

Overcoming the Rare Word Problem for Low-Resource Language Pairs in Neural Machine Translation

Thi-Vinh Ngo

Thai Nguyen University
ntvinh@ictu.edu.vn

Phuong-Thai Nguyen
Vietnam National University
thainp@vnu.edu.vn

Thanh-Le Ha

Karlsruhe Institute of Technology
thanh-le.ha@kit.edu

Le-Minh Nguyen
JAIST, Japan
nguyenml@jaist.ac.jp

Abstract

Among the six challenges of neural machine translation (NMT) coined by (Koehn and Knowles, 2017), rare-word problem is considered the most severe one, especially in translation of low-resource languages. In this paper, we propose three solutions to address the rare words in neural machine translation systems. First, we enhance source context to predict the target words by connecting directly the source embeddings to the output of the attention component in NMT. Second, we propose an algorithm to learn morphology of unknown words for English in supervised way in order to minimize the adverse effect of rare-word problem. Finally, we exploit synonymous relation from the WordNet to overcome out-of-vocabulary (OOV) problem of NMT. We evaluate our approaches on two low-resource language pairs: English-Vietnamese and Japanese-Vietnamese. In our experiments, we have achieved significant improvements of up to roughly +1.0 BLEU points in both language pairs.

1 Introduction

NMT systems have achieved better performance compared to statistical machine translation (SMT) systems in recent years not only on available data language pairs (Sennrich et al., 2016a; Cho et al., 2016), but also on low-resource language pairs (Nguyen and Chiang, 2017; Cettolo et al., 2016). Nevertheless, NMT still exists many challenges which have adverse effects on its effectiveness (Koehn and Knowles, 2017). One of these challenges is that NMT has biased tend in translating high-frequency words, thus words which have lower frequencies are often translated incorrectly. This challenge has also been confirmed again in (Nguyen and Chiang, 2017), and they have proposed two strategies to tackle this problem with modifications on the model’s output distribution:

one for normalizing some matrices by fixing them to constants after several training epochs and another for adding a direct connection from source embeddings through a simple feed forward neural network (FFNN). These approaches increase the size and the training time of their NMT systems. In this work, we follow their second approach but simplify the computations by replacing FFNN with two single operations.

Despite above approaches can improve the prediction of rare words, however, NMT systems often use limited vocabularies in their sizes, from 30K to 80K most frequent words of the training data, in order to reduce computational complexity and the sizes of the models (Bahdanau et al., 2015; Luong et al., 2015b), so the rare-word translation are still problematic in NMT. Even when we use a larger vocabulary, this situation still exists (Jean et al., 2015). A word which has not seen in the vocabulary of the input text (called *unknown word*) are presented by the *unk* symbol in NMT systems. Inspired by alignments and phrase tables in phrase-based machine translation (SMT) as suggested by (Koehn et al., 2007), (Luong et al., 2015b) proposed to address OOV words using an annotated training corpus. They then used a dictionary generated from alignment model or maps between source and target words to determine the translations of *unks* if translations are not found. (Sennrich et al., 2016b) proposed to reduce unknown words using Gage’s Byte Pair Encoding (BPE) algorithm (Gage, 1994), but NMT systems are less effective for low-resource language pairs due to the lack of data and also for other languages that sub-word are not the optimal translation unit. In this paper, we employ several techniques inspired by the works from NMT and the traditional SMT mentioned above. Instead of a loosely unsupervised approach, we suggest a supervised approach to solve this trouble using syn-

onymous relation of word pairs from WordNet on Japanese→Vietnamese and English→Vietnamese systems. To leverage effectiveness of this relation in English, we transform variants of words in the source texts to their original forms by separating their affixes collected by hand.

Our contributes in this work are:

- We release the state-of-the-art for Japanese-Vietnamese NMT systems.
- We proposed the approach to deal with the rare word translation by integrating source embeddings to the attention component of NMT.
- We present a supervised algorithm to reduce the number of unknown words for the English→Vietnamese translation system.
- We demonstrate the effectiveness of leveraging linguistic information from WordNet to alleviate the rare-word problem in NMT.

2 Neural Machine Translation

Our NMT system use a bidirectional recurrent neural network (biRNN) as an encoder and a single-directional RNN as a decoder with input feeding of (Luong et al., 2015a) and the attention mechanism of (Bahdanau et al., 2015). The Encoder’s biRNN are constructed by two RNNs with the hidden units in the LSTM cell, one for forward and the other for backward of the source sentence $\mathbf{x} = (x_1, \dots, x_n)$. Every word x_i in sentence is first encoded into a continuous representation $E_s(x_i)$, called the source embedding. Then \mathbf{x} is transformed into a fixed-length hidden vector \mathbf{h}_i representing the sentence at the time step i , which called the annotation vector, combined by the states of forward $\vec{\mathbf{h}}_i$ and backward $\overleftarrow{\mathbf{h}}_i$:

$$\begin{aligned}\vec{\mathbf{h}}_i &= f(E_s(x_i), \vec{\mathbf{h}}_{i-1}) \\ \overleftarrow{\mathbf{h}}_i &= f(E_s(x_i), \overleftarrow{\mathbf{h}}_{i+1})\end{aligned}$$

The decoder generates the target sentence $\mathbf{y} = (y_1, \dots, y_m)$, and at the time step j , the predicted probability of the target word y_j is estimated as follows:

$$p(y_j|y_{<j}, \mathbf{x}) \propto \text{softmax}(\mathbf{W}\mathbf{z}_j + \mathbf{b})$$

where \mathbf{z}_j is the output hidden states of the attention mechanism and computed by the previous output hidden states \mathbf{z}_{j-1} , the embedding of previous target word $E_t(y_{j-1})$ and the context \mathbf{c}_j :

$$\mathbf{z}_j = g(E_t(y_{j-1}), \mathbf{z}_{j-1}, \mathbf{c}_j)$$

The source context \mathbf{c}_j is the weighted sum of the encoder’s annotation vectors \mathbf{h}_i :

$$\mathbf{c}_j = \sum_{i=1}^n \alpha_{ij} \mathbf{h}_i$$

where α_{ij} are the alignment weights, denoting the relevance between the current target word y_j and all source annotation vectors \mathbf{h}_i .

3 Rare Word translation

In this section, we present the details about our approaches to overcome the rare word situation. While the first strategy augments the source context to translate low-frequency words, the remaining strategies reduce the number of OOV words in the vocabulary.

3.1 Low-frequency Word Translation

The attention mechanism in RNN-based NMT maps the target word into source context corresponding through the annotation vectors \mathbf{h}_i . In the recurrent hidden unit, \mathbf{h}_i is computed from the previous state \mathbf{h}_{t-1} . Therefore, the information flow of the words in the source sentence may be diminished over time. This leads to the accuracy reduction when translating low-frequency words, since there is no direct connection between the target word and the source word. To alleviate the adverse impact of this problem, (Nguyen and Chiang, 2017) combined the source embeddings with the predictive distribution over the output target word in several following steps:

Firstly, the weighted average vector of the source embeddings is computed as follows:

$$\mathbf{l}_t = \tanh \sum_e \alpha_j(e) \mathbf{f}_e$$

where $\alpha_j(e)$ are alignment weights in the attention component and $f_e = E_s(x)$, are the embeddings of the source words.

Then \mathbf{l}_j is transformed through one-hidden-layer FFNN with residual connection proposed by (He et al., 2015):

$$\mathbf{t}_j = \tanh(\mathbf{W}_l \mathbf{l}_j) + \mathbf{l}_j$$

Finally, the output distribution over the target word is calculated by:

$$p(y_j|y_{<j}, \mathbf{x}) = \text{softmax}(\mathbf{W}\mathbf{z}_j + \mathbf{b} + \mathbf{W}_t \mathbf{t}_j + \mathbf{b}_t)$$

The matrices \mathbf{W}_l , \mathbf{W}_t and \mathbf{b}_t are trained together with other parameters of the NMT model.

This approach improves the performance of the NMT systems but introduces more computations as the model size increase due to the additional parameters \mathbf{W}_l , \mathbf{W}_t and \mathbf{b}_t . We simplify this method by using the weighted average of source embeddings directly in the softmax output layer:

$$p(y_j|y_{<j}, \mathbf{x}) = \text{softmax}(\mathbf{W}(\mathbf{z}_j + \mathbf{l}_j) + \mathbf{b})$$

Our method does not learn any additional parameters. Instead, it requires the source embedding size to be compatible with the decoder’s hidden states. With the additional information provided from the source embeddings, we achieve similar improvements compared to the more expensive method described in (Nguyen and Chiang, 2017).

3.2 Reducing Unknown Words

In our previous experiments for English→Vietnamese, BPE algorithm (Sennrich et al., 2016b) applied to the source side does not significantly improves the systems despite it is able to reduce the number of unknown English words. We speculate that it might be due to the morphological differences between the source and the target languages (English and Vietnamese in this case). The unsupervised way of BPE while learning sub-words in English thus might be not explicit enough to provide the morphological information to the Vietnamese side. In this work, we would like to attempt a more explicit, supervised way. We collect 52 popular affixes (prefixes and suffixes) in English and then apply the separating affixes algorithm (called *SAA*) to reduce the number of unknown words as well as to force our NMT systems to learn better morphological mappings between two languages.

The main ideal of our *SAA* is to separate affixes of unknown words while ensuring that the rest of them still exists in the vocabulary. Let the vocabulary V containing K most frequency words from the training set $T1$, a set of prefixes P , a set of suffixes S , we call word w' is the rest of an unknown word or rare word w after delimiting its affixes. We iteratively pick a w from N words (including unknown words and rare words) of the source text $T2$ to consider if w starts with a prefix p in P or ends with a suffix s in S , we then determine splitting its affixes if w' in V . A rare word in V also can be separated its affixes if its frequency is less than the given threshold. We set this threshold by 2 in our experiments. Similarly to BPE approach,

we also employ a pair of the special symbol @ for separating affixes from the word. Listing 3.2 shows our *SAA* algorithm.

```
Input: T1, T2, P, S, threshold=1
Output: the output text T

V = get_most_frequency_K_words(T1)
N = get_words_from_the_source_text(T2)
T = T2
```

```
for each word w in N:
  if w not in V or freq(w) <= threshold:
    for each prefix p in P:
      w1 = separate_prefix(p)
      if w1 != w and w1 in V:
        T = replace(T, w, w1, p)
        break
    for each suffix s in S:
      w2 = separate_suffix(s)
      if w2 != w1 and w2 in V:
        T = replace(T, w2, w1, s)
        break
return T
```

```
Example: intercepted -> intercept @@ed
         impulsively -> impulsive @@ly
         overlooks -> over@@ look @@s
         disowned -> dis@@ own @@ed
```

The proposed *SAA* for separating affixes from words.

3.3 Dealing with OOV using WordNet

WordNet is a lexical database grouping words into sets which share some semantic relations. Its version for English is proposed for the first time by (Fellbaum, 1998). It becomes a useful resource for many tasks of natural language processing (Kolte and Bhirud, 2008; Méndez O., 2013; Gao et al., 2014). WordNet are available mainly for English and German, the version for other languages are being developed including some Asian languages in such as Japanese, Chinese, Indonesian and Vietnamese. Several works have employed WordNet in SMT systems (Khodak et al., 2017; Arcan et al., 2019) but to our knowledge, none of the work exploits the benefits of WordNet in order to ease the rare word problem in NMT. In this work, we propose the learning synonymous algorithm (called *LSW*) from the WordNet of English and Japanese to handle unknown words in our NMT systems.

In WordNet, synonymous words are organized in groups which are called synsets. Our aim is to replace an OOV word by its synonym which appears in the vocabulary of the translation system. From the training set of the source language $T1$, we extract the vocabulary V in size of K most frequent words. For each OOV word from $T1$, we learn its synonyms which exist in

the V from the WordNet W . The synonyms are then arranged in the descending order of their frequencies to facilitate selection of the n best words which have the highest frequencies. The output file C of the algorithm contains OOV words and its corresponding synonyms and then it is applied to the input text $T2$. We also utilize a frequency threshold for rare words in the same way as in SAA algorithm. In practice, we set this threshold as 0, meaning no words on V is replaced by its synonym. If a source sentence has m unknown words and each of them has n best synonyms, it would generate m^n sentences. Translation process allow us to select the best hypothesis based on their scores. Because of each word in the WordNet can belong to many synsets with different meanings, thus an inappropriate word can be placed in the current source context. We will solve this situation in the further works. Our systems only use 1-best synonym for each OOV word. Listing 3.3 presents the *LSW* algorithm.

```

Input: T1, T2,  $W_s$ , threshold=1
Output: - C: The list contains
         synonymous words for OOV words.
        - T: The input of the
         translation systems

def learn_synonym()
V=get_most_frequency_K_words(T1)
N=get_words_from_the_source_text(T2)
C={}
for each word w in N:
    if w not in V or freq(w) <= threshold:
        I=get_synonyms_from_WordNet(w,  $W_s$ )
        for each i in I:
            if i not in V:
                I=I \ {i} #remove i from I
        sort_words_by_descend_of_frequency(I)
        C = C ∪ {w, I}
return C

n_best=3
apply_to_input_file(C, T2, n_best)

```

The LSW learns synonymous words from WordNet.

4 Experiments

We evaluate our approaches on the English-Vietnamese and the Japanese-Vietnamese translation systems. Translation performance is measured in BLEU (Kishore Papineni and Zhu, 2012) by the multi-BLEU scripts from Moses¹.

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

4.1 Datasets

We consider two low-resource language pairs: Japanese-Vietnamese and English-Vietnamese. For Japanese-Vietnamese, we use the TED data provided by WIT3 (Cettolo et al., 2012) and compiled by (Ngo et al., 2018). The training set includes 106758 sentence pairs, the validation and test sets are *dev2010* (568 pairs) and *tst2010* (1220 pairs). For English→Vietnamese, we use the dataset from IWSLT 2015 (Mauro Cettolo and Federico, 2015) with around 133K sentence pairs for the training set, 1553 pairs in *tst2012* as the validation and 1268 pairs in *tst2013* as the test sets.

For *LSW* algorithm, we crawled pairs of synonymous words from Japanese-English WordNet² and achieved 315850 pairs for English and 1419948 pairs for Japanese.

4.2 Preprocessing

For English and Vietnamese, we tokenized the texts and then true-cased the tokenized texts using Moses script. We do not use any word segmentation tool for Vietnamese. For comparison purpose, Sennrich’s BPE algorithm is applied for English texts. Following the same preprocessing steps for Japanese (*JPBPE*) in (Ngo et al., 2018), we use KyTea³ (Neubig et al., 2011) to tokenize texts and then apply BPE on those texts. The number of BPE merging operators are 50k for both Japanese and English.

4.3 Systems and Training

We implement our NMT systems using *OpenNMT-py* framework⁴ (Klein et al., 2017) with the same settings as in (Ngo et al., 2018) for our baseline systems. Our system are built with two hidden layers in both encoder and decoder, each layer has 512 hidden units. In the encoder, a BiLSTM architecture is used for each layer and in the decoder, each layer are basically an LSTM layer. The size of embedding layers in both source and target sides is also 512. Adam optimizer is used with the initial learning rate of 0.001 and then we apply learning rate annealing. We train our systems for 16 epochs with the batch size of 32. Other parameters are the same as the default settings of *OpenNMT-py*.

²<http://compling.hss.ntu.edu.sg/wnja/>

³<http://www.phontron.com/kytea/>

⁴<https://github.com/OpenNMT/OpenNMT-py>

No.	Systems	Japanese→Vietnamese	
		dev2010	tst2010
(1)	Baseline	7.91	9.42
(2)	+ Source Embedding	7.77	9.96
(3)	+ LSW	8.37	10.34
(4)	JPBPE+VNBPE at Ngo et al (2018)	7.77	9.04
(5)	JPBPE+VNBPE + BT + Mixsource at Ngo et al (2018)	8.56	9.64

No.	Systems	Vietnamese→Japanese	
		dev2010	tst2010
(1)	Baseline	9.53 (9.53)	10.95 (10.99)
(2)	+ Source Embedding	10.51 (10.51)	11.37 (11.39)
(3)	JPBPE+VNBPE at Ngo et al (2018)	9.74	11.13

Table 1: Results of Japanese-Vietnamese NMT systems

We then modify the baseline architecture with the alternative proposed in Section 3.1 in comparison to our baseline systems. All settings are the same as the baseline systems.

4.4 Results

In this section, we show the effectiveness of our methods on two low-resource language pairs and compare them to the other works. The empirical results are shown in Table 1 for Japanese-Vietnamese and in Table 3 for English-Vietnamese. Note that, the Multi-BLEU is only measured in the Japanese→Vietnamese direction and the standard BLEU points are written in brackets.

4.4.1 Japanese-Vietnamese Translation

We conduct two out of the three proposed approaches for Japanese-Vietnamese translation systems and the results are given in the Table 1.

Baseline Systems. We find that our translation systems which use Sennrich’s BPE method for Japanese texts and do not use word segmentation for Vietnamese texts are neither better or insignificant differences compare to those systems used word segmentation in (Ngo et al., 2018). Particularly, we obtained +0.38 BLEU points between (1) and (4) in the Japanese→Vietnamese and -0.18 BLEU points between (1) and (3) in the Vietnamese→Japanese.

Our Approaches. On the systems trained with the modified architecture mentioned in the section 3.1, we obtained an improvements of +0.54 BLEU points in the Japanese→Vietnamese and +0.42 BLEU points on the Vietnamese→Japanese compared to the baseline systems.

Due to the fact that Vietnamese WordNet is not available, we only exploit WordNet to tackle unknown words of Japanese texts in our

Japanese→Vietnamese translation system. After using Kytea, Japanese texts are applied *LSW* algorithm to replace OOV words by their synonyms. We choose 1-best synonym for each OOV word. Table 2 shows the number of OOV words replaced by their synonyms. The replaced texts are then BPEd and trained on the proposed architecture. The largest improvement is +0.92 between (1) and (3). We observed an improvement of +0.7 BLEU points between (3) and (5) without using data augmentation described in (Ngo et al., 2018).

	Train	dev2010	tst2010
Number of words	1015	36	25

Table 2: The number of Japanese OOV words replaced by their synonyms.

4.4.2 English-Vietnamese Translation

We examine the effect of all approaches presented in Section 3 for our English-Vietnamese translation systems. Table 3 summarizes those results and the scores from other systems (Nguyen and Chiang, 2017; Huang et al., 2018).

Baseline systems. After preprocessing data using Moses scripts, we train the systems of English↔Vietnamese on our baseline architecture. Our translation system obtained +0.82 BLEU points compared to (Nguyen and Chiang, 2017) in the English→Vietnamese and this is lower than the system of (Huang et al., 2018) with neural phrase-based translation architecture.

Our approaches. The datasets from the baseline systems are trained on our modified NMT architecture. The improvements can be found as +0.55 BLEU points between (1) and (2) in the English→Vietnamese and +0.45 BLEU points (in *tst2012*) between (1) and (2) in the Vietnamese→English.

No.	Systems	English→Vietnamese	
		tst2012	tst2013
(1)	Baseline	26.91 (24.39)	29.86 (27.52)
(2)	+ Source Embedding	27.41 (24.92)	30.41 (28.05)
(3)	+ Sennrich’s BPE	26.96 (24.46)	30.10 (27.84)
(4)	+ SAA	27.16 (24.67)	30.60 (28.34)
(5)	+ LSW	27.46 (24.99)	30.85 (28.54)
(6)	Nguyen and Chiang (2017)	-	26.7
(7)	Huang et al (2018)	-	28.07

No.	Systems	Vietnamese→English	
		tst2012	tst2013
(1)	Baseline	27.97 (28.52)	30.07 (29.89)
(2)	+ Source Embedding	28.42 (29.04)	30.12 (29.93)

Table 3: Results of English-Vietnamese NMT systems

For comparison purpose, English texts are split into sub-words using Sennrich’s BPE methods. We observe that, the achieved BLEU points are lower. Therefore, we then apply the *SAA* algorithm on the English texts from (2) in the English→Vietnamese. The number of applied words are listed in Table 4. The improvement in BLEU are +0.74 between (4) and (1).

	Train	tst2012	tst2013
Number of words	5342	84	93

Table 4: The number of rare words in which their affixes are detached from the English texts in the *SAA* algorithm.

Similarly to the Japanese→Vietnamese system, we apply *LSW* algorithm on the English texts from (4) while selecting 1-best synonym for each OOV word. The number of replaced words on English texts are indicated in the Table 5. Again, we obtained a bigger gain of +0.99 (+1.02) BLEU points in English→Vietnamese direction. Compared to the most recent work (Huang et al., 2018), our system reports an improvement of +0.47 standard BLEU points on the same dataset.

	Train	tst2012	tst2013
Number of words	1889	37	41

Table 5: The number of English OOV words are replaced by their synonyms.

We investigate some examples of translations generated by the English→Vietnamese systems with our proposed methods in the Table 6. The bold texts in red color present correct or approximate translations while the italic texts in gray color denote incorrect translations. The first example, we consider two words: *presentation* and

the unknown word *applauded*. The word *presentation* is predicted correctly as "*bài thuyết trình*" in most cases when we combined source context through embeddings. The unknown word *applauded* which has not seen in the vocabulary is ignored in the first two cases (baseline and source embedding) but it is roughly translated as "*hoan nghênh*" in the *SAA* because it is separated into *applaud* and *ed*. In the second example, we observe the translations of the unknown word *tryout*, they are mistaken in the first three cases but in the *LSW*, it is predicted with a closer meaning as "*bài kiểm tra*" due to the replacement by its synonymous word as *test*.

5 Related Works

Addressing unknown words was mentioned early in the Statistical Machine Translation (SMT) systems. Some typical studies as: (Habash, 2008) proposed four techniques to overcome this situation by extend the morphology and spelling of words or using a bilingual dictionary or transliterating for names. These approaches are difficult when manipulate to different domains. (Trieu, 2016) trained word embedding models to learn word similarity from monolingual data and an unknown word are then replaced by a its similar word. (Madhyastha and España Bonet, 2017) used a linear model to learn maps between source and target spaces base on a small initial bilingual dictionary to find the translations of source words. However, in NMT, there are not so many works tackling this problem. (Jean et al., 2015) use a very large vocabulary to solve unknown words. (Luong et al., 2015b) generate a dictionary from alignment data based on annotated corpus to decide the hypotheses of unknown words. (Nguyen and Chiang, 2017) have introduced the solutions for dealing

Source	which presentation have you applauded the most this morning ?
Reference	bài thuyết trình nào bạn vỗ tay nhiều nhất trong sáng nay ?
Baseline	<i>điều này</i> có thể diễn ra trong buổi sáng hôm nay ?
+Source Embedding	bài thuyết trình nào có thể tạo ra buổi sáng hôm nay ?
+SAA	bài thuyết trình này có hoan nghênh buổi sáng hôm nay không ?
+LSW	<i>điều gì</i> đã diễn ra với bạn buổi sáng hôm nay ?
Source	I started this as a tryout in Esperance , in Western Australia .
Reference	tôi đã bắt đầu như một sự thử nghiệm tại Esperance , tây Úc .
Baseline	tôi bắt đầu như thế này như là một <i>người đàn ông</i> , ở phương Tây Úc .
+Source Embedding	tôi đã bắt đầu điều này như là một <i>người đàn ông</i> ở ven biển ở Tây Úc .
+SAA	tôi đã bắt đầu như thế này với tư cách là một <i>người đàn ông</i> trong lĩnh vực này, ở Tây Úc .
+LSW	tôi bắt đầu thí nghiệm này như một bài kiểm tra ở Quảng trường , ở Tây Úc .

Table 6: Examples of outputs from the English→Vietnamese translation systems with the proposed methods.

with the rare word problem, however, their models require more parameters, thus, decreasing the overall efficiency.

In another direction, (Sennrich et al., 2016b) exploited the BPE algorithm to reduce number of unknown words in NMT and achieved significant efficiency on many language pairs. The second approach presented in this works follows this direction when instead of using an unsupervised method to split rare words and unknown words into sub-words that are able to translate, we use a supervised method. Our third approach using WordNet can be seen as a smoothing way, when we use the translations of the synonymous words to approximate the translation of an OOV word. Another work followed this direction is worth to mention is (Niehues et al., 2016), when they use the morphological and semantic information as the factors of the words to help translating rare words.

6 Conclusion

In this study, we have proposed three difference strategies to handle rare words in NMT, in which the combination of methods brings significant improvements to the NMT systems on two low-resource language pairs. In future works, we will consider selecting some appropriate synonymous words for the source sentence from n-best synonymous words to further improve the performance of the NMT systems and leverage more unsupervised methods based on monolingual data to address rare word problem.

7 Acknowledgments

This work is supported by the project "Building a machine translation system to support translation of documents between Vietnamese and Japanese to help managers and businesses in Hanoi approach to Japanese market", No. TC.02-2016-03.

References

- Mihael Arcan, John P. McCrae, and Paul Buitelaar. 2019. [Polylingual wordnet](#). *CoRR*, abs/1903.01411.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). *Proceedings of International Conference on Learning Representations*.
- M Cettolo, J Niehues, S Stüker, L Bentivogli, R Cattoni, and M Federico. 2016. The IWSLT 2016 Evaluation Campaign. In *Proceedings of the 13th International Workshop on Spoken Language Translation (IWSLT 2016)*, Seattle, WA, USA.
- 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.
- Eunah Cho, Jan Niehues, Thanh-Le Ha, Matthias Sperber, Mohammed Mediani, and Alex Waibel. 2016. Adaptation and combination of nmt systems: The kit translation systems for iwslt 2016. In *Proceedings of the ninth International Workshop on Spoken Language Translation (IWSLT)*, Seattle, WA.
- Christiane Fellbaum. 1998. Wordnet: An electronic lexical database. In *Bradford Books*.
- Philip Gage. 1994. A new algorithm for data compression. In *C Users J.*, 12(2):23–38, February.
- Ningning Gao, Wanli Zuo, Yaokang Dai, and Wei Lv. 2014. [Word sense disambiguation using wordnet semantic knowledge](#). *Advances in Intelligent Systems and Computing*, 278:147–156.
- Nizar Habash. 2008. Four techniques for online handling of out-of-vocabulary words in arabic english statistical machine translation. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics*, pages 57–60.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.

- Po-Sen Huang, Chong Wang, Sitao Huang, Dengyong Zhou, and Li Deng. 2018. [Towards neural phrase-based machine translation](#). In *International Conference on Learning Representations*.
- Sebastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On using very large target vocabulary for neural machine translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 1–10. Association for Computational Linguistics.
- Mikhail Khodak, Andrej Risteski, Christiane Fellbaum, and Sanjeev Arora. 2017. [Extending and improving wordnet via unsupervised word embeddings](#). *CoRR*, abs/1705.00217.
- Todd Ward Kishore Papineni, Salim Roukos and Weijing Zhu. 2012. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. Opennmt: Open-source toolkit for neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics-System Demonstrations*, pages 67–72, Vancouver, Canada, July 30 - August 4, 2017. Association for Computational Linguistics.
- 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.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. *Proceedings of the First Workshop on Neural Machine Translation*, abs/1706.03872:28–39.
- S. G. Kolte and S. G. Bhirud. 2008. Word sense disambiguation using wordnet domains. In *2008 First International Conference on Emerging Trends in Engineering and Technology*, pages 1187–1191.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015a. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421.
- Thang Luong, Ilya Sutskever, Quoc V. Le, Oriol Vinyals, and Wojciech Zaremba. 2015b. Addressing the rare word problem in neural machine translation. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, abs/1410.8206:11–19.
- P.S. Madhyastha and C. España Bonet. 2017. Learning bilingual projections of embeddings for vocabulary expansion in machine translation. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*. pp.139-145. Association for Computational Linguistics, Vancouver, Canada (Aug 2017).
- Sebastian Stüker Luisa Bentivogli Roldano Cattoni Mauro Cettolo, Jan Niehues and Marcello Federico. 2015. The iwslt 2015 evaluation campaign. In *International Conference on Spoken Language*.
- Moreno-Armendáriz M.A. Méndez O., Calvo H. 2013. A reverse dictionary based on semantic analysis using wordnet. In *Advances in Artificial Intelligence and Its Applications, MICAI 2013*. Lecture Notes in Computer Science, vol 8265. Springer, Berlin, Heidelberg.
- Graham Neubig, Yosuke Nakata, and Shinsuke Mori. 2011. [Pointwise prediction for robust, adaptable japanese morphological analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers - Volume 2, HLT '11*, pages 529–533, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Thi-Vinh Ngo, Thanh-Le Ha, Phuong-Thai Nguyen, and Le-Minh Nguyen. 2018. Combining advanced methods in japanese-vietnamese neural machine translation. *2018 10th International Conference on Knowledge and Systems Engineering (KSE)*, pages 318–322.
- Toan Q. Nguyen and David Chiang. 2017. Improving lexical choice in neural machine translation. *Proceedings of NAACL-HLT 2018*, pages 334–343.
- Jan Niehues, Thanh-Le Ha, Eunah Cho, and Alex Waibel. 2016. Using factored word representation in neural network language models. In *Proceedings of the First Conference on Machine Translation: Volume 1, Research Papers*, pages 74–82.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Edinburgh Neural Machine Translation Systems for WMT'16. *arXiv preprint arXiv:1606.02891*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. Neural Machine Translation of Rare Words with Subword Units. In *Association for Computational Linguistics (ACL 2016)*.
- Nguyen L. M. Nguyen P. T. Trieu, H. L. 2016. Dealing with out-of-vocabulary problem in sentence alignment using word similarity. In *Proceedings of the 30th Pacific Asia Conference on Language, Information and Computation: Oral Papers (pp. 259-266)*.