

Robust Machine Translation with Domain Sensitive Pseudo-Sources: Baidu-OSU WMT19 MT Robustness Shared Task System Report

Renjie Zheng^{*‡} Hairong Liu^{*†} Mingbo Ma[†] Baigong Zheng[†] Liang Huang^{†,‡}

[†]Baidu Research, Sunnyvale, CA

[‡]School of EECS, Oregon State University, Corvallis, OR

{zrenj11, lhrbss, cosmb, zbgzbg2007, liang.huang.sh}@gmail.com

Abstract

This paper describes the machine translation system developed jointly by Baidu Research and Oregon State University for WMT 2019 Machine Translation Robustness Shared Task. Translation of social media is a very challenging problem, since its style is very different from normal parallel corpora (e.g. News) and also include various types of noises. To make it worse, the amount of social media parallel corpora is extremely limited. In this paper, we use a domain sensitive training method which leverages a large amount of parallel data from popular domains together with a little amount of parallel data from social media. Furthermore, we generate a parallel dataset with pseudo noisy source sentences which are back-translated from monolingual data using a model trained by a similar domain sensitive way. We achieve more than 10 BLEU improvement in both En-Fr and Fr-En translation compared with the baseline methods.

1 Introduction

Translation of social media is very challenging. First, there are various types of noises, such as abbreviations, spelling errors, obfuscated profanities, inconsistent capitalization, Internet slang and emojis (Michel and Neubig, 2018). Second, the amount of parallel data is limited. These characteristics of social media make existing neural machine translation systems extremely vulnerable.

The noise issue of social media has been investigated in some previous work (Baldwin et al., 2013; Eisenstein, 2013). Most recently, Belinkov and Bisk (2017) demonstrated the vulnerability of neural machine translation system to both synthetic and natural noises. However, the noises tested in (Belinkov and Bisk, 2017) are not real noises in social media. To our best knowledge, there seems

to be a lack of translation methods systematically targeting noises in social media.

Existing neural machine translation systems are famous for their hungry of data. However, the amount of parallel data in social media domain is very limited. Just recently, a dataset collected from Reddit has been published and attracted a lot of attention (Michel and Neubig, 2018). The amount of data in this dataset is still very small, compared to the large amount of data from News domain. Naturally, how to utilize the large amount of parallel data from the News domain become a central problem in improving the translation of social media.

In this paper, inspired by the success of back-translation technique (Sennrich et al., 2015a), we propose to learn a model to generate “social-media-style” translation in source language from clean sentences in target language. Since the amount of parallel data in social media domain is limited, we utilize the large amount of parallel data in News domain to help the training. With this model, large amount of parallel data for back-translation can be generated from monolingual data in target language. In the final translation model, a special “domain” symbol is added to indicate which domain the source sentence belonging to.

The contributions of this paper are multifold, and some important ones are highlighted below:

1. We found that “social-media-style” sentences can be generated by training a translation model with different “start-of-sentence” symbols for sentences in different domains in the decoder side. The model is trained with data from all domains, especially News domain, which has a large amount of parallel data, but also adapted to the style in the domain of social media, even the amount of

* Equal contribution

parallel data in social media is limited. As demonstrated by our experiments, generating “social-media-style” sentences is crucial in the effectiveness of back-translation for training a translation model suitable for translating social media.

2. We illustrated that adding a domain symbol in source sentence improves the robustness of the model. This may be because the encoder learns some domain-specific features from input sentences.

2 Methods

Noisy text translation is short of in-domain training data. In this section, we present approaches to leverage a large amount of out-of-domain (e.g. News) dataset and monolingual data paired with pseudo noisy source data from back-translation.

2.1 Domain Sensitive Data Mixing

To improve the translation model from limited parallel data, we want to make the use of larger amount of out-of-domain data. However, simply mixing the clean and noisy data will make the whole training set unbalanced. To differentiate the data from different domain, we use different start symbol in source side.

The intuition of injecting domain label in source side is based on the noise occurrence statistics from (Michel and Neubig, 2018), which shows much more spelling and grammar errors in the source side of noisy text translation dataset. Thus the clean and noisy start symbols work as a meaningful sign of source text style for encoder. Compared with the source side sentences, the human translation of target side sentences are less noisier with less spelling and grammar errors.

2.2 Noisy Pseudo-Sources Generation with Back-Translation

To further make the use of monolingual data, we regard them as target data and generate its corresponding source data by back-translation (Sennrich et al., 2015a). However, different from Sennrich et al. (2015a) who uses this method in both clean source and target sentences, the source side sentences in our test set is much noisier than target side (as mentioned in previous subsection). Therefore, we reverse the source and target sentences where the noisy source sentences becomes target and cleaner target sentences becomes source. For

example, to generate noisy pseudo French source sentences for English monolingual data, we train a En-Fr translation model which takes the noisy French source sentences in Fr-En noisy dataset as target, and the corresponding paralleled English target sentences as source. In this way, the model will learned how to inject noise into the target side. Similar to previous domain sensitive method, we include out-of-domain clean data during the training of this noisy translation model and differentiate them by different start symbol int target side.

2.3 Ensemble

In our experiments with relatively small training dataset, the translation qualities of models with different initializations can vary notably. To make the performance much more stable and improve the translation quality, we ensemble different models during decoding to achieve better translation.

To ensemble, we take the average of all model outputs:

$$\hat{y}_t = \sum_{i=1}^N \frac{\hat{y}_t^i}{N} \quad (1)$$

where \hat{y}_t^i denotes the output distribution of i th model at position t . Similar to Zhou et al. (2017) and Zheng et al. (2018c), we can ensemble models trained with different architectures and training algorithms.

3 Experiments

To investigate the empirical performances of our proposed methods, we conduct experiments on MTNT dataset (Michel and Neubig, 2018) using Transformer (Vaswani et al., 2017).

We first apply BPE (Sennrich et al., 2015b) on both sides in order to reduce the vocabulary for both source and target sides. We then exclude the sentences pairs whose length are longer than 256 words or subwords. We use length reward (Huang et al., 2017) to find the optimal target length.

Our implementation is adapted from PyTorch-based OpenNMT (Klein et al., 2017). Our Transformer’s parameters are as the same as the base model’s parameter settings in the original paper (Vaswani et al., 2017).

In all experiments, our evaluation uses sacreBLEU ¹, a standardized BLEU score evaluation

¹<https://github.com/mjpost/sacreBLEU>

	Training	Validation	Test
Clean	2,207,962	-	-
Monolingual	26,485	-	-
Noisy	36,058	852	1,020

Table 1: Statistics of En2Fr Dataset. Monolingual data is French only.

	Training	Validation	Test
Clean	2,207,962	-	-
Monolingual	2,244,020	-	-
Noisy	19,161	886	1,022

Table 2: Statistics of Fr2En Dataset, Monolingual data is English only.

	En2Fr	Fr2En
Domain Insensitive	31.3	34.6
Domain Sensitive	35.7	39.5

Table 3: Results of noisy data generation. We reverse the source and target direction of MTNT Fr2En (En2Fr) dev-set to evaluate the ability of noisy data generation for En2Fr (Fr2En).

tool by Post (2018). We specify the `intl` tokenization option during BLEU evaluation. We also use detokenization and normalization tools in Moses.

Table 1 and 2 show statistics of En2Fr and Fr2En datasets. For both En-Fr and Fr-En dataset, the clean parallel data is from WMT15 news translation task. The noisy data is from (Michel and Neubig, 2018) collected from social network. Except the French and English monolingual data from WMT15 news translation task, we also make the use of English portion of parallel data from KFTT, TED and JESC used in (Michel and Neubig, 2018).

3.1 Noisy Data Generation

To make use of monolingual target data, we want to generate the corresponding parallel pseudo noisy source data and put them into training set. Table 3 shows the performance of our noisy data generation models. In this experiment, we mix the clean and noisy dataset as the training set, but use the target sentences in reversed direction of noisy dataset (training, validation, test) set as source and source sentences as target. The domain insensitive method simply mix the clean and noisy dataset in training while the domain sensitive method differentiate the clean and noisy dataset in target side by

starting with different symbol (e.g. `< clean_s >`, `< noisy_s >`). The experiment shows that the domain sensitive method can outperform the domain insensitive method with a large margin.

3.2 Methods Comparison

Table 4 shows the final results of different methods on test set. Similar with the previous experiments, the domain insensitive methods simply mix all the clean, noisy training data. The performance has a little improvement in En-Fr by adding the monolingual data paired with the pseudo source data generated by the model trained in previous experiments. To differentiate the clean and noisy dataset, we assign different label at the start of them and the performance is thus boosted about 3 to 4 BLEU score. We further generate pseudo noisy source data from the monolingual target with the model using the domain sensitive method in previous experiment. By adding these noisy back-translation data, we achieve more than 2 BLEU improvement. Our final submission ensembles 5 models trained with the domain sensitive method and including the noisy back translation data.

3.3 Final Results

Table 5 and Table 6 show the final results of our submission in Fr-En and En-Fr. Our system ranks third in both directions. Table 7 shows the human judgments over all submitted systems which are done by Li et al. (2019) who also analyze and discuss all submitted systems.

4 Related Work

The method proposed in this paper is a kind of domain adaptation technique. There are many previous work on domain adaptation for machine translation (Britz et al., 2017; Wang et al., 2017; Chu et al., 2017; Chu and Wang, 2018), which leverages out-of-domain parallel corpora and in-domain monolingual corpora to improve translation. The difference between our method and previous work lies in that we use back-translation (Sennrich et al., 2015a) for domain adaptation. Different from some previous work using adversarial training (Liu et al., 2017) or different attention (Zheng et al., 2018a) to differentiate multiple tasks, we simply assign different starting symbol for multiple tasks (Lample et al., 2018).

A similar method was proposed in (Xie et al.,

	Methods	En-Fr	Fr-En
Baseline	MTNT †	21.8	23.3
	+ tuning †	29.7	30.3
Domain Insensitive	Mix training	33.4	34.5
	+ Back translation	33.7	34.3
Domain Sensitive	Mix training	36.3	38.7
	+ Noisy back translation	38.4	41.0
	+ Ensemble	40.4	42.3

Table 4: Results of different methods on test-set. †(Michel and Neubig, 2018)

	BLEU	BLEU-cased	BEER	CharacTER
NLE	48.8	47.9	0.676	0.364
CUNI	45.8	44.8	0.654	0.395
BD-OSU*	44.5	43.6	0.641	0.499
JHU	41.2	40.2	0.624	-
CMU	32.8	32.2	0.573	0.514
FOKUS†	30.8	29.9	0.530	0.574
MTNT	26.2	25.6	0.529	0.550
IITP-MT	25.5	20.8	0.499	0.594

Table 5: Semi-blind test results of Fr-En. *Our submission. †Unconstrained.

	BLEU	BLEU-cased	BEER	CharacTER
NLE	42.0	41.4	0.626	0.446
CUNI	39.1	38.5	0.605	0.483
BD-OSU*	37.0	36.4	0.599	0.512
FOKUS†	24.8	24.2	0.515	0.619
MTNT	22.5	22.1	0.498	0.621
CMU	20.8	20.4	0.488	0.622
IITP-MT	20.7	19.2	0.492	0.619
SFU	19.4	19.1	0.491	0.614

Table 6: Semi-blind test results of En-Fr. *Our submission. †Unconstrained.

	En-Fr	En-Fr	Fr-En	Fr-En
	Score	Rank	BLEU	Rank
BD-OSU*	71.5	2	80.6	3
CMU	-	-	58.2	6
CUNI	66.3	3	82.0	2
JHU	-	-	76.3	4
NaverLabs	75.5	1	85.3	1
FOKUS†	52.5	4	62.6	5

Table 7: Human judgments over all submitted systems (the higher the better) *Our submission. †Unconstrained.

2018) in the context of grammar correction, where a model is trained to add noises on original sentences to produce noisy sentences. However, instead of learn how to generate arbitrary “noises”, our goal is to learn “social-media-style” translations. Singh et al. (2019) injects artificial noise

in the clean data according to the distribution of noisy data. Liu et al. (2019a) propose to leverage phonetic information to reduce the noises in data.

Another group of work related to this paper is data augmentation in machine translation. Although data augmentation is very popular in gen-

eral learning tasks, such as image processing, it is non-trivial to do so in machine translation because even slight modifications of sentences can make huge difference in semantics. To our best knowledge, there are two categories of successful data augmentation approaches for machine translation. The first one is based on back-translation ((Sennrich et al., 2015a)) which augments monolingual data into training set. The second one is based on word replacement, such as (Sennrich et al., 2016) and (Wang et al., 2018). Zheng et al. (2018b) make the use of multiple references and generates even more pseudo-references and achieve improvement in both machine translation and image captioning.

5 Conclusions and Future Work

In this paper, we proposed a method to improve the translation of social media. The style of social media is very unique, and is very different from the style of widely researched News sentences. The core part of our method is to generate useful parallel data for back-translation, that is, generating synthetic in-domain parallel data. To achieve this goal, we proposed a method to generate “social-media-style” source sentences from monolingual target sentences. We also distinguish the domain of source sentences by inserting a domain symbol into source sentences. Both techniques are proven to be extremely useful in the scenario of translating social media. Finally, we utilized the ensemble to further boosts the translation performance.

The noises in social media are mostly introduced by human mistakes. There are some other cases that noises in source side are introduced by systems, such as ASR in speech-to-text translation (Liu et al., 2019b). We plan to further investigate this domain sensitive method on these tasks, even on speech-to-text simultaneous translation (Ma et al., 2018; Zheng et al., 2019).

References

Timothy Baldwin, Paul Cook, Marco Lui, Andrew MacKinlay, and Li Wang. 2013. How noisy social media text, how diffrent social media sources? In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 356–364.

Yonatan Belinkov and Yonatan Bisk. 2017. Synthetic and natural noise both break neural machine translation. *arXiv preprint arXiv:1711.02173*.

Denny Britz, Quoc Le, and Reid Pryzant. 2017. Effective domain mixing for neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pages 118–126.

Chenhui Chu, Raj Dabre, and Sadao Kurohashi. 2017. An empirical comparison of simple domain adaptation methods for neural machine translation. *arXiv preprint arXiv:1701.03214*.

Chenhui Chu and Rui Wang. 2018. A survey of domain adaptation for neural machine translation. *arXiv preprint arXiv:1806.00258*.

Jacob Eisenstein. 2013. What to do about bad language on the internet. In *Proceedings of the 2013 conference of the North American Chapter of the association for computational linguistics: Human language technologies*, pages 359–369.

Liang Huang, Kai Zhao, and Mingbo Ma. 2017. When to finish? optimal beam search for neural text generation (modulo beam size). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2134–2139.

G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. M. Rush. 2017. *OpenNMT: Open-Source Toolkit for Neural Machine Translation*. *ArXiv e-prints*.

Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, et al. 2018. Phrase-based & neural unsupervised machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5039–5049.

Xian Li, Paul Michel, Antonios Anastasopoulos, Yonatan Belinkov, Nadir K. Durrani, Orhan Firat, Philipp Koehn, Graham Neubig, Juan M. Pino, and Hassan Sajjad. 2019. Findings of the first shared task on machine translation robustness. In *Proceedings of the Fourth Conference on Machine Translation, Volume 2: Shared Task Papers*, Florence, Italy. Association for Computational Linguistics.

Hairong Liu, Mingbo Ma, Liang Huang, Hao Xiong, and Zhongjun He. 2019a. Robust neural machine translation with joint textual and phonetic embedding. *ACL*.

Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2017. Adversarial multi-task learning for text classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1–10.

Yuchen Liu, Hao Xiong, Zhongjun He, Jiajun Zhang, Hua Wu, Haifeng Wang, and Chengqing Zong. 2019b. End-to-end speech translation with knowledge distillation. *arXiv preprint arXiv:1904.08075*.

Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, Hua Wu, and Haifeng Wang. 2018. Stacl: Simultaneous translation with integrated anticipation and controllable latency. *ACL*.

- Paul Michel and Graham Neubig. 2018. Mtn: A testbed for machine translation of noisy text. *arXiv preprint arXiv:1809.00388*.
- Matt Post. 2018. A call for clarity in reporting bleu scores. *arXiv preprint arXiv:1804.08771*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015a. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015b. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Edinburgh neural machine translation systems for wmt 16. *arXiv preprint arXiv:1606.02891*.
- Sumeet Singh, Craig Stewart, Graham Neubig, et al. 2019. Improving robustness of machine translation with synthetic noise. *arXiv preprint arXiv:1902.09508*.
- 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 *Advances in Neural Information Processing Systems 30*.
- Rui Wang, Masao Utiyama, Lemao Liu, Kehai Chen, and Eiichiro Sumita. 2017. Instance weighting for neural machine translation domain adaptation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1482–1488.
- Xinyi Wang, Hieu Pham, Zihang Dai, and Graham Neubig. 2018. Switchout: an efficient data augmentation algorithm for neural machine translation. *arXiv preprint arXiv:1808.07512*.
- Ziang Xie, Guillaume Genthial, Stanley Xie, Andrew Ng, and Dan Jurafsky. 2018. Noising and denoising natural language: Diverse backtranslation for grammar correction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 619–628.
- Baigong Zheng, Renjie Zheng, Mingbo Ma, and Liang Huang. 2019. Simultaneous translation with flexible policy via restricted imitation learning. *ACL*.
- Renjie Zheng, Junkun Chen, and Xipeng Qiu. 2018a. Same representation, different attentions: shareable sentence representation learning from multiple tasks. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 4616–4622. AAAI Press.
- Renjie Zheng, Mingbo Ma, and Liang Huang. 2018b. Multi-reference training with pseudo-references for neural translation and text generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3188–3197.
- Renjie Zheng, Yilin Yang, Mingbo Ma, and Liang Huang. 2018c. Ensemble sequence level training for multimodal mt: Osu-baidu wmt18 multimodal machine translation system report. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 632–636.
- Long Zhou, Wenpeng Hu, Jiajun Zhang, and Chengqing Zong. 2017. Neural system combination for machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 378–384.