

# A Deep Reinforced Model for Zero-Shot Cross-Lingual Summarization with Bilingual Semantic Similarity Rewards

Zi-Yi Dou Sachin Kumar Yulia Tsvetkov

Language Technologies Institute

Carnegie Mellon University

{zdou, sachink, ytsvetko}@cs.cmu.edu

## Abstract

Cross-lingual text summarization aims at generating a document summary in one language given input in another language. It is a practically important but under-explored task, primarily due to the dearth of available data. Existing methods resort to machine translation to synthesize training data, but such pipeline approaches suffer from error propagation. In this work, we propose an end-to-end cross-lingual text summarization model. The model uses reinforcement learning to directly optimize a bilingual semantic similarity metric between the summaries generated in a target language and gold summaries in a source language. We also introduce techniques to pre-train the model leveraging monolingual summarization and machine translation objectives. Experimental results in both English–Chinese and English–German cross-lingual summarization settings demonstrate the effectiveness of our methods. In addition, we find that reinforcement learning models with bilingual semantic similarity as rewards generate more fluent sentences than strong baselines.<sup>1</sup>

## 1 Introduction

Cross-lingual text summarization (XLS) is the task of compressing a long article in one language into a summary in a different language. Due to the dearth of training corpora, standard sequence-to-sequence approaches to summarization cannot be applied to this task. Traditional approaches to XLS thus follow a pipeline, for example, summarizing the article in the source language followed by translating the summary into the target language or vice-versa (Wan et al., 2010; Wan, 2011). Both of these approaches require separately trained summarization and translation models, and suffer from error propagation (Zhu et al., 2019).

<sup>1</sup>[https://github.com/zdou0830/crosslingual\\_summarization\\_semantic](https://github.com/zdou0830/crosslingual_summarization_semantic).




Figure 1: Along with minimizing the XLS cross-entropy loss  $L_{xls}$ , we also apply reinforcement learning to optimize the model by directly comparing the outputs with gold references in the source language.

Prior studies have attempted to train XLS models in an end-to-end fashion, through knowledge distillation from pre-trained machine translation (MT) or monolingual summarization models (Ayana et al., 2018; Duan et al., 2019), but these approaches have been only shown to work for short outputs. Alternatively, Zhu et al. (2019) proposed to automatically translate source-language summaries in the training set thereby generating pseudo-reference summaries in the target language. With this parallel dataset of source documents and target summaries, an end-to-end model is trained to simultaneously summarize and translate using a multi-task objective. Although the XLS model is trained end-to-end, it is trained on MT-generated reference translations and is still prone to compounding of translation and summarization errors.

In this work, we propose to train an end-to-end XLS model to directly generate target language summaries given the source articles by matching the semantics of the predictions with the semantics of the source language summaries. To achieve this, we use reinforcement learning (RL) with a bilingual semantic similarity metric as a reward (Wieting et al., 2019b). This metric is computed between the machine-generated summary in the target language and the gold summary in the source language. Additionally, to better initialize our XLS




Figure 2: Illustration of the supervised pre-training stage. The model is trained with cross-lingual summarization, machine translation and distillation objectives. The parameters of bottom layers of the decoders are shared across tasks.

model for RL, we propose a new multi-task pre-training objective based on machine translation and monolingual summarization to encode common information available from the two tasks. To enable the model to still differentiate between the two tasks, we add task specific tags to the input (Wu et al., 2016).

We evaluate our proposed method on English–Chinese and English–German XLS test sets. These test corpora are constructed by first using an MT-system to translate source summaries to the target language, and then being post-edited by human annotators. Experimental results demonstrate that just using our proposed pre-training method without fine-tuning with RL improves the best-performing baseline by up to 0.8 ROUGE-L points. Applying reinforcement learning yields further improvements in performance by up to 0.5 ROUGE-L points. Through extensive analyses and human evaluation, we show that when the bilingual semantic similarity reward is used, our model generates summaries that are more accurate, longer, more fluent, and more relevant than summaries generated by baselines.

## 2 Model

In this section, we describe the details of the task and our proposed approach.

### 2.1 Problem Description

We first formalize our task setup. We are given  $N$  articles and their summaries in the source language  $\{(x_{\text{src}}^{(1)}, y_{\text{src}}^{(1)}), \dots, (x_{\text{src}}^{(N)}, y_{\text{src}}^{(N)})\}$  as a training set. Our goal is to train a summarization model  $f(\cdot; \theta)$

which takes an article in the source language  $x_{\text{src}}$  as input and generates its summary in a pre-specified target language  $\hat{y}_{\text{tgt}} = f(x_{\text{src}}; \theta)$ . Here,  $\theta$  are the learnable parameters of  $f$ . During training, no gold summary  $y_{\text{tgt}}^{(i)}$  is available.

Our model consists of one encoder, denoted as  $E$ , which takes  $x_{\text{src}}$  as input and generates its vector representation  $\mathbf{h}$ .  $\mathbf{h}$  is fed as input to two decoders. The first decoder  $D_1$  predicts the summary in the target language ( $\hat{y}_{\text{tgt}}$ ) one token at a time. The second decoder  $D_2$  predicts the translation of the input text ( $\hat{v}_{\text{tgt}}$ ). While both  $D_1$  and  $D_2$  are used during training, only  $D_1$  is used for XLS at test time. Intuitively, we want the model to select parts of the input article which might be important for the summary and also translate them into the target language. To bias our model to encode this behavior, we propose the following algorithm for pre-training:

- Use a machine translation (MT) model to generate pseudo reference summaries ( $\tilde{y}_{\text{tgt}}$ ) by translating  $y_{\text{src}}$  to the target language. Then, translate  $\tilde{y}_{\text{tgt}}$  back to the source language using a target-to-source MT model and discard the examples with high reconstruction errors, which are measured with ROUGE (Lin, 2004) scores. The details of this step can be found in Zhu et al. (2019).
- Pre-train the model parameters  $\theta$  using a multi-task objective based on MT and monolingual summarization objectives with some simple yet effective techniques as described in §2.2.
- Further fine-tune the model using reinforcement learning with bilingual semantic similarity metric (Wieting et al., 2019b) as reward, which is described in §2.3.

### 2.2 Supervised Pre-Training Stage

Here, we describe the second step of our algorithm (Figure 2). The pre-training loss we use is a weighted combination of three objectives. Similarly to Zhu et al. (2019), we use an XLS pre-training objective and an MT pre-training objective as described below with some simple but effective improvements. We also introduce an additional objective based on distilling knowledge from a monolingual summarization model.

**XLS Pre-training Objective ( $L_{\text{xls}}$ )** This objective computes the cross-entropy loss of the predictions from  $D_1$ , considering the machine-generated

summaries in the target language,  $\tilde{y}_{\text{tgt}}^{(i)}$  as references, given  $x_{\text{src}}^{(i)}$  as inputs. Per sample, this loss can be formally written as:

$$L_{\text{xls}} = \sum_{j=1}^M \log p(\tilde{y}_{\text{tgt},j}^{(i)} | \tilde{y}_{\text{tgt},<j}^{(i)}, x_{\text{src}}^{(i)})$$

where  $M$  is the number of tokens in the summary  $i$ .

**Joint Training with Machine Translation** Zhu et al. (2019) argue that machine translation can be considered a special case of XLS with a compression ratio of 1:1. In line with Zhu et al. (2019), we train  $E$  and  $D_2$  as the encoder and decoder of a translation model using an MT parallel corpus  $\{(u_{\text{src}}^{(i)}, v_{\text{tgt}}^{(i)})\}$ . The goal of this step is to make the encoder have an inductive bias towards encoding information specific to translation. Similar to  $L_{\text{xls}}$ , the machine translation objective per training sample  $L_{\text{mt}}$  is:

$$L_{\text{mt}} = \sum_{j=1}^K \log p(v_{\text{tgt},j}^{(i)} | v_{\text{tgt},<j}^{(i)}, u_{\text{src}}^{(i)})$$

where  $K$  is the number of tokens in  $v_{\text{tgt}}^{(i)}$ . The  $L_{\text{xls}}$  and  $L_{\text{mt}}$  objectives are inspired by Zhu et al. (2019). We propose the following two enhancements to the model to leverage better the two objectives:

1. We share the parameters of bottom layers of the two decoders, namely  $D_1$  and  $D_2$ , to share common high-level representations while the parameters of the top layers more specialized to decoding are separately trained.
2. We append an artificial task tag  $\langle \text{SUM} \rangle$  (during XLS training) and  $\langle \text{TRANS} \rangle$  (during MT training) at the beginning of the input document to make the model aware of which kind of input it is dealing with.

We show in §4.1 that such simple modifications result in noticeable performance improvements.

**Knowledge Distillation from Monolingual Summarization** To bias the encoder to identify sentences which can be most relevant to the summary, first, we use an extractive monolingual summarization method to predict the probability  $q_i$  of each sentence or keyword in the input article being relevant to the summary. We then distill knowledge from this model into the encoder  $E$  by making it predict these probabilities. Concretely, we append

an additional output layer to the encoder of our model and it predicts the probability  $p_i$  of including the  $i$ -th sentence or word in the summary. The objective is to minimize the difference between  $p_i$  and  $q_i$ . We use the following loss (for each sample) for the model encoder:<sup>2</sup>

$$L_{\text{dis}} = \frac{1}{L} \sum_{j=1}^L (\log q_j - \log p_j)^2, \quad (1)$$

where  $L$  is the number of sentences or keywords in each article.

Our final pre-training objective during the supervised pre-training stage is:

$$L_{\text{sup}} = L_{\text{xls}} + L_{\text{mt}} + \lambda L_{\text{dis}} \quad (2)$$

where  $\lambda$  is a hyper-parameter and is set to 10 in our experiments. Training with  $L_{\text{mt}}$  requires an MT parallel corpus whereas the other two objectives utilize the cross-lingual summarization dataset. Pre-training algorithm alternates between the two parts of the objective using mini-batches from the two datasets as follows until convergence:

1. Sample a minibatch from the MT corpus  $\{(u_{\text{src}}^{(i)}, v_{\text{tgt}}^{(i)})\}$  and train the parameters of  $E$  and  $D_2$  with  $L_{\text{mt}}$ .
2. Sample a minibatch from the XLS corpus,  $\{(x_{\text{src}}^{(i)}, \tilde{y}_{\text{tgt}}^{(i)})\}$  and train the parameters of  $E$  and  $D_1$  with  $L_{\text{xls}} + \lambda L_{\text{dis}}$ .

### 2.3 Reinforcement Learning Stage

For XLS, the target language reference summaries ( $\tilde{y}_{\text{tgt}}$ ) used during pre-training are automatically generated with MT models and thus they may contain errors. In this section, we describe how we further fine-tune the model using only human-generated source language summaries ( $y_{\text{src}}$ ) with reinforcement learning (RL). Specifically, we first feed the article  $x_{\text{src}}$  as an input to the encoder  $E$ , and generate the target language summary  $\hat{y}_{\text{tgt}}$  using  $D_1$ . We then compute a cross-lingual similarity metric between  $\hat{y}_{\text{tgt}}$  and  $y_{\text{src}}$  and use it as a reward to fine-tune  $E$  and  $D_1$ .

Following Paulus et al. (2018), we adopt two different strategies to generate  $\hat{y}_{\text{tgt}}$  at each training iteration, (a)  $\hat{y}_{\text{tgt}}^s$  obtained by sampling from the softmax layer at each decoding step, and (b)  $\hat{y}_{\text{tgt}}^g$

<sup>2</sup>We also experimented with a common distillation objective based on minimizing KL divergence,  $\frac{1}{n} \sum_{i=1}^n q_i \log p_i$ , but it did not perform as well.

obtained by greedy decoding. The RL objective per sample is given by:

$$L_{\text{rl}} = (r(\hat{y}_{\text{tgt}}^g) - r(\hat{y}_{\text{tgt}}^s)) \sum_{j=1}^M \log p(\hat{y}_{\text{tgt},i}^s | \hat{y}_{\text{tgt},<j}^s, x), \quad (3)$$

where  $r(\cdot)$  is the reward function. To fine-tune the model, we use the following hybrid training objective:  $\gamma L_{\text{rl}} + (1 - \gamma)L_{\text{xls}}$ , where  $\gamma$  is a scaling factor.

We train a cross-lingual similarity model (XSIM) with the best performing model in [Wieting et al. \(2019b\)](#). This model is trained using an MT parallel corpus. Using XSIM, we obtain sentence representations for both  $\hat{y}_{\text{tgt}}$  and  $y_{\text{src}}$  and treat the cosine similarity between the two representations as the reward  $r(\cdot)$ .

### 3 Experimental Setup

#### 3.1 Datasets

We evaluate our models on English–Chinese and English–German article-summary datasets. The English–Chinese dataset is created by [Zhu et al. \(2019\)](#), constructed using the CNN/DailyMail monolingual summarization corpus ([Hermann et al., 2015](#)). The training, validation and test sets consist of about 364K, 3K and 3K samples, respectively. The English–German dataset is our contribution, constructed from the Gigaword dataset ([Rush et al., 2015](#)). We sample 2.48M training, 2K validation and 2K test samples from the dataset. Pseudo-parallel corpora for both language pairs are constructed by translating the summaries to the target language (and filtered after back-translation; see §2). This is done for training, validation as well as test sets. These two pseudo-parallel training sets are used for pre-training with  $L_{\text{xls}}$ . Translated Chinese and German summaries of the test articles are then post-edited by human annotators to construct the test set for evaluating XLS. We refer the readers to ([Zhu et al., 2019](#)) for more details. For the English–Chinese dataset, we use word-based segmentation for the source (articles in English) and character-based segmentation for the target (summaries in Chinese) as in ([Zhu et al., 2019](#)). For the English–German dataset, byte-pair encoding is used ([Sennrich et al., 2016](#)) with 60K merge operations. For machine translation and training the XSIM model, we sub-sample 5M sentences from the WMT2017 Chinese–English and

WMT2014 German–English training dataset ([Bojar et al., 2014, 2017](#)).

#### 3.2 Implementation Details

We use the Transformer-BASE model ([Vaswani et al., 2017](#)) as the underlying architecture for our model ( $E$ ,  $D_1$ ,  $D_2$ , extractive summarization model for distillation and baselines). We refer the reader to [Vaswani et al. \(2017\)](#) for hyperparameter details. In the input article, a special token  $\langle \text{SEP} \rangle$  is added at the beginning of each sentence to mark sentence boundaries. For the CNN/DailyMail corpus, the monolingual extractive summarization used in the distillation objective has the same architecture as the encoder  $E$  and is trained the CNN/DailyMail corpus constructed by ([Liu and Lapata, 2019](#)). To train the encoder with  $L_{\text{dis}}$ , we take the final hidden representation of each  $\langle \text{SEP} \rangle$  token and apply a 2-layer feed-forward network with ReLU activation in the middle layer and sigmoid at the final layer to get  $q_i$  for each sentence  $i$  (see §2.2).

For the Gigaword dataset, because the inputs and outputs are typically short, we choose keywords rather than sentences as the prediction unit. Specifically, we first use TextRank ([Mihalcea and Tarau, 2004](#)) to extract all the keywords from the source document. Then, for each keyword  $i$  that appears in the target summary, the gold label  $q_i$  in equation 1 is assigned to 1, and  $q_i$  is assigned to 0 for keywords that do not appear in the target side.

We share the parameters of the bottom four layers of the decoder in the multi-task setting. We use the TRIGRAM model in ([Wieting et al., 2019b,a](#)) to measure the cross-lingual sentence semantic similarities. As pointed out in §2, after the pre-training stage, we only use  $D_1$  for XLS. The final results are obtained using only  $E$  and  $D_1$ . We use two metrics for evaluating the performance of models: ROUGE (1, 2 and L) ([Lin, 2004](#)) and XSIM ([Wieting et al., 2019b](#)).

Following [Paulus et al. \(2018\)](#), we select  $\gamma$  in equation 3 to 0.998 for the Gigaword Corpus and  $\gamma = 0.9984$  for the CNN/DailyMail dataset.

#### 3.3 Baselines

We compare our proposed method with the following baselines:

**Pipeline Approaches** We report results of summarize-then-translate (SUM-TRAN) and

Method	English–Chinese				English–German			
	ROUGE-1	ROUGE-2	ROUGE-L	XSIM	ROUGE-1	ROUGE-2	ROUGE-L	XSIM
<i>Pipeline-Based Methods</i>								
TRAN-SUM (Zhu et al., 2019)	28.19	11.40	25.77	-	-	-	-	-
SUM-TRAN (Zhu et al., 2019)	32.17	13.85	29.43	-	-	-	-	-
<i>End-to-End Training Methods</i>								
MLE-XLS	37.38	17.96	33.85	45.17	23.06	8.40	21.28	43.41
MLE-XLS+MT (Zhu et al., 2019)	40.23	22.32	36.59	-	-	-	-	-
MLE-XLS+MT (Reimplemented)	41.25	22.40	37.93	48.77	36.83	17.62	35.54	49.21
MLE-XLS+MT+DIS	42.19*	22.91*	38.74*	49.20*	37.74*	18.40*	36.34*	49.53*
RL-ROUGE	42.51*	22.96	38.98*	49.65	38.32*	18.46	36.86*	49.66
RL-XSIM	<b>42.83*</b>	<b>23.30*</b>	<b>39.29*</b>	<b>50.85*</b>	<b>38.69*</b>	<b>18.76*</b>	<b>37.20*</b>	<b>50.17*</b>
RL-ROUGE+XSIM	42.49	23.29*	38.95	49.88*	38.19*	18.17	36.72*	49.64

Table 1: Performance of different models. The highest scores are in **bold** and statistical significance compared with the best baseline is indicated with \* ( $p < 0.05$ , computed using *compare-mt* (Neubig et al., 2019)). XSIM is computed between the target language system outputs and the source language reference summaries.

translate-then-summarize (TRAN-SUM) pipelines. These results are taken from Zhu et al. (2019).

**MLE-XLS** We pre-train  $E$  and  $D_1$  with only  $L_{xls}$  without any fine-tuning.

**MLE-XLS+MT** We pre-train  $E$ ,  $D_1$  and  $D_2$  with  $L_{xls} + L_{mt}$  without using  $L_{dis}$ . This is the best performing model in (Zhu et al., 2019). We show their reported results as well as results from our re-implementation.

**MLE-XLS+MT+DIS** We pre-train the model using (2) without fine-tuning with RL. We also share the decoder layers and add task specific tags to the input as described in §2.2.

**RL-ROUGE** Using ROUGE score as a reward function has been shown to improve summarization quality for monolingual summarization models (Paulus et al., 2018). In this baseline, we fine-tune the pre-trained model in the above baseline using ROUGE-L as a reward instead of the proposed XSIM. The ROUGE-L score is computed between the output of  $D_1$  and the machine-generated summary  $\tilde{y}_{tgt}$ .

**RL-ROUGE+XSIM** Here, we use the average of ROUGE score and XSIM score as a reward function to fine-tune the pre-trained model (MLE-XLS+MT+DIS).

## 4 Results

The main results of our experiments are summarized in Table 1. Pipeline approaches, as expected, show the weakest performance, lagging behind even the weakest end-to-end approach by more than 5 ROUGE-L points. TRAN-SUM performs even worse than SUM-TRAN, likely because the

Method	ROUGE-1	ROUGE-2	ROUGE-L
MLE-XLS	37.38	17.96	33.85
+EXTRACT	39.19	19.58	35.68
+DIS	<b>40.46</b>	<b>20.47</b>	<b>36.93</b>
MLE-XLS+MT	41.25	22.40	37.93
+EXTRACT	40.04	21.58	36.74
+DIS	<b>41.68</b>	<b>22.48</b>	<b>38.29</b>

Table 2: Effect of using hard (EXTRACT) vs soft (DIS) extraction of summary sentences from the input article

translation model is trained on sentences and not long articles. First translating the article with many sentences introduces way more errors than translating a short summary with fewer sentences would. Using just our pre-training method as described in 2.2 (MLE-XLS+MT+DIS), our proposed model outperforms the strongest baseline (MLE-XLS+MT) in both ROUGE-L (by 0.8) and XSIM (by 0.5). Applying reinforcement learning to fine-tune the model with both ROUGE (RL-ROUGE), XSIM (RL-XSIM) or their mean (RL-ROUGE+XSIM) as rewards results in further improvements. Our proposed method, RL-XSIM performs the best overall, indicating the importance of using cross-lingual similarity as a reward function. RL-ROUGE uses a machine-generated reference to compute the rewards since target language summaries are unavailable, which might be a reason for its worse performance.

### 4.1 Analysis

In this section, we conduct experiments on the CNN/DailyMail dataset to establish the importance of every part of the proposed method and gain further insights into our model.

**Soft Distillation vs. Hard Extraction** The results in table 1 already show that adding the knowl-

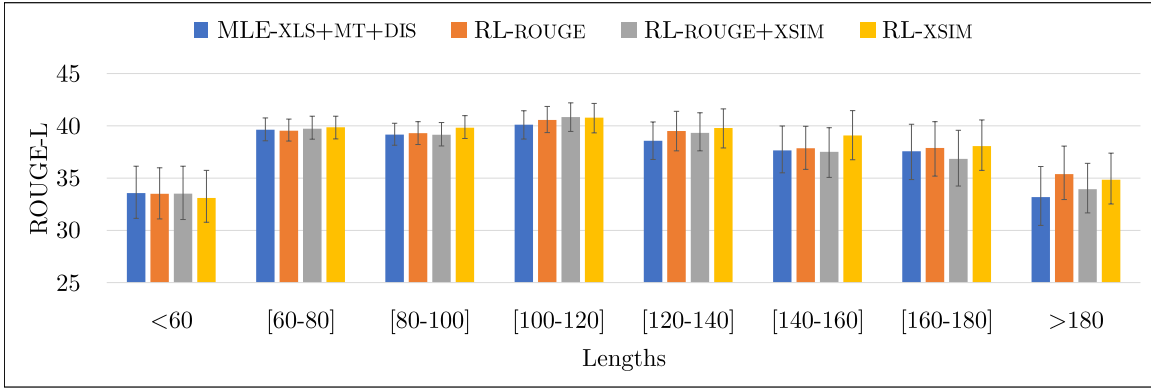


Figure 3: Reinforcement learning can make the model better at generating long summaries. We use the *compare-mt* tool (Neubig et al., 2019) to get these statistics.

Method	ROUGE-1	ROUGE-2	ROUGE-L
MLE+XLS+MT	41.25	22.40	37.93
+SHARE	41.36	22.43	37.95
+TAG	41.45	22.47	38.04

Table 3: Effect of sharing decoder layers and adding task-specific tags

edge distillation objective  $L_{\text{dis}}$  to the pre-training leads to an improvement in performance. The intuition behind using  $L_{\text{dis}}$  is to bias the model to (softly) select sentences in the input article that might be important for the summary. Here, we replace this soft selection with a hard selection. That is, using the monolingual extractive summarization model (as described in §3.2), we extract top 5 sentences from the input article and use them as the input to the encoder instead. We compare this method with  $L_{\text{dis}}$  as shown in Table 2. With just MLE-XLS as the pre-training objective, EXTRACT shows improvement (albeit with lower overall numbers) in performance but leads to a decrease in performance of MLE-XLS+MT. On the other hand, using the distillation objective helps in both cases.

### Effect of the Sharing and Tagging Techniques

In Table 3, we demonstrate that introducing simple enhancements like sharing the lower-layers of the decoder (share) and adding task-specific tags (tags) during multi-task pre-training also helps in improving the performance while at the same using fewer parameters and hence a smaller memory footprint.

**Effect of Summary Lengths** Next, we study how different baselines and our model performs with respect to generating summaries (in Chinese) of different lengths, in terms of number of characters. As shown in Figure 3, after fine-tuning the model with RL, our proposed model becomes

better at generating longer summaries than the one with only pre-training (referred to as MLE-XLS+MT+DIS in the figure) with RL-XSIM performing the best in most cases. We posit that this improvement is due to RL based fine-tuning reducing the problem of exposure bias introduced during teacher-forced pre-training, which especially helps longer generations.

**Human Evaluation** In addition to automatic evaluation, which can sometimes be misleading, we perform human evaluation of summaries generated by our models. We randomly sample 50 pairs of the model outputs from the test set and ask three human evaluators to compare the pre-trained supervised learning model and reinforcement learning models in terms of *relevance* and *fluency*. For each pair, the evaluators are asked to pick one out of: first model (MLE-XLS+MT+DIS; lose), second model (RL models; win) or say that they prefer both or neither (tie). The results are summarized in table 4. We observe that the outputs of model trained with ROUGE-L rewards are more favored than the ones generated by only pre-trained model in terms of relevance but not fluency. This is likely because the RL-ROUGE model is trained using machine-generated summaries as references which might lack fluency. Figure 4 displays one such example. On the other hand, cross-lingual semantic similarity as a reward results in generations which are more favored both in terms of relevance and fluency.

## 5 Related Work

Most previous work on cross-lingual text summarization utilize either the summarize-then-translate or translate-then-summarize pipeline (Wan et al., 2010; Wan, 2011; Yao et al., 2015; Ouyang et al.,

Output	
Ref	<p>市议会周三投票通过了一项提高购买香烟法定年龄的法案。纽约是美国最大的将购买年龄提高到18岁以上的城市。这项法律预计明年初生效。</p> <p>A bill to raise the legal age to buy cigarettes was voted into law Wednesday by the City Council. New York is the largest US city to raise the purchase age above the federal limit of 18-years-old. The law is expected to go into effect early next year.</p>
Sup	<p>纽约成为美国最大的购买年龄。纽约不是第一个提高法定饮酒年龄的城市。</p> <p>New York has become the largest purchase age in the United States. New York is not the first city to raise the legal drinking age.</p>
RL-ROUGE	<p>纽约成为美国最大的购买年龄，法律年龄从18岁提高到21岁。市议会周三批准了一项法案，将烟草购买年龄从18岁提高到21岁。纽约不是第一个提高法定饮酒年龄的城市。</p> <p>New York has become the largest purchase age in the United States, and the legal age has increased from 18 to 21. The City Council approved a law on Wednesday to increase the age of tobacco purchases from 18 to 21. New York is not the first city to raise the legal drinking age.</p>
RL-XSIM	<p>纽约成为美国最大的购买香烟的城市。市议会周三批准了一项法案，将烟草购买年龄从18岁提高到21岁。纽约不是第一个提高法定饮酒年龄的城市。</p> <p>New York has become the largest city in the United States for buying cigarettes. The City Council approved a law on Wednesday to increase the age of tobacco purchases from 18 to 21. New York is not the first city to raise the legal drinking age.</p>

Figure 4: Example outputs. The bilingual semantic similarity rewards can make the output more fluent than using ROUGE-L as rewards. “Sup” refers to the MLE-XLS+MT+DIS baseline.

Metric	Model v. MLE	Win (%)	Lose (%)	Tie (%)
Relevance	RL-ROUGE	25.3	15.3	59.3
	RL-XSIM	36.0	31.3	32.7
Fluency	RL-ROUGE	13.3	17.3	69.3
	RL-XSIM	37.3	28.7	34.0

Table 4: Results showing preferences of human evaluators towards the summaries generated by the mentioned RL methods vs ones from the pre-trained model (MLE-XLS+MT+DIS referred in short as MLE)

2019). These methods suffer from error propagation and we have demonstrated their sub-optimal performance in our experiments. Recently, there has been some work on training models for this task in an end-to-end fashion (Ayana et al., 2018; Duan et al., 2019; Zhu et al., 2019), but these models are trained with cross-entropy using machine-generated summaries as references which have already lost some information in the translation step.

Prior work in monolingual summarization have explored hybrid extractive and abstractive summarization objectives which inspires our distillation objective (Gehrmann et al., 2018; Hsu et al., 2018; Chen and Bansal, 2018). This line of research mainly focus on either compressing sentences extracted by a pre-trained model or biasing the prediction towards certain words.

Language generation models trained with cross-entropy using teacher-forcing suffer from exposure bias and a mismatch between training and evaluation objective. To solve these issues, using reinforcement learning to fine-tune such models

have been explored for monolingual summarization where ROUGE rewards is typically used (Paulus et al., 2018; Liu et al., 2018; Pasunuru and Bansal, 2018). Other rewards such as BERT score (Zhang et al., 2019) have also been explored (Li et al., 2019). Computing such rewards requires access to the gold summaries which are typically unavailable for cross-lingual summarization. This work is the first to explore using cross-lingual similarity as a reward to work around this issue.

## 6 Conclusion

In this work, we propose to use reinforcement learning with a bilingual semantic similarity metric as rewards for cross-lingual document summarization. We demonstrate the effectiveness of the proposed approach in a resource-deficient setting, where target language gold summaries are not available. We also propose simple strategies to better initialize the model towards reinforcement learning by leveraging machine translation and monolingual summarization. In future work, we plan to explore methods for stabilizing reinforcement learning as well to extend our methods to other datasets and tasks, such as using the bilingual similarity metric as a reward to improve the quality of machine translation.

## Acknowledgements

We are grateful to anonymous reviewers for their helpful suggestions and Chunting Zhou, Shuyan

Zhou for proofreading the paper. We also thank Ruihan Zhai, Zhi-Hao Zhou for the help with human evaluation and Anurag Katakhar for post-editing the dataset. This material is based upon work supported by NSF grants IIS1812327 and by Amazon MLRA award. We also thank Amazon for providing GPU credits.

## References

- Ayana, Shi-qi Shen, Yun Chen, Cheng Yang, Zhi-yuan Liu, and Mao-song Sun. 2018. Zero-shot cross-lingual neural headline generation. *IEEE/ACM Transactions on Audio, Speech and Language Processing*.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, et al. 2014. Findings of the 2014 workshop on statistical machine translation. In *Proc. WMT*.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, et al. 2017. Findings of the 2017 conference on machine translation. In *Proc. WMT*.
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. In *Proc. ACL*.
- Xiangyu Duan, Mingming Yin, Min Zhang, Boxing Chen, and Weihua Luo. 2019. Zero-shot cross-lingual abstractive sentence summarization through teaching generation and attention. In *Proc. ACL*.
- Sebastian Gehrmann, Yuntian Deng, and Alexander Rush. 2018. Bottom-up abstractive summarization. In *Proc. EMNLP*.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proc. NeurIPS*.
- Wan-Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun. 2018. A unified model for extractive and abstractive summarization using inconsistency loss. In *Proc. ACL*.
- Siyao Li, Deren Lei, Pengda Qin, and William Yang Wang. 2019. Deep reinforcement learning with distributional semantic rewards for abstractive summarization. In *Proc. EMNLP*.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*.
- Linqing Liu, Yao Lu, Min Yang, Qiang Qu, Jia Zhu, and Hongyan Li. 2018. Generative adversarial network for abstractive text summarization. In *Proc. AAAI*.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. In *Proc. EMNLP*.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proc. EMNLP*.
- Graham Neubig, Zi-Yi Dou, Junjie Hu, Paul Michel, Danish Pruthi, and Xinyi Wang. 2019. compare-mt: A tool for holistic comparison of language generation systems. In *Proc. NAACL Demo*.
- Jessica Ouyang, Boya Song, and Kathleen McKeown. 2019. A robust abstractive system for cross-lingual summarization. In *Proc. NAACL*.
- Ramakanth Pasunuru and Mohit Bansal. 2018. Multi-reward reinforced summarization with saliency and entailment. In *Proc. NAACL*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. In *Proc. ICLR*.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proc. EMNLP*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proc. ACL*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proc. NeurIPS*.
- Xiaojun Wan. 2011. Using bilingual information for cross-language document summarization. In *Proc. ACL*.
- Xiaojun Wan, Huiying Li, and Jianguo Xiao. 2010. Cross-language document summarization based on machine translation quality prediction. In *Proc. ACL*.
- John Wieting, Taylor Berg-Kirkpatrick, Kevin Gimpel, and Graham Neubig. 2019a. Beyond bleu: Training neural machine translation with semantic similarity. In *Proc. ACL*.
- John Wieting, Kevin Gimpel, Graham Neubig, and Taylor Berg-Kirkpatrick. 2019b. Simple and effective paraphrastic similarity from parallel translations. In *Proc. ACL*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Jin-ge Yao, Xiaojun Wan, and Jianguo Xiao. 2015. Phrase-based compressive cross-language summarization. In *Proc. EMNLP*.



Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Junnan Zhu, Qian Wang, Yining Wang, Yu Zhou, Jiajun Zhang, Shaonan Wang<sup>1</sup>, and Chengqing Zong. 2019. NCLS: Neural cross-lingual summarization. In *Proc. EMNLP*.