

# Improving Multi-Head Attention with Capsule Networks

Shuhao GU<sup>1,2</sup> and Yang FENG<sup>1,2\*</sup>

<sup>1</sup> Institute of Computing Technology, Chinese Academy of Sciences (ICT/CAS)

<sup>2</sup> University of Chinese Academy of Sciences  
{gushuhao17g, fengyang}@ict.ac.cn

**Abstract.** Multi-head attention advances neural machine translation by working out multiple versions of attention in different subspaces, but the neglect of semantic overlapping between subspaces increases the difficulty of translation and consequently hinders the further improvement of translation performance. In this paper, we employ capsule networks to comb the information from the multiple heads of the attention so that similar information can be clustered and unique information can be reserved. To this end, we adopt two routing mechanisms of Dynamic Routing and EM Routing, to fulfill the clustering and separating. We conducted experiments on Chinese-to-English and English-to-German translation tasks and got consistent improvements over the strong Transformer baseline.

**Keywords:** Neural machine translation · Transformer · Capsule network · Multi-head attention

## 1 Introduction

Neural machine translation (NMT) [2,4,7,11,24,26] has made great progress and drawn much attention recently. Although NMT models may have different structures for encoding and decoding, most of them employ an attention function to collect source information to translate at each time step. Multi-head attention proposed by [26] has shown its superiority in different translation tasks and been accepted as an advanced technique to improve the existing attention functions [12,15,17].

In contrast to conventional attention, the multi-head attention mechanism extends attention from one unique space to different representation subspaces. It works by first projecting the queries, keys, and values to different subspaces, then performing dot products to work out the corresponding attention in each subspace, and finally concatenating all these attentions to get the multi-head attention. This projecting process explores possible representations in different subspaces independently and hence can mitigate the all-in risk caused by one unique space. Different attention heads may carry different features of the target sequences in different subspaces. However, the subspaces are not always orthogonal to each other and the overlapping will lead to redundant semantic. Concatenating the attentions of different heads directly neglects the redundancy

---

\*Corresponding Author

and may bring about wrong subsequent operations by treating them as different semantics, resulting in degraded performance. This also results in that only some important individual heads play consistent and often linguistically-interpretable roles and others can be pruned directly without harming the overall performance too much[33]. Therefore, it is desirable to design a separate component to arrange and fuse the semantic and spatial information from different heads to help boost the translation quality of the model.

To address this problem, we propose a method to utilize capsule networks [10,19] to model the relationship of the attention from different heads explicitly. Capsule networks provide an effective way to cluster the information of the input capsules via an iterative dynamic routing process and store the representative information of each cluster in an output capsule. Our method inserts a capsule network layer right after the multi-head attention so that the information from all the heads can be combed. We adopt two routing mechanisms, Dynamic Routing, and EM routing, for the capsule network to decide the flow of information. Then the output is feed into a fully connected feed-forward neural network. We also employed a residual connection around the input and final output layer. In our experiments, we gradually replaced the multi-head attention of the original model with ours at different positions. The experiments on the Chinese-to-English and English-to-German translation tasks show that EM routing works better than Dynamic Routing and our method with either routing mechanism can outperform the strong transformer baseline.

## 2 Background

The attention mechanism was first introduced for machine translation task by [2]. The core part of the attention mechanism is to map a sequence of  $K$ , the keys, to the distribution of weights  $a$  by computing its relevance with  $q$ , the queries, which can be described as:

$$\mathbf{a} = f(\mathbf{q}, \mathbf{K})$$

where the keys and the queries are all vectors. In most cases,  $K$  is the word embeddings or the hidden states of the model which encode the data features whereupon attention is computed.  $q$  is a reference when computing the attention distribution. The attention mechanism will emphasize the input elements considered to be inherently relevant to the query. The attention mechanism in Transformer is the so-called scaled dot product attention which uses the dot-product of the query and keys to present the relevance of the attention distribution:

$$\mathbf{a} = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)$$

where the  $d_k$  is the dimensions of the keys. Then the weighted values are summed together to get the final results:

$$\mathbf{u} = \sum \mathbf{a} \odot \mathbf{V}$$

Instead of performing a single attention function with a single version of a query, key, and value, multi-head attention mechanism gets  $h$  different versions of queries,

keys, and values with different projection functions:

$$Q^i, K^i, V^i = QW_i^Q, KW_i^K, VW_i^V, i \in [1, h]$$

where  $Q^i, K^i, V^i$  are the query, key and value representations of the  $i$ -th head respectively.  $W_i^Q, W_i^K, W_i^V$  are the transformation matrices.  $h$  is the number of attention heads.  $h$  attention functions are applied in parallel to produce the output states  $\mathbf{u}_i$ . Finally, the outputs are concatenated to produce the final attention:

$$\mathbf{u} = \text{Concat}(\mathbf{u}_1, \dots, \mathbf{u}_h)$$

### 3 Related Work

**Attention Mechanism** Attention was first introduced in for machine translation tasks by [2] and it already has become an essential part in different architectures [7,13,26] though that they may have different forms. Many works are trying to modify the attention part for different purposes [3,14,16,22,23,25,29]. Our work is mainly related to the work which tries to improve the multi-head attention mechanism in the Transformer model.

[34] analyze different aspects of the attention part of the transformer model. It shows that the multi-head attention mechanism can only bring limited improvement compared to the 1-head model. [33] evaluate the contribution made by the individual attention heads in the encoder to the overall performance of the model and then analyze the roles played by them. They find that the most important and confident heads play consistent roles while others can be pruned without harming the performance too much. We believe that because of the overlapping of the subspaces between different attention heads, the function of some attention heads can be replaced by other heads. [12] share the same motivation with ours. They add three kinds of L2-norm regularization methods, which are the subspace, the attended positions, and the output representation, to the loss function to encourage each attention head to be different from other heads. This is a straightforward approach, but it may ignore some semantic information. [35] is similar to our work, they use the routing-by-agreement algorithm, which is from the capsule network, to improve the information aggregation for multi-head attention. We did our work independently, without drawing on their work. Besides, the main structure of our model is different from theirs. [1] learn different weights for each attention head, then they sum the weighted attention heads up to get the attention results rather than just concatenating them together. [21] states that the original attention mechanisms do not explicitly model relative or absolute position information in its structure, thus they add a relative position representation in the attention function.

**Capsule Networks in NLP** Capsule network was first introduced by [19] for the computer vision task which aims to improve the representational limitations of the CNN structure. Then [10] replace the dynamic routing method with the Expectation-Maximization method to better estimate the agreement between capsules.

There are also some researchers trying to apply the capsule network to NLP tasks. [30] explored capsule networks with dynamic routing for text classification and achieved

competitive results over the compared baseline methods on 4 out of 6 data sets. [31] explored the capsule networks used for relation extraction in a multi-instance multi-label learning framework. [8] designed two dynamic routing policies to aggregate the outputs of the RNN/CNN encoding layer into a final encoding layer. [28] uses an aggregation mechanism to map the source sentence into a matrix with pre-determined size and then decode the target sequence from the source representation which can ensure the whole model runs in time that is linear in the length of the sequences.

## 4 The Proposed Method

Our work is based on the multi-head attention mechanism of the Transformer model:

$$\mathbf{u}_i = \text{softmax}\left(\frac{\mathbf{Q}_i \mathbf{K}_i^T}{\sqrt{d_k}}\right) \mathbf{V}_i; \quad i \in [1, h] \quad (1)$$

where  $\mathbf{Q}_i$ ,  $\mathbf{K}_i$  and  $\mathbf{V}_i$  are computed by different versions of projection functions:

$$\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i = \mathbf{QW}_i^Q, \mathbf{KW}_i^K, \mathbf{VW}_i^V, i \in [1, h] \quad (2)$$

We aim to find a proper representation  $\mathbf{v}$  based on these attention heads  $\mathbf{u}$ . These attention heads can be regarded as the different observations from different viewpoints on the same entity in the sequence.

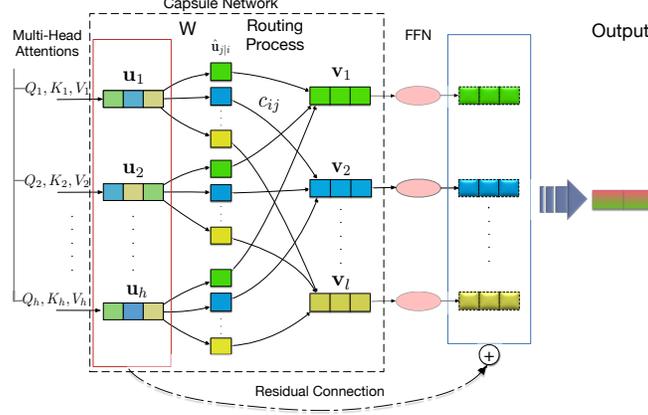
Capsule network was first proposed by [19] for the computer vision tasks. A capsule is a group of neurons whose outputs represent different properties of the same entity. The activities of the neurons within an active capsule represent the various properties of a particular entity. A part produces a vote by multiplying its pose matrix which is a learned transformation matrix that represents the viewpoint invariant relationship between the part and the whole. In the multi-head attention mechanism, different attention heads can be regarded as the different observations from different viewpoints on the same entity in the sequence. The input capsule layer represents different linguistic properties of the same input. The iterative routing process can better decide what and how much information flow to the output capsules. Ideally, each output capsule represents a distinct property of the input and carry all the deserved information when they are combined.

The overall architecture is given in Figure 1. First, the capsule computes a vote by multiplying the input capsules  $\mathbf{u}_i$  by a learned transformation matrix  $\mathbf{W}_{ij}$  that represents the viewpoint invariant relationship between the part and the whole:

$$\hat{\mathbf{u}}_{j|i} = \mathbf{W}_{ij} \mathbf{u}_i \quad (3)$$

Then we compute and update the output capsules  $\mathbf{v}$ , the vote  $\hat{\mathbf{u}}$ , and the assignment probabilities  $\mathbf{c}$  between them by a specific routing process iteratively to ensure the input to be sent to an appropriate output capsule:

$$\begin{aligned} \mathbf{v} &= f(\hat{\mathbf{u}}, \mathbf{c}) \\ \mathbf{c} &= \text{Update}(\hat{\mathbf{u}}, \mathbf{v}) \end{aligned} \quad (4)$$



**Fig. 1.** The architecture of our method

Last, the output capsules  $\mathbf{v}$  are concatenated together and fed into a feed-forward network (FFN) which consists of two linear transformations with a ReLU activation in between:

$$\text{FFN}(x) = \max(0, \mathbf{x}\mathbf{W}_1 + \mathbf{b}_1)\mathbf{W}_2 + \mathbf{b}_2 \quad (5)$$

We also add a residual connection between the layer  $u$  and  $v$  [9]. Thus the final output is:

$$\mathbf{O} = \mathbf{u} + \text{FFN}(\mathbf{v}) \quad (6)$$

where

$$\begin{aligned} \mathbf{u} &= \text{Concat}(\mathbf{u}_1, \dots, \mathbf{u}_h) \\ \mathbf{v} &= \text{Concat}(\mathbf{v}_1, \dots, \mathbf{v}_l) \end{aligned} \quad (7)$$

More specifically, we have tried the Dynamic Routing and EM Routing in our method.

**Dynamic Routing** In this method, we sum up all these weighted vote vectors to get the origin output capsule vectors:

$$\mathbf{s}_j = \sum_i c_{ij} \hat{\mathbf{u}}_{j|i} \quad (8)$$

where

$$c_{ij} = \frac{\exp(b_{ij})}{\sum_k \exp(b_{ik})} \quad (9)$$

the  $c_{ij}$  are determined by computing the "routing softmax" of the initial logits  $b_{ij}$  which are initialized to zero.

Next, the origin output capsule vectors  $\mathbf{s}_j$  is applied with a squashing function to bring non-linearity to the whole model:

$$\mathbf{v}_j = \frac{\|\mathbf{s}_j\|^2}{1 + \|\mathbf{s}_j\|^2} \frac{\mathbf{s}_j}{\|\mathbf{s}_j\|} \quad (10)$$

The initial coupling coefficients  $b_{ij}$  are iteratively refined by measuring the agreement between the current output  $\mathbf{v}_j$  by the dot-product of the input capsules and each output capsule:

$$b_{ij} \leftarrow b_{ij} + \hat{\mathbf{u}}_{j|i} \cdot \mathbf{v}_j \quad (11)$$

Zh→En	# Para.	MT03	MT04	MT05	MT06	MT08	AVE.
[27]	-	46.60	47.73	45.35	43.97	-	-
[32]	-	48.28	-	46.24	46.14	38.07	-
Transformer-Base	84.6M	46.70	47.68	47.04	46.16	37.19	44.95
+ Dynamic Routing	92.2M	47.60	48.04	47.30	46.56	37.80	45.46 (+0.51)
+ EM Routing	91.7M	47.62	48.07	47.97*	46.80	38.01*	45.70 (+0.75)

**Table 1.** Case-insensitive BLEU scores for Zh→En translation. “# Para” denotes the number of parameters. “\*” is used to indicate the improvement is statistically significant with  $\rho < 0.05$  [5].

**EM Routing** In this method, each capsule becomes a combination of a  $n \times n$  pose matrix,  $M_i$ , and an activation probability,  $\alpha_i$ . Each output capsule corresponds to a Gaussian distribution and the pose matrix of each active capsule in the lower-layer corresponds to a data-point. The iterative routing process is a version of the Expectation-Maximization procedure which iteratively adjusts the means, variances and activation probabilities of the output capsules and the assignment probabilities between the two layers. The whole procedure can be divided into two parts:

**M-Step** Keep the assignment probabilities between the two layers fixed and compute the mean  $\mu_j$  and variance  $\sigma_j$  of the output capsules:

$$\begin{aligned} \mathbf{v}_j &= \frac{\sum_i c_{ij} \hat{\mathbf{u}}_{j|i}}{\sum_i c_{ij}} \\ (\sigma_j)^2 &= \frac{\sum_i c_{ij} (\hat{\mathbf{u}}_{j|i} - \mu_j)^2}{\sum_i c_{ij}} \end{aligned} \quad (12)$$

Then we compute the incremental cost and the activation probability:

$$\begin{aligned} cost_j &= \left( \log(\sigma_j) + \frac{1 + \ln 2\pi}{2} \right) \sum_i c_{ij} \\ \alpha_j &= \text{logistic}(\lambda(\beta_\alpha - \beta_\mu \sum_i c_{ij} - \sum_h cost_j^h)) \end{aligned} \quad (13)$$

where  $\sum_i c_{ij}$  is the amount of data assigned to  $j$ . We learn  $\beta_\alpha$  and  $\beta_\mu$  discriminatively and set a fixed schedule for  $\lambda$  as a hyper-parameter.

**E-Step** Keeping the Gaussian distributions of the output capsules fixed, we need to calculate the incremental cost of explaining a whole data-point  $i$  by using an active capsule  $j$ :

$$p_j = \frac{1}{\sqrt{2\pi}(\sigma_j)^2} \exp - \frac{(\hat{\mathbf{u}}_{j|i} - \mathbf{v}_j)^2}{2(\sigma_j)^2} \quad (14)$$

and then adjust the assignment probabilities based on this:

$$c_{ij} = \frac{\alpha_j p_j}{\sum_k \alpha_k p_k} \quad (15)$$

Please refer to [10] for more details. The output capsules are then reshaped into vectors and also fed into the feed-forward network. We just make use of the higher capsule as the representation of the attention results and abandon its activation probabilities to take advantage of the information aggregation way of the capsule network.

## 5 Experiments

We evaluated our method on the NIST Chinese→English (Zh→En) and WMT14 English→German (En→De) translation tasks.

### 5.1 Setup

**Chinese→English** The training data consists of about 1.25M sentence pairs from LDC corpora with 27.9M Chinese words and 34.5M English words respectively<sup>3</sup>. We used NIST 02 data set as the development set and NIST 03, 04, 05, 06, 08 sets are used as the test sets. We tokenized and lowercased the English sentences using the Moses scripts<sup>4</sup>. For the Chinese data, we performed word segmentation using the Stanford Segmentor<sup>5</sup>. Besides 30K merging operations were performed to learn byte-pair encoding(BPE) [20] on both sides.

**English→German** For this task, we used the WMT14 corpora pre-processed and released by Google<sup>6</sup> which consists of about 4.5M sentences pairs with 118M English words and 111M German words. We chose the newstest2013 as our development set and newsstet2014 as our test set.

We evaluate the proposed approaches on the Transformer model and implement it on the top of an open-source toolkit - Fairseq-py [6]. We follow [26] to set the configurations and have reproduced their reported results on the En→De task with both of the Base and Big model. All the models were trained on a single server with eight NVIDIA TITAN Xp GPUs where each was allocated with a batch size of 4096 tokens. The routing iterations are set to 3 and the number of output capsules is set to equal to the number of input capsules if there is no other statement.

During decoding, we set beam size to 4, and length penalty  $\alpha=0.6$ . Other training parameters were the same as the default configuration of the Transformer model.

<sup>3</sup> The corpora include LDC2002E18, LDC2003E07, LDC2003E14, Hansards portion of LDC2004T07, LDC2004T08 and LDC2005T06.

<sup>4</sup> <http://www.statmt.org/moses/>

<sup>5</sup> <https://nlp.stanford.edu/>

<sup>6</sup> [https://drive.google.com/uc?export=download&id=0B\\_bZck-ksdkpM25jRUN2X2UxMm8](https://drive.google.com/uc?export=download&id=0B_bZck-ksdkpM25jRUN2X2UxMm8)

En→De	# Para.	BLEU
Transformer-Base	60.9M	27.34
+ Dynamic Routing	62.0M	27.67
+ EM Routing	61.6M	27.77
Transformer-Big	209.9M	28.43
+Dynamic Routing	216.2M	28.65
+EM Routing	214.2M	28.71

**Table 2.** Case-sensitive BLEU scores for En→De translation.

## 5.2 Main Results

We reported the case-insensitive and case-sensitive 4-gram NIST BLEU score [18] on the Zh→En and En→De tasks, respectively. During the experiments, we found that our proposed method achieved the best performance when we only insert the capsule network after the last multi-head attention sub-layer in the decoder and in the attention layer between source and target. We will analyze this phenomenon in detail in the next subsection.

The Zh→En results of the Transformer-Base model are shown in the Table 1. Both of our models (Row 4,5) with the proposed capsule network attention mechanism can not only outperform the vanilla Transformer(Row 3) but also achieve a competitive performance compared to the state-of-the-art systems(Row 1,2, we use the results from the related paper directly), indicating the necessity and effectiveness of the proposed method. It shows that our method can get the information well combed and preserve all the deserved information.

Among them, the '+EM Routing' method is slightly better than the '+Dynamic Routing' method by 0.24 which because of better estimating the agreement during the routing. Besides, it requires fewer parameters and runs much faster. Considering the training speed and performance, the '+EM Routing' method is used as the default multi-head aggregation method in subsequent analysis experiments.

The En→De results are shown in the Table 2. In this experiment, we have applied our proposed methods both on the Base and Big model. The results show that our model can still outperform the baseline model, indicating the universality of the proposed approach.

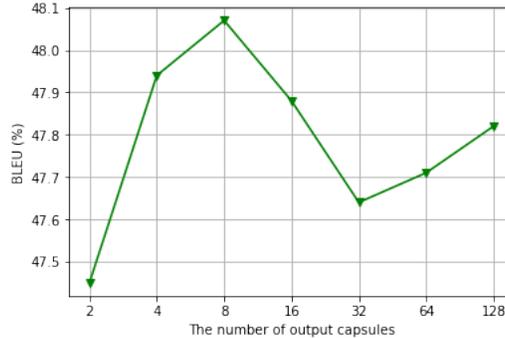
## 5.3 Impact of Different Ways to Integrate Capsule Networks

The Transformer model consists of three kinds of attention, including encoder self-attention, encoder-decoder attention and decoder self-attention at every sublayer of the encoder and decoder. We gradually insert the capsule network in different places and measured the BLEU scores on the NIST 04 test set based on the "+EM Routing" Transformer-Base model. The results are shown in the Table 3. It shows that not all of the changes are positive to the results.

First, any changes to the encoder self-attention, no matter to the top sublayer( $Enc_5, Enc_6$ ) or the bottom sublayer( $Enc_1, Enc_2$ ) of the encoder is harmful to the performance of the

# Model Variations	NIST 04
Transformer-Base	47.68
Enc <sub>1</sub>	47.25 (-0.43)
Enc <sub>1</sub> , Enc <sub>2</sub>	47.19 (-0.49)
Enc <sub>5</sub> , Enc <sub>6</sub>	47.64 (-0.04)
ED <sub>1</sub>	47.38 (-0.30)
ED <sub>6</sub>	47.98 (+0.30)
ED <sub>5</sub> , ED <sub>6</sub>	47.43 (+0.11)
Dec <sub>1</sub>	47.45 (-0.23)
Dec <sub>6</sub>	48.03 (+0.35)
Dec <sub>5</sub> , Dec <sub>6</sub>	47.83 (+0.15)
ED <sub>6</sub> , Dec <sub>6</sub>	48.07 (+0.39)
ED <sub>5</sub> , ED <sub>6</sub> , Dec <sub>5</sub> , Dec <sub>6</sub>	47.99 (+0.31)

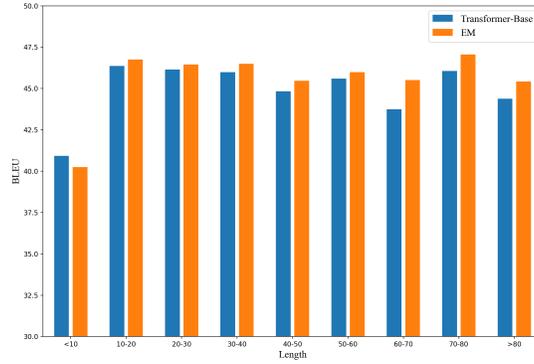
**Table 3.** Case-insensitive BLEU scores for different ways of integrating capsule networks. Enc<sub>i</sub>, Dec<sub>i</sub>, ED<sub>i</sub> mean the capsule network is inserted after i-th multi-head attention sub-layer in the encoder, in the decoder and in the attention layer between source and target, respectively. For example, Dec<sub>6</sub> means the capsule network is inserted between the multi-head attention and the FFN in the 6th layer of the decoder.



**Fig. 2.** Impact of the number of output capsules.

whole model. One possible reason for this may be that the routing part should be close to the supervisory signals to be well trained. Without its help, the capsule network only extracts internal features regardless of whether these features are helpful to the translation quality. Another reason for this may be that although we add a residual connection between the input capsule layer and the output capsule layer to ensure preserve all the information, we don't add the reconstruction procedure of the origin work [19] which may make the output leave out some information inevitable.

Then, the changes to the bottom sublayer of the encoder-decoder attention(ED<sub>1</sub>) and the bottom sublayer of the decoder self-attention(Dec<sub>1</sub>) also degrade the performance, which is also far from the supervisory signals.



**Fig. 3.** Effect of source sentence lengths.

Last, the changes to the top sublayer of the encoder-decoder attention( $ED_5, ED_6$ ) and the top sublayer of the decoder self-attention( $Dec_5, Dec_6$ ) are beneficial for the final results because that they are more close to the output layer which supports our hypothesis above.

#### 5.4 Impact of the Number of Output Capsules

The number of output capsules  $l$  is a key parameter of our model. We assumed that the capsule network can capture and extract high-level semantic information. But it is not obvious how much high-level information and what kind of information can be aggregated. Therefore we varied the number of the output capsules and also measured the BLEU scores on the NIST 04 test set based on the "+EM Routing" Transformer-Base model. The results are shown in Figure 2. It should be mentioned that the dimension of each output capsule is set to  $d/l$  to keep the final output be consistent with the hidden layer. The results show that our proposed method achieves the best performance when the number of output capsules is equal to the number of input capsules.

#### 5.5 Effect of Source Sentence Length

We also evaluated the performance of the best version of our proposed method '+EM Routing' and the baseline on the combined NIST 03-08 test set with different source sentence lengths. The results are shown in Figure 3. In the bins holding sentences no longer than 60, the BLEU scores of the two systems are close to each other. When the sentence length surpasses 60, our method shows its superiority over the Transformer base model. As the sentence length grows, the difference becomes increasingly large. That is because our method provides an effective way to cluster the information of the multi-head results so that it can get information well aggregated especially when the sentence lengths increase and handle more information.

## 6 Conclusion

In this work, we argue that the neglect of semantic overlapping between subspaces of the different attention heads increases the difficulty of translation. We adopt the Dynamic Routing and EM Routing and evaluated our methods on popular translation tasks of different language pairs and the results showed that our method can outperform the strong baselines. The extensive analysis further suggests that it can help to improve the translation performance only when we set the capsule part close to the supervisory signals.

## References

1. Karim Ahmed, Nitish Shirish Keskar, and Richard Socher. 2017. Weighted transformer network for machine translation. *arXiv preprint arXiv:1711.02132*.
2. Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
3. Kehai Chen, Rui Wang, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2018. Syntax-directed attention for neural machine translation. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
4. 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. *arXiv preprint arXiv:1406.1078*.
5. Michael Collins, Philipp Koehn, and Ivona Kučerová. 2005. Clause restructuring for statistical machine translation. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 531–540. Association for Computational Linguistics.
6. Sergey Edunov, Myle Ott, and Sam Gross. 2017. <https://github.com/pytorch/fairseq>.
7. Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*.
8. Jingjing Gong, Xipeng Qiu, Shaojing Wang, and Xuanjing Huang. 2018. Information aggregation via dynamic routing for sequence encoding. *arXiv preprint arXiv:1806.01501*.
9. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
10. Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. 2018. Matrix capsules with em routing.
11. Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent continuous translation models. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1700–1709.
12. Jian Li, Zhaopeng Tu, Baosong Yang, Michael R Lyu, and Tong Zhang. 2018. Multi-head attention with disagreement regularization. *arXiv preprint arXiv:1810.10183*.
13. Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
14. Fandong Meng, Zhengdong Lu, Hang Li, and Qun Liu. 2016. Interactive attention for neural machine translation. *arXiv preprint arXiv:1610.05011*.
15. Fandong Meng and Jinchao Zhang. 2018. Dtm: A novel deep transition architecture for neural machine translation. *arXiv preprint arXiv:1812.07807*.
16. Haitao Mi, Zhiguo Wang, and Abe Ittycheriah. 2016. Supervised attentions for neural machine translation. *arXiv preprint arXiv:1608.00112*.
17. Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. Scaling neural machine translation. *arXiv preprint arXiv:1806.00187*.

18. Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
19. Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. 2017. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, pages 3856–3866.
20. Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
21. Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. *arXiv preprint arXiv:1803.02155*.
22. Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, Shirui Pan, and Chengqi Zhang. 2018a. Disan: Directional self-attention network for rnn/cnn-free language understanding. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
23. Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, and Chengqi Zhang. 2018b. Bi-directional block self-attention for fast and memory-efficient sequence modeling. *arXiv preprint arXiv:1804.00857*.
24. Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
25. Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling coverage for neural machine translation. *arXiv preprint arXiv:1601.04811*.
26. 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*, pages 5998–6008.
27. Mingxuan Wang, Jun Xie, Zhixing Tan, Jinsong Su, Deyi Xiong, and Chao Bian. 2018a. Neural machine translation with decoding history enhanced attention. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1464–1473.
28. Mingxuan Wang, Jun Xie, Zhixing Tan, Jinsong Su, et al. 2018b. Towards linear time neural machine translation with capsule networks. *arXiv preprint arXiv:1811.00287*.
29. Baosong Yang, Zhaopeng Tu, Derek F Wong, Fandong Meng, Lidia S Chao, and Tong Zhang. 2018a. Modeling localness for self-attention networks. *arXiv preprint arXiv:1810.10182*.
30. Min Yang, Wei Zhao, Jianbo Ye, Zeyang Lei, Zhou Zhao, and Soufei Zhang. 2018b. Investigating capsule networks with dynamic routing for text classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3110–3119.
31. Ningyu Zhang, Shumin Deng, Zhanlin Sun, Xi Chen, Wei Zhang, and Huajun Chen. 2018a. Attention-based capsule networks with dynamic routing for relation extraction. *arXiv preprint arXiv:1812.11321*.
32. Zhirui Zhang, Shuangzhi Wu, Shujie Liu, Mu Li, Ming Zhou, and Enhong Chen. 2018b. Regularizing neural machine translation by target-bidirectional agreement. *arXiv preprint arXiv:1808.04064*.
33. Voita, Elena and Talbot, David and Moiseev, Fedor and Sennrich, Rico and Titov, Ivan. 2019a. Analyzing Multi-Head Self-Attention: Specialized Heads Do the Heavy Lifting, the Rest Can Be Pruned *arXiv preprint arXiv:1905.09418*
34. Domhan, Tobias 2018b. How much attention do you need? a granular analysis of neural machine translation architectures *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1799–1808.
35. Li, Jian and Yang, Baosong and Dou, Zi-Yi and Wang, Xing and Lyu, Michael R and Tu, Zhaopeng 2019a. Information Aggregation for Multi-Head Attention with Routing-by-Agreement *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 3566–3575.