

# Multi-Source Neural Machine Translation with Data Augmentation

Yuta Nishimura<sup>1</sup>, Katsuhito Sudoh<sup>1</sup>, Graham Neubig<sup>2,1</sup>, Satoshi Nakamura<sup>1</sup>

<sup>1</sup>Nara Institute of Science and Technology, 8916-5 Takayama-cho, Ikoma, Nara 630-0192, Japan

<sup>2</sup>Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

{nishimura.yuta.nn9, sudoh, s-nakamura}@is.naist.jp

gneubig@cs.cmu.edu

## Abstract

*Multi-source* translation systems translate from multiple languages to a single target language. By using information from these multiple sources, these systems achieve large gains in accuracy. To train these systems, it is necessary to have corpora with parallel text in multiple sources and the target language. However, these corpora are rarely complete in practice due to the difficulty of providing human translations in *all* of the relevant languages. In this paper, we propose a data augmentation approach to fill such incomplete parts using multi-source neural machine translation (NMT). In our experiments, results varied over different language combinations but significant gains were observed when using a source language similar to the target language.

## 1. Introduction

Machine Translation (MT) systems usually translate one source language to one target language. However, in many real situations, there are multiple languages in the corpus of interest. Examples of this situation include the multilingual official document collections of the European parliament [1] and the United Nations [2]. These documents are manually translated into all official languages of the respective organizations. Many methods have been proposed to use these multiple languages in translation systems to improve the translation accuracy [3, 4, 5, 6]. In almost all cases, multilingual machine translation systems output better translations than one-to-one systems, as the system has access to multiple sources of information to reduce ambiguity in the target sentence structure or word choice.

However, in contrast to the more official document collections mentioned above where it is mandated that all translations in all languages, there are also more informal multilingual captions such as those of talks [7] and movies [8]. Because these are based on voluntary translation efforts, large portions of them are not translated, especially into languages with a relatively small number of speakers.

Nishimura *et al.* [9] have recently proposed a method for multi-source NMT that is able to deal with the case of missing source data encountered in these corpora. The implementation is simple: missing source translations are replaced with a special symbol  $\langle \text{NULL} \rangle$  as shown in Figure 1(a). This

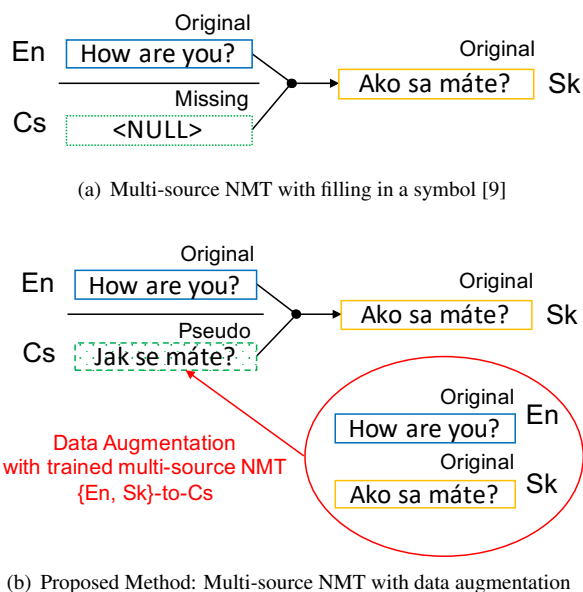


Figure 1: Example of multi-source NMT with an incomplete corpus; The language pair is {English, Czech}-to-Slovak and the translation of Czech is missing.

method allows us to use incomplete corpora both at training time and test time, and multi-source NMT with this method was shown to achieve higher translation accuracy. If the model is trained on corpora with a large number of  $\langle \text{NULL} \rangle$  symbols on the source side, a large number of training examples will be different from test time, when we actually have multiple sources. Thus, these examples will presumably be less useful in training a model intended to do multi-source translation. In this paper, we propose an improved method for utilizing multi-source examples with missing data: using a pseudo-corpus whose missing translations are filled up with machine translation outputs using a trained multi-source NMT system as shown in Figure 1(b). Experimental results show that the proposed method is a more effective method to incorporate incomplete multilingual corpora, achieving improvements of up to about 2 BLEU over the previous method where each missing input sentence is replaced by  $\langle \text{NULL} \rangle$ .

## 2. Related Work

### 2.1. Multi-source NMT

There are two major approaches to multi-source NMT; multi-encoder NMT [10] and mixture of NMT Experts [11]. In this work, we focus on the multi-encoder NMT that showed better performance in most cases in Nishimura *et al.* [9].

Multi-encoder NMT [10] is similar to the standard attentional NMT framework [12] but uses multiple encoders corresponding to the source languages and a single decoder.

Suppose we have two LSTM-based encoders and their hidden and cell states at the end of the inputs are  $h_1$ ,  $h_2$  and  $c_1$ ,  $c_2$ , respectively. Multi-encoder NMT initializes its decoder hidden state  $h$  and cell state  $c$  using these encoder states as follows:

$$h = \tanh(W_c[h_1; h_2]) \quad (1)$$

$$c = c_1 + c_2 \quad (2)$$

Attention is then defined over encoder states at each time step  $t$  and resulting context vectors  $d_t^1$  and  $d_t^2$  are concatenated together with the corresponding decoder hidden state  $h_t$  to calculate the final context vector  $\tilde{h}_t$ .

$$\tilde{h}_t = \tanh(W_c[h_t; d_t^1; d_t^2]) \quad (3)$$

Our multi-encoder NMT implementation is basically similar to the original one [10] but has a difference in its attention mechanism. We use global attention used in Nishimura *et al.* [9], while Zoph and Knight used local-p attention. The global attention allows the decoder to look at everywhere in the input, while the local-p attention forces to focus on a part of the input [13].

### 2.2. Data Augmentation for NMT

Sennrich *et al.* proposed a method to use monolingual training data in the target language for training NMT systems, with no changes to the network architecture [14]. It first trains a seed target-to-source NMT model using a parallel corpus and then translates the monolingual target language sentences into the source language to create a *synthetic* parallel corpus. It finally trains a source-to-target NMT model using the seed and synthetic parallel corpora. This very simple method called *back-translation* makes effective use of available resources, and achieves substantial gains. Imamura *et al.* proposed a method that enhances the encoder and attention using target monolingual corpora by generating multiple source sentences via sampling as an extension of the back-translation [15].

There are also other approaches for data augmentation other than back-translation. Wang *et al.* proposed a method of randomly replacing words in both the source sentence and the target sentence with other random words from their corresponding vocabularies [16]. Kim and Rush proposed a

sequence-level knowledge distillation in NMT that uses machine translation results by a large teacher model to train a small student model as well as ground-truth translations [17].

Our work is an extension of the back-translation approach in multilingual situations by generating pseudo-translations using multi-source NMT.

## 3. Proposed Method

We propose three types of data augmentation for multi-encoder NMT; “fill-in”, “fill-in and replace” and “fill-in and add.” Firstly, we explain about the data requirements and overall framework using Figure 1(b). We used three languages; English, Czech and Slovak. Our goal is to get the Slovak translation, and to do so we take three steps. There are not any missing data in English translations, but Slovak and Czech translations have some missing data. In the first step, we train a multi-encoder NMT model (Source: English and Slovak, Target: Czech) to get Czech pseudo-translations using the baseline method, which is to replace a missing input sentence with a special symbol  $\langle \text{NULL} \rangle$ . In the second step, we create Czech pseudo-translations using multi-encoder NMT which was trained on the first step. We conducted three types of augmentation, which we introduce later. Finally in the third step, we switch the role of Czech and Slovak, in other words, we train a new multi-encoder NMT model (Source: English and Czech, Target: Slovak). At this time, we use Slovak pseudo-translations in the source language side. This method is similar to back-translation but taking advantage of the fact that we have an additional source of knowledge (Czech or Slovak) when trying to augment the other language (Slovak or Czech respectively).

We next introduce three types of augmentation. Figure 2 illustrates their examples in {English, Czech}-to-{Slovak} case where one Czech sentence is missing.

(a) **fill-in**: where only missing parts in the corpus are filled up with pseudo-translations.

(b) **fill-in and replace**: where we both augment the missing part and replace original translations with pseudo-translations in the source language except English whose translations has not any missing data. The motivation behind this method is not to use unreliable translation. Morishita *et al.* [18] demonstrated the effectiveness of applying back-translation for an unreliable part of a provided corpus. Translations of TED talks are from many independent volunteers, so there may be some differences between translations other than original English, or even they may include some free or over-simplified translations. We aim to fill such a gap using data augmentation.

(c) **fill-in and add**: where we both augment the missing part and added pseudo-translations from original translations in the source language except English. This helps prevent introduction of too much noise due to the complete replacement of original translations with pseudo-translations in the second method.

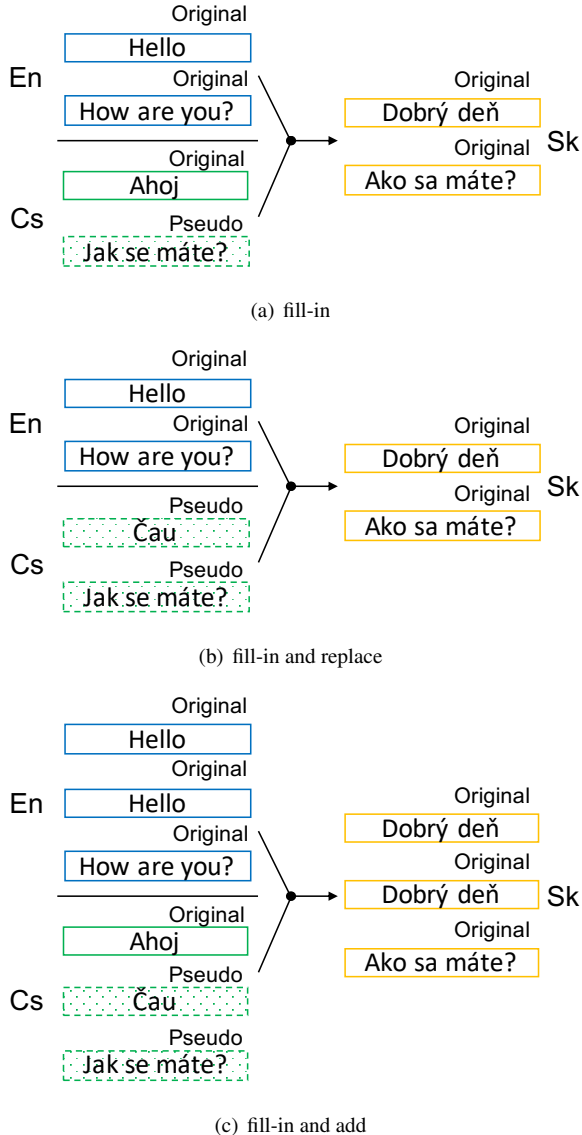


Figure 2: Example of three types of augmentation; Language Pair is {English, Czech}-to-{Slovak} and Czech translation corresponding to “How are you?” is missing. In this example, the dotted background indicates the pseudo-translation produced from multi-source NMT and the white background means the original translation.

## 4. Experiment

We conducted MT experiments to examine the performance of the proposed method using actual multilingual corpora of TED Talks.

### 4.1. Data

We used a collection of transcriptions of TED Talks and their multilingual translations. The numbers of these voluntary translations differs significantly by language. We chose

Table 1: “train” shows the number of available training sentences, and “missing” shows the number and the fraction of missing sentences in comparison with English ones.

Pair	Trg	train	missing
en-hr/sr	hr	118949	35564 (29.9%)
	sr	133558	50203 (37.6%)
en-sk/cs	sk	100600	58602 (57.7%)
	cs	59918	17380 (29.0%)
en-vi/id	vi	160984	87816 (54.5%)
	id	82592	9424 (11.4%)

three different language sets for the experiments: {English (en), Croatian (hr), Serbian (sr)}, {English (en), Slovak (sk), Czech (cs)}, and {English (en), Vietnamese (vi), Indonesian (id)}. Since the great majority of TED talks are in English, the experiments were designed for the translation from English to another language with the help of the other language in the language set, with no missing portions in the English sentences. Table 1 shows the number of training sentences for each language set. At test time, we experiment with a complete corpus with both source sentences represented, as this is the sort of multi-source translation setting that we are aiming to create models for.

### 4.2. Baseline Methods

We compared the proposed methods with the following three baseline methods.

**One-to-one NMT:** a standard NMT model from one source language to another target language. The source language is fixed to English in the experiments. If the target language part is missing in the parallel corpus, such sentences pairs cannot be used in training so they are excluded from the training set.

**Multi-encoder NMT with back-translation:** a multi-encoder NMT system using English-to-X NMT to fill up the missing parts in the other source language X.<sup>1</sup>

**Multi-encoder NMT with  $\langle \text{NULL} \rangle$ :** a multi-encoder NMT system using a special symbol  $\langle \text{NULL} \rangle$  to fill up the missing parts in the other source language X [9].

### 4.3. NMT settings

NMT settings are the same for all the methods in the experiments. We use bidirectional LSTM encoders [12], and global attention and input feeding for the NMT model [13]. The number of dimensions is set to 512 for the hidden and embedding layers. Subword segmentation was applied using SentencePiece [19]. We trained one subword segmentation model for English and another shared between the other two

<sup>1</sup>This is not exactly *back-translation* because the pseudo-translations are not from the target language but from the other source language (English) in our multi-source condition. But we use this familiar term here for simplicity.

languages in the language set because the amount of training data for the languages other than English was small. For parameter optimization, we used Adam [20] with gradient clipping of 5. We performed early stopping, saving parameter values that had the best log likelihoods on the validation data and used them when decoding test data.

#### 4.4. Results

Table 2 shows the results in BLEU [21]. We can see that our proposed methods demonstrate larger gains in BLEU than baseline methods in two language sets: {English, Croatian, Serbian}, {English, Slovak, Czech}. On these pairs, we can say that our proposed method is an effective way for using incomplete multilingual corpora, exceeding other reasonably strong baselines. However, in {English, Vietnamese, Indonesian}, our proposed methods obtained lower scores than the baseline methods. We observed that the improvement by the use of multi-encoder NMT against one-to-one NMT in the baseline was significantly smaller than the other language sets, so multi-encoder NMT was not as effective compared to one-to-one NMT in the first place. Our proposed method is affected by which languages to use, and the proposed method is likely more effective for similar language pairs because the expected accuracy of the pseudo-translation gets better by the help of lexical and syntactic similarity including shared subword entries.

### 5. Discussion

#### 5.1. Different Types of Augmentation

We examined three types of augmentation: “fill-in”, “fill-in and replace”, “fill-in and add”. In Table 2, we can see that there were no significant differences among them, despite the fact that their training data were very different from each other. We conducted additional experiments using incomplete corpora with lower quality augmentation by one-to-one NMT to investigate the differences of the three types of augmentation. We created three types of pseudo-multilingual corpora using back-translation from one-to-one NMT and trained multi-encoder NMT models using them. Our expectation here was that the aggressive use of low quality pseudo-translations caused to contaminate the training data and to decrease the translation accuracy.

Table 3 shows the results. In {English, Croatian, Serbian} and {English, Slovak, Czech}, we obtained significant drop in BLEU scores with the aggressive strategies (“fill-in and replace” and “fill-in and add”), while there are few differences in {English, Vietnamese, Indonesian}. One possible reason is that the quality of pseudo-translations by one-to-one NMT in Indonesian and Vietnamese was better than the other languages; in other words, the BLEU from one-to-one NMT in Table 2 was sufficiently good without multi-source NMT. Thus the translation performance for Croatian, Serbian, Slovak and Czech could not improve in the experiments here due to *noisy* pseudo-translations of those languages.

Contrary, the BLEU from “fill-in and add” was the highest when the target language was Indonesian. We hypothesize that this is due to much smaller fraction of the missing parts in Indonesian corpus as shown in Table 1, so there should be little room for improvement if we fill in only the missing parts even if the accuracy of the pseudo-translations is relatively high.

#### 5.2. Iterative Augmentation

It can be noted that if we have a better multi-source NMT system, it can be used to produce better pseudo-translations. This leads to a natural iterative training procedure where we alternatively update the multi-source NMT systems into the two target languages.

Table 4 shows the results of {English, Croatian, Serbian}. We found that this produced negative results; BLEU decreased gradually in every step. We observed very similar results in the other language pairs, while we omit the actual numbers here. This indicates that the iterative training may be introducing more noise than it is yielding improvements, and thus may be less promising than initially hypothesized.

#### 5.3. Non-parallelism

A problem in the use of multilingual corpora is non-parallelism. In case of TED multilingual captions, they are translated from English transcripts independently by many volunteers, which may cause some differences in details of the translation in the various target languages. For example in {English, Croatian, Serbian}, Croatian and Serbian translations may not be completely parallel. Table 5 shows such an example where the Serbian translation does not have a phrase corresponding to “let me.” This kind of non-parallelism may be resolved by overriding such translations with pseudo-translations with “fill-in and replace” and “fill-in and add”. Here, the Serbian pseudo-translation includes the corresponding phrase “Dozvolite mi” and can be used to compensate for the missing information. This would be one possible reason of the improvements by “fill-up and replace” or “fill-up and add”.

### 6. Conclusions

In this paper, we examined data augmentation of incomplete multilingual corpora in multi-source NMT. We proposed three types of augmentation; fill-in, fill-in and replace, fill-in and add. Our proposed methods proved better than baseline system using the corpus where missing part was filled up with “⟨NULL⟩”, although results depended on the language pair. One limitation in the current experiments with a set of three languages was that missing parts in the test sets could not be filled in. This can be resolved if we use more languages, and we will investigate this in future work.

Table 2: Main results in BLEU for English-Croatian/Serbian (en-hr/sr), English-Slovak/Czech (en-sk/cs), and English-Vietnamese/Indonesian (en-vi/id).

Pair	Trg	baseline method			proposed method		
		one-to-one (En-to-Trg)	multi-encoder NMT (fill up with symbol)	multi-encoder NMT (back translation)	fill-in	fill-in and replace	fill-in and add
en-hr/sr	hr	20.21	28.18	27.57	29.17	29.37	<b>29.40</b>
	sr	16.42	23.85	22.73	24.41	<b>24.96</b>	24.15
en-sk/cs	sk	13.79	20.27	19.83	20.26	20.43	<b>20.59</b>
	cs	14.72	19.88	19.54	20.78	<b>20.90</b>	20.61
en-vi/id	vi	24.60	25.70	26.66	<b>26.73</b>	26.48	26.32
	id	24.89	<b>26.89</b>	26.34	26.40	25.73	26.21

Table 3: The difference of three types of augmentation in BLEU for English-Croatian/Serbian (en-hr/sr), English-Slovak/Czech (en-sk/cs), and English-Vietnamese/Indonesian (en-vi/id). We used one-to-one model to produce pseudo-translations.

Pair	Trg	multi-encoder NMT (back-translation)		
		fill-in	fill-in and replace	fill-in and add
en-hr/sr	hr	<b>27.57</b>	24.05	24.79
	sr	<b>22.73</b>	17.77	22.02
en-sk/cs	sk	<b>19.83</b>	16.75	18.16
	cs	<b>19.54</b>	17.04	18.40
en-vi/id	vi	<b>26.66</b>	26.39	26.65
	id	26.34	23.90	<b>26.67</b>

## 7. Acknowledgements

Part of this work was supported by JSPS KAKENHI Grant Numbers and JP16H05873 and JP17H06101.

## 8. References

- [1] P. Koehn, “Europarl: A Parallel Corpus for Statistical Machine Translation,” in *Conference Proceedings: the tenth Machine Translation Summit*, AAMT. Phuket, Thailand: AAMT, 2005, pp. 79–86.
- [2] M. Ziemski, M. Junczys-Dowmunt, and B. Pouliquen, “The United Nations Parallel Corpus v1.0,” in *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, N. C. C. Chair), K. Choukri, T. Declerck, S. Goggi, M. Grobelnik, B. Maegaard, J. Mariani, H. Mazo, A. Moreno, J. Odijk, and S. Piperidis, Eds. Paris, France: European Language Resources Association (ELRA), May 2016.
- [3] D. Dong, H. Wu, W. He, D. Yu, and H. Wang, “Multi-Task Learning for Multiple Language Translation,” in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*. Beijing, China: Association for Computational Linguistics, July 2015, pp. 1723–1732.
- [4] O. Firat, K. Cho, and Y. Bengio, “Multi-Way, Multilingual Neural Machine Translation with a Shared Attention Mechanism,” in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California: Association for Computational Linguistics, June 2016, pp. 866–875.
- [5] M. Johnson, M. Schuster, Q. V. Le, M. Krikun, Y. Wu, Z. Chen, N. Thorat, F. Vidas, M. Wattenberg, G. Corrado, M. Hughes, and J. Dean, “Google’s Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation,” *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 339–351, 2017.
- [6] T.-L. Ha, J. Niehues, and A. Waibel, “Toward Multilingual Neural Machine Translation with Universal Encoder and Decoder,” in *Proceedings of the 13th International Workshop on Spoken Language Translation*, Seattle, Washington, December 2016.
- [7] M. Cettolo, C. Girardi, and M. Federico, “WIT<sup>3</sup>: Web Inventory of Transcribed and Translated Talks,” in *Proceedings of the 16th EAMT Conference*, May 2012, pp. 261–268.
- [8] J. Tiedemann, “News from OPUS - A Collection of Multilingual Parallel Corpora with Tools and Interfaces,” in *Recent Advances in Natural Language Processing*, N. Nicolov, K. Bontcheva, G. Angelova, and R. Mitkov, Eds. Borovets, Bulgaria: John Benjamins, Amsterdam/Philadelphia, 2009, vol. V, pp. 237–248.
- [9] Y. Nishimura, K. Sudoh, G. Neubig, and S. Nakamura, “Multi-Source Neural Machine Translation with Missing Data,” in *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*. Association for Computational Linguistics, July 2018, pp. 92–99.

Table 4: BLEU (and BLEU gains compared to step 1) in each step of iterative augmentation.

Pair	Trg	step 1	step 2	step 3	step 4
en-hr/sr	hr	29.17 (+0.00)	29.03 (-0.14)	29.10 (-0.07)	29.05 (-0.12)
	sr	24.41 (+0.00)	24.18 (-0.23)	24.17 (-0.24)	23.95 (-0.46)

Table 5: Example of the Serbian pseudo-translation. This pseudo-translation is the output of {English, Croatian}-to-Serbian translation.

Type	Sentence
Original (En)	So <b>let me</b> conclude with just a remark to bring it back to the theme of choices.
Original (Sr)	Da zakljuim jednom konstatacijom kojom se vraam na temu izbora.
Pseudo (Sr)	<b>Dozvolite mi</b> da zakljuim samo jednom opaskom, da se vratim na temu izbora.

- [10] B. Zoph and K. Knight, “Multi-Source Neural Translation,” in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California: Association for Computational Linguistics, June 2016, pp. 30–34.
- [11] E. Garmash and C. Monz, “Ensemble Learning for Multi-Source Neural Machine Translation,” in *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. Osaka, Japan: The COLING 2016 Organizing Committee, December 2016, pp. 1409–1418.
- [12] D. Bahdanau, K. Cho, and Y. Bengio, “Neural Machine Translation by Jointly Learning to Align and Translate,” in *Proceedings of the 3rd International Conference on Learning Representations*, May 2015.
- [13] T. Luong, H. Pham, and C. D. Manning, “Effective Approaches to Attention-based Neural Machine Translation,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, September 2015, pp. 1412–1421.
- [14] R. Sennrich, B. Haddow, and A. Birch, “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)*. Association for Computational Linguistics, 2016, pp. 86–96.
- [15] K. Imamura, A. Fujita, and E. Sumita, “Enhancement of encoder and attention using target monolingual corpora in neural machine translation,” in *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*. Association for Computational Linguistics, 2018, pp. 55–63. [Online]. Available: <http://aclweb.org/anthology/W18-2707>
- [16] X. Wang, H. Pham, Z. Dai, and G. Neubig, “Switchout: an efficient data augmentation algorithm for neural machine translation,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.
- [17] Y. Kim and A. M. Rush, “Sequence-level knowledge distillation,” in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, November 2016, pp. 1317–1327. [Online]. Available: <https://aclweb.org/anthology/D16-1139>
- [18] M. Morishita, J. Suzuki, and M. Nagata, “NTT Neural Machine Translation Systems at WAT 2017,” in *Proceedings of the 4th Workshop on Asian Translation (WAT2017)*, 2017, pp. 89–94.
- [19] T. Kudo, “Subword regularization: Improving neural network translation models with multiple subword candidates,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2018, pp. 66–75.
- [20] D. P. K. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” in *Proceedings of the 3rd International Conference on Learning Representations*, May 2015.
- [21] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “BLEU: a Method for Automatic Evaluation of Machine Translation,” in *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, Philadelphia, July 2002, pp. 311–318.