

Enhancing Context Modeling with a Query-Guided Capsule Network for Document-level Translation

Zhengxin Yang¹²³, Jinchao Zhang³, Fandong Meng³
Shuhao Gu¹², Yang Feng^{12*}, Jie Zhou³

¹ University of Chinese Academy of Sciences

² Key Laboratory of Intelligent Information Processing
Institute of Computing Technology, Chinese Academy of Sciences (ICT/CAS)

³ Pattern Recognition Center, WeChat AI, Tencent Inc, China

{yangzhengxin17z, gushuhao17g, fengyang}@ict.ac.cn
{dayerzhang, fandongmeng, withtomzhou}@tencent.com

Abstract

Context modeling is essential to generate coherent and consistent translation for Document-level Neural Machine Translations. The widely used method for document-level translation usually compresses the context information into a representation via hierarchical attention networks. However, this method neither considers the relationship between context words nor distinguishes the roles of context words. To address this problem, we propose a query-guided capsule networks to cluster context information into different perspectives from which the target translation may concern. Experiment results show that our method can significantly outperform strong baselines on multiple data sets of different domains.

1 Introduction

The encoder-decoder based Neural machine translation (NMT) models (Sutskever et al., 2014; Bahdanau et al., 2014; Wu et al., 2016; Vaswani et al., 2017; Zhang et al., 2019) have made great progresses and drawn much attention in recent years. In practical applications, NMT systems are often fed with a document-level input which requires reference resolution, the consistency of tenses and noun expressions and so on. Many researchers have proven that contextual information is essential to generate coherent and consistent translation for document-level translation (Hardmeier, 2012; Meyer and Webber, 2013; Sim Smith, 2017; Jean et al., 2017; Maruf and Haffari, 2018; Miculicich et al., 2018; Zhang et al., 2018a; Voita et al., 2018; Wang et al., 2017; Tu et al., 2018; Maruf et al., 2019). Despite the great success of the above models, they are designed for sentence-level transla-

tion tasks and exclude contextual information in the model architecture.

On these grounds, the widely used hierarchical attention networks (HAN) was proposed to integrate contextual information in document-level translation (Miculicich et al., 2018). In this method, the context sentences are considered in the form of hierarchical attentions retrieved by the current generation. That is, it utilizes a word-level attention to represent a sentence and then a sentence-level attention to represent all the involved context. In this way, the final attention representation has to encode all the information needed for coherent and consistent translation, including reference information, tenses, expressions and so on. To get the multi-perspective information, it is necessary to distinguish the role of each context word and model their relationship especially when one context word could take on multiple roles (Zhang et al., 2018b). However, this is difficult to realize for the HAN model as its final representation for the context is produced with an isolated relevance with the query word which ignores relations with other context words.

To address the problem, we introduce Capsule Networks into document-level translation which have proven good at modelling the *parts-wholes* relations between low-level capsules and high-level capsules (Hinton et al., 2011; Xiao et al., 2018; Sabour et al., 2017; Hinton et al., 2018; Gu and Feng, 2019). With capsule networks, the words in a context source sentence is taken as low-level capsules and the information of different perspectives is treated as high-level capsules. Then in the dynamic routing process of capsule networks, all the low-level capsules trade off against each other and consider over all the high-level capsules and drop themselves at a proper proportion to the high-level capsules. In this way, the relation among low-level capsules and that between low-

Joint work with Pattern Recognition Center, WeChat AI, Tencent Inc, China.

* Corresponding Author

level capsules and high-level capsules are both explored. In order to make sure high-level capsules indeed cluster information needed by the target translation, we apply capsule networks to both sides of the current sentence and add a regularization layer using Pearson Correlation Coefficients to force the high-level capsules on the two sides to approach to each other. In addition, we still need to ensure the final output of capsule networks is relevant to the current sentence. Therefore we propose a Query-guided Capsule Network (QCN) to have the current source sentence to take part in the routing process so that high-level capsules can retain information related to the current source sentence.

To the best of our knowledge, this is the first work which applies capsule networks to document-level translation tasks and QCN is also the first attempt to customize attention for capsule networks in translation tasks. We conducted experiments on three English-German translation data sets in different domains and the results demonstrate that our method can significantly improve the performance of document-level translation compared with strong baselines.

2 Background

2.1 Sentence-level NMT and Transformer Model

Sentence-level NMTs are generally based on an encoder-decoder framework (Sutskever et al., 2014; Bahdanau et al., 2014; Wu et al., 2016; Vaswani et al., 2017; Zhang et al., 2019; Meng and Zhang, 2019). In this framework, the encoder encodes the source sentence $\mathbf{x} = \{x_1, x_2, \dots, x_M\}$ into a sequence of continuous representations $\mathbf{z} = \{z_1, z_2, \dots, z_M\}$. Given \mathbf{z} , the decoder predicts target words in order and returns the target translation $\mathbf{y} = \{y_1, y_2, \dots, y_N\}$. The objective function of NMT is to maximize the log-likelihood of a set of source-target language sentence pairs as

$$\mathcal{L}(\Theta; \mathbf{x}, \mathbf{y}) = \frac{1}{|\mathcal{S}|} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{S}} \log P(\mathbf{y} | \mathbf{x}; \Theta) \quad (1)$$

where \mathcal{S} is the training set.

As we implement our approaches based on the Transformer architecture (Vaswani et al., 2017), which is a strong sentence-level NMT baseline, we give a brief description of the Transformer model. The encoder and decoder are composed of similar layers which consists of two general types of

sub-layers: multi-head attention mechanism and point-wise fully connected feed-forward network (FFN).

Encoder: The encoder of the Transformer is composed of N identical layers. Each of these layers includes a multi-head self-attention mechanism that allows each position of the output of previous encoder layer to attend to all other positions, as well as a position-wise FFN, which is stacked on top of the multi-head self-attention, composed of two linear transformations and a ReLU activation function.

Decoder: The architecture of the decoder is similar to the encoder, however, it employs an additional multi-head attention sub-layer over the encoder output between the multi-head self-attention sub-layer and position-wise FFN sub-layer. The multi-head self-attention sub-layer needs to mask the input target tokens in the future.

2.2 Document-level NMT

The document-level translation task is to translate each source sentence with consideration of the previous context in the document. Formally, the translation of a document \mathbf{D} containing $|\mathbf{D}| = J$ sentence pairs can be defined as given the source document \mathbf{X} in order and the translation system generates each translation $\mathbf{y}^j \in \mathbf{Y}$ in order. The document translation probability can be defined as:

$$P(\mathbf{Y} | \mathbf{X}; \Theta) = \prod_{j=1}^J P(\mathbf{y}^j | \mathbf{x}^j, \mathbf{D}_{<j}; \Theta) \quad (2)$$

where $\mathbf{x}^j, \mathbf{y}^j$ denote the j th source and target sentence respectively, and $\mathbf{D}_{<j}$ denotes all of the previous sentence pairs in document. For each sentence \mathbf{x}^j , each target word is generated according to the source representation and the generated target hypothesis, therefore the Eq. (2) can be formulated as:

$$P(\mathbf{Y} | \mathbf{X}; \Theta) = \prod_{j=1}^J \prod_{i=1}^{I^j} P(y_i^j | \mathbf{y}_{<i}^j, \mathbf{x}^j, \mathbf{D}_{<j}; \Theta) \quad (3)$$

where y_i^j denotes the i th word of the j th translation \mathbf{y}^j with the length as I^j and $\mathbf{y}_{<i}^j$ denotes the generated target hypothesis.

The training objective of document-level NMTs is to maximize the log-likelihood of translations in

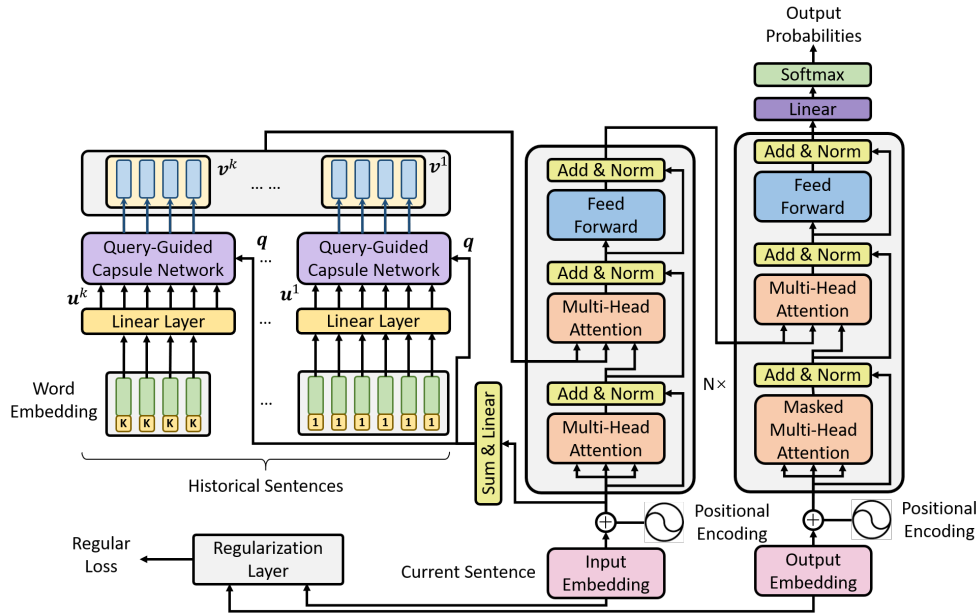


Figure 1: The overall architecture consists of three modules: the *Query-guided Capsule Network* in the upper left of this figure, the *Regularization Layer* in the lower left of this figure and the *Sub-layer-expanded Transformer* in the right of this figure.

document context as following:

$$\mathcal{L}(\Theta; \mathbf{X}, \mathbf{Y}) = \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{X}, \mathbf{Y}) \in \mathcal{D}} \log P(\mathbf{Y} | \mathbf{X}; \Theta) \quad (4)$$

where \mathcal{D} is the training set of the DocNMT.

Compared with the sentence-level NMTs, the critical part of document-level NMTs is to effectively capture and utilize the related contextual information when translating the to-be-translated source sentence.

3 Our Approach

In this section, we introduce the proposed Query-guided Capsule Network (QCN) for enhancing the document-level NMT. First, we present the overall architecture of the network, and then we describe the QCN in detail.

3.1 Overall Architecture

We aim to enhance the document-level NMT performance through effectively capturing and employing the contextual features in each historical sentence that related to the current sentence. We integrating a novel Query-guided Capsule Network into the sentence-level Transformer-based NMT (Vaswani et al., 2017) to capture the document-level contextual information for translating the current source sentence.

As shown in Figure 1, the overall architecture of our translation model is composed of three modules:

- **The Query-guided Capsule Network** takes the to-be-translated source sentence as the query to guide the procedure of retrieving related and helpful contextual features from historical sentences with a novel dynamic routing algorithm.
- **Sub-layer-expanded Transformer** contains a new sub-layer that attending the contextual features extracted from the QCN to effectively utilize them for translation.
- **Regularization Layer** contains two conventional Capsule Networks to unify the source sentence and the target sentence into an identical semantic space through computing an extra PCCs loss item at the training stage.

3.2 Query-guided Capsule Network

The Capsule Network (CapsNet) (Sabour et al., 2017) was proposed to build parts-wholes relationships in the iterative routing procedure, which can be used to capture features in historical sentences from low level to high level. Capsules in the lower layer vote for those in the higher layer by aggregating their transformations with iteratively updated




Figure 2: Query-guided Capsule Network, the green blocks in the middle of the figure are lower-level capsules u_i , right-side blocks indicate the higher-level capsules v_j and the query vector q is initially tiled to length $|v|$ and updated with corresponding higher-level capsules. Here shows a Query-guided Capsule Network with 3 high-level (feature) capsules. In each iteration, the PPCs is computed according to the input vector u_i and the query vector q_j , then $p_{ij}\hat{u}_{j|i}$ and $c_{ij}\hat{u}_{j|i}$ are added to obtain the higher-level capsules which will be used to update corresponding query later.

coupling coefficients. However, there exists an obvious drawback of directly applying the Capsule Network into the document-level NMT for capturing contextual features. The reason is that the CapsNet can only extract internal features without considering whether features are related to the to-be-translated source sentence.

To address this issue, we proposed the Query-guided Capsule Network (QCN), which employ the representation of the to-be-translated source sentence as the query to guide the feature extraction procedure of the Capsule Network. In this way, contextual features that generated from the QCN are based on the internal semantic relations among each historical sentence and the external semantic relations between historical sentences and the to-be-translated source sentence. The QCN is based on an improved dynamic routing algorithm which will be detailed introduced in the next section.

Improved Dynamic Routing of QCN

Given a query vector q and a set of input capsules $u = \{u_1, u_2, \dots, u_n\}$, the dynamic routing algorithm iteratively calculates the correlation between the query vector and each input capsule and updates output capsules. Specifically, query vector q is the representation of source to-be-translated sentence and input capsules u are all word embeddings in previous sentences, which can be for-

Algorithm 1 Improved Dynamic Routing

Input: r, q and $u = \{u_1, u_2, \dots, u_n\}$
Output: $v = \{v_1, v_2, \dots, v_m\}$

- 1: // Initialization
- 2: **for** each output capsule v_j in higher-layer **do**
- 3: **for** each input capsule u_i in lower-layer **do**
- 4: $\hat{u}_{j|i} \leftarrow W_{ij}u_i$
- 5: $\alpha_{ij} \leftarrow 0$
- 6: $q_j \leftarrow q$
- 7: $p_{ij} \leftarrow \tanh(\text{PCCs}(u_i, q_j))$
▷ PCCs computes Eq 7
- 8: **end for**
- 9: **end for**
- 10: **end for**
- 11:
- 12: // Iteration
- 13: $itr \leftarrow 0$
- 14: **repeat**
- 15: **for** each input capsule u_i in lower-layer **do**
- 16: $c_i \leftarrow \text{softmax}(\alpha_i)$
- 17: **end for**
- 18:
- 19: **for** each output capsule v_j in higher-layer **do**
- 20: $s_j \leftarrow \sum_{i=1}^n (c_{ij} + p_{ij})\hat{u}_{j|i}$
- 21: $v_j \leftarrow \text{squash}(s_j)$
▷ squash computes Eq 8
- 22: **end for**
- 23: **end for**
- 24:
- 25: **for** each output capsule v_j in higher-layer **do**
- 26: **for** each input capsule u_i in lower-layer **do**
- 27: $\alpha_{ij} \leftarrow \alpha_{ij} + p_{ij}\hat{u}_{j|i}v_j$
- 28: **end for**
- 29: $q_j \leftarrow \frac{q_j + v_j}{2}$
- 30: **for** each input capsule u_i in lower-layer **do**
- 31: $p_{ij} \leftarrow \tanh(\text{PCCs}(u_i, q_j))$
- 32: **end for**
- 33: **end for**
- 34: $itr \leftarrow itr + 1$
- 35: **until** $itr = r$

mulated as:

$$q = g\left(\sum_{x \in x} \text{embedding}(x)\right) \quad (5)$$

$$u_i^{-k} = f([\text{embedding}(x_i^{-k}); \text{onehot}(k)]) \quad (6)$$

where f and g are both linear transformation functions and k indicates the distance of historical sentence from the to-be-translated sentence. Each word embedding in historical sentences is concatenated with a distance-determined one-hot vector to provide positional markers. Compared to the dynamic routing method proposed by (Sabour et al., 2017), our improved dynamic routing method of QCN can model information that related to the query among input capsules.

Algorithm 1 shows details of the algorithm and Figure 2 illustrates the interaction of the various components in the QCN. Initially, the improved dynamic routing algorithm of QCN receives a sequence of lower-level capsules $u = \{u_1, u_2, \dots, u_n\}$ and a query vector q and

then calculates a Pearson Correlation Coefficients (PCCs) between \mathbf{q} and each input capsule \mathbf{u}_i (line 7). The PCCs is a measure of the linear correlation between two variables. When PCCs is close to +1, it means they have very strong positive linear correlation, and close to -1 means total negative correlation. Given a pair of variables (\mathbf{A}, \mathbf{B}) , the formula for computing PCCs is

$$\begin{aligned} \text{PCCs}(\mathbf{A}, \mathbf{B}) &= \frac{\text{Cov}(\mathbf{A}, \mathbf{B})}{\sigma_A \sigma_B} \\ &= \frac{(\mathbf{A} - \frac{1}{n} \sum_{i=1}^n a_i)^T (\mathbf{B} - \frac{1}{n} \sum_{i=1}^n b_i)}{\|\mathbf{A}\| \cdot \|\mathbf{B}\|} \quad (7) \end{aligned}$$

where \mathbf{A} and \mathbf{B} are both n -dimension vectors, Cov is the covariance and σ_A, σ_B is the standard deviation of \mathbf{A} and \mathbf{B} respectively.

The routing iteration process then computes coupling coefficients, denoted as c_i (line 16), with regard to a input capsule \mathbf{u}_i and all the higher-level capsules \mathbf{v} . In the original dynamic routing algorithm (Sabour et al., 2017), coupling coefficients are only determined by the cumulative ‘‘agreement’’, which are the prior probabilities that capsule \mathbf{u}_i should be coupled to capsule \mathbf{v}_j . However, in DocNMT situation, it is far from enough to cluster the high-level information from the capsules in lower-level that related to query vector according to a naive ‘‘agreement’’. To address this issue, we reduce the ‘‘agreement’’, when PCCs show a negative linear correlation between the query vector \mathbf{q} and the input vector \mathbf{u}_i . Instead, we increase the ‘‘agreement’’, when PCCs are positive (line 27). The query vector \mathbf{q} is initially tiled to length $|\mathbf{v}|$ and updated with the corresponding higher-level capsules \mathbf{v}_j in each iteration (line 29).

Our routing iteration updates higher-level capsules by adding $\sum_{i=1}^n p_{ij} \hat{\mathbf{u}}_{j|i}$ to \mathbf{s}_j . This step can add more information related to the query vector and cut off the unrelated features (line 20). It is necessary to get different length of output capsules shrunk into the 0 to 1 interval using ‘‘squash’’ function (line 21) proposed by Sabour et al. (2017) which is shown in Eq.(8).

$$\text{squash}(\mathbf{t}) = \frac{\|\mathbf{t}\|^2}{1 + \|\mathbf{t}\|^2} \frac{\mathbf{t}}{\|\mathbf{t}\|} \quad (8)$$

where \mathbf{t} can be the initial input capsule \mathbf{u}_i or the vector \mathbf{s}_j for predicting the output capsule.

3.3 Sub-layer-expanded Transformer

To effectively utilize contextual features extracted from each historical sentence by QCN, we intro-

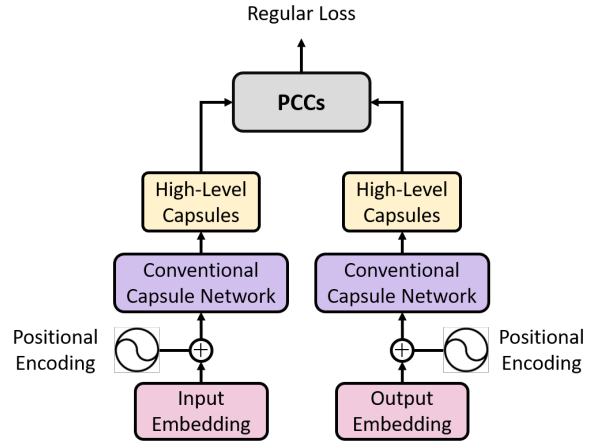


Figure 3: Regularization Layer. In the training stage, this layer takes the inputs of the encoder (left part of the figure) and decoder (right part of the figure) as the input capsules of two capsule networks separately, and computes the PCCs between the high-level capsules of two networks.

duce an additional context-aware multi-head attention sub-layer in each layer of the Transformer encoder. The multi-head attention can attend to all the positions of the contextual features with outputs of the previous sub-layer as a query. Then the output of the context-aware multi-head attention sub-layer is fed into the point-wise feed-forward sub-layer in each layer of the encoder. The right part in Figure 1 shows details of the Sub-layer-expanded Transformer. Specifically, the equation of multi-head attention computing procedure is as following:

$$\begin{aligned} \text{MultiHead}(Q, K, V) &= \text{Concat}(H_1, \dots, H_h) W^O \\ H_i &= \text{Attention}(QW_i^Q, KW_i^K, VW_i^V) \quad (9) \end{aligned}$$

where W_i^Q, W_i^K and W_i^V are the parameter matrices, Q, K and V indicates the query, key and value representations. The computation of the attention function is as following:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (10)$$

where d_k indicates the dimension of queries Q and keys K .

3.4 Regularization Layer

To better modeling the document-level translation task, we incorporate a regularization layer (Figure 3) into the whole architecture to restrict the source sentence and target sentence to an identical semantic space. This layer separately feeds the inputs

of encoder and decoder into two capsule networks Caps_{enc} and Caps_{dec} , and computes the PCCs between the outputs of two networks. We regard the PCCs as an extra regularization term in the final objective at the training stage. The loss function of our model can be formulated as:

$$\begin{aligned} \mathcal{L}(\Theta; \mathbf{X}, \mathbf{Y}) = & \frac{1}{|\mathcal{D}|} \\ & \times \sum_{j=1}^J \sum_{i=1}^{I^j} \{ \log P(y_i^j | \mathbf{y}_{<i}^j, \mathbf{x}^j, \mathbf{D}_{<j}; \Theta) \\ & + \text{PCCs}(\text{Caps}_{enc}(\mathbf{x}^j), \text{Caps}_{dec}(\mathbf{y}^j)) \} \end{aligned} \quad (11)$$

where Θ are parameters of the model, $\mathbf{D}_{<j}$ are historical sentences of the to-be-translated source sentence, \mathbf{x}^j is the to-be-translated sentence and $\mathbf{y}_{<i}^j$ denotes the generated target hypothesis.

4 Experiments

4.1 Settings

Datasets and Evaluation Metrics

We carry out experiments on English-German translation tasks in three different domains: talks, news, and speeches. The corpora statistics are shown in Table 1.

- **TED.** This corpus is a Machine Translation part of the IWSLT 2017 (Cettolo et al., 2012) evaluation campaigns¹, each TED talk is considered to be a document. we take the *tst2016-207* as the test set, and other as our development set.
- **News.** We take the sentence-aligned document-delimited News Commentary v11 corpus² as our training set. The WMT’16 *news-test2015* and *news-test2016* are used for development and testing respectively.
- **Europarl.** The corpus are extracted from the Europarl v7 (Koehn, 2005) according to the method mentioned in Maruf and Haffari (2018). The training, development and test sets are obtained through randomly splitting the corpus.

We download all of above extracted corpora³ from Maruf et al. (2019). The tokenization and

¹<https://wit3.fbk.eu/>

²<http://www.casmat.eu/corpus/news-commentary.html>

³<https://github.com/sameenmaruf/selective-attn/tree/master/data>

truecase pre-processing are implemented on all datasets using the scripts of the Moses Toolkit⁴ (Koehn et al., 2007). We also apply segmentation into BPE subword units⁵ (Sennrich et al., 2016) with 30K merge operations.

We use two metrics: BLEU (Papineni et al., 2002) and Meteor (Lavie and Agarwal, 2007) to evaluate the translation quality.

Models and Baselines

We use the Transformer architecture as our context-agnostic baseline and adopt three context-aware baselines (Zhang et al., 2018a; Miculicich et al., 2018; Maruf and Haffari, 2018).

We performed the same configuration on our models according to the settings of the Maruf and Haffari (2018). Specifically, for the Transformer, we set the hidden size and point-wise FFN size as 512 and 2048 respectively. We use 4 layers and 8 attention heads in both encoder and decoder. All dropout rates are set to 0.1 for context-agnostic model and 0.2 for context-aware model.

In the training phase, we use the default Adam optimizer (Kingma and Ba, 2014) with a fixed learning rate of 0.0001. The batch size is 1500 on TED dataset and 900 on both News and Europarl datasets.

4.2 Results and Analysis

Main Results

Table 2 shows that our model surpasses all the context-agnostic (Vaswani et al., 2017) and context-aware (Zhang et al., 2018a; Miculicich et al., 2018; Maruf and Haffari, 2018) baselines on TED and Europarl datasets. For TED dataset, the performance of our model greatly exceeds that of all other baselines, and is better than Miculicich et al. (2018) with a gain of +0.59 BLEU and +0.61 Meteor. For Europarl dataset, our model got improvements with a gain of +0.07 on BLEU metric, but the Meteor score is +0.64 higher than Maruf et al. (2019) which utilize the whole document as the contextual information, whereas we only using 3 previous sentences.

Results on the sequence-level Transformer and our DocNMTs show that the captured contextual features provide helpful semantic information for enhancing the translation quality. The regularization term that we proposed can effectively

⁴<https://github.com/moses-smt/mosesdecoder/tree/master/scripts>

⁵<https://pytorch.org/project/subword-nmt/>

	TED		News		Europarl	
	Sent No.	Doc len avg	Sent No.	Doc len avg	Sent No.	Doc len avg
Training	206,126	121.39	236,287	38.93	1,666,904	14.14
Development	8,967	96.42	2,169	26.78	3,587	14.95
Test	2,271	98.74	2,999	19.35	5,134	14.26

Table 1: The statistics of the training/development/test corpora in number of sentence pairs and the average document length (in sentences).

Model	TED		News		Europarl	
	BLEU	Meteor	BLEU	Meteor	BLEU	Meteor
Transformer-DocNMT (Zhang et al., 2018a)	24.00	44.69	23.08	42.40	29.32	46.72
HAN-DocNMT (Miculicich et al., 2018)	24.58	45.48	25.03	44.02	28.60	46.09
SAN-DocNMT (Maruf et al., 2019)	24.42	45.38	24.84	44.27	29.75	47.22
Transformer	23.28	44.17	21.67	41.11	28.72	46.22
Transformer + Regularization Term	24.55	45.57	22.09	41.77	29.42	47.59
Transformer + QCN	24.41	46.09	22.22	41.90	29.48	47.49
Transformer + QCN + Regularization Term	25.19	45.91	22.37	41.88	29.82	47.86

Table 2: BLEU and Meteor scores of models. There exists three context-aware baseline models and a context-agnostic model. ‘‘Regularization Term’’ denotes integrating the regularization layer into the Transformer model and ‘‘QCN’’ denotes incorporating Query-guided Capsule Network in the Transformer model.




Figure 4: Effect of Contextual Information Scope. It is not that the more historical sentences we utilize, the better translation performance is. Effect of scope on BLEU and Meteor scores are inconsistent to some degree.

further improve the model performance on TED and Europarl datasets. For the restriction of the GPU memory, we have to filter long sentences to keep our model running. Although, it hurts the model performance on the ‘‘News’’ dataset (contains many long sentences), the QCN module and regularization term still bring improvements.

Effect of Contextual Information Scope

To investigate the effect of contextual information scope, we carry on the number of historical sentences hyper-parameter experiments on the TED talk dataset. We fix the hyper-parameters of the QCN by setting both the number of higher-level

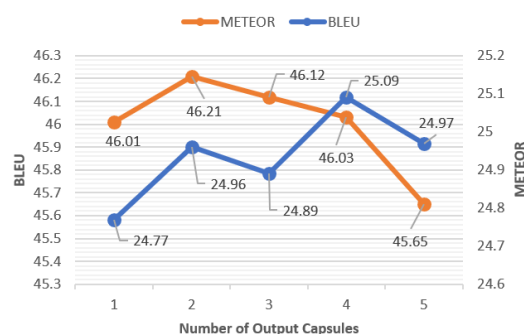


Figure 5: Effect of Feature Capsule Number. Appropriate capsule number is important for model performance.

capsules and routing iteration to four, and investigating the impact of changes in the number of historical sentences on BLEU and Meteor scores. Figure 4 shows that using one historical sentence in QCN can obtain the best Meteor score while the highest BLEU score is presented when we utilize two historical sentences. We found that the growth of Meteor and BLEU scores are opposite. Therefore, the choice of utilizing how many historical sentences is a trade-off between both scores. We choose three historical sentences as our final setting. Through experimentation, we also found it is not that the more historical sentences we utilize, the better translation performance is.




Figure 6: Agreement distribution of a sentence (26 words) as inputs of the QCN and 4 contextual features as output. Blue color means higher agreement and the red means lower. From top to bottom are four heat maps in 0_{th} to 3_{th} iterations.

Effect of Feature Capsule Number

QCN is the crucial part of our overall architecture and the positive and negative impact depends on the configuration of the QCN. Therefore, We also investigate the effect of the hyper-parameter of QCN: the number of feature capsules.

We set the number of historical sentence as 3 according to the previously experimental results, and the number of routing iteration is set to 4. Figure 5 shows that Meteor score become highest when the number of higher-level capsules is set as 2, but BLEU score can obtain best score at 4. We finally choose 4 as the final setting because both BLEU and Meteor can obtain relatively good results.

Visualization of Agreement and PCCs

Coupling Coefficients (CCs) can indirectly reflect the variation of the ‘‘agreement’’, so we visualize the coefficients in each routing iteration at the stage of decoding as shown in Figure 6. In the first iteration, all coupling coefficients are initialized in a uniform distribution, all higher-level capsules are voted by lower-level capsules equally. Then, lower-level capsules are iteratively trained to send more information to the proper higher-level capsule. Figure 6 shows that most of input capsules are tend to vote the 2_{ed} feature capsule finally.

Different from the CCs, PCCs show the linear correlation between query and inputs. Initially, the query vector is tiled with the number of feature capsules and is used to calculate the PCCs with each input capsules, as the Figure 7 shows that the color of every column is identical. As the iterative routing begins, each query is updated according to the higher-level capsules. The function of the PPCs is to increase or decrease the coupling coefficients with the positive or negative

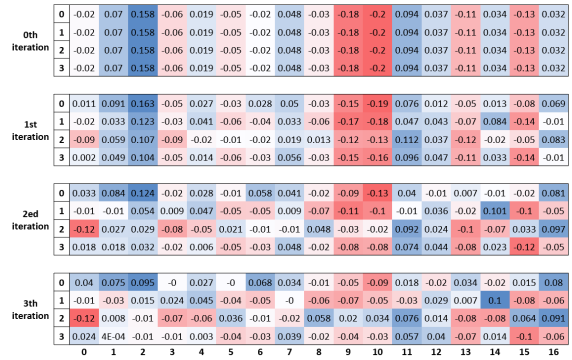


Figure 7: PCCs distribution of a sentence has 17 words as inputs of the QCN and 4 duplicated query. Blue color denotes positive correlation and red means negative. From top to bottom are four heat maps in 0_{th} to 3_{th} iterations.

value. See Figure 7, we can find that PCCs varies as iteration changes.

5 Related Work

Document-level Machine Translation

Document-level machine translation became a hot research direction in the later stage of statistical machine translation era. [Hardmeier and Federico \(2010\)](#) represented the links between word pairs in the context using a word dependency model for SMT to improve the translation of anaphoric pronouns. [Hardmeier et al. \(2012, 2013\)](#) first proposed a new document-level SMT paradigm that translates whole documents as units. However, in this period, most of the work has not achieved too many compelling results or has been only focused on a part of difficulties.

With the coming of the era of Neural Machine Translation, many works began to focus on Document-level NMT tasks. [Xiong et al. \(2019\)](#) trained a reward teacher to refine the translation quality from a document perspective. [Tiedemann and Scherrer \(2017\)](#) simply concatenated sentences in one document as models’ input or output. [Jean et al. \(2017\)](#) used additional context encoder to capture larger-context information. [Kuang et al. \(2017\); Tu et al. \(2018\)](#) used a cache to memorize most relevant words or features in previous sentences or translations.

Recently, several studies integrated additional modules into the Transformer-based NMTs for modeling contextual information. ([Voita et al., 2018; Zhang et al., 2018a](#)), [Maruf and Haffari \(2018\)](#) proposed a document-level NMT using a

memory-networks, and Wang et al. (2017) and Micolich et al. (2018) integrated hierarchical attention network in RNN-based NMT or Transformer to model the document-level information. Maruf et al. (2019) used the whole document as the contextual information and firstly divided document-level translation tasks into two types: offline and online.

Capsule Networks

Hinton et al. (2011) proposed the capsule conception to use vector for describing the pose of an object. The dynamic routing algorithm was proposed by Sabour et al. (2017) to build the part-whole relationship through the iterative routing procedure. Hinton et al. (2018) designed a new routing style based on the EM algorithm. Some researchers investigated to apply the capsule network for various tasks. Wang et al. (2018) investigated a novel capsule network with dynamic routing for linear time NMT. Yang et al. (2018) explored capsule networks for text classification with strategies to stabilize the dynamic routing process. Gu and Feng (2019) introduces capsule networks into Transformer to model the relations between different heads in multi-head attention. We specifically investigated dynamic routing algorithms for the document-level NMT.

6 Conclusion

We have proposed a novel Query-guided Capsule Network with an improved dynamic routing algorithm for enhancing context modeling for the document-level Neural Machine Translation Model. Experiments on English-German in different domains showed our model significantly outperforms sentence-level NMTs and achieved state-of-the-art performance on two of three datasets, which proved the effectiveness of our approaches.

Acknowledgments

We thank the anonymous reviewers for their insightful comments. This work was supported by the National Natural Science Foundation of China (NO.61662077, NO.61876174) and National Key R&D Program of China (NO.2017YFE9132900).

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly

learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. Wit³: Web inventory of transcribed and translated talks. In *Proceedings of the 16th Conference of the European Association for Machine Translation (EAMT)*, pages 261–268, Trento, Italy.

Shuhao Gu and Yang Feng. 2019. Improving multi-head attention with capsule networks. In *Proceedings of the 8th CCF International Conference on Natural Language Processing and Chinese Computing*.

Christian Hardmeier. 2012. Discourse in statistical machine translation. a survey and a case study. *Discours. Revue de linguistique, psycholinguistique et informatique. A journal of linguistics, psycholinguistics and computational linguistics*, (11).

Christian Hardmeier and Marcello Federico. 2010. Modelling pronominal anaphora in statistical machine translation. In *IWSLT (International Workshop on Spoken Language Translation); Paris, France; December 2nd and 3rd, 2010.*, pages 283–289.

Christian Hardmeier, Joakim Nivre, and Jörg Tiedemann. 2012. Document-wide decoding for phrase-based statistical machine translation. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1179–1190, Jeju Island, Korea. Association for Computational Linguistics.

Christian Hardmeier, Sara Stymne, Jörg Tiedemann, and Joakim Nivre. 2013. Docent: A document-level decoder for phrase-based statistical machine translation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 193–198, Sofia, Bulgaria. Association for Computational Linguistics.

Geoffrey E Hinton, Alex Krizhevsky, and Sida D Wang. 2011. Transforming auto-encoders. In *International Conference on Artificial Neural Networks*, pages 44–51. Springer.

Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. 2018. Matrix capsules with EM routing. In *International Conference on Learning Representations*.

Sébastien Jean, Stanislas Lauly, Orhan Firat, and Kyunghyun Cho. 2017. Does neural machine translation benefit from larger context? *CoRR*, abs/1704.05135.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86.

- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th annual meeting of the association for computational linguistics companion volume proceedings of the demo and poster sessions*, pages 177–180.
- Shaohui Kuang, Deyi Xiong, Weihua Luo, and Guodong Zhou. 2017. Cache-based document-level neural machine translation. *arXiv preprint arXiv:1711.11221*.
- Alon Lavie and Abhaya Agarwal. 2007. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments. In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 228–231. Association for Computational Linguistics.
- Sameen Maruf and Gholamreza Haffari. 2018. Document context neural machine translation with memory networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1275–1284, Melbourne, Australia. Association for Computational Linguistics.
- Sameen Maruf, André FT Martins, and Gholamreza Haffari. 2019. Selective attention for context-aware neural machine translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Minneapolis, Minnesota. Association for Computational Linguistics.
- Fandong Meng and Jinchao Zhang. 2019. Dtmmt: A novel deep transition architecture for neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 224–231.
- Thomas Meyer and Bonnie Webber. 2013. Implication of discourse connectives in (machine) translation. In *Proceedings of the Workshop on Discourse in Machine Translation*, pages 19–26.
- Lesly Miculicich, Dhananjay Ram, Nikolaos Pappas, and James Henderson. 2018. Document-level neural machine translation with hierarchical attention networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2947–2954, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. 2017. Dynamic routing between capsules. In *Advances in neural information processing systems*, pages 3856–3866.
- 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, Berlin, Germany. Association for Computational Linguistics.
- Karin Sim Smith. 2017. On integrating discourse in machine translation. In *Proceedings of the Third Workshop on Discourse in Machine Translation*, pages 110–121, Copenhagen, Denmark. 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.
- Jörg Tiedemann and Yves Scherrer. 2017. Neural machine translation with extended context. In *Proceedings of the Third Workshop on Discourse in Machine Translation*, pages 82–92, Copenhagen, Denmark. Association for Computational Linguistics.
- Zhaopeng Tu, Yang Liu, Shuming Shi, and Tong Zhang. 2018. Learning to remember translation history with a continuous cache. *Transactions of the Association for Computational Linguistics*, 6:407–420.
- 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.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. Context-aware neural machine translation learns anaphora resolution. In *ACL*.
- Longyue Wang, Zhaopeng Tu, Andy Way, and Qun Liu. 2017. Exploiting cross-sentence context for neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2826–2831, Copenhagen, Denmark. Association for Computational Linguistics.
- Mingxuan Wang, Jun Xie, Zhixing Tan, Jinsong Su, Deyi Xiong, and Chao Bian. 2018. Towards linear time neural machine translation with capsule networks. *CoRR*, abs/1811.00287.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between

human and machine translation. *arXiv preprint arXiv:1609.08144*.

Liqiang Xiao, Honglun Zhang, Wenqing Chen, Yongkun Wang, and Yaohui Jin. 2018. Mcapsnet: Capsule network for text with multi-task learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4565–4574.

Hao Xiong, Zhongjun He, Hua Wu, and Haifeng Wang. 2019. Modeling coherence for discourse neural machine translation. *CoRR*, abs/1811.05683.

Min Yang, Wei Zhao, Jianbo Ye, Zeyang Lei, Zhou Zhao, and Soufei Zhang. 2018. [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, Brussels, Belgium. Association for Computational Linguistics.

Jiacheng Zhang, Huanbo Luan, Maosong Sun, Feifei Zhai, Jingfang Xu, Min Zhang, and Yang Liu. 2018a. [Improving the transformer translation model with document-level context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 533–542, Brussels, Belgium. Association for Computational Linguistics.

Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019. Bridging the gap between training and inference for neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4334–4343. Association for Computational Linguistics.

Wen Zhang, Jiawei Hu, Yang Feng, and Qun Liu. 2018b. Refining source representations with relation networks for neural machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1292–1303.