

Explicit Sentence Compression for Neural Machine Translation

Zuchao Li,^{1,2,3} Rui Wang,^{4,*} Kehai Chen,⁴ Masao Utiyama,⁴ Eiichiro Sumita,⁴
Zhuosheng Zhang,^{1,2,3} Hai Zhao^{1,2,3,*}

¹Department of Computer Science and Engineering, Shanghai Jiao Tong University

²Key Laboratory of Shanghai Education Commission for Intelligent Interaction and Cognitive Engineering, Shanghai Jiao Tong University, Shanghai, China

³MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

⁴National Institute of Information and Communications Technology (NICT), Kyoto, Japan

{charlee, zhangzs}@sjtu.edu.cn, {wangrui, khchen, mutiyama, eiichiro.sumita}@nict.go.jp, zhaohai@cs.sjtu.edu.cn

Abstract

State-of-the-art Transformer-based neural machine translation (NMT) systems still follow a standard encoder-decoder framework, in which source sentence representation can be well done by an encoder with self-attention mechanism. Though Transformer-based encoder may effectively capture general information in its resulting source sentence representation, the backbone information, which stands for the gist of a sentence, is not specifically focused on. In this paper, we propose an explicit sentence compression method to enhance the source sentence representation for NMT. In practice, an explicit sentence compression goal used to learn the backbone information in a sentence. We propose three ways, including backbone source-side fusion, target-side fusion, and both-side fusion, to integrate the compressed sentence into NMT. Our empirical tests on the WMT English-to-French and English-to-German translation tasks show that the proposed sentence compression method significantly improves the translation performances over strong baselines.

1 Introduction

Neural machine translation (NMT) is popularly implemented as an encoder-decoder framework (Vaswani et al. 2017), in which the encoder is right in charge of source sentence representation. Typically, the input sentence is implicitly represented as a contextualized source representation through deep learning networks. By further feeding the decoder, the source representation is used to learn dependent time-step context vectors for predicting target translation (Bojar et al. 2018).

In state-of-the-art Transformer-based encoder, self-attention mechanisms are good at capturing the general

information in a sentence (Dou et al. 2018a; Wang et al. 2019; Yang et al. 2019a). However, it is difficult to distinguish which kind of information lying deeply under the language is really salient for learning source representation. Intuitively, when a person reads a source sentence, he/she often selectively focuses on the basic sentence meaning, and re-reads the entire sentence to understand its meaning completely. Take the English sentence in Table 1 as an example. We manually annotate its basic meaning as a shorter sequence of words than in the original sentence, called backbone information. Obviously, these words with the basic meaning contain more important information for human understanding than the remaining words in the sentence. We argue that such backbone information is also helpful for learning source representation, and is not explicitly considered by the existing NMT system to enrich the source sentence representation.

In this paper, we propose a novel explicit sentence compression approach to enhance the source representation for NMT. To this end, we first design three sentence compression models to accommodate the needs of various languages and scenarios, including supervised, unsupervised, and semi-supervised ways, to learn a backbone information words sequence (as shown in Table 1) from the source sentence. We then propose three translation models, including backbone source-side fusion based NMT (**BSFNMT**), backbone target-side fusion (**BTFNMT**), and both-side fusion based NMT (**BBFNMT**), to introduce this backbone knowledge into the existing Transformer NMT system for improving translation predictions. Empirical results on the WMT14 English-to-German and English-to-French translation tasks show that the proposed approach significantly improves the translation performance over the strong even state-of-the-art NMT baselines.

2 Explicit Sentence Compression

Generally, sentence compression¹ is a typical sequence generation task which aims to maximize the absorption and long-term retention of large amounts of data over a relatively

¹There are many types of sentence compression. In this paper, we focus on abstract sentence summarization.

* Corresponding author. Zuchao Li and Zhuosheng Zhang were internship research fellows at NICT when conducting this work. Hai Zhao was partially supported by National Key Research and Development Program of China (No. 2017YFB0304100) and Key Projects of National Natural Science Foundation of China (No. U1836222 and No. 61733011). Rui Wang was partially supported by JSPS grant-in-aid for early-career scientists (19K20354): “Unsupervised Neural Machine Translation in Universal Scenarios” and NICT tenure-track researcher startup fund “Toward Intelligent Machine Translation”. Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Sentence	Both the <i>US authorities and the Mexican security forces</i> are engaged in an ongoing <i>battle against the drug cartels</i> .
Basic Meaning	US authorities and Mexican forces battle against drug cartels
Backbone supervised ESC	US and Mexican fight drug cartels
Backbone unsupervised ESC	US authorities and Mexican security forces battle drug cartels
Backbone semi-supervised ESC	US authorities and Mexican security forces battle against drug cartels

Table 1: An example of sentence compression.

short sequence for text understanding (Knight and Marcu 2002; Che et al. 2015). To distinguish the importance of words in the sentence and, more importantly, to dig out the most salient part in the sentence representation, we utilize the sentence compression method to explicitly distill the key knowledge that can retain the key meaning of the sentence, termed explicit sentence compression (**ESC**) in this paper. Depending on whether or not the sentence compression is trained using **human annotated data**, the proposed method can be implemented in three ways: supervised ESC, unsupervised ESC, and semi-supervised ESC.

2.1 Supervised ESC

Sentence compression usually relies on large-scale raw data together with their human-labeled data, which can be viewed as supervision, to train a sentence compression model (Rush, Chopra, and Weston 2015a; Hu, Chen, and Zhu 2015; Chopra, Auli, and Rush 2016; Cheng and Lapata 2016; Nallapati et al. 2016; Duan et al. 2019). For example, Nallapati et al. (2016) proposed an attentive encoder-decoder recurrent neural network (RNN) to model abstractive text summarization. Song et al. (2019) further proposed MAsked Sequence to Sequence pre-training (MASS) for the encoder-decoder sentence compression framework which reported state-of-the-art performance on both the Gigaword Corpus and DUC Corpus².

Sentence compression can be conducted by a typical sequence-to-sequence model. The encoder represents the input sentence S as a sequence of annotation vectors, and the decoder depends on the attention mechanism to learn the context vector for generating a compressed version S' with the key meaning of the input sentence. Recently, the new Transformer architecture proposed by Vaswani et al. (2017), which fully relies on self-attention networks, has exhibited state-of-the-art translation performance for several language pairs. We follow this practice and attempt to apply the Transformer architecture to such a compression task.

2.2 Unsupervised ESC

A major challenge in supervised sentence compression is the scarce high quality human annotated parallel data. In practice, due to the lack of parallel annotated data, the supervised sentence compression model cannot be trained or the annotated data domain is different, resulting in the sentence compression model trained on the in-domain performing poorly on the out-of-domain.

Supervised sentence compression models have achieved impressive performances based on large corpora containing pairs of verbose and compressed sentences with human annotation (Nallapati et al. 2016; Song et al. 2019). However, the effectiveness relies heavily on the availability of large amounts of parallel original and human-annotated compressed sentences. This hinders the sentence compression approach from further improvements for many low-resource scenarios. Recently, motivated by recent progress in unsupervised cross-lingual embeddings, the unsupervised NMT (Artetxe et al. 2018; Lample et al. 2018a; 2018b) opened the door to solving the problem of sequence-to-sequence learning without any parallel sentence pairs. It takes advantage of the lossless (ideal situation) nature of machine translation between languages; i.e., it can translate language L_1 to language L_2 and back translate L_2 to language L_1 . However, sentence compression does not have this feature. It is lossy from sentence S to sentence S' , which makes it difficult to restore from the compressed sentence S' to the original sentence S .

Fevry and Phang (2018) added noises to extend the original sentences and trained a denoising auto-encoder to recover the original, constructing an end-to-end training network without any examples of compressed sentences in sequence to sequence framework. In doing so, the model has to exclude and reorder the noisy sentence input, and hence learns to output more semantic important, shorter but grammatically correct sentences. There are two types of noise used in the model: **Additive Sampling Noise** and **Shuffle Noise**.

Additive Sampling Noise: To extend the original sentence, we sample additional sentence from the training dataset randomly, and then sub-sample a subset of words from each without replacement. The newly sampled words are appended to the original sentence.

Shuffle Noise: In order for the model to learn to rephrase the input sentence to make the output shorter, we shuffle the resultant additive noisy sentence.

To gain a better quality for the compressed sentences, we transfer the method of Fevry et al. (2018) into the Transformer architecture instead of their suggested RNN architecture, which makes it conducive to deeper network training and a larger corpus.

2.3 Semi-supervised ESC

As pointed out in Song et al. (2019), sequence to sequence framework has attracted much attention recently due to the advances of deep learning by using large-scale data.

²<https://duc.nist.gov/duc2004/tasks.html>

Many language generation tasks have only a small scale of pair data which can't support to train a deep model with good generalization ability. In comparison, there is a lot of unpaired data which is easier to obtain.

We observe a performance degradation caused by different domains in the supervised ESC. According to the experimental results of Fevry and Phang (2018), the accuracy of the unsupervised ESC is currently lower than the supervised one. Therefore, we have further adopted the semi-supervised explicit sentence compression model to alleviate this problem. Specifically, the unsupervised training (often referred to as pre-training) is performed on the unpaired data first and fine-tuning with the small scale paired data (supervised training) to obtain the ESC model with good performance and generalization ability.

2.4 Compression Rate Control

Explicit compression rate (length) control is a common method which has been used in previous sentence compression works. Kikuchi et al. (2016) examined several methods of introducing target output length information, and found that they were effective without negatively impacting summarization quality. Fan et al. (2018) introduced a length marker token that induces the model to target an output of a desired length, coarsely divided into discrete bins. Fevry et al. (2018) augmented the decoder with an additional length countdown input which is a single scalar that ticks down to 0 when the generation reached the desired length.

Different with the length marker or length countdown input, to induce our model to output the compression sequence with desired length, we use beam search during generation to find the sequence S' that maximizes a score function $s(S', S)$ given a trained ESC model. The length normalization is introduced to account for the fact that we have to compare hypotheses of different length. Without some form of length-normalization regular ln , beam search will favor shorter sequences over longer ones on average since a negative log-probability is added at each step, yielding lower (more negative) scores for longer sentences. Moreover, a coverage penalty cp is also added to favor the sequence that cover the source sentence meaning as much as possible according to the attention weights (Wu et al. 2016).

$$s(S', S) = \log(P(S'|S))/\ln(S') + cp(S; S'), \quad (1)$$

$$\ln(S') = (5 + |S'|)^\alpha / (5 + 1)^\alpha, \quad (2)$$

$$cp(S; S') = \beta \times \sum_{i=1}^{|S|} \log(\min(\sum_{j=1}^{|S'|} p_{i,j}, 1.0)), \quad (3)$$

where $p_{i,j}$ is the attention probability of the j -th target word on the i -th source word. Parameters α and β control the strength of the length normalization and the coverage penalty. Although α can be used to control the compression ratio softly, we use the compression ratio γ to control the maximum length of decoding generation by hard requirements. When the decoding length $|S'|$ is greater than $\gamma|S|$, the decoding stops.

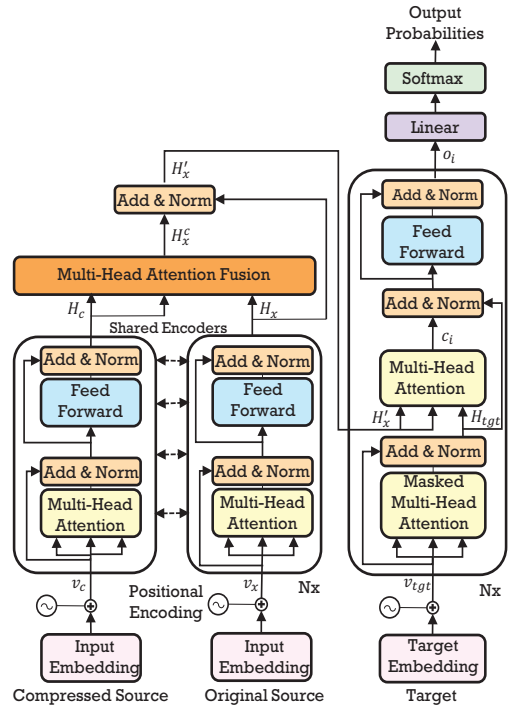


Figure 1: The architecture of proposed BSFNMT model.

3 NMT with ESC

In this section, we first introduce the Transformer networks for machine translation. Then based on the fusion position of the backbone knowledge sequence, we propose three novel translation models: the backbone source-side fusion based NMT model (as shown in Figure 1), the backbone target-side based NMT model (as shown in Figure 2), and the backbone both-side based NMT. All of these models can make use of the source backbone knowledge generated by our sentence compression models.

3.1 Transformer Networks

A Transformer NMT model consists of an encoder and a decoder, which fully rely on self-attention networks (SANs), to translate a sentence in one language into another language with equivalent meaning. Formally, one input sentence $x = \{x_1, \dots, x_J\}$ of length J is first mapped into a sequence of word vectors. Then the sequence and its position embeddings add up to form the input representation $v_x = \{v_1^x, \dots, v_J^x\}$. The sequence $\{v_1^x, \dots, v_J^x\}$ is then packed into a query matrix \mathbf{Q}_x , a key matrix \mathbf{K}_x , and a value matrix \mathbf{V}_x . For the SAN-based encoder, the self-attention sub-layer is first performed over \mathbf{Q} , \mathbf{K} , and \mathbf{V} to the matrix of outputs as:

$$\text{SelfAtt}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_{\text{model}}}}\right)\mathbf{V}, \quad (4)$$

where d_{model} represents the dimensions of the model. Similarly, the translated target words are used to generate the decoder hidden state \mathbf{s}_i at the current time-step i . Generally, the self-attention function is further refined as multi-head

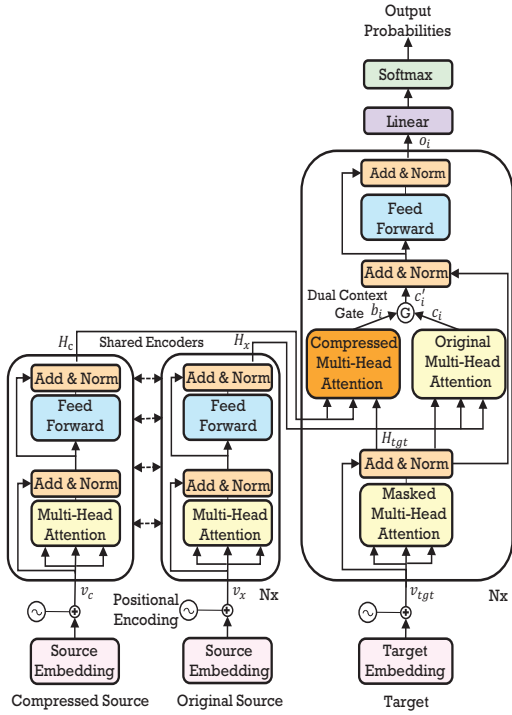


Figure 2: The architecture of proposed BTFNMT model.

self-attention to jointly consider information from different representation subspaces at different positions:

$$\begin{aligned} \text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{Concat}(\text{head}_1, \dots, \text{head}_H) \mathbf{W}^O, \\ \text{head}_h &= \text{SelfAtt}(\mathbf{Q} \mathbf{W}_h^Q, \mathbf{k} \mathbf{W}_h^K, \mathbf{V} \mathbf{W}_h^V), \end{aligned} \quad (5)$$

where the projections are parameter matrices $\mathbf{W}_h^Q \in \mathbb{R}^{d_{model} \times d_k}$, $\mathbf{W}_h^K \in \mathbb{R}^{d_{model} \times d_k}$, $\mathbf{W}_h^V \in \mathbb{R}^{d_{model} \times d_v}$, and $\mathbf{W}^O \in \mathbb{R}^{hd_v \times d_{model}}$. For example, there are $H=8$ heads, d_{model} is 512, and $d_k=d_v=512/8=64$. A position-wise feed-forward network (FFN) layer is applied over the output of multi-head self-attention, and then is added with the matrix \mathbf{V} to generate the final source representation $H_x = \{H_1^x, \dots, H_J^x\}$:

$$H_x = \text{FFN}(\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V})) + \mathbf{V}. \quad (6)$$

The SAN of decoder then uses both H_x and target context hidden state H_{tgt} to learn the context vector o_i by ‘‘encoder-decoder attention’’:

$$c_i = \text{FFN}(\text{MultiHead}(H_{tgt}, H_x, H_x)), \quad (7)$$

$$o_i = c_i + H_{tgt}. \quad (8)$$

Finally, the context vector o_i is used to compute translation probabilities of the next target word y_i by a linear, potentially multi-layered function:

$$P(y_i | y_{<i}, x) \propto \text{Softmax}(\mathbf{L}_o \text{GeLU}(\mathbf{L}_w o_i)), \quad (9)$$

where \mathbf{L}_o and \mathbf{L}_w are projection matrices.

3.2 Backbone Source-side Fusion based NMT

In the backbone source-side fusion based NMT (**BSFNMT**) model, given an input sentence $x = \{x_1, \dots, x_J\}$, there is an additional compressed sequence $x_c = \{x_1^c, \dots, x_K^c\}$ of length K generated by the proposed sentence compression model. This compressed sequence is also input to the SAN shared with the original encoder with word vectors $v_c = \{v_1^c, \dots, v_K^c\}$ in shared vocabulary to learn its final representation $H_c = \{H_1^c, \dots, H_K^c\}$. In the proposed SFNMT model, we introduce an additional multi-head attention layer to fuse the compressed sentence and the original input sentence for learning a more effective source representation.

Specifically, for the multi-head attention-fusion layer, a compressed sentence-specific context representation H_x^c is computed by the multi-head attention on the original sentence representation H_x and the compressed sentence representation H_c :

$$H_x^c = \text{FFN}(\text{MultiHead}(H_x, H_c, H_c)). \quad (10)$$

H_x^c and H_x are added to form a fusion source representation H'_x :

$$H'_x = H_x + H_x^c. \quad (11)$$

Finally, the H'_x instead of H_x is input to the Eq. (7) in turn for predicting the target translations word by word.

3.3 Backbone Target-side Fusion based NMT

In the backbone target-side fusion based NMT (**BTFNMT**) model, both the original sentence and its compressed version are also represented as H_x and H_c respectively by the shared SANs. We then use a tuple (H_x, H_c) instead of the source-side fusion representation H'_x as the input to the decoder. Specifically, we introduce an additional ‘‘encoder-decoder attention’’ module into the decoder to learn the compressed sequence context b_i at the current time-step i :

$$b_i = \text{FFN}(\text{MultiHead}(H_{tgt}, H_c, H_c)). \quad (12)$$

Since we are here to treat the original sentence and the compressed sentence as two independent source contexts when encoding at the source side, we use a context gate g_i for integrating two independent contexts of the source: original context c_i and compressed context b_i . The gate g_i is calculated by:

$$g_i = \sigma(\text{MLP}([c_i; b_i])). \quad (13)$$

Therefore, the final target fusion context c'_i is:

$$c'_i = g_i \otimes c_i + (1 - g_i) \otimes b_i, \quad (14)$$

where σ is the logistic sigmoid function, \otimes is the point-wise multiplication, and $[\cdot]$ represent the concatenation operation.

The context c'_i is input to replace the c_i the Eq. (8) to compute the probabilities of next target word.

3.4 Backbone Both-side Fusion based NMT

In the backbone both-side fusion based NMT (**BBFNMT**) model, we combine **BSFNMT** and **BTFNMT**. Both the original representation H_x and its compressed enhanced representation H'_x are as the input to the decoder. Similarly, we introduce an additional “encoder-decoder attention” module into the decoder to learn the compressed sequence enhanced context b'_i at the current time-step i :

$$b'_i = \text{FFN}(\text{MultiHead}(H_{tgt}, H'_x, H'_x)). \quad (15)$$

Then, the context gate g_i consistent with **BTFNMT** is applied to combine the two context information c_i and b'_i .

4 Experiments

4.1 Setup

Sentence Compression To evaluate the quality of our sentence compression model, we used the Annotated Gigaword corpus (Napoles, Gormley, and Van Durme 2012) as the benchmark (Rush, Chopra, and Weston 2015b). The data includes approximately 3.8 M training samples, 400 K validation samples, and 2 K test samples. The byte pair encoding (BPE) algorithm (Sennrich, Haddow, and Birch 2016) was adopted for subword segmentation, and the vocabulary size was set at 40 K for our supervised, unsupervised and semi-supervised settings (Zhang et al. 2019).

Baseline systems include AllText and F8W (Rush, Chopra, and Weston 2015b; Wang and Lee 2018). F8W is simply the first 8 words of the input, and AllText uses the whole text as the compression output. The F_1 score of ROUGE-1 (R-1), ROUGE-2 (R-2), and ROUGE-L (R-L) was used to evaluate this task (Lin 2004). We use beam search with a beam size of 5, the length length normalization of 0.5, and the coverage penalty of 0.2.

For the semi-supervised setting, in order to make the results comparable to Song et al. (2019), we used the same 190M English monolingual unpaired data from WMT News Crawl datasets for pre-training (unsupervised training). We included the other pretraining methods: masked language modeling (MLM, BERT) (Devlin et al. 2018), denoising auto-encoder (DAE) (Vincent et al. 2008), and masked sequence to sequence (MASS) (Song et al. 2019) to compare with our unsupervised pretraining method in the semi-supervised setting.

Machine Translation The proposed NMT model was evaluated on the WMT14 English-to-German (EN-DE) and English-to-French (EN-FR) tasks, which are both standard large-scale corpora for NMT evaluation. For the EN-DE translation task, 4.43 M bilingual sentence pairs from the WMT14 dataset were used as training data, including Common Crawl, News Commentary, and Europarl v7. The newstest2013 and newstest2014 datasets were used as the dev set and test set, respectively. For the EN-FR translation task, 36 M bilingual sentence pairs from the WMT14 dataset were used as training data. Newstest12 and newstest13 were combined for validation and the newstest14 was the

test set, following the setting of Gehring et al. (2017). The BPE algorithm (Sennrich, Haddow, and Birch 2016) was also adopted, and the joint vocabulary size was set at 40 K. For the hyper-parameters of our Transformer (base/large) models, we followed the settings used in Vaswani et al. (2017)’s work.

In addition, we also reported the state-of-the-art results in recent literatures, including modelling local dependencies (**Localness**) (Yang et al. 2018), fusing multiple-layer representations in SANs (**Context-Aware**) (Yang et al. 2019b), and fusing all global context representations in SANs (**global-deep context**) (Dou et al. 2018b). MultiBLEU was used to evaluate the translation task.

4.2 Main Results

Model	R-1	R-2	R-L
<i>Baselines:</i>			
All text	28.91	10.22	25.08
F8W	26.90	9.65	25.19
<i>Unsupervised:</i>			
(Fevry and Phang 2018)	28.42	7.82	24.95
ESC (This work)	31.37	8.25	28.01
<i>Supervised:</i>			
RNN-based Seq2seq	35.50	15.54	32.45
(Nallapati et al. 2016)	34.97	17.17	32.70
ESC (This work)	37.53	18.48	34.79
<i>Semi-supervised:</i>			
MLM Pretraining	37.75	18.45	34.85
DAE Pretraining	35.97	17.17	33.14
(Song et al. 2019)	38.73	19.71	35.96
ESC (This work)	39.54	20.35	36.79

Table 2: Performance on the sentence compression task

Sentence Compression To evaluate the quality of our sentence compression model, we conducted a horizontal comparison between the proposed sentence compression model and other sentence compression models in different settings. Table 2 shows the comparison results. We observed that the proposed unsupervised ESC model performed substantially better than Fevry and Phang (2018)’s unsupervised method. The proposed supervised ESC model also substantially outperformed the RNN-based Seq2seq and Nallapati et al. (2016)’s baseline method. That is, our supervised model gave +2.0 improvements on R-1, R-2, and R-L scores over the RNN-based Seq2seq. This means that the proposed Transformer-based approaches can generate compressed sentences of high quality.

We further compared our semi-supervised model with the semi-supervised pretraining methods of MLM (Devlin et al. 2018), DAE (Vincent et al. 2008), and MASS (Song et al. 2019). Our unsupervised pretraining method outperformed the other unsupervised pretraining ones on the sentence compression task consistently.

Machine Translation According to the results in Table 2, we chose the **semi-supervised ESC model** (which

System	EN-DE	#Speed	#Params	EN-FR	#Speed	#Params
<i>Existing NMT systems</i>						
Transformer (base) (Vaswani et al. 2017)	27.3	N/A	65.0M	38.1	N/A	N/A
+Localness (Yang et al. 2018)	28.11	N/A	88.8M	N/A	N/A	N/A
+Context-Aware SANs (Yang et al. 2019b)	28.26	N/A	194.9M	N/A	N/A	N/A
+global-deep context (Dou et al. 2018b)	28.58	N/A	111M	N/A	N/A	N/A
Transformer (big) (Vaswani et al. 2017)	28.4	N/A	213.0M	41.0	N/A	N/A
+Localness (Yang et al. 2018)	28.89	N/A	267.4M	N/A	N/A	N/A
+Context-Aware SANs (Yang et al. 2019b)	28.89	N/A	339.6M	N/A	N/A	N/A
+global-deep context (Dou et al. 2018b)	29.21	N/A	396M	N/A	N/A	N/A
<i>Our NMT systems</i>						
Transformer (base)	27.24	131k	66.5M	38.21	130k	85.7M
BSFNMT	27.75++	121k	72.1M	39.09++	120k	89.0M
BTFNMT	28.14+	120k	72.7M	39.22++	119k	89.8M
BBFNMT	28.35++	119k	78.6M	39.40++	116k	91.4M
Transformer (big)	28.23	11k	221.0M	41.15	11k	222.3M
BSFNMT	28.52+	10k	225.2M	41.92+	9k	227.1M
BTFNMT	29.16++	9k	225.7M	42.22++	8k	227.5M
BBFNMT	29.37++	8k	228.9M	42.52++	8k	230.3M

Table 3: Comparison with existing NMT systems on WMT14 EN-DE and EN-FR Translation Tasks. “++/+” after the BLEU score indicate that the proposed method was significantly better than the corresponding baseline Transformer (base or big) at significance level $p < 0.01/0.05$. “#Speed” denotes the decoding speed measured in target tokens per second.

performed the best) to generate compressed sentences for the machine translation task. The main results on the WMT14 EN-DE and EN-FR translation tasks are shown in Table 3. In the EN-DE task, we made the following observations:

1) The baseline Transformer (base) in this work achieved a performance comparable to the original Transformer (base) (Vaswani et al. 2017). This indicates that it is a strong baseline NMT system.

2) All **BSFNMT**, **BTFNMT**, and **BBFNMT** significantly outperformed the baseline Transformer (base/big). This indicates that the learned compressed backbone information was beneficial for the Transformer translation system.

3) Among the proposed three methods, **BTFNMT** performed better than **BSFNMT**. This indicates that the backbone fusion at the target-side is better than at the source-side. In addition, **BBFNMT** (base/big) outperformed the comparison systems +Localness and +Context-Aware SANs. This indicates that the compression knowledge as an additional context can enhance NMT better.

4) **BBFNMT** (based) is comparable to the +global-deep context, the best comparison system, while **BBFNMT** (big) slightly outperformed +global-deep context by 0.16 BLEU scores. In particular, the parameters of **BBFNMT** (base/big) model, which just increased 12.1/7.9M over the Transformer (base/big), were only 70% of the +global-deep context model. This denotes that the **BBFNMT** model is more efficient than the +global-deep context model. In addition, the training speed of the proposed models slightly decreased (8%), compared to the corresponding baselines.

5) The proposed **BBFNMT** (base) slightly outperformed the Transformer (big) which contains much more parameters than **BBFNMT** (base). This indicates that our improvement is not likely to be due to the increased number of parameters.

For the EN-FR translation task, the proposed models gave similar improvements over the baseline systems and comparing methods (except that the Transformer (big) performed much more better than Transformer (base)). These results show that our method is robust for improving the translation of other language pairs.

4.3 Ablation Study

Evaluating Sentence Compression To demonstrate the effectiveness of sentence compression, we compared the compressed sentences ($\gamma = 0.6$) generated in the Transformer translation system (**BBFNMT**) under different settings: AllText, F8W, RandSample (random sampling), supervised ESC, Unsupervised ESC and semi-supervised ESC. Table 4 shows the results on newstest2014 for the EN-DE translation task.

Model	BLEU on EN-DE
Baseline	27.24
+AllText	27.24
+F8W	27.40
+RandSample	26.53
+Supervised ESC	27.80
+Unsupervised ESC	27.97
+Semi-supervised ESC	28.35

Table 4: The effect of our ESC methods.

We made the following observations: 1) Simply introducing AllText and F8W achieved few improvement, and RandSample is lower than the baseline. In comparison, all the +supervised ESC, +unsupervised ESC, and +semi-supervised ESC models substantially improved the performance over the baseline Transformer (base). This means that

our ESC method provides a richer source information for machine translation tasks.

2) +Unsupervised ESC can gain better improvements over the +supervised ESC although supervised ESC model can achieve higher quality than the unsupervised ESC model in the benchmark test dataset. This may be due to that the annotated sentence compression training data is in different domain with the WMT EN-DE training data. Meanwhile, +Semi-supervised ESC with annotated data fine-tuning outperformed both +Unsupervised and +supervised ESC.

Effect of Encoder Parameters In our model, representations of the original sentence and its compressed version were learned by a shared encoder. To explore the effect of the encoder parameters, we also designed a **BBFNMT** with two independent encoders to learn representations of the original sentence and its compressed version, respectively. Table 5 shows results on the newstest2014 test set for the WMT14 EN-DE translation task.

Model	BLEU	#Params
Transformer (base)	27.24	66.4M
BBFNMT w/ Shared encoder	28.35	78.6M
BBFNMT w/ Independent encoders	28.50	91.6M

Table 5: The effect of encoder parameters.

The **BBFNMT** (w/ independent params) slightly outperformed the proposed shared encoder model by a BLEU score of 0.15, but its parameters increased by approximately 30%. In contrast, the parameters in our model are comparable to the baseline Transformer (base). Considering the parameter scale, we took a shared encoder to learn source representation, which makes it easy to verify the effectiveness of the additional translation knowledge, such as our backbone knowledge.

Evaluating Compression Ratio In order to verify the impact of different compression ratios on translation quality, we conducted experiments on EN-DE translation task with semi-supervised sentence compression in **BBFNMT** model.




Figure 3: Performances on EN-DE newstest2014 with different sentence compression ratios.

We controlled the compression ratio γ from 0 to 1.0. Consider two boundary conditions, when the compression ratio $\gamma = 0$, it means no compression sequence generated, which is the same as the vanilla Transformer. When the compression ratio $\gamma = 1.0$, it is equivalent to re-paraphrasing the source sentence using the sentence compression model (maintaining the same length) as the additional input for **BBFNMT**.

The experimental results are shown in Fig. 3. As can be seen from the results, in our experiments, sentence compression (re-paraphrasing) can bring performance improvement, even when the compression ratio $\gamma = 1.0$ and the sentence length is not shortened, re-paraphrasing can still bring slight improvement of translation quality.

5 Related Work

To let the translation have more focus over the source sentence information, efforts have been initiated on exploiting sentence segmentation, sentence simplification, and sentence compression for machine translation. Mellebeek et al. (2006) described an early method of skeleton-based translation that breaks down input sentences into grammatically meaningful chunks. The central part of the sentence is identified and remains unaltered, while the rest of the sentence is simplified. Sudoh et al. (2010) described the “divide and translate” approach to deal with complex input sentences. They parsed the input sentences, replaced the clauses with placeholders, and replaced them with the separately translated clauses. Xiao et al. (2014; 2016) built on the work by Mellebeek et al. (2006) but proposed a simpler approach based on source skeletons used in a single decoding step. Pouget-Abadie et al. (2014) experimented with automatically segmenting the source sentence to overcome problems with overly long sentences. Hasler et al. (2017) showed that the spaces of original and simplified translations can be effectively combined using translation lattices.

Different from these work, our proposed sentence compression model does not rely on any known linguistics motivated (such as syntax) skeleton simplification, but directly trains a computation motivated sentence compression model to learn to compress sentences and re-paraphrase them directly in seq2seq model. Though with a pure computation source, our sentence compression model can surprisingly generate more grammatically correct and refined sentences, and the words in the compressed sentence do not have to be the same as the original sentence. In the meantime, our sentence compression model can stably give source backbone representation exempt from unstable performance of a syntactic parser which is essential for syntactic skeleton simplification. Our sentence compression model can perform unsupervised training on large-scale data sets, and then use the supervised data for finetune, which is more promising from the results.

6 Conclusion

To give a more focused source representation, this paper makes the first attempt to propose an explicit

sentence compression method to enhance state-of-the-art Transformer-based NMT. The experimental results on WMT14 EN-DE and EN-FR translation tasks show that our proposed NMT model can yield significantly improved results over strong baseline translation systems.

References

- Artetxe, M.; Labaka, G.; Agirre, E.; and Cho, K. 2018. Unsupervised neural machine translation. In *ICLR*.
- Bojar, O.; Federmann, C.; Fishel, M.; Graham, Y.; Haddow, B.; Koehn, P.; and Monz, C. 2018. Findings of the 2018 conference on machine translation (WMT18). In *WMT*.
- Che, W.; Zhao, Y.; Guo, H.; Su, Z.; and Liu, T. 2015. Sentence compression for aspect-based sentiment analysis. *TASLP*.
- Cheng, J., and Lapata, M. 2016. Neural summarization by extracting sentences and words. In *ACL*.
- Chopra, S.; Auli, M.; and Rush, A. M. 2016. Abstractive sentence summarization with attentive recurrent neural networks. In *NAACL:HLT*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dou, Z.-Y.; Tu, Z.; Wang, X.; Shi, S.; and Zhang, T. 2018a. Exploiting deep representations for neural machine translation. In *EMNLP*.
- Dou, Z.-Y.; Tu, Z.; Wang, X.; Shi, S.; and Zhang, T. 2018b. Exploiting deep representations for neural machine translation. In *EMNLP*.
- Duan, X.; Yin, M.; Zhang, M.; Chen, B.; and Luo, W. 2019. Zero-shot cross-lingual abstractive sentence summarization through teaching generation and attention. In *ACL*.
- Fan, A.; Grangier, D.; and Auli, M. 2018. Controllable abstractive summarization. In *WNMT*.
- Fevry, T., and Phang, J. 2018. Unsupervised sentence compression using denoising auto-encoders. In *CoNLL*.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. N. 2017. Convolutional sequence to sequence learning. In *ICML*.
- Hasler, E.; de Gispert, A.; Stahlberg, F.; Waite, A.; and Byrne, B. 2017. Source sentence simplification for statistical machine translation. *CSL*.
- Hu, B.; Chen, Q.; and Zhu, F. 2015. LCSTS: A large scale Chinese short text summarization dataset. In *EMNLP*.
- Kikuchi, Y.; Neubig, G.; Sasano, R.; Takamura, H.; and Okumura, M. 2016. Controlling output length in neural encoder-decoders. In *EMNLP*.
- Knight, K., and Marcu, D. 2002. Summarization beyond sentence extraction: A probabilistic approach to sentence compression. *AI*.
- Lample, G.; Conneau, A.; Denoyer, L.; and Ranzato, M. 2018a. Unsupervised machine translation using monolingual corpora only.
- Lample, G.; Ott, M.; Conneau, A.; Denoyer, L.; and Ranzato, M. 2018b. Phrase-based & neural unsupervised machine translation. In *EMNLP*.
- Lin, C.-Y. 2004. ROUGE: A package for automatic evaluation of summaries. In *ACL*.
- Mellebeek, B.; Owczarzak, K.; Groves, D.; Van Genabith, J.; and Way, A. 2006. A syntactic skeleton for statistical machine translation.
- Nallapati, R.; Zhou, B.; dos Santos, C.; Gulcehre, C.; and Xiang, B. 2016. Abstractive text summarization using sequence-to-sequence rns and beyond. In *CoNLL*.
- Napoles, C.; Gormley, M.; and Van Durme, B. 2012. Annotated Gigaword. In *AKBC-WEKEX*.
- Pouget-Abadie, J.; Bahdanau, D.; van Merriënboer, B.; Cho, K.; and Bengio, Y. 2014. Overcoming the curse of sentence length for neural machine translation using automatic segmentation. In *SSST-8*.
- Rush, A. M.; Chopra, S.; and Weston, J. 2015a. A neural attention model for abstractive sentence summarization. In *EMNLP*.
- Rush, A. M.; Chopra, S.; and Weston, J. 2015b. A neural attention model for abstractive sentence summarization. In *EMNLP*.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Neural machine translation of rare words with subword units. In *ACL*.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2019. Mass: Masked sequence to sequence pre-training for language generation. In *ICML*.
- Sudoh, K.; Duh, K.; Tsukada, H.; Hirao, T.; and Nagata, M. 2010. Divide and translate: improving long distance reordering in statistical machine translation. In *WMT*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS*.
- Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P.-A. 2008. Extracting and composing robust features with denoising autoencoders. In *ICML*.
- Wang, Y., and Lee, H.-y. 2018. Learning to encode text as human-readable summaries using generative adversarial networks. In *EMNLP*.
- Wang, X.; Tu, Z.; Wang, L.; and Shi, S. 2019. Exploiting sentential context for neural machine translation. In *ACL*.
- 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 preprint arXiv:1609.08144*.
- Xiao, T.; Zhu, J.; Zhang, C.; and Liu, T. 2016. Syntactic skeleton-based translation. In *AAAI*.
- Xiao, T.; Zhu, J.; and Zhang, C. 2014. A hybrid approach to skeleton-based translation. In *ACL*.
- 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. 2019a. Context-aware self-attention networks.
- Yang, B.; Li, J.; Wong, D. F.; Chao, L. S.; Wang, X.; and Tu, Z. 2019b. Context-aware self-attention networks. *CoRR* abs/1902.05766.
- Zhang, Z.; Zhao, H.; Ling, K.; Li, J.; Li, Z.; He, S.; and Fu, G. 2019. Effective subword segmentation for text comprehension. In *TASLP*.