

Dynamic Layer Aggregation for Neural Machine Translation with Routing-by-Agreement

Zi-Yi Dou
Carnegie Mellon University
zdou@andrew.cmu.edu

Zhaopeng Tu*
Tencent AI Lab
zptu@tencent.com

Xing Wang
Tencent AI Lab
brightxwang@tencent.com

Longyue Wang
Tencent AI Lab
vinnylywang@tencent.com

Shuming Shi
Tencent AI Lab
shumingshi@tencent.com

Tong Zhang
Tencent AI Lab
bradymzhang@tencent.com

Abstract

With the promising progress of deep neural networks, layer aggregation has been used to fuse information across layers in various fields, such as computer vision and machine translation. However, most of the previous methods combine layers in a static fashion in that their aggregation strategy is independent of specific hidden states. Inspired by recent progress on capsule networks, in this paper we propose to use *routing-by-agreement* strategies to aggregate layers dynamically. Specifically, the algorithm learns the probability of a part (individual layer representations) assigned to a whole (aggregated representations) in an iterative way and combines parts accordingly. We implement our algorithm on top of the state-of-the-art neural machine translation model TRANSFORMER and conduct experiments on the widely-used WMT14 English⇒German and WMT17 Chinese⇒English translation datasets. Experimental results across language pairs show that the proposed approach consistently outperforms the strong baseline model and a representative static aggregation model.

Introduction

Deep neural networks have advanced the state of the art in various communities, from computer vision to natural language processing. Researchers have directed their efforts into designing patterns of modules that can be assembled systematically, which makes neural networks deeper and wider. However, one key challenge of training such huge networks lies in how to transform and combine information across layers. To encourage gradient flow and feature propagation, researchers in the field of computer vision have proposed various approaches, such as residual connections (He et al. 2016), densely connected network (Huang et al. 2017) and deep layer aggregation (Yu et al. 2018).

State-of-the-art neural machine translation (NMT) models generally implement encoder and decoder as multiple layers (Wu et al. 2016; Gehring et al. 2017; Vaswani et al. 2017; Chen et al. 2018), in which only the top layer is

exploited in the subsequent processes. Fusing information across layers for deep NMT models, however, has received substantially less attention. A few recent studies reveal that simultaneously exposing all layer representations outperforms methods that utilize just the top layer for natural language processing tasks (Peters et al. 2018; Shen et al. 2018; Wang et al. 2018; Dou et al. 2018). However, their methods mainly focus on *static* aggregation in that the aggregation mechanisms are the same across different positions in the sequence. Consequently, useful context of sequences embedded in the layer representations are ignored, which could be used to further improve layer aggregation.

In this work, we propose dynamic layer aggregation approaches, which allow the model to aggregate hidden states across layers for each position dynamically. We assign a distinct aggregation strategy for each symbol in the sequence, based on the corresponding hidden states that represent both syntax and semantic information of this symbol. To this end, we propose several strategies to model the dynamic principles. First, we propose a simple *dynamic combination* mechanism, which assigns a distinct set of aggregation weights, learned by a feed-forward network, to each position. Second, inspired by the recent success of *iterative routing* on assigning parts to wholes for computer vision tasks (Sabour, Frosst, and Hinton 2017; Hinton, Sabour, and Frosst 2018), here we apply the idea of *routing-by-agreement* to layer aggregation. Benefiting from the high-dimensional coincidence filtering, *i.e.* the agreement between every two internal neurons, the routing algorithm has the ability to extract the most active features shared by multiple layer representations.

We evaluated our approaches upon the standard TRANSFORMER model (Vaswani et al. 2017) on two widely-used WMT14 English⇒German and WMT17 Chinese⇒English translation tasks. We show that although static layer aggregation strategy indeed improves translation performance, which indicates the necessity and effectiveness of fusing information across layers for deep NMT models, our proposed dynamic approaches outperform their static counterpart. Also, our models consistently improve translation performance over the vanilla TRANSFORMER model across language pairs. It is worth mentioning that TRANSFORMER-

* Zhaopeng Tu is the corresponding author. Work was done when Zi-Yi Dou was interning at Tencent AI Lab. Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

BASE with dynamic layer aggregation outperforms the vanilla TRANSFORMER-BIG model with only less than half of the parameters.

Contributions. Our key contributions are:

- Our study demonstrates the necessity and effectiveness of dynamic layer aggregation for NMT models, which benefits from exploiting useful context embedded in the layer representations.
- Our work is among the few studies (*cf.* (Gong et al. 2018; Zhao et al. 2018)) which prove that the idea of capsule networks can have promising applications on natural language processing tasks.

Background

Deep Neural Machine Translation

Deep representations have a noticeable effect on neural machine translation (Meng et al. 2016; Zhou et al. 2016; Wu et al. 2016). Generally, multiple-layer encoder and decoder are employed to perform the translation task through a series of nonlinear transformations from the representation of input sequences to final output sequences.

Specifically, the encoder is composed of a stack of L identical layers with the bottom layer being the word embedding layer. Each encoder layer is calculated as

$$\mathbf{H}_e^l = \text{LAYER}_e(\mathbf{H}_e^{l-1}) + \mathbf{H}_e^{l-1} \quad (1)$$

where a residual connection (He et al. 2016) is employed around each of the two layers. $\text{LAYER}(\cdot)$ is the layer function, which can be implemented as RNN (Cho et al. 2014), CNN (Gehring et al. 2017), or self-attention network (SAN) (Vaswani et al. 2017). In this work, we evaluate the proposed approach on the standard Transformer model, while it is generally applicable to any other type of NMT architectures.

The decoder is also composed of a stack of L layers:

$$\mathbf{H}_d^l = \text{LAYER}_d(\mathbf{H}_d^{l-1}, \mathbf{H}_e^L) + \mathbf{H}_d^{l-1} \quad (2)$$

which is calculated based on both the lower decoder layer and the top encoder layer \mathbf{H}_e^L . The top layer of the decoder \mathbf{H}_d^L is used to generate the final output sequence.

As seen, both the encoder and decoder stack layers in sequence and only utilize the information in the top layer. While studies have shown deeper layers extract more semantic and more global features (Zeiler and Fergus 2014; Peters et al. 2018), these do not prove that the last layer is the ultimate representation for any task. Although residual connections have been incorporated to combine layers, these connections have been “shallow” themselves, and only fuse by simple, one-step operations (Yu et al. 2018).

Exploiting Deep Representations

Recently, aggregating layers to better fuse semantic and spatial information has proven to be of profound value in computer vision tasks (Huang et al. 2017; Yu et al. 2018). For machine translation, Shen et al. (2018) and Dou et al.(2018)

have proven that simultaneously exposing all layer representations outperforms methods that utilize just the top layer on several generation tasks. Specifically, one of the methods proposed by Dou et al.(2018) is to linearly combine the outputs of all layers:

$$\tilde{\mathbf{H}} = \sum_{l=1}^L \mathbf{W}_l \mathbf{H}^l \quad (3)$$

where $\{\mathbf{W}_1, \dots, \mathbf{W}_L\} \in \mathbb{R}^d$ are trainable parameter matrices, where d is the dimensionality of hidden layers. The linear combination strategy is applied to both the encoder and decoder. The combined layer $\tilde{\mathbf{H}}$ that embeds all layer representations instead of only the top layer \mathbf{H}^L , is used in the subsequent processes.

As seen, the linear combination is encoded in a static set of weights $\{\mathbf{W}_1, \dots, \mathbf{W}_L\}$, which ignores the useful context of sentences that could further improve layer aggregation. In this work, we introduce the dynamic principles into layer aggregation mechanisms.

Approach

Dynamic Combination

An intuitive extension of static linear combination is to generate different weights for each layer combination rather than apply the same weights all the time. To this end, we calculate the weights of the linear combination as

$$\mathbf{W}_l = \text{FFN}_l(\mathbf{H}^1, \dots, \mathbf{H}^L) \in \mathbb{R}^{J \times d} \quad (4)$$

where J is the length of the hidden layer \mathbf{H}^l , and $\text{FFN}_l(\cdot)$ is a distinct feed-forward network associated with the l -th layer \mathbf{H}^l . Specifically, we use all the layer representations as the context, based on which we output a weight matrix that shares the same dimensionality with \mathbf{H}^l . Accordingly, the weights are adapted during inference depending on the input layer combination.

Our approach has two strengths. First, it is a more flexible strategy to dynamically combine layers by capturing contextual information among them, which is ignored by the conventional version. Second, the transformation matrix $\text{FFN}_l(\cdot)$ offers the ability to assign a distinct weight to each state in the layers, while its static counterpart fails to exploit such strength since the length of input layers J varies across sentences thus cannot be pre-defined.

Layer Aggregation as Capsule Routing

The goal of layer aggregation is to find a whole representation of the input from partial representations captured by different layers. This is identical to the aims of capsule network, which becomes an appealing alternative to solving the problem of *assigning parts to wholes* (Hinton, Krizhevsky, and Wang 2011). Capsule network employs a fast iterative process called *routing-by-agreement*. Concretely, the basic idea is to iteratively update the proportion of how much a part should be assigned to a whole, based on the agreement between parts and wholes. An important difference between

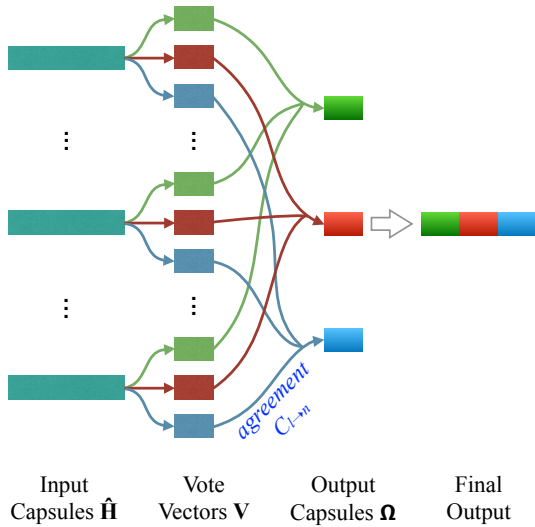


Figure 1: Illustration of the dynamic routing algorithm.

iterative routing and layer aggregation is that the former provides a new way to aggregate information according to the representation of the final output.

A capsule is a group of neurons whose outputs represent different properties of the same entity from the input (Hinton, Sabour, and Frosst 2018). Similarly, a layer consists of a group of hidden states that represent different linguistic properties of the same input (Peters et al. 2018; Anastasopoulos and Chiang 2018), thus each hidden layer can be viewed as a capsule. Given the layers as input capsules, we introduce an additional layer of output capsules and then perform iterative routing between these two layers of capsules. Specifically, in this work we explore two representative routing mechanisms, namely *dynamic routing* and *EM routing*, which differ at how the iterative routing procedure is implemented. We expect layer aggregation can benefit greatly from advanced routing algorithms, which allow the model to allow the model to directly learn the part-whole relationships.

Dynamic Routing Dynamic routing is a straightforward implementation of routing-by-agreement. To illustrate, the information of L input capsules is dynamically routed to N output capsules, which are concatenated to form the final output $\hat{\mathbf{H}} = [\Omega_1, \dots, \Omega_N]$, as shown in Figure 1. Each vector output of capsule n is calculated with a non-linear “squashing” function (Sabour, Frosst, and Hinton 2017):

$$\Omega_n = \frac{\|\mathbf{S}_n\|^2 \mathbf{S}_n}{1 + \|\mathbf{S}_n\|^2 \|\mathbf{S}_n\|} \quad (5)$$

$$\mathbf{S}_n = \sum_{l=1}^L C_{l \rightarrow n} V_{l \rightarrow n} \quad (6)$$

where \mathbf{S}_n is the total input of capsule Ω_n , which is a weighted sum over all “vote vectors” $V_{* \rightarrow n}$ transformed from the input capsules $\hat{\mathbf{H}}$:

$$V_{l \rightarrow n} = \mathbf{W}_{l \rightarrow n} \hat{\mathbf{H}}^l \quad (7)$$

Algorithm 1 Iterative Dynamic Routing. Input: input capsules $\hat{\mathbf{H}} = \{\hat{\mathbf{H}}^1, \dots, \hat{\mathbf{H}}^L\}$, iterations T ; Output: capsules $\Omega = \{\Omega_1, \dots, \Omega_N\}$.

```

1: procedure ROUTING( $\hat{\mathbf{H}}, T$ ):
2:    $\forall (\hat{\mathbf{H}}^l, \Omega_n): B_{l \rightarrow n} = 0$ 
3:   for  $T$  iterations do
4:      $\forall (\hat{\mathbf{H}}^l, \Omega_n): C_{l \rightarrow n} = \text{softmax}(B_{l \rightarrow n})$ 
5:      $\forall \Omega_n$ : compute  $\Omega_n$  by Eq. 5
6:      $\forall (\hat{\mathbf{H}}^l, \Omega_n): B_{l \rightarrow n} += \Omega_n \cdot V_{l \rightarrow n}$ 
   return  $\Omega$ 

```

where $\mathbf{W}_{l \rightarrow n}(\cdot)$ is a trainable transformation matrix, and $\hat{\mathbf{H}}^l$ is an input capsule associated with input layer \mathbf{H}^l :

$$\hat{\mathbf{H}}^l = F_l(\mathbf{H}^1, \dots, \mathbf{H}^L) \quad (8)$$

where $F_l(\cdot)$ is a distinct transformation function.¹ $C_{l \rightarrow n}$ is the assignment probability (i.e. *agreement*) that is determined by the iterative dynamic routing.

Algorithm 1 lists the algorithm of iterative dynamic routing. The assignment probabilities associated with each input capsule $\hat{\mathbf{H}}^l$ sum to 1: $\sum_n C_{l \rightarrow n} = 1$, and are determined by a “routing softmax” (Line 4):

$$C_{l \rightarrow n} = \frac{\exp(B_{l \rightarrow n})}{\sum_{n'=1}^N \exp(B_{l \rightarrow n'})} \quad (9)$$

where $B_{l \rightarrow n}$ measures the degree that $\hat{\mathbf{H}}^l$ should be coupled to capsule n (similar to energy function in the attention model (Bahdanau, Cho, and Bengio 2015)), which is initialized as all 0 (Line 2). The initial assignment probabilities are then iteratively refined by measuring the agreement between the vote vector $V_{l \rightarrow n}$ and capsule n (Lines 4-6), which is implemented as a simple scalar product $\alpha_{l \rightarrow n} = \Omega_n \cdot V_{l \rightarrow n}$ in this work (Line 5).

With the iterative routing-by-agreement mechanism, an input capsule prefers to send its representation to output capsules, whose activity vectors have a big scalar product with the vote V coming from the input capsule. Benefiting from the high-dimensional coincidence filtering, capsule neurons are able to ignore all but the most active feature from the input capsules. Ideally, each capsule output represents a distinct property of the input. To make the dimensionality of the final output be consistent with that of hidden layer (i.e. d), the dimensionality of each capsule output is set to d/N .

EM Routing Dynamic routing uses the cosine of the angle between two vectors to measure their agreement: $\Omega_n \cdot V_{l \rightarrow n}$. The cosine saturates at 1, which makes it insensitive to the difference between a quite good agreement and a very good agreement. In response to this problem, Hinton, Sabour, and Frosst (2018) propose a novel Expectation-Maximization routing algorithm.

¹Note that we calculate each input capsule with $F_l(\mathbf{H}^1, \dots, \mathbf{H}^L)$ instead of $F_l(\mathbf{H}^l)$, since the former achieves better performance on translation task by exploiting more context as shown in our experiment section.

Algorithm 2 Iterative EM Routing returns **activation** A^Ω of the output capsules, given the **activation** A^H and **vote** V of the input capsule.

```

1: procedure EM ROUTING( $A^H, V$ ):
2:    $\forall \hat{\mathbf{H}}^l, \Omega_n$ :  $C_{l \rightarrow n} = 1/N$ 
3:   for  $T$  iterations do
4:      $\forall \Omega_n$ : M-Step( $C, A^H, V$ )
5:      $\forall \hat{\mathbf{H}}^l$ : E-Step( $\mu, \sigma, A^\Omega, V$ )
6:      $\forall \Omega_n$ :  $\Omega_n = A_n^\Omega * \mu_n$ 
   return  $\Omega$ 

1: procedure M-STEP( $C, A^H, V$ )
2:    $\triangleright$  hold  $C$  constant, adjust  $(\mu_n, \sigma_n, A_n^\Omega)$  for  $\Omega_n$ 
3:    $\forall \hat{\mathbf{H}}^l$ :  $C_{l \rightarrow n} = C_{l \rightarrow n} * A_i^H$ 
4:   Compute  $\mu_n, \sigma_n$  by Eq. 11 and 12
5:   Compute  $A_n^\Omega$  by Eq. 14

1: procedure E-STEP( $\mu, \sigma, A^\Omega, V$ )
2:    $\triangleright$  hold  $(\mu, \sigma, A^\Omega)$  constant, adjust  $C_{l \rightarrow *}$  for  $\hat{\mathbf{H}}^l$ 
3:    $\forall \Omega_n$ : compute  $C_{l \rightarrow n}$  by Eq. 16

```

Specifically, the routing process fits a mixture of Gaussians using Expectation-Maximization (EM) algorithm, where the output capsules play the role of Gaussians and the means of the activated input capsules play the role of the datapoints. It iteratively adjusts the means, variances, and activation probabilities of the output capsules, as well as the assignment probabilities C of the input capsules, as listed in Algorithm 2. Comparing with the dynamic routing described above, the EM routing assigns means, variances, and activation probabilities for each capsule, which are used to better estimate the agreement for routing.

The activation probability A_i^H of the input capsule $\hat{\mathbf{H}}^l$ is calculated by

$$A_i^H = \mathbf{W}_i^H \hat{\mathbf{H}}^l \quad (10)$$

where \mathbf{W}_i^H is a trainable transformation matrix, and $\hat{\mathbf{H}}^l$ is calculated by Equation 8. The activation probabilities A^H and votes V of the input capsules are fixed during the EM routing process.

M-Step for each Gaussian associated with Ω_n consists of finding the mean μ_n of the votes from input capsules and the variance σ_n about that mean for each dimension h :

$$\mu_n^h = \frac{\sum_l C_{l \rightarrow n} V_{l \rightarrow n}^h}{\sum_l C_{l \rightarrow n}} \quad (11)$$

$$(\sigma_n^h)^2 = \frac{\sum_l C_{l \rightarrow n} (V_{l \rightarrow n}^h - \mu_n^h)^2}{\sum_l C_{l \rightarrow n}} \quad (12)$$

The incremental cost of using an active capsule Ω_n is

$$cost_n^h = \left(\log(\sigma_n^h) + \frac{1 + \log(2\pi)}{2} \right) \sum_l C_{l \rightarrow n} \quad (13)$$

The activation probability of capsule Ω_n is calculated by

$$A_n^\Omega = \text{logistic}(\lambda(\beta_A - \beta_\mu \sum_l C_{l \rightarrow n} - \sum_h cost_n^h)) \quad (14)$$

where β_A is a fixed cost for coding the mean and variance of Ω_n when activating it, β_μ is another fixed cost per input capsule when not activating it, and λ is an inverse temperature parameter set with a fixed schedule. We refer the readers to (Hinton, Sabour, and Frosst 2018) for more details.

E-Step adjusts the assignment probabilities $C_{l \rightarrow *}$ for each input $\hat{\mathbf{H}}^l$. First, we compute the negative log probability density of the vote $V_{l \rightarrow n}$ from $\hat{\mathbf{H}}^l$ under the Gaussian distribution fitted by the output capsule Ω_n it gets assigned to:

$$p_n = \frac{1}{\sqrt{2\pi}(\sigma_n)^2} \exp\left(-\frac{(V_{l \rightarrow n} - \mu_n)^2}{2(\sigma_n)^2}\right) \quad (15)$$

Accordingly, the assignment probability is re-normalized by

$$C_{l \rightarrow n} = \frac{A_n^\Omega p_n}{\sum_{n'} A_{n'}^\Omega p_{n'}} \quad (16)$$

As has been stated above, EM routing is a more powerful routing algorithm, which can better estimate the agreement by allowing active capsules to receive a cluster of similar votes. In addition, it assigns an additional activation probability A to represent the probability of whether each capsule is present, rather than the length of vector.

Experiment

Setting

We conducted experiments on two widely-used WMT14 English \Rightarrow German (En \Rightarrow De) and WMT17 Chinese \Rightarrow English (Zh \Rightarrow En) translation tasks and compared our model with results reported by previous work (Gehring et al. 2017; Vaswani et al. 2017; Hassan et al. 2018). For the En \Rightarrow De task, the training corpus consists of about 4.56 million sentence pairs. We used newstest2013 as the development set and newstest2014 as the test set. For the Zh \Rightarrow En task, we used all of the available parallel data, consisting of about 20 million sentence pairs. We used newsdev2017 as the development set and newstest2017 as the test set. All the data had been tokenized and segmented into subword symbols using byte-pair encoding with 32K merge operations (Sennrich, Haddow, and Birch 2016). We used 4-gram NIST BLEU score (Papineni et al. 2002) as the evaluation metric, and *sign-test* (Collins, Koehn, and Kucerova 2005) for statistical significance test.

We evaluated the proposed approaches on the Transformer model (Vaswani et al. 2017). We followed the configurations in (Vaswani et al. 2017), and reproduced their reported results on the En \Rightarrow De task. The parameters of the proposed models were initialized by the pre-trained model. All the models were trained on eight NVIDIA P40 GPUs where each was allocated with a batch size of 4096 tokens. In consideration of computation cost, we studied model variations with TRANSFORMER-BASE model on En \Rightarrow De task, and evaluated overall performance with TRANSFORMER-BASE and TRANSFORMER-BIG model on both Zh \Rightarrow En and En \Rightarrow De tasks.

#	Model	# Para.	Train	Decode	BLEU	Δ
1	TRANSFORMER-BASE	88.0M	1.79	1.43	27.31	-
2	+ Linear Combination (Dou et al. 2018)	+14.7M	1.57	1.36	27.73	+0.42
3	+ Dynamic Combination	+25.2M	1.50	1.30	28.33	+1.02
4	+ Dynamic Routing	+37.8M	1.37	1.24	28.22	+0.91
5	+ EM Routing	+56.8M	1.10	1.15	28.81	+1.50

Table 1: Translation performance on WMT14 English \Rightarrow German translation task. “# Para.” denotes the number of parameters, and “Train” and “Decode” respectively denote the training (steps/second) and decoding (sentences/second) speeds.

System	Architecture	En \Rightarrow De		Zh \Rightarrow En	
		# Para.	BLEU	# Para.	BLEU
<i>Existing NMT systems</i>					
(Wu et al. 2016)	RNN with 8 layers	N/A	26.30	N/A	N/A
(Gehring et al. 2017)	CNN with 15 layers	N/A	26.36	N/A	N/A
(Vaswani et al. 2017)	TRANSFORMER-BASE	65M	27.3	N/A	N/A
	TRANSFORMER-BIG	213M	28.4	N/A	N/A
(Hassan et al. 2018)	TRANSFORMER-BIG	N/A	N/A	N/A	24.2
<i>Our NMT systems</i>					
<i>this work</i>	TRANSFORMER-BASE	88M	27.31	108M	24.13
	+ EM Routing	144M	28.81 [†]	164M	24.81 [†]
	TRANSFORMER-BIG	264M	28.58	304M	24.56
	+ EM Routing	490M	28.97 [†]	530M	25.00 [†]

Table 2: Comparing with existing NMT systems on WMT14 English \Rightarrow German (“En \Rightarrow De”) and WMT17 Chinese \Rightarrow English (“Zh \Rightarrow En”) tasks. “[†]” indicates statistically significant difference ($p < 0.01$) from the TRANSFORMER baseline.

Results

Model Variations Table 1 shows the results on WMT14 En \Rightarrow De translation task. As one would expect, the linear combination (Row 2) improves translation performance by +0.42 BLEU points, indicating the necessity of aggregating layers for deep NMT models.

All dynamic aggregation models (Rows 3-5) consistently outperform its static counterpart (Row 2), demonstrating the superiority of the dynamic mechanisms. Among the model variations, the simplest strategy – dynamic combination (Row 3) surprisingly improves performance over the baseline model by up to +1.02 BLEU points. Benefiting from the advanced routing-by-agreement algorithm, the dynamic routing strategy can achieve similar improvement. The EM routing further improves performance by better estimating the agreement during the routing. These findings suggest potential applicability of capsule networks to natural language processing tasks, which has not been fully investigated yet.

All the dynamic aggregation strategies introduce new parameters, ranging from 25.2M to 56.8M. Accordingly, the training speed would decrease due to more efforts to train the new parameters. Dynamic aggregation mechanisms only marginally decrease decoding speed, with EM routing being the slowest one, which decreases decoding speed by 19.6%.

Main Results Table 2 lists the results on both WMT17 Zh \Rightarrow En and WMT14 En \Rightarrow De translation tasks. As seen, dynamically aggregating layers consistently improves translation performance across NMT models and language pairs, which demonstrating the effectiveness and universal-

ity of the proposed approach. It is worth mentioning that TRANSFORMER-BASE with EM routing outperforms the vanilla TRANSFORMER-BIG model, with only less than half of the parameters, demonstrating our model could utilize the parameters more efficiently and effectively.

Analysis of Iterative Routing

We conducted extensive analysis from different perspectives to better understand the iterative routing process. All results are reported on the development set of En \Rightarrow De task with “TRANSFORMER-BASE + EM routing” model.

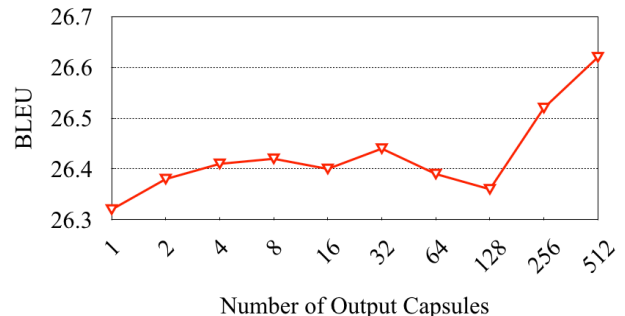


Figure 2: Impact of number of output capsules.

Impact of the Number of Output Capsules The number of output capsules N is a key parameter for our model, as shown in Figure 1. We plot in Figure 2 the BLEU score with

different number of output capsules. Generally, the BLEU score goes up with the increase of the capsule numbers. As aforementioned, to make the dimensionality of the final output be consistent with hidden layer (i.e. d), the dimensionality of each capsule output is d/N . When n increases, the dimensionality of capsule output d/N decreases (the minimum value is 1), which may lead to more subtle representations of different properties of the input.




Figure 3: Impact of routing iterations.

Impact of Routing Iterations Another key parameter is the iteration of the iterative routing T , which affects the estimation of the agreement. As shown in Figure 3, the BLEU score typically goes up with the increase of the iterations T , while the trend does not hold when $T > 3$. This indicates that more iterations may over-estimate the agreement between two capsules, thus harms the performance. The optimal iteration 3 is also consistent with the findings in previous work (Sabour, Frosst, and Hinton 2017; Hinton, Sabour, and Frosst 2018).

Model	Construct $\hat{\mathbf{H}}^l$ with	BLEU
BASE	N/A	25.84
OURS	$F_l(\mathbf{H}^l)$	26.18
	$F_l(\mathbf{H}^1, \dots, \mathbf{H}^L)$	26.62

Table 3: Impact of functions to construct input capsules.

Impact of Functions to Construct Input Capsules For the iterative routing models, we use $\hat{\mathbf{H}}^l = F_l(\mathbf{H}^1, \dots, \mathbf{H}^L)$ instead of $\hat{\mathbf{H}}^l = F_l(\mathbf{H}^l)$ to construct each input capsule $\hat{\mathbf{H}}^l$. Table 3 lists the comparison results, which shows that the former indeed outperforms the latter. We attribute this to that $F_l(\mathbf{H}^1, \dots, \mathbf{H}^L)$ is more representative by extracting features from the concatenation of the original layer representations.

Visualization of Agreement Distribution The assignment probability $C_{l \rightarrow n}$ before M step with $\sum_n C_{l \rightarrow n} = 1$ denotes the agreement between the input capsule $\hat{\mathbf{H}}^l$ and the




Figure 4: Agreement distribution with 6 input capsules (y-axis) and 512 output capsules (x-axis). Darker color denotes higher agreement. The three heatmaps from top to bottom are respectively the 1st to 3rd iterations.

output capsule Ω_n , which is determined by the iterative routing. A higher agreement $C_{l \rightarrow n}$ denotes that the input capsule $\hat{\mathbf{H}}^l$ prefers to send its representation to the output capsule Ω_n . We plot in Figure 4 the alignment distribution in different routing iterations. In the first iteration (top panel), the initialized uniform distribution is employed as the agreement distribution, and each output capsule equally attends to all the input capsules. As the iterative routing goes, the input capsules learn to send their representations to proper output capsules, and accordingly output capsules are more likely to capture distinct features. We empirically validate our claim from the following two perspectives.

We use the entropy to measure the skewness of the agreement distributions:

$$entropy = \frac{1}{N \cdot L} \sum_{n=1}^N \sum_{l=1}^L C_{l \rightarrow n} \log C_{l \rightarrow n} \quad (17)$$

A lower entropy denotes a more skewed distribution, which indicates that the input capsules are more certain about which output capsules should be routed more information. The entropies of the three iterations are respectively 6.24, 5.93, 5.86, which indeed decreases as expected.

To validate the claim that different output capsules focus on different subsets of input capsules, we measure the diversity between each two output capsules. Let $\mathbf{C}_n = \{C_{1 \rightarrow n}, \dots, C_{L \rightarrow n}\}$ be the agreement probabilities assigned to the output capsule Ω_n , we calculate the diversity among all the output capsules as

$$diversity = \frac{1}{\binom{N}{2}} \sum_{i=1}^N \sum_{j=i+1}^N \cos(\mathbf{C}_i, \mathbf{C}_j) \quad (18)$$

A higher diversity score denotes that output capsules attend to different subsets of input capsules. The diversity scores

of the three iterations are respectively 0.0, 0.09, and 0.18, which reconfirm our observations.

Effect on Encoder and Decoder

Model	Applied to		BLEU
	Encoder	Decoder	
BASE	N/A	N/A	25.84
OURS	✓	×	26.33
	×	✓	26.34
	✓	✓	26.62

Table 4: Effect of EM routing on encoder and decoder.

Both encoder and decoder are composed of a stack of L layers, which may benefit from the proposed approach. In this experiment, we investigate how our model affects the two components, as shown in Table 4. Aggregating layers of encoder or decoder individually consistently outperforms the vanilla baseline model, and exploiting both components further improves performance. These results provide support for the claim that aggregating layers is useful for both understanding input sequence and generating output sequence.

Length Analysis

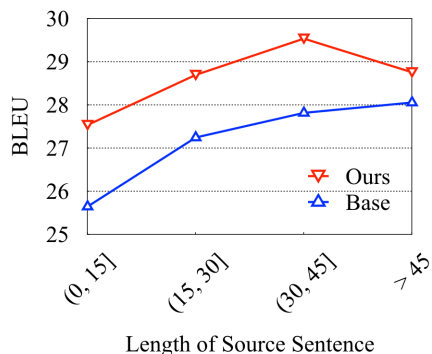


Figure 5: BLEU scores on the En=>De test set with respect to various input sentence lengths.

Following Bahdanau, Cho, and Bengio (2015) and Tu et al. (2016), we grouped sentences of similar lengths together and computed the BLEU score for each group, as shown in Figure 5. Generally, the performance of TRANSFORMER-BASE goes up with the increase of input sentence lengths. We attribute this to the strength of self-attention mechanism to model global dependencies without regard to their distance. Clearly, the proposed approaches outperform the baseline in all length segments.

Related Work

Our work is inspired by research in the field of *exploiting deep representation* and *capsule networks*.

Exploiting Deep Representation Exploiting deep representations have been studied by various communities, from computer vision to natural language processing. He et al. (2016) propose a residual learning framework, combining layers and encouraging gradient flow by simple shortcut connections. Huang et al. (2017) extend the idea by introducing densely connected layers which could better strengthen feature propagation and encourage feature reuse. Deep layer aggregation (Yu et al. 2018) designs architecture to fuse information iteratively and hierarchically.

Concerning natural language processing, Peters et al. (2018) have found that combining different layers is helpful and their model significantly improves state-of-the-art models on various tasks. Researchers have also explored fusing information for NMT models and demonstrate aggregating layers is also useful for NMT (Shen et al. 2018; Wang et al. 2018; Dou et al. 2018). However, all of these works mainly focus on static aggregation in that their aggregation strategy is independent of specific hidden states. In response to this problem, we introduce dynamic principles into layer aggregation. In addition, their approaches are a fixed policy without considering the representation of the final output, while the routing-by-agreement mechanisms are able to aggregate information according to the final representation.

Capsule Networks The idea of dynamic routing is first proposed by Sabour, Frosst, and Hinton (2017), which aims at addressing the representational limitations of convolutional and recurrent neural networks for image classification. The iterative routing procedure is further improved by using Expectation-Maximization algorithm to better estimate the agreement between capsules (Hinton, Sabour, and Frosst 2018). In computer vision community, Xi, Bing, and Jin (2017) explore its application on CIFAR data with higher dimensionality. LaLonde and Bagci (2018) apply capsule networks on object segmentation task.

The applications of capsule networks in natural language processing tasks, however, have not been widely investigated to date. Zhao et al. (2018) testify capsule networks on text classification tasks and Gong et al. (2018) propose to aggregate a sequence of vectors via dynamic routing for sequence encoding. To the best of our knowledge, this work is the first to apply the idea of dynamic routing to NMT.

Conclusion

In this work, we propose several methods to dynamically aggregate layers for deep NMT models. Our best model, which utilizes EM-based iterative routing to estimate the agreement between inputs and outputs, has achieved significant improvements over the baseline model across language pairs. By visualizing the routing process, we find that capsule networks are able to extract most active features shared by different inputs. Our study suggests potential applicability of capsule networks across computer vision and natural language processing tasks for aggregating information of multiple inputs.

Future directions include validating our approach on other

NMT architectures such as RNN (Chen et al. 2018) and CNN (Gehring et al. 2017), as well as on other NLP tasks such as dialogue and reading comprehension. It is also interesting to combine with other techniques (Shaw, Uszkoreit, and Vaswani 2018; Li et al. 2018; Dou et al. 2018; Yang et al. 2018; 2019; Kong et al. 2019) to further boost the performance of Transformer.

References

- Anastasopoulos, A., and Chiang, D. 2018. Tied multitask learning for neural speech translation. In *NAACL*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Chen, M. X.; Firat, O.; Bapna, A.; Johnson, M.; Macherey, W.; Foster, G.; Jones, L.; Niki, P.; Schuster, M.; Chen, Z.; Wu, Y.; and Hughes, M. 2018. The Best of Both Worlds: Combining Recent Advances in Neural Machine Translation. In *ACL*.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*.
- Collins, M.; Koehn, P.; and Kucerova, I. 2005. Clause restructuring for statistical machine translation. In *ACL*.
- Dou, Z.-Y.; Tu, Z.; Wang, X.; Shi, S.; and Zhang, T. 2018. Exploiting deep representations for neural machine translation. In *EMNLP*.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. N. 2017. Convolutional sequence to sequence learning. In *ICML*.
- Gong, J.; Qiu, X.; Wang, S.; and Huang, X. 2018. Information aggregation via dynamic routing for sequence encoding. In *COLING*.
- Hassan, H.; Aue, A.; Chen, C.; Chowdhary, V.; Clark, J.; Federmann, C.; Huang, X.; Junczys-Dowmunt, M.; Lewis, W.; Li, M.; et al. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Hinton, G. E.; Krizhevsky, A.; and Wang, S. D. 2011. Transforming auto-encoders. In *ICANN*.
- Hinton, G. E.; Sabour, S.; and Frosst, N. 2018. Matrix capsules with em routing. In *ICLR*.
- Huang, G.; Liu, Z.; van der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *CVPR*.
- Kong, X.; Tu, Z.; Shi, S.; Hovy, E.; and Zhang, T. 2019. Neural machine translation with adequacy-oriented learning. In *AAAI*.
- LaLonde, R., and Bagci, U. 2018. Capsules for object segmentation. *arXiv*.
- Li, J.; Tu, Z.; Yang, B.; Lyu, M. R.; and Zhang, T. 2018. Multi-head attention with disagreement regularization. In *EMNLP*.
- Meng, F.; Lu, Z.; Tu, Z.; Li, H.; and Liu, Q. 2016. A deep memory-based architecture for sequence-to-sequence learning. In *ICLR Workshop*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*.
- Peters, M. E.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep contextualized word representations. In *NAACL*.
- Sabour, S.; Frosst, N.; and Hinton, G. E. 2017. Dynamic routing between capsules. In *NIPS*.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Neural machine translation of rare words with subword units. In *ACL*.
- Shaw, P.; Uszkoreit, J.; and Vaswani, A. 2018. Self-Attention with Relative Position Representations. In *NAACL*.
- Shen, Y.; Tan, X.; He, D.; Qin, T.; and Liu, T.-Y. 2018. Dense information flow for neural machine translation. In *NAACL*.
- Tu, Z.; Lu, Z.; Liu, Y.; Liu, X.; and Li, H. 2016. Modeling coverage for neural machine translation. In *ACL*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*.
- Wang, Q.; Li, F.; Xiao, T.; Li, Y.; Li, Y.; and Zhu, J. 2018. Multi-layer representation fusion for neural machine translation. In *COLING*.
- Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv*.
- Xi, E.; Bing, S.; and Jin, Y. 2017. Capsule network performance on complex data. *arXiv*.
- Yang, B.; Tu, Z.; Wong, D. F.; Meng, F.; Chao, L. S.; and Zhang, T. 2018. Modeling localness for self-attention networks. In *EMNLP*.
- Yang, B.; Li, J.; Wong, D. F.; Chao, L. S.; Wang, X.; and Tu, Z. 2019. Context-aware self-attention networks. In *AAAI*.
- Yu, F.; Wang, D.; Shelhamer, E.; and Darrell, T. 2018. Deep layer aggregation. In *CVPR*.
- Zeiler, M. D., and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *ECCV*.
- Zhao, W.; Ye, J.; Yang, M.; Lei, Z.; Zhang, S.; and Zhao, Z. 2018. Investigating capsule networks with dynamic routing for text classification. In *EMNLP*.
- Zhou, J.; Cao, Y.; Wang, X.; Li, P.; and Xu, W. 2016. Deep recurrent models with fast-forward connections for neural machine translation. *TACL*.