

Learning Source Phrase Representations for Neural Machine Translation

Hongfei Xu^{1,2} Josef van Genabith^{1,2*} Deyi Xiong³ Qiuhui Liu⁴ Jingyi Zhang²

¹Saarland University / Saarland, Germany

²German Research Center for Artificial Intelligence / Saarland, Germany

³Tianjin University / Tianjin, China

⁴China Mobile Online Services / Henan, China

hfxunlp@foxmail.com, Josef.Van_Genabith@dfki.de, dyxiong@tju.edu.cn,
liuqhano@foxmail.com, Jingyi.Zhang@dfki.de

Abstract

The Transformer translation model (Vaswani et al., 2017) based on a multi-head attention mechanism can be computed effectively in parallel and has significantly pushed forward the performance of Neural Machine Translation (NMT). Though intuitively the attentional network can connect distant words via shorter network paths than RNNs, empirical analysis demonstrates that it still has difficulty in fully capturing long-distance dependencies (Tang et al., 2018). Considering that modeling phrases instead of words has significantly improved the Statistical Machine Translation (SMT) approach through the use of larger translation blocks (“phrases”) and its reordering ability, modeling NMT at phrase level is an intuitive proposal to help the model capture long-distance relationships. In this paper, we first propose an attentive phrase representation generation mechanism which is able to generate phrase representations from corresponding token representations. In addition, we incorporate the generated phrase representations into the Transformer translation model to enhance its ability to capture long-distance relationships. In our experiments, we obtain significant improvements on the WMT 14 English-German and English-French tasks on top of the strong Transformer baseline, which shows the effectiveness of our approach. Our approach helps Transformer Base models perform at the level of Transformer Big models, and even significantly better for long sentences, but with substantially fewer parameters and training steps. The fact that phrase representations help even in the big setting further supports our conjecture that they make a valuable contribution to long-distance relations.

1 Introduction

NMT is a new approach to machine translation that has achieved great success in the last a few

years (Sutskever et al., 2014; Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017). Compared to plain SMT (Brown et al., 1993; Koehn et al., 2003; Chiang, 2005), a neural language model decoder (Sutskever et al., 2014) is better at long-distance re-ordering, and attention mechanisms (Bahdanau et al., 2015; Vaswani et al., 2017) have been proven effective in modeling long-distance dependencies, while these two issues were both challenging for SMT.

The Transformer (Vaswani et al., 2017), which has outperformed previous RNN/CNN based translation models (Bahdanau et al., 2015; Gehring et al., 2017), is based on multi-layer multi-head attention networks and can be trained in parallel very efficiently. Though attentional networks can connect distant words via shorter network paths than RNNs, empirical results show that its ability in capturing long-range dependencies does not significantly outperform RNNs, and it is still a problem for the Transformer to fully model long-distance dependencies (Tang et al., 2018).

Using phrases instead of words enables conventional SMT to condition on a wider range of context, and results in better performance in re-ordering and modeling long-distance dependencies. It is intuitive to let the NMT model additionally condition on phrase level representations to capture long-distance dependencies better, but there are two main issues which prevent NMT from directly using phrases:

- There are more phrases than tokens, and the phrase table is much larger than the word vocabulary, which is not affordable for NMT;
- Distribution over phrases is much sparser than that over words, which may lead to data sparsity and hurt the performance of NMT.

Instead of using phrases directly in NMT, in

* Corresponding author.

this work, we address the issues above with the following contributions:

- To address the large phrase table issue, we propose an attentive feature extraction model and generate phrase representation based on token representations. Our model first summarizes the representation of a given token sequence with mean or max-over-time pooling, then computes the attention weight of each token based on the token representation and the summarized representation, and generates the phrase representation by a weighted combination of token representations;
- To help the Transformer translation model better model long-distance dependencies, we let both encoder layers and decoder layers of the Transformer attend the phrase representation sequence which is shorter than the token sequence, in addition to the original token representation. Since the phrase representations are produced and attended at each encoder layer, the encoding of each layer is also enhanced with phrase-level attention computation;
- To the best of our knowledge, our work is the first to model phrase representations and incorporating them into the Transformer.

Our approach empirically brings about significant and consistent improvements over the strong Transformer model (both base and big settings). We conducted experiments on the WMT 14 English-German and English-French news translation task, and obtained +1.29 and +1.37 BLEU improvements respectively on top of the strong Transformer Base baseline, which demonstrates the effectiveness of our approach. Our approach helps Transformer Base models perform at the level of Transformer Big models, and even significantly better for long sentences, but with substantially fewer parameters and training steps. It also shows effectiveness with the Transformer Big setting. We also conducted length analysis with our approach, and the results show how our approach improves long-distance dependency capturing, which supports our conjecture that phrase representation sequences can help the model capture long-distance relations better.

2 Background and Related Work

In this section, we first review previous work which utilizes phrases in recurrent sequence-to-

sequence models, then give a brief introduction to the stronger Transformer translation model that our work is based on.

2.1 Utilizing Phrases in RNN-based NMT

Most previous work focuses on utilizing phrases from SMT in NMT to address its coverage (Tu et al., 2016) problem.

Dahlmann et al. (2017) suggested that SMT usually performs better in translating rare words and profits from using phrasal translations, even though NMT achieves better overall translation quality. They introduced a hybrid search algorithm for attention-based NMT which extended the beam search of NMT with phrase translations from SMT. Wang et al. (2017a) proposed that while NMT generally produces fluent but often inadequate translations, SMT yields adequate translations though less fluent. They incorporate SMT into NMT through utilizing recommendations from SMT in each decoding step of NMT to address the coverage issue and the unknown word issue of NMT. Wang et al. (2017b) suggested that phrases play a vital role in machine translation, and proposed to translate phrases in NMT by integrating target phrases from an SMT system with a phrase memory given that it is hard to integrate phrases into NMT which reads and generates sentences in a token-by-token way. The phrase memory is provided by the SMT model which dynamically picks relevant phrases with the partial translation from the NMT decoder in each decoding step.

2.2 The Transformer Translation Model

Our research is based on the Transformer translation model (Vaswani et al., 2017) shown in Figure 1, which significantly outperforms the previous recurrent sequence-to-sequence approach and can be efficiently computed in parallel.

The Transformer includes an encoder and a decoder. Both encoder and decoder are a stack of 6 layers. Besides the embedding matrix and positional embedding matrix in both encoder and decoder, the decoder also has a softmax classifier layer to produce translated tokens. The weights of the softmax classifier are normally tied to the target embedding matrix.

Both encoder layers and decoder layers make use of the multi-head attention mechanism. The multi-head attention mechanism calculates attention results of given queries on corresponding keys

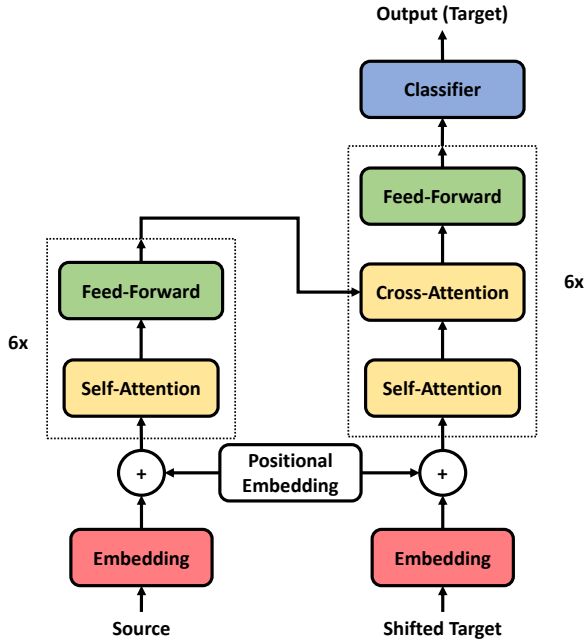


Figure 1: The Transformer Translation Model. Residual connection and Layer normalization are omitted for simplicity.

and values. It first projects queries, keys and values with 3 independent linear transformations, then splits the transformed key, query and value embeddings into several chunks of d_k dimension vectors, each chunk is called a head,¹ and scaled dot-product attention is independently applied in each head:

$$Attn(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (1)$$

where Q , K and V stand for the query vectors, key vectors and value vectors. Finally, the network concatenates the outputs of all heads and transforms it into the target space with another linear layer. The self-attention network uses the query sequence also as the key sequence and the value sequence in computation, while the cross-attention feeds another vector sequence to attend as queries and values.

Comparing the computation of the attentional network with RNNs, it is obvious that the attention computation connects distant words with a shorter network path, and intuitively it should perform better in capturing long-distance dependencies. However, empirical results show that its ability in modeling long-range dependencies does not significantly outperform RNNs.

¹ d_k is 64 for both the Transformer Base and the Transformer Big, and the numbers of heads for them are 8 and 16 respectively.

2.3 Comparison with Previous Works

Compared to previous works using RNN-based NMT (He et al., 2016; Wang et al., 2017a,b; Dahlmann et al., 2017), our proposed approach is based on the Transformer model, with the following further important differences:

- Our approach aims to improve the long-distance dependency modeling ability of NMT instead of coverage (Tu et al., 2016);
- Our approach does not require to train an SMT system or to extract aligned phrase translation from the training corpus, which makes it efficient and avoids suffering from potential error propagation from the SMT system. The phrase representation learning model is a neural model, and is deeply integrated in the translation model, and the whole neural model is end-to-end trainable;
- We iteratively and dynamically generate phrase representations with token vectors. Previous work does not use SMT phrases in this way.

In more recent work, Wang et al. (2019) augment self attention with structural position representations to model the latent structure of the input sentence; Hao et al. (2019) propose multi-granularity self-attention which performs phrase-level attention with several attention heads.

3 Transformer with Phrase Representation

For the segmentation of phrases, given that N-gram phrases are effective for tensor libraries, we first try to cut a token sequence into a phrase sequence with a fixed phrase length which varies with the sequence length.² We pad the last phrase in case it does not have sufficient tokens, thus we can transform the whole sequence into a tensor.

The N-gram phrase segmentation is efficient and simple, and we suggest the drawbacks of such “casual” segmentation boundaries can be alleviated with self-attention computation across the whole sequence and the attention mechanism applied in the generation of phrase representation which values tokens differently to a large extent, given that

²We implement this as: $ntok = \max(\min(8, seq/6), 3)$, where $ntok$ and seq stand for the number of tokens in each phrase and the length of a sentence respectively.

Algorithm 1 Extracting Phrases from a Parse Tree.
Input: A parse tree T , maximum tokens allowed in a phrase n ; Output: Extracted phrase sequence S .

```

1: while  $T$  is not empty do
2:   Initialize a phrase sequence  $p = []$ , maximum tokens allowed in this phrase  $mt = n$ ;
3:   Find the largest sub-tree  $ST$  with  $nst$  tokens ( $nst < n$ ) and depth  $dst$  from the right side of  $T$ ;
4:   Add the token sequence in  $ST$  into  $p$ ;
5:   Remove  $ST$  from  $T$ ;
6:   while  $mt > 0$  do
7:     Find the adjacent sub-tree  $STA$  of depth  $dst$  with  $nsta$  tokens from the right side of  $T$ ;
8:     if  $STA$  exists and  $nsta \leq mt$  then
9:       Insert the token sequence of  $STA$  to the beginning of  $p$ ;
10:      Remove  $STA$  from  $T$ ;
11:       $mt = mt - nsta$ ;
12:     else
13:       Break;
14:     end if
15:   end while
16:   Append  $p$  to  $S$ ;
17: end while
18: Reverse  $S$ ;
19: return  $S$ 

```

neural models have been proven good at learning competitively effective representations with gate or attention mechanism even without modeling linguistic structures (Cho et al., 2014; Hochreiter and Schmidhuber, 1997; Vaswani et al., 2017; Devlin et al., 2019).

In our experiments we also explore phrases extracted from the Stanford Parser (Socher et al., 2013) as an alternative to our simple segmentation strategy. The maximum number of tokens allowed is consistent with the simple segmentation approach, and we try to use the tokens from the largest sub-tree that complies with the maximum token limitation or from several adjacent sub-trees of the same depth as a phrase for efficiency. Our algorithm to extract phrases from parse trees is shown in Algorithm 1.

To efficiently parallelize parser-based phrases of various length in a batch of data, we pad short phrases to the same length of the longest phrases in the batch of sentences, thus a batch of sequences of

phrases can be saved into a tensor. But significantly more “<pad>” tokens will be introduced, and the model is slightly slower than the simple approach.

3.1 Attentive Phrase Representation Generation

Merging several token vectors into one is very likely to incur information loss, and introducing an importance evaluation mechanism is better than treating tokens equally. To highlight the most important features in a segmented phrase chunk, we introduce an attentive phrase representation generation model to value tokens differently according to their importance in the phrase. The model first roughly extracts features from all tokens into a vector, then assigns a score to each token by comparing each token vector with the extracted feature vector, and produces the weighted accumulation of all token vectors according to their scores.

Phrase representations are generated in every encoder layer, for the k_{th} encoder layer, we generate phrase representation $R_{e_{phrase}}^k$ from its input representation. Assume the phrase contains m tokens $\{t_1, \dots, t_m\}$, and $\{R_{e_{t_1}}^k, R_{e_{t_2}}^k, \dots, R_{e_{t_m}}^k\}$ are the corresponding input vectors to the encoder layer, we first generate a summary representation by:

$$R_{e_{all}}^k = F_{glance}(R_{e_{t_1}}^k, \dots, R_{e_{t_m}}^k) \quad (2)$$

where F_{glance} is a function to extract features of the vector sequence into a fixed-dimension vector; We explore both element-wise mean operation and max-over-time pooling operation in our work.

After the summarized representation is produced, we calculate a score for each token in the phrase, the score of the i_{th} token s_i^k is calculated as:

$$s_i^k = W_2^k \sigma(W_1^k [R_{e_{t_i}}^k | R_{e_{all}}^k] + b_1^k) + b_2^k \quad (3)$$

where σ is the sigmoid activation function, and “|” means concatenation of vectors. The rationale for designing this approach is further explained below.

Then we normalize the score vector to weights with the softmax function, and the probability of the i_{th} token p_i^k is:

$$p_i^k = \frac{e^{s_i^k}}{\sum_{i=1}^m e^{s_i^k}} \quad (4)$$

Finally, the representation of the phrase in the k_{th} encoder layer $R_{e_{phrase}}^k$ is generated by a weighted combination of all vectors:

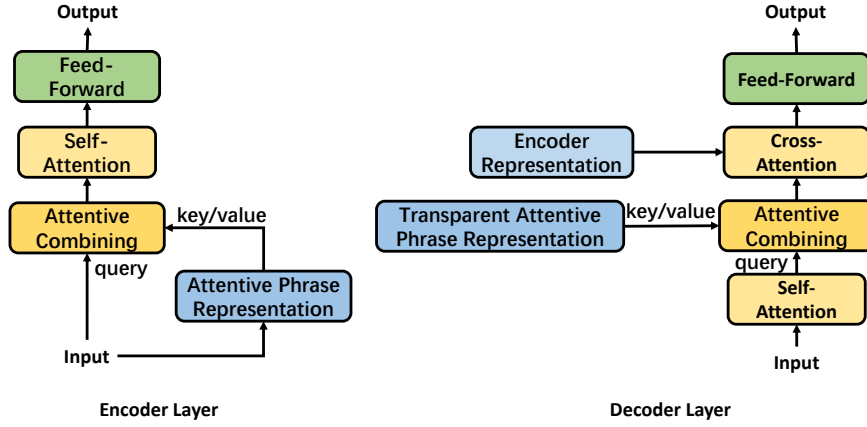


Figure 2: The Encoder/Decoder Layer of the Transformer Model with Phrase Representation. Residual connection and Layer normalization are omitted for simplicity.

$$R_{e_{phrase}}^k = \sum_{i=1}^m p_i^k R_{e_{t_i}}^k \quad (5)$$

The representation of the phrase sequence can be computed efficiently in parallel. Each encoder layer will produce a vector sequence as the phrase representation. We do not use the multi-head attention in the computation of the phrase-representation attention because of two reasons:

- The multi-head attention calculates weights through dot-product, we suggest that a 2-layer neural network might be more powerful at semantic level feature extraction, and it is less likely to be affected by positional embeddings which are likely to vote up adjacent vectors;
- Though we employ a 2-layer neural network, it only has one linear transformation and a vector to calculate attention weights, which contains fewer parameters than the multi-head attention model that has 4 linear transformations.

Recent studies show that different encoder layers capture linguistic properties of different levels (Peters et al., 2018), and aggregating layers is of profound value to better fuse semantic information (Shen et al., 2018; Dou et al., 2018; Wang et al., 2018; Dou et al., 2019). We assume that different decoder layers may value different levels of information i.e. the representation of different encoder layers differently, thus we weighted combined phrase representations from every encoder layer for each decoder layer with the Transparent Attention (TA) mechanism (Bapna et al., 2018).

For the decoder layer j , the phrase representation $R_{d_{phrase}}^j$ fed into that layer is calculated by:

$$R_{d_{phrase}}^j = \sum_{i=0}^d w_i^j R_{e_{phrase}}^i \quad (6)$$

where w_i^j are softmax normalized parameters trained jointly with the full model to learn the importance of encoder layers for the j th decoder layer. d is the number of encoder layers, and 0 corresponds to the embedding layer.

3.2 Incorporating Phrase Representation into NMT

After the phrase representation sequence for each encoder layer and decoder layer is calculated with the approach described above, we propose an attentive combination network to incorporate the phrase representation for each layer into the Transformer translation model to aid it modeling long-distance dependencies. The attentive combination network is inserted in each encoder layer and each decoder layer to bring in information from the phrase representation. The structures of the encoder layer and the decoder layer of the Transformer model with phrase representation are shown in Figure 2.

For an encoder layer, the new computation order is: cross-attention to phrases \rightarrow self-attention over tokens \rightarrow feed-forward neural network to process collected features, while for a decoder layer it is: self-attention over decoded tokens \rightarrow cross-attention to source phrases \rightarrow cross-attention to source tokens \rightarrow feed-forward neural network to process collected features. Compared to the computation order of the standard Transformer, the new computation order performs additional attending at

phrase level before attending source token representations at token level. We conjecture that attending at phrase level should be easier than at token level, and attention results at phrase level may aid the attention computation at the token-level.

For a given input sequence x and a phrase vector sequence R_{phrase} , the attentive combination network first attends the phrase representation sequence and computes the attention output out_{phrase} as follows:

$$out_{phrase} = \text{Attn}_{\text{MH}}(x, R_{phrase}) \quad (7)$$

where Attn_{MH} is a multi-head cross-attention network with x as keys and R_{phrase} as corresponding queries and values.

The attention result is then combined again with the original input sequence x with a 2-layer neural network which aims to make up for potential information loss in the phrase representation with the original token representation:

$$out = W_4\sigma(W_3[x|out_{phrase}] + b_3) + b_4 \quad (8)$$

We also employ a residual connection around the attentive combination layer, followed by layer normalization to stabilize the training.

Since the phrase representation is produced inside the Transformer model and utilized as the input of layers, and all related computations are differentiable, the attentive phrase representation model is simply trained as part of the whole model through backpropagation effectively.

4 Experiments

To compare with Vaswani et al. (2017), we conducted our experiments on the WMT 14 English to German and English to French news translation tasks.

4.1 Settings

We implemented our approaches based on the Neutron implementation (Xu and Liu, 2019) of the Transformer translation model. We applied joint Byte-Pair Encoding (BPE) (Sennrich et al., 2016) with $32k$ merge operations on both data sets to address the unknown word problem. We only kept sentences with a maximum of 256 subword tokens for training. Training sets were randomly shuffled in every training epoch. The concatenation of newstest 2012 and newstest 2013 was used for

Models	En-De	En-Fr
Transformer Base	27.38	39.34
+PR	28.67 [†]	40.71 [†]
Transformer Big	28.49	41.36
+PR	29.60 [†]	42.45 [†]

Table 1: Results on WMT 14 En-De and En-Fr.

validation and newstest 2014 as test sets for both tasks.

The number of warm-up steps was set to $8k$, and each training batch contained at least $25k$ target tokens. Our experiments run on 2 GTX 1080 Ti GPUs, and a large batch size was achieved through gradient accumulation. We used a dropout of 0.1 for all experiments except for the Transformer Big on the En-De task which was 0.3. The training steps for Transformer Base and Transformer Big were $100k$ and $300k$ respectively following Vaswani et al. (2017). The other settings were the same as (Vaswani et al., 2017) except that we did not bind the embedding between the encoder and the decoder for efficiency.

We used a beam size of 4 for decoding, and evaluated tokenized case-sensitive BLEU³ with the averaged model of the last 5 checkpoints for Transformer Base and 20 checkpoints for Transformer Big saved with an interval of 1,500 training steps (Vaswani et al., 2017). We also conducted significance tests (Koehn, 2004).

4.2 Main Results

We applied our approach to both the Transformer Base setting and the Transformer Big setting, and conducted experiments on both tasks to validate the effectiveness of our approach. Since parsing a large training set (specifically, the En-Fr dataset) is slow, we did not use phrases from parse results in this experiment (reported in Table 1). Results are shown in Table 1. † indicates $p < 0.01$ compared to the baseline for the significance test.

Table 1 shows that modeling phrase representation can bring consistent and significant improvements on both tasks, and benefit both the Transformer Base model and the stronger Transformer Big model. “+PR” is the Transformer with Phrase Representation, corresponding to the

³<https://github.com/moses-smt/mosesdecoder/blob/master/scripts/generic/multi-bleu.perl>

Models	BLEU	Δ	Para. (M)	Time	
				Train	Decode
Transformer Base	27.38		88.1	1.00x	1.00x
+Mean	27.99	0.61		1.64x	1.45x
+Max	28.13	0.75	129.0	1.60x	1.40x
+Max+Attn	28.52	1.14		1.74x	1.52x
+Max+Attn+TA	28.67	1.29	173.0	1.75x	1.53x
+Max+Attn+TA+Parsing Phrase	28.76	1.38		1.83x	1.60x
Transformer Big	28.49	1.11	264.1	7.73x	2.68x

Table 2: Ablation Study. Δ indicates the BLEU improvements compared to the Transformer Base. Time represents the time consumption compared to the Transformer Base (in training and decoding). The Transformer Big consumes 3 times training steps of the Transformer Base.

“+Max+Attn+TA” setting in Table 2.

The En-Fr task used a larger dataset ($\sim 36M$ sentence pairs) and achieved a higher baseline BLEU than the En-De task, we suggest significant improvements obtained by our approach on the En-Fr task with the Transformer Big supports the effectiveness of our approach in challenging settings.

4.3 Ablation Study

We also conducted a Transformer Base based ablation study on the WMT 14 En-De task to assess the influence of phrase representation, attention mechanism in phrase representation generation, transparent attention and phrases from parser output on performance. Results are shown in Table 2.

“+Mean” and “+Max” are only using element-wise mean operation and max-over-time pooling to generate an initial rough phrase representation of a given token sequence. “+Attn” indicates generating phrase representations with our attentive approach, on top of the max-over-time pooling as F_{glance} in Equation 2. “+TA” indicates use of the Transparent Attention mechanism to fuse information generated from every encoder layer for different decoder layers,⁴ otherwise only outputs of the last encoder layer are fed into all decoder layers. “+Parse” means using phrases extracted from parse results with Algorithm 1.

Table 2 shows that introducing phrase representation can significantly improve the strong Transformer Base baseline, even only with a simple element-wise mean operation over token repre-

⁴This only introduces an additional $7 * 6$ parameter matrix, which does not show significant influence in view of the amount of parameters.

sentations brings about a +0.61 BLEU improvement ($p < 0.01$). Summarizing representations with max-over-time pooling performs slightly better than with the element-wise mean operation. Our attentive phrase representation generation approach can bring further improvements over the max-over-time pooling approach. Though utilizing phrases from the parser can make use of linguistic knowledge and obtains most improvements, our simple and effective segmenting approach performs competitively, and we interpret these comparisons to show the positive effects of collapsing token sequences into shorter phrase sequences on the modeling of long-distance dependencies.

Though a significant amount of parameters are introduced for incorporating phrase representation into the Transformer model, our approach (“+Max+Attn+TA”) improved the performance of the Transformer Base model by +1.29 BLEU on the WMT 14 En-De news task, and the proposed Transformer model with phrase representation still performs competitively compared to the Transformer Big model with only about half the number of parameters and 1/3 of the training steps. Thus, we suggest our improvements are not only because of introducing parameters, but also due to the modeling and utilization of phrase representation.

4.4 Length Analysis

To analyze the effects of our phrase representation approach on performance with increasing input length, we conducted a length analysis on the news test set of the WMT 14 En-De task. Following Bahdanau et al. (2015) and Tu et al. (2016), we grouped sentences of similar lengths together and computed BLEU scores of Transformers and Transformers




Figure 3: BLEU scores with respect to various input sentence lengths.

with phrase representations for each group. Results are shown in Figure 3.

Figure 3 shows that our approach incorporating phrase representation into the Transformer significantly improves its performance in all length groups, and longer sentences show significantly more improvements than shorter sentences. In the Transformer Base setting, our approach improved the group with sentences of more than 45 tokens by +1.72 BLEU, almost twice of the improvements for sentences with less than 15 tokens which was +0.93 BLEU.

The effects of incorporating phrase representations into the Transformer is more significant especially when compared to the Transformer Big which has about twice the number of parameters than our approach and consumes 3 times the training steps. According to Tang et al. (2018), the number of attention heads in Transformers impacts their ability to capture long-distance dependencies, and specifically, many-headed multi-head attention is essential for modeling long-distance phenomena with only self-attention. The Transformer Big model with twice the number of heads in the multi-head attention network compared to those in the Transformer Base model, should be better at capturing long-distance dependencies. However, comparing with the Transformer Base, the improvement of the Transformer Big on long sentences (+1.20 BLEU for sentences with more than 45 tokens) was similar to that on short sentences (+1.14 BLEU for sentences with no more than 15 tokens), while our approach to model phrases in the Transformer model even brings significantly ($p < 0.01$) more improvements (+1.72 BLEU) on the performance of longer sentences with the Transformer Base setting (8 heads) than the Transformer Big with 16

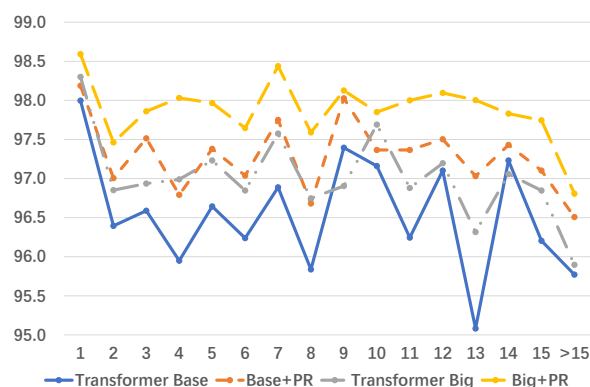


Figure 4: Subject-Verb Agreement Analysis. X-axis and y-axis represent subject-verb distance in words and the accuracy respectively.

heads (+1.20 BLEU).

The length analysis result is consistent with our conjecture to some extent given that there are likely to be more long-distance dependencies in longer source sentences. We suggest that phrase sequences which are shorter than corresponding token sequences can help the model capture long-distance dependencies better, and modeling phrase representations for the Transformer can enhance its performance on long sequences.

4.5 Subject-Verb Agreement Analysis

Intuitively, in translating longer sentences we should encounter more long-distance dependencies than in short sentences. To verify whether our method can improve the capability of the NMT model to capture long-distance dependencies, we also conducted a linguistically-informed verb-subject agreement analysis on the *Lingeval97* dataset (Sennrich, 2017) following Tang et al. (2018).

In German, subjects and verbs must agree with one another in grammatical number and person. In *Lingeval97*, each contrastive translation pair consists of a correct reference translation, and a contrastive example that has been minimally modified to introduce one translation error. The accuracy of a model is the number of times it assigns a higher score to the reference translation than to the contrastive one, relative to the total number of predictions. Results are shown in Figure 4.

Figure 4 shows that our approach can improve the accuracy of long-distance subject-verb dependencies, especially for cases where there are more than 10 tokens between the verb and the corresponding subject when comparing the “Base+PR”

with the “Transformer Big”.

5 Conclusion

Considering that the strong Transformer translation model still has difficulty in fully capturing long-distance dependencies (Tang et al., 2018), and that using a shorter phrase sequence (in addition to the original token sequence) is an intuitive approach to help the model capture long-distance features, in this paper, we first propose an attention mechanism to generate phrase representations by merging corresponding token representations. In addition, we incorporate the generated phrase representations into the Transformer translation model to help it capture long-distance relationships. We obtained statistically significant improvements on the WMT 14 English-German and English-French tasks over the strong Transformer baseline, which demonstrates the effectiveness of our approach. Our further analysis shows that the Transformer with phrase representation empirically improves its performance especially in long-distance dependency learning.

Acknowledgments

We thank anonymous reviewers for their insightful comments and helpful advice. Hongfei Xu acknowledges the support of China Scholarship Council ([2018]3101, 201807040056). Deyi Xiong is supported by the National Natural Science Foundation of China (Grant No. 61861130364), the Natural Science Foundation of Tianjin (Grant No. 19JCZDJC31400) and the Royal Society (London) (NAF\R1\180122). Hongfei Xu, Josef van Genabith and Jingyi Zhang are supported by the German Federal Ministry of Education and Research (BMBF) under the funding code 01IW17001 (Deeplee).

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*.

Ankur Bapna, Mia Chen, Orhan Firat, Yuan Cao, and Yonghui Wu. 2018. [Training deeper neural machine translation models with transparent attention](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages

3028–3033. Association for Computational Linguistics.

- Peter F. Brown, Stephen A. Della Pietra, Vincent J. Della Pietra, and Robert L. Mercer. 1993. [The mathematics of statistical machine translation: Parameter estimation](#). *Computational Linguistics*, 19(2):263–311.
- David Chiang. 2005. [A hierarchical phrase-based model for statistical machine translation](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 263–270, Ann Arbor, Michigan. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder-decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Leonard Dahlmann, Evgeny Matusov, Pavel Petrushkov, and Shahram Khadivi. 2017. [Neural machine translation leveraging phrase-based models in a hybrid search](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1411–1420, Copenhagen, Denmark. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zi-Yi Dou, Zhaopeng Tu, Xing Wang, Shuming Shi, and Tong Zhang. 2018. [Exploiting deep representations for neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4253–4262, Brussels, Belgium. Association for Computational Linguistics.
- Zi-Yi Dou, Zhaopeng Tu, Xing Wang, Longyue Wang, Shuming Shi, and Tong Zhang. 2019. [Dynamic layer aggregation for neural machine translation with routing-by-agreement](#). In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, pages 86–93.
- 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 of *Proceedings of Machine*

- Learning Research*, pages 1243–1252, International Convention Centre, Sydney, Australia. PMLR.
- Jie Hao, Xing Wang, Shuming Shi, Jinfeng Zhang, and Zhaopeng Tu. 2019. [Multi-granularity self-attention for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 887–897, Hong Kong, China. Association for Computational Linguistics.
- Wei He, Zhongjun He, Hua Wu, and Haifeng Wang. 2016. [Improved neural machine translation with SMT features](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 151–157.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9(8):1735–1780.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*.
- Philipp Koehn, Franz J. Och, and Daniel Marcu. 2003. [Statistical phrase-based translation](#). In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 127–133.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Rico Sennrich. 2017. [How grammatical is character-level neural machine translation? assessing MT quality with contrastive translation pairs](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 376–382, Valencia, Spain. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [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. Association for Computational Linguistics.
- Yanyao Shen, Xu Tan, Di He, Tao Qin, and Tie-Yan Liu. 2018. [Dense information flow for neural machine translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1294–1303, New Orleans, Louisiana. Association for Computational Linguistics.
- Richard Socher, John Bauer, Christopher D. Manning, and Andrew Y. Ng. 2013. [Parsing with compositional vector grammars](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 455–465, Sofia, Bulgaria. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Gongbo Tang, Mathias Müller, Annette Rios, and Rico Sennrich. 2018. [Why self-attention? A targeted evaluation of neural machine translation architectures](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4263–4272, Brussels, Belgium. Association for Computational Linguistics.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. [Modeling coverage for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 76–85, 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 *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Qiang Wang, Fuxue Li, Tong Xiao, Yanyang Li, Yinqiao Li, and Jingbo Zhu. 2018. [Multi-layer representation fusion for neural machine translation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3015–3026, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Xing Wang, Zhengdong Lu, Zhaopeng Tu, Hang Li, Deyi Xiong, and Min Zhang. 2017a. [Neural machine translation advised by statistical machine translation](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 3330–3336.
- Xing Wang, Zhaopeng Tu, Longyue Wang, and Shuming Shi. 2019. [Self-attention with structural position representations](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1403–1409, Hong Kong, China. Association for Computational Linguistics.

Xing Wang, Zhaopeng Tu, Deyi Xiong, and Min Zhang. 2017b. [Translating phrases in neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1421–1431, Copenhagen, Denmark. Association for Computational Linguistics.

Hongfei Xu and Qiuhui Liu. 2019. [Neutron: An Implementation of the Transformer Translation Model and its Variants](#). *arXiv preprint arXiv:1903.07402*.