

Improving Neural Machine Translation Robustness via Data Augmentation: Beyond Back Translation

Zhenhao Li

Department of Computing
Imperial College London, UK
zhenhao.li18@imperial.ac.uk

Lucia Specia

Department of Computing
Imperial College London, UK
l.specia@imperial.ac.uk

Abstract

Neural Machine Translation (NMT) models have been proved strong when translating clean texts, but they are very sensitive to noise in the input. Improving NMT models robustness can be seen as a form of “domain” adaption to noise. The recently created Machine Translation on Noisy Text task corpus provides noisy-clean parallel data for a few language pairs, but this data is very limited in size and diversity. The state-of-the-art approaches are heavily dependent on large volumes of back-translated data. This paper has two main contributions: Firstly, we propose new data augmentation methods to extend limited noisy data and further improve NMT robustness to noise while keeping the models small. Secondly, we explore the effect of utilizing noise from external data in the form of speech transcripts and show that it could help robustness.

1 Introduction

Neural Machine Translation (NMT) models trained on large, clean parallel corpora have reached impressive performance in translating clean texts following various architectures (Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017).

Despite this success, NMT models still lack robustness when applied to noisy sentences. [Belyukov and Bisk \(2017\)](#) show that perturbations in characters could cause a significant decrease in translation quality. They point out that training on noisy data, which can be seen as adversarial training, might help to improve model robustness. [Michel and Neubig \(2018\)](#) propose the Machine Translation on Noisy Text (MTNT) dataset, which contains parallel sentence pairs with comments crawled from Reddit¹ and manual translations. This dataset contains user-generated text

with different kinds of noise, e.g., typos, grammatical errors, emojis, spoken languages, etc. for two language pairs.

In the WMT19 Robustness Task² ([Li et al., 2019](#)), improving NMT robustness is treated as a domain adaption problem. The MTNT dataset is used as in-domain data, where models are trained with clean data and adapted to noisy data. Domain adaption is conducted in two main methods: fine tuning on in-domain data ([Dabre and Sumita, 2019](#); [Post and Duh, 2019](#)) and mixed training with domain tags ([Berard et al., 2019](#); [Zheng et al., 2019](#)). The size of the noisy data provided by the shared task is small, with only thousands of noisy sentence pairs on each direction. Hence most approaches participating in the task performed noisy data augmentation using back translation ([Berard et al., 2019](#); [Helcl et al., 2019](#); [Zheng et al., 2019](#)), with some approaches also directly adding synthetic noise ([Berard et al., 2019](#)). The robustness of an NMT model can be seen as denoising source sentences (e.g. dealing with typos, etc.) while keeping a similar level of informal language in the translations (e.g. keeping emojis/emoticons). Based on this assumption, we believe that back translation of clean texts, although providing a large volume of extra data, is limited since it removes most types of noise from the translations. In addition to adapting models on noisy parallel data, other techniques have been used to improve performance, generally measured according to BLEU ([Papineni et al., 2002](#)) against clean references. For example, [Berard et al. \(2019\)](#) apply inline-casing by adding special tokens before each word to represent word casing. In ([Murakami et al., 2019](#)), placeholders are used to help to translate sentences with emojis.

In this paper, we also explore data augmenta-

¹www.reddit.com

²<http://www.statmt.org/wmt19/robustness.html>

tion for robustness but focus on techniques other than back translation. We follow the WMT19 Robustness Task and conduct experiments under constrained and unconstrained data settings on Fr \leftrightarrow En as language pairs. Under the constrained setting, we only use datasets provided by the shared task, and propose new data augmentation methods to generate noise from this data. We compare back translation (BT) (Sennrich et al., 2016a) with forward translation (FT) on noisy texts and find that pseudo-parallel data from forward translation can help improve more robustness. We also adapt the idea of fuzzy matches from Bulte and Tezcan (2019) to the MTNT case by finding similar sentences in a parallel corpus to augment the limited noisy data. Results show that the fuzzy match method can extend noisy parallel data and improve model performance on both noisy and clean texts. The proposed techniques substantially outperform the baseline. While they still lag behind the winning submission in the WMT19 shared task, the resulting models are trained on much smaller clean data but augmented noisy data, leading to faster and more efficient training. Under the unconstrained setting, we propose for the first time the use of speech datasets, in two forms: (a) the IWSLT (Cettolo et al., 2012) and MuST-C (Di Gangi et al., 2019) human transcripts as a source of spontaneous, clean speech data, and (b) automatically generated transcripts for the MuST-C dataset as another source of noise. We show that using informal language from spoken language datasets can also help to increase NMT robustness.

This paper is structured as follows: Section 2 introduces the data augmentation methods for noisy texts, including the previously proposed methods and our approaches to data augmentation. Section 3 describes our experimental settings, including the datasets we used, the augmented data and the baseline model. Section 4 shows the results of models built from different training and evaluated on both clean and noisy test sets.

2 Noisy Data Augmentation

2.1 Previous Work

Considering the limited size of noisy parallel data, data augmentation methods are commonly used to generate more noisy training materials.

Previous methods include back translation, injecting synthetic noise, and adversarial attacks. In the WMT19 Robustness Task, back translation

on monolingual data was used to generate noisy parallel sentences (Murakami et al., 2019; Zhou et al., 2019; Berard et al., 2019; Helcl et al., 2019). Zheng et al. (2019) proposed an extension of back translation that generates noisy translations from clean monolingual data. Therefore, after reversing the direction, the noisy translations become the source, which would simulate the noisy source sentences from the MTNT parallel data. Synthetic noise is injected into clean data to form noisy parallel data in Belinkov and Bisk (2017); Karpukhin et al. (2019). However, rule-based synthetic noise injection is limited to certain types of noise. Adversarial methods are proposed to inject random noise into clean training data in Cheng et al. (2018, 2019).

We explore the following new methods as alternative ways to augment noisy parallel data.

2.2 Fuzzy Matches

We adapted the method to augment data from parallel corpus from Bulte and Tezcan (2019). The original method aims to find similar source sentences to those in a parallel corpus (S, T) using a monolingual corpus, and then reuse the translation of the original source sentences as translations for the similar source sentences found. We adapted it to use only on the provided noisy training corpus. For each source sentence $s_i \in S$ in the training set, all other source sentences $s_j \in S (s_i \neq s_j)$ are compared with this sentence by measuring string similarity $Sim(s_i, s_j)$. If the similarity of the two sentences is above a threshold λ , the two sentences are mapped to each other’s corresponding target sentence and the two new sentence pairs $(s_i, t_j), (s_j, t_i)$ are added into our augmented training data. The similarity is measured with Levenshtein distance (Levenshtein, 1966) on the token level. The similarity score is calculated as the edit distance divided by the minimum length of the two sentences (Equation 1). In addition to fuzzy matches in the parallel corpus, we experimented with the monolingual corpus by mapping sentence m_i in monolingual corpus to its fuzzy match’s target sentence t_j (If $Sim(m_i, s_j) > \lambda$, we add new sentence pair (m_i, t_j) to the training augmented data).

$$Sim(S_i, S_j) = \frac{editdistance(s_i, s_j)}{\min(len(s_i), len(s_j))} \quad (1)$$

To boost the speed of finding matches, we followed the approach in Bulte and

Tezcan (2019) and used a Python library `SetSimilaritySearch`³ to select similar candidates before calculating edit distance. For each source sentence, only the top 10 similar candidates are selected to calculate the edit distance score.

2.3 Forward Translation

Back translation (Sennrich et al., 2016a) is a very popular technique for in-domain data augmentation. In the experiment, we back-translated MTNT monolingual data using a model fine-tuned on MTNT parallel corpus. However, considering that the task of improving robustness is to produce less noisy output, data generated with back translation might have noisy target translations (from monolingual data) and less noisy source texts (from back translation). Since this might increase the noise level of the output texts, we also experimented with forward translation using models fine-tuned on the noisy parallel corpus. Pseudo parallel data generated by forward translation is used for fine tuning models on the same language direction. To avoid overfitting, we merged the noisy parallel data of both language directions to produce noisy forward translations. The pseudo parallel data generated by back translation and forward translation is combined with noisy parallel data and fine-tuned on the baseline model.

2.4 Automatic Speech Recognition

We used ASR systems and transcribed audio files into texts. In this case, we would expect noise to be generated during the process of automatic speech recognition. We selected a dataset with both audio and human transcripts, namely the MuST-C dataset. In this dataset, audio files A , human transcripts S and transcript translated into another language T are provided. We used Google Speech-to-Text API⁴ and transcribed the audio files into automatic transcripts S' . The human and ASR transcripts of the audio (S and S') are treated as the source texts while the translations T are target texts. We formed a new set of parallel data (S', T) with ASR generated texts and the corresponding gold translations by humans.

Looking into the ASR transcripts, we found that the ASR system tends to skip some sentences due

³<https://github.com/ekzhu/SetSimilaritySearch>

⁴<https://cloud.google.com/speech-to-text/>

to the fast speaking speed, therefore we did some filtering based on the length ratio of the human transcripts and the ASR transcripts. We set a ratio threshold λ . For each pair s_i, t_i in the ASR parallel data (S', T), we removed this sentence pair if the length of t_i dividing by s_i is larger than the threshold λ .

3 Experiments

3.1 Corpora

We used all parallel corpora from the WMT19 Robustness Task on Fr \leftrightarrow En. For out-of-domain training, we used the WMT15 Fr \leftrightarrow En News Translation Task data, including Europarl v7, Common Crawl, UN, News Commentary v10, and Gigaword Corpora. In the following sections, we represent the combination of these corpora as "clean data". The MTNT dataset is used as our in-domain data for fine tuning. We also experimented with external corpora, namely the IWSLT2017⁵ and MuST-C⁶ corpora⁷, to explore the effect of informal spoken languages in human transcripts from speech. We also utilized monolingual data in the MTNT dataset. The size of parallel and monolingual data is shown in Table 1.

We used the development set in MTNT and the *newsdiscussdev2015* for validation. Models with best performance on the validation set are evaluated on both noisy (MTNT and MTNT2019) and clean (*newstest2014* and *newsdiscusstest2015*) test sets.

For preprocessing, we first tokenized the data with Moses tokenizer (Koehn et al., 2007) and applied Byte Pair Encoding (BPE) (Sennrich et al., 2016b) to segment words into subwords. We experimented with a large vocabulary to include noisy as well as clean tokens and applied 50k merge operations for BPE. Upon evaluation, we detokenized our hypothesis files with the Moses detokenizer. We used `multi-bleu-detok.perl` to evaluate the BLEU score on the test sets.

⁵<https://wit3.fbk.eu/mt.php?release=2017-01-trnted>

⁶<https://ict.fbk.eu/must-c/>

⁷The data from IWSLT has the same sentences in both translation directions, so we reversed the En \rightarrow Fr data on the Fr \rightarrow En direction. The MuST-C data is only used on En \rightarrow Fr direction.

Corpus	Sentences	Words	
		source	target
Gigaword	22.52M	575.67M	672.07M
UN Corpus	12.89M	316.22M	353.92M
Common Crawl	3.24M	70.73M	76.69M
Europarl	2.01M	50.26M	52.35M
News Commentary	200k	4.46M	5.19M
IWSLT	236k	4.16M	4.34M
MuST-C(en-fr)	275k	5.09M	5.30M
MTNT(en-fr)	36k	841k	965k
MTNT(fr-en)	19k	661k	634k
MTNT(en)	81k	3.41M	
MTNT(fr)	26k	1.27M	

Table 1: Size of parallel and monolingual data. Source language represents English unless specified in brackets. The last two rows are noisy monolingual datasets in the MTNT dataset.

3.2 Augmented Data

Fuzzy Match Considering the large size of our clean data, we only looked for noisy matches within the MTNT parallel and monolingual data. The parallel data on both directions was merged to get more combinations. With the similarity threshold setting to 50%, we have 7,290 new sentence pairs in the En→Fr and 7,154 in the Fr→En language directions.

Forward Translation We used noisy monolingual data to produce forward translations. Therefore, the number of parallel samples is the same as the monolingual data in Table 1. Apart from tokenization and BPE, we did not do any other preprocessing to the forward translation.

ASR We filtered the ASR parallel data by different length ratio thresholds (1.5 and 1.2). We measured the noise level of ASR transcripts S' by evaluating Word Error Rate (WER) and Word Recognition Rate (WRR) with `asr_evaluation`⁸ library, comparing it to the gold human transcripts S .

Data	WER	WRR
ASR($\lambda = 1.5$)	36.41%	65.38%
ASR($\lambda = 1.2$)	31.70%	70.54%

Table 2: Noise level of ASR generated data filtered with different length ratio thresholds.

⁸<https://github.com/belambert/asr-evaluation>

3.3 Model

Our baseline model, which only uses clean data for training, is the standard Transformer model (Vaswani et al., 2017) with default hyperparameters. The batch size is 4096 tokens (subwords). Our models are trained with OpenNMTpy (Klein et al., 2017) on a single GTX 1080 Ti for 5 epochs. We experimented with fine tuning on noisy data and mixed training with “domain” tags (Caswell et al., 2019; Kobus et al., 2016) indicating where the sentences are sourced from. We used different tags for clean data, MTNT parallel data, forward translation, back translation, ASR data, and fuzzy match data. Tags are added at the beginning of each source sentences.

During the fine tuning on in-domain data, we continued the learning rate and model parameters. We stop the fine tuning when the perplexity on noisy validation set does not improve after 3 epochs. Best fine-tuned checkpoints are evaluated on the test sets.

The MTNT dataset provides noisy parallel data in specific language pairs. We used models with two fine tuning strategies: Tune-S and Tune-B. The Tune-S model is fine-tuned only with the noisy parallel data in the same direction while the Tune-B model is fine-tuned with the combination of both language directions (Fr→En and En→Fr).

4 Results

We evaluated the models fine-tuned on different datasets in terms of BLEU on both noisy and clean test sets. We note that although both MTNT and MTNT2019 test sets are noisy, the MTNT2019 is less noisy and contains fewer occurrences of noise such as emojis. Similarly, since the *newsdiscuss* test set contains informal language, it is slightly noisier than *newstest* test set. The evaluation results for both directions are shown in Tables 3 and 4.

4.1 Fine Tuning on Noisy Text

It can be seen that for both directions fine tuning on noisy data gives better performance for the noisy test set. Although the size of training data in MTNT is only 19k and 36k sentences, by simply fine tuning on it, the BLEU scores of Tune-S model increase by +5.65 and +6.03 on MTNT test set, +5.44 and +2.68 on MTNT2019 test set (See the second row in the tables). It is also worth noticing that although fine-tuned on noisy data, the

Models	Fine-tune	MTNT	MTNT2019	newstest2014	newsdiscusstest2015
Baseline	None	34.41	36.14	36.51	34.43
Constrained	Tune-S	40.16	41.58	35.94	36.75
	+BT	37.93	39.19	34.03	34.00
	Tune-B	38.25	40.12	35.36	35.15
	+FM	38.91	40.85	35.62	34.58
	+FT	39.38	41.82	35.56	35.00
	+FT+FM	40.13	42.80	35.82	35.88
Unconstrained	+double tune	40.57	42.55	36.06	36.53
	IWSLT	34.22	38.28	35.96	32.95
	MTNT+IWSLT	37.52	41.22	36.33	34.07

Table 3: BLEU scores of models fine-tuned on different data in the Fr→En direction. The Tune-B model is fine-tuned with MTNT data merging both Fr→En and En→Fr directions while the Tune-S is fine-tuned only with Fr→En data. BT and FT stand for back and forward translation, respectively. FM means fuzzy matching data. By double tuning we mean we fine-tune the model a second time on the MTNT training data.

Models	Fine-tune	MTNT	MTNT2019	newstest2014	newsdiscusstest2015
Baseline	None	30.12	29.57	35.51	35.93
Constrained	MTNT(Tune-S)	36.15	32.25	36.77	37.16
	+fix punctuation	—	35.49	—	—
	MTNT*(Tune-B)	35.64	31.27	36.67	37.21
	+FM	36.03	31.37	36.58	37.13
	+FT	35.98	31.37	36.30	37.91
	+FT+FM	36.21	31.25	36.37	37.76
	+double tune	36.78	32.10	36.72	37.84
Unconstrained	+fix punctuation	—	36.46	—	—
	IWSLT	33.66	30.58	36.89	37.34
	MTNT+IWSLT	35.44	31.24	36.73	37.90
	ASR($\lambda = 1.5$)	30.53	28.66	36.46	35.61
	ASR($\lambda = 1.2$)	31.09	29.48	36.59	35.54
Mixed training	None	35.32	31.14	35.5	36.05

Table 4: BLEU scores of models fine-tuned on different data in the En→Fr direction. λ in ASR data represents the filtering threshold, as mentioned in Section 2.4. The mixed-training model combines all data available and adds domain tags in front of each sentence. Other notations are same as in Table 3.

performance on clean test sets increases as well. This shows that noisy parallel data could improve model robustness on both noisy and clean texts.

4.2 Data Augmentation

As it is common in the field, we experimented with back translation (third row in Table 3). We used the target-to-source Tune-S model to back-translate monolingual data in MTNT corpus. The back-translated data is combined with the noisy parallel data and used to fine-tune the baseline model. For Fr→En, by introducing back-translated data, the model performance drops by over 2 BLEU scores compared the simply tuning on parallel data. This would suggest that the back

translation data might break the noise level gap between source and target texts, and hence the model fine-tuned on back-translated data tends to output noisier translations and performs worse.

Since the size of the MTNT dataset is too small, we tried merging the data in both language directions, resulting in 55k sentence pairs for fine tuning. For both directions, the models tuned on the merged MTNT data (Tune-B) show worse performance than the models tuned on single direction data (Tune-S). This is due to the introduction of opposite direction data would increase the noise in target texts. We added the forward translation and fuzzy matches data separately and fine-tuned with the merged MTNT data. Results show that the in-

roduction of either forward translation or fuzzy match data would improve model performance on noisy test sets, compared to the Tune-B model. However, with only forward translation or fuzzy match data added, the model still lags behind the Tune-S performance. Therefore, we mix the FT, FM, and merged MTNT data. After we use the mixed data for fine tuning, models in both directions scored better with the augmented data. The Fr→En model with forward translation and fuzzy matches data achieved a performance of 42.80 BLEU score on the MTNT2019 test set, an improvement of +1.32 BLEU points over the Tune-S model. The forward translation data is generated using the Tune-B model, which includes information on the opposite direction, and this might benefit forward translation and prevent the model from overfitting. Compared to back translation, forward translation could keep the noise level difference between the source and target sentences⁹.

4.3 Double Fine Tuning

Considering that the opposite direction data from that of the MTNT dataset would harm the model performance, we applied double fine tuning to compensate. We used the model which had already been tuned on the combination of forward translation data, fuzzy match data and merged MTNT data, and fine-tuned it with the MTNT data on the corresponding direction (e.g. Fr→En data for Fr→En model). In this case, the MTNT data in the same language direction would fine-tune the model twice, thus adapting the model to the specific language direction domain. The second fine tuning was able to further improve model robustness to noisy data and keep a similar performance on clean data. In the En→Fr direction, the second fine tuning improves +0.57 and +0.85 BLEU points on MTNT and MTNT2019.

4.4 Punctuation Fixing

The MTNT2019 test set uses a different set of punctuation in French text as the MTNT dataset and clean training data. In the MTNT2019 test set, the French references use apostrophes (') and angle quotes (« and »), instead of the single quotes (') and double quotes (") used in the MTNT training data (Berard et al., 2019). Therefore, mod-

⁹We experimented with fuzzy matches plus MTNT data (not merged), but it does not improve performance, because without the opposite direction information the model overfits to the tuning data.

els fine-tuned with MTNT training data would show an inconsistent performance for punctuation when evaluating on the MTNT2019 test set. We fixed the punctuation in the En→Fr direction as a postprocessing step. This single replacement improves +4.36 BLEU score over the double fine-tuned model. For comparison, we also postprocessed the output from the Tune-S model. The punctuation fix results in an increase of +3.24 BLEU score.

4.5 External Data

To explore the effect of other types of noise, we fine-tuned our baseline model on different external datasets (see the "Unconstrained" rows in Table 3 and 4). We experimented with human transcript and translation in IWSLT dataset. The BLEU score (Fr→En) on MTNT2019 increases by +2.14 over the baseline, while the results on the other three test sets decrease. In the En→Fr direction, fine tuning on IWSLT improves the model performance on all four test sets, and with MTNT data added, the BLEU score on noisy data performs even better than the Tune-B model. The benefit of speech transcripts might come from informal languages such as slang, spoken language, and domain-related words. Apart from this, we also kept the indicating words (e.g. "[laughter]" and "[applause]") in the transcripts, which could also play a role of noise.

When using ASR data generated from the audio files in the MuST-C dataset, we first filtered ASR data by removing sentences where the original transcript length is over 1.5 times that of the ASR transcript. The model fine-tuned on ASR data shows a slight decrease in BLEU scores. We found that the ASR transcript often skips some phrases in a sentence. Therefore, we reduced the length ratio threshold to 1.2, and with that the model achieves similar performance as the baseline model. Evaluated on *newstest*, the ASR-tuned model improves +1.08 BLEU score over the baseline. Finally, we tried with the parallel texts from human transcript and translation in MuST-C corpus, similar to IWSLT, the performance increases on all test data. This suggests that the introduction of external data with different types of noise could improve model robustness, even without the use of in-domain noisy data.

Finally, we conducted a domain-sensitive training experiment by adding tags for different data.

We mixed all available data, including the MTNT data (two directions), IWSLT, MuST-C, ASR, forward translation, and fuzzy match data. Tags are added at the beginning of the source sentences. As shown in the last row in Table 4, mixed training could improve the performance over baseline on noisy texts. However, the model does not outperform the models with fine tuning. This might result from the introduction of ASR generated data, which can contain more low quality training samples.

4.6 WMT19 Robustness Leaderboard

We submitted our best constrained systems to WMT19 Robustness Leaderboard¹⁰, as shown in Table 5 and Table 6. In the Fr→En direction, we submitted the model fine-tuned on merged MTNT data, forward translation and fuzzy match data (row 5 in Table 3). In the En→Fr direction, the double-tuned model with punctuation fixed was submitted (row 6 in Table 4).

System	BLEU-uncased
(Berard et al., 2019)	48.8
(Helcl et al., 2019)	45.8
(Zheng et al., 2019)	44.5
Ours	43.8
(Post and Duh, 2019)	41.8
(Zhou et al., 2019)	36.0
(Grozea, 2019)	30.8
MTNT paper baseline	26.2

Table 5: WMT19 Robustness Leaderboard on Fr→En.

System	BLEU-uncased
(Berard et al., 2019)	42.0
(Helcl et al., 2019)	39.1
Ours	37.1
(Zheng et al., 2019)	37.0
(Grozea, 2019)	24.8
MTNT paper baseline	22.5

Table 6: WMT19 Robustness Leaderboard on En→Fr.

Our systems would have achieved the 4th and 3rd place on Fr→En and En→Fr directions. The leading systems use back translation on a large volume of clean monolingual data, therefore could benefit from the size of clean data. Although our system does not utilize clean monolingual data,

¹⁰<http://matrix.statmt.org>

we find an alternative way to extend noisy parallel data, which might be more efficient for training. The results show that our systems could achieve a competitive position.

5 Conclusions

In this paper we use data augmentation strategies to improve neural machine translation models robustness. We experiment under the setting of the WMT19 Robustness Task for the Fr↔En language directions. We propose the use of forward translation and fuzzy matches as alternatives to back translation to augment noisy data. Our best models with augmented noisy data could improve +1.32 and +0.97 BLEU scores for Fr→En and En→Fr over models fine-tuned with noisy parallel data. We also explore the effect of external noisy data in the form of speech transcripts and show that models could benefit from data injected with noise through manual transcriptions of spoken language. The ASR generated data does not help improving robustness as it contains low quality training samples that break the sentences, while the human transcripts from speech proved helpful to translate noisy texts, even without in-domain data. Future work might include training a domain-related speech recognition model and generate better ASR parallel data instead of using the off-the-shelf system.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Yonatan Belinkov and Yonatan Bisk. 2017. Synthetic and natural noise both break neural machine translation. *ArXiv*, abs/1711.02173.
- Alexandre Berard, Ioan Calapodescu, and Claude Roux. 2019. [Naver labs europe’s systems for the wmt19 machine translation robustness task](#). In *Proceedings of the Fourth Conference on Machine Translation*, pages 725–731, Florence, Italy. Association for Computational Linguistics.
- Bram Bulte and Arda Tezcan. 2019. [Neural fuzzy repair: Integrating fuzzy matches into neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1800–1809, Florence, Italy. Association for Computational Linguistics.

- Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. [Tagged back-translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 53–63, Florence, Italy. Association for Computational Linguistics.
- Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. Wit³: Web inventory of transcribed and translated talks. In *Proceedings of the 16th Conference of the European Association for Machine Translation (EAMT)*, pages 261–268, Trento, Italy.
- 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.
- 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.
- Raj Dabre and Eiichiro Sumita. 2019. [Nict’s supervised neural machine translation systems for the wmt19 translation robustness task](#). In *Proceedings of the Fourth Conference on Machine Translation*, pages 732–735, Florence, Italy. Association for Computational Linguistics.
- Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. [MuST-C: a Multilingual Speech Translation Corpus](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2012–2017, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1243–1252. JMLR. org.
- Cristian Grozea. 2019. [System description: The submission of fokus to the wmt 19 robustness task](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 537–538, Florence, Italy. Association for Computational Linguistics.
- Jindřich Helcl, Jindřich Libovický, and Martin Popel. 2019. [Cuni system for the wmt19 robustness task](#). In *Proceedings of the Fourth Conference on Machine Translation*, pages 738–742, Florence, Italy. 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. *arXiv preprint arXiv:1902.01509*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proc. ACL*.
- Catherine Kobus, Josep Crego, and Jean Senellart. 2016. Domain control for neural machine translation. *arXiv preprint arXiv:1612.06140*.
- 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.
- Vladimir I Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710.
- Xian Li, Paul Michel, Antonios Anastasopoulos, Yonatan Belinkov, Nadir Durrani, Orhan Firat, Philipp Koehn, Graham Neubig, Juan Pino, and Hassan Sajjad. 2019. [Findings of the first shared task on machine translation robustness](#). In *Proceedings of the Fourth Conference on Machine Translation*, pages 218–229, 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.
- Soichiro Murakami, Makoto Morishita, Tsutomu Hira, and Masaaki Nagata. 2019. [Ntt’s machine translation systems for wmt19 robustness task](#). In *Proceedings of the Fourth Conference on Machine Translation*, pages 743–750, Florence, Italy. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post and Kevin Duh. 2019. [Jhu 2019 robustness task system description](#). In *Proceedings of the Fourth Conference on Machine Translation*, pages 751–757, Florence, Italy. 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.

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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.

Renjie Zheng, Hairong Liu, Mingbo Ma, Baigong Zheng, and Liang Huang. 2019. [Robust machine translation with domain sensitive pseudo-sources: Baidu-osu wmt19 mt robustness shared task system report](#). In *Proceedings of the Fourth Conference on Machine Translation*, pages 758–763, Florence, Italy. 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.