
Aligned Cross Entropy for Non-Autoregressive Machine Translation

Marjan Ghazvininejad¹ Vladimir Karpukhin¹ Luke Zettlemoyer¹ Omer Levy¹

Abstract

Non-autoregressive machine translation models significantly speed up decoding by allowing for parallel prediction of the entire target sequence. However, modeling word order is more challenging due to the lack of autoregressive factors in the model. This difficulty is compounded during training with cross entropy loss, which can highly penalize small shifts in word order. In this paper, we propose aligned cross entropy (AXE) as an alternative loss function for training of non-autoregressive models. AXE uses a differentiable dynamic program to assign loss based on the best possible monotonic alignment between target tokens and model predictions. AXE-based training of conditional masked language models (CMLMs) substantially improves performance on major WMT benchmarks, while setting a new state of the art for non-autoregressive models.

1. Introduction

Non-autoregressive machine translation models can significantly improve decoding speed by predicting every word in parallel (Gu et al., 2018; Libovický & Helcl, 2018). This advantage comes at a cost to performance since modeling word order is trickier when the model cannot condition on its previous predictions. A range of semi-autoregressive models (Lee et al., 2018; Stern et al., 2019; Gu et al., 2019; Ghazvininejad et al., 2019) have shown there is a speed-accuracy tradeoff that can be optimized with limited forms of autoregression. However, increasing performance of the purely non-autoregressive models without sacrificing decoding speed remains an open challenge. In this paper, we present a new training loss for non-autoregressive machine translation that softens the penalty for word order errors, and significantly improves performance with no modification to the model or to the decoding algorithm.

¹Facebook AI Research. Correspondence to: Marjan Ghazvininejad <ghazvini@fb.com>.

Target Y	<i>it</i>	<i>tastes</i>	<i>pretty</i>	<i>good</i>	<i>though</i>
	but	it	tastes	delicious	.
Model	however	that	makes	good	,
Predictions	for	this	looks	tasty	so
P (Top 5)	and	for	taste	fine	though
	though	the	feels	exquisite	!

Figure 1. The model predictions are quite similar to the target, but misaligned by one token. The first and second target tokens (*it tastes*) are predicted in the second and third positions, respectively, leaving only the predictions in the fourth and fifth positions aligned with the target. The cross-entropy loss will heavily penalize the predictions in the first, second, and third positions.

Existing models (both autoregressive and non-autoregressive) are typically trained with cross entropy loss. Cross entropy is a strict loss function, where a penalty is incurred for every word that is predicted out of position, even for output sequences with small edit distances (see Figure 1). Autoregressive models learn to avoid such penalties, since words are generated conditioned on the sentence prefix. However, non-autoregressive models do not know the exact sentence prefix, and should (intuitively) focus more on root errors (e.g. a missing word) while allowing more partial credit for cascading errors (the right word in the wrong place).

To achieve this more relaxed loss, we introduce aligned cross entropy (AXE), a new objective function that computes the cross entropy loss based on an alignment between the sequence of token labels and the sequence of token distribution predictions. AXE uses dynamic programming to find the monotonic alignment that minimizes the cross entropy loss. It provides non-autoregressive models with a more accurate training signal by ignoring absolute positions and focusing on relative order and lexical matching. We efficiently implement AXE via matrix operations, and use it to train conditional masked language models (CMLM; Ghazvininejad et al., 2019) for machine translation. AXE only slightly increases training time compared to cross entropy, and requires no changes to parallel argmax decoding.

Extensive experiments on machine translation benchmarks demonstrate that AXE substantially boosts the performance of CMLMs, while having the same decoding speed. In WMT’14 EN-DE, training CMLMs with AXE (instead of

Algorithm 1 Aligned Cross Entropy

Input: tokens Y , predictions P
 $A_{0,0} = 0$
for $i = 1$ **to** n **do**
 $A_{i,0} = A_{i-1,0} - \delta \cdot \log P_1(Y_i)$
end for
for $j = 1$ **to** m **do**
 $A_{0,j} = A_{0,j-1} - \log P_j(\varepsilon)$
end for
for $i = 1$ **to** n **do**
 for $j = 1$ **to** m **do**
 align = $A_{i-1,j-1} - \log P_j(Y_i)$
 skip_prediction = $A_{i,j-1} - \log P_j(\varepsilon)$
 skip_target = $A_{i-1,j} - \delta \cdot \log P_j(Y_i)$
 $A_{i,j} = \min\{\text{align}, \text{skip_prediction}, \text{skip_target}\}$
 end for
end for
return $A_{n,m}$

the regular cross entropy loss) increases performance by 5 BLEU points; we observe similar trends in WMT’16 EN-RO and WMT’17 EN-ZH. Moreover, AXE CMLMs significantly outperform state-of-the-art non-autoregressive models, such as FlowSeq (Ma et al., 2019), as well as the recent CRF-based *semi-autoregressive* model with bigram LM decoding (Sun et al., 2019). Our detailed analysis suggests that training with AXE makes models more confident in their predictions, thus reducing multimodality, and alleviating a key problem in non-autoregressive machine translation.

2. Aligned Cross Entropy

Let Y be a target sequence of n tokens Y_1, \dots, Y_n , and P be the model predictions, a sequence of m token probability distributions P_1, \dots, P_m . Our goal is to find a monotonic alignment between Y and P that will minimize the cross entropy loss, and thus focus the penalty on lexical errors (predicting the wrong token) rather than positional errors (predicting the right token in the wrong place).

We define an alignment α to be a function that maps target positions to prediction positions, i.e. $\alpha : \{1, \dots, n\} \rightarrow \{1, \dots, m\}$. We further assume that this alignment is monotonic, i.e. $i \leq j$ iff $\alpha(i) \leq \alpha(j)$. Given a specific alignment α , we define a conditional AXE loss as:

$$\begin{aligned}
 AXE(Y_1, \dots, Y_n, P_1, \dots, P_m | \alpha) = & \\
 & - \sum_{i=1}^n \log P_{\alpha(i)}(Y_i) \\
 & - \sum_{k \in \{1 \dots m\} \setminus \{\alpha(1), \dots, \alpha(n)\}} \log P_k(\varepsilon)
 \end{aligned} \tag{1}$$

The first term of this loss function is an aligned cross entropy between Y and P , and the second term is a penalty for unaligned predictions. Epsilon (ε) is a special “blank” token in our vocabulary that appears in the probability distributions, but that does not appear in the final output string.

Now, the final AXE loss is the minimum over all possible monotonic alignments of the conditional loss:

$$\begin{aligned}
 AXE(Y_1, \dots, Y_n, P_1, \dots, P_m) = & \\
 \min_{\alpha} AXE(Y_1, \dots, Y_n, P_1, \dots, P_m | \alpha) = & \\
 \min_{\alpha(1) \dots \alpha(n)} \left(- \sum_{i=1}^n \log P_{\alpha(i)}(Y_i) \right. & \\
 \left. - \sum_{k \in \{1 \dots m\} \setminus \{\alpha(1), \dots, \alpha(n)\}} \log P_k(\varepsilon) \right) & \tag{2}
 \end{aligned}$$

$$\text{s.t. } 1 \leq \alpha(1) \leq \alpha(2) \leq \alpha(3) \dots \leq \alpha(n) \leq m$$

Finding the optimal monotonic alignment between two sequences is a well studied problem. For instance, dynamic time warping (DTW) (Sakoe & Chiba, 1978) is a well-known algorithm for finding the optimal alignment between two different time series. Here we have extended the idea to compute the optimal alignment between a sequence of target tokens and a sequence of prediction probability distributions. We use a simple dynamic program to find the optimal alignment while calculating the AXE loss.

Dynamic Programming Given a sequence of target tokens $Y = Y_1 \dots Y_n$ and a sequence of predictions $P = P_1 \dots P_m$ we propose a method to find the score of the optimal alignment between any prefix of these two sequences $Y_{1:i} = Y_1 \dots Y_i$ and $P_{1:j} = P_1 \dots P_j$, for any i and j . The score of the optimal alignment for the full sequences is obtained at $i = n$ and $j = m$.

We start by defining a matrix A of $n + 1$ by $m + 1$ dimensions, respectively corresponding to Y and P , where $A_{i,j}$ represents the minimum loss value for aligning $Y_{1:i}$ to $P_{1:j}$ as defined in Equation 2. We initialize $A_{0,0}$ to be 0 and then proceed to fill the matrix by taking the local minimum at each cell $A_{i,j}$ from three possible operators: *Align*, *Skip Prediction*, and *Skip Target*. Table 1 describes each operation and its update formula. Once the matrix is full, the cell $A_{n,m}$ will contain the cross entropy loss of the optimal alignment. Algorithm 1 lays out a straightforward implementation of AXE’s dynamic program.

According to Equation 2, the optimal alignment can be many-to-one, where multiple target positions can be mapped to a single prediction. This would be computed by *align*-

Align	Aligns the current target Y_i with the current prediction P_j , updating A along the diagonal.	$A_{i,j} = A_{i-1,j-1} - \log P_j(Y_i)$
Skip Prediction	Skips the current prediction P_j by predicting an empty token (ε), updating A along the P axis. This operation is akin to inserting an empty token to the target sequence at the i -th position.	$A_{i,j} = A_{i,j-1} - \log P_j(\varepsilon)$
Skip Target	Skips the current target Y_i by predicting it without incrementing the prediction iterator j , updating A along the Y axis. This operation is akin to duplicating the prediction P_j . The hyperparameter δ controls how expensive this operation is; high values of δ will discourage alignments that skip too many target tokens.	$A_{i,j} = A_{i-1,j} - \delta \cdot \log P_j(Y_i)$

Table 1. The three local update operators in AXE’s dynamic program.

Target Y	<i>it</i>	<i>tastes</i>	<i>pretty</i>	<i>good</i>	<i>though</i>
Alignment $\alpha : Y \rightarrow P$	2	3	3	4	5
Model Predictions P (Top 5)	but	it	tastes	delicious	ε
	however	ε	makes	good	.
	ε	that	looks	tasty	,
	for	this	taste	fine	so
	and	for	feels	exquisite	though

Figure 2. An example illustrating how AXE aligns model predictions with the target sequence. The operations that created the optimal alignment in this example are: (1) skip prediction P_1 (cost: $\log P_1(\varepsilon)$), (2) align Y_1 to P_2 (cost: $\log P_2(\text{it})$), (3) align Y_2 to P_3 (cost: $\log P_3(\text{tastes})$), (4) skip target Y_3 (cost: $\delta \log P_3(\text{pretty})$), (5) align Y_4 to P_4 (cost: $\log P_4(\text{good})$), (6) align Y_5 to P_5 (cost: $\log P_5(\text{though})$).

ing the first mapped token and *skipping* the rest of target tokens. To discourage skipping too many target tokens, we penalize skip target operators separately with a parameter δ as described in Table 1. Setting $\delta = 1$ will result in the loss function defined in Equation 2, but as we show in our ablation study (Section 4.3), higher δ values yield better performance in practice.

Efficient Implementation The implementation in Algorithm 1 has $O(n \cdot m)$ time complexity. However, multiple updates of the matrix A can be parallelized on GPUs and other tensor-processing architectures. Rather than iterating over each cell, we iterate over each *anti-diagonal*, computing all the values along the anti-diagonal in parallel. In other words, we first compute the values of $[A_{0,1}, A_{1,0}]$, followed by $[A_{0,2}, A_{1,1}, A_{2,0}]$, etc. Since the number of anti-diagonals is $n + m + 1$, we arrive at a time complexity of $O(n + m)$. Since m is typically on the same order of magnitude as n , the linear cost of computing AXE during training becomes negligible compared to forward and backward passes through the model.¹

¹Batch implementation of this algorithm is straightforward. By doing so, we are able to achieve training times similar to (about 1.2 times slower than) training with cross entropy loss.

Example Figure 2 depicts an example application of AXE. We see that the predictions are generally good, but start with a shift with respect to the target. This misalignment would cause the regular cross entropy loss to severely penalize the first three predictions, even though P_2 and P_3 are correct when aligned with Y_1 and Y_2 . AXE, on the other hand, finds an alignment between the target and the predictions, which allows it to focus the penalty on the redundant prediction in P_1 and the missing token Y_3 , i.e. the root errors.

3. Training Non-Autoregressive Models

We use AXE to train conditional masked language models (CMLMs) for non-autoregressive machine translation (Ghazvininejad et al., 2019).²

3.1. Conditional Masked Language Models

A conditional masked language model takes a source sequence X and a partially-observed target sequence Y_{obs} as input, and predicts the probabilities of the masked (un-

²While in this work we apply AXE to CMLMs, the loss function can be used to train other models as well. We leave further investigation of this direction to future work.

observed) target sequence tokens Y_{mask} . The underlying architecture is an encoder-decoder transformer (Vaswani et al., 2017).

In the original paper, CMLMs are used for machine translation where a random subset of Y tokens are masked at training time. However, at inference all target tokens are masked ($Y = Y_{\text{mask}}$) and the length of Y (the number of masked tokens) is unknown. To estimate the length of Y , an auxiliary task is introduced to predict the target length based on the source sequence X .³

3.2. Adapting CMLMs to AXE

In our case, the model can also produce blank tokens (ε), which effectively shorten the predicted sequence’s length. To account for potentially skipped tokens during inference, we multiply the predicted length by a hyperparameter λ (which is tuned on the validation set) before applying argmax decoding.

3.3. Adapting the Training Objectives to AXE

Since this work focuses on the purely non-autoregressive setting, the entire target sequence will be masked at inference time ($Y_{\text{mask}} = Y$). The same does not have to hold for training; we can utilize partially observed sequences in order to provide the learner with easier and more focused training examples. We experiment with three variations:

Unobserved Input, Predict All All the tokens in the target sequence are masked, and the model is expected to predict all of them. This is a direct replication of the task at inference time. While AXE allows for the number of masked tokens m to be different from the length of the gold target sequence n , we found that setting $m = n$ produced better models in preliminary experiments.

Partially-Observed Input, Predict All As in the original CMLM training process, a random subset of the target sequence is masked before being passed onto the model as input.⁴ We then apply AXE on the entire sequence, regardless of which tokens were observed. When training on partially-observed inputs, we always set $m = n$ to avoid further alterations of the gold target sequence beyond masking.

Partially-Observed Input, Predict Masks The straightforward application of AXE to CMLM training (which ignores whether each token was masked or observed) works well in practice. However, we can also allow AXE to skip the observed tokens when computing cross entropy, and focus the training signal on the actual task. We do so by setting

$P_i(Y_i) = 1$ for every observed token Y_i ; i.e. if the i -th token is observed and is aligned with the prediction corresponding to the same position (P_i), there is no penalty. Our ablation studies show that this modification provides a modest but consistent boost in performance (see Section 4.3). As a result, we use this setting for training our model.

4. Experiments

We evaluate CMLMs trained with AXE on 6 standard machine translation benchmarks, and demonstrate that AXE significantly improves performance over cross entropy trained CMLMs and over recently-proposed non-autoregressive models as well.

4.1. Setup

Translation Benchmarks We evaluate our method on both directions of three standard machine translation datasets with various training data sizes: WMT’14 English-German (4.5M sentence pairs), WMT’16 English-Romanian (610k pairs), and WMT’17 English-Chinese (20M pairs). The datasets are tokenized into subword units using BPE (Sennrich et al., 2016).⁵ We use the same data and preprocessing as Vaswani et al. (2017), Lee et al. (2018), and Wu et al. (2019) for WMT’14 EN-DE, WMT’16 EN-RO, and WMT’17 EN-ZH respectively. We evaluate performance with BLEU (Papineni et al., 2002) for all language pairs, except for translating English to Chinese where we use SacreBLEU (Post, 2018).⁶

Hyperparameters We generally follow the transformer *base* hyperparameters (Vaswani et al., 2017): 6 layers for the encoder and decoder, 8 attention heads per layer, 512 model dimensions, and 2048 hidden dimensions. We follow the weight initialization schema from BERT (Devlin et al., 2018), and sample weights from $\mathcal{N}(0, 0.02)$, set biases to zero, and set layer normalization parameters to $\beta = 0$ and $\gamma = 1$. For regularization, we set dropout to 0.3, and use 0.01 L_2 weight decay and label smoothing with $\varepsilon = 0.1$. We train batches of 128k tokens using Adam (Kingma & Ba, 2015) with $\beta = (0.9, 0.999)$ and $\varepsilon = 10^{-6}$. The learning rate warms up to $5 \cdot 10^{-4}$ within 10k steps, and then decays with the inverse square-root schedule. We train all models for 300k steps. We measure the validation loss at the end of each epoch, and average the 5 best checkpoints based on their validation loss to create the final model. We train all models with mixed precision floating point arithmetic on 16 Nvidia V100 GPUs. For autoregressive decoding, we use a beam size of $b = 5$ (Vaswani et al., 2017) and tune

⁵We run joint BPE for all language pairs except English-Chinese.

⁶SacreBLEU hash: BLEU+case.mixed+lang.en-zh +num-refs.1+smooth.exp+test.wmt17+tok.zh+version.1.3.7

³See (Ghazvininejad et al., 2019) for further detail.

⁴The number of masked input tokens is distributed uniformly between 1 and n .

Aligned Cross Entropy for Non-Autoregressive Machine Translation

Model	WMT'14		WMT'16		WMT'17	
	EN-DE	DE-EN	EN-RO	RO-EN	EN-ZH	ZH-EN
Cross Entropy CMLM (Ghazvininejad et al., 2019)	18.05	21.83	27.32	28.20	24.23	13.64
AXE CMLM (Ours)	23.53	27.90	30.75	31.54	30.88	19.79

Table 2. The performance (test set BLEU) of AXE CMLM compared to cross entropy CMLM on all of our benchmarks. Both models are purely non-autoregressive, using a single forward pass during argmax decoding.

Model	Decoding Iterations	WMT'14		WMT'16	
		EN-DE	DE-EN	EN-RO	RO-EN
Autoregressive					
Transformer Base	<i>N</i>	27.61	31.38	34.28	33.99
+ Knowledge Distillation	<i>N</i>	27.75	31.30	—	—
Non-Autoregressive					
Iterative Refinement (Lee et al., 2018)	1	13.91	16.77	24.45	25.73
CTC Loss (Libovický & Helcl, 2018)	1	17.68	19.80	19.93	24.71
NAT w/ Fertility (Gu et al., 2018)	1	17.69	21.47	27.29	29.06
Cross Entropy CMLM (Ghazvininejad et al., 2019)	1	18.05	21.83	27.32	28.20
Auxiliary Regularization (Wang et al., 2019)	1	20.65	24.77	—	—
Bag-of-ngrams Loss (Shao et al., 2019)	1	20.90	24.61	28.31	29.29
Hint-based Training (Li et al., 2019)	1	21.11	25.24	—	—
FlowSeq (Ma et al., 2019)	1	21.45	26.16	29.34	30.44
Bigram CRF (Sun et al., 2019)	1	23.44	27.22	—	—
AXE CMLM (Ours)	1	23.53	27.90	30.75	31.54

Table 3. The performance (test set BLEU) of CMLMs trained with AXE, compared to other non-autoregressive methods. The standard (autoregressive) transformer results are also reported for reference.

the length penalty on the validation set. Similarly we use $\ell = 5$ length candidates for CMLM models, tune the length multiplier ($\lambda \in \{1.05 \dots 1.1\}$),⁷ and the target skipping penalty ($\delta \in \{1 \dots 5\}$) on the validation set.

Knowledge Distillation Similar to previous work on non-autoregressive translation (Gu et al., 2018; Lee et al., 2018; Ghazvininejad et al., 2019; Stern et al., 2019), we use sequence-level knowledge distillation (Kim & Rush, 2016) by training CMLMs on translations generated by a standard left-to-right transformer model (transformer *large* for WMT'14 EN-DE and WMT'17 EN-ZH, transformer *base* for WMT'16 EN-RO). We report the performance of standard autoregressive base transformers trained on distilled data for WMT'14 EN-DE and WMT'17 EN-ZH.

4.2. Main Results

AXE vs Cross Entropy We first compare the performance of AXE-trained CMLMs to that of CMLMs trained with the original cross entropy loss. Table 2 shows that training with AXE substantially increases the performance CMLMs across all benchmarks. On average, we gain 5.2

⁷Our preliminary analysis shows that AXE selects Skip Prediction in 5 – 10% of the time, roughly suggesting that five to ten percent of generated tokens are epsilons. Hence, we search the same range for the length multiplier.

BLEU by replacing cross entropy with AXE, with gains of up to 6.65 BLEU in WMT'17 EN-ZH.

State of the Art We compare the performance of CMLMs with AXE against nine strong baseline models: the fertility-based sequence-to-sequence model (Gu et al., 2018), transformers trained with CTC loss (Libovický & Helcl, 2018), the iterative refinement approach (Lee et al., 2018), transformers trained with auxiliary regularization (Wang et al., 2019), CMLMs trained with (regular) cross entropy loss (Ghazvininejad et al., 2019), Flowseq: a latent variable model based on generative flow (Ma et al., 2019), hint-based training (Li et al., 2019), bag-of-ngrams training (Shao et al., 2019), and the CRF-based semi-autoregressive model (Sun et al., 2019). All of these models except the last one are purely non-autoregressive, while the CRF-based model uses bigram statistics during decoding, which deviates from the purely non-autoregressive setting.⁸

Table 3 shows that our system yields the highest BLEU scores of all non-autoregressive models. AXE-trained CMLMs outperform the best purely non-autoregressive

⁸CMLMs (Ghazvininejad et al., 2019) and the iterative refinement method (Lee et al., 2018) are presented as semi-autoregressive models that run in multiple decoding iterations. However, the first decoding iteration of these models is purely non-autoregressive, which is what we use as our baselines.

Aligned Cross Entropy for Non-Autoregressive Machine Translation

Model	Decoding Iterations	WMT'14		WMT'16	
		EN-DE	DE-EN	EN-RO	RO-EN
Knowledge Distillation					
<i>AXE CMLM</i> (Ours)	1	23.53	27.90	30.75	31.54
Raw Data					
<i>Cross Entropy CMLM</i> (Ghazvininejad et al., 2019)	1	10.64	— —	21.22	— —
<i>CTC Loss</i> (Libovický & Helcl, 2018)	1	17.68	19.80	19.93	24.71
<i>FlowSeq</i> (Ma et al., 2019)	1	18.55	23.36	29.26	30.16
<i>AXE CMLM</i> (Ours)	1	20.40	24.90	30.47	31.42

Table 4. The performance (test set BLEU) of AXE CMLM, compared to other non-autoregressive methods on raw data. The result of AXE CMLM trained with distillation is also reported as a reference.

Training Objective		WMT'14	
Input Tokens	Loss Function	EN-DE	DE-EN
<i>Unobserved</i>	<i>All Tokens</i>	21.97	26.32
<i>Partially-Observed</i>	<i>All Tokens</i>	22.80	27.59
<i>Partially-Observed</i>	<i>Only Masks</i>	23.13	28.01

Table 5. The effect of different training objectives on performance, measured on WMT'14 DE-EN and EN-DE (validation set BLEU).

model (FlowSeq) on both directions of WMT'14 EN-DE and WMT'16 EN-RO by 1.6 BLEU on average. Moreover, our approach achieves higher BLEU scores than the semi-autoregressive CRF decoder across all available benchmarks.

Raw Data Finally, we compare the performance of AXE to other methods that train on raw data *without* knowledge distillation. Table 4 shows that AXE CMLMs still significantly outperform other non-autoregressive models in the raw data scenario. In addition, comparing raw data to knowledge distillation training follows previously-published results that demonstrate the importance of knowledge distillation for non-autoregressive approaches (Gu et al., 2018; Ghazvininejad et al., 2019; Zhou et al., 2019), although the gap is much smaller for WMT'16 EN-RO.

4.3. Ablation Study

In this section, we consider several variations of our proposed method to investigate the effect of each component. We test the performance of AXE CMLMs with these variations on the WMT'14 DE-EN and EN-DE datasets. To prevent overfitting, we evaluate on the validation set using $\ell = 5$ length candidates.

Different Training Objectives Table 5 shows the effects of different training objectives (Section 3.3), in which all or part of the target tokens are masked and the loss function is calculated on all tokens or masked tokens only. We find that simulating the inference scenario, where all tokens are unobserved, is actually less effective than revealing a

Skip Target Penalty	WMT'14 EN-DE		WMT'14 DE-EN	
	BLEU	Skip Target	BLEU	Skip Target
1	22.60	17.57%	26.84	16.87%
2	23.01	10.91%	27.77	10.53%
3	22.85	9.56%	27.87	9.04%
4	22.90	8.14%	28.01	7.83%
5	23.13	7.40%	27.79	6.95%

Table 6. The effect of changing the skip target penalty coefficient δ on performance (BLEU) and the percentage of target words that were skipped, using the validation sets of WMT'14 DE-EN and EN-DE.

subset of the target tokens as input during training. We speculate that partially-observed inputs add easier examples to the training set, allowing for better optimization as in curriculum learning (Bengio et al., 2009). We also see that including only the masked tokens in the loss function gives us a modest but consistent boost in performance, possibly because the training signal is focused on the actual task.

Skip Target Penalty The hyperparameter δ acts as a coefficient for the penalty associated with skipping a target token (see Table 1 for a definition). We experiment with different values of δ , and report our findings in Table 6. We observe that tuning δ can significantly improve performance with respect to the default of $\delta = 1$. As intended, high values of δ discourage alignments that skip too many target tokens.

Length Multiplier The length multiplier λ inflates the length predicted by a CMLM to account for extra blank tokens (ϵ) that the model could potentially generate (see Section 3.2 for more detail). Table 7 compares the effect of different length multiplier λ values. Using the best length multiplier increases the performance by 0.53 BLEU on average for WMT'14 EN-DE and WMT'16 EN-RO.

5. Analysis

We provide a qualitative analysis to provide some insight where AXE improves over cross entropy, and potential direc-

Length Multiplier	WMT'14		WMT'16	
	EN-DE	DE-EN	EN-RO	RO-EN
$\lambda = 1$	22.96	27.50	30.43	32.22
$\lambda = 1.01$	23.06	27.56	30.43	32.25
$\lambda = 1.02$	23.09	27.70	30.66	32.50
$\lambda = 1.03$	23.11	27.81	30.75	32.69
$\lambda = 1.04$	23.13	27.85	30.88	32.83
$\lambda = 1.05$	23.13	27.93	30.88	32.94
$\lambda = 1.06$	23.06	28.01	31.01	32.84
$\lambda = 1.07$	23.09	27.93	31.10	32.61
$\lambda = 1.08$	23.06	27.71	31.14	32.45
$\lambda = 1.09$	23.07	27.68	31.06	32.14
$\lambda = 1.10$	22.92	27.49	30.85	32.01

Table 7. The effect of tuning the length multiplier λ on performance (BLEU), using the validation set.


		Cross Entropy	AXE	
WMT'14	EN-DE	$1 \leq N < 10$	18.75	20.48
		$10 \leq N < 20$	21.69	23.92
		$20 \leq N < 30$	18.64	24.21
		$30 \leq N < 40$	15.37	22.65
		$40 \leq N < 50$	14.04	23.04
	DE-EN	$1 \leq N < 10$	22.57	24.39
		$10 \leq N < 20$	25.28	27.86
		$20 \leq N < 30$	22.43	28.78
		$30 \leq N < 40$	19.03	27.18
		$40 \leq N < 50$	16.16	27.55
	$50 \leq N$	12.23	27.64	

Table 8. The performance (test set BLEU) of cross entropy CMLM and AXE CMLM on WMT'14 EN-DE and DE-EN, bucketed by target sequence length (N).


tions for future research on non-autoregressive generation.

AXE Handles Long Sequences Better We first measure performance of cross entropy versus AXE-trained CMLMs for different sequence lengths. We use `compare-mt` (Neubig et al., 2019) to split the test sets of WMT'14 EN-DE and DE-EN into different buckets based on target sequence length and calculate BLEU for each bucket. Table 8 shows that the performance of models trained with cross entropy drops drastically as the sequence length increases, while the performance of AXE-trained models remains relatively stable. One explanation for this result is that the longer the sequence, the more likely we are to observe misalignments between the model’s predictions and the target; AXE realigns these cases, providing the model with a cleaner signal for modeling long sequences.

AXE Increases Position Confidence We also study how confident each model is about the position of each generated token. Ideally, we would like each predicted token to have a high probability at the position in which it was predicted



(a) Short sequences (less than 10 tokens).



(b) Long sequences (more than 30 tokens).

Figure 3. The average prediction probability assigned to a token as a function of its relative distance from where it was generated in the sequence. Each plot shows the average probabilities for CMLMs trained with cross entropy (dashed blue line) and AXE (solid red line).

and a very low probability in the neighboring positions. After applying argmax decoding, we compute the probability assigned to each generated token in all positions of the sequence and average these probabilities based on the relative distance (positive or negative) to the generated position. Figure 3 plots these averaged probabilities for both short (< 10 tokens) and long (> 30 tokens) target sequences.

Both models are rather confident in their predictions for short sequences (Figure 3(a)): the probability has a high peak at the generated position and drops rapidly as we move further away. However, for longer sentences (Figure 3(b)), we observe that the plot for cross entropy has lost its sharpness. Specifically, the immediate neighbors of the prediction position (± 1) receive about 0.14 probability on average, almost a third of the peak probability. Meanwhile, the probabilities predicted by the AXE-trained model are significantly sharper, assigning negligible probabilities to the generated token in neighboring positions when compared to the center.

One way to explain this result is that cross entropy training encourages predictions to have some probability mass of their neighbors, in order to “hedge their bets” in case the predictions are misaligned with the target. Since AXE finds

Model	WMT'14	
	EN-DE	DE-EN
Cross Entropy CMLM	16.72%	12.31%
AXE CMLM	1.41%	1.03%

Table 9. The percentage of repeated tokens on the test sets of WMT'14 EN-DE and DE-EN.

the best alignment before computing the actual loss, spreading the probability mass of a token among its neighbors is no longer necessary.

AXE Reduces Multimodality We further argue that AXE reduces the multimodality problem in non-autoregressive machine translation (Gu et al., 2018). Due to minimal coordination between predictions in many non-autoregressive models, a model might consider many possible translations at the same time. In this situation, the model might merge two or more different translations and generate an inconsistent output that is typically characterized by token repetitions. We therefore use the frequency of repeated tokens as a proxy for measuring multimodality in a model.

Table 9 shows the repetition rate for cross entropy and AXE-trained CMLMs. Replacing cross entropy with AXE drastically reduces multimodality, decreasing the number of repetitions by a multiplicative factor of 12.

6. Related Work

Advances in neural machine translation techniques in recent years has brought an increasing interest in breaking the autoregressive generation bottleneck in translation models.

Semi-autoregressive models introduce partial parallelism into the decoding process. Some of these techniques include iterative refinement of translations based on previous predictions (Lee et al., 2018; Ghazvininejad et al., 2019; 2020; Gu et al., 2019; Kasai et al., 2020) and combining a lighter autoregressive decoder with a non-autoregressive one (Sun et al., 2019).

Building a fully non-autoregressive machine translation model is a much more challenging task. One branch of prior work approaches this problem by modeling with latent variables. Gu et al. (2018) introduces word fertility as a latent variable to model the number of generated tokens per each source word. Ma et al. (2019) uses generative flow to model complex distribution of latent variables for parallel decoding of target. Shu et al. (2019) proposes a latent-variable non-autoregressive model with continuous latent variables and a deterministic inference procedure.

There is also work that develops other alternative loss functions for non-autoregressive machine translation. Libovický & Helcl (2018) use the Connectionist Temporal Classifi-

cation (CTC) training objective, a loss function from the speech recognition literature that is designed to eliminating repetitions. Concurrent work by Saharia et al. (2020) shows that a different application of CTC can achieve significantly stronger results, particularly when trained with knowledge distillation. Li et al. (2019) uses the learning signal provided by hidden states and attention distributions of an autoregressive teacher. Yang et al. (2019) improves the decoder hidden representations by adding the reconstruction error of source sentence from these representations as an auxiliary regularization term to the loss function. Finally, Shao et al. (2019) introduce the bag-of-ngrams training objective to encourage the model to capture target-side sequential dependencies.

7. Conclusion

We introduced Aligned Cross Entropy (AXE) as an alternative loss function for training non-autoregressive models. AXE focuses on relative order and lexical matching instead of relying on absolute positions. We showed that, in the context of machine translation, a conditional masked language model (CMLM) trained with AXE significantly outperforms cross entropy trained models, setting a new state-of-the-art for non-autoregressive models.

Acknowledgements

We thank Abdelrahman Mohamed for sharing his expertise on non-autoregressive models, and our colleagues at FAIR for valuable feedback.

References

- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 41–48, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Ghazvininejad, M., Levy, O., Liu, Y., and Zettlemoyer, L. Mask-predict: Parallel decoding of conditional masked language models. In *Proc. of EMNLP-IJCNLP*, 2019. URL <https://www.aclweb.org/anthology/D19-1633>.
- Ghazvininejad, M., Levy, O., and Zettlemoyer, L. Semi-autoregressive training improves mask-predict decoding. *arXiv preprint arXiv:2001.08785*, 2020.
- Gu, J., Bradbury, J., Xiong, C., Li, V. O., and Socher, R. Non-autoregressive neural machine translation. In *Proc. of ICLR*, 2018.

- Gu, J., Wang, C., and Zhao, J. Levenshtein transformer. In *Proc. of NeurIPS*, 2019. URL <https://arxiv.org/abs/1905.11006>.
- Kasai, J., Cross, J., Ghazvininejad, M., and Gu, J. Parallel machine translation with disentangled context transformer. *arXiv preprint arXiv:2001.05136*, 2020.
- Kim, Y. and Rush, A. M. Sequence-level knowledge distillation. In *Proc. of EMNLP*, 2016. URL <https://arxiv.org/abs/1606.07947>.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference for Learning Representations*, 2015.
- Lee, J. D., Mansimov, E., and Cho, K. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *Proc. of EMNLP*, 2018. URL <https://arxiv.org/abs/1802.06901>.
- Li, Z., Lin, Z., He, D., Tian, F., Qin, T., Wang, L., and Liu, T.-Y. Hint-based training for non-autoregressive machine translation. *arXiv preprint arXiv:1909.06708*, 2019.
- Libovický, J. and Helcl, J. End-to-end non-autoregressive neural machine translation with connectionist temporal classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3016–3021, Brussels, Belgium, 2018. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/D18-1336>.
- Ma, X., Zhou, C., Li, X., Neubig, G., and Hovy, E. Flowseq: Non-autoregressive conditional sequence generation with generative flow. *arXiv preprint arXiv:1909.02480*, 2019.
- Neubig, G., Dou, Z.-Y., Hu, J., Michel, P., Pruthi, D., and Wang, X. compare-mt: A tool for holistic comparison of language generation systems. In *Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL) Demo Track*, Minneapolis, USA, June 2019. URL <http://arxiv.org/abs/1903.07926>.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, July 2002. URL <https://www.aclweb.org/anthology/P02-1040>.
- Post, M. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*. Association for Computational Linguistics, October 2018. URL <https://www.aclweb.org/anthology/W18-6319>.
- Saharia, C., Chan, W., Saxena, S., and Norouzi, M. Non-autoregressive machine translation with latent alignments. *arXiv preprint arXiv:2004.07437*, 2020.
- Sakoe, H. and Chiba, S. Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing*, 26(1): 43–49, 1978.
- 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 (Volume 1: Long Papers)*, August 2016. URL <https://www.aclweb.org/anthology/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. *arXiv preprint arXiv:1911.09320*, 2019.
- Shu, R., Lee, J., Nakayama, H., and Cho, K. Latent-variable non-autoregressive neural machine translation with deterministic inference using a delta posterior. *arXiv preprint arXiv:1908.07181*, 2019.
- Stern, M., Chan, W., Kiros, J. R., and Uszkoreit, J. Insertion transformer: Flexible sequence generation via insertion operations. In *Proc. of ICML*, 2019. URL <https://arxiv.org/abs/1902.03249>.
- Sun, Z., Li, Z., Wang, H., He, D., Lin, Z., and Deng, Z. Fast structured decoding for sequence models. In *Advances in Neural Information Processing Systems*, pp. 3011–3020, 2019.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- Wang, Y., Tian, F., He, D., Qin, T., Zhai, C., and Liu, T.-Y. Non-autoregressive machine translation with auxiliary regularization. In *Proc. of AAAI*, 2019. URL <https://arxiv.org/abs/1902.10245>.
- Wu, F., Fan, A., Baevski, A., Dauphin, Y. N., and Auli, M. Pay less attention with lightweight and dynamic convolutions. *International Conference on Learning Representations*, 2019.
- Yang, B., Liu, F., and Zou, Y. Non-autoregressive video captioning with iterative refinement. *arXiv preprint arXiv:1911.12018*, 2019.
- Zhou, C., Neubig, G., and Gu, J. Understanding knowledge distillation in non-autoregressive machine translation. *arXiv preprint arXiv:1911.02727*, 2019.