

On the Learning of Non-Autoregressive Transformers

Fei Huang^{*1†} Tianhua Tao^{*1} Hao Zhou^{2†} Lei Li³ Minlie Huang¹

Abstract

Non-autoregressive Transformer (NAT) is a family of text generation models, which aims to reduce the decoding latency by predicting the whole sentences in parallel. However, such latency reduction sacrifices the ability to capture left-to-right dependencies, thereby making NAT learning very challenging. In this paper, we present theoretical and empirical analyses to reveal the challenges of NAT learning and propose a unified perspective to understand existing successes. First, we show that simply training NAT by maximizing the likelihood can lead to an approximation of marginal distributions but drops all dependencies between tokens, where the dropped information can be measured by the dataset’s *conditional total correlation*. Second, we formalize many previous objectives in a unified framework and show that their success can be concluded as maximizing the likelihood on a proxy distribution, leading to a reduced information loss. Empirical studies show that our perspective can explain the phenomena in NAT learning and guide the design of new training methods.

1. Introduction

Non-Autoregressive Transformers (NATs, Gu et al., 2018; Gu et al., 2019; Ma et al., 2019; Ding et al., 2021a; Gu & Kong, 2021) have received growing attention due to their significantly lower decoding latency and approaching accuracy compared to the autoregressive Transformers (ATs)

^{*}Equal contribution [†]This work is done while Fei Huang was a research intern and Hao Zhou was a research scientist at ByteDance AI Lab. ¹The CoAI group, Tsinghua University. Institute for Artificial Intelligence, State Key Lab of Intelligent Technology and Systems, Beijing National Research Center for Information Science and Technology, DCST, Tsinghua University. ²Institute for AI Industry Research, Tsinghua University. ³University of California Santa Barbara. Correspondence to: Hao Zhou <haozhou0806@gmail.com>, Minlie Huang <ai-huang@tsinghua.edu.cn>.




Figure 1: The log-likelihood and the BLEU score on a 2D section in NAT’s weight space. GLAT+KD (★) and MLE (▲) are trained via different objectives starting from an initial checkpoint (◆) for 10k steps, where MLE achieve higher log-likelihood but lower BLEU, and GLAT is the opposite. Each point of the contour map is linearly interpolated from the three checkpoints. All values are evaluated on one translation benchmark, the validation set of WMT14 En-De.

in text generation (Qian et al., 2021b; Huang et al., 2022b). NATs generate the whole sequence in parallel based on the assumption that each token can be predicted independently. However, unlike ATs that can be easily trained via Maximum Likelihood Estimation (MLE), NAT learning is very challenging because it drops the left-to-right dependencies. Gu et al. (2018) show that directly training NATs via MLE leads to implausible outputs with repeated tokens, revealing their inability to preserve the consistency in generated texts.

To address the problem, many training methods have been proposed. For example, knowledge distillation (KD, Kim & Rush, 2016; Gu et al., 2018) supervises NATs with target sentences distilled from an AT teacher model. GLAT (Qian et al., 2021a) improves the training by utilizing a masked language model objective. These methods only change the training objectives without modifying the model, but they demonstrate significant improvements in generation quality.

Despite the empirical successes in NAT learning, there still exists a surprising characteristic not well studied: the objectives leading to a good generation quality actually result in a very low likelihood.¹ As shown in Fig.1, we finetune two NATs with different objectives from an initial checkpoint and track the changes in the log-likelihood and the BLEU score. The optimal training directions under the two metrics

¹In this paper, we mainly discuss the NATs that use the same architecture and without iterative refinement unless otherwise specified. The likelihood is obtained on the validation set.

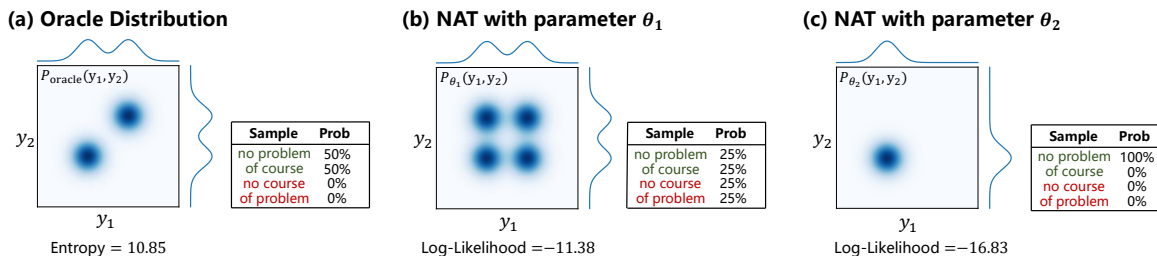


Figure 2: An example explaining why directly training NATs towards higher likelihood does not lead to better generation quality. We illustrate continuous distributions with two variables, i.e., $P(y_1, y_2)$, which are analogous to two-token sentences shown in the table. Given (a) a two-mode oracle distribution, (b) P_{θ_1} achieves higher likelihood than (c) P_{θ_2} but generates undesired outputs that mix the two modes. NATs have to follow the independent assumption, i.e., the joint distribution satisfies $P_{\theta}(y_1, y_2) = P_{\theta}(y_1)P_{\theta}(y_2)$. All distributions are conditioned on a same X , which is omitted.

are inconsistent, where GLAT+KD improves the generation quality despite that its perplexity is about 10 times of that of the counterpart trained via MLE.

Based on the phenomenons, we raise two questions:

- *Q1*: Why is NAT learning so challenging that the MLE training does not work well?
- *Q2*: Why are previously proposed objectives successful despite they lead to a low likelihood?

In this paper, we present theoretical and empirical analyses to answer the two questions. For *Q1*, we investigate the challenges of NAT learning from intuitive and theoretical perspectives. Specifically, we show that directly training NATs towards high likelihood prevents them from learning correct dependencies between target tokens, thereby degrading the generation performance. The lost information can be measured by a property of the data distribution, namely, the *conditional total correlation* (Conditional TC, \mathcal{C}), which also measures the difficulties of NAT learning.

For *Q2*, we revisit many previous training objectives and explain their success in a unified framework. Generally, we find that previous success on NAT can be concluded as adopting a training objective different from the vanilla MLE. Specifically, instead of maximizing the likelihood on the original dataset $\langle X, Y \rangle$, they in fact approximate a revised distribution, namely, the *proxy distribution* $\langle X + Z, T \rangle$, where Z and T are designed to enhance the inputs and simplify the targets. With carefully designed Z and T , the proxy distribution has a lower \mathcal{C} than the original distribution, thereby alleviating the information loss in the NAT learning. Based on the above analysis, we formulate previous training objectives in a unified framework, named Maximum Proxy-Likelihood Estimation (MPLE). We further derive a general objective to reveal the connections between the proxy distribution and the real distribution, which empirically correlates well with the generation performance and further guides the design of new training methods.

Our contributions are as follows:

- We present empirical and theoretical analyses showing

that NAT learning is challenging due to the information loss in dependencies, which can be measured by a dataset’s property, *conditional total correlation* \mathcal{C} .

- We revisit the existing successes in NAT learning and propose to understand previous training objectives in a unified perspective. We reveal that these objectives construct a proxy distribution with a reduced \mathcal{C} , thereby alleviating the information loss.

2. Challenges of NAT Learning

2.1. Preliminary

Maximum Likelihood Estimation (MLE) is a widely-used method in training text generation model, which finds a model with the closest distribution to the data distribution in terms of KL divergence (Akaike, 1998). Given a source sentence $X = [x_1, x_2, \dots, x_N]$ and a target sentence $Y = [y_1, y_2, \dots, y_M]$, MLE training minimizes

$$\begin{aligned} \mathcal{L}_{\text{MLE}} &= \mathcal{D}_{\text{KL}} [P_{\text{data}}(Y|X) || P_{\theta}(Y|X)] \\ &= -H_{\text{data}}(Y|X) - \mathbb{E}_{P_{\text{data}}(Y|X)} [\log P_{\theta}(Y|X)], \end{aligned} \quad (1)$$

where H_{data} is a constant representing the Shannon Entropy, and the second term is the log-likelihood. For autoregressive Transformers (ATs), the log-likelihood is defined as

$$\log P_{\theta}^{\text{AT}}(Y|X) = \sum_{i=1}^M \log P_{\theta}^{\text{AT}}(y_i | y_{<i}, X), \quad (2)$$

where y_i is predicted based on the prefix $y_{<i}$.

The vanilla NAT makes a *conditional independent assumption* where each token is independent of each other when X is given. Formally, we have

$$\log P_{\theta}^{\text{NAT}}(Y|X) = \sum_{i=1}^M \log P_{\theta}^{\text{NAT}}(y_i | X). \quad (3)$$

Such assumption makes the NAT a poor approximator of the real data distribution, thereby bringing many challenges in NAT learning. In the next sections, we present an intuitive explanation of the challenges and provide a quantitative method to evaluate the difficulties of NAT learning.

2.2. Challenges from Intuitive Perspective

A major challenge in NAT learning is that directly training NATs towards higher likelihood cannot lead to good generation performance. We show an intuitive example in Fig.2, which contains a continuous distribution with two variables as the real distribution (analogous to a sentence with two tokens) and two NATs with different parameters.

Comparing P_{θ_1} and P_{θ_2} , we find that P_{θ_1} perfectly approximates the marginal distributions $P(y_1)$ and $P(y_2)$, thereby achieving a higher likelihood. However, P_{θ_1} drops the dependency between y_1 and y_2 , leading to wrong outputs by mixing two sentences, previously known as the multimodality problem (Gu et al., 2018). In contrast, although P_{θ_2} has a low likelihood due to the poor approximations of the marginal distributions, it captures one of the real modes while preserving the correct lexical collocation, i.e., *no* is followed by *problem* but not *course*.

This example intuitively shows that directly training NAT to maximize the likelihood cannot capture correct lexical collocation due to the severe dependency dropping in target tokens. In the next section, we quantify the dropped dependencies based on information theory and further evaluate the difficulties of NAT learning.

2.3. Challenges from Theoretical Perspective

With the autoregressive decomposition of Eq.2, ATs can achieve zero KL divergence theoretically.² However, we show that NATs’ KL divergence is bounded by a non-negative constant, which corresponds to the information loss in approximating the data distribution.

Theorem 1. *For a NAT model $P_{\theta}(Y|X)$, we have $\min_{\theta} \mathcal{D}_{KL}[P_{data}(Y|X)||P_{\theta}(Y|X)] \geq \mathcal{C}$, where $\mathcal{C} = \sum_{i=1}^M H_{data}(y_i|X) - H_{data}(Y|X)$, and $H_{data}(\cdot|X)$ is the Shannon Entropy.*

$$\begin{aligned} \text{Proof. } & \mathcal{D}_{KL}[P_{data}(Y|X)||P_{\theta}(Y|X)] \\ &= -H_{data}(Y|X) - \mathbb{E}_{P_{data}(Y|X)} \left[\sum_{i=1}^M \log P_{\theta}(y_i|X) \right] \\ & \quad \text{(Conditional Independent Assumption of Eq.3)} \\ &= -H_{data}(Y|X) - \sum_{i=1}^M \mathbb{E}_{P_{data}(y_i|X)} [\log P_{\theta}(y_i|X)] \\ &\geq -H_{data}(Y|X) + \sum_{i=1}^M H_{data}(y_i|X) \quad \text{(Gibbs' Inequality)} \end{aligned}$$

The equality is achieved when $P_{\theta}(y_i|X) = P_{data}(y_i|X)$. Note that \mathcal{C} is a non-negative constant called *conditional total correlation* (Conditional TC, Watanabe, 1960) or *multi-information* (Studený & Vejnarová, 1998), which measures the information of dependencies between the target tokens

²It is achieved when $P_{\theta}(y_i|y_{<i}, X) = P_{data}(y_i|y_{<i}, X)$.

Table 1: Estimated \mathcal{C} and the gap of BLEU between AT and NAT on various datasets.³ A large \mathcal{C} leads to significant performance gap between AT and NAT. All models are trained via MLE. $\Delta\text{BLEU} = \text{BLEU}_{\text{AT}} - \text{BLEU}_{\text{NAT}}$. The dataset details are presented in Appendix B.

Dataset	\mathcal{C}	ΔBLEU	BLEU_{AT}	BLEU_{NAT}
WMT14 En-De	2.50	15.32	27.11	11.79
WMT16 En-Ro	2.20	9.98	33.70	23.72
Synthetic B	1.51	5.66	20.97	15.31
Synthetic A	0.92	0.35	26.96	26.61

when X is known. We make two remarks on Theorem 1:

Remark 1. *A well-trained NAT (in terms of KL divergence) achieves perfect approximations on marginal distributions but drops all the dependency information between target tokens, which can be measured by \mathcal{C} .*

Remark 2. *\mathcal{C} is a property of data distribution representing the difficulties in NAT learning. Given the data distribution, an NAT cannot achieve an information loss less than \mathcal{C} regardless of its parameters or training methods.*

Conditional TC and Performance Gap To better understand how \mathcal{C} affects the NAT performance, we estimate \mathcal{C} and compare the generation performance of AT and NAT models trained via MLE on four datasets. Since \mathcal{C} is usually high for most datasets, besides two translation benchmarks, we further construct two synthetic datasets that have a lower \mathcal{C} . Please refer to Appendix B for more details.

As shown in Table 1, large \mathcal{C} indicates strong dependencies between target tokens, leading to a serious performance gap between NAT and AT models. When \mathcal{C} is small, NAT can achieve a similar performance with AT, verifying that the large Conditional TC is the main obstacle in NAT learning.

3. Understanding NAT Learning via Maximum Proxy-Likelihood Estimation

Sec.2 shows that MLE-trained NAT drops the dependencies between tokens, where \mathcal{C} measures the difficulties in NAT learning. In this section, we investigate previous successes in training NATs and propose a unified perspective to understand them.

Specifically, we revisit existing training objectives and find that many of them improve the MLE training by simplifying the target sentences or enhancing the training inputs. Such modifications significantly change the training directions, where they actually maximize the likelihood on a *proxy*

³We report \mathcal{C} with the *V-entropy* (Xu et al., 2020) instead of the Shannon entropy because the latter is intractable due to the unknown data distribution. Note that \mathcal{C} measures the amount of information, whose unit is *bit*, which is comparable across datasets. However, ΔBLEU is not strictly comparable, where we present a more rigorous comparison in Appendix B.

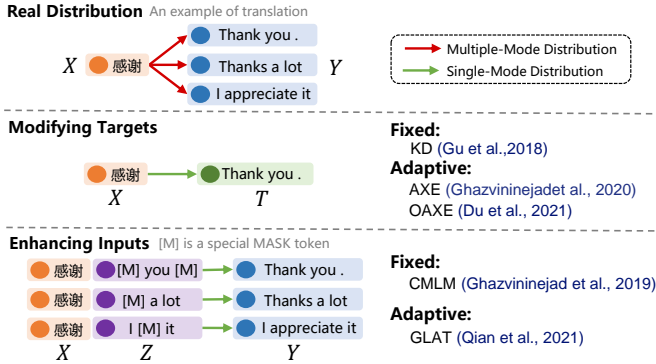


Figure 3: An overview of methods to construct the proxy distribution Q , which fall into two categories: *Modifying Targets* (replacing Y by a proxy target T) and *Enhancing Inputs* (with a proxy input Z). *Fixed* and *Adaptive* indicate whether the proxy distribution is adjusted through the training.

distribution instead of the original distribution. The proxy distribution with modified targets or inputs usually has a low \mathcal{C} , thereby reducing the information loss in NAT learning. Based on the above analysis, we formulate these methods in a unified framework, named Maximum Proxy-Likelihood Estimation (MPLE). Intuitively, MPLE’s objective can be expressed as

$$\mathcal{L} = \mathcal{D}_{\text{KL}}(Q||P_{\theta}) + \mathcal{R}(Q, P_{\text{data}}). \quad (4)$$

The first term is similar to the MLE objective, which trains the model towards the proxy distribution Q instead of P_{data} . The second term is a regularizer controlling the distortion between Q and P_{data} .

3.1. Revisiting Previous Successes

Considering the severe challenges in NAT learning, many training methods are proposed to improve the generation performance. For example, *Aligned Cross Entropy* (AXE, Ghazvininejad et al., 2020) finds that the cross-entropy loss highly penalizes small shifts in word order, which deviates from the evaluation of generation quality and thus hinders the NAT training. They propose an aligned-based objective that allows small target shifts to alleviate the problem. *GLAT* (Qian et al., 2021a) proposes to promote representation learning by utilizing curriculum learning. Specifically, they train NATs similar to the masked language model, which feeds a masked target as the decoder input and adjusts the training difficulties by annealing the masking ratio.

Although these methods are proposed with different motivations, we find that they generally share a similar objective that can be interpreted as the MLE training. Specifically, they still use the cross-entropy loss between the NAT predictions and the target tokens, except that the target labels or model inputs are changed.⁴ Then we can interpret the loss

⁴They also do not change the inference process, e.g., no extra inputs are introduced for decoder in generation.

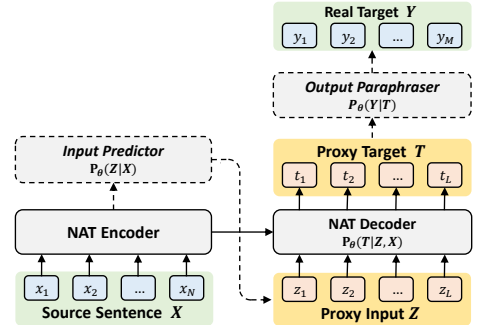


Figure 4: The latent variable model used in the derivation. The source sentence X and the real target Y are observable, whereas the proxy input Z and the proxy target T are latent.

as an objective of MLE, but the target distribution is actually replaced by a new distribution with their new inputs and outputs, where we call it a *proxy distribution* Q . In these methods, NATs are trained on the proxy distribution Q to maximize the likelihood, which explains why they have a low likelihood on the original validation set.

By examining these methods closely, we find that the proxy distribution Q is an essential key to their success. Specifically, we divide existing methods into two categories: *Modifying Targets* or *Enhancing Inputs*. As shown in Fig.3, both approaches try to preserve a one-to-one mapping between the new inputs and outputs, which intuitively reduces \mathcal{C} by limiting the possible modes in the proxy distribution,⁵ thereby alleviating the information loss in NAT learning.

Formally, we denote the proxy distribution by $Q(T|Z, X)$, where the original Y is replaced by a *proxy target* T , and the original X is enhanced with a *proxy input* Z . Next, we revisit existing methods of NAT learning to study how they construct the proxy distribution Q .

Constructing Q by Modifying Targets ($Y \rightarrow T$) Sequence-level knowledge distillation (KD, Gu et al., 2018) is a direct method to simplify the targets. For a given input X , an autoregressive teacher generates the proxy target T by beam search, which replaces the diverse references and thus reduces the possible outputs in the data distribution. The KD data are usually generated in advance and does not change during NAT training.

Some methods construct T through the training, which are adaptively adjusted according to the NAT model. AXE (Ghazvininejad et al., 2020) and OaXE (Du et al., 2021) use alignment-based objectives, which match each prediction with a reference token and calculate the cross-entropy loss. The two losses are equivalent to obtaining the

⁵A distribution with multiple modes requires dependency information to recover the joint distribution, as shown in Fig.2.

MLE objective with a new target T , where T is a permutation of Y but closer to the model prediction.

Constructing Q by Enhancing Inputs ($X \rightarrow Z, X$)
 CMLM⁶ (Ghazvininejad et al., 2019) uses a masked language model objective, where a randomly masked target sentence is fed into the NAT decoder. Intuitively, if $P(Y|X)$ has multiple possible outputs, $Q(Y|Z, X)$ can reduce the number of candidates with the constraint of Z , which again leads to a simplified distribution with reduced \mathcal{C} .

Unlike CMLM that samples Z from a predefined distribution by random masking, GLAT (Qian et al., 2021a) proposes to sample Z adaptively according to the NAT performance. Specifically, if the NAT well approximates $P(Y|X)$ without Z , GLAT will mask most tokens in the proxy input. Since the NAT uses full masks in inference, GLAT improves CMLM by reducing the training and inference gap.

3.2. A Unified Objective of MPLE

Existing methods simply train NAT by maximizing the likelihood on the proxy distribution $Q(T|Z, X)$. However, they do not answer *when the performance on the proxy distribution can generalize to the real distribution*. For example, a good approximation of Q may not guarantee good generation performance on P_{data} since there can be a substantial distortion between the two distributions.

In MPLE, we propose considering \mathcal{C} and the data distortion together in a unified objective. Specifically, we regard T and Z as latent variables and build a latent variable model that connects Z, T and X, Y , as shown in Fig.4. Formally,

$$P_{\theta}(Y|X) = \sum_Z \sum_T P_{\theta}(Y|T)P_{\theta}(T|Z, X)P_{\theta}(Z|X), \quad (5)$$

where $P_{\theta}(T|Z, X)$ is the NAT decoder, and the other two modules bridge Z with X , and T with Y , respectively. Then we derive our objective from the likelihood on $P_{\text{data}}(Y|X)$:

$$\begin{aligned} & -\mathbb{E}_{P_{\text{data}}(Y|X)} \log P_{\theta}(Y|X) \\ &= -\mathbb{E}_{P_{\text{data}}(Y|X)} \log \left[\mathbb{E}_{Q(T, Z|X)} \frac{P_{\theta}(Y, T, Z|X)}{Q(T, Z|X)} \right] \\ &\leq -\mathbb{E}_{P_{\text{data}}(Y|X)} \mathbb{E}_{Q(T, Z|X)} \left[\log \frac{P_{\theta}(Y, T, Z|X)}{Q(T, Z|X)} \right] \quad (6) \\ &= -\mathbb{E}_{P_{\text{data}}(Y|X)} \mathbb{E}_{Q(T, Z|X)} \left[\log P_{\theta}(Y|T) + \right. \\ &\quad \left. \log \frac{P_{\theta}(T|Z, X)}{Q(T|Z, X)} + \log \frac{P_{\theta}(Z|X)}{Q(Z|X)} \right] \quad (7) \end{aligned}$$

In Eq.6, we apply *variational principle* (Fox & Roberts, 2012) by introducing $Q(T, Z|X)$, which specifies how we obtain Z, T and can be decomposed into the proxy distribution $Q(T|Z, X)$ and $Q(Z|X)$.

⁶We discuss the non-iterative version of CMLM here, following Ghazvininejad et al. (2020); Du et al. (2021).

Eq.7 is our unified objective $\mathcal{L}_{\text{MPLE}}$, which can be simplified and recovers our intuition in Eq.4:

$$\mathcal{L}_{\text{MPLE}} = \underbrace{\mathcal{L}_{\text{NAT}}}_{\mathcal{D}_{\text{KL}}(Q||P_{\theta})} + \underbrace{\mathcal{L}_{\text{target}} + \mathcal{L}_{\text{input}}}_{\mathcal{R}(Q, P_{\text{data}})}, \quad (8)$$

$$\mathcal{L}_{\text{NAT}} = \mathbb{E}_{Q(Z|X)} \mathcal{D}_{\text{KL}} [Q(T|Z, X)||P_{\theta}(T|Z, X)], \quad (9)$$

$$\mathcal{L}_{\text{target}} = \mathbb{E}_{P_{\text{data}}(Y|X)} \mathbb{E}_{Q(T|X)} [-\log P_{\theta}(Y|T)], \quad (10)$$

$$\mathcal{L}_{\text{input}} = \mathcal{D}_{\text{KL}} [Q(Z|X)||P_{\theta}(Z|X)]. \quad (11)$$

In Eq.8, \mathcal{L}_{NAT} supervises the decoder $P_{\theta}(T|Z, X)$ to maximize the likelihood on the proxy distribution. $\mathcal{L}_{\text{target}}$ and $\mathcal{L}_{\text{input}}$ measure the cost in bridging T, Z with X, Y , and act as regularizers to avoid large distortions between the proxy and original variables.

Moreover, since $P_{\theta}(T|Z, X)$ still follows the independent assumption, we can derive a lower bound of \mathcal{L}_{NAT} in a similar way of Theorem 1. Specifically, we have

$$\mathcal{C}' := \mathbb{E}_{Q(Z|X)} \left[\sum_{i=1}^L H_Q(t_i|Z, X) - H_Q(T|Z, X) \right], \quad (12)$$

that satisfies $\mathcal{L}_{\text{NAT}} \geq \mathcal{C}'$, where \mathcal{C}' is the Conditional TC of the proxy distribution $Q(T|Z, X)$.

3.3. Understanding Existing Methods in MPLE

$\mathcal{L}_{\text{MPLE}}$ seems a bit complex because it includes both the likelihood term to train the NAT model and the objective for selecting the proxy inputs and targets. To understand previous methods in MPLE, we describe the training process as an *Expectation Maximization* algorithm including two steps: (1) find optimal proxy distribution Q by adjusting the proxy variables Z and T ; (2) optimize model parameter θ .

In **E-step**, we fix the model parameter θ and update proxy variables to reduce $\mathcal{L}_{\text{MPLE}}$, which aims to find good proxy distribution Q to balance \mathcal{L}_{NAT} and the data distortion. Since θ is fixed, adjusting proxy variables for lower \mathcal{L}_{NAT} does not affect the NAT model but actually optimizes \mathcal{C}' , where \mathcal{L}_{NAT} is the upper bound of \mathcal{C}' defined in Eq.12.

However, such optimization is non-trivial, where existing methods utilize some heuristic rules. For example, KD obtains T by distilling sentences from a pre-trained AT teacher, which efficiently alleviates the information loss by reducing the modes in the dataset. AXE and OaXE obtain T by aligning the NAT prediction with Y , where they have hyper-parameters for controlling the distortion within an acceptable range. Please refer to Appendix E for more details about the heuristic rules in existing methods.

As introduced in Sec.3.1, these heuristic rules utilize either fixed or adaptive strategies. Fixed strategies obtain the proxy distribution before the training, where \mathcal{C}' is lower than the original \mathcal{C} , but not further optimized. In contrast, adaptive strategies adjust the proxy distribution through the training, which usually outperforms the fixed ones.

In **M-step**, we fix the proxy distribution Q and optimize the model θ . Since all Q 's entropies are constants and thus ignored, the three losses in Eq.8 can be easily calculated based on Z and T previously obtained in the E-step. Specially, \mathcal{L}_{NAT} recovers the objectives of existing methods by maximizing the likelihood on the proxy distribution.

Quantifying Data Distortion Existing methods heuristically obtain Z and T to balance the training difficulties and the data distortion, which does not involve a measurement of the distortion. MPLE provides a method to quantify the distortion, allowing for comparisons between different methods in constructing the proxy distribution.

Specifically, we use the output paraphraser $P_\theta(Y|T)$ and the input predictor $P_\theta(Z|X)$ to measure the data distortion $\mathcal{L}_{\text{target}}$ and $\mathcal{L}_{\text{input}}$, respectively. For $\mathcal{L}_{\text{target}}$, we define the output paraphraser as a simple non-trainable distribution related to the similarity between Y and T :

$$P_\theta(Y|T) = \exp(\beta S(Y, T)) / \zeta, \quad (13)$$

where β is a hyper-parameter, $S(Y, T)$ is the sentence BLEU, and $\zeta = \sum_Y \exp(\beta S(Y, T))$. However, the normalization term ζ is intractable, so we drop it and empirically use $\hat{\mathcal{L}}_{\text{target}}$ instead:

$$\hat{\mathcal{L}}_{\text{target}} = \mathbb{E}_{P_{\text{data}}(Y|X)} \mathbb{E}_{Q(T|X)} [-\beta S(Y, T)]. \quad (14)$$

Intuitively, Eq.14 measures the distortion between proxy and real targets by the average BLEU score.

For $\mathcal{L}_{\text{input}}$, we design a trainable input predictor specially for GLAT and CMLM, where the other methods without an extra input Z always have $\mathcal{L}_{\text{input}} = 0$. Specifically, we define the input predictor as a classifier, which predicts z_i from the vocabulary including a special mask token. We predict Z non-autoregressively (See Appendix E.6 for details):

$$\log P_\theta(Z|X) = \sum_{i=1}^M \log P_\theta(z_i|X). \quad (15)$$

Then, $\mathcal{L}_{\text{input}}$ can be calculated according to Eq.11.

Discussing More Work from MPLE Perspective Besides the methods discussed above, MPLE can also explain many other objectives proposed for NAT learning, including (1) the methods introducing continuous or discrete latent variables (Kaiser et al., 2018; Shu et al., 2020; Bao et al., 2021; 2022); (2) enhancing NAT decoder with order information (Bao et al., 2019; Ran et al., 2021), POS taggings (Yang et al., 2021), or tokens sampled from target sentences (Huang et al., 2022a); (3) KD variants like reverse distillation (Ding et al., 2021b) or repeated distillation (Zhou et al., 2020; Sun & Yang, 2020).

Notably, CTC-based methods (Libovický & Helcl, 2018; Saharia et al., 2020) and DA-Transformer (Huang et al.,

2022b) have been shown very effective in NAT learning, where they also utilize alignment-based objectives. Unlike AXE, these methods predict a sequence longer than the real target, and then remove useless tokens by rules or model predictions. Their success show that the proxy target T does not necessarily have similar length with Y , where a longer T can be more flexible and efficient in reducing the token dependencies. Moreover, they introduce a different $P_\theta(Y|T)$ from Eq.13, which predicts Y from a longer T with reconstruction of dependency information, e.g., by transitions predicted in DA-Transformer.⁷

Finally, MPLE also connects with iterative NATs (Lee et al., 2018; Ghazvininejad et al., 2019; Kasai et al., 2020; Guo et al., 2020). Although iterative NATs do not satisfy the independent assumption in Eq.3, they still predict tokens independently in each iterative step. Specifically, we point out that (1) \mathcal{C} measures the information loss of iterative NAT in each refinement step; (2) some iterative NATs are special cases of MPLE with parameter sharing in Input Predictor and NAT decoder. Please refer to Appendix A for details.

3.4. A KD Variant from MPLE

Existing methods heuristically obtain Z and T to construct the proxy distribution. We propose a new variant of KD that improves the proxy distribution by explicitly balancing \mathcal{L}_{NAT} and the data distortion, named dynamic KD.

For a source sentence X , we obtain a target candidate set Γ , which contains the raw data and distilled data from AT teachers of different sizes, i.e., Transformer-*tiny/small/base/big*. Then we choose a best target $T \in \Gamma$ that minimizes $\mathcal{L}_{\text{NAT}} + \hat{\mathcal{L}}_{\text{target}}$. Noticing that Eq.14 is intractable due to the sampling from P_{data} , we use the pairwise BLEU between the candidates instead. More details are presented in Appendix F.

Previous work (Zhou et al., 2020) finds that the KD data from a larger AT teacher is closer to the real data but more difficult to predict, where they suggest choosing the teacher size according to NAT's capacity. Our method dynamically selects the best proxy target from multiple KD candidates for each sample, which achieves substantial improvement over NATs trained on any single KD data.

4. Experiments

Dataset We use two translation benchmarks, WMT14 En-De (4.5M) and WMT17 Zh-En (20M), and follow Zhou et al. (2020); Kasai et al. (2020) for preprocessing.

⁷Both methods do not directly fit in Eq.8 because they maximize the logarithm of probability sum on all alignments instead of a single T . However, we refer the reader to Sec.3.2 of Huang et al. (2022b), which shows that their objectives can be regarded as utilizing multiple proxy targets with different weights.

Knowledge Distillation We use Transformer-*base* with the same settings as Vaswani et al. (2017) and generate the distilled data with beam size 5. All models are based on KD unless otherwise specified.

Implementation Details We implement Raw Data, KD, AXE, OaXE for obtaining proxy targets, and Vanilla (i.e., no extra input), CMLM, GLAT for obtaining proxy inputs. We generally follow the hyper-parameters in Qian et al. (2021a). For fair comparisons, we only modify the heuristic rules to obtain Z and T , which may be different from their original implementations. For example, we do not use iterative refinement for CMLM, or combine OaXE with CMLM. Unless otherwise specified, we do not utilize reranking methods or other decoding tricks. More details are in Appendix G.

Metrics We utilize tokenized BLEU (Papineni et al., 2002) to evaluate the translation performance. $\mathcal{L}_{\text{input}}$ and \mathcal{L}_{NAT} are averaged per token on validation set. $\hat{\mathcal{L}}_{\text{target}}$ in Eq.14 requires multiple real targets Y from P_{data} , so we utilize multi-reference annotations from Ott et al. (2018); Hassan et al. (2018), where each sample has 10(2) extra human-annotated references for En-De(Zh-En). To calculate $\hat{\mathcal{L}}_{\text{target}}$ by Eq.14, we use $\beta = 0.2$ for En-De, $\beta = 0.25$ for Zh-En. $\hat{\mathcal{L}}_{\text{MPLE}} := \mathcal{L}_{\text{NAT}} + \mathcal{L}_{\text{input}} + \hat{\mathcal{L}}_{\text{target}}$. We measure the speedup of decoding latency on WMT14 En-De with batch size 1.

4.1. Verification of Reduced Information Loss

Theorem 1 implies that any NAT approximating the real distribution cannot achieve less information loss than dataset’s \mathcal{C} . We argue that existing methods approximate a proxy distribution Q instead, thereby achieving reduced information loss. To verify the proposition, we compare \mathcal{L}_{NAT} of different methods against the the original dataset’s \mathcal{C} .

As shown in Fig.5, most methods except Raw Data achieve lower \mathcal{L}_{NAT} than the dataset’s \mathcal{C} . Note that \mathcal{L}_{NAT} evaluates the information loss in approximating the proxy distribution, which is the upper bound of Q ’s Conditional TC, i.e., \mathcal{C}' defined Eq.12. The results empirically verify that (1) training on the proxy distribution alleviates the information loss in NAT learning; (2) the proxy distribution has a reduced Conditional TC.




Figure 5: \mathcal{L}_{NAT} of different methods on WMT14 En-De. All methods except Raw Data achieve lower \mathcal{L}_{NAT} than the dataset’s \mathcal{C} , verifying that training on the proxy distribution alleviates the information loss.

Table 2: Comparison of proxy targets on WMT14 En-De. All methods use Vanilla for Z and $\mathcal{L}_{\text{input}} = 0$. $\hat{\mathcal{L}}_{\text{MPLE}}$ and BLEU are strongly correlated (Pearson’s $|r|=0.99$). AXE’s τ and OaXE’s numbers indicate the skip penalty and pre-training steps, which are hyper-parameters in choosing T .

Models	\mathcal{L}_{NAT}	$\hat{\mathcal{L}}_{\text{target}}$	$\hat{\mathcal{L}}_{\text{MPLE}}$	BLEU
Raw Data	4.41	-6.42	-2.01	11.79
KD	2.42	-7.08	-4.66	20.87
+ AXE ($\tau=1$)	0.78	-5.13	-4.35	18.56
+ AXE ($\tau=5$)	1.09	-6.34	-5.25	22.22
+ AXE ($\tau=10$)	1.25	-6.50	-5.26	22.35
+ OaXE (10k)	1.03	-4.41	-3.38	15.00
+ OaXE (50k)	0.79	-5.84	-5.06	21.37
+ OaXE (300k)	0.83	-6.28	-5.44	22.76

However, lower \mathcal{L}_{NAT} does not promise higher BLEU because they do not control the data distortion. In the next sections, we will analyze how different methods affect the performance by balancing \mathcal{L}_{NAT} and the data distortion.

4.2. Effects of Proxy Target

In this section, we compare different methods of obtaining proxy targets with varying hyper-parameters. We present the results on En-De in Table 2 and Zh-En in Appendix C.

Strong Correlation. $\hat{\mathcal{L}}_{\text{MPLE}}$ is strongly correlated with BLEU, where \mathcal{L}_{NAT} and $\hat{\mathcal{L}}_{\text{target}}$ are both important. For example, AXE($\tau=1$) achieves low \mathcal{L}_{NAT} with high $\hat{\mathcal{L}}_{\text{target}}$, indicating that T is easy to predict but heavily distorted from the real target. On the contrary, KD’s proxy target is less distorted but hard to predict. OaXE(300k) balances the two losses well and thus achieves the best BLEU.

β in Eq.14 will affect the scale of $\hat{\mathcal{L}}_{\text{target}}$, where we choose $\beta = 0.2$ to maximize the correlation. However, the choice of β is not sensitive that $|r| \geq 0.8$ for all $\beta \in [0.1, 0.5]$.

Secret Advantage of KD. Previous work (Gu et al., 2018) has shown that KD can simplify the training data and thus reduce \mathcal{L}_{NAT} . However, our results show a secret advantage that KD also achieves the lowest $\hat{\mathcal{L}}_{\text{target}}$, indicating that the KD data are even closer to the multiple human references on average than the raw data. This result is caused by the diversity of human annotations, as shown in Fig.6.




Figure 6: KD data are even closer to the multiple references on average than Raw Data on WMT14 En-De, thereby achieving lower $\hat{\mathcal{L}}_{\text{target}}$ and improving the performance.

Table 3: Comparison of proxy inputs on WMT14 En-De. All methods use KD for T and $\hat{\mathcal{L}}_{\text{target}} = -7.08$. Sample and Default indicate the BLEU score with Input Sampling and Default Decoding. $\hat{\mathcal{L}}_{\text{MPLE}}$ is strongly correlated with Sample BLEU ($|r|=0.99$) but less correlated with Default BLEU ($|r|=0.37$). The variants of CMLM and GLAT use different strategies for masking, detailed in Appendix E.6.

Models	$\mathcal{L}_{\text{input}}$	\mathcal{L}_{NAT}	$\hat{\mathcal{L}}_{\text{MPLE}}$	Sample	Default
Vanilla	0	2.42	-4.66	20.87	20.87
CMLM	0.99	0.46	-5.63	23.48	19.39
+ fixed masking ratio	0.48	0.55	-6.05	24.28	19.35
GLAT	0.45	0.66	-5.96	23.98	25.12
+ Levenshtein dist.	0.41	0.73	-5.94	24.03	24.84
+ mask by P_{ref}	0.25	1.24	-5.59	22.98	24.22
+ mask by $1 - P_{\text{ref}}$	0.57	0.50	-6.01	24.35	25.19

Although the KD data may not belong to any modes of the data distribution, it still has higher similarity on average.

Hyper-parameters and Trade-off. AXE and OaXE utilize tricks to avoid large distortion between the proxy target and the real target. For example, AXE tunes the skip penalty τ , and OaXE tunes the pre-training step. MPLE provides a quantifiable method to measure the trade-off between the likelihood loss \mathcal{L}_{NAT} and the distortion $\hat{\mathcal{L}}_{\text{target}}$, which improves the interpretability of hyper-parameter selection.

4.3. Effects of Proxy Input

We compare methods that obtain proxy inputs including several variants of CMLM and GLAT, which are also used in Qian et al. (2021a). These variants use different strategies for masking, whose details are presented in Appendix E.6.

In the inference of CMLM and GLAT, they use a full mask as the proxy input by default (Default Decoding), leading to a large gap between train and inference. We propose to sample the latent input (Input Sampling) based our latent variable model in Eq.5: We first sample Z according to the input predictor $P_{\theta}(Z|X)$,⁸ and then choose the most likely tokens predicted by the NAT decoder. We present the results on En-De in Table 3 and Zh-En in Appendix C.

Strong Correlation with Sample BLEU. Our objective is strongly correlated with BLEU when using Input Sampling, where $\mathcal{L}_{\text{input}}$ and \mathcal{L}_{NAT} should be balanced to achieve the best performance. For example, Vanilla NAT does not introduce extra inputs, leading to large \mathcal{L}_{NAT} ; CMLM introduces too many tokens in Z , bringing a large distortion from the original input. However, $\hat{\mathcal{L}}_{\text{MPLE}}$ is less correlated with BLEU of Default Decoding, which can be caused by the decoding strategy as discussed below.

⁸More precisely, we first decide whether a token is masked according to the predicted distribution. If it is not masked, we directly use the most likely non-mask token, which empirically leads to better performance. Please see Appendix E.6 for details.




Figure 7: Decoding Confidence I and the BLEU score with two decoding strategies. CMLM with Input Sampling is more confident and thus achieves better BLEU than Default Decoding. GLAT is the opposite.

Potentials for Decoding Strategies. Previous work (Qian et al., 2021a) showed that CMLM performs poorly with a full masked decoder input, but we find that it can be improved by utilizing the input predictor to generate a better proxy input Z in inference. Specifically, Input Sampling brings about 4 BLEU points improvement on CMLM. This idea is connected with the iterative NATs, where their refined sentence can be interpreted as a proxy input to improve the generation quality.

We also find that CMLM and GLAT prefer different decoding strategies, which can be explained by the decoding confidence $P_{\theta}(Z|X)$ and $P_{\theta}(T|Z, X)$. As shown in Fig.7, CMLM is more confident with Input Sampling than Default Decoding, whereas GLAT is the opposite.

4.4. Results of Dynamic KD

We combine Dynamic KD with GLAT and further utilize a reranking method following our baselines. As shown in Table 4, Dynamic KD brings about 0.6 ~ 0.7 BLEU improvement against the single KD distilled from Transformer-base. Moreover, our best results achieve competitive translation quality with ATs with the modest cost in reranking.

We further compare Dynamic KD against single KD data distilled from different AT teachers. As shown in Fig.8,

Table 4: Comparing Dynamic KD against AT and previous NATs. †: Reported by Qian et al. (2021a) and Du et al. (2021). LPD (Wei et al., 2019a) and NPD (Gu et al., 2018) indicate reranking methods with the number of candidates. NPD are slower than LPD due to the use of an external AT reranker. Please see Table 9 for more results.

	Models	En-De	Zh-En	Speedup
AT	Transformer	27.11	23.89	1.0x
NAT	MLE	11.79	8.69	15.3x
	GLAT (NPD=7)†	26.55	/	7.9x
	OaXE (LPD=5)†	26.1	22.1	14.2x
Ours	GLAT + KD	25.12	22.51	15.3x
	+ Dynamic KD	25.88	23.07	15.3x
	+ LPD=3	26.89	24.42	14.6x




Figure 8: Comparing Dynamic KD against NATs trained on single KD data. Tiny/Small/Base/Big indicates the size of *AT teacher*. Raw represents the raw data. Dynamic KD utilizes all KD data and outperforms any single of them.

Dynamic KD outperforms the best result on any single KD data, verifying that explicitly balancing \mathcal{L}_{NAT} and the data distortion leads to better performance. Notably, applying both proxy inputs and targets (GLAT+Dynamic KD) is better than simply applying one of them (Vanilla+Dynamic KD), showing that unifying the two methods of constructing proxy distributions is effective.

The results suggest that our perspective effectively guides the design of new training methods. Explicitly optimizing $\mathcal{L}_{\text{MPLE}}$ provides a promising way to find better proxy distributions, which outperforms existing heuristic methods.

5. Related Work

NATs are proposed to reduce the decoding latency but suffer from poor generation quality. Many studies are devoted to solving the problem. Besides the methods discussed in our analysis, some studies are also helpful in improving the NAT performance, mainly including (1) objectives not based on cross-entropy (Wei et al., 2019a; Sun et al., 2019; Shao et al., 2020; 2021); (2) iteratively refining the generated outputs (Lee et al., 2018; Ghazvininejad et al., 2019; Gu et al., 2019; Kasai et al., 2020; Guo et al., 2020). Although the iterative approaches usually lead to better quality, Kasai et al. (2021) find that these models are much slower and may not have advantages against ATs. Moreover, recent works show that the non-iterative methods can also achieve competitive quality with AT models and have substantial lower latency than iterative methods (Gu & Kong, 2021; Qian et al., 2021b; Huang et al., 2022b).

Notably, a previous study (Zhou et al., 2020) also analyzes the NAT learning but mainly focuses on the KD method. They propose metrics to evaluate the complexity of the KD data and explain how KD improves NAT generation. Unlike their analysis that only considers the KD data, our perspective is more general in understanding many SoTA methods and better supported by the information theory.⁹

⁹Specifically, their proposed metric may not correctly reflect

6. Conclusion

In this paper, we investigate the challenges in NAT learning. From intuitive and theoretical perspectives, we show that the problem roots in the large information loss in capturing dependencies between tokens, where the dropped information can be measured by the dataset’s *conditional total correlation* \mathcal{C} . Furthermore, we revisit the existing successes in NAT learning and find that many previous studies alleviate the problem by maximizing the likelihood on a proxy distribution, which is designed to have a lower \mathcal{C} . Based on the analysis, we propose a unified framework named Maximum Proxy-Likelihood Estimation (MPLE), which provides a unified objective revealing how the choice of proxy distribution contributes to the final performance. This framework improves our understanding of a wide range of NAT learning methods, including the SOTA ones like alignment-based objectives and glancing training. Empirical analyses show that our perspective can well explain the phenomena in NAT learning, where the proposed objective highly correlates with the generation performance and can further guide the design of better training methods.

Acknowledgement

We would like to thank Yuxuan Song for the help in the derivation of Theorem 1 and the MPLE framework. This work was supported by the National Science Foundation for Distinguished Young Scholars (with No. 62125604) and the NSFC projects (Key project with No. 61936010 and regular project with No. 61876096). This work was also supported by the Guoqiang Institute of Tsinghua University, with Grant No. 2019GQG1 and 2020GQG0005, and sponsored by Tsinghua-Toyota Joint Research Fund.

References

- Akaike, H. Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike*, pp. 199–213. Springer, 1998.
- Bao, Y., Zhou, H., Feng, J., Wang, M., Huang, S., Chen, J., and Li, L. Non-autoregressive transformer by position learning. *CoRR*, abs/1911.10677, 2019. URL <http://arxiv.org/abs/1911.10677>.
- Bao, Y., Huang, S., Xiao, T., Wang, D., Dai, X., and Chen, J. Non-autoregressive translation by learning target categorical codes. In Toutanova, K., Rumshisky, A., Zettlemoyer, L., Hakkani-Tür, D., Beltagy, I., Bethard, S., Cotterell, R., Chakraborty, T., and Zhou, Y. (eds.), *Proceedings of*

the difficulties of NAT learning in some cases. For a dataset $y_1, \dots, y_M \stackrel{iid}{\sim} p(\cdot|X)$, their proposed metric is high if $p(\cdot|X)$ has a large entropy. However, our $\mathcal{C} = 0$, correctly showing that the data satisfy the independent assumption.

- the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021, pp. 5749–5759. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.naacl-main.458. URL <https://doi.org/10.18653/v1/2021.naacl-main.458>.
- Bao, Y., Zhou, H., Huang, S., Wang, D., Qian, L., Dai, X., Chen, J., and Li, L. latent-glat: Glancing at latent variables for parallel text generation. *CoRR*, abs/2204.02030, 2022. doi: 10.48550/arXiv.2204.02030. URL <https://doi.org/10.48550/arXiv.2204.02030>.
- Ding, L., Wang, L., Liu, X., Wong, D. F., Tao, D., and Tu, Z. Understanding and improving lexical choice in non-autoregressive translation. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021a. URL <https://openreview.net/forum?id=ZTFeSBIX9C>.
- Ding, L., Wang, L., Liu, X., Wong, D. F., Tao, D., and Tu, Z. Rejuvenating low-frequency words: Making the most of parallel data in non-autoregressive translation. In Zong, C., Xia, F., Li, W., and Navigli, R. (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pp. 3431–3441. Association for Computational Linguistics, 2021b. doi: 10.18653/v1/2021.acl-long.266. URL <https://doi.org/10.18653/v1/2021.acl-long.266>.
- Du, C., Tu, Z., and Jiang, J. Order-agnostic cross entropy for non-autoregressive machine translation. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 2849–2859. PMLR, 2021. URL <http://proceedings.mlr.press/v139/du21c.html>.
- Fox, C. W. and Roberts, S. J. A tutorial on variational bayesian inference. *Artif. Intell. Rev.*, 38(2):85–95, 2012. doi: 10.1007/s10462-011-9236-8. URL <https://doi.org/10.1007/s10462-011-9236-8>.
- Ghazvininejad, M., Levy, O., Liu, Y., and Zettlemoyer, L. Mask-predict: Parallel decoding of conditional masked language models. In Inui, K., Jiang, J., Ng, V., and Wan, X. (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pp. 6111–6120. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1633. URL <https://doi.org/10.18653/v1/D19-1633>.
- Ghazvininejad, M., Karpukhin, V., Zettlemoyer, L., and Levy, O. Aligned cross entropy for non-autoregressive machine translation. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3515–3523. PMLR, 2020. URL <http://proceedings.mlr.press/v119/ghazvininejad20a.html>.
- Gu, J. and Kong, X. Fully non-autoregressive neural machine translation: Tricks of the trade. In Zong, C., Xia, F., Li, W., and Navigli, R. (eds.), *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pp. 120–133. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.findings-acl.11. URL <https://doi.org/10.18653/v1/2021.findings-acl.11>.
- Gu, J., Bradbury, J., Xiong, C., Li, V. O. K., and Socher, R. Non-autoregressive neural machine translation. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=B118Bt1Cb>.
- Gu, J., Wang, C., and Zhao, J. Levenshtein transformer. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 11179–11189, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/675f9820626f5bc0afb47b57890b466e-Abstract.html>.
- Guo, J., Xu, L., and Chen, E. Jointly masked sequence-to-sequence model for non-autoregressive neural machine translation. In Jurafsky, D., Chai, J., Schluter, N., and Tetreault, J. R. (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pp. 376–385. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.acl-main.36. URL <https://doi.org/10.18653/v1/2020.acl-main.36>.
- Hassan, H., Aue, A., Chen, C., Chowdhary, V., Clark, J., Federmann, C., Huang, X., Junczys-Dowmunt, M., Lewis,

- W., Li, M., Liu, S., Liu, T., Luo, R., Menezes, A., Qin, T., Seide, F., Tan, X., Tian, F., Wu, L., Wu, S., Xia, Y., Zhang, D., Zhang, Z., and Zhou, M. Achieving human parity on automatic chinese to english news translation. *CoRR*, abs/1803.05567, 2018. URL <http://arxiv.org/abs/1803.05567>.
- Huang, C., Zhou, H., Zaïane, O. R., Mou, L., and Li, L. Non-autoregressive translation with layer-wise prediction and deep supervision. *The Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022*, 2022a. URL <https://arxiv.org/abs/2110.07515>.
- Huang, F., Zhou, H., Liu, Y., Li, H., and Huang, M. Directed acyclic transformer for non-autoregressive machine translation. In *Proceedings of the 39th International Conference on Machine Learning, ICML 2022*, 2022b. URL <https://arxiv.org/abs/2205.07459>.
- Kaiser, L., Bengio, S., Roy, A., Vaswani, A., Parmar, N., Uszkoreit, J., and Shazeer, N. Fast decoding in sequence models using discrete latent variables. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2395–2404. PMLR, 2018. URL <http://proceedings.mlr.press/v80/kaiser18a.html>.
- Kasai, J., Cross, J., Ghazvininejad, M., and Gu, J. Non-autoregressive machine translation with disentangled context transformer. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 5144–5155. PMLR, 2020. URL <http://proceedings.mlr.press/v119/kasai20a.html>.
- Kasai, J., Pappas, N., Peng, H., Cross, J., and Smith, N. A. Deep encoder, shallow decoder: Reevaluating non-autoregressive machine translation. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=KpfasTaLUpq>.
- Kim, Y. and Rush, A. M. Sequence-level knowledge distillation. In Su, J., Carreras, X., and Duh, K. (eds.), *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pp. 1317–1327. The Association for Computational Linguistics, 2016. doi: 10.18653/v1/d16-1139. URL <https://doi.org/10.18653/v1/d16-1139>.
- Lee, J., Mansimov, E., and Cho, K. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J. (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pp. 1173–1182. Association for Computational Linguistics, 2018. doi: 10.18653/v1/d18-1149. URL <https://doi.org/10.18653/v1/d18-1149>.
- Libovický, J. and Helcl, J. End-to-end non-autoregressive neural machine translation with connectionist temporal classification. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J. (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pp. 3016–3021. Association for Computational Linguistics, 2018. doi: 10.18653/v1/d18-1336. URL <https://doi.org/10.18653/v1/d18-1336>.
- Ma, X., Zhou, C., Li, X., Neubig, G., and Hovy, E. H. Flowseq: Non-autoregressive conditional sequence generation with generative flow. In Inui, K., Jiang, J., Ng, V., and Wan, X. (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pp. 4281–4291. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1437. URL <https://doi.org/10.18653/v1/D19-1437>.
- Ott, M., Auli, M., Grangier, D., and Ranzato, M. Analyzing uncertainty in neural machine translation. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3953–3962. PMLR, 2018. URL <http://proceedings.mlr.press/v80/ott18a.html>.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. fairseq: A fast, extensible toolkit for sequence modeling. In Ammar, W., Louis, A., and Mostafazadeh, N. (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Demonstrations*, pp. 48–53. Association for Computational Linguistics, 2019. doi: 10.18653/v1/n19-4009. URL <https://doi.org/10.18653/v1/n19-4009>.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of*

- the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pp. 311–318. ACL, 2002. doi: 10.3115/1073083.1073135. URL <https://aclanthology.org/P02-1040/>.
- Qian, L., Zhou, H., Bao, Y., Wang, M., Qiu, L., Zhang, W., Yu, Y., and Li, L. Glancing transformer for non-autoregressive neural machine translation. In Zong, C., Xia, F., Li, W., and Navigli, R. (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pp. 1993–2003. Association for Computational Linguistics, 2021a. URL <https://aclanthology.org/2021.acl-long.155>.
- Qian, L., Zhou, Y., Zheng, Z., Zhu, Y., Lin, Z., Feng, J., Cheng, S., Li, L., Wang, M., and Zhou, H. The volctrans GLAT system: Non-autoregressive translation meets WMT21. *CoRR*, abs/2109.11247, 2021b. URL <https://arxiv.org/abs/2109.11247>.
- Ran, Q., Lin, Y., Li, P., and Zhou, J. Guiding non-autoregressive neural machine translation decoding with reordering information. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pp. 13727–13735. AAAI Press, 2021. URL <https://ojs.aaai.org/index.php/AAAI/article/view/17618>.
- Saharia, C., Chan, W., Saxena, S., and Norouzi, M. Non-autoregressive machine translation with latent alignments. In Webber, B., Cohn, T., He, Y., and Liu, Y. (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pp. 1098–1108. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.emnlp-main.83. URL <https://doi.org/10.18653/v1/2020.emnlp-main.83>.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics, 2016. doi: 10.18653/v1/p16-1162. URL <https://doi.org/10.18653/v1/p16-1162>.
- Shao, C., Zhang, J., Feng, Y., Meng, F., and Zhou, J. Minimizing the bag-of-ngrams difference for non-autoregressive neural machine translation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 198–205. AAAI Press, 2020. URL <https://aaai.org/ojs/index.php/AAAI/article/view/5351>.
- Shao, C., Feng, Y., Zhang, J., Meng, F., and Zhou, J. Sequence-level training for non-autoregressive neural machine translation. *Comput. Linguistics*, 47(4):891–925, 2021. doi: 10.1162/coli_a_00421. URL https://doi.org/10.1162/coli_a_00421.
- Shu, R., Lee, J., Nakayama, H., and Cho, K. Latent-variable non-autoregressive neural machine translation with deterministic inference using a delta posterior. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 8846–8853. AAAI Press, 2020. URL <https://aaai.org/ojs/index.php/AAAI/article/view/6413>.
- Studený, M. and Vejnárová, J. The multiinformation function as a tool for measuring stochastic dependence. In Jordan, M. I. (ed.), *Learning in Graphical Models*, volume 89 of *NATO ASI Series*, pp. 261–297. Springer Netherlands, 1998. doi: 10.1007/978-94-011-5014-9_10. URL https://doi.org/10.1007/978-94-011-5014-9_10.
- Sun, Z. and Yang, Y. An EM approach to non-autoregressive conditional sequence generation. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9249–9258. PMLR, 2020. URL <http://proceedings.mlr.press/v119/sun20c.html>.
- Sun, Z., Li, Z., Wang, H., He, D., Lin, Z., and Deng, Z. Fast structured decoding for sequence models. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 3011–3020, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/74563ba21a90da13dacf2a73e3ddefa7-Abstract.html>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I.

- Attention is all you need. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5998–6008, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- Watanabe, M. S. Information theoretical analysis of multivariate correlation. *IBM J. Res. Dev.*, 4(1):66–82, 1960. doi: 10.1147/rd.41.0066. URL <https://doi.org/10.1147/rd.41.0066>.
- Wei, B., Wang, M., Zhou, H., Lin, J., and Sun, X. Imitation learning for non-autoregressive neural machine translation. In Korhonen, A., Traum, D. R., and Màrquez, L. (eds.), *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pp. 1304–1312. Association for Computational Linguistics, 2019a. doi: 10.18653/v1/p19-1125. URL <https://doi.org/10.18653/v1/p19-1125>.
- Wei, B., Wang, M., Zhou, H., Lin, J., and Sun, X. Imitation learning for non-autoregressive neural machine translation. In Korhonen, A., Traum, D. R., and Màrquez, L. (eds.), *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pp. 1304–1312. Association for Computational Linguistics, 2019b. doi: 10.18653/v1/p19-1125. URL <https://doi.org/10.18653/v1/p19-1125>.
- Xu, Y., Zhao, S., Song, J., Stewart, R., and Ermon, S. A theory of usable information under computational constraints. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=r1eBeyHFDH>.
- Yang, K., Lei, W., Liu, D., Qi, W., and Lv, J. Pos-constrained parallel decoding for non-autoregressive generation. In Zong, C., Xia, F., Li, W., and Navigli, R. (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pp. 5990–6000. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.acl-long.467. URL <https://doi.org/10.18653/v1/2021.acl-long.467>.
- Zhou, C., Gu, J., and Neubig, G. Understanding knowledge distillation in non-autoregressive machine trans-

A. Relation to Iterative NATs

Although MPLE provides a unified perspective to understand many previous methods, we do not discuss an important branch of the NAT model, i.e., the iterative NATs. The reason lies in the conditional independent assumption of Eq.3, which is the basic assumption of our analyses but not satisfied in the iterative methods. However, our perspective can also improve the understanding of iterative NATs, and we point out some important relations.

\mathcal{C} measures the information loss in each iterative step. Although iterative NATs do not satisfy Eq.3, they still predict tokens independently in each refinement step, so the minimum information loss can be measured by \mathcal{C} . For example, in some popular iterative NATs (Lee et al., 2018; Ghazvininejad et al., 2019), the i -th refinement step’s log-likelihood is defined as

$$\log P_{\theta}^{\text{NAT}}(Y|X, Y^{(i-1)}) = \sum_{k=1}^M \log P_{\theta}^{\text{NAT}}(y_k|X, Y^{(i-1)}), \quad (16)$$

where Y is the target sentence, $Y^{(i-1)}$ is the refined result of previous steps, and M is the target length. Similar to Theorem 1, we can prove that the minimal KL divergence is \mathcal{C}_i , i.e., the conditional TC of the target distribution when $Y^{(i-1)}$ is given, where

$$\mathcal{C}_i = \sum_{k=1}^M H_{\text{data}}(y_k|X, Y^{(i-1)}) - H_{\text{data}}(Y|X, Y^{(i-1)}), \quad (17)$$

The result implies that iterative NATs also suffer from the information loss due to the dependency dropping and explains why they can benefit from methods that reduces \mathcal{C} , e.g., knowledge distillation.

Iterative Masked Prediction are special cases of MPLE with shared Input Predictor and NAT decoder. If we regard the output of the $(i - 1)$ -th refinement as the proxy input, iterative NATs actually construct a proxy distribution $P_{\text{data}}(Y|X, Y^{(i-1)})$, which reduces \mathcal{C} by providing an extra decoder input. Unlike the input predictor defined in our MPLE framework, iterative NATs predict the proxy input $Y^{(i-1)}$ by the NAT decoder itself with $i - 1$ refinement steps. In Table.2, we propose Input Sampling method that generates Z from the input predictor, which is similar to a single step refinement but with a separate NAT decoder.

B. Conditional TC and Performance Gap

Table 1 aims to show that the large \mathcal{C} is the main obstacle in NAT learning, and we provide more details here.

Dataset We first choose WMT14 En-De and WMT16 En-Ro, which contains 4.5M pairs and 610k pairs in the training set, respectively. Since natural datasets usually have a large \mathcal{C} , we further construct two synthetic datasets for comparison. Both synthetic datasets use the English corpus in WMT14 En-De as targets, and the source sentences are modified from the targets by word replacement or word dropping. In other words, the synthetic dataset trains the NAT to generate clean English sentences from corrupted English inputs. Specifically, *Synthetic A* replaces 50% of tokens by randomly sampled tokens from the vocabulary. *Synthetic B* further drops 10% of tokens in the source sentences based on Synthetic A.

Estimation of Conditional TC To estimate \mathcal{C} , we use *V-entropy* (Xu et al., 2020) instead of the Shannon entropy because the latter is intractable due to the unknown data distribution. The V-entropy is comparable only when the function family for estimation is fixed. In our implementation, we use *Transformers-base* as the function family.

More Rigorous Comparison Since BLEU is not strictly comparable across datasets, we present a more rigorous comparison by estimating the parameter size required for an autoregressive Transformer to achieve a similar performance with NAT. The comparison is based on the assumption that a smaller AT will suffer from more information loss than a larger AT. We use *Transformer-base* for NATs, and the AT architecture is choose from Table 6.

As shown in Table.5, we find that an AT only requires about 2.1% ~ 3.2% of parameters to achieve similar performance with the NAT on WMT14 En-De. However, on Synthetic A, an AT requires at least 27.3% of parameters to compete with the NAT. The results verify that large \mathcal{C} brings much information loss, making AT easily outperform NAT with much less parameters.

Table 5: Estimated \mathcal{C} and parameter size for an AT to achieve similar performance with NAT on various dataset. *Parameter Ratio* is the ratio of AT parameter size and NAT parameter size. For example, ATs with 2.1% ~ 3.2% of parameters achieve 4.60 ~ 11.88 BLEU on WMT14 En-De, which is similar to the NAT performance (11.79). The large \mathcal{C} leads to serious information loss in NAT learning, therefore a small AT can easily outperforms the NAT.

Dataset	\mathcal{C}	Parameter Ratio	BLEU _{NAT}	BLEU _{AT}
WMT14 En-De	2.50	2.1% ~ 3.2%	11.79	4.60 ~ 11.88
WMT16 En-Ro	2.20	3.2% ~ 4.3%	23.72	20.50 ~ 24.75
Synthetic B	1.51	4.3% ~ 7.0%	15.31	14.10 ~ 16.10
Synthetic A	0.92	27.3% ~ 100%	26.61	23.99 ~ 26.96

Table 6: AT Architectures used in searching the parameter size.

	d_{model}	d_{hidden}	n_{layers}	n_{heads}	# Param	Parameter Ratio
1	32	128	2	2	1.3M	2.1%
2	48	192	2	2	2.0M	3.2%
3	64	256	2	2	2.7M	4.3%
4	96	384	2	2	4.4M	7.0%
5	128	512	3	4	6.6M	10.5%
6	256	1024	3	4	17.1M	27.3%
7	512	2048	6	8	62.6M	100%

C. Results on WMT17 Zh-En

We repeat the experiments in Sec.4.2 and Sec.4.3 on WMT17 Zh-En. As shown in Table 7 and Table 8, our objective is strongly correlated with the translation quality, which supports our claim well.

Table 7: Comparison of methods that obtain proxy targets on WMT17 Zh-En. All methods use Vanilla and $\mathcal{L}_{input}=0$. $\hat{\mathcal{L}}_{MPL}$ and BLEU are strongly correlated (Pearson’s $|r|=0.96$). AXE’s τ and OaXE’s numbers indicate the skip penalty and the pre-training step, which are hyper-parameters in choosing proxy targets.

Models	\mathcal{L}_{NAT}	$\hat{\mathcal{L}}_{target}$	$\hat{\mathcal{L}}_{MPL}$	BLEU
Raw Data	4.43	-6.25	-1.82	8.69
KD	2.85	-5.72	-2.87	15.53
+ AXE($\tau=1$)	1.02	-2.95	-1.93	9.68
+ AXE($\tau=5$)	1.93	-5.00	-3.07	18.39
+ AXE($\tau=10$)	2.31	-5.20	-2.90	18.25
+ OaXE(10k)	1.46	-3.38	-1.92	12.31
+ OaXE(50k)	1.19	-4.50	-3.31	18.79
+ OaXE(300k)	1.15	-4.66	-3.50	19.46

Table 8: Comparison of methods and variants that obtain proxy inputs on WMT17 Zh-En. All methods use KD and $\hat{\mathcal{L}}_{target} = -5.72$. Sample and Default indicate the BLEU score in Input Sampling and Default Decoding. $\hat{\mathcal{L}}_{MPL}$ is strongly correlated with Sample BLEU (Pearson’s $|r|=0.99$) but less correlated with Default BLEU ($|r|=0.35$).

Models	\mathcal{L}_{input}	\mathcal{L}_{NAT}	$\hat{\mathcal{L}}_{MPL}$	Sample	Default
Vanilla	0	2.85	-2.87	15.53	15.53
CMLM	1.13	0.76	-3.83	19.74	14.12
+ Fixed	0.29	1.13	-4.30	20.81	14.49
GLAT	0.33	1.26	-4.13	20.73	22.51
+ Levenshtein	1.08	0.44	-4.19	20.71	21.70
+ P_{ref}	0.40	1.60	-3.73	18.98	21.22
+ $1 - P_{ref}$	0.73	0.81	-4.17	20.79	21.51

D. Details and Full Results of Dynamic KD

Decoding Tricks In Table 4, we apply some decoding tricks for the results on the last row:

- We use length parallel decoding (LPD, Wei et al., 2019b). We use a candidate set of 3. Since all candidates can be generated simultaneously, LPD is still much fast in inference. It is worth noting that LPD is faster than NPD (Gu et al., 2018) since it does not need an external reranker.
- We use the de-duplication trick (Lee et al., 2018), i.e., removing the repeated tokens in generated sentences.
- We adjust the predicted length by a factor λ (Ghazvininejad et al., 2020). The factor is tuned on the validation set. We use $\lambda = 1$ (i.e., the predicted length is not changed) for WMT14 En-De, and $\lambda = 1.05$ for WMT17 Zh-En.

Full Results In Table 9, we compare Dynamic KD against strong baselines including non-iterative and iterative NATs. Moreover, we justify the necessity of the regularizer in Eq.8 by removing $\hat{\mathcal{L}}_{\text{target}}$ in choosing the proxy target (i.e., $T^* = \arg \min_T \mathcal{L}_{\text{NAT}}$) as an ablation study.

Table 9: Full results of comparing Dynamic KD against AT and previous NATs. The best results of Non-iterative NATs are bolded. NPD (Gu et al., 2018) and LPD (Wei et al., 2019a) indicate reranking methods with the number of candidates. NPD are slower than LPD due to the use of an external AT reranker. The ablation of regularizer in Dynamic KD indicates removing $\hat{\mathcal{L}}_{\text{target}}$ in choosing the proxy target. †: Results reported by previous studies. Speed up of Imputer are re-evaluated in our implementation. §: Use adaptive iteration numbers.

	Models	# Iters	En-De	Zh-En	Speedup
AT	Transformer	L	27.11	23.89	1.0x
Iterative NATs	CMLM† (Ghazvininejad et al., 2019)	4	25.94	21.90	3.0x
		10	27.03	23.21	1.3x
		4	25.83	22.42	4.3x
	DisCo† (Kasai et al., 2020)	10	27.06	23.68	3.2x
		≈ 4 §	27.34	23.83	/
		1	25.8	/	14.9x
Non-iterative NATs	Imputer† (Saharia et al., 2020)	2	27.5	/	7.5x
		8	28.2	/	2.7x
	MLE	1	11.79	8.69	15.3x
	GLAT (NPD=7)† (Qian et al., 2021a)	1	26.55	/	7.9x
	OaXE (LPD=5)† (Du et al., 2021)	1	26.1	22.1	14.2x
Ours	Vanilla + KD	1	20.98	15.53	15.3x
	+ Dynamic KD w/o Regularizer	1	20.51	18.10	15.3x
	+ Dynamic KD	1	22.82	18.29	15.3x
	+ LPD=3 + Decoding Tricks	1	24.83	19.97	14.6x
	GLAT + KD	1	25.12	22.51	15.3x
	+ Dynamic KD w/o Regularizer	1	22.66	22.14	15.3x
	+ Dynamic KD	1	25.88	23.07	15.3x
	+ LPD=3 + Decoding Tricks	1	26.89	24.42	14.6x

E. Formalization of Existing Methods in MPLE

In the main paper, we briefly describe how existing methods obtain Z and T . In this section, we present detailed formalization of these methods by describing the heuristic rules and their objectives in the framework of MPLE. Specifically, we formalize each method in two steps:

First, we define the variational distribution following their heuristic rule. Existing methods use heuristic rule to obtain Z and T , which builds the variational distribution $Q(T, Z|X)$ used in the derivation of MPLE (Eq.6). For all methods in our analysis, their variational distribution is defined by

$$Q(T, Z|X) := Q(T|X)Q(Z|T, X), \quad (18)$$

where $Q(T|X)$ is defined by the methods that obtain the proxy target (including Raw Data, KD, AXE, and OaXE), and $Q(Z|T, X)$ is defined by the methods that obtain the proxy input (including Vanilla, CMLM, and GLAT).

Second, we prove that their original objective is equivalent to minimizing \mathcal{L}_{NAT} of Eq.9. Notably, in M-step, Q is unchanged when optimizing the model θ as discussed in Sec.3.3, so we only prove that their objective is equivalent to minimizing the NLL:

$$\mathcal{L}_{\text{NAT}} = -\mathbb{E}_{Q(Z,T|X)} [\log P_{\theta}(T|Z, X)] + \text{Constant} \quad (19)$$

For some methods, the proof is trivial and thus omitted in the following sections.

E.1. Raw Data

Raw Data uses the original target sentence as the proxy target. Formally, it defines $Q(T|X)$ as a one-point distribution that $Q(T = Y^*|X) = 1$, where Y^* is the original target in the dataset.

E.2. Knowledge Distillation (KD, Gu et al., 2018)

KD first trains an autoregressive model P_{AR} on the raw data, and then uses beam search to obtain $T^* = \arg \max_Y P_{AR}(Y|X)$. Formally, $Q(T|X)$ is defined as a one-point distribution at T^* .

E.3. Aligned Cross Entropy (AXE, Ghazvininejad et al., 2020)

$Q(T|X)$ is defined as a one-point distribution at T^* , where $T^* = \arg \min_{T \in \mathcal{S}(R)} \mathcal{L}_{\text{NAT}}$. The reference $R = [r_1, \dots, r_L]$ is picked from Raw Data or KD. Any $T \in \mathcal{S}(R)$ is a subsequence of R with empty tokens ϵ inserted.¹⁰ An example is shown in Fig.9.

Original Objective AXE introduces a monotonic alignment $\alpha = [\alpha_1, \dots, \alpha_L]$, where the i -th token of the reference R is aligned to the α_i -th token of the NAT prediction. Formally, the AXE loss is defined as

$$\mathcal{L}_{\text{AXE}} = \min_{\alpha} \left[-\sum_{i=1}^L \log P_{\alpha_i}(r_i) - \sum_{k \notin \alpha} \log P_k(\epsilon) \right],$$

s.t. $1 \leq \alpha_1 \leq \dots \leq \alpha_L \leq L$.

The first term is the cross entropy between aligned targets and predictions, and the second term is a penalty for unaligned predictions.

In the AXE loss, a single prediction may be aligned to multiple target tokens. In their original paper, aligning the prediction to the first target token is called the ‘‘align’’ operation, and aligning the prediction to later tokens is called the ‘‘skip target’’ operation. However, a one-to-many alignment will damage the performance, so they penalize the ‘‘skip target’’ operations with a factor δ . This trick is called the *skip penalty*.

Proof of Equivalence To connect their definition with ours, we convert the alignment to an adjacency list, as shown in Figure 9, where β_i is a list containing all aligned tokens for the i -th prediction. Specially, if the i -th prediction is not aligned, we set $\beta_i = [0]$ and $r_0 = \epsilon$. Then, \mathcal{L}_{AXE} can be reformulated as

$$\min_{\beta} \left[-\sum_{i=1}^L \log P_i(r_{\beta_{i,1}}) - \delta \sum_{i=1}^L \sum_{j=2}^{|\beta_i|} \log P_i(r_{\beta_{i,j}}) \right], \quad (20)$$

where $\beta_{i,j}$ indicates the j -th element of β_i . The first term is the cross entropy between the prediction and a new target $T^* = [r_{\beta_{1,1}}, \dots, r_{\beta_{L,1}}]$, and the second term is the penalty for ‘‘skipping target’’ operations.

When $\delta = 0$, Eq.20 is equivalent to finding an optimal T^* to minimize \mathcal{L}_{NAT} in Eq.19. Since α is a monotonic alignment, T^* is constrained and should be a subsequence of R with some empty tokens inserted, which recover our definition. When $\delta \neq 0$, the second term can be regarded as a regularizer to control the distortion between proxy targets and real targets.

E.4. Order-agnostic Cross Entropy (OaXE, Du et al., 2021)

OaXE is similar to AXE despite the constraint $\mathcal{S}(R)$. Any $T \in \mathcal{S}(R)$ is a permutation of R . An example is shown in Fig.9.

¹⁰The NAT model may learn to predict empty tokens, which will be removed after generation.




Figure 9: Examples of the alignment α , the adjacency list β , and the proxy target T^* in AXE and OaXE. Red dotted line in AXE indicates the *skip target* operation.

Original Objective Different from AXE, OaXE’s α is a non-monotonic alignment, and each predicted token can only be used once. The loss is defined as

$$\mathcal{L}_{\text{OaXE}} = \min_{\alpha \in \text{Perm}(L)} \left[- \sum_{i=1}^L \log P_{\alpha_i}(r_i) \right],$$

where $\text{Perm}(L)$ indicates the permutations of sequences containing 1 to L .

Proof of Equivalence Similar to the derivation for AXE, we can reformulate $\mathcal{L}_{\text{OaXE}}$ as

$$\min_{\beta} \left[- \sum_{i=1}^L \log P_i(r_{\beta_{i,1}}) \right].$$

The above formulation recovers our definition: It finds an optimal T^* to minimize \mathcal{L}_{NAT} in Eq.19, where T^* can be an arbitrary permutation of R .

However, without the monotonic constraints, T^* in OaXE may be heavily distorted from the real target Y . To alleviate the problem, OaXE first pretrains the NAT using the vanilla MLE and then finetunes it to minimize $\mathcal{L}_{\text{OaXE}}$. This trick is based on an intuition that the optimal T^* in a well-trained NAT will be close to the real target.

E.5. Vanilla

Many NATs use a full masked sequence as Z or predict Z by Uniform Copy (Gu et al., 2018) or attention (Qian et al., 2021a). We regard them as vanilla decoder inputs because they do not introduce any hints from the target. Formally, they can be formulated as a one-point distribution $Q(Z = Z^* | T, X) = 1$, where Z^* is obtained from a deterministic function $f(X)$.

E.6. CMLM (Ghazvininejad et al., 2019) / GLAT (Qian et al., 2021a)

CMLM and GLAT sample the proxy input by randomly masking the target sentence. Specifically, they first sample $l \in [1, L]$ as the number of unmasked tokens, and then obtain the proxy input by randomly masking $L - l$ tokens.

In Table 3, we compare CMLM, GLAT and their variants. Here we list their differences:

- **CMLM** sets the number of unmasked tokens $l = \lambda L$, where λ is uniformly sampled from 0 to 1.
- **CMLM + fixed masking ratio** uses $l = 0.2L$ instead of random sampling.
- **GLAT** uses an adaptive sampling strategy according to the NAT prediction accuracy. Specifically, $l = \lambda \sum_{i=1}^L [T_i \neq \arg \max(t_i | X)]$. We follow their original paper and anneal λ from 0.5 to 0.3.
- **GLAT + mask by P_{ref}** use the same l as GLAT, but chooses the unmasked tokens according to the difficulties in predicting them, where the probability of an unmasked z_i is proportional to the prediction probability $P_{\theta}(t_i | X)$.
- **GLAT + mask by $1 - P_{\text{ref}}$** chooses the unmasked tokens proportional to $1 - P_{\theta}(t_i | X)$.

Implementation Details of Input Predictor In Eq.15, we mention that Z is predicted non-autoregressively. Concretely, $P_\theta(z_i|X)$ is composed of two modules: $P_\theta(z_i \text{ is masked}|X)$ predicts whether z_i is a masked token, and $P_\theta(t_i|X)$ predicts the target token t_i from the vocabulary if z_i is not masked. Formally,

$$P_\theta(z_i|X) = \begin{cases} P_\theta(z_i \text{ is masked}|X), & \text{if } z_i \text{ is masked;} \\ (1 - P_\theta(z_i \text{ is masked}|X))P_\theta(t_i|X), & \text{if } z_i = t_i; \\ 0, & \text{otherwise.} \end{cases}$$

Therefore, $\mathcal{L}_{\text{input}}$ can be formulated as

$$\mathcal{L}_{\text{input}} = \mathbb{E}_{Q(Z|X)} \left[- \sum_{i=1}^L \log P_\theta(z_i|X) + \log Q(Z|X) \right],$$

where $Q(Z|X)$ can be obtained according to the definition of heuristic rules.

For the first module $P_\theta(z_i \text{ is masked}|X)$, we reuse the Transformer encoder and the NAT decoder and further add a binary classification layer on top of the NAT decoder. For the second module $P_\theta(t_i|X)$, we use a pre-trained vanilla NAT and freeze its parameters during the training of CMLM or GLAT. In this way, $P_\theta(t_i|X)$ can be computed offline to speed up the training.

For Input Sampling used in Table 3, we only do sampling from $P_\theta(z_i \text{ is masked}|X)$. If z_i is not masked, we directly use $z_i = \arg \max_{t_i} P_\theta(t_i|X)$ because it empirically leads to better performance.

Original Objective In the original implementation, CMLM and GLAT use a masked language model objective, where the unmasked tokens are not included in the loss \mathcal{L}_{NAT} . Formally,

$$\mathcal{L}_{\text{NAT}} = \mathbb{E}_{Q(Z,T|X)} \left[- \sum_{i \notin \mathcal{G}} \log P_\theta(t_i|Z, X) \right],$$

where \mathcal{G} is the set of the unmasked token.

Proof of Equivalence To reach a same formulation of Eq.19, we add a *copy mechanism* in the NAT decoder. The decoder directly copies the unmasked token as the prediction if available. As a result, for an unmasked token t_i , $\log P_\theta(t_i|Z, X) = 0$ because the prediction of t_i is always correct. Therefore, the masked language model objective recovers our objective:

$$\begin{aligned} \mathcal{L}_{\text{NAT}} &= \mathbb{E}_{Q(Z,T|X)} \left[- \sum_{i \notin \mathcal{G}} \log P_\theta(t_i|Z, X) \right] + 0 \\ &= \mathbb{E}_{Q(Z,T|X)} \left[- \sum_{i \notin \mathcal{G}} \log P_\theta(t_i|Z, X) - \sum_{i \in \mathcal{G}} \log P_\theta(t_i|Z, X) \right] \\ &= \mathbb{E}_{Q(Z,T|X)} [- \log P_\theta(T|Z, X)] \end{aligned}$$

Note that the *copy mechanism* does not require modifications to the network architecture.

E.7. VAE (Shu et al., 2020)

Although not discussed in our main analysis, VAE and its variants (Kaiser et al., 2018; Bao et al., 2021; 2022) can also be formulated as a method to provide proxy input in MPLE. VAE uses two trainable networks, the prior and posterior networks, to model $P_\theta(Z|X)$ and $Q(Z|T, X)$, respectively. Specially, the posterior network $Q(Z|T, X)$ can be trained together with θ .

F. Implementation Details of Dynamic KD

Candidate Generation. Dynamic KD chooses the proxy target from a candidate set Γ , which contains Raw Data and four distilled targets. We generate the distilled targets with beam size 5 from four AT teachers, whose hyper-parameters and performance are shown in Table 10. For WMT14 En-De, we train the AT teachers for 100k updates with a batch of approximately 64k tokens. For WMT17 Zh-En, we raise the step to 300k to match the size of training data, and tune the length penalty in the beam search on the validation set.

Table 10: Hyper-parameters and performance of AT teachers, which generate the target candidates in Dynamic KD.

Model	<i>tiny</i>	<i>small</i>	<i>base</i>	<i>big</i>
d_{model}	128	256	512	1024
d_{hidden}	512	1024	2048	4096
n_{layers}	3	3	6	6
n_{heads}	4	4	8	8
Dropout	0.1	0.1	0.3	0.3
WMT14 En-De	20.46	24.29	27.11	28.49
WMT17 Zh-En	19.38	22.47	23.89	24.84

Candidate Selection Criterion. Dynamic KD chooses the proxy target T by minimizing $\mathcal{L}_{\text{NAT}} + \hat{\mathcal{L}}_{\text{target}}$. However, $\hat{\mathcal{L}}_{\text{target}}$ requires samples from $P_{\text{data}}(Y|X)$ as defined in Eq.14, which is intractable on the training set. To tackle the issue, we approximate $\hat{\mathcal{L}}_{\text{target}}$ by the pairwise BLEU between the candidates: ¹¹

$$\begin{aligned} \hat{\mathcal{L}}_{\text{target}} &= -\beta \mathbb{E}_{P_{\text{data}}(Y|X)} [S(Y, T)] \\ &\approx -\beta \sum_{i=1}^5 \gamma_i S(\Gamma_i, T) \triangleq \hat{\mathcal{L}}'_{\text{target}}, \end{aligned} \quad (21)$$

where S is the sentence BLEU, Γ_i is the target distilled from the i -th teacher model, and γ_i is hyper-parameters to bias the candidates from different teachers ($i = 5$ indicates Raw Data). We use $\hat{\mathcal{L}}'_{\text{target}}$ instead of $\hat{\mathcal{L}}_{\text{target}}$ in selecting the proxy target.

Hyper-parameter Selection. To find the optimal value of γ_i , we introduce the multi-reference dataset (Ott et al., 2018; Hassan et al., 2018), making it possible to adjust the value of $\hat{\mathcal{L}}'_{\text{target}}$ according to the real $\hat{\mathcal{L}}_{\text{target}}$. Intuitively, if \mathcal{L}_{NAT} is the same for all target candidates, we should choose a proxy target that minimize the data distortion. Therefore, we obtain $T = \arg \min \hat{\mathcal{L}}'_{\text{target}}$ as the current proxy target with a specific γ_i , and then evaluate the real data distortion $\hat{\mathcal{L}}_{\text{target}}$. We tune γ_i to minimize $\hat{\mathcal{L}}_{\text{target}}$. Notably, tuning γ_i only involves calculating the BLEU score, which does not need to train a NAT model. We do a manual search from 1 to 3 with the step of 0.1 and finally choose $\gamma = [0.7, 0.6, 1.3, 1.5, 2.3]$ for WMT14 En-De and $\gamma = [0.9, 1.4, 1.3, 0.9, 2.0]$ for WMT17 Zh-En. Then we train a NAT with the dynamic KD and further tune β according to the generation performance on the validation set, where we finally choose $\beta = 0.2$ for WMT14 En-De and $\beta = 0.1$ for WMT17 Zh-En.

Applying Dynamic KD to GLAT. When combining GLAT with Dynamic KD, \mathcal{L}_{NAT} may suffer from high variance because Z is sampled from $Q(Z|T)$ following the rule of GLAT. In our implementation, we simply ignore Z when choosing the proxy target. Specifically, we obtain the proxy target by $T^* = \arg \min_T \hat{\mathcal{L}}_{\text{NAT}} + \hat{\mathcal{L}}'_{\text{target}}$, where

$$\hat{\mathcal{L}}_{\text{NAT}} = \mathbb{E}_{Q(T|X)} [-\log P_{\theta}(T|Z^*, X)],$$

and Z^* is a full masked sequence.

G. Details of Experiment Settings

For WMT14 En-De, we follow Zhou et al. (2020) to use a joint BPE (Sennrich et al., 2016) with 32K merge operations, which leads to a vocabulary of 40k tokens. For WMT17 Zh-En, we follow Kasai et al. (2020) to use a BPE with 32K merge operations, which leads to vocabularies of 48k tokens in Chinese and 33k tokens in English.

All our models are implemented with Fairseq (Ott et al., 2019) and generally follow the hyper-parameter of transformer-*base* (Vaswani et al., 2017). For regularization, we set dropout to 0.1, weight decay to 0.01, and label smoothing to 0.1. Except for OaXE, all models are trained for 300k updates with a batch of approximately 64k tokens. The learning rate warms up to $5 \cdot 10^{-4}$ within 10k steps and then decays with the inverse square-root schedule. For OaXE, we choose a pre-trained vanilla NAT and finetune the model for 100k steps with a fixed learning rate of 10^{-5} . We evaluate the BLEU scores on the validation set every epoch and average the best 5 checkpoints for the final model. All models are trained with mixed precision floating point arithmetic on 8 Nvidia V100-32G GPUs. It costs approximately 20 hours for a vanilla NAT and 30 hours for Dynamic KD + GLAT.

¹¹We assume $Q(T|X)$ is a one-point distribution on the selected proxy target T .

On the Learning of Non-Autoregressive Transformers

For fair comparisons in Table 2 and 3, we do not use any decoding tricks and only modify the methods for obtaining Z and T . Taking OaXE as an example, our implementation differs from their original paper (Du et al., 2021) in: (1) Our OaXE is finetuned on a vanilla NAT, not a CMLM. (2) They use Transformer-*big* for KD whereas we use Transformer-*base*. (3) They use Length Parallel Decoding (Wei et al., 2019a) of beam 5 and the de-duplication trick (Lee et al., 2018) for decoding. We do not use any reranking methods here. (4) We do not use the truncation trick because it is incompatible with \mathcal{L}_{NAT} in our formulation.