

# Mixed Cross Entropy Loss for Neural Machine Translation

Haoran Li<sup>\* 1</sup> Wei Lu<sup>\* 1</sup>

## Abstract

In neural machine translation, cross entropy (CE) is the standard loss function in two training methods of auto-regressive models, i.e., *teacher forcing* and *scheduled sampling*. In this paper, we propose *mixed cross entropy loss* (mixed CE) as a substitute for CE in both training approaches. In teacher forcing, the model trained with CE regards the translation problem as a one-to-one mapping process, while in mixed CE this process can be relaxed to one-to-many. In scheduled sampling, we show that mixed CE has the potential to encourage the training and testing behaviours to be similar to each other, more effectively mitigating the *exposure bias* problem. We demonstrate the superiority of mixed CE over CE on several machine translation datasets, WMT’16 Ro-En, WMT’16 Ru-En, and WMT’14 En-De in both teacher forcing and scheduled sampling setups. Furthermore, in WMT’14 En-De, we also find mixed CE consistently outperforms CE on a multi-reference set as well as a challenging paraphrased reference set. We also found the model trained with mixed CE is able to provide a better probability distribution defined over the translation output space. Our code is available at <https://github.com/haorannlp/mix>.

## 1. Introduction

Conditional language generation tasks, e.g., machine translation, text summarization, are all text-to-text problems. The most popular models to solve these tasks include RNNs (Elman, 1990; Hochreiter & Schmidhuber, 1997; Cho et al., 2014), Transformer (Vaswani et al., 2017), etc., which are usually arranged in an encoder-decoder architecture (Sutskever et al., 2014; Cho et al., 2014). Given a training example  $(x, y) \in \mathcal{D}$ , the encoder first compresses the source sequence  $x = (x_1, \dots, x_m)$  into a vector  $h$  and then

<sup>\*</sup>Equal contribution <sup>1</sup>StatNLP Research Group, Singapore University of Technology and Design, Singapore. Correspondence to: Wei Lu <luwei@sutd.edu.sg>.




Figure 1. Example of machine translation’s one-to-many property. “Source” and “Target” belong to the training set while “Other Translations” denotes other plausible translations which are not present in the training data.

the decoder will produce a target sequence  $\hat{y} = (\hat{y}_1, \dots, \hat{y}_n)$  from  $h$ <sup>1</sup>. However, due to information loss in the compression process and the limited expressiveness power of the model, it can be hard for the decoder to recover a good target sequence  $\hat{y}$  from  $h$  alone. To solve this problem, people often use *teacher forcing* (Williams & Zipser, 1989) where both  $h$  and the gold target  $y = (y_1, \dots, y_n)$  are fed into the decoder during training. Next, the aim is typically to minimize the *cross entropy* (CE) loss, which can be written as  $-\sum_{t=1}^n \log p_{\theta}(y_t|y_{<t}, x)$ .<sup>2</sup> The general idea is that, by optimizing the CE, we hope the output probability distribution  $p_{\theta}(\cdot|y_{<t}, x) \in \mathbb{R}^{|V|}$  ( $|V|$  is the vocabulary size) can approximate the target one-hot encoding of  $y_t$  (a vector with only one position being 1, others being 0). In practice, empirical success of various neural machine translation (NMT) models trained with CE (Bahdanau et al., 2015; Vaswani et al., 2017) has demonstrated CE’s effectiveness.

However, it is worth noting that NMT is inherently a *one-to-many* mapping problem where a source sentence has multiple plausible translations. In Fig. 1, suppose we want to predict the last word  $y_6$  “decline” conditioned on the source sentence and the prefix  $y_{<6}$ . Although “decline” is the gold target that our model should assign the most probability mass to, other synonyms of “decline” such as “drop”, “decrease” are also plausible translations in this context, namely the corresponding values of these synonyms in the vector

<sup>1</sup>With potential reference to  $x$ , e.g., through the attention mechanism. We ignore this for brevity in our discussion.

<sup>2</sup>Here  $y_{<t} = (y_0, y_1, \dots, y_{t-1})$  is the partial target sequence, where  $y_0$  is the special start token, and  $\theta$  denotes model parameters.

representation of  $y_6$  should be non-trivial. Ignoring these synonyms and simply fitting the one-hot encoding may limit the model’s generalization ability because  $p_\theta(\cdot|y_{<t}, x)$  deviates from the ground truth  $p^*(\cdot|y_{<t}, x)$  where the test data is drawn (Norouzi et al., 2016; Szegedy et al., 2016; Xiao et al., 2019). Nevertheless, there is no doubt that when we train our model with CE, the predictions  $p_\theta(\cdot|y_{<t}, x)$  still contain useful information about the ground truth  $p^*(\cdot|y_{<t}, x)$ . In this paper, we make a simple assumption: given a well-trained model, if the predicted token (the one with the largest probability given by the model) does not match the ground truth, this token is very likely to be a synonym or part of a synonym of the ground truth. This assumption is supported by the fact that in a typical parallel corpus, the same source word can have multiple translations in different training instances. With a comparable number of occurrences of different translations given the same source word, the model tends to evenly allocate the probability mass to them after learning, thereby having a chance to predict a synonym of the gold token as the most probable one during decoding. Based on this assumption, in teacher forcing, we use mixed CE to incrementally exploit the information regarding synonyms so as to boost model’s generalization capability, which can be better demonstrated in a multi-reference test set (see Section 4.2).

Despite its simplicity, teacher forcing suffers from *exposure bias* (Ranzato et al., 2016), which refers to the discrepancy that during training time, gold target  $y$  is observable by the decoder whereas at test time  $y$  is unknown. Thus, at test time, the model has to sample from its own predictions to decode auto-regressively. The major solution to exposure bias is to train the model on its own predictions at training stage such that the mismatch between training/testing input distributions can be reduced (Daumé et al., 2009; Ross et al., 2011; Venkatraman et al., 2015; Bengio et al., 2015; Ranzato et al., 2016; Bahdanau et al., 2017; Leblond et al., 2018; Zhang et al., 2019). One of the those methods is *scheduled sampling* (Bengio et al., 2015) (see Section 2), which mixes gold input<sup>3</sup> with model-generated input (model predictions) and then maps this mixed input to the gold target.

We argue that besides making changes to the input distribution during the training time, so that it simulates the *test-time inputs*, we can also mitigate exposure bias by simulating the *test-time behaviour*, making the model more robust to the difference between training/testing inputs. That is, we should design our model in such a way that it always produces similar results no matter whether the input is gold or model-generated. In scheduled sampling, this can be done with mixed CE, which maps the mixed input not only to the gold target, but also to the output which is obtained from

<sup>3</sup>The input here refers to *decoder input*. Unless otherwise specified, the *encoder input* remains the same.




Figure 2. Scheduled sampling for Transformer. We apply argmax to the output logits, or equivalently the log likelihood  $\log p_\theta(\cdot|y_{<t}, x)$ , at each step  $t$  to obtain  $\hat{y}$  in the first pass. *eos* is the end-of-sequence symbol.

the gold input.

In this paper, we propose mixed CE to substitute CE in both teacher forcing and scheduled sampling training of NMT models. In teacher forcing, mixed CE guides the model to learn a one-to-many mapping by incrementally exploiting useful synonym information in the model predictions. In scheduled sampling, mixed CE encourages the model to output similar results regardless of if the model is fed with the gold input or the mixed input which consists of gold target tokens and model predictions. We demonstrate the effectiveness of mixed CE not only on the standard test set of WMT’16 Ro-En, WMT’16 Ru-En, WMT’14 En-De but also on a multi-reference set (Ott et al., 2018) as well as a challenging paraphrased reference set (Freitag et al., 2020).

## 2. Background

In this section, we review some background knowledge about *exposure bias* and *scheduled sampling*. In teacher forcing, the model uses the gold target  $y$  as input during training but  $y$  is not available at test time. Hence, at test time, the decoder has to sample from its own predictions  $\hat{y}_t$  in the last time step  $t - 1$  and feed the sampled token as input to the current step  $t$ . Note that the model has only been trained on the empirical data distribution instead of its own predictions and this may lead to arbitrary decoding decisions at test time. To alleviate exposure bias, Bengio et al. (2015) proposed scheduled sampling for RNN training. In scheduled sampling, the model randomly decides to use gold data  $y_t$  with probability  $\epsilon \in [0, 1]$  or the model’s own prediction  $\hat{y}_t$  (the one with the largest probability) with probability  $1 - \epsilon$  as input at each time step  $t$ .  $\epsilon$  is initialized to a large value close to 1 and then decays as training proceeds (see Section 4.1), which assures the model can converge better and faster. Consequently, the model is exposed to part of its predictions during training and this reduces the risk of having a bad generation at test time.

For Transformer (Vaswani et al., 2017), scheduled sampling

is typically done in a bit different way because we want to avoid sequential decoding and make full use of Transformer decoder’s parallel computation mechanism (masked self-attention) during training. It has been proposed and widely adopted by the community (Mihaylova & Martins, 2019; Zhang et al., 2019; Duckworth et al., 2019) that we could run the Transformer once without accumulating the gradients and store the greedily predicted sequence  $\hat{y}$  for the second pass (see Fig. 2). Then we randomly replace each token in the gold target sequence  $y$  with the corresponding token in the predicted sequence  $\hat{y}$  with probability  $1 - \epsilon$ , obtaining a mixed sequence  $y^{\text{mix}}$ . Next we feed  $y^{\text{mix}}$  to the Transformer decoder again and compute the loss:

$$\mathcal{L}_{CE} = - \sum_{t=1}^n \log p_{\theta}(y_t | y_{<t}, x). \quad (1)$$

Even though the model involves two forward passes (i.e., perform decoding twice), it is still faster than sequential decoding. Zhang et al. (2019) proposed a variant of scheduled sampling called *word oracle*, which adds some Gumbel noise  $g_t \in \mathbb{R}^{|V|}$  to the log likelihood produced in the first forward pass, i.e.,  $s_t = \log p_{\theta}(\cdot | y_{<t}, x) + g_t$ , before taking  $\text{argmax}$ . All the elements in  $g_t$  are i.i.d. samples drawn from  $\text{Gumbel}(0, 1)$ :  $g_{t,i} = -\log(-\log(u))$ ,  $u \sim \text{Uniform}(0, 1)$  (Gumbel, 1954; Jang et al., 2017). The  $\text{argmax}$  results of  $\log p_{\theta}(\cdot | y_{<t}, x)$  and  $s_t$  may be different and thus enable the decoder to observe a wider variety of input combinations  $y^{\text{mix}}$ .

In this paper, we only study mixed CE in teacher forcing and scheduled sampling training under the state-of-the-art Transformer framework due to its wide adoption and empirical success (Vaswani et al., 2017).

### 3. Approach

In this section, we start with the formulation of mixed cross entropy (Mixed CE) and argue its superiority over CE in teacher forcing training and scheduled sampling training.

#### 3.1. Mixed Cross Entropy (Mixed CE)

Given a training instance  $(x, y)$ , in teacher forcing, mixed CE can be written as:

$$\begin{aligned} \mathcal{L}_{\text{mix}} = & - \left[ (1 - \alpha_i) \cdot \sum_{t=1}^n \log p_{\theta}(y_t | y_{<t}, x) \right. \\ & \left. + \alpha_i \cdot \sum_{t=1}^n \log p_{\theta}(\hat{y}_t | y_{<t}, x) \right] \end{aligned} \quad (2)$$

where  $\hat{y}_t = \text{argmax}_{1 \leq k \leq |V|} \log p_{\theta}(w_k | y_{<t}, x)$ <sup>4</sup>

<sup>4</sup>Here with slight abuse of notation,  $y_t, \hat{y}_t$  can be used to denote the token or the token’s index in the vocabulary.

In scheduled sampling, mixed CE is a bit different since the model needs to forward twice (see Section 2):

$$\begin{aligned} \mathcal{L}_{\text{mix}} = & - \left[ (1 - \alpha_i) \cdot \sum_{t=1}^n \log p_{\theta}(y_t | y_{<t}^{\text{mix}}, x) \right. \\ & \left. + \alpha_i \cdot \sum_{t=1}^n \log p_{\theta}(\hat{y}_t | y_{<t}^{\text{mix}}, x) \right]. \end{aligned} \quad (3)$$

Here,  $\hat{y}_t$  is still  $\text{argmax}_{1 \leq k \leq |V|} \log p_{\theta}(w_k | y_{<t}, x)$  obtained in the first forward pass.  $y^{\text{mix}}$  consists of gold tokens in  $y$  and the greedily-generated tokens in  $\hat{y}$ , as discussed in the previous section.  $\alpha_i$  is a scalar that is related to the  $i$ -th training iteration, which can be computed as follows:

$$\alpha_i = m \cdot \frac{i}{\text{total\_iter}}, 1 \leq i \leq \text{total\_iter} \quad (4)$$

where “total\_iter” denotes the total training iterations and  $m = 0.5$  such that the max value of  $\alpha_i$  is limited to 0.5.

The first part of Eq. 2 and 3 is the standard CE in both training methods, while the second part is model-dependent CE. When  $\hat{y}_t = y_t$ , mixed CE degenerates to CE, and when  $\hat{y}_t \neq y_t$ , we justify mixed CE in teacher forcing and scheduled sampling separately. Specifically, in teacher forcing, we relate machine translation to a noisy label problem, while in scheduled sampling, we argue mixed CE can make the model less sensitive to input variations, by better simulating test-time behaviours.

#### 3.2. Mixed CE in Teacher Forcing: Machine Translation as a Noisy Label Problem

Noisy labels are unreliable labels that are corrupted from the ground truth (Song et al., 2020). In a one-hot representation, a noisy label places “1” on a different index from the ground truth. This problem often occurs in classification tasks due to untrue annotations at the data collection stage. In Section 1, we have argued that in CE,  $p_{\theta}(\cdot | y_{<t}, x)$  approximating an inappropriate one-hot encoding instead of the underlying ground truth  $p^*(\cdot | y_{<t}, x)$  might impede generalization. Ideally, the (soft) label of each token in a sentence should be provided by  $p^*(\cdot | y_{<t}, x)$  which is context-dependent and the probability mass is scattered over different tokens, especially synonyms. In that sense, all the  $y_t$  (one-hot encoding) in the training set can be deemed as “noisy” labels. Thus, we should not fully trust such noisy labels. On the other hand, the empirical success of using noisy labels as ground truth (Bahdanau et al., 2015; Vaswani et al., 2017) suggests that the learned model still preserves information about  $p^*(\cdot | y_{<t}, x)$ . That is to say,  $p_{\theta}(\cdot | y_{<t}, x)$  is closer to  $p^*(\cdot | y_{<t}, x)$  than the one-hot encoding of  $y_t$ . So it is reasonable to exploit  $p_{\theta}(\cdot | y_{<t}, x)$  during training. But how much information in  $p_{\theta}(\cdot | y_{<t}, x)$  can be used?

Previously, we assume that if the model is well-trained and  $\hat{y}_t \neq y_t$ ,  $\hat{y}_t$  is very likely to be a synonym of  $y_t$ . Under this assumption,  $\hat{y}_t$  shall be the useful information provided by  $p_\theta(\cdot|y_{<t}, x)$  that can be exploited. Therefore, mixed CE also chooses to maximize  $\hat{y}_t$ 's probability with a dynamic weight  $\alpha_i$  besides the gold target  $y_t$ . It is worth noting that in Eq. 4,  $1 - \alpha_i \geq \alpha_i$ . This is because we want our model to be trained on gold data more in the early stage when the model is not yet informative. As training progresses, we gradually shift more weights to the model's predictions and finally we treat the gold target  $y_t$  and the model's predictions  $\hat{y}_t$  equally.

The above learning process essentially assumes there exists a collection of plausible target translations assigned to each input sentence to be translated, where each such target translation serve as a "label". This is essentially a learning problem involving soft or noisy labels. This label set collapse to a trivial set with only one element, which is the gold target sentence, in the case of standard CE.

Our mixed CE formulation may be reminiscent of other approaches that also adopt soft labels, such as the label smoothing approach (Szegedy et al., 2016). Label smoothing assigns  $1 - \gamma$  ( $0 < \gamma < 1$ ) probability mass to the gold token, with the remaining  $\gamma$  uniformly distributed among all the tokens in the vocabulary, i.e.,  $y_t^{LS} = (1 - \gamma) \cdot y_t + \frac{\gamma}{|V|} \cdot u$ , where  $y_t$  is the one-hot encoding and  $u$  is a discrete uniform distribution over the whole vocabulary. In Section 4, we empirically show that the models trained with mixed CE and label smoothing have totally different impacts on the output probability distribution.

### 3.3. Mixed CE in Scheduled Sampling: Better Simulation of Test-time Behaviours

Now let us focus on the learning objective for mixed CE in the presence of scheduled sampling. In Eq. 3, the first part of mixed CE is the objective commonly used in scheduled sampling (i.e., Eq. 1), which encourages the model to map the mixed input  $y^{\text{mix}}$  to the gold target  $y$ . The motivation underlying this is to encourage the model to generate a good output sequence even with its own predictions as input, which is approximated with  $y^{\text{mix}}$ . This is essentially a process that is simulating the test-time inputs.

The second part of the mixed CE objective, however, encourages the output conditioned on  $y^{\text{mix}}$  to approximate the greedily-generated sequence  $\hat{y}$ . How do we understand such an objective? Note that from Fig. 2, we can see that  $\hat{y}$  is in turn the output from the first decoder, which is also parameterized by  $\theta$ . In other words,  $\hat{y}_t = \arg\max_{1 \leq k \leq |V|} \log p_\theta(w_k|y_{<t}, x)$ . Putting things together, this means with this objective, we are essentially encouraging the model, parameterized by  $\theta$ , to produce the same output ( $\hat{y}$ ) regardless of whether the input is the gold

input  $y$  or the mixed input  $y^{\text{mix}}$ . In other words, we are not simply interested in simulating the test-time inputs now, but we would also like to make  $y^{\text{mix}}$  and  $y$  indistinguishable when serving as inputs to the model. This effectively requires the model to share the same internal *behaviour* (e.g., similar internal neural states) when the input is either  $y$  or  $y^{\text{mix}}$ , while the former is related to the training phase, and the latter is related to the testing phase.

Overall, our mixed CE approach is essentially simulating the test-time behaviour (with the second term), while encouraging the model to learn to predict the gold output  $y$  with the simulated test-time input  $y^{\text{mix}}$  (with the first term).

## 4. Experiments

In this section, we conducted experiments to verify the effectiveness of mixed CE in teacher forcing and scheduled sampling on several benchmark datasets with different sizes, WMT'16 Romanian-English (Ro-En, 610K pairs), WMT'16 Russian-English (Ru-En, 2.1M pairs) and WMT'14 English-German (En-De, 4.5M pairs). We begin with some training details of all experiments and then we study mixed CE in teacher forcing and scheduled sampling separately.

### 4.1. Experimental Setup

We used the preprocessed WMT'16 Ro-En dataset from Lee et al. (2018) with vocabulary sizes of Romanian and English being 27,591 and 21,175, respectively. For WMT'14 En-De, we used the script from Fairseq (Ott et al., 2019)<sup>5</sup> for preprocessing (we used newstest2013 as the validation set instead following Zhang et al. (2019)). The vocabulary size of English is 40,511 while for German it is 42,735. We used a script similar to WMT'14 En-De to preprocess WMT'16 Ru-En but separate BPE codes (Sennrich et al., 2016b) for Russian and English with 24K merge operations, resulting in 26,327 and 26,319 tokens in each vocabulary were adopted. We used separate token embeddings for source and target languages in all 3 datasets. Standard base Transformer (Vaswani et al., 2017) architecture was used in the experiment. We trained the model for totally 8,000/45,000/80,000 iterations for Ro-En/Ru-En/De-En datasets with each batch containing  $12,288 \times 4 / 12,288 \times 4 / 12,288 \times 8$  tokens. All models were pre-trained with CE for 5 epochs. We used the Adam (Kingma & Ba, 2015) optimizer with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.98$ . Learning rate is 0.0007 and will be reduced by half if the BLEU (Papineni et al., 2002) score on validation set does not increase in the last 4 epochs. Unless otherwise specified, we also used label smoothing ( $\gamma = 0.1$ ) in our experiments. The decay strategy for scheduled sampling (see Section 2) that we use follows Leblond et al. (2018):

$$\epsilon_i = d^{i/\text{total.iter}}, 0 < d < 1, 1 \leq i \leq \text{total.iter} \quad (5)$$

<sup>5</sup><https://github.com/pytorch/fairseq>

where  $\epsilon_i$  is related to  $i$ -th training iteration and decays as training proceeds. For Ro-En,  $d = 0.7$  while for Ru-En/En-De,  $d = 0.8$ . We saved a checkpoint after training the model for each epoch and we selected the best checkpoint based on the performance on the validation set, which we refer to as “Single”. Besides, we also reported the performance of the average models obtained by averaging either the last 5 checkpoints or top 5 checkpoints, depending on the performance on the validation set. We refer to this average model as “Average”.

## 4.2. Mixed CE in Teacher Forcing

We use a base Transformer trained with CE as a baseline and compare it with mixed CE. We also compare mixed CE with a loss function that is originally designed for neural machine translation, Dual Skew Divergence (DSD) loss (Xiao et al., 2019), which minimizes the forward and reverse  $\alpha$ -skew KL divergence between empirical data distribution and model prediction. The best-performing hyperparameters in the original DSD paper were used in our experiments. Besides, according to Xiao et al. (2019), DSD only works after the model has been pre-trained with CE, otherwise the performance would drop quickly. The number of pre-training iterations was chosen based on the performance on the validation set. Moreover, we also compare mixed CE with a sequence level self-distillation method (self-dist) (Kim & Rush, 2016; Furlanello et al., 2018).<sup>6</sup> We first train a Transformer with CE and re-generate the target sequence  $y$  in the training data using greedy search (we choose greedy search because it corresponds to the argmax operation in Eq. 2), obtaining new  $\hat{y}$ . Note that during greedy search, we force the length of  $\hat{y}$  to be the same as  $y$ . Then we use the same loss function as Eq. 2 but replace the original  $\hat{y}$  with our distilled  $\hat{y}$  to train the model. Results are shown in Table 1. All results are averaged over 3 runs. Self-dist does not perform well and we conjecture this is because the model has difficulty fitting two different modes at the same time. Mixed CE outperforms CE in most settings even though the performance gap is smaller on larger data sets. Compared to DSD, mixed CE generally gives better performance and mixed CE is much easier to train. It can be observed that mixed CE typically brings larger improvement in single model testing. In Ru-En, the single model trained with mixed CE even approaches the average model trained with CE. Besides, the improvements on average models seem to be more significant when we apply greedy decoding (beam size = 1). All the test sets used so far are single-reference sets, which may

<sup>6</sup>There is another approach is also called self-knowledge distillation (Hahn & Choi, 2019) where the  $\alpha$  is computed from the scaled Euclidean distance between the embeddings of the gold token and the model’s prediction. We tried several scale factors but failed to find one that is able to outperform the baseline in our training settings. We discuss this in the appendix.

Table 1. BLEU scores on test sets of Transformers trained with CE and mixed CE. The results of beam search decoding with beam size 1/5 are presented. All results are averaged over 3 runs.

DATA SET	LOSS	SINGLE	AVERAGE
RO-EN	CE	30.63/31.42	32.07/32.59
	DSD	<b>31.17</b> /31.80	32.03/32.74
	SELF-DIST	28.65/31.45	31.66/32.61
	MIXED CE	<b>31.17</b> / <b>32.02</b>	<b>32.63</b> / <b>33.25</b>
RU-EN	CE	28.87/30.24	29.48/30.79
	DSD	28.89/30.30	29.69/30.90
	SELF-DIST	28.76/30.34	29.32/30.63
	MIXED CE	<b>29.59</b> / <b>30.74</b>	<b>30.14</b> / <b>31.05</b>
EN-DE	CE	26.23/26.91	26.67/27.41
	DSD	26.10/26.84	26.66/27.30
	SELF-DIST	24.15/25.98	24.23/25.91
	MIXED CE	<b>26.32</b> / <b>27.28</b>	<b>26.72</b> / <b>27.61</b>

not demonstrate mixed CE’s ability to exploit synonyms. So we further experiment with a multi-reference set as well as a paraphrased single-reference set.

**Additional Reference** We also conducted experiments to see how CE-based, mixed CE-based models perform when we use a new set of references. We chose two additional reference sets: 1) a multiple-reference test set of WMT’14 En-De (Ott et al., 2018)<sup>7</sup> where there are 10 additional references (the original reference is excluded) for each of the 500 test sentences taken from the original test set; 2) a paraphrased as-much-as-possible version of the original WMT’19 En-De reference (Freitag et al., 2020)<sup>8</sup>. In the first set, 10 human reference translations of the same source sentence cover a broader reference space and exhibit certain amount of diversity in lexical choices. This will help us verify the effectiveness of mixed CE since it is expected to maximize the probability of synonyms of training tokens. Nevertheless, these human translations are often influenced by source sentences and thus tend to have a monotonic alignment to the source side and a relatively simple vocabulary (Koppel & Ordan, 2011; Freitag et al., 2020). Therefore, we further selected the second set in which each reference translation is paraphrased from the original reference by human experts and differs significantly from the original one in word choices (more advanced synonyms) and sentence structures (non-monotonic alignment) but keeps the same meaning. According to Freitag et al. (2020), BLEU scores on this paraphrased test set correlates better with human ratings than the original reference and hence we should be more confident about the superiority of mixed CE if it yields

<sup>7</sup><https://github.com/facebookresearch/analyzing-uncertainty-nmt>

<sup>8</sup><https://github.com/google/wmt19-paraphrased-references>

Table 2. BLEU improvement of mixed CE over CE on 10 additional references of WMT’14 En-De test set. All results are averaged over 3 runs.

REF	AVG		TOP	
	CE	MIXED CE	CE	MIXED CE
REF 1	36.73	<b>37.32 (+0.59)</b>	38.61	<b>39.13 (+0.52)</b>
REF 2	47.48	<b>48.50 (+1.02)</b>	50.08	<b>51.36 (+1.28)</b>
REF 3	42.59	<b>43.25 (+0.66)</b>	44.89	<b>45.89 (+1.00)</b>
REF 4	28.93	<b>29.78 (+0.85)</b>	30.29	<b>30.98 (+0.69)</b>
REF 5	31.75	<b>32.53 (+0.78)</b>	33.48	<b>34.18 (+0.70)</b>
REF 6	26.41	<b>26.83 (+0.42)</b>	27.60	<b>27.96 (+0.36)</b>
REF 7	42.18	<b>42.89 (+0.71)</b>	44.37	<b>44.90 (+0.53)</b>
REF 8	32.36	<b>33.05 (+0.69)</b>	33.77	<b>34.55 (+0.78)</b>
REF 9	28.51	<b>29.03 (+0.52)</b>	29.65	<b>30.27 (+0.62)</b>
REF 10	33.75	<b>33.94 (+0.19)</b>	35.23	<b>35.68 (+0.45)</b>
MEAN	35.07	<b>35.71 (+0.64)</b>	36.80	<b>37.49 (+0.69)</b>

higher BLEU. Note that when evaluated on the WMT’19 En-De test set, the models were still trained with the WMT’14 training data.

We listed the BLEU score improvement of the model trained with mixed CE over the model with CE on 10 additional references in Table 2. We used beam search (beam size 10) to generate 10 hypotheses for each source sentence. We reported the average (AVG) and the largest (TOP) BLEU scores of the 10 hypotheses with respect to each reference. Mixed CE consistently outperforms CE across all additional references with the average improvement being 0.64/0.69 BLEU. This suggests that mixed CE can assist in generating hypotheses that align better with human translations in a general sense.

The results for the second reference set is shown in Table 3. We used beam search (beam size 1/10) and sampling (sampled 100 times and selected the sentence with the highest average log likelihood) to generate hypotheses. As a reference, the BLEU score of the machine translation system trained with WMT’19 En-De training data (it has much more sentence pairs than WMT’14 En-De, 38.8M vs 4.5M) on this paraphrased test set is 12.5 (Freitag et al., 2020). In addition, it is worth to note that BLEU score on this reference set is much lower than on the original one due to less  $n$ -gram matches. Despite a significant change in reference, mixed CE still surpasses CE (+0.34/+0.27/+1.01 BLEU), especially when using sampling decoding. These results on the two additional references again confirm that mixed CE exploits useful information in the model predictions.

Last, we give a concrete example of English-German translation produced by CE-trained and mixed CE-trained models in Fig. 3. We plotted each token’s  $\log_2$  score given by the models. For the synonyms “*sinkt*” and “*geht zurück*” (both mean “*decline*”), CE-trained model gives them scores  $-1.5$  and  $-2.06$  while mixed CE-trained model produces  $-1.4$

English	The birth rate continues to decline .						
German 1.	Die Geburtenrate geht weiter zurück .						
CE	-0.29	-0.56	-1.46	-0.58	-0.04	-0.15	(-3.08)
Mixed CE	-0.22	-0.53	-1.35	-0.51	-0.05	-0.14	(-2.80)
German 2.	Die Geburtenrate sinkt weiter .						
CE	-0.29	-0.56	-2.06	-1.49	-0.21		(-4.61)
Mixed CE	-0.22	-0.53	-1.48	-1.17	-0.17		(-3.57)

Figure 3. Multiple translations with their positional scores. We used beam search to translate the English sentence into German and we selected the top-2 candidates. Words in the same color have the same meaning.

Table 3. BLEU scores of beam search/sampling results on WMT’19 En-De paraphrased test set. As a reference, Freitag et al. (2020) reported that the BLEU score improvement of the machine translation system augmented with Automatic-Post-Editing/Back-Translation (Freitag et al., 2019; Sennrich et al., 2016a) on this paraphrased set was 0.2/0.4 BLEU.

LOSS	BEAM 1	BEAM 10	SAMPLING
CE	11.26	11.67	8.89
MIXED CE	<b>11.60</b>	<b>11.94</b>	<b>9.90</b>

and  $-1.48$  respectively. The above example shows that: 1) CE-based models contain useful synonym information since they can find synonyms in the top-2 candidates through beam search; 2) mixed CE-based models can allocate more probability mass to those synonyms. It can also be observed that mixed CE-based model concentrates more (in terms of probability mass) on the top-2 candidate translations than CE-based model ( $-2.8$  vs  $-3.08$  ;  $-3.57$  vs  $-4.61$ ), even though both models produce the same candidates.

**Comparison with Label Smoothing** Since label smoothing and mixed CE have a similar formulation (see Section 3.2), we also studied the different impacts these two techniques have on the model. Label smoothing improves generalization by penalizing confident predictions (Pereyra et al., 2017; Müller et al., 2019). To figure out whether mixed CE works in the same way, we first trained 4 different models on WMT’14 En-De: 1) without label smoothing and mixed CE; 2) only with label smoothing; 3) only with mixed CE; 4) with both label smoothing and mixed CE. Then each model generated 5 hypotheses for each sentence in the validation set using sampling decoding and we calculated the Pairwise-BLEU (PB) (Shen et al., 2019) among these hypotheses. PB is used to measure the diversity of the generated hypotheses. The more diverse the hypotheses are, the lower the PB is. Furthermore, a flat probability distribution tends to generate more diverse hypotheses if we use sampling decoding and thus we can measure the sharpness of the output distribution

with PB. We also computed the BLEU score of the 4 models on the validation set using beam search (beam size 5).

The results are shown in Table 4. We can see that label smoothing leads to lower PB than the baseline (the first row), indicating a more flat distribution. Mixed CE, however, gives much higher PB which suggests a more peaked distribution. When combining label smoothing and mixed CE, the resulting PB lies in the middle ground. Thus, mixed CE has totally different impacts on the output distribution as can be measured from PB. Based on the BLEU score, we can see that both label smoothing and mixed CE boost model performance and the combination of them yields even better results. These facts show the two approaches work differently (according to PB) and are able to complement each other (according to BLEU).

To study the distribution properties of 4 models quantitatively, we calculated the cumulative sequence probability<sup>9</sup> of all the hypotheses generated by beam search (Ott et al., 2018). The results are shown in Fig. 4. We can see that label smoothing smears the probability mass evenly in the whole hypothesis space according to the linear increasing trends in cumulative probability. Mixed CE tends to assign the probability mass to the top-scoring candidates as revealed by the sharp increase of the cumulative probability in the top-50 candidates. This is also consistent with the findings in Fig. 3. As we increase the number of hypotheses to  $> 50$ , the cumulative sequence probability increases linearly (approximately).

All the evidence above have proved that mixed CE is different from label smoothing despite a similar formulation. For generation tasks, it is desirable to have more probability mass assigned to tokens that are relevant to the context only, while for irrelevant tokens their probabilities of appearing in a specific context shall be very low, if at all possible (e.g., the word “chicken” is likely an irrelevant word with respect to the text discussing aircraft maintenance). On the other hand, it is also necessary to avoid overfitting by penalizing confident predictions (Pereyra et al., 2017; Müller et al., 2019). Therefore, using both label smoothing and mixed CE together may allow us to have the best of both worlds, arriving at the optimal results in practice.

### 4.3. Mixed CE in Scheduled Sampling

In scheduled sampling, we tested mixed CE on two different baselines, standard SS (Bengio et al., 2015) and *word oracle* (Zhang et al., 2019).<sup>10</sup> Note that in word oracle experiments, the  $\hat{y}_t$  in the loss function (see Eq. 3) is still obtained from

<sup>9</sup>The sum of the probabilities of the generated hypotheses.

<sup>10</sup>Zhang et al. (2019) compared their approach with other reinforcement learning based baselines and thus we did not include those comparisons here.

Table 4. PB, BLEU on WMT’14 En-De validation set. Pairwise-BLEU is obtained using sampling decoding while the BLEU score is obtained using beam search. LS is short for label smoothing.

Loss	PB ( $\downarrow$ )	BLEU ( $\uparrow$ )
NO LS, NO MIXED CE	17.52	25.81
+ LS	5.22	26.48
+ MIXED CE	25.99	26.26
+ LS, MIXED CE	7.79	<b>26.75</b>




Figure 4. Cumulative sequence probability of generated hypotheses using beam search with beam size 200 on WMT’14 En-De validation set.

$\operatorname{argmax}_{1 \leq k \leq |V|} \log p_{\theta}(w_k | y_{< t}, x)$  while  $y^{\text{mix}}$  in Fig. 2 is obtained by mixing  $y_t$  with  $\hat{y}_t = \operatorname{argmax}_{1 \leq k \leq |V|} s_{t,k}$  ( $k$ -th element of  $s_t$ ; for  $s_t$ , see Sec. 2). The results are shown in Table 5. We can see that mixed CE surpasses CE on small and medium data sets by a large margin but this gap is smaller on a larger data set.

But the BLEU improvements alone are not enough to demonstrate the necessity of the second part in Eq. 3. Hence, we further modify the way to calculate  $\hat{y}_t$  which is produced in the first pass:

$$\hat{y}_t = \operatorname{Rand}\left(\operatorname{Top-2}_{1 \leq k \leq |V|}(\log p_{\theta}(w_k | y_{< t}, x))\right) \quad (6)$$

We randomly sampled the highest or the second highest scoring token to substitute the  $\hat{y}_t$  in Eq. 3. The results are shown in Table 6. We can see that in the third row, top-2 mixed CE significantly impairs the model’s performance. We surmise that if the model can correctly predict the gold token as the highest-scoring token and yet we still maximize the likelihood of the second highest-scoring token, this will confuse the model and lead to a drop in performance. So we further modified Eq. 6:

$$\hat{y}_t = \begin{cases} y_t, & \text{if } y_t = \operatorname{argmax}_{1 \leq k \leq |V|} \log p_{\theta}(w_k | y_{< t}, x) \\ \operatorname{Rand}(V), & \text{otherwise} \end{cases} \quad (7)$$

$\operatorname{Rand}(V)$  denotes a random token in the vocabulary. The results can be found in the 4-th row in Table 6. Random mixed CE is better than top-2 mixed CE but still underperforms

Table 5. BLEU scores on test sets of Transformers trained with CE and mixed CE. The results of beam search decoding with beam size 1/5 are presented. All results are averaged over 3 runs.

DATA SET	LOSS	SCHEDULED SAMPLING	
		SINGLE	AVERAGE
RO-EN	CE	30.71/31.72	32.29/33.05
	MIXED CE	<b>31.71/32.53</b>	<b>32.88/33.45</b>
RU-EN	CE	29.28/30.63	29.62/30.83
	MIXED CE	<b>30.19/31.23</b>	<b>30.47/31.39</b>
EN-DE	CE	26.36/27.29	26.84/27.56
	MIXED CE	<b>26.75/27.57</b>	<b>26.99/27.71</b>
DATA SET	LOSS	WORD ORACLE	
		SINGLE	AVERAGE
RO-EN	CE	31.71/32.37	33.05/33.76
	MIXED CE	<b>32.43/33.06</b>	<b>33.66/34.14</b>
RU-EN	CE	29.40/30.61	29.87/31.00
	MIXED CE	<b>30.24/31.09</b>	<b>30.72/31.50</b>
EN-DE	CE	26.66/27.45	26.94/27.71
	MIXED CE	<b>26.81/27.80</b>	<b>26.94/27.88</b>

Table 6. BLEU scores of Transformers trained with different loss functions on the WMT’16 Ro-En validations sets.

LOSS	SS	WORD ORACLE
CE	32.66	33.82
MIXED CE	<b>33.64</b>	<b>34.51</b>
TOP-2 MIXED CE	32.17	32.76
RANDOM MIXED CE	33.26	34.18
SOFT MIXED CE	32.03	33.08

mixed CE, indicating the importance of approximating the outputs conditioned on gold inputs. It is also worth to note that random mixed CE outperforms CE for which we conjecture that random mixed CE provides a form of regularization. We also tried a soft version of mixed CE in scheduled sampling as shown in Eq. 8, which is denoted as soft mixed CE. Here  $q_\theta(w_k|y_{<t}, x)$  denotes the model’s prediction in the first pass for the  $k$ -th token in the vocabulary at time step  $t$  conditioned on gold prefixes. The results can be found in the 5-th row in Table 6. The results indicate that we do not need to simulate the outputs in the first pass that are considered unlikely by the model in order to homogenize train/test behaviors.

$$\mathcal{L}_{\text{mix}} = - \sum_{t=1}^n [(1 - \alpha_i) \cdot \log p_\theta(y_t|y_{<t}^{\text{mix}}, x) + \alpha_i \cdot \sum_{k=1}^{|V|} q_\theta(w_k|y_{<t}, x) \cdot \log p_\theta(w_k|y_{<t}^{\text{mix}}, x)] \quad (8)$$

Table 7. BLEU scores of Transformers trained with *double mixed CE* and *mixed CE-2nd pass* on validations sets.

LOSS	RO-EN	RU-EN	EN-DE
CE	33.82	29.83	26.51
MIXED CE	<b>34.51</b>	<b>30.46</b>	26.88
DOUBLE MIXED CE	34.23	<b>30.46</b>	<b>27.06</b>
MIXED CE-2ND PASS	33.84	30.16	26.83

#### 4.4. Combining Two Mixed CE

In scheduled sampling training, we also considered making use of the model’s predictions in the second forward pass as what we do in teacher forcing by modifying mixed CE in scheduled sampling as follows:

$$\mathcal{L}_{\text{mix}} = - \sum_{t=1}^n [(1 - \alpha_i) \cdot \log p_\theta(y_t|y_{<t}^{\text{mix}}, x) + \frac{\alpha_i}{2} \cdot (\log p_\theta(\hat{y}_t|y_{<t}^{\text{mix}}, x) + \log p_\theta(\tilde{y}_t|y_{<t}^{\text{mix}}, x))] \quad (9)$$

Here,  $\tilde{y}_t$  is the greedy decision in the second forward pass, i.e.,  $\tilde{y}_t = \operatorname{argmax}_{1 \leq k \leq |V|} \log p_\theta(w_k|y_{<t}^{\text{mix}}, h)$ . We denote this new loss as double mixed CE. The results are shown in Table 7. The performance does not vary much on all data sets. But does this mean  $\tilde{y}_t$  and  $\hat{y}_t$  are equally informative? So we further modified Eq. 9:

$$\mathcal{L}_{\text{mix}} = - \sum_{t=1}^n [(1 - \alpha_i) \cdot \log p_\theta(y_t|y_{<t}^{\text{mix}}, x) + \alpha_i \cdot \log p_\theta(\tilde{y}_t|y_{<t}^{\text{mix}}, x)] \quad (10)$$

We denote this loss as mixed CE-2nd pass. It turns out that mixed CE-2nd is better than CE and this is because  $\tilde{y}$  resembles  $\hat{y}$  (consider that  $y^{\text{mix}}$  and  $y$  have common parts). However, it is still worse than mixed CE, indicating the importance of approximating  $\hat{y}$ .

#### 4.5. The Effect of $m$ in $\alpha_i$

Here we tried different values of the constant  $m = \{0.3, 0.4, 0.5, 0.6, 0.7\}$  in Eq. 4 in both teacher forcing and scheduled sampling training on WMT’16 Ro-En. The best BLEU scores on validation set with different  $m$  are shown in Fig. 5. In teacher forcing, larger  $m$  values favor the model’s inconsistent predictions over the gold data whereas smaller values do not trust model predictions that much. In scheduled sampling, larger  $m$  values highlight the importance of reconciling two outputs in two passes while smaller  $m$  values emphasize approximating the gold target. In general,  $m = 0.5$  makes a good trade-off and the gap between different  $m$  is smaller in scheduled sampling than in teacher forcing. We also tried to fix  $\alpha_i$  to be 0.5 throughout the






Figure 5. BLEU scores on the WMT’16 Ro-En validation set with different  $m$  values. The blue and orange dotted lines denote the BLEU scores of the model with  $\alpha_i = 0.5$  while the dashed lines denote the result of training with CE loss.

training process and the best validation BLEU is shown in dotted lines in Fig. 5. It turns out that this fixed strategy is more harmful in teacher forcing since it discourages the model from learning at the early stage when the model is still uninformative. We also plotted the validation score of the models with CE loss in Fig. 5 (dashed lines). It can be seen that the fixed strategy still outperforms CE loss in scheduled sampling even though in teacher forcing CE exceeds the fixed strategy by a large margin.

## 5. Related Work

Neural machine translation has made significant progress in recent year (Cho et al., 2014; Sutskever et al., 2014; Bahdanau et al., 2015; Vaswani et al., 2017) But many autoregressive models face exposure bias (Ranzato et al., 2016). The mainstream solution to exposure bias is to train the model on its own predictions, (Daumé et al., 2009; Ross et al., 2011; Venkatraman et al., 2015). Later, Bengio et al. (2015) proposed scheduled sampling for RNN training. There are also many variants of scheduled sampling (Goyal et al., 2017; Mihaylova & Martins, 2019; Zhang et al., 2019; Duckworth et al., 2019). In addition, Ranzato et al. (2016); Bahdanau et al. (2017) incorporated the ideas in reinforcement learning into sequence prediction problem while Wiseman & Rush (2016); Zhang et al. (2019) integrated beam search into the sequence-to-sequence training procedure. Leblond et al. (2018) linked RNNs with SEARN (Daumé et al., 2009) and proposed SEARNN. There are also some other interesting works addressing exposure bias, e.g., minimum risk training (Shen et al., 2016), adversarial training (Lamb et al., 2016). More recently, Schmidt (2019), He et al. (2019) provided a new perspective to understand exposure bias.

Real noisy data problem in neural machine translation has been well studied (Koehn et al., 2018; Wang et al., 2018; Belinkov & Bisk, 2018; Dakwale & Monz, 2019), but we’ve argued that NMT with clean data can also be treated as a noisy label problem in a sense (see Section 3.2). Song et al. (2020) provides a thorough review of approaches to handling noisy

labels and one of the approaches called Bootstrapping (Reed et al., 2015) is similar to mixed CE. However, there are several differences: 1) mixed CE is used in machine translation with clean labels; 2) when we treat those clean texts as “noisy” ones, the ratio of “noisy” labels is nearly 100% while in Reed et al. (2015) this number is much smaller; 3) mixed CE assigns linearly decay coefficients to the two log likelihood with human priors while Bootstrapping selects a fixed value via cross-validation. Another similar approach to mixed CE in teacher forcing is called self-knowledge distillation (Hahn & Choi, 2019) which uses the scaled Euclidean distance between the word embeddings of the target token and model’s greedy prediction to compute the  $\alpha$  values. We re-implemented their approach but failed to find one proper scaled factor that could outperform the baseline in our setting. More details can be found in the Appendix.

## 6. Conclusion

In this paper, we propose mixed CE which can be used in teacher forcing training and scheduled sampling training in neural machine translation. In teacher forcing training, mixed CE can make full use of the model’s own predictions during training and tends to assign probability mass to the tokens related to the gold targets. We systematically analyze the output distribution’s properties of mixed CE and make a comparison with label smoothing. In scheduled sampling, mixed CE forces the model to approximate not only the gold targets but also the greedy predictions in the first forward pass conditioned on gold inputs and thus mitigate exposure bias more effectively. We demonstrate the effectiveness of mixed CE on several standard machine translation data sets at different scales, namely WMT’16 Ro-En, WMT’16 Ru-En, WMT’14 En-De as well as two sets of additional challenging references. Specifically, in a multi-reference set, mixed CE consistently outperforms CE across 10 additional references. Such results demonstrate the effectiveness of our proposed mixed CE objective for neural machine translation. In the future, it would be also interesting to explore the use of mixed CE in non-autoregressive machine translation and domain robustness problems.

## Acknowledgements

We would like to thank the anonymous reviewers and the meta-reviewer for their insightful comments. We would also like to thank Wen Zhang, Yang Feng, Zhanming Jie, Perry Lam for their help. This research is supported by Ministry of Education, Singapore, under its Academic Research Fund (AcRF) Tier 2 Programme (MOE AcRF Tier 2 Award No: MOE2017-T2-1-156). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of the Ministry of Education, Singapore.

## References

- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations*, 2015.
- Bahdanau, D., Brakel, P., Xu, K., Goyal, A., Lowe, R., Pineau, J., Courville, A. C., and Bengio, Y. An actor-critic algorithm for sequence prediction. In *5th International Conference on Learning Representations*, 2017.
- Belinkov, Y. and Bisk, Y. Synthetic and natural noise both break neural machine translation. In *6th International Conference on Learning Representations*, 2018.
- Bengio, S., Vinyals, O., Jaitly, N., and Shazeer, N. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, volume 28, pp. 1171–1179. 2015.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, 2014.
- Dakwale, P. and Monz, C. Improving neural machine translation using noisy parallel data through distillation. In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, pp. 118–127, 2019.
- Daumé, H., Langford, J., and Marcu, D. Search-based structured prediction. *Machine Learning*, 75:297–325, 2009.
- Duckworth, D., Neelakantan, A., Goodrich, B., Kaiser, L., and Bengio, S. Parallel scheduled sampling. *ArXiv*, abs/1906.04331, 2019.
- Elman, J. Finding structure in time. *Cogn. Sci.*, 14:179–211, 1990.
- Freitag, M., Caswell, I., and Roy, S. APE at scale and its implications on MT evaluation biases. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pp. 34–44, 2019.
- Freitag, M., Grangier, D., and Caswell, I. Bleu might be guilty but references are not innocent. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 61–71, 2020.
- Furlanello, T., Lipton, Z., Tschannen, M., Itti, L., and Anandkumar, A. Born again neural networks. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1607–1616. PMLR, 2018.
- Goyal, K., Dyer, C., and Berg-Kirkpatrick, T. Differentiable scheduled sampling for credit assignment. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pp. 366–371, 2017.
- Gumbel, E. *Statistical Theory of Extreme Values and Some Practical Applications: A Series of Lectures*. Applied mathematics series. U.S. Government Printing Office, 1954.
- Hahn, S. and Choi, H. Self-knowledge distillation in natural language processing. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pp. 423–430, 2019.
- He, T., Zhang, J., Zhou, Z., and Glass, J. R. Quantifying exposure bias for neural language generation. *CoRR*, abs/1905.10617, 2019.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9:1735–1780, 1997.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations*. OpenReview.net, 2017.
- Kim, Y. and Rush, A. M. Sequence-level knowledge distillation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1317–1327, 2016.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2015.
- Koehn, P., Khayrallah, H., Heafield, K., and Forcada, M. L. Findings of the WMT 2018 shared task on parallel corpus filtering. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pp. 726–739, 2018.
- Koppel, M. and Ordan, N. Translationese and its dialects. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 1318–1326, 2011.
- Lamb, A. M., ALIAS PARTH GOYAL, A. G., Zhang, Y., Zhang, S., Courville, A. C., and Bengio, Y. Professor forcing: A new algorithm for training recurrent networks. In *Advances in Neural Information Processing Systems*, volume 29, pp. 4601–4609. 2016.
- Leblond, R., Alayrac, J.-B., Osokin, A., and Lacoste-Julien, S. SEARNN: Training RNNs with global-local losses. In *6th International Conference on Learning Representations*, 2018.

- Lee, J., Mansimov, E., and Cho, K. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1173–1182, 2018.
- Mihaylova, T. and Martins, A. F. T. Scheduled sampling for transformers. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pp. 351–356, 2019.
- Müller, R., Kornblith, S., and Hinton, G. E. When does label smoothing help? In *Advances in Neural Information Processing Systems*, volume 32, pp. 4694–4703, 2019.
- Norouzi, M., Bengio, S., Chen, z., Jaitly, N., Schuster, M., Wu, Y., and Schuurmans, D. Reward augmented maximum likelihood for neural structured prediction. In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 29, pp. 1723–1731. Curran Associates, Inc., 2016.
- Ott, M., Auli, M., Grangier, D., and Ranzato, M. Analyzing uncertainty in neural machine translation. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3956–3965. PMLR, 2018.
- 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 *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pp. 48–53, 2019.
- 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*, pp. 311–318, 2002.
- Pereyra, G., Tucker, G., Chorowski, J., Kaiser, L., and Hinton, E. G. Regularizing neural networks by penalizing confident output distributions. *5th International Conference on Learning Representations*, 2017.
- Ranzato, M., Chopra, S., Auli, M., and Zaremba, W. Sequence level training with recurrent neural networks. In *4th International Conference on Learning Representations*, 2016.
- Reed, S. E., Lee, H., Anguelov, D., Szegedy, C., Erhan, D., and Rabinovich, A. Training deep neural networks on noisy labels with bootstrapping. In *3rd International Conference on Learning Representations*, 2015.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Schmidt, F. Generalization in generation: A closer look at exposure bias. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pp. 157–167, 2019.
- Sennrich, R., Haddow, B., and Birch, A. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pp. 86–96, 2016a.
- 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*, pp. 1715–1725, 2016b.
- Shen, S., Cheng, Y., He, Z., He, W., Wu, H., Sun, M., and Liu, Y. Minimum risk training for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pp. 1683–1692, 2016.
- Shen, T., Ott, M., Auli, M., and Ranzato, M. Mixture models for diverse machine translation: Tricks of the trade. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5719–5728. PMLR, 2019.
- Song, H., Kim, M., Park, D., and Lee, J.-G. Learning from noisy labels with deep neural networks: A survey, 2020.
- Sutskever, I., Vinyals, O., and Le, Q. V. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, volume 27, pp. 3104–3112. 2014.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pp. 5998–6008. 2017.
- Venkatraman, A., Hebert, M., and Bagnell, J. A. Improving multi-step prediction of learned time series models. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI, pp. 3024–3030, 2015.

- Wang, W., Watanabe, T., Hughes, M., Nakagawa, T., and Chelba, C. Denoising neural machine translation training with trusted data and online data selection. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 133–143, 2018.
- Williams, R. J. and Zipser, D. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2):270–280, 1989.
- Wiseman, S. and Rush, A. M. Sequence-to-sequence learning as beam-search optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1296–1306, 2016.
- Xiao, F., Wu, Y., Zhao, H., Wang, R., and Jiang, S. Dual skew divergence loss for neural machine translation. *CoRR*, abs/1908.08399, 2019.
- Zhang, W., Feng, Y., Meng, F., You, D., and Liu, Q. Bridging the gap between training and inference for neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4334–4343, 2019.