

Hint-Based Training for Non-Autoregressive Machine Translation

Zhuohan Li¹ Zi Lin^{2,3} Di He² Fei Tian⁴ Tao Qin⁵ Liwei Wang² Tie-Yan Liu⁵

¹UC Berkeley ²Peking University ³Google AI ⁴Facebook ⁵Microsoft Research

zhuohan@cs.berkeley.edu {zi.lin, di.he, wanglw}@pku.edu.cn

feitia@fb.com {taoqin, tyliu}@microsoft.com

Abstract

Due to the unparallelizable nature of the autoregressive factorization, AutoRegressive Translation (ART) models have to generate tokens sequentially during decoding and thus suffer from high inference latency. Non-AutoRegressive Translation (NART) models were proposed to reduce the inference time, but could only achieve inferior translation accuracy. In this paper, we proposed a novel approach to leveraging the hints from hidden states and word alignments to help the training of NART models. The results achieve significant improvement over previous NART models for the WMT14 En-De and De-En datasets and are even comparable to a strong LSTM-based ART baseline but one order of magnitude faster in inference.

1 Introduction

Neural machine translation has attracted much attention in recent years (Bahdanau et al., 2014, 2016; Kalchbrenner et al., 2016; Gehring et al., 2016). Given a sentence $x = (x_1, \dots, x_{T_x})$ from the source language, the straight-forward way for translation is to generate the target words $y = (y_1, \dots, y_{T_y})$ one by one from left to right. This is also known as the *AutoRegressive* Translation (ART) models, in which the joint probability is decomposed into a chain of conditional probabilities:

$$P(y|x) = \prod_{t=1}^{T_y} P(y_t|y_{<t}, x), \quad (1)$$

While the ART models have achieved great success in terms of translation quality, the time consumption during inference is still far away from satisfactory. During training, the predictions at different positions can be estimated in parallel since the ground truth pair (x, y) is exposed to the model. However, during inference, the model

has to generate tokens sequentially as $y_{<t}$ must be inferred on the fly. Such autoregressive behavior becomes the bottleneck of the computational time (Wu et al., 2016).

In order to speed up the inference process, a line of works begin to develop non-autoregressive translation models. These models break the autoregressive dependency by decomposing the joint probability with

$$P(y|x) = P(T_y|x) \prod_{t=1}^{T_y} P(y_t|x). \quad (2)$$

The lost of autoregressive dependency largely hurt the consistency of the output sentences, increase the difficulty in the learning process and thus lead to a low quality translation. Previous works mainly focus on adding different components into the NART model to improve the expressiveness of the network structure to overcome the loss of autoregressive dependency (Gu et al., 2017; Lee et al., 2018; Kaiser et al., 2018). However, the computational overhead of new components will hurt the inference speed, contradicting with the goal of the NART models: to parallelize and speed up neural machine translation models.

To tackle this, we proposed a novel hint-based method for NART model training. We first investigate the causes of the poor performance of the NART model. Comparing with the ART model, we find that: (1) the positions where the NART model outputs incoherent tokens will have very high hidden states similarity; (2) the attention distributions of the NART model are more ambiguous than those of ART model. Therefore, we design two kinds of hints from the hidden states and attention distributions of the ART model to help the training of the NART model. The experimental results show that our model achieves significant improvement over the NART baseline models and is even comparable to a strong ART baseline in Wu et al. (2016).

The work was performed at Microsoft Research Asia.

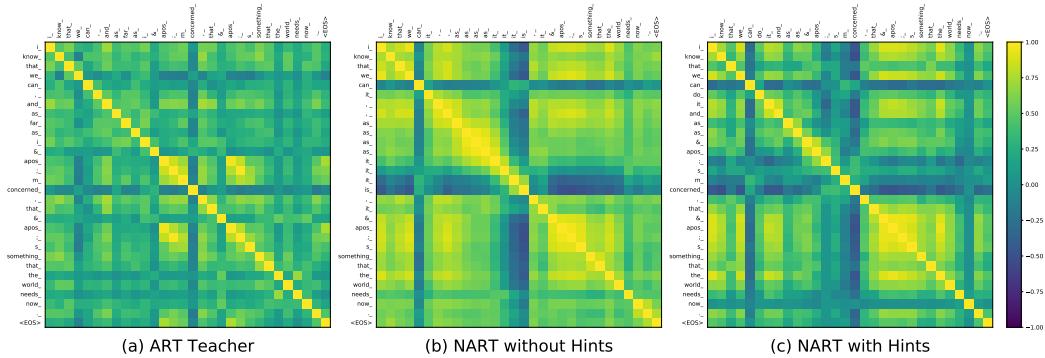


Figure 1: Case study: the above three figures visualize the hidden state cosine similarities of different models. The axes correspond to the generated target tokens. Each pixel shows the cosine similarities \cos_{ij} between the last layer hidden states of the i -th and j -th generated tokens, where the diagonal pixel will always be 1.0.

2 Approach

In this section, we first describe the observations on the ART and NART models, and then discuss what kinds of information can be used as hints to help the training of the NART model. We follow the network structure in Vaswani et al. (2017), use a copy of the source sentence as decoder input, remove the attention masks in decoder self-attention layers and add a positional attention layer as suggested in Gu et al. (2017). We provide a visualization of ART and NART models we used in Figure 3 and a detailed description of the model structure in Appendix.

2.1 Observation: Illed States and Attentions

According to the case study in Gu et al. (2017), the translations of the NART models contain incoherent phrases (e.g. repetitive words) and miss meaningful tokens on the source side, while these patterns do not commonly appear in ART models. After some empirical study, we find two non-obvious facts that lead to this phenomenon.

First, we visualize the cosine similarities between decoder hidden states of a certain layer in both ART and NART models for sampled cases. Mathematically, for a set of hidden states r_1, \dots, r_T , the pairwise cosine similarity can be derived by $\cos_{ij} = \langle r_i, r_j \rangle / (\|r_i\| \cdot \|r_j\|)$. We then plot the heatmap of the resulting matrix \cos . A typical example is shown in Figure 1, where the cosine similarities in the NART model are larger than those of the ART model, indicating that the hidden states across positions in the NART model are “similar”. Positions with highly-correlated hidden states tend to generate the same word and make the NART model output repetitive tokens,

e.g., the yellow area on the top-left of Figure 1(b), while this does not happen in the ART model (Figure 1(a)). According to our statistics, 70% of the cosine similarities between hidden states in the ART model are less than 0.25, and 95% are less than 0.5.

Second, we visualize the encoder-decoder attentions for sampled cases, shown in Figure 2. Good attentions between the source and target sentences are usually considered to lead to accurate translation while poor ones may cause wrong output tokens (Bahdanau et al., 2014). In Figure 2(b), the attentions of the ART model almost covers all source tokens, while the attentions of the NART model do not cover “farm” but with two “morning”. This directly makes the translation result worse in the NART model. These phenomena inspire us to use the intermediate hidden information in the ART model to guide the learning process of the NART model.

2.2 Hints from the ART teacher Model

Our study motivates us to leverage the intermediate hidden information from an ART model to improve the NART model. We focus on how to define hints from a well-trained ART *teacher* model and use it to guide the training process of a NART *student* model. We study layer-to-layer hints and assume both the teacher and student models have an M -layer encoder and an N -layer decoder, despite the difference in stacked components.

Without the loss of generality, we discuss our method on a given paired sentence (x, y) . In real experiments, losses are averaged over all training data. For the teacher model, we use $a_{t,l,h}^{tr}$ as the encoder-to-decoder attention distribution of h -th head in the l -th decoder layer at position t , and

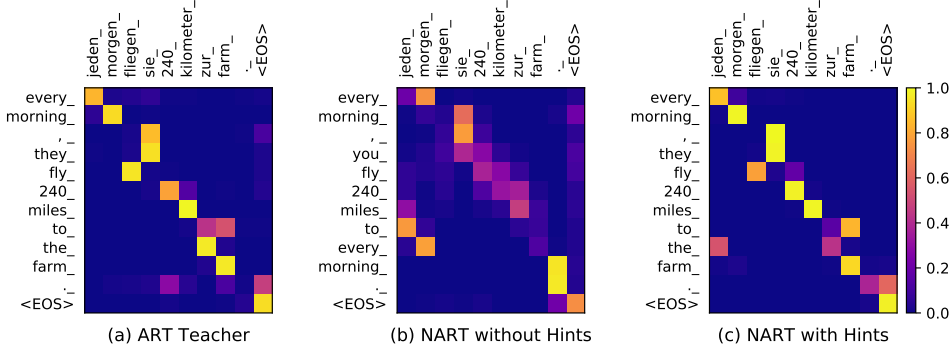


Figure 2: Case study: the above three figures visualize the encoder-decoder attention weights of different models. The x-axis and y-axis correspond to the source and generated target tokens respectively. The attention distribution is from a single head of the third layer encoder-decoder attention, which is the most informative one according to our observation. Each pixel shows attention weights α_{ij} between the i -th source token and j -th target token.

use $r_{t,l}^{tr}$ as the output of the l -th decoder layer after feed forward network at position t . Correspondingly, $a_{t,l,h}^{st}$ and $r_{t,l}^{st}$ are used for the student model. We propose a hint-based training framework that contains two kinds of hints:

Hints from hidden states The discrepancy of hidden states motivates us to use hidden states of the ART model as a hint for the learning process of the NART model. One straight-forward method is to regularize the L_1 or L_2 distance between each pair of hidden states in ART and NART models. However, since the network components are quite different in ART and NART models, applying the straight-forward regression on hidden states hurts the learning process and fails. Therefore, we design a more implicit loss to help the student refrain from the incoherent translation results by acting towards the teacher in the hidden-state level:

$$\mathcal{L}_{hid} = \frac{2}{(T_y - 1)T_y N} \sum_{s=1}^{T_y-1} \sum_{t=s+1}^{T_y} \sum_{l=1}^N \phi(d_{st}, d_{tr}),$$

where $d_{st} = \cos(r_{s,l}^{st}, r_{t,l}^{st})$, $d_{tr} = \cos(r_{s,l}^{tr}, r_{t,l}^{tr})$, and ϕ is a penalty function. In particular, we let

$$\phi(d_{st}, d_{tr}) = \begin{cases} -\log(1 - d_{st}), & \text{if } d_{st} \geq \gamma_{st} \\ & \text{and } d_{tr} \leq \gamma_{tr}; \\ 0, & \text{else,} \end{cases}$$

where $-1 \leq \gamma_{st}, \gamma_{tr} \leq 1$ are two thresholds controlling whether to penalize or not. We design this loss since we only want to penalize hidden states that are highly similar in the NART model, but not similar in the ART model. We have tested several choices of $-\log(1 - d_{st})$, e.g., $\exp(d_{st})$, from which we find similar experimental results.

Hints from word alignments We observe that meaningful words in the source sentence are sometimes untranslated by the NART model, and the corresponding positions often suffer from ambiguous attention distributions. Therefore, we use the word alignment information from the ART model to help the training of the NART model.

In particular, we minimize KL-divergence between the per-head encoder-to-decoder attention distributions of the teacher and the student to encourage the student to have similar word alignments to the teacher model, i.e.

$$\mathcal{L}_{align} = \frac{1}{T_y N H} \sum_{t=1}^{T_y} \sum_{l=1}^N \sum_{h=1}^H D_{KL}(a_{t,l,h}^{tr} \| a_{t,l,h}^{st}).$$

Our final training loss \mathcal{L} is a weighted sum of two parts stated above and the negative log-likelihood loss \mathcal{L}_{nll} defined on bilingual sentence pair (x, y) , i.e.

$$\mathcal{L} = \mathcal{L}_{nll} + \lambda \mathcal{L}_{hid} + \mu \mathcal{L}_{align}, \quad (3)$$

where λ and μ are hyperparameters controlling the weight of different loss terms.

3 Experiments

3.1 Experimental Settings

The evaluation is on two widely used public machine translation datasets: IWSLT14 German-to-English (De-En) (Huang et al., 2017; Bahdanau et al., 2016) and WMT14 English-to-German (En-De) dataset (Wu et al., 2016; Gehring et al., 2017). To compare with previous works, we also reverse WMT14 English-to-German dataset and obtain WMT14 German-to-English dataset.

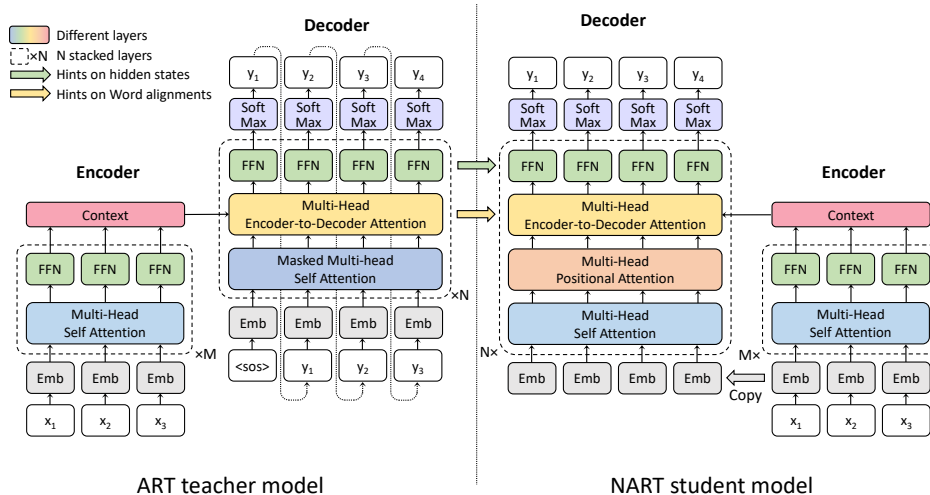


Figure 3: Hint-based training from ART model to NART model.

We pretrain Transformer (Vaswani et al., 2017) as the teacher model on each dataset, which achieves 33.26/27.30/31.29 in terms of BLEU (Papineni et al., 2002) in IWSLT14 De-En, WMT14 En-De and De-En test sets. The student model shares the same number of layers in encoder/decoder, size of hidden states/embeddings and number of heads as the teacher models (Figure 3). Following Gu et al. (2017); Kim and Rush (2016), we replace the target sentences by the decoded output of the teacher models.

Hyperparameters (γ_{st} , γ_{tr} , λ , μ) for hint-based learning are determined to make the scales of three loss components similar after initialization. We also employ label smoothing of value $\epsilon_{ls} = 0.1$ (Szegedy et al., 2016) in all experiments. We use Adam optimizer and follow the setting in Vaswani et al. (2017). Models for WMT14/IWSLT14 tasks are trained on 8/1 NVIDIA M40 GPUs respectively. The model is based on the open-sourced `tensor2tensor` (Vaswani et al., 2018).¹ More settings can be found in Appendix.

3.2 Inference

During training, T_y does not need to be predicted as the target sentence is given. During testing, we have to predict the length of the target sentence for each source sentence. In many languages, the length of the target sentence can be roughly estimated from the length of the source sentence. We choose a simple method to avoid the computational overhead, which uses input length to deter-

mine target sentence length: $T_y = T_x + C$, where C is a constant bias determined by the average length differences between the source and target training sentences. We can also predict the target length ranging from $[(T_x + C) - B, (T_x + C) + B]$, where B is the halfwidth. By doing this, we can obtain multiple translation results with different lengths. Note that we choose this method only to show the effectiveness of our proposed method and a more advanced length estimation method can be used to further improve the performance.

Once we have multiple translation results, we additionally use our ART teacher model to evaluate each result and select the one that achieves the highest probability. As the evaluation is fully parallelizable (since it is identical to the parallel training of the ART model), this rescoring operation will not hurt the non-autoregressive property of the NART model.

3.3 Experimental Results

We compare our model with several baselines, including three ART models, the fertility based (FT) NART model (Gu et al., 2017), the deterministic iterative refinement based (IR) NART model (Lee et al., 2018), and the Latent Transformer (LT; Kaiser et al., 2018) which is not fully non-autoregressive by incorporating an autoregressive sub-module in the NART model architecture.

The results are shown in the Table 1.² Across

²† and ‡ indicate that the latency is measured on our own platform or by previous works, respectively. Note that the latencies may be affected by hardware settings and such absolute values are not fair for direct comparison, so we also list the speedup of the works compared to their ART baselines.

¹Open-source code can be found at <https://github.com/zhuohan123/hint-nart>

Models	WMT14		IWSLT14	Latency	Speedup
	En-De	De-En	De-En		
<i>Autoregressive models</i>					
LSTM-based S2S (Wu et al., 2016; Bahdanau et al., 2016)	24.60	/	28.53	/	/
ConvS2S (Gehring et al., 2017; Edunov et al., 2017)	26.43	/	32.84	/	/
Transformer (Vaswani et al., 2017)	27.30	31.29	33.26	784 ms [‡]	1.00×
<i>Non-autoregressive models</i>					
FT (Gu et al., 2017)	17.69	20.62	/	39 ms [†]	15.6×
FT (rescoring 10 candidates)	18.66	22.41	/	79 ms [†]	7.68×
FT (rescoring 100 candidates)	19.17	23.20	/	257 ms [†]	2.36×
IR (Lee et al., 2018, adaptive refinement steps)	21.54	25.43	/	/	2.39×
LT (Kaiser et al., 2018)	19.8	/	/	105 ms [†]	5.78×
LT (rescoring 10 candidates)	21.0	/	/	/	/
LT (rescoring 100 candidates)	22.5	/	/	/	/
NART w/ hints	21.11	25.24	25.55	26 ms[‡]	30.2×
NART w/ hints ($B = 4, 9$ candidates)	25.20	29.52	28.80	44 ms[‡]	17.8×

Table 1: Performance on WMT14 En-De, De-En and IWSLT14 De-En tasks. “/” means non-reportable.

different datasets, our method achieves significant improvements over previous non-autoregressive models. Specifically, our method outperforms fertility based NART model with 6.54/7.11 BLEU score improvements on WMT En-De and De-En tasks in similar settings and achieves comparable results with state-of-the-art LSTM-based model on WMT En-De task. Furthermore, our model achieves a speedup of 30.2 (output a single sentence) or 17.8 (teacher rescoring) times over the ART counterparts. Note that our speedups significantly outperform all previous works, because of our lighter design of the NART model: without any computationally expensive module trying to improve the expressiveness.

We also visualize the hidden state cosine similarities and attention distributions for the NART model with hint-based training, as shown in Figure 1(c) and 2(c). With hints from hidden states, the hidden states similarities of the NART model decrease in general, and especially for the positions where the original NART model outputs incoherent phrases. The attention distribution of the NART model after hint-based training is more similar to the ART teacher model and less ambiguous comparing to the NART model without hints.

According to our empirical analysis, the percentage of repetitive words drops from 8.3% to 6.5% by our proposed methods on the IWSLT14 De-En test set, which is a 20%+ reduction. This shows that our proposed method effectively improve the quality of the translation outputs. We also provide several case studies in Appendix.

Finally, we conduct an ablation study on IWSLT14 De-En task. As shown in Table 2, the

Model	\mathcal{L}_{nll}	$\mathcal{L}_{nll} + \mathcal{L}_{align}$	$\mathcal{L}_{nll} + \mathcal{L}_{align} + \mathcal{L}_{hid}$
BLEU	23.08	24.76	25.55
Long-sentence BLEU	17.48	19.24	20.63

Table 2: Ablation studies on IWSLT14 De-En. Results are BLEU scores without teacher rescoring.

hints from word alignments provide an improvement of about 1.6 BLEU points, and the hints from hidden states improve the results by about 0.8 BLEU points. We also test these models on a subsampled set whose source sentence lengths are at least 40. Our model outperforms the baseline model by more than 3 BLEU points (20.63 v.s. 17.48).

4 Conclusion

In this paper, we proposed to use hints from a well-trained ART model to enhance the training of NART models. Our results on WMT14 En-De and De-En significantly outperform previous NART baselines, with one order of magnitude faster in inference than ART models. In the future, we will focus on designing new architectures and training methods for NART models to achieve comparable accuracy as ART models.

Acknowledgment

This work is supported by National Key R&D Program of China (2018YFB1402600), NSFC (61573026) and BJNSF (L172037) and a grant from Microsoft Research Asia. We would like to thank the anonymous reviewers for their valuable comments on our paper.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2016. An actor-critic algorithm for sequence prediction. *arXiv preprint arXiv:1607.07086*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Sergey Edunov, Myle Ott, Michael Auli, David Grangier, and Marc'Aurelio Ranzato. 2017. Classical structured prediction losses for sequence to sequence learning. *arXiv preprint arXiv:1711.04956*.
- Jonas Gehring, Michael Auli, David Grangier, and Yann N Dauphin. 2016. A convolutional encoder model for neural machine translation. *arXiv preprint arXiv:1611.02344*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2017. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Minghao Hu, Yuxing Peng, Furu Wei, Zhen Huang, Dongsheng Li, Nan Yang, and Ming Zhou. 2018. Attention-guided answer distillation for machine reading comprehension. *arXiv preprint arXiv:1808.07644*.
- Po-Sen Huang, Chong Wang, Dengyong Zhou, and Li Deng. 2017. Neural phrase-based machine translation. *arXiv preprint arXiv:1706.05565*.
- Łukasz 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 preprint arXiv:1803.03382*.
- Nal Kalchbrenner, Lasse Espeholt, Karen Simonyan, Aaron van den Oord, Alex Graves, and Koray Kavukcuoglu. 2016. Neural machine translation in linear time. *arXiv preprint arXiv:1610.10099*.
- Yoon Kim and Alexander M Rush. 2016. Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947*.
- Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. Deterministic non-autoregressive neural sequence modeling by iterative refinement. *arXiv preprint arXiv:1802.06901*.
- 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 on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826.
- Ashish Vaswani, Samy Bengio, Eugene Brevdo, Francois Chollet, Aidan N. Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, Ryan Sepassi, Noam Shazeer, and Jakob Uszkoreit. 2018. [Tensor2tensor for neural machine translation](#). *CoRR*, abs/1803.07416.
- 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 6000–6010.
- 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*.

Appendix

A Related Works

A.1 AutoRegressive Translation

Given a sentence $x = (x_1, \dots, x_{T_x})$ from the source language, the straight-forward way for translation is to generate the words in the target language $y = (y_1, \dots, y_{T_y})$ one by one from left to right. This is also known as the autoregressive factorization in which the joint probability is decomposed into a chain of conditional probabilities, as in the Eqn. (1). Deep neural networks are widely used to model such conditional probabilities based on the encoder-decoder framework. The encoder takes the source tokens (x_1, \dots, x_{T_x}) as input and encodes x into a set of context states $c = (c_1, \dots, c_{T_x})$. The decoder takes c and subsequence $y_{<t}$ as input and estimates $P(y_t|y_{<t}, c)$ according to some parametric function.

There are many design choices in the encoder-decoder framework based on different types of layers, e.g., recurrent neural network(RNN)-based (Bahdanau et al., 2014), convolution neural network(CNN)-based (Gehring et al., 2017) and recent self-attention based (Vaswani et al., 2017) approaches. We show a self-attention based network (Transformer) in the left part of Figure 3. While the ART models have achieved great success in terms of translation quality, the time consumption during inference is still far away from satisfactory. During training, the ground truth pair (x, y) is exposed to the model, and thus the prediction at different positions can be estimated in parallel based on CNN or self-attention networks. However, during inference, given a source sentence x , the decoder has to generate tokens sequentially, as the decoder inputs $y_{<t}$ must be inferred on the fly. Such autoregressive behavior becomes the bottleneck of the computational time (Wu et al., 2016).

A.2 Non-AutoRegressive Translation

In order to speed up the inference process, a line of works begin to develop non-autoregressive translation models. These models follow the encoder-decoder framework and inherit the encoder structure from the autoregressive models. After generating the context states c by the encoder, a separate module will be used to predict the target sentence length T_y and decoder inputs $z = (z_1, \dots, z_{T_y})$ by a parametric function: $(T_y, z) \sim f_z(x, c; \theta)$,

which is either deterministic or stochastic. The decoder will then predict y based on following probabilistic decomposition

$$P(y|x, T_y, z) = \prod_{t=1}^{T_y} P(y_t|z, c). \quad (4)$$

Different configurations of T_y and z enable the decoder to produce different target sentence y given the same input sentence x , which increases the output diversity of the translation models.

Previous works mainly pay attention to different design choices of f_z . Gu et al. (2017) introduce *fertilities*, corresponding to the number of target tokens occupied by each of the source tokens, and use a non-uniform copy of encoder inputs as z according to the fertility of each input token. The prediction of fertilities is done by a separated neural network-based module. Lee et al. (2018) define z by a sequence of generated target sentences $y^{(0)}, \dots, y^{(L)}$, where each $y^{(i)}$ is a refinement of $y^{(i-1)}$. Kaiser et al. (2018) use a sequence of autoregressively generated discrete latent variables as inputs of the decoder.

While the expressiveness of z improved by different kinds of design choices, the computational overhead of z will hurt the inference speed of the NART models. Comparing to the more than $15\times$ speed up in Gu et al. (2017), which uses a relatively simpler design choice of z , the speedup of Kaiser et al. (2018) is reduced to about $5\times$, and the speedup of Lee et al. (2018) is reduced to about $2\times$. This contradicts with the design goal of the NART models: to parallelize and speed up neural machine translation models.

A.3 Knowledge Distillation and Hint-Based Training

Knowledge Distillation (KD) was first proposed by Hinton et al. (2015), which trains a small *student network* from a large (possibly ensemble) *teacher network*. The training objective of the student network contains two parts. The first part is the standard classification loss, e.g, the cross entropy loss defined on the student network and the training data. The second part is defined between the output distributions of the student network and the teacher network, e.g, using KL-divergence . Kim and Rush (2016) introduces the KD framework to neural machine translation models. They replace the ground truth target sentence by the generated sentence from a well-trained teacher model. Sentence-level KD is also proved helpful

for non-autoregressive translation in multiple previous works (Gu et al., 2017; Lee et al., 2018).

However, knowledge distillation only uses the outputs of the teacher model, but ignores the rich hidden information inside a teacher model. Romero et al. (2014) introduced *hint-based training* to leverage the intermediate representations learned by the teacher model as hints to improve the training process and final performance of the student model. Hu et al. (2018) used the attention weights as hints to train a small student network for reading comprehension.

B Network Architecture

Encoder and decoder Same as the ART model, the encoder of the NART model takes the embeddings of source tokens as inputs,³ and generates a set of context vectors. As discussed in the main paper, the NART model needs to predict z given length T_y and source sentence x . We use a simple and efficient method to predict $z = (z_1, \dots, z_{T_y})$. Given source sentence $x = (x_1, \dots, x_{T_x})$ and target length T_y , we denote $e(x_i)$ as the embedding of x_i . We linearly combine the embeddings of all the source tokens to generate z as follows:

$$z_j = \sum_i w_{ij} \cdot e(x_i), \quad (5)$$

where w_{ij} is the normalized weight that controls the contribution of $e(x_i)$ to z_j according to

$$w_{ij} \propto \exp(-(j - j'(i))^2 / \tau), \quad (6)$$

where $j = 1, \dots, T_y$ and $j'(i) = (T_y/T_x) \cdot i$. τ is a hyperparameter to control the ‘‘sharpness’’ of the weight distribution. We use $f_z(x, T_y, \tau)$ for this weighted average function to be consistent as in the non-autoregressive decomposition.

Three types of multi-head attention The ART and NART models share two types of multi-head attentions: multi-head *self* attention and multi-head *encoder-to-decoder* attention. The NART model specifically uses multi-head *positional* attention to model local word orders within the

³Following (Vaswani et al., 2017; Gu et al., 2017) we also use *positional embedding* to model relative correlation between positions and add it to word embedding in both source and target sides. The positional embedding is represented by a sinusoidal function of different frequencies to encode different positions. Specifically, the positional encoding e_{pos} is computed as $e_{pos}(j, k) = \sin(j/10000^{k/d})$ (for even k) or $\cos(j/10000^{k/d})$ (for odd k), where j is the position index and k is the dimension index of the embedding vector.

sentence (Vaswani et al., 2017; Gu et al., 2017). A general attention mechanism can be formulated as querying a dictionary with key-value pairs (Vaswani et al., 2017), e.g.,

$$\begin{aligned} & \text{Attention}(Q, K, V) \\ &= \text{softmax}\left(\frac{QK^T}{\sqrt{d_{model}}}\right) \cdot V, \end{aligned} \quad (7)$$

where d_{model} is the dimension of hidden representations and Q (Query), K (Key), V (Value) differ among three types of attentions. For self attention, Q , K and V are hidden representations of the previous layer. For encoder-to-decoder attention, Q is hidden representations of the previous layer, whereas K and V are context vectors from the encoder. For positional attention, positional embeddings are used as Q and K , and hidden representations of the previous layer are used as V . The multi-head variant of attention allows the model to jointly attend to information from different representation subspaces, and is defined as

$$\begin{aligned} & \text{Multi-head}(Q, K, V) \\ &= \text{Concat}(\text{head}_1, \dots, \text{head}_H)W^O, \end{aligned} \quad (8)$$

$$\text{head}_h = \text{Attention}(QW_h^Q, KW_h^K, VW_h^V), \quad (9)$$

where $W_h^Q \in \mathbb{R}^{d_{model} \times d_k}$, $W_h^K \in \mathbb{R}^{d_{model} \times d_k}$, $W_h^V \in \mathbb{R}^{d_{model} \times d_k}$, and $W_h^O \in \mathbb{R}^{d_{model} \times Hd_v}$ are project parameter matrices, H is the number of heads, and d_k and d_v are the numbers of dimensions.

In addition to multi-head attentions, the encoder and decoder also contain fully connected feed-forward network (FFN) layers with ReLU activations, which are applied to each position separately and identically. Compositions of self attention, encoder-to-decoder attention, positional attention, and position-wise feed-forward network are stacked to form the encoder and decoder of the ART model and the NART model, with residual connections (He et al., 2016) and layer normalization (Ba et al., 2016).

C Extra Experimental Settings

Dataset specifications The split of the training/validation/test sets of the IWSLT14 dataset⁴ contain about 153K/7K/7K sentence pairs, respectively. The training set of the WMT14 dataset⁵

⁴<https://wit3.fbk.eu/>

⁵<http://www.statmt.org/wmt14/translation-task>

<i>Source:</i>	ich weiß , dass wir es können , und soweit es mich betrifft ist das etwas , was die welt jetzt braucht .
<i>Target:</i>	i know that we can , and as far as i 'm concerned , that 's something the world needs right now .
<i>ART:</i>	i know that we can , and as far as i 'm concerned , that 's something that the world needs now .
<i>NART w/o Hints:</i>	i know that we can it , , as as as it it is , it 's something that the world needs now .
<i>NART w/ Hints:</i>	i know that we can do it and as as 's m concerned , that 's something that the world needs now .
<i>Source:</i>	jeden morgen fliegen sie 240 kilometer zur farm .
<i>Target:</i>	every morning , they fly 240 miles into the farm .
<i>ART:</i>	every morning , they fly 240 miles to the farm .
<i>NART w/o Hints:</i>	every morning , you fly 240 miles to every morning .
<i>NART w/ Hints:</i>	every morning , they fly 240 miles to the farm .
<i>Source:</i>	aber bei youtube werden mehr als 48 stunden video pro minute hochgeladen .
<i>Target:</i>	but there are over 48 hours of video uploaded to youtube every minute .
<i>ART:</i>	but on youtube , more than 48 hours of video are uploaded per minute .
<i>NART w/o Hints:</i>	but on youtube , uploaded than 48 hours hours of video per minute .
<i>NART w/ Hints:</i>	but on youtube , more than 48 hours video are uploaded per minute .

Table 3: Cases on IWSLT14 De-En.

contains 4.5M parallel sentence pairs. Newstest2014 is used as the test set, and Newstest2013 is used as the validation set. In both datasets, tokens are split into a 32000 word-piece dictionary (Wu et al., 2016) which is shared in source and target languages.

Model specifications For the WMT14 dataset, we use the default network architecture of the `base` Transformer model in Vaswani et al. (2017), which consists of a 6-layer encoder and 6-layer decoder. The size of hidden nodes and embeddings are set to 512. For the IWSLT14 dataset, we use a smaller architecture, which consists of a 5-layer encoder, and a 5-layer decoder. The size of hidden states and embeddings are set to 256 and the number of heads is set to 4.

Hyperparameter specifications Hyperparameters ($\tau, \gamma_{st}, \gamma_{tr}, \lambda, \mu$) are determined to make the scales of three loss components similar after initialization. Specifically, we use $\tau = 0.3, \gamma_{st} = 0.1, \gamma_{tr} = 0.9, \lambda = 5.0, \mu = 1.0$ for IWSLT14 De-En, $\tau = 0.3, \gamma_{st} = 0.5, \gamma_{tr} = 0.9, \lambda = 5.0, \mu = 1.0$ for WMT14 De-En and WMT14 En-De.

BLEU scores We use the BLEU score (Papineni et al., 2002) as our evaluation measure. During inference, we set C to 2, -2, 2 for WMT14 En-De, De-En and IWSLT14 De-En datasets respectively, according to the average lengths of different languages in the training sets. When using the teacher to rescore, we set $B = 4$ and thus have 9 candidates in total. We also evaluate the average per-sentence decoding latencies on one NVIDIA

TITAN Xp GPU card by decoding on WMT14 En-De test sets with batch size 1 for our ART teacher model and NART models, and calculate the speedup based on them.

D Case Study

We provide some case studies for the NART models with and without hints in Table 3. From the first case, we can see that the model without hints translates the meaning of “*as far as I’m concerned*” to a set of meaningless tokens. In the second case, the model without hints omits the phrase “*the farm*” and replaces it with a repetitive phrase “*every morning*”. In the third case, the model without hints mistakenly puts the word “*uploaded*” to the beginning of the sentence, whereas our model correctly translates the source sentence. In all cases, hint-based training helps the NART model to generate better target sentences.