

Robust Unsupervised Neural Machine Translation with Adversarial Denoising Training

Haipeng Sun^{1*}, Rui Wang², Kehai Chen², Xugang Lu²,
Masao Utiyama², Eiichiro Sumita², and Tiejun Zhao^{1†}

¹Harbin Institute of Technology, Harbin, China

²National Institute of Information and Communications Technology (NICT), Kyoto, Japan
sunhaipeng.nlp@gmail.com, tjzhao@hit.edu.cn
{wangrui, khchen, xugang.lu, mutiyama, eiichiro.sumita}@nict.go.jp

Abstract

Unsupervised neural machine translation (UNMT) has recently attracted great interest in the machine translation community. The main advantage of the UNMT lies in its easy collection of required large training text sentences while with only a slightly worse performance than supervised neural machine translation which requires expensive annotated translation pairs on some translation tasks. In most studies, the UNMT is trained with clean data without considering its robustness to the noisy data. However, in real-world scenarios, there usually exists noise in the collected input sentences which degrades the performance of the translation system since the UNMT is sensitive to the small perturbations of the input sentences. In this paper, we first time explicitly take the noisy data into consideration to improve the robustness of the UNMT based systems. First of all, we clearly defined two types of noises in training sentences, i.e., word noise and word order noise, and empirically investigate its effect in the UNMT, then we propose adversarial training methods with denoising process in the UNMT. Experimental results on several language pairs show that our proposed methods substantially improved the robustness of the conventional UNMT systems in noisy scenarios.

1 Introduction

Recently, unsupervised neural machine translation (UNMT) has attracted great interest in the machine translation community (Artetxe et al., 2018; Lample et al., 2018a; Yang et al., 2018; Lample et al., 2018b; Sun et al., 2019; Sun et al., 2020b). Typically, UNMT relies solely on monolingual corpora rather than bilingual parallel data in supervised neural machine translation (SNMT) to model translations between the source language and target language and has achieved remarkable results on several translation tasks (Conneau and Lample, 2019). However, previous work only focus on how to build state-of-the-art UNMT systems on the clean data and ignore the robustness of UNMT on the noisy data. In the real-world scenario, there often exists noises or perturbations in the input sentences, for example, word character misspelling, replacement, or word position misordering, etc. The translation model is sensitive to these perturbations, leading to various errors even the perturbations are small. The existing neural translation system, which lacks of robustness, is difficult to be widely applied to the noisy-data scenario (denoted as noisy scenario in the following sections). Therefore, the robustness of neural translation system is not only worthy of being studied, but also very essential in the real-world scenarios.

The robustness of SNMT (Belinkov and Bisk, 2018; Cheng et al., 2018; Cheng et al., 2019; Karpukhin et al., 2019) has been well-studied. However, most previous work only focus on the effect of the word substitution for translation performance, and ignore the effect of word order for translation performance. Moreover, the noisy robustness of UNMT is much more difficult since the noisy input data may be relieved in some degree by the SNMT due to its supervised check in training. Currently, there is no study considering the noisy robustness of the UNMT. In this paper, we first define two types of noises which cover the noise types mentioned above, i.e., word noise and word order noise. Then we empirically

* Most of this work was done when Haipeng Sun was an internship research fellow at NICT.

† Corresponding author.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

investigate the performance of UNMT in these noisy scenarios. Our empirical results show that the performance of UNMT model decreased substantially, regardless of any noisy scenario. To improve the robustness, we proposed adversarial training methods to alleviate the poor performance in these noisy scenarios. To the best of our knowledge, this paper is the first work to explore the robustness of UNMT. Experimental results on several language pairs show that the proposed strategies substantially outperform conventional UNMT systems in the noisy scenarios.

Our main contributions are summarized as follows:

- We explicitly defined two types of noises, i.e., word noise and word order noise, and empirically investigate the performance of UNMT in the noisy scenarios.
- We propose adversarial training methods with denoising process in UNMT training to improve the robustness of the UNMT systems.
- Our proposed adversarial training methods achieve improvement up to 10 BLEU scores in the noisy scenarios, compared with an UNMT based baseline system.

2 Background of UNMT

There are four primary components of the state-of-the-art UNMT (Conneau and Lample, 2019): cross-lingual language model pre-training, denoising auto-encoder, back-translation, and sharing latent representations. Consider monolingual data $\{X_i\}$ in language L_1 and $\{Y_i\}$ in another language L_2 . $|X|$ and $|Y|$ are the number of sentences in monolingual corpora $\{X_i\}$ and $\{Y_i\}$ respectively. The encoders and decoders of L_1, L_2 are trained through denoising and back-translation. The objective function \mathcal{L}_{all} of the entire UNMT model would be optimized as:

$$\mathcal{L}_{all} = \mathcal{L}_D + \mathcal{L}_B, \quad (1)$$

where \mathcal{L}_D is the objective function for denoising, and \mathcal{L}_B is the objective function for back-translation.

Cross-lingual language model pre-training: It aims at building a universal cross-lingual encoder that can encode two monolingual sentences into a shared embedding space. The pre-trained cross-lingual encoder is then used to initialize the UNMT model.

Denoising auto-encoder: In contrast with the normal auto-encoder, denoising auto-encoder (Vincent et al., 2010) could improve the model learning ability by introducing noise in the form of random token deleting and swapping in this input sentence. The denoising auto-encoder, which encodes a noisy version and reconstructs it with the decoder in the same language, acts as a language model during UNMT training. It is optimized by minimizing the objective function:

$$\mathcal{L}_D = \sum_{i=1}^{|X|} -\log P_{L_1 \rightarrow L_1}(X_i | C(X_i)) + \sum_{i=1}^{|Y|} -\log P_{L_2 \rightarrow L_2}(Y_i | C(Y_i)), \quad (2)$$

where $\{C(X_i)\}$ and $\{C(Y_i)\}$ are noisy sentences. $P_{L_1 \rightarrow L_1}$ and $P_{L_2 \rightarrow L_2}$ denote the reconstruction probability in the language L_1 and L_2 , respectively.

Back-translation: It (Sennrich et al., 2016a) is adapted to train a translation system across different languages based on monolingual corpora. The pseudo-parallel sentence pairs $\{(Y_M(X_i), X_i)\}$ and $\{(X_M(Y_i), Y_i)\}$ produced by the model at the previous iteration would be used to train the new translation model. The UNMT model would be improved through iterative back-translation. Therefore, the back-translation probability would be optimized by minimizing

$$\mathcal{L}_B = \sum_{i=1}^{|X|} -\log P_{L_2 \rightarrow L_1}(X_i | Y_M(X_i)) + \sum_{i=1}^{|Y|} -\log P_{L_1 \rightarrow L_2}(Y_i | X_M(Y_i)), \quad (3)$$

where $P_{L_1 \rightarrow L_2}$ and $P_{L_2 \rightarrow L_1}$ denote the translation probability across the two languages.

Sharing latent representations: The same vocabulary is used for both languages. Encoders and decoders are shared for both languages, to help UNMT model to translate more fluently with synthetic source sentences benefiting from denoising training.

3 Preliminary Experiments toward Noisy Input to UNMT System

In this section, we first introduce the two primary types of noises in the corpus of UNMT. Then we empirically analyze the effect of these noises on UNMT.

3.1 Synthetic Noise Generation

A few studies of the SNMT robustness (Belinkov and Bisk, 2018; Karpukhin et al., 2019) focus on character-level noise, which affects the spelling of a single word. In this paper, we study the word-level noise, which affects the meaning of a word in one sentence. We refer to this noise as word noise in this paper. Moreover, we study the sentence-level noise, which affects the order of a whole sentence. We refer to this noise as word order noise in this paper.

Word Noise: We replace every word in the source sentence by an arbitrary word with a probability a . A larger probability a results in that more words are replaced by arbitrary words.

Word Order Noise: Motivated by the input shuffling strategy (Lample et al., 2018a), we apply a random permutation γ to the source sentence to change the word order of the original sentence, to meet the condition:

$$|\gamma(i) - i| \leq b, \forall i \in \{1, n\}, \quad (4)$$

where n denotes the length of the source sentence, and b is a hyper-parameter to control the magnitude of the word order adjustment. A larger b results in worse order in the source sentence.

To generate a random permutation verifying the above condition for a sentence of size n , we generate a random array Q of size n :

$$Q_i = i + U(0, b), \quad (5)$$

where U is the uniform distribution function in the range from 0 to b . Then, γ is defined to be the permutation that sorts the array Q . We apply this permutation to adjust word order in order to generate synthetic noise. Note that the order will be changed only when $b > 1$.

3.2 Noisy Scenario

As the cross-lingual language model pre-training, which needs large-scale additional monolingual data, the UNMT system performed comparable with SNMT system which only relies on parallel data. To investigate the performance of UNMT in the noisy scenario, we empirically choose English (En)–French (Fr) as the language pair to do stimulated experiments. The detailed experimental settings for UNMT are given in Section 5. The training data is clean and the test data contains synthetic noises described in Section 3.1.

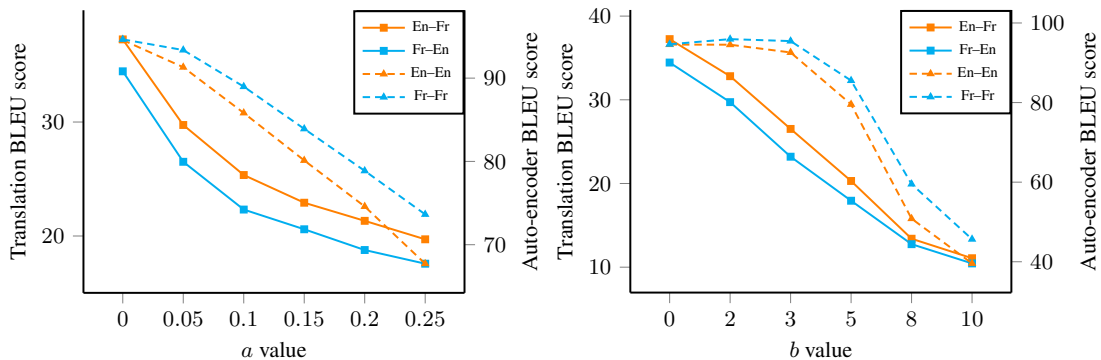


Figure 1: The UNMT performance as the word noise (a value) and word order noise (b value) increases on the noisy En-Fr newstest2014 set.

Figure 1 shows the BLEU scores of UNMT system with different settings of word noise (left panel) and word order noise (right panel). As shown in Figure 1, we can see that as the ratio of noise in the source language input increased, the performance of both the translation direction (En–Fr and Fr–En) and

auto-encoder direction (En–En and Fr–Fr) of UNMT system decreased. In addition, when a slight noise ($a \leq 0.1$ or $b \leq 3$) was added into the input sentence, the translation direction of UNMT still decreased rapidly while a slight downward trend could be maintained in the auto-encoder direction. From these results, we conformed that even a slight input noise could drastically degrade the translation performance of a UNMT system. It is necessary for us to figure out robust solutions to deal with these two types of input noise.

4 Proposed Methods

Based on the previous empirical findings and analysis, we propose two adversarial training methods during denoising training to improve the robustness of UNMT in the two noisy scenarios as defined in Section 3.1. The frameworks are illustrated in Figure 2. The Figure 2 (a) is the original denoising training framework in which an encoder-decoder structure is applied. Based on this original framework, two adversarial training frameworks are proposed (Figure 2 (b) and (c)) in which the word noise and word-order noise blocks are explicitly modeled. Although the original structure has the denoising ability to replace some wrong words and adjust word order, the denoising ability is implicitly modeled without considering the embedding representation effect which is important in translation. As an adversarial training, our two proposed frameworks explicitly model the two noise effects on the word embedding and position embedding which is expected to improve the translation performance. The details are explained in the following subsections.

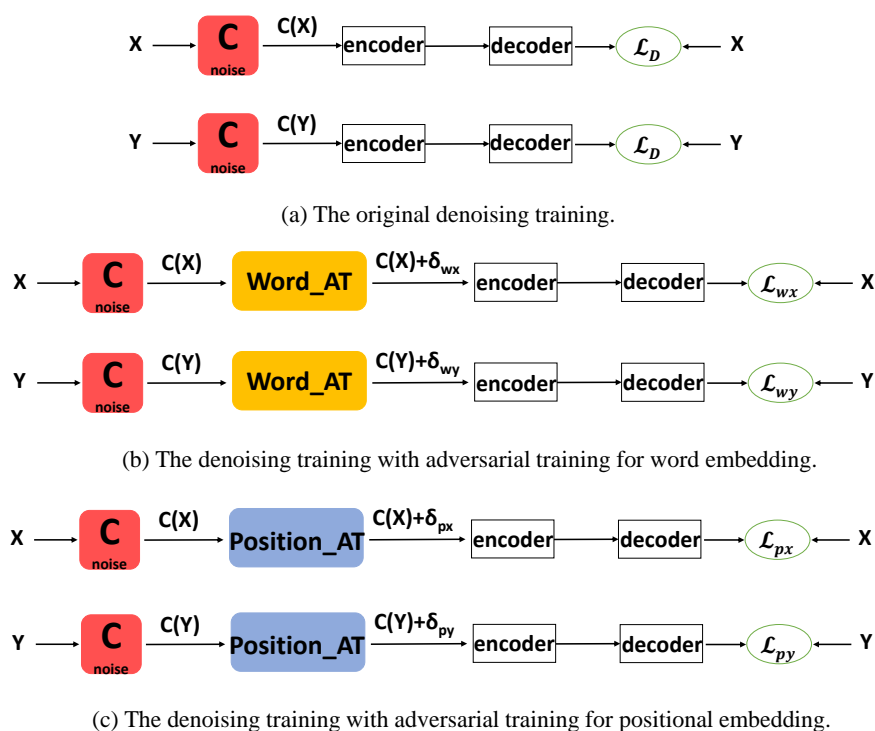


Figure 2: Illustration of (a) the original denoising process; (b) the denoising process with adversarial training for word embedding (Word_AT); (c) the denoising process with adversarial training for positional embedding (Position_AT).

4.1 Adversarial Training for Word Embedding

The adversarial training strategy (Miyato et al., 2017) has been applied to text classification when the input text is noisy. Inspired by this strategy, we apply adversarial training method to the denoising process of UNMT. Regarding to the original denoising process, adversarial perturbation would be added to enhance the learning ability of the denoising auto-encoder. Adversarial perturbation is added to the word embedding as a combined word embedding before combining with the positional embedding, compared to the original transformer based architecture.

As a denoising auto-encoder, for an input of language L_1 with the worst case perturbation δ_{wx} on word embedding, the purpose is to recover the clean version. It is realized by minimizing the reconstruction error defined as:

$$\mathcal{L}_{wx} = \sum_{i=1}^{|X|} -\log P_{L_1 \rightarrow L_1}(X_i | C(X_i) + \delta_{wx}), \quad (6)$$

$$\delta_{wx} = \arg \max_{\|\delta\| \leq \epsilon} -\log P_{L_1 \rightarrow L_1}(X_i | C(X_i) + \delta), \quad (7)$$

where δ is a small perturbation in the source side. Actually, it is intractable to calculate the maximization of objective function as shown in Eq. 7. Following Miyato et al. (2017)'s method, the word adversarial perturbation δ_{wx} is approximated via the gradient of objective function as

$$\delta_{wx} = \epsilon g_x / \|g_x\|_2, \quad (8)$$

$$g_x = \nabla_x -\log P_{L_1 \rightarrow L_1}(X_i | C(X_i)), \quad (9)$$

where g_x denotes the gradient of objective function, calculated by back-propagation algorithm. ϵ is a hyper-parameter to control the magnitude of adversarial perturbation.

The word adversarial perturbation objective function for language L_2 is similarly optimized as

$$\mathcal{L}_{wy} = \sum_{i=1}^{|Y|} -\log P_{L_2 \rightarrow L_2}(Y_i | C(Y_i) + \delta_{wy}), \quad (10)$$

$$\delta_{wy} = \epsilon g_y / \|g_y\|_2, \quad (11)$$

$$g_y = \nabla_y -\log P_{L_2 \rightarrow L_2}(Y_i | C(Y_i)). \quad (12)$$

Typically, to improve UNMT robustness, objective function \mathcal{L}_{wx} and \mathcal{L}_{wy} would be added during the UNMT denoising training process. The entire UNMT objective function is reformulated as follows:

$$\mathcal{L}_{all} = \mathcal{L}_{D'} + \mathcal{L}_B, \quad (13)$$

$$\mathcal{L}_{D'} = \mathcal{L}_D + \mathcal{L}_{wx} + \mathcal{L}_{wy}. \quad (14)$$

4.2 Adversarial Training for Positional Embedding

Word order is very important in translation. In the state-of-the art NMT framework, the positional embedding has been used to encode order information for the source sentence in the transformer based architecture. To capture the order robustness, we propose adversarial training method based on the original positional embedding.

The adversarial perturbation is added to the original positional embedding as a new positional embedding before combining with the word embedding, compared to the original transformer architecture. The positional adversarial perturbation δ_{px} is then used to penalize the existing positional embedding of the source sentence to generate a new order positional embedding during UNMT denoising training for both languages. Similarly as defined in denoising training for word embedding, the objective function for denoising of the positional perturbation is defined as

$$\mathcal{L}_{px} = \sum_{i=1}^{|X|} -\log P_{L_1 \rightarrow L_1}(X_i | C(X_i) + \delta_{px}), \quad (15)$$

$$\mathcal{L}_{py} = \sum_{i=1}^{|Y|} -\log P_{L_2 \rightarrow L_2}(Y_i | C(Y_i) + \delta_{py}), \quad (16)$$

To improve UNMT robustness for word order noise, the objective function \mathcal{L}_{px} and \mathcal{L}_{py} would be added during the UNMT denoising training process. The entire denoising objective function is reformulated as follows:

$$\mathcal{L}_{D'} = \mathcal{L}_D + \mathcal{L}_{px} + \mathcal{L}_{py}. \quad (17)$$

4.3 UNMT with Adversarial Training Mechanism

Based on our proposed adversarial training methods, we design three UNMT systems with adversarial training: adversarial training for word embedding (Word_AT), adversarial training for positional embedding (Position_AT), and the combination of word and positional adversarial training (Both_AT), all of which enrich robust information via adversarial perturbation.

5 Experiments

5.1 Datasets

We considered two language pairs to do simulated experiments on the Fr \leftrightarrow En and German(De) \leftrightarrow En translation tasks. We used 50 million sentences from WMT monolingual news crawl datasets for each language. To make our experiments comparable with previous work (Conneau and Lample, 2019), we reported results on newstest2014 for Fr \leftrightarrow En and newstest2016 for De \leftrightarrow En.

For preprocessing, we used Moses tokenizer (Koehn et al., 2007)¹ for all languages. For cleaning, we only applied the Moses script `clean-corpus-n.perl` to remove lines in the monolingual data containing more than 50 tokens. For BPE (Sennrich et al., 2016b), we used a shared vocabulary for every language pair with 60K subword tokens based on BPE.

5.2 UNMT Settings

We used a transformer-based XLM toolkit² and followed settings of Conneau and Lample (2019) for UNMT: 6 layers for the encoder and the decoder. The dimension of hidden layers was set to 1024. The Adam optimizer (Kingma and Ba, 2015) was used to optimize the model parameters. The initial learning rate was 0.0001, $\beta_1 = 0.9$, and $\beta_2 = 0.98$. The cross-lingual language model was used to pretrain the encoder and decoder of the whole UNMT model. We used 8 NVIDIA V100 GPUs with a batch size of 2,000 tokens per GPU for UNMT training. We used the case-sensitive 4-gram BLEU score computed by `multi-bleu.perl` script from Moses (Koehn et al., 2007) to evaluate the performance on test sets.

5.3 Main Results

Table 1 shows the detailed BLEU scores of all UNMT systems in different level of noisy scenarios on the Fr \leftrightarrow En and De \leftrightarrow En test sets. Our observations are as follows:

1) Our re-implemented baseline in this work outperformed the state-of-the-art method (Conneau and Lample, 2019), using clean input on the Fr \leftrightarrow En test set, and achieved performance comparable to the original method on the De \leftrightarrow En test sets. This indicates that it is a strong UNMT baseline system.

2) In the scenario that there is only word noise in the input, our proposed Word_AT method substantially outperformed the original baseline by approximately 7.4 BLEU scores. Moreover, in the scenario that there is only word order noise in the input, our proposed Position_AT method substantially outperformed the original baseline by approximately 8.3 BLEU scores.

3) In the noisy scenario containing the word noise and word order noise, the performance of the original UNMT system decreased drastically. Our proposed Word_AT and Position_AT method achieved average improvements of 4 and 3.3 BLEU scores, respectively. Moreover, Our proposed Both_AT method could further improve UNMT performance, achieving an average improvement of 10 BLEU scores. This demonstrates that our proposed methods effectively alleviate the noisy input issue.

4) Although our designed adversarial training frameworks tries to remove the effect of noise with explicit inserting word noise and word order noise processing blocks, the frameworks could improve the

¹<https://github.com/moses-smt/mosesdecoder>

²<https://github.com/facebookresearch/XLM>

Word Noise	Word Order Noise	Method	En-Fr	Fr-En	En-De	De-En
No	No	Lample et al. (2018a)	15.05	14.31	9.64	13.33
		Artetxe et al. (2018)	15.13	15.56	n/a	n/a
		Lample et al. (2018b)	27.60	27.68	20.23	25.19
		Conneau and Lample (2019)	33.40	33.30	26.40	34.30
No	No	UNMT	37.25	34.32	26.61	34.24
		+Word_AT	37.68	34.88	27.55	34.39
		+Position_AT	37.72	34.75	27.48	34.41
		+Both_AT	37.81	34.90	27.61	34.43
Yes	No	UNMT	25.35	22.32	14.40	26.29
		+Word_AT	32.90	30.46	24.09	30.42
		+Position_AT	25.85	23.11	15.82	27.16
		+Both_AT	33.52	31.19	24.95	30.58
No	Yes	UNMT	26.52	23.20	15.83	27.52
		+Word_AT	27.48	24.02	16.33	28.48
		+Position_AT	35.44	33.29	25.72	31.86
		+Both_AT	36.29	33.78	25.86	32.58
Yes	Yes	UNMT	21.78	19.03	11.35	22.89
		+Word_AT	26.02	22.49	16.07	26.14
		+Position_AT	25.07	22.02	15.39	25.77
		+Both_AT	32.30	30.75	22.59	29.08

Table 1: Performance (BLEU score) of UNMT in different level of noisy scenarios. $a = 0.1$ for word noise; $b = 3$ for word order noise.

UNMT performance even in clean scenarios, achieving an average improvement of 0.6 BLEU scores on all clean test sets. This suggests that our proposed methods may facilitate the training efficiency of the network.

5.4 Analysis

We empirically investigated the performance of UNMT with Both_AT framework, using the noisy input with different level of word noise and word order noise, respectively. Figure 3 shows the trend measured in BLEU score of UNMT baseline system (orange line) and UNMT system with our proposed Both_AT framework (purple line) for the translation direction (solid line) and auto-encoder direction (dashed line). As shown in Figure 3, we can see that our proposed adversarial training method performed significantly better than the UNMT baseline system in the noisy scenario, regardless of any direction. In particular, as the ratios of word noise and word order noise increase, the performance gap between the baseline UNMT system and our proposed system is increased. These demonstrate that our proposed adversarial training method is robust and effective.

a value	UNMT	+Both_AT	b value	UNMT	+Both_AT
0	100.00	100.00	0	100.00	100.00
0.05	52.61	73.79	2	60.28	78.22
0.1	41.09	66.45	3	43.43	74.07
0.15	35.40	60.40	5	31.56	62.99
0.2	32.10	54.42	8	20.15	43.02
0.25	28.48	48.41	10	16.18	33.75

Table 2: Average similarity (BLEU score) across the translation generated by the clean input and noisy input on the En-Fr newstest2014 set.

In order to get a more complete picture of the robustness of our proposed adversarial training methods, we empirically evaluated the similarity across the translation generated by the clean input and noisy input with different level of noise on the En-Fr newstest2014 set as shown in Table 2. As it can be seen, our

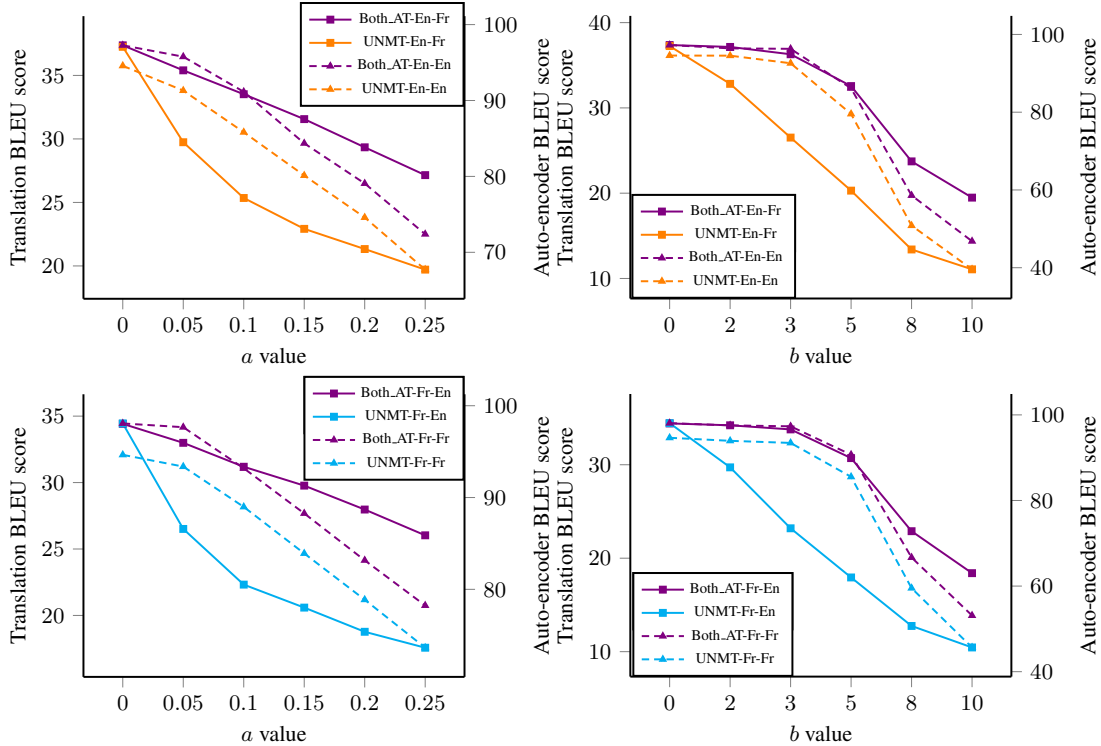


Figure 3: The performance of baseline system and UNMT with Both_AT mechanism as the word noise (a value) and word order noise (b value) increases on the noisy En-Fr newstest2014 set.

proposed Both_AT mechanism produced substantially more similar translations generated by the noisy input and clean input, compared with the UNMT baseline. As Table 2 reports, the similarity of our proposed Both_AT mechanism across the translation generated by the clean input and noisy input with different level of noise was 22.76 BLEU scores more in average for the word noise scenario, and 24.09 BLEU scores more in average for the word order noise scenario, compared with the UNMT system. These further demonstrate that our proposed Both_AT mechanism is robust and can effectively alleviate the impact of two types of noise on translation performance.

5.5 Evaluation on MTNT dataset

To better assess the effectiveness of our proposed adversarial training methods, we investigated the performance of UNMT with Both_AT framework on the MTNT dataset, which is a noisy dataset proposed by Michel and Neubig (2018). The detailed statistics of MTNT data set is presented as shown in Table 3. To make our experiments comparable with previous work (Michel and Neubig, 2018; Zhou et al., 2019), we used the same MTNT parallel training data to fine-tune our proposed +Both_AT system and used sacreBLEU (Post, 2018) to evaluate the translation performance.

Corpus	en-fr	fr-en
Training set	36,058	19,161
Valid set	852	886
Test set	1,020	1,022

Table 3: Statistics of MTNT data set.

Methods	en-fr	fr-en
Michel and Neubig (2018)	21.77	23.27
+Fine-tuning	29.73	30.29
Zhou et al. (2019)	n/a	24.50
+Fine-tuning	n/a	31.70
+Both_AT	31.60	33.80
+Fine-tuning	39.00	41.30

Table 4: BLEU score on the En-Fr MTNT test set.

As shown in Table 4, Our proposed +Both_AT system significantly outperformed the previous work (Michel and Neubig, 2018; Zhou et al., 2019) by approximately 10 BLEU scores. The performance of our proposed system without fine-tuning was even better than that of these previous work with fine-tuning. After fine-tuning, our proposed system achieved promising performance, obtaining 39.00

BLEU scores on the en-fr testset and 41.30 BLEU scores on the fr-en testset. These further demonstrate that our proposed +Both_AT system is a robust system.

5.6 Case Study

Moreover, we analyze translation examples to further analyze the effectiveness of our proposed adversarial training method in the noisy scenario. Table 5 shows two translation examples, which were generated by UNMT baseline system and +Both_AT system, using clean and noisy input on the Fr-En dataset, respectively. For the first example, +Both_AT could adjust the wrong word order to the natural English word order. For the second example, +Both_AT could be better at mitigating the impact of missing words and noisy words. These examples indicate that our proposed +Both_AT system could be widely applied to the real-world noisy scenario.

Clean Input	On est très excités, mais très tendus aussi.
Noisy Input	On est très excités, très mais tendus aussi.
Reference	We were very excited, but also very tense.
Baseline on Clean Input	We're very excited, but very tense too .
+Both_AT on Clean Input	We're very excited, but also very tense.
Baseline on Noisy Input	We're very excited, very tense but also .
+Both_AT on Noisy Input	We're very excited, but also very tense.
Clean Input	Avec la crise de l' euro, le Projet Europe est officiellement mort .
Noisy Input	la crise de l' euro, le Projet Europe est officiellement décès .
Reference	With the euro crisis, Project Europe is officially dead .
Baseline on Clean Input	With the euro crisis, the Europe Project is officially dead .
+Both_AT on Clean Input	With the euro crisis, Project Europe is officially dead .
Baseline on Noisy Input	Amid the euro crisis, the Europe Project is officially a death .
+Both_AT on Noisy Input	With the euro crisis, Project Europe is officially at its death .

Table 5: Comparison of translation results of baseline and +Both_AT system for clean and noisy input.

6 Related Work

Recently, UNMT (Artetxe et al., 2018; Lample et al., 2018a; Yang et al., 2018; Lample et al., 2018b; Sun et al., 2019) that relies solely on monolingual corpora in each language via bilingual word embedding initialization, denoising auto-encoder, back-translation and sharing latent representations. More recently, Conneau and Lample (2019) and Song et al. (2019) introduced the pretrained cross-lingual language model to achieve state-of-the-art UNMT performance. Sun et al. (2020a) extended UNMT to the multilingual UNMT training on a large scale of European languages. Marie et al. (2019) won the first place in the unsupervised translation task of WMT19 by combining UNMT and unsupervised statistical machine translation. However, previous work only focuses on how to build state-of-the-art UNMT systems and ignore the robustness of UNMT on the noisy data. In this paper, we propose adversarial training methods with denoising process in UNMT training to improve the robustness of the UNMT systems. Moreover, our proposed methods could improve the UNMT performance even in clean scenarios.

Actually, Belinkov and Bisk (2018) pointed out that synthetic and natural noise both influenced the translation performance. Belinkov and Bisk (2018), Ebrahimi et al. (2018), and Karpukhin et al. (2019) designed character-level noise, which affects the spelling of a single word, to improve the model robustness. Meanwhile, both textual and phonetic embeddings were used to improve the robustness of SNMT to homophone noises (Liu et al., 2019). Adversarial examples, generated by gradient-based method, attacked the translation model to improve the robustness of SNMT (Cheng et al., 2019). In contrast with this work, we applied adversarial perturbation to the denoising training of UNMT, instead of translation training, to enhance the learning ability of UNMT model.

Adversarial training method, first proposed in the computer vision (Goodfellow et al., 2015; Moosavi-Dezfooli et al., 2016), was applied to several natural language processing tasks (Miyato et al., 2017; Jia and Liang, 2017; Belinkov and Bisk, 2018; Ebrahimi et al., 2018).

7 Conclusion

As a data driven machine learning method, neural translation model is sensitive to small data perturbations which may result in poor performance in noisy scenarios. In this paper, we focused on the noisy robust problem in unsupervised machine translation. We first explicitly defined two types of noise which are frequently appeared in real applications. Then we proposed denoising process on word embedding and word positional embedding based adversarial training in the UNMT framework. Experimental results confirmed that the proposed adversarial training significantly improved the robustness of the UNMT systems in the noisy scenarios. Moreover, even for clean scenario, our proposed framework could slightly improve the performance. From the results, we can conclude that our proposed framework could improve the noisy robustness of the UNMT without sacrificing the performance for clean condition.

In this study, our proposed adversarial training was implemented with the transformer based neural architectures for UNMT. In the future, we will examine the proposed method in other natural language processing tasks and integrate other machine learning methods for the robustness of the UNMT systems.

Acknowledgments

Tiejun Zhao was partially supported by National Key Research and Development Program of China via grant 2017YFB1002102.

References

- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018. Unsupervised neural machine translation. In *Proceedings of the 6th International Conference on Learning Representations*, Vancouver, Canada. OpenReview.net.
- Yonatan Belinkov and Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine translation. In *Proceedings of the 6th International Conference on Learning Representations*, Vancouver, Canada. OpenReview.net.
- Yong Cheng, Zhaopeng Tu, Fandong Meng, Junjie Zhai, and Yang Liu. 2018. Towards robust neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1756–1766, Melbourne, Australia. Association for Computational Linguistics.
- Yong Cheng, Lu Jiang, and Wolfgang Macherey. 2019. Robust neural machine translation with doubly adversarial inputs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4324–4333, Florence, Italy. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems 32*, pages 7059–7069, Vancouver, Canada. Curran Associates, Inc.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. HotFlip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *Proceedings of the 3rd International Conference on Learning Representations*, San Diego, CA, USA.
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Vladimir Karpukhin, Omer Levy, Jacob Eisenstein, and Marjan Ghazvininejad. 2019. Training on synthetic noise improves robustness to natural noise in machine translation. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 42–47, Hong Kong, China. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations*, San Diego, CA, USA.

- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018a. Unsupervised machine translation using monolingual corpora only. In *Proceedings of the 6th International Conference on Learning Representations*, Vancouver, Canada. OpenReview.net.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018b. Phrase-based & neural unsupervised machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5039–5049, Brussels, Belgium. Association for Computational Linguistics.
- Hairong Liu, Mingbo Ma, Liang Huang, Hao Xiong, and Zhongjun He. 2019. Robust neural machine translation with joint textual and phonetic embedding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3044–3049, Florence, Italy. Association for Computational Linguistics.
- Benjamin Marie, Haipeng Sun, Rui Wang, Kehai Chen, Atsushi Fujita, Masao Utiyama, and Eiichiro Sumita. 2019. NICT’s unsupervised neural and statistical machine translation systems for the WMT19 news translation task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 294–301, Florence, Italy. Association for Computational Linguistics.
- Paul Michel and Graham Neubig. 2018. MTNT: A testbed for machine translation of noisy text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 543–553, Brussels, Belgium. Association for Computational Linguistics.
- Takeru Miyato, Andrew M. Dai, and Ian J. Goodfellow. 2017. Adversarial training methods for semi-supervised text classification. In *Proceedings of the 5th International Conference on Learning Representations*, Toulon, France. OpenReview.net.
- Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. 2016. Deepfool: A simple and accurate method to fool deep neural networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2574–2582, Las Vegas, NV, USA. IEEE Computer Society.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. 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.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. Mass: Masked sequence to sequence pre-training for language generation. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5926–5936, Long Beach, CA, USA. PMLR.
- Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2019. Unsupervised bilingual word embedding agreement for unsupervised neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1235–1245, Florence, Italy. Association for Computational Linguistics.
- Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2020a. Knowledge distillation for multilingual unsupervised neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3525–3535, Online, July. Association for Computational Linguistics.
- Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2020b. Unsupervised neural machine translation with cross-lingual language representation agreement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:1170–1182.

- Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11:3371–3408.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2018. Unsupervised neural machine translation with weight sharing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 46–55, Melbourne, Australia, July. Association for Computational Linguistics.
- Shuyan Zhou, Xiangkai Zeng, Yingqi Zhou, Antonios Anastasopoulos, and Graham Neubig. 2019. Improving robustness of neural machine translation with multi-task learning. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 565–571, Florence, Italy. Association for Computational Linguistics.