

Improving Scheduled Sampling with Elastic Weight Consolidation for Neural Machine Translation

Michalis Korakakis
University of Cambridge
mk2008@cam.ac.uk

Andreas Vlachos
University of Cambridge
av308@cam.ac.uk

Abstract

Despite strong performance in many sequence-to-sequence tasks, autoregressive models trained with maximum likelihood estimation suffer from exposure bias, i.e. the discrepancy between the ground-truth prefixes used during training and the model-generated prefixes used at inference time. Scheduled sampling is a simple and empirically successful approach which addresses this issue by incorporating model-generated prefixes into training. However, it has been argued that it is an inconsistent training objective leading to models ignoring the prefixes altogether. In this paper, we conduct systematic experiments and find that scheduled sampling, while it ameliorates exposure bias by increasing model reliance on the input sequence, worsens performance when the prefix at inference time is correct, a form of catastrophic forgetting. We propose to use Elastic Weight Consolidation to better balance mitigating exposure bias with retaining performance. Experiments on four IWSLT'14 and WMT'14 translation datasets demonstrate that our approach alleviates catastrophic forgetting and significantly outperforms maximum likelihood estimation and scheduled sampling baselines.

1 Introduction

Autoregressive models trained with maximum likelihood estimation (MLE) constitute the dominant approach in several sequence-to-sequence tasks, such as machine translation (Bahdanau et al., 2015), text summarization (See et al., 2017), and conversational modeling (Vinyals and Le, 2015). However, this paradigm suffers from exposure bias (Bengio et al., 2015; Ranzato et al., 2016), i.e. during training the model generates tokens by conditioning on the ground-truth prefixes, while at inference time model-generated prefixes are used instead. Since the model is never exposed to its own errors, if a token is mistakenly generated during inference the error will be propagated along

the sequence (Ross et al., 2011). Prior work has attributed to exposure bias various forms of text degeneration, such as repetitiveness, incoherence, and hallucinations (Holtzman et al., 2020; Wang and Sennrich, 2020), i.e. outputs which contain information irrelevant to the input sequence.

Bengio et al. (2015) introduced scheduled sampling to address exposure bias in MLE-trained autoregressive models. Scheduled sampling uses a stochastic mixture of ground-truth and model-generated prefixes during training, thereby allowing the model to learn how to recover from its own errors. While various alternatives to MLE have been proposed to mitigate exposure bias (Ranzato et al., 2016; Wiseman and Rush, 2016; Shen et al., 2016; Bahdanau et al., 2017, *inter alia*), scheduled sampling remains one of the most popular due to its simplicity and performance improvements in many conditional sequence generation tasks (Bengio et al., 2015; Du and Ji, 2019; Zhang et al., 2019; Li and Lu, 2021).

Conversely, other studies have reported that using scheduled sampling may hurt performance (Leblond et al., 2018; Mihaylova and Martins, 2019). The dominant hypothesis for these negative results is that scheduled sampling creates models that are more likely to recover from their own mistakes by training them to ignore the prefixes entirely (Huszar, 2015). However, no attempt has been made to empirically assess this hypothesis. Thus it is still unclear how scheduled sampling affects training.

In this paper, we provide insights into the working mechanisms of scheduled sampling. Following Voita et al. (2021), we apply Layerwise Relevance Propagation (LRP) (Bach et al., 2015) to quantify the contributions of the input sequence and the prefix during inference, and empirically evaluate the hypothesis of Huszar (2015). We find that models trained with scheduled sampling increase their reliance on the input sequence, and

therefore mitigate exposure bias by depending less on the potentially incorrect model-generated prefix. However, we also observe that this has the side-effect of worsening the predictions when the model-generated prefix is correct. To address this form of catastrophic forgetting (French, 1999), we propose using Elastic Weight Consolidation (EWC) (Kirkpatrick et al., 2016) to regularize the parameter updates in scheduled sampling so that the performance is not affected when the prefix is correct.

Experiments on the commonly-used IWSLT’14 German-English, IWSLT’14 Vietnamese-English, WMT’14 English-German, and WMT’14 English-French translation datasets show that our proposed method mitigates catastrophic forgetting and significantly improves translation performance, in terms of BLEU (Papineni et al., 2002), over MLE, standard scheduled sampling, and a recently proposed scheduled sampling variant (Zhang et al., 2019). Importantly, performance gains occur in both long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) and Transformer (Vaswani et al., 2017) models, and across different annealing schedules, demonstrating that our proposed EWC-regularized scheduled sampling is more robust and easier to tune. Finally, human evaluation shows that our method can improve both translation adequacy and fluency.

2 Scheduled Sampling

Autoregressive models estimate the conditional probability of the output \mathbf{y} given the input \mathbf{x} one token at a time in a monotonic fashion:

$$P(\mathbf{y} | \mathbf{x}) = \prod_{t=1}^T p(y_t | \mathbf{y}_{<t}, \mathbf{x}; \theta), \quad (1)$$

where y_t is the t -th token in \mathbf{y} , $\mathbf{y}_{<t}$ denotes all previous tokens, and θ is the set of model parameters.

Given a dataset $\mathcal{D} = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^N$ of input-output pairs, the standard approach to optimize the parameters θ of an autoregressive model entails maximizing the conditional log-likelihood:

$$\hat{\theta}_{\text{MLE}} = \underset{\theta}{\operatorname{argmax}} \mathcal{L}(\theta; \mathcal{D}), \quad (2)$$

$$\mathcal{L}(\theta; \mathcal{D}) = \sum_{i=1}^N \sum_{t=1}^{L^{(i)}} \log p(y_t^{(i)} | \mathbf{y}_{<t}^{(i)}, \mathbf{x}^{(i)}; \theta). \quad (3)$$

Here i indicates the i -th output sequence in the dataset and $L^{(i)}$ is the length of the i -th output

Algorithm 1 Scheduled Sampling

```

1: Input: Dataset  $\mathcal{D}$ 
2: Initialize  $a = 1$ 
3: for  $i = 1, \dots, I$  do
4:   repeat
5:     for  $t = 1, \dots, T$  do
6:        $\tilde{y}_t = \begin{cases} y_t & \text{with } a \\ \hat{y}_t \sim p_\theta(y_t | \tilde{y}_{1:t-1}, x) & \text{with } 1 - a \end{cases}$ 
7:        $L(\theta) = L(\theta) + \log p_\theta(y_t | \tilde{y}_{1:t-1}, x)$ 
8:     until  $B$  times
9:      $\theta = \theta + \eta \cdot \nabla_\theta L(\theta)$ 
10:     $a = \text{SCHEDULE}(a)$ 

```

sequence. This training objective is known as teacher-forcing (Williams and Zipser, 1989), since the model conditions on the ground-truth prefix $\mathbf{y}_{<t}^{(i)}$ to generate the token $y_t^{(i)}$. However, at inference time, the model generates the token \hat{y}_t by conditioning on its own outputs, i.e. $\hat{\mathbf{y}}_{<t}$ instead of $\mathbf{y}_{<t}$, which creates a discrepancy between training and inference known as the exposure bias problem (Bengio et al., 2015; Ranzato et al., 2016).

Bengio et al. (2015) introduced scheduled sampling to mitigate the above-mentioned discrepancy between MLE training and inference. Scheduled sampling uses the same training objective as teacher-forcing (Equation 3), the only difference being that the conditioning prefixes $\tilde{\mathbf{y}}_{<t}^{(i)}$ are a stochastic mixture of ground-truth $\mathbf{y}_{<t}^{(i)}$ and model-generated prefixes $\hat{\mathbf{y}}_{<t}^{(i)}$:

$$\mathcal{L}(\theta; \mathcal{D}) = \sum_{i=1}^N \sum_{t=1}^{L^{(i)}} \log p(y_t^{(i)} | \tilde{\mathbf{y}}_{<t}^{(i)}, \mathbf{x}^{(i)}; \theta). \quad (4)$$

An annealing schedule is used to gradually decrease the probability a of conditioning using the ground-truth prefix during training. Typically, for each mini-batch b , a is decreased using the following annealing schedules:

- **Linear:** $a = \max(a - kb, 0)$
- **Exponential:** $a = k^b$
- **Inverse sigmoid:** $a = k / (k + \exp(b/k))$

Here k is a hyperparameter which controls the speed of decay of a in each schedule. Algorithm 1 summarizes training with scheduled sampling.

3 Analysis of Scheduled Sampling

In this section, we propose two ways to investigate how scheduled sampling works. First, we use Layerwise Relevance Propagation (LRP) (Bach et al.,

2015; Voita et al., 2021) to examine the contributions of the input sequence and the prefix during inference, and thus assess the hypothesis of Huszar (2015). Then we treat training with scheduled sampling as a form of progressive fine-tuning, by assuming that the model-generated prefixes can be viewed as a downstream task we adapt the model to after it is trained with the ground-truth prefixes. To this end, we apply teacher-forcing at inference time to quantify the impact of catastrophic forgetting on models trained with scheduled sampling.

3.1 Prefixes under Scheduled Sampling

Huszar (2015) argued that scheduled sampling is an inappropriate training objective since it learns models that ignore the prefixes. This limitation arises because the model-generated outputs correspond to a distribution that is different from the ground-truth sequences the model is trained to generate. Therefore training might not converge to the correct model even as the dataset and the capacity of the model increase indefinitely.

To empirically verify this hypothesis, we apply LRP to quantify the influence of the input sequence and the prefix during prediction when the model conditions using (i) random prefixes that are valid in the target language but semantically unrelated to the input sequence, and (ii) model-generated prefixes. The intuition behind this experiment is that if scheduled sampling creates models that ignore the prefixes, then the input sequence and the prefix contributions should remain relatively stable across the two prefix types.

LRP quantifies the proportion of each token’s influence during prediction, and thus it has been used to analyze the effect of source and target contexts (Ding et al., 2017; Voita et al., 2021). In terms of explainability and reliability, it has been shown to outperform related methods in text classification (Arras et al., 2017; Poerner et al., 2018; Arras et al., 2019). Concretely, given an input sequence token x_i and a prefix token y_j , LRP computes the relevance scores $r_t(x_i)$ and $r_t(y_j)$ at every output generation step t by back-propagating from the output to the input. Importantly, the total relevance scores for each generated token are equal to 1:

$$\sum_i r_t(x_i) + \sum_{j=1}^{t-1} r_t(y_j) = 1. \quad (5)$$

In the general case, the relevance score $r_j^{(l)}$ for the

Algorithm 2 Scheduled Sampling as Fine-Tuning

```

1: Input: Dataset  $\mathcal{D}$ 
2: Initialize  $a = 1$ 
3: for  $i = 1, \dots, I$  do
4:   Initialize  $\tilde{\mathcal{D}}, \tilde{\mathcal{D}}_{cond} = \emptyset$ 
5:   repeat
6:     Sample example  $(x, y)$  from  $\mathcal{D}$ 
7:     for  $t = 1, \dots, T$  do
8:        $\tilde{y}_t = \begin{cases} y_t & \text{with } a \\ \hat{y}_t \sim p_\theta(y_t | \tilde{y}_{1:t-1}, x) & \text{with } 1 - a \end{cases}$ 
9:       Add  $(x, y)$  to  $\tilde{\mathcal{D}}$  and  $(x, \tilde{y})$  to  $\tilde{\mathcal{D}}_{cond}$ 
10:    until  $B$  times
11:    Compute loss on  $\tilde{\mathcal{D}}$  conditioning with  $\tilde{\mathcal{D}}_{cond}$ 
12:     $a \leftarrow \text{SCHEDULE}(a)$ 

```

j -th neuron in layer l is computed as:¹

$$r_j^{(l)} = \sum_{k=1}^K \left(\alpha \frac{z_{jk}^+}{\sum_i z_{ji}^+} - \beta \frac{z_{jk}^-}{\sum_i z_{ji}^-} \right) r_k^{(l+1)}, \quad (6)$$

where $z_{jk} = x_j w_{jk}$. Here K denotes the total number of neurons at each layer, j and k are neurons at two consecutive layers l and $l + 1$, x_j is the activation of neuron j at layer l , w_{jk} is the weight connecting neurons j and k , $r^{(l+1)}$ is the relevance score at layer $(l + 1)$, $+$ and $-$ indicate positive and negative contributions to $r^{(l+1)}$, and finally, α and β are hyperparameters that determine the importance of the positive (+) and negative (-) contributions.

3.2 Scheduled Sampling as Fine-Tuning

Our hypothesis is that scheduled sampling can be viewed as progressively fine-tuning a model using its own predictions as prefixes to address exposure bias (Algorithm 2). A commonly occurring issue is that during fine-tuning the model’s ability to retain previously learned knowledge is negatively impacted, a phenomenon known as catastrophic forgetting (French, 1999).

To this end, inspired by Wu et al. (2018) who applied teacher-forcing during inference to investigate the influence of exposure bias in NMT, we use the ground-truth prefix $\mathbf{y}_{<t}$ instead of the model-generated prefix $\hat{\mathbf{y}}_{<t}$ to generate the output token $\hat{\mathbf{y}}_t$. If there is no catastrophic forgetting, then the output quality of models trained with scheduled sampling is expected to be similar to that of MLE-trained models when ground-truth tokens are used during inference.

¹Since we use implementation of Voita et al. (2021), we adopt the LRP- $\alpha\beta$ rule (Bach et al., 2015; Binder et al., 2016) for computing relevance scores.

4 Experiments

In this section we apply the methods discussed in Section 3 to examine if training with scheduled sampling leads to models ignoring the prefixes, and also determine whether scheduled sampling causes catastrophic forgetting.

4.1 Experimental Setup

Data Following prior work exploring scheduled sampling in NMT (Goyal et al., 2017; Xu et al., 2019), we conduct experiments on two IWSLT’14 translation datasets, the German-English (DE-EN) and Vietnamese-English (VI-EN). We use byte pair encoding (BPE) (Sennrich et al., 2016) to encode sentences. We report cased, tokenized BLEU (Papineni et al., 2002).²

Models We use LSTM (Hochreiter and Schmidhuber, 1997) and Transformer (Vaswani et al., 2017) models trained with default architectures and hyperparameters from fairseq (Ott et al., 2019). Models are first trained with MLE and then fine-tuned using scheduled sampling. For scheduled sampling, we use the inverse sigmoid annealing schedule for all experiments unless otherwise stated, since we found it to perform better than the linear and exponential schedules. We set all hyperparameters according to the performance on the validation set of each task. We report the mean and standard deviation over 5 runs with different random seeds.

4.2 Results

Prefixes under Scheduled Sampling Figure 1 shows the source contributions for the random and model-generated prefixes. Following the experimental setup of Voita et al. (2021), we randomly choose 100 source-target pairs from each test set with the same length in the source and target and average all results, as this renders the analyses of the contributions more reliable. Furthermore, we use the Transformer model to quantify contributions since the implementation of LRP by Voita et al. (2021) is tailored to this particular neural architecture. Figure 1 illustrates that at the beginning, MLE-trained models tend to use the source more, however, as the target prefix becomes longer, they rely less on the source and more on the prefix (especially for the random prefixes), given that:

²<https://github.com/moses-smt/mosesdecoder/blob/master/scripts/generic/multi-bleu.perl>

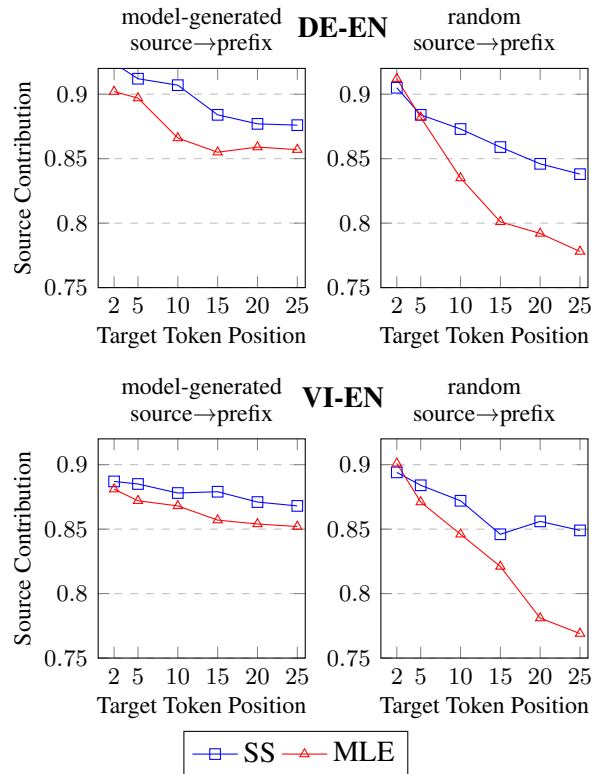


Figure 1: Source contributions with model-generated and random prefixes.

$$\sum_{j=1}^{t-1} r_t(y_j) = 1 - \sum_i r_t(x_i). \quad (7)$$

This behavior is in agreement with prior work that has attributed the increasing reliance on the prefix to teacher-forcing (Wang and Sennrich, 2020; Voita et al., 2021). Conversely, scheduled sampling, which mitigates exposure bias, promotes the usage of source information to generate a target token, as the decrease in the influence of source contributions is less drastic compared to MLE-trained models across the two prefix types. This suggests that scheduled sampling results in models that tend to ignore the prefix more than MLE-trained models, thus confirming the hypothesis of Huszar (2015). However, in contrast to Huszar (2015), we argue that this is a positive outcome in the context of NMT, i.e. the model-generated prefix is often incorrect, but the source that is being translated does not change. Therefore, learning to rely on it more is a reasonable approach to alleviate exposure bias.

Scheduled Sampling as Fine-Tuning Table 1 empirically confirms the occurrence of catastrophic

Method	DE-EN			VI-EN		
	MP	TF	Δ	MP	TF	Δ
LSTM						
MLE	26.72	28.67	+1.95	22.64	24.98	+2.34
SS	27.35	23.59	-3.76	23.58	18.92	-4.66
Transformer						
MLE	34.08	35.96	+1.88	27.13	29.14	+2.01
SS	34.65	29.84	-4.81	27.57	22.42	-5.15

Table 1: BLEU scores by conducting inference with two approaches, namely, (i) by conditioning using model-generated prefixes (MP), and (ii) with teacher-forcing (TF), i.e. feeding the ground-truth as an input to generate the output token.

forgetting in models trained with scheduled sampling. We observe that when models condition using the ground-truth prefixes, the quality of generations drops heavily in terms of BLEU across both languages and neural architectures, suggesting that scheduled sampling results in forgetting how to predict when the prefix is correct. In particular, on DE-EN the BLEU score reduction is 3.76 for the LSTM model and 4.81 for the Transformer. Similarly, on VI-EN the BLEU score degrades by 4.66 and by 5.15. On the other hand, we see that in MLE-trained models, using teacher-forcing at inference time leads to consistent BLEU score improvements over conditioning with model-generated prefixes. Specifically, on DE-EN the BLEU score gain is 1.95 for the LSTM model and 1.88 for the Transformer, and on VI-EN the BLEU score increases by 2.34 and by 2.01.

5 Elastic Weight Consolidation for Scheduled Sampling

Section 4 shows that even though scheduled sampling addresses exposure bias by increasing model reliance on the input sequence, it also leads to output degradation due to catastrophic forgetting (French, 1999) when the model-generated prefix is correct at inference time. Elastic Weight Consolidation (EWC) (Kirkpatrick et al., 2016) is an effective regularization-based approach for mitigating catastrophic forgetting during fine-tuning with successful applications in NMT (Thompson et al., 2019; Saunders et al., 2019).

To this end, we use EWC to regularize conditioning with the model-generated prefixes $\hat{y}_{<t}$ to avoid forgetting using the ground-truth $y_{<t_j}$ as prefixes. EWC penalizes changes to the parameters which are important for the ground-truth prefixes,

while adapting more the parameters that are less important:

$$\mathcal{L}_{\text{EWC}}(\theta^{\hat{y}_{<t}}; \mathcal{D}^{\hat{y}_{<t}}) = \mathcal{L}(\theta^{\hat{y}_{<t}}; \mathcal{D}^{\hat{y}_{<t}}) + \lambda \sum_j F_j (\theta_j^{\hat{y}_{<t}} - \theta_j^{y_{<t}})^2, \quad (8)$$

where $\mathcal{L}(\theta^{\hat{y}_{<t}}; \mathcal{D}^{\hat{y}_{<t}})$ is the loss (Equation 4) over the model-generated prefixes, λ is a hyperparameter which determines the importance of conditioning with $y_{<t}$, $\theta_j^{y_{<t}}$ are the parameters of the initial model trained with $y_{<t}$, $\theta_j^{\hat{y}_{<t}}$ are the parameters of the model trained with $\hat{y}_{<t}$, and $F_j = E[\nabla^2 L(\theta_j^{y_{<t}})]$ is an estimate of how important the parameters $\theta_j^{y_{<t}}$ are to $y_{<t}$. This can be approximated via the empirical Fisher information matrix (Martens, 2014), which requires samples from $y_{<t}$. Concretely, using the EWC regularizer, we modify the Lines 7 and 9 of Algorithm 1 to penalize changes in model parameters associated with conditioning using the ground-truth.

6 Experiments with EWC

In this section, we perform experiments with the proposed EWC-regularized scheduled sampling variant following the setup described in Section 4. To further demonstrate the effectiveness of our method, we also conduct experiments on two large-scale WMT’14 translation datasets, the English-German (EN-DE) and English-French (EN-FR). In addition, we report results using the sentence-oracle scheduled sampling variant proposed by Zhang et al. (2019), which conditions on model-generated prefixes obtained via beam search.

Automatic Evaluation Table 2 shows that by using the EWC-regularized scheduled sampling variant we obtain better BLEU scores across all translation datasets and neural architectures. Specifically, for the LSTM model, our method outperforms standard scheduled sampling by 0.62, 0.73, 0.60, and 0.42 BLEU points on DE-EN, VI-EN, EN-DE, and EN-FR, respectively. Similarly, for the Transformer model, it outperforms standard scheduled sampling by 0.54, 0.59, 0.52, and 0.47 BLEU points. Furthermore, we also observe that our method improves over the sentence oracle scheduled sampling variant. For the LSTM model, it outperforms sentence oracle by 0.46, 0.42, 0.31, and 0.34 BLEU points, and for the Transformer

Method	IWSLT'14		WMT'14	
	DE-EN	VI-EN	EN-DE	EN-FR
LSTM				
MLE	26.72±0.21	22.64±0.16	25.43±0.28	32.43±0.33
SS	27.35±0.32	23.58±0.33	26.08±0.22	32.76±0.26
+ EWC	27.97±0.24	24.31±0.27	26.68±0.25	33.18±0.31
Sent.	27.68±0.23	23.91±0.25	26.52±0.19	33.14±0.25
+ EWC	28.14±0.19	24.33±0.17	26.83±0.27	33.48±0.29
Transformer				
MLE	34.08±0.15	27.13±0.18	27.26±0.16	38.92±0.19
SS	34.65±0.19	27.57±0.26	27.92±0.11	39.86±0.24
+ EWC	35.19±0.21	28.16±0.26	28.44±0.15	40.33±0.17
Sent.	34.89±0.15	27.82±0.17	28.34±0.16	40.51±0.20
+ EWC	35.24±0.13	28.11±0.16	28.61±0.10	40.83±0.25

Table 2: Average BLEU scores and standard deviation on test sets for models trained on DE-EN, VI-EN, EN-DE, and EN-FR.

Model	Adequacy	Fluency
MLE	4.23	4.69
SS	4.35	4.78
+ EWC	4.41	4.82

Table 3: Human evaluation on adequacy and fluency for DE-EN.

model, by 0.35, 0.29, 0.27 and 0.32 BLEU points. Improvements are significant with $p < 0.01$.³

Human Evaluation To further examine the differences between MLE, scheduled sampling, and our proposed EWC-regularized variant, we conduct a human evaluation. We ask annotators to rate from 1 to 5 in terms of fluency and adequacy 100 randomly sampled DE-EN translations generated by the Transformer model. Table 3 shows the results of the human evaluation. Compared with the MLE baseline, both scheduled sampling and our EWC-regularized variant improve adequacy by 0.12 and 0.18, respectively. For fluency, scheduled sampling achieves a 0.09 score improvement, and our variant a 0.13 score increase.

Catastrophic Forgetting We investigate whether EWC addresses the catastrophic forgetting problem in models trained with scheduled sampling by repeating the teacher-forcing experiment discussed in Section 3. Table 4 shows less catastrophic for-

³We conduct statistical significance testing following Collins et al. (2005).


Method	DE-EN			VI-EN		
	MP	TF	Δ	MP	TF	Δ
LSTM						
MLE	26.72	28.67	+1.95	22.64	24.98	+2.34
SS	27.35	23.59	-3.76	23.58	18.92	-4.66
+ EWC	27.97	26.90	-1.07	24.31	22.93	-1.38
Transformer						
MLE	34.08	35.96	+1.88	27.13	29.14	+2.01
SS	34.65	29.84	-4.81	27.57	22.42	-5.15
+ EWC	35.19	33.17	-2.02	28.16	26.05	-2.11

Table 4: BLEU scores by conducting inference with two approaches, namely, (i) by conditioning using model-generated prefixes (MP), and (ii) with teacher-forcing (TF). Results for MLE and SS repeat numbers from Table 1.

getting occurring compared to the results obtained with standard scheduled sampling. In particular, on DE-EN the reduction in BLEU for standard scheduled sampling is 3.76 and 4.81, whereas with EWC the reduction is 1.07 and 2.02. Similarly, on VI-EN for standard scheduled sampling the BLEU score decreases by 4.66 and by 5.15, while for scheduled sampling with EWC by 1.38 and 2.11. This suggests that using EWC explicitly addresses the catastrophic forgetting phenomenon.

Performance at Different Output Lengths Figure 3 shows BLEU scores on the DE-EN test set for different output sequence lengths by grouping them into bins of width 20. We observe that the output quality degrades as the sequence length in-

LSTM



Transformer




Figure 2: Hyperparameter analysis using heatmaps. Better performance is achieved with different hyperparameter pairs over standard scheduled sampling (bottom row where $\lambda = 0$), thus showing the robustness of our proposed variant. We set λ to $\{0, 0.0005, 0.005, 0.1, 0.5\}$, k for the linear schedule to $\{-0.000003, -0.00003, -0.0003, 0.00003, 0.0003\}$, k for the exponential schedule to $\{0.999997, 0.99999, 0.9999, 0.999, 1\}$, and k for the sigmoid schedule to $\{6, 000, 8, 000, 10, 000, 12, 000, 15, 000\}$.




Figure 3: BLEU scores versus output sequence length (bins of width 20).

creases across all training objectives. In particular, MLE-trained models suffer the most from this due to exposure bias. Conversely, our proposed scheduled sampling variant improves the BLEU score across all bins (compared to MLE and standard scheduled sampling), even for long sequences.

Effect of λ Figure 2 shows the BLEU scores on the DE-EN test set under different values for

k (for the annealing schedule) and λ (for EWC). We observe that scheduled sampling with EWC outperforms standard scheduled sampling (bottom row where $\lambda = 0$) across all hyperparameter pairs. Thus we conclude that the introduction of λ does not increase the complexity of hyperparameter tuning as it improves the robustness of scheduled sampling across annealing schedules and their parameterizations. Huszár (2018) pointed out that EWC is an approximation of the joint optimization of the loss on two datasets, via encapsulating the information from the initial dataset in the model parameters that are fine-tuned. Scheduled sampling (without EWC) also aims at achieving the same goal, i.e. optimizing the parameters for generating outputs using both ground-truth prefixes and model-generated ones. In principle, this should be possible to achieve via a well-tuned annealing schedule, however, as Figure 2 shows, it is substantially easier to tune the hyperparameter controlling the EWC regularizer. Finally, in the theoretical analysis of SEARN which also uses an annealing schedule in a similar manner, Daumé III et al. (2009) argued that

the theoretically guaranteed annealing rate would be too slow to be applied in practice.

7 Related Work

Imitation Learning for NMT Scheduled sampling is an adaptation to recurrent neural networks (RNNs) and autoregressive models more broadly of DAGGER (Ross et al., 2011), a well-known imitation learning (IL) technique for mitigating exposure bias. Since Bengio et al. (2015), a number of research works have adapted standard IL algorithms (Daumé III et al., 2009; Ross and Bagnell, 2014; Chang et al., 2015) to autoregressive models to address exposure bias. SEARNN (Leblond et al., 2018) is a neural adaptation of SEARN (Daumé III et al., 2009) that computes a local loss for each generated token. Goyal et al. (2017) proposed a differentiable scheduled sampling variant to mitigate the credit assignment problem, while Xu et al. (2019) aimed to fix the time step alignment issue between the oracle and the model-generated tokens. Zhang et al. (2019) introduced a sequence-level scheduled sampling variant which conditions on model-generated prefixes with the highest BLEU score. A separate family of methods focus on using IL algorithms for non-autoregressive NMT (Wei et al., 2019; Gu et al., 2019; Welleck et al., 2019; Xu and Carpuat, 2020), and knowledge distillation (Liu et al., 2018; Lin et al., 2020).

Non-MLE Training Methods The discrepancy between MLE training and inference has prompted the development of many alternatives training algorithms. Approaches based on reinforcement learning (RL), such as Actor-Critic (Bahdanau et al., 2017) and MIXER (Ranzato et al., 2016) optimize the task-level loss directly. GOLD (Pang and He, 2020) uses off-policy RL to learn from human demonstrations through importance weighting (Hastings, 1970). Reward augmented maximum-likelihood (RAML) (Norouzi et al., 2016) maximizes the likelihood of sequences with respect to the exponentiated reward distribution. Minimum Risk Training (MRT) (Och, 2003; Smith and Eisner, 2006; Shen et al., 2016) optimizes model parameters by minimizing the expected loss directly with respect to the task-level loss. Other non-MLE training methods include energy-based models (Deng et al., 2020), beam-search optimization (BSO) (Wiseman and Rush, 2016), and unlikelihood training (Welleck et al., 2020).

Mitigating Catastrophic Forgetting Despite progress, various state-of-the-art neural architectures in NLP fail to retain previously learned knowledge due to catastrophic forgetting (Yogatama et al., 2019). To this end, many approaches have been proposed to overcome this limitation in a plethora of NLP tasks, such as NMT (Miceli Barone et al., 2017; Khayrallah et al., 2018; Saunders et al., 2019; Thompson et al., 2019), reading comprehension (Xu et al., 2020), fact verification (Thorne and Vlachos, 2020), and visual question answering (Perez et al., 2018; Greco et al., 2019). Another line of research works investigates techniques which address catastrophic forgetting during the fine-tuning stage of deep pretrained language models (Arora et al., 2019; Chronopoulou et al., 2019; Chen et al., 2019, 2020; Pilault et al., 2020).

8 Conclusion

In this work, we conduct systematic analyses to examine the strengths and weaknesses of scheduled sampling. By applying LRP to quantify the influence of the input sequence and the prefix during prediction, we show that models trained with scheduled sampling increase their reliance on the former to address exposure bias. However, we also demonstrate that as a side-effect, this leads to output degradation due to catastrophic forgetting when the prefix generated by the model is correct. Accordingly, we propose using EWC to better balance mitigating exposure bias with retaining performance. Our EWC-regularized scheduled sampling variant alleviates catastrophic forgetting and significantly improves performance over MLE and scheduled sampling baselines on four IWSLT’14 and WMT’14 translation datasets.

Limitations

Even though we show that it is easy to tune the hyperparameter controlling the EWC regularizer across different annealing schedules and their parameterization (Figure 2), in terms of runtime, we empirically observe that estimating the Fisher matrix diagonal for EWC, incurs additional training costs on top of running scheduled sampling. Furthermore, we did not assess whether the improvements we report are in any way dependent on properties of the language pairs we considered in our experiments. Thus, further testing is needed to establish whether our approach would transfer to other translation datasets.

Acknowledgements

The authors wish to thank Ferenc Huszár and Bill Byrne for their helpful comments and feedback. Michalis Korakakis is supported by the Cambridge Commonwealth, European and International Trust and the ESRC Doctoral Training Partnership for Social Sciences.

References

- Gaurav Arora, Afshin Rahimi, and Timothy Baldwin. 2019. [Does an LSTM forget more than a CNN? an empirical study of catastrophic forgetting in NLP](#). In *Proceedings of the The 17th Annual Workshop of the Australasian Language Technology Association*, pages 77–86, Sydney, Australia. Australasian Language Technology Association.
- Leila Arras, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2017. [Explaining recurrent neural network predictions in sentiment analysis](#). In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 159–168, Copenhagen, Denmark. Association for Computational Linguistics.
- Leila Arras, Ahmed Osman, Klaus-Robert Müller, and Wojciech Samek. 2019. [Evaluating recurrent neural network explanations](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 113–126, Florence, Italy. Association for Computational Linguistics.
- Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. 2015. [On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation](#). *PLOS ONE*, 10(7):1–46.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2017. [An actor-critic algorithm for sequence prediction](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. [Scheduled sampling for sequence prediction with recurrent neural networks](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1171–1179.
- Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, Klaus-Robert Müller, and Wojciech Samek. 2016. [Layer-wise relevance propagation for neural networks with local renormalization layers](#). In *Artificial Neural Networks and Machine Learning - ICANN 2016 - 25th International Conference on Artificial Neural Networks, Barcelona, Spain, September 6-9, 2016, Proceedings, Part II*, volume 9887 of *Lecture Notes in Computer Science*, pages 63–71. Springer.
- Kai-Wei Chang, Akshay Krishnamurthy, Alekh Agarwal, Hal Daumé III, and John Langford. 2015. [Learning to search better than your teacher](#). In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2058–2066. JMLR.org.
- Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. 2020. [Recall and learn: Fine-tuning deep pretrained language models with less forgetting](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7870–7881, Online. Association for Computational Linguistics.
- Xinyang Chen, Sinan Wang, Bo Fu, Mingsheng Long, and Jianmin Wang. 2019. [Catastrophic forgetting meets negative transfer: Batch spectral shrinkage for safe transfer learning](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 1906–1916.
- Alexandra Chronopoulou, Christos Baziotis, and Alexandros Potamianos. 2019. [An embarrassingly simple approach for transfer learning from pretrained language models](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2089–2095, Minneapolis, Minnesota. Association for Computational Linguistics.
- Michael Collins, Philipp Koehn, and Ivona Kučerová. 2005. [Clause restructuring for statistical machine translation](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 531–540, Ann Arbor, Michigan. Association for Computational Linguistics.
- Hal Daumé III, John Langford, and Daniel Marcu. 2009. [Search-based structured prediction](#). *Mach. Learn.*, 75(3):297–325.
- Yuntian Deng, Anton Bakhtin, Myle Ott, Arthur Szlam, and Marc’Aurelio Ranzato. 2020. [Residual energy-based models for text generation](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

- Yanzhuo Ding, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. [Visualizing and understanding neural machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1150–1159, Vancouver, Canada. Association for Computational Linguistics.
- Wanyu Du and Yangfeng Ji. 2019. [An empirical comparison on imitation learning and reinforcement learning for paraphrase generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6012–6018, Hong Kong, China. Association for Computational Linguistics.
- Robert M. French. 1999. [Catastrophic forgetting in connectionist networks](#). *Trends in Cognitive Sciences*, 3(4):128–135.
- Kartik Goyal, Chris Dyer, and Taylor Berg-Kirkpatrick. 2017. [Differentiable scheduled sampling for credit assignment](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 366–371, Vancouver, Canada. Association for Computational Linguistics.
- Claudio Greco, Barbara Plank, Raquel Fernández, and Raffaella Bernardi. 2019. [Psycholinguistics meets continual learning: Measuring catastrophic forgetting in visual question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3601–3605, Florence, Italy. Association for Computational Linguistics.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. [Levenshtein transformer](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 11179–11189.
- W. K. Hastings. 1970. [Monte carlo sampling methods using markov chains and their applications](#). *Biometrika*, 57(1):97–109.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9(8):1735–1780.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Ferenc Huszar. 2015. [How \(not\) to train your generative model: Scheduled sampling, likelihood, adversary?](#) *CoRR*, abs/1511.05101.
- Ferenc Huszár. 2018. [Note on the quadratic penalties in elastic weight consolidation](#). *Proceedings of the National Academy of Sciences*, 115(11):E2496–E2497.
- Huda Khayrallah, Brian Thompson, Kevin Duh, and Philipp Koehn. 2018. [Regularized training objective for continued training for domain adaptation in neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 36–44, Melbourne, Australia. Association for Computational Linguistics.
- James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2016. [Overcoming catastrophic forgetting in neural networks](#). *CoRR*, abs/1612.00796.
- Rémi Leblond, Jean-Baptiste Alayrac, Anton Osokin, and Simon Lacoste-Julien. 2018. [SEARNN: training rnns with global-local losses](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Haoran Li and Wei Lu. 2021. [Mixed cross entropy loss for neural machine translation](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 6425–6436. PMLR.
- Alexander Lin, Jeremy Wohlwend, Howard Chen, and Tao Lei. 2020. [Autoregressive knowledge distillation through imitation learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6121–6133, Online. Association for Computational Linguistics.
- Yijia Liu, Wanxiang Che, HuaiPeng Zhao, Bing Qin, and Ting Liu. 2018. [Distilling knowledge for search-based structured prediction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1393–1402, Melbourne, Australia. Association for Computational Linguistics.
- James Martens. 2014. [New perspectives on the natural gradient method](#). *CoRR*, abs/1412.1193.
- Antonio Valerio Miceli Barone, Barry Haddow, Ulrich Germann, and Rico Sennrich. 2017. [Regularization techniques for fine-tuning in neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1489–1494, Copenhagen, Denmark. Association for Computational Linguistics.
- Tsvetomila Mihaylova and André F. T. Martins. 2019. [Scheduled sampling for transformers](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 351–356, Florence, Italy. Association for Computational Linguistics.
- Mohammad Norouzi, Samy Bengio, Zhifeng Chen, Navdeep Jaitly, Mike Schuster, Yonghui Wu, and

- Dale Schuurmans. 2016. [Reward augmented maximum likelihood for neural structured prediction](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1723–1731.
- Franz Josef Och. 2003. [Minimum error rate training in statistical machine translation](#). In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 160–167, Sapporo, Japan. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [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)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Richard Yuanzhe Pang and He He. 2020. [Text generation by learning from off-policy demonstrations](#). *CoRR*, abs/2009.07839.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. 2018. [Film: Visual reasoning with a general conditioning layer](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 3942–3951. AAAI Press.
- Jonathan Pilault, Amine Elhattami, and Christopher J. Pal. 2020. [Conditionally adaptive multi-task learning: Improving transfer learning in NLP using fewer parameters & less data](#). *CoRR*, abs/2009.09139.
- Nina Poerner, Hinrich Schütze, and Benjamin Roth. 2018. [Evaluating neural network explanation methods using hybrid documents and morphosyntactic agreement](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 340–350, Melbourne, Australia. Association for Computational Linguistics.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. [Sequence level training with recurrent neural networks](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Stéphane Ross and J. Andrew Bagnell. 2014. [Reinforcement and imitation learning via interactive no-regret learning](#). *CoRR*, abs/1406.5979.
- Stéphane Ross, Geoffrey J. Gordon, and Drew Bagnell. 2011. [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, AISTATS 2011, Fort Lauderdale, USA, April 11-13, 2011*, volume 15 of *JMLR Proceedings*, pages 627–635. JMLR.org.
- Danielle Saunders, Felix Stahlberg, Adrià de Gispert, and Bill Byrne. 2019. [Domain adaptive inference for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 222–228, Florence, Italy. Association for Computational Linguistics.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [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)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. [Minimum risk training for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1683–1692, Berlin, Germany. Association for Computational Linguistics.
- David A. Smith and Jason Eisner. 2006. [Minimum risk annealing for training log-linear models](#). In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pages 787–794, Sydney, Australia. Association for Computational Linguistics.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. [Overcoming catastrophic forgetting during domain adaptation of neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2062–2068, Minneapolis, Minnesota. Association for Computational Linguistics.
- J. Thorne and A. Vlachos. 2020. [Elastic weight consolidation for better bias inoculation](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz

- Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Oriol Vinyals and Quoc V. Le. 2015. [A neural conversational model](#). *CoRR*, abs/1506.05869.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2021. [Analyzing the source and target contributions to predictions in neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1126–1140, Online. Association for Computational Linguistics.
- Chaojun Wang and Rico Sennrich. 2020. [On exposure bias, hallucination and domain shift in neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3544–3552, Online. Association for Computational Linguistics.
- Bingzhen Wei, Mingxuan Wang, Hao Zhou, Junyang Lin, and Xu Sun. 2019. [Imitation learning for non-autoregressive neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1304–1312, Florence, Italy. Association for Computational Linguistics.
- Sean Welleck, Kianté Brantley, Hal Daumé III, and Kyunghyun Cho. 2019. [Non-monotonic sequential text generation](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6716–6726. PMLR.
- Sean Welleck, Iliia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Ronald J. Williams and David Zipser. 1989. [A learning algorithm for continually running fully recurrent neural networks](#). *Neural Comput.*, 1(2):270–280.
- Sam Wiseman and Alexander M. Rush. 2016. [Sequence-to-sequence learning as beam-search optimization](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1296–1306, Austin, Texas. Association for Computational Linguistics.
- Lijun Wu, Xu Tan, Di He, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. [Beyond error propagation in neural machine translation: Characteristics of language also matter](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3602–3611, Brussels, Belgium. Association for Computational Linguistics.
- Weijia Xu and Marine Carpuat. 2020. [EDITOR: an edit-based transformer with repositioning for neural machine translation with soft lexical constraints](#). *CoRR*, abs/2011.06868.
- Weijia Xu, Xing Niu, and Marine Carpuat. 2019. [Differentiable sampling with flexible reference word order for neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2047–2053, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ying Xu, Xu Zhong, Antonio José Jimeno-Yepes, and Jey Han Lau. 2020. [Forget me not: Reducing catastrophic forgetting for domain adaptation in reading comprehension](#). In *2020 International Joint Conference on Neural Networks, IJCNN 2020, Glasgow, United Kingdom, July 19-24, 2020*, pages 1–8. IEEE.
- Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomás Kociský, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, and Phil Blunsom. 2019. [Learning and evaluating general linguistic intelligence](#). *CoRR*, abs/1901.11373.
- Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. [Bridging the gap between training and inference for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4334–4343, Florence, Italy. Association for Computational Linguistics.