, Table 3), hence we expect DePA to also improve DA-Transformer (Huang et al., 2022) and DDRS w/ NMLA (Shao and Feng, 2022) and will verify this w/ and w/o KD in future work. **Applying DePA to fully NAT models retains the inference speed-up advantages of fully NAT models.** Applying DePA to vanilla NAT, GLAT, and SOTA CTC-DSLP-MT obtain 15.4 \times , 15.1 \times , and 14.7 \times speed-up over the autoregressive Transformer-base (teacher) (Row1). Overall Table 2 shows that **DePA achieves greater BLEU gains with less speed-up loss than DSLP on all baselines.** These results demonstrate superiority of DePA over DSLP on improving other fully NAT models.

4.3 Analysis

Ablation Study We analyze the respective efficacy of IT and FBD in DePA on the IWSLT16 DE-EN validation and the WMT and SP EN-JA test sets. Table 3 shows that FBD and IT improve GLAT by **+1.26 BLEU/+1.5 ChrF** and **+0.34 BLEU/+1.0 ChrF** on IWSLT16 DE-EN validation set, respectively. Considering that GLAT w/FBD has more training steps than GLAT, we also train GLAT (400K steps) which has the same training steps as GLAT w/FBD for fair comparison. **Similar to findings on WMT datasets, we observe plateaus of accuracy on IWSLT and EN-JA datasets from more training steps than the original 100K.** Just training more steps hardly improves the baseline

(only +0.07 BLEU gain) on IWSLT16 DE-EN, whereas GLAT w/FBD brings **+1.19 BLEU/+1.2 ChrF** gains over GLAT (400K steps).

Table 4 shows our IT outperforms Linear Mapping (Guo et al., 2019) by **+2.31 BLEU** gain on IWSLT14 DE-EN test set. IT has the same number of extra parameters as Linear Mapping. Hence, the large gain proves that improvements from IT are not just from additional layers. The number of extra parameters of IT, as from W_q in Eq.2, is quite small: $512*512=262144$ for Transformer-base on WMT datasets and $256*256=65536$ for Transformer-small on IWSLT datasets. The large BLEU gain **+3.18** from applying IT to vanilla NAT proves vanilla transformer decoder cannot achieve similar transformation effectiveness as IT. Table 3 shows that for language pairs with different levels of source-target vocabulary sharing, such as WMT EN-DE and DE-EN, IWSLT DE-EN, EN-RO, and EN-JA, our IT method can achieve consistent improvements over GLAT and CTC-DSLP-MT. Applying IT consistently improves GLAT and CTC-DSLP-MT although these gains are smaller than gain on vanilla NAT. This is because decoder input of vanilla NAT only replicates source embedding, whereas GLAT and CTC-DSLP-MT already transform decoder input by replacing selected positions in decoder input with target embedding, hence reducing improvements of IT. Still, gains from w/IT+FBD over w/FBD confirms our hypothesis that IT can enhance effectiveness of FBD. On GLAT, IT+FBD yields **+1.4 BLEU/+2.7 ChrF** gains on IWSLT16 DE-EN and **+1.43 BLEU**

Models	IWSLT16 DE-EN		WMT'14 EN-DE DE-EN		WMT'16 EN-RO RO-EN	
	BLEU	ChrF	BLEU		BLEU	
CTC-DSLP-MT	31.04	56.7	27.02	31.61	34.17	34.60
CTC-DSLP-MT w/ IT	31.29	57.1	27.21	31.78	34.32	34.71
CTC-DSLP-MT w/ FBD	31.73	57.5	27.44	31.90	34.60	34.92
CTC-DSLP-MT w/ IT+FBD	31.89	57.8	27.51	31.96	34.68	34.98

Models	IWSLT16 DE-EN		EN-JA
	BLEU	ChrF	BLEU
GLAT	29.61	51.8	27.67
GLAT (400K step)	29.68	52.1	–
GLAT w/ IT	29.95	52.8	27.95
GLAT w/ FBD	30.87	53.3	28.87
GLAT w/ IT+FBD	31.01	54.5	29.10

Table 3: Effect of **IT** and **FBD** and **IT+FBD** (i.e., **DePA**) on the IWSLT16 DE-EN validation set, the WMT and SP EN-JA test sets. We report **mean of BLEU/ChrF** from 3 runs with different random seeds. BLEU gains from DePA on SOTA CTC-DSLP-MT on each set, [**0.85, 0.49, 0.51**], are larger than std (≤ 0.17).

Models	BLEU
Vanilla NAT (Guo et al., 2019)	22.95
Vanilla NAT w/ Linear Mapping (Guo et al., 2019)	24.13
Vanilla NAT (our implementation)	23.26
Vanilla NAT w/ IT	26.44

Table 4: Compare **IT** and Linear Mapping (Guo et al., 2019) on vanilla NT on the IWSLT14 DE-EN test set.

on EN-JA and on SOTA CTC-DSLP-MT, **+0.85 BLEU/+1.1 ChrF** gain on IWSLT16 DE-EN.

To further analyze IT, we compare cosine similarity between the target embedding against the original decoder input and the transformed decoder input, respectively. For each sample in the IWSLT16 DE-EN validation set, we average all its token embeddings as the decoder input representation and the same for the target representation and then compute cosine similarity. We average similarities of all samples as the final similarity. We find that **IT significantly improves similarity between the decoder input and the target representation**, $0.04951 \rightarrow 0.14521$ for GLAT and $0.04837 \rightarrow 0.14314$ for vanilla NAT.

Impact of Different Dependency Curricula in FBD Table 5 presents results from applying different forward-backward dependency modeling curricula (Figure 1) on GLAT on the IWSLT16 DE-EN validation and the SP EN-JA test sets. Compared with modeling backward dependency in Phase 1 (B-NAT and BF-NAT), modeling forward dependency in Phase 1 (F-NAT, FB-NAT, and FBF-NAT) performs notably better. FB-NAT outperforms BF-NAT by **+3.04 BLEU** on IWSLT16

DE-EN and **+2.08 BLEU** on EN-JA. It seems that forward dependency modeling achieves good initialization for subsequent training phases, while backward dependency modeling cannot. We observe the best curriculum as **FBF-NAT**, i.e., first learn forward dependency, next learn backward dependency, then another round of forward dependency training before NAT training. Table 5 shows the same trend of curricula on SP EN-JA as on IWSLT16 DE-EN, with FBF-NAT performing best, demonstrating that this trend of forward-backward dependency modeling curricula is consistent for both right-branching (English) and left-branching (Japanese) target languages. All these observations confirm our hypotheses in Section 3.2. Our FBF-NAT consistently outperforms baseline GLAT (denoted by NAT in Table 5) by **+1.58** on IWSLT16 DE-EN and **+1.56** on SP EN-JA and outperforms prior works modeling forward dependency only (Guo et al., 2020a) on GLAT (denoted by F-NAT in Table 5) by **+1.15** on DE-EN and **+1.08** on EN-JA.

DePA on Raw Data We evaluate DePA on raw data by training models on the original training set without KD (Section 4.1). DePA improves GLAT on the IWSLT16 DE-EN validation set by **+1.57 BLEU** ($26.57 \rightarrow 28.14$), proving that DePA effectively enhances the dependency modeling ability of fully NAT models hence **reduces dependence of NAT training on AT models**.

Effectiveness of Target-side Embedding Compression We propose a linear compression module to reduce the selection candidates of the target

IWSLT16 DE-EN validation set				SP EN-JA test set			
Models	BLEU	Models	BLEU	Models	BLEU	Models	BLEU
NAT	29.61	BF-NAT	27.83	NAT	27.67	BF-NAT	26.79
F-NAT	30.04	FB-NAT	30.87	F-NAT	28.15	FB-NAT	28.87
B-NAT	27.05	FBF-NAT	31.19	B-NAT	25.83	FBF-NAT	29.23

Table 5: BLEU from different dependency modeling curricula on GLAT. Best results for each set are in bold. **NAT** denotes GLAT baseline. **F** and **B** denote forward dependency and backward dependency phase respectively (Figure 1). For example, F-NAT denotes forward dependency training then NAT training.




Figure 3: Visualization of the decoder self-attention distribution in NAT models on IWSLT16 DE-EN validation set. Definitions of model names are the same as in Table 5.

embedding for IT (Section 3.3). We use the dichotomy to determine the compression dimension interval [1000, 2000] and evaluate GLAT w/ IT using different dimensions with step size 200 in this interval for IT on the IWSLT16 DE-EN validation set. As shown in Table 6, applying IT on GLAT improves BLEU up to **+0.78** (29.61 \rightarrow 30.39) with compressed dimension 1800. We also experiment with target-side embedding compression on a larger model on WMT16 EN-RO but find no gains. We assume that for relatively small models and data, this approach helps filter out some redundant target information, hence refines the target representation space and improves the translation accuracy.

4.4 Case Study and Visualization

Table 1 presents case studies of GLAT w/ FBD (“FB-NAT”) and with only forward modeling (“F-NAT”) on IWSLT16 DE-EN validation set. Some typical multi-modality errors in F-NAT predictions

Compressed Dimension	w/o IT	1000	1200	1400	1600	1800	2000
BLEU	29.61	29.45	29.56	29.77	29.85	30.39	29.14

Table 6: BLEU from GLAT w/ IT on the IWSLT16 DE-EN validation set with Target-side Embedding Compression described in Section 3.3.

are corrected by incorporating both forward and backward dependency modeling through FBD.

For a more intuitive analysis of FBD, we present a visualization of the decoder self-attention distribution of different NAT models in Figure 3. All models are based on GLAT and model names conform to those in Table 5. In the baseline GLAT (Figure 3a), the self-attention distribution of each position is scattered in adjacent positions, indicating that the NAT model lacks dependency and has high confusion during decoding, causing multi-modality errors. In F-NAT and B-NAT models, significant forward and backward dependencies can be observed in Figure 3b and 3c, indicating that these two models can better use information in previous or future positions. Encouragingly, forward and backward dependencies are fused in the FB-NAT model (Figure 3d), which can focus on future information while modeling forward dependency, capable of alleviating problems shown in Table 1.

5 Conclusion

We propose a novel and general Dependency-Aware Decoder (DePA) to enhance target dependency modeling for fully NAT models, with forward-backward dependency modeling and decoder input transformation. Extensive experiments show that DePA improves the translation accuracy of highly competitive and SOTA fully NAT models while preserving their inference latency. In future work, we will evaluate DePA on iterative NAT models such as Imputer, CMLM, and Levenshtein Transformer and incorporate ranking approaches into DePA.

6 Limitations

Apart from all the advantages that our work achieves, some limitations still exist. Firstly, in this work, we investigate the efficacy of applying our proposed DePA approach on the representative vanilla NAT, the highly competitive fully NAT model GLAT and current SOTA CTC-DSLP-MT for fully NAT models, but we have yet to apply DePA to iterative NAT models, such as Imputer (Saharia et al., 2020), CMLM (Ghazvininejad et al., 2019), and Levenshtein Transformer (Gu et al., 2019). Hence, the effectiveness of DePA on iterative NAT models still needs to be verified. Secondly, we have not yet incorporated re-ranking approaches such as Noisy Parallel Decoding (NPD) (Gu et al., 2018) into DePA. Thirdly, our proposed method FBD requires multiple additional training phases before NAT training, resulting in longer training time and using more GPU resources. Reducing the computational cost of FBD training is one future work that will be beneficial for energy saving. Last but not least, NAT models have limitations on handling long text. They suffer from worse translation quality when translating relatively long text. We plan to investigate all these topics in future work.

References

- Yu Bao, Hao Zhou, Shujian Huang, Dongqi Wang, Lihua Qian, Xinyu Dai, Jiajun Chen, and Lei Li. 2022. **GLAT: glancing at latent variables for parallel text generation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 8398–8409. Association for Computational Linguistics.
- Nanxin Chen, Shinji Watanabe, Jesús Villalba, and Najim Dehak. 2019. Listen and fill in the missing letters: Non-autoregressive transformer for speech recognition. *arXiv preprint arXiv:1911.04908*.
- Katsuki Chousa, Katsuhito Sudoh, and Satoshi Nakamura. 2019. Simultaneous neural machine translation using connectionist temporal classification. *arXiv preprint arXiv:1911.11933*.
- Andrew M. Finch and Eiichiro Sumita. 2009. **Bidirectional phrase-based statistical machine translation**. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, 6-7 August 2009, Singapore, A meeting of SIG-DAT, a Special Interest Group of the ACL*, pages 1124–1132. ACL.
- Junlong Gao, Xi Meng, Shiqi Wang, Xia Li, Shanshe Wang, Siwei Ma, and Wen Gao. 2019. Masked non-autoregressive image captioning. *arXiv preprint arXiv:1906.00717*.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. Mask-predict: Parallel decoding of conditional masked language models. *arXiv preprint arXiv:1904.09324*.
- Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. 2006. **Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks**. In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148 of *ACM International Conference Proceeding Series*, pages 369–376. ACM.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2018. Non-autoregressive neural machine translation. In *International Conference on Learning Representations*.
- Jiatao Gu and Xiang Kong. 2021. **Fully non-autoregressive neural machine translation: Tricks of the trade**. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 120–133, Online. Association for Computational Linguistics.
- Jiatao Gu, Changan Wang, and Jake Zhao. 2019. Levenshtein transformer. *arXiv preprint arXiv:1905.11006*.
- Junliang Guo, Xu Tan, Di He, Tao Qin, Linli Xu, and Tie-Yan Liu. 2019. Non-autoregressive neural machine translation with enhanced decoder input. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3723–3730.
- 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. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 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.
- Junliang Guo, Zhirui Zhang, Linli Xu, Hao-Ran Wei, Boxing Chen, and Enhong Chen. 2020c. Incorporating bert into parallel sequence decoding with adapters. In *NeurIPS*.
- Chenyang Huang, Hao Zhou, Osmar R Zaiane, Lili Mou, and Lei Li. 2021. Non-autoregressive translation with layer-wise prediction and deep supervision. *arXiv preprint arXiv:2110.07515*.

- Fei Huang, Hao Zhou, Yang Liu, Hang Li, and Minlie Huang. 2022. [Directed acyclic transformer for non-autoregressive machine translation](#). *CoRR*, abs/2205.07459.
- 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 *International Conference on Machine Learning*, pages 2390–2399. PMLR.
- Jungo Kasai, James Cross, Marjan Ghazvininejad, and Jiatao Gu. 2020. Non-autoregressive machine translation with disentangled context transformer. In *International Conference on Machine Learning*, pages 5144–5155. 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.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR (Poster)*.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- 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.
- 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.
- Jinglin Liu, Yi Ren, Xu Tan, Chen Zhang, Tao Qin, Zhou Zhao, and Tie-Yan Liu. 2020. [Task-level curriculum learning for non-autoregressive neural machine translation](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 3861–3867. ijcai.org.
- 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.
- Aaron Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, et al. 2018. Parallel wavenet: Fast high-fidelity speech synthesis. In *International conference on machine learning*, pages 3918–3926. PMLR.
- Maja Popovic. 2015. [chrF: character n-gram f-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation, WMT@EMNLP 2015, 17-18 September 2015, Lisbon, Portugal*, pages 392–395. The Association for Computer Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers, WMT 2018, Belgium, Brussels, October 31 - November 1, 2018*, pages 186–191. 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.
- 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.
- Chenze Shao and Yang Feng. 2022. Non-monotonic latent alignments for ctc-based non-autoregressive machine translation. *arXiv preprint arXiv:2210.03953*.
- 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 *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8846–8853.
- Mitchell Stern, William Chan, Jamie Kiros, and Jakob Uszkoreit. 2019. Insertion transformer: Flexible sequence generation via insertion operations. In *International Conference on Machine Learning*, pages 5976–5985. PMLR.
- Zhiqing Sun, Zhuohan Li, Haoqing Wang, Di He, Zi Lin, and Zhi-Hong Deng. 2019. Fast structured decoding for sequence models. In *NeurIPS*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.

Yiren Wang, Fei Tian, Di He, Tao Qin, ChengXiang Zhai, and Tie-Yan Liu. 2019. Non-autoregressive machine translation with auxiliary regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5377–5384.

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.

Xiangwen Zhang, Jinsong Su, Yue Qin, Yang Liu, Rongrong Ji, and Hongji Wang. 2018. [Asynchronous bidirectional decoding for neural machine translation](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5698–5705. AAAI Press.

Case study for the proposed IT NAT models generally suffer from the multi-modality problem, which shows as *over-translation* (repetition), *under-translation* (missing information), and *wrong lexical choice* (incorrect translations caused by polysemy) (Ran et al., 2020). As shown in Table 7, Vanilla NAT and GLAT tend to generate repetitive tokens which are highlighted in red (over-translation). Additionally, Vanilla NAT omits the translation of “schließlich” which is in bold and underlined (under-translation). By applying our method IT, the decoder input is closer to the target representation space and the model has a better perception for the target-side information, so that the repetition and under-translation problems can be effectively alleviated. As for incorrect translations caused by polysemy, as shown in Case #1 in Table 7, “Drucks” means both “printing” and “pressure” in German. GLAT mistakenly translates “Drucks” into “printing”, but our method can help the model correctly translate it into “pressure”. Particularly, in Case #2, “Bauplan” means “blueprint” in German. Although both baseline models Vanilla NAT and GLAT generate the correct words, they also generate the redundant word “plan” which is also a subword of “Bauplan”. These examples demonstrate that the baseline models may confuse the source representation space with the target representation space during generation, but our method IT effectively remedies this problem.

	Case #1	Case #2
Source	obwohl sie erwischt wurden , wurden sie <u>schließlich</u> freigelassen aufgrund immensen internationalen Drucks .	das ist ein Bauplan für Länder wie China und den Iran .
Target Reference	even though they were caught , they were eventually released after heavy international pressure .	this is a blueprint for countries like China and Iran .
Vanilla NAT	although they were caught , they were released released because because of huge drug .	this is a blueprint plan for countries like China and and Iran .
Vanilla NAT w/ IT	although they were caught , they were finally released because huge international pressure .	this is a blueprint for countries like China and Iran .
GLAT	although they were caught , they finally were released because of of international printing .	this is a blueprint plan for countries like China and Iran .
GLAT w/ IT	although they were caught , they were finally released after huge international pressure .	this is a blueprint for countries like China and Iran .

Table 7: Case studies of our method **IT** on the IWSLT16 DE-EN validation set by comparing the translations from the two baseline models Vanilla NAT and GLAT and from them after applying IT (models in bold). Repetitive tokens are in red. Source words that are not semantically translated are marked in bold and underlined (under-translation). Wrong lexical choice (incorrect translations caused by polysemy) and redundant words are in blue.