

Synchronous bidirectional inference for neural sequence generation



Jiajun Zhang^{a,b,*}, Long Zhou^{a,b}, Yang Zhao^{a,b}, Chengqing Zong^{a,b,c}

^a National Laboratory of Pattern Recognition, CASIA, Beijing, China

^b University of Chinese Academy of Sciences, Beijing, China

^c CAS Center for Excellence in Brain Science and Intelligence Technology, Beijing, China

ARTICLE INFO

Article history:

Received 20 March 2019

Received in revised form 25 October 2019

Accepted 30 December 2019

Available online 8 January 2020

Keywords:

Sequence to sequence learning

Bidirectional inference

Beam search

Machine translation

Summarization

ABSTRACT

In sequence to sequence generation tasks (e.g. machine translation and abstractive summarization), inference is generally performed in a left-to-right manner to produce the result token by token. The neural approaches, such as LSTM and self-attention networks, are now able to make full use of all the predicted history hypotheses from left side during inference, but cannot meanwhile access any future (right side) information and usually generate unbalanced outputs (e.g. left parts are much more accurate than right ones in Chinese-English translation). In this work, we propose a synchronous bidirectional inference model to generate outputs using both left-to-right and right-to-left decoding simultaneously and interactively. First, we introduce a novel beam search algorithm that facilitates synchronous bidirectional decoding. Then, we present the core approach which enables left-to-right and right-to-left decoding to interact with each other, so as to utilize both the history and future predictions simultaneously during inference. We apply the proposed model to both LSTM and self-attention networks. Furthermore, we propose a novel fine-tuning based parameter optimization algorithm in addition to the simple two-pass strategy. The extensive experiments on machine translation and abstractive summarization demonstrate that our synchronous bidirectional inference model can achieve remarkable improvements over the strong baselines.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Many tasks in natural language processing, such as machine translation, abstractive summarization and chatbot, can be formalized as a sequence to sequence (seq2seq) generation problem which takes a sequence as input (e.g. source language sentence) and produces another sequence as output (e.g. target language translation). Generally, the seq2seq framework performs inference in a left-to-right (L2R) manner and predicts the current output token conditioned on previously generated tokens. Existing methods mainly focus on how to fully exploit the already predicted outputs on the left. And the recently proposed neural architectures for sequence generation including recurrent networks [1,2], convolutional networks [3] and self-attention ones (known as Transformer) [4] facilitate the exploration of all the history information during inference.

However, conventional seq2seq models cannot access the future predictions on the right and usually produce unbalanced outputs. Take Chinese-to-English translation as an example. Left parts are much more accurate than right ones when ap-

* Corresponding author at: National Laboratory of Pattern Recognition, CASIA, Beijing, China.

E-mail addresses: jjzhang@nlpr.ia.ac.cn (J. Zhang), long.zhou@nlpr.ia.ac.cn (L. Zhou), yang.zhao@nlpr.ia.ac.cn (Y. Zhao), cqzong@nlpr.ia.ac.cn (C. Zong).

Table 1

Matching accuracy of the first and last four tokens between model predictions and references in NIST Chinese-English machine translation tasks. L2R denotes conventional left-to-right inference while R2L denotes right-to-left inference.

Architecture	Direction	First four	Last four
LSTM	L2R	36.35%	31.64%
	R2L	31.22%	34.01%
Transformer (self-attention network)	L2R	40.21%	35.10%
	R2L	35.67%	39.47%




Fig. 1. A simple illustration of our synchronous bidirectional inference model in which L2R and R2L models interact with each other. The left part is the input sequence and the right part denotes the scheme of synchronous bidirectional inference. Take generating \vec{y}_2 as an example. We cannot only utilize the past predictions (\vec{y}_0, \vec{y}_1) of the L2R inference, but also could leverage the future contexts (\vec{y}_0, \vec{y}_1) which have been already predicted by the R2L inference. It is similar when predicting \vec{y}_2 . Note that \vec{y}_{n-1} is not necessary the same as \vec{y}_0 and n may be different from n' . The final output sequence will be $\vec{y}_0 \vec{y}_1 \vec{y}_2 \dots \vec{y}_{n-1}$ if L2R inference wins. It will be $\vec{y}_0 \vec{y}_1 \vec{y}_2 \dots \vec{y}_{n'-1}$ otherwise.

plying left-to-right inference. The phenomenon is similar for right-to-left (R2L) inference where the right parts are much better.¹ In order to have a more intuitive understanding, we have investigated both L2R and R2L inferences using LSTM [6] and self-attention networks (SAN) [4] on the typical sequence generation task, namely machine translation. Table 1 shows the matching accuracy of the first and last four tokens between model predictions and references for Chinese-to-English translation. It is obvious that left-to-right inference performs much better on predictions of head tokens while right-to-left inference excels in tail token predictions. Intuitively, it is a promising direction to combine the merits of bidirectional inferences and make full use of both history and future contexts.

Researchers have made great efforts to take advantages of both L2R and R2L inferences. [7,8] enforce the agreement between L2R and R2L predictions during training, and then L2R inference will be improved accordingly. [9,10] employ R2L model to rerank the n -best hypotheses of the L2R model. [11] first obtains the R2L outputs and optimizes the L2R inference model based on both of the original input and the R2L outputs. Specifically, [11] aims to learn a better L2R inference model. They first train a R2L model with the bilingual training data. Then, the optimized R2L decoder translates the input of each parallel training instance and generates the outputs (hidden states) which serve as the additional context for L2R prediction when optimizing the L2R inference model. Despite the performance improvement, these approaches suffer from two issues. On one hand, they have to train two separate seq2seq models for L2R and R2L inferences respectively. And the two-pass decoding strategy makes the decoding efficiency much lowered down. On the other hand, the two models cannot interact with each other during inference, which limits the potential of performance improvement.

In this article, we generalize the previous work [12] and propose a synchronous bidirectional inference model that produces outputs using both L2R and R2L decoding simultaneously and interactively. This work has three new contributions. First, we generalize the decoding model into synchronous bidirectional inference method for the general sequence-to-sequence architectures (e.g. LSTM and Transformer) and general sequence generation tasks (e.g. translation and summarization). Second, in addition to the two-pass training method, we further propose and investigate a new fine-tuning based algorithm for efficient parameter optimization. Third, we conduct deep analysis about the characteristics of the proposed synchronous bidirectional inference model.

Specifically, we first introduce a novel beam search algorithm to accommodate L2R and R2L inferences at the same time. At each timestep during inference, each half beam retain the hypotheses from L2R and R2L inferences respectively and each hypothesis is generated by utilizing already predicted outputs from both directions. The interaction between L2R and R2L inferences is realized through a synchronous attention model that attempts to leverage both the history and future sequential predictions simultaneously during inference. Fig. 1 gives a simple illustration of the proposed synchronous bidirectional inference model. The middle part dotted arrows on the right of Fig. 1 is the core of our model. L2R and R2L inferences interact with each other in an implicit way illustrated by the dotted arrows. The arrows indicate the information passing flow. Solid arrows show the conventional history context dependence while dotted arrows

¹ It should be noted that the unbalance phenomenon is different for different language pairs. For example, [5] reported that right parts are much better translated than left parts when performing left-to-right decoding for Japanese-English translation.

introduce the future context dependence on the other inference direction. For example, besides the past predictions (\vec{y}_0, \vec{y}_1) , L2R inference can also utilize the future contexts $(\overleftarrow{y}_0, \overleftarrow{y}_1)$ generated by the R2L inference when predicting \vec{y}_2 .

As we mentioned above, there are many sequence to sequence models. To test the generalization capacity of our model, we apply the proposed synchronous bidirectional inference model into two representative seq2seq frameworks using LSTM and self-attention networks. Furthermore, we propose a novel fine-tuning based optimization algorithm for parameter training, in addition to a relatively simple two-pass strategy. We choose machine translation and abstractive sentence summarization as the testbed to verify the effectiveness of the models. The extensive experiments demonstrate that our proposed model remarkably outperforms the strong baselines.

2. Synchronous bidirectional inference

The task of sequence to sequence learning is to find the most probable output sequence $y = y_0 y_1 \cdots y_{n-1}$ which maximizes the following conditional probability given the input sequence $x = x_0 x_1 \cdots x_{m-1}$.

$$P(y|x) = P(y_0 y_1 \cdots y_{n-1} | x_0 x_1 \cdots x_{m-1}) \quad (1)$$

Unlike sequential labeling tasks in which y shares the same length as x , the output length of y is unknown until the inference process ends and in most cases the token numbers are different between output and input sequences in seq2seq learning tasks. For instance, the result summary (output sequence) should be much shorter than the original text (input sequence) in the summarization task.

2.1. Unidirectional inference

Conventionally, Equation (1) is decomposed in a left-to-right manner as follows:

$$P(y|x) = \prod_{i=0}^{n-1} p(y_i | y_0 \cdots y_{i-1}, x) \quad (2)$$

Since the search space V ($y_i \in V$) is very large and contains tens of thousands of entries in most cases, a beam search algorithm is usually employed to approximately find the most probable output sequence according to the history predictions $y_0 \cdots y_{i-1}$ and the input sequence x . Currently, neural methods such as LSTM, convolutional networks and self-attention ones can model the conditional probability $p(y_i | y_0 \cdots y_{i-1}, x)$ more and more accurately, but leaves the future contexts unexplored.

In order to leverage the right hand information, Equation (1) can also be decomposed in a right-to-left manner as follows:

$$P(y|x) = \prod_{i=n-1}^0 p(y_i | y_{i+1} \cdots y_{n-1}, x) \quad (3)$$

Using this decomposition, the right-side hypotheses $y_{i+1} \cdots y_{n-1}$ are available when predicting y_i , while the left-side predictions are still missing.

2.2. Synchronous bidirectional beam search

Ideally, we expect to utilize both the past and future contexts ($y_0 \cdots y_{i-1}$ and $y_{i+1} \cdots y_{n-1}$) when determining the best prediction of y_i . However, it is contradictory to some extent. Predicting y_i needs y_{i+1} on the right, while determining y_{i+1} requires y_i on the left. Obviously, it is impractical to use the whole contexts of both sides ($y_0 \cdots y_{i-1}$ and $y_{i+1} \cdots y_{n-1}$) in a single inference model. We take a step back and attempt to explore the bidirectional contexts as many as possible if not all.

We propose a synchronous bidirectional inference model in which left-to-right and right-to-left inferences perform in parallel while keeping interaction with each other. In this way, Equation (1) is decomposed as follows:

$$P(y|x) = \begin{cases} \prod_{i=0}^{n-1} p(\vec{y}_i | \vec{y}_0 \cdots \vec{y}_{i-1}, x, \overleftarrow{y}_0 \cdots \overleftarrow{y}_{i-1}) & \text{if L2R} \\ \prod_{i=0}^{n-1} p(\overleftarrow{y}_i | \overleftarrow{y}_0 \cdots \overleftarrow{y}_{i-1}, x, \vec{y}_0 \cdots \vec{y}_{i-1}) & \text{if R2L} \end{cases} \quad (4)$$

Equation (4) says that the bidirectional inference model accommodates L2R and R2L decoding at the same time. At timestep i , we have already generated the left $i-1$ hypotheses $\vec{y}_0 \cdots \vec{y}_{i-1}$ with L2R inference and the right $i-1$ predictions $\overleftarrow{y}_0 \cdots \overleftarrow{y}_{i-1}$ with R2L inference. Thus, different from Equation (2) and Equation (3), both-side predictions can be utilized as contexts in the above bidirectional composition.

Input: Input sequence $x = \langle s \rangle x_0 \dots x_j \dots x_{m-1} \langle /s \rangle$, beam size K , maximum length of output sequence $MaxLen$
Output: Optimal output sequence $y = y_0 \dots y_i \dots y_{n-1}$

- 1 Initialize complete hypothesis list $B = \Phi$, left-to-right temporary hypothesis list $B_{l-tmp} = \Phi$ and partial hypothesis list $B_{l-part}^0 = \{\langle l2r \rangle\}$, right-to-left temporary hypothesis list $B_{r-tmp} = \Phi$ and partial hypothesis list $B_{r-part}^0 = \{\langle r2l \rangle\}$; $\langle l2r \rangle$ and $\langle r2l \rangle$ are tags indicating inference direction;
- 2 **for** $i = 1; i \leq MaxLen$ **do**
- 3 $B_{l-tmp} = \text{ExpandHypo}(B_{l-part}^{i-1}, B_{r-part}^{i-1})$;
- 4 $B_{r-tmp} = \text{ExpandHypo}(B_{r-part}^{i-1}, B_{l-part}^{i-1})$;
- 5 $[B_{l-part}^i, B] = \text{UpdateHypo}(B_{l-tmp}, B_{l-part}^{i-1}, B)$;
- 6 $[B_{r-part}^i, B] = \text{UpdateHypo}(B_{r-tmp}, B_{r-part}^{i-1}, B)$;
- 7 $B_{l-tmp} = \Phi$;
- 8 $B_{r-tmp} = \Phi$;
- 9 **end**
- 10 **if** $B \neq \Phi$ **then**
- 11 sort B in descending order;
- 12 $y = B[0]$;
- 13 **else**
- 14 **if** $p(B_{l-part}^{MaxLen}[0]) \geq p(B_{r-part}^{MaxLen}[0])$ **then**
- 15 $y = B_{l-part}^{MaxLen}[0]$;
- 16 **else**
- 17 $y = B_{r-part}^{MaxLen}[0]$;
- 18 **end**
- 19 **end**
- 20 **if** $y[0] = \langle l2r \rangle$ **then**
- 21 return y ;
- 22 **else**
- 23 return reversed y ;
- 24 **end**

Algorithm 1: Synchronous Bidirectional Beam Search Algorithm.

Input: Partial hypothesis list B_{l-part}^{i-1} for current decoding direction and B_{r-part}^{i-1} for the opposite decoding direction
Output: temporary hypothesis list $B_{tmp} = \Phi$

- 1 **for** $cand$ in B_{l-part}^{i-1} **do**
- 2 **for** y_i^* in $TargetVocab$ **do**
- 3 $cand = cand + y_i^*$;
- 4 $p(cand) = \text{SBInfer}(cand, B_{l-part}^{i-1}, B_{r-part}^{i-1})$;
- 5 $B_{tmp} = B_{tmp} \cup \{cand\}$;
- 6 **end**
- 7 **end**
- 8 sort B_{tmp} in a descending order;
- 9 $B_{tmp} = B_{tmp}[0 : K/2]$;
- 10 return B_{tmp} ;

Algorithm 2: $\text{ExpandHypo}(B_{l-part}^{i-1}, B_{r-part}^{i-1})$.

Algorithm 1 shows the beam search procedure of the synchronous bidirectional inference model. The working flow is similar to the unidirectional beam search. We keep three kinds of lists. B is employed to store complete hypotheses. (B_{l-tmp}, B_{r-tmp}) and (B_{l-part}, B_{r-part}) are used to maintain the temporary and partial hypotheses at each decoding timestep for L2R and R2L inferences respectively. Lines 2-9 in **Algorithm 1** is the main part of the beam search algorithm. At timestep i , L2R and R2L inferences perform in parallel but interactively to expand the partial hypotheses B_{l-part}^{i-1} and B_{r-part}^{i-1} from the previous timestep (lines 3-4). Then, the complete hypothesis list B and the partial hypothesis list $(B_{l-part}^i, B_{r-part}^i)$ will be updated according to the temporary hypothesis list $(B_{l-tmp}^i, B_{r-tmp}^i)$ (lines 5-6).

Algorithm 2 and **Algorithm 3** respectively detail the hypothesis expansion procedure and hypothesis update process. **Algorithm 3** is trivial and is the same as the conventional unidirectional beam search. The algorithm $\text{ExpandHypo}(B_{l-part}^{i-1}, B_{r-part}^{i-1})$ is the key for synchronous bidirectional inference (**SBInfer**). In **Algorithm 2**, for a partial hypothesis in B_{l-part}^{i-1} , we calculate the probability of each candidate token in the target vocabulary by utilizing both of the history context B_{l-part}^{i-1} and the future information B_{r-part}^{i-1} through the function $\text{SBInfer}(cand, B_{l-part}^{i-1}, B_{r-part}^{i-1})$.

Obviously, the function $\text{SBInfer}(cand, B_{l-part}^{i-1}, B_{r-part}^{i-1})$ is the most important part and requires specific design for different seq2seq architectures. Next, we introduce how to implement the function $\text{SBInfer}(cand, B_{l-part}^{i-1}, B_{r-part}^{i-1})$ for both LSTM-based and self-attention based seq2seq networks.

```

Input: Temporary hypothesis list  $B_{tmp} = \Phi$ , partial hypothesis list  $B_{part}$ , and complete hypothesis list  $B$ 
Output: Partial hypothesis list  $B_{part}$  and complete hypothesis list  $B$ 
1 for  $cand$  in  $B_{tmp}$  do
2   if  $\langle /s \rangle = cand[-1]$  then
3      $B = B \cup \{cand\}$ ;
4     if  $K \geq |B|$  then
5       break;
6     end
7   else
8      $B_{part} = B_{part} \cup \{cand\}$ ;
9   end
10 end
11 return  $[B_{part}, B]$ ;

```

Algorithm 3: UpdateHypo(B_{tmp}, B_{part}, B).

3. Synchronous bidirectional inference for LSTM-based Seq2Seq framework

No matter what kind of network architecture is used, most of Seq2Seq frameworks consist of an encoder and a decoder.² Given an input sequence $x = (x_0, x_2, \dots, x_{m-1})$, the encoder transforms x into a sequence of abstract context representations $C = (\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_{m-1})$ whose size is the same as the length of the input text. Then, from the context vectors C the decoder generates the output sequence $y = (\vec{y}_0, \vec{y}_1, \dots, \vec{y}_{n-1})$ one token each time by maximizing the probability of $p(\vec{y}_i | \vec{y}_{<i}, C)$ with a left-to-right inference model.

Hereafter, we leverage \mathbf{x}_j and \mathbf{y}_i to denote the word embeddings corresponding to the input and output tokens x_j and y_i . Next, we briefly review the encoder introducing how to obtain C and the decoder addressing how to calculate $p(\vec{y}_i | \vec{y}_{<i}, C)$ for the conventional LSTM-based Seq2Seq architecture. Then, we propose to enable synchronous bidirectional inference $SBInfer(cand, B_{l-part}^{i-1}, B_{r-part}^{i-1})$ in the LSTM-based architecture.

3.1. LSTM-based Seq2Seq framework

The **encoder** employs L stacked LSTM layers to learn the context vectors $C = (\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_{m-1})$. In the l -th layer ($l > 1$), \mathbf{h}_j^l is calculated as follows:

$$\mathbf{h}_j^l = \text{LSTM}(\mathbf{h}_{j-1}^l, \mathbf{h}_j^{l-1}) \quad (5)$$

In the first layer ($l = 1$), \mathbf{h}_j^1 is obtained through a bidirectional LSTM:

$$\vec{\mathbf{h}}_j^1 = \text{LSTM}(\vec{\mathbf{h}}_{j-1}^1, \mathbf{x}_j) \quad (6)$$

$$\overleftarrow{\mathbf{h}}_j^1 = \text{LSTM}(\overleftarrow{\mathbf{h}}_{j+1}^1, \mathbf{x}_j) \quad (7)$$

Given $\vec{\mathbf{h}}_j^1$ and $\overleftarrow{\mathbf{h}}_j^1$, \mathbf{h}_j^1 is calculated with a feed-forward neural network $\mathbf{h}_j^1 = \tanh(\mathbf{W}_h^l \cdot \vec{\mathbf{h}}_j^1 + \mathbf{W}_h^r \cdot \overleftarrow{\mathbf{h}}_j^1 + \mathbf{b}_h)$.

The **decoder** computes the conditional probability $p(\vec{y}_i | \vec{y}_{<i}, C)$ with the help of **attention** mechanism [2] that leverages different input context \mathbf{c}_i at different decoding time step:

$$p(\vec{y}_i | \vec{y}_{<i}, C) = p(\vec{y}_i | \vec{y}_{<i}, \mathbf{c}_i) = \text{softmax}(\mathbf{W} \vec{\mathbf{z}}_i) \quad (8)$$

where $\vec{\mathbf{z}}_i$ is the attention output:

$$\vec{\mathbf{z}}_i = \tanh(\mathbf{W}_c[\vec{\mathbf{z}}_i^l; \mathbf{c}_i]) \quad (9)$$

in which $\vec{\mathbf{z}}_i^l$ is the top hidden state of the decoder network and $\vec{\mathbf{z}}_i^l$ in the l -th layer is computed using the following formula:

$$\vec{\mathbf{z}}_i^l = \text{LSTM}(\vec{\mathbf{z}}_{i-1}^l, \vec{\mathbf{z}}_i^{l-1}) \quad (10)$$

If $l = 1$, $\vec{\mathbf{z}}_i^1$ will be calculated by combining $\vec{\mathbf{z}}_{i-1}$ as feed input [14]:

$$\vec{\mathbf{z}}_i^1 = \text{LSTM}(\vec{\mathbf{z}}_{i-1}^1, y_{i-1}, \vec{\mathbf{z}}_{i-1}) \quad (11)$$

² There is an exception that [13] proposed a new Seq2Seq model that utilizes a single 2D convolutional neural network across both sequences.




Fig. 2. Bidirectional inference model for Seq2Seq framework with LSTM architecture.

The dynamic context \mathbf{c}_i is the weighted sum of the source-side context vectors and is calculated by the attention model:

$$\mathbf{c}_i = \sum_{j=0}^{m-1} \alpha_{ij} \mathbf{h}_j \quad (12)$$

where α_{ij} is a normalized item calculated as follows:

$$e_{ij} = \mathbf{v}_a^T \tanh(\mathbf{W}_a \vec{\mathbf{z}}_i^L + \mathbf{U}_a \mathbf{h}_j) \quad (13)$$

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{j'} \exp(e_{ij'})} \quad (14)$$

The greater the value of the variable α_{ij} , the more contribution of the j -th input token to the generation of the i -th output word. The left part in Fig. 2 gives the overall illustration of this unidirectional inference model for LSTM-based Seq2Seq framework. Note that residual connections and layer normalizations are employed as well and they are neglected in the description for simplicity.

3.2. Synchronous bidirectional inference for LSTM-based architecture

In synchronous bidirectional inference, $p(\vec{\mathbf{y}}_i)$ is calculated with both history and future contexts according to Equation (4): $p(\vec{\mathbf{y}}_i | \vec{\mathbf{y}}_0 \cdots \vec{\mathbf{y}}_{i-1}, x, \vec{\mathbf{y}}_0 \cdots \vec{\mathbf{y}}_{i-1})$. The previous section introduces the way to use input x and history contexts $\vec{\mathbf{y}}_0 \cdots \vec{\mathbf{y}}_{i-1}$ in Equation (8): $p(\vec{\mathbf{y}}_i | \vec{\mathbf{y}}_{<i}, C)$. The synchronous bidirectional inference adopts the same mechanism as follows:

$$p(\vec{\mathbf{y}}_i | \vec{\mathbf{y}}_{<i}, C, \overleftarrow{\mathbf{y}}_{<i}) = p(\vec{\mathbf{y}}_i | \vec{\mathbf{y}}_{<i}, \mathbf{c}_i, \overleftarrow{\mathbf{y}}_{<i}) = \text{softmax}(\mathbf{W} \vec{\mathbf{z}}_i) \quad (15)$$

Different from unidirectional inference, the synchronous bidirectional inference calculates the attention output $\vec{\mathbf{z}}_i$ with both L2R and R2L predictions:

$$\vec{\mathbf{z}}_i = \tanh(\mathbf{W}_c [\vec{\mathbf{z}}_i^L; \mathbf{c}_i; \overleftarrow{\mathbf{z}}_i]) \quad (16)$$

where the future context $\overleftarrow{\mathbf{z}}_i$ is obtained using another attention model as illustrated with dotted arrows from right part to the left part in Fig. 2.

$$\overleftarrow{\mathbf{z}}_i = \sum_{k=0}^{i-1} \overleftarrow{\gamma}_{ik} \overleftarrow{\mathbf{z}}_k^L \quad (17)$$

in which $\overleftarrow{\gamma}_{ik}$ is a normalized coefficient:

$$e_{ik} = \mathbf{v}_z^T \tanh(\mathbf{W}_z \vec{\mathbf{z}}_i^L + \mathbf{U}_z \overleftarrow{\mathbf{z}}_k^L) \quad (18)$$

$$\overleftarrow{\gamma}_{ik} = \frac{\exp(e_{ik})}{\sum_{k'} \exp(e_{ik'})} \quad (19)$$

It should be noted that L2R and R2L inferences perform simultaneously in parallel. Thus, when calculating $p(\vec{\mathbf{y}}_i | \vec{\mathbf{y}}_{<i}, C, \overleftarrow{\mathbf{y}}_{<i})$, we can as well compute $p(\overleftarrow{\mathbf{y}}_i | \overleftarrow{\mathbf{y}}_{<i}, C, \vec{\mathbf{y}}_{<i})$ at the same time in a similar way.

$$p(\overleftarrow{y}_i | \overleftarrow{y}_{<i}, C, \overrightarrow{y}_{<i}) = p(\overleftarrow{y}_i | \overleftarrow{y}_{<i}, \mathbf{c}_i, \overrightarrow{y}_{<i}) = \text{softmax}(W \overleftarrow{\mathbf{z}}_i) \quad (20)$$

$$\overleftarrow{\mathbf{z}}_i = \tanh(W_c[\overleftarrow{\mathbf{z}}_i^L; \mathbf{c}_i; \overrightarrow{\mathbf{z}}_i]) \quad (21)$$

Where the left context $\overleftarrow{\mathbf{z}}_i$ is obtained using a similar attention model as illustrated with dotted arrows from the left part to the right part in Fig. 2.

$$\overrightarrow{\mathbf{z}}_i = \sum_{k=0}^{i-1} \overrightarrow{y}_{ik} \overrightarrow{\mathbf{z}}_k^L \quad (22)$$

4. Synchronous bidirectional inference for self-attention based framework

The self-attention based Seq2Seq framework is known as **Transformer** [4]. In this section, we first give an overview of Transformer and then propose the implementation of synchronous bidirectional inference $SBInfer(cand, B_{l-part}^{i-1}, B_{r-part}^{i-1})$ in Transformer.

4.1. Transformer

The Transformer also follows the encoder-decoder architecture. The encoder includes L identical layers and each layer is composed of two sub-layers: the self-attention sub-layer followed by the feed-forward sub-layer.

The decoder also consists of L identical layers. Each layer has three sub-layers. The first one is the masked self-attention mechanism. The second one is the decoder-encoder attention sub-layer and the third one is the feed-forward sub-layer.³

Obviously, the key component is the attention mechanism.⁴ The three kinds of attention mechanisms can be formalized into the same formula.

$$\text{Attention}(\mathbf{q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V} \quad (23)$$

where \mathbf{q} , \mathbf{K} and \mathbf{V} stand for a query, the key list and the value list respectively. d_k is the dimension of the key.

For the self-attention in encoder, the queries, keys and values are from the same layer. For example, if we calculate the output of the first layer in the encoder at the j -th position. The query is vector \mathbf{x}_j .⁵ The keys and values are the same and both are the embedding matrix $\mathbf{x} = [\mathbf{x}_0 \cdots \mathbf{x}_{m-1}]$. Using Equation (23) followed by a feed-forward network, we can get the representation of the second layer. After L layers, we obtain the input contexts $C = [\mathbf{h}_0, \cdots, \mathbf{h}_{m-1}]$.

The masked self-attention in decoder is similar to that of encoder except that the query at the i -th position can only attend to positions before i , since the predictions after i -th position are not available in the auto-regressive unidirectional inference.

$$\overrightarrow{\mathbf{z}}_i^{\text{past}} = \text{Attention}(\overrightarrow{\mathbf{q}}_i, \overrightarrow{\mathbf{K}}_{\leq i}, \overrightarrow{\mathbf{V}}_{\leq i}) = \text{softmax}\left(\frac{\overrightarrow{\mathbf{q}}_i \overrightarrow{\mathbf{K}}_{\leq i}^T}{\sqrt{d_k}}\right) \overrightarrow{\mathbf{V}}_{\leq i} \quad (24)$$

The decoder-encoder attention mechanism is the same as that of LSTM-based Seq2Seq architecture. The query is the output of the masked self-attention sub-layer $\overrightarrow{\mathbf{z}}_i^{\text{past}}$. The keys and values are the same encoder contexts C . The feed-forward sub-layer is then applied to yield the output of a whole layer. After L such layers, we obtain the final hidden state $\overrightarrow{\mathbf{z}}_i$. Softmax function (Equation (8)) is then employed to predict the output \overrightarrow{y}_i . Left part in Fig. 3 depicts the overall architecture of Transformer.

4.2. Synchronous bidirectional inference for transformer

In synchronous bidirectional inference, the essential difference lies in the improvement over the masked self-attention mechanism for decoder. In standard Transformer, the masked self-attention model calculates the output $\overrightarrow{\mathbf{z}}_i$ ($\overrightarrow{\mathbf{z}}_i^{\text{past}}$) using only the history contexts. In contrast, synchronous bidirectional inference performs L2R and R2L decoding in parallel and interactively. At the i -th timestep, L2R and R2L inferences have already generated $i-1$ outputs $\overrightarrow{\mathbf{z}}_{\leq i}^{\text{past}} = (\overrightarrow{\mathbf{z}}_0^{\text{past}} \cdots \overrightarrow{\mathbf{z}}_{i-1}^{\text{past}})$ and $\overleftarrow{\mathbf{z}}_{\leq i}^{\text{past}} = (\overleftarrow{\mathbf{z}}_0^{\text{past}} \cdots \overleftarrow{\mathbf{z}}_{i-1}^{\text{past}})$. Therefore, both $\overrightarrow{\mathbf{z}}_{\leq i}^{\text{past}}$ and $\overleftarrow{\mathbf{z}}_{\leq i}^{\text{past}}$ can be employed to compute $\overrightarrow{\mathbf{z}}_i$.

Accordingly, we design two self-attention mechanisms to handle history contexts $\overrightarrow{\mathbf{z}}_{\leq i}^{\text{past}}$ and future contexts $\overleftarrow{\mathbf{z}}_{\leq i}^{\text{past}}$ respectively. In addition to Equation (24) that utilizes history information, we propose another self-attention mechanism to leverage the future information generated by the opposite inference direction.

³ Residual connection and layer normalization are performed for each sub-layer in both encoder and decoder.

⁴ In fact, multi-head attention is employed and we just introduce basic attention for simplicity.

⁵ Suppose \mathbf{x}_j is the sum vector of input token embedding and the positional embedding.




Fig. 3. Bidirectional inference model for seq2seq framework with self-attention architecture.

$$\vec{z}_i^{future} = \text{Attention}(\vec{q}_i, \vec{K}_{\leq i}, \vec{V}_{\leq i}) = \text{softmax}\left(\frac{\vec{q}_i \vec{K}_{\leq i}^T}{\sqrt{d_k}}\right) \vec{V}_{\leq i} \quad (25)$$

where \vec{q}_i is the query (i.e. the embedding of the $(i-1)$ -th output \vec{y}_{i-1}) from the L2R decoder. $\vec{K}_{\leq i}$ and $\vec{V}_{\leq i}$ are keys and values (i.e. the embeddings of the previous i predictions $\vec{y}_0, \dots, \vec{y}_{i-1}$) from the R2L decoder. Fig. 3 illustrates how synchronous bidirectional inference performs. When producing \vec{z}_i , the solid lines denote the original masked self-attention with history while the dotted lines indicate the self-attention with future contexts.

Finally, we introduce a function to combine \vec{z}_i^{past} and \vec{z}_i^{future} to obtain a new representation \vec{z}_i that encodes both past and future contexts.

$$\vec{z}_i = f(\vec{z}_i^{past}, \vec{z}_i^{future}) = \vec{z}_i^{past} + \lambda \times \tanh(\vec{z}_i^{future}) \quad (26)$$

In which λ is a hyper-parameter to balance the contribution of the history and future contexts. For R2L decoding, \overleftarrow{z}_i can be calculated similarly in parallel.

$$\overleftarrow{z}_i^{future} = \text{Attention}(\overleftarrow{q}_i, \overleftarrow{K}_{\leq i}, \overleftarrow{V}_{\leq i}) = \text{softmax}\left(\frac{\overleftarrow{q}_i \overleftarrow{K}_{\leq i}^T}{\sqrt{d_k}}\right) \overleftarrow{V}_{\leq i} \quad (27)$$

$$\overleftarrow{z}_i = \overleftarrow{z}_i^{past} + \lambda \times \tanh(\overleftarrow{z}_i^{future}) \quad (28)$$

5. Training

Since our synchronous bidirectional inference performs L2R and R2L decoding in parallel, L2R decoder aims to generate the gold reference \vec{y} during training while R2L decoder attempts to produce the reversed gold reference \overleftarrow{y} at the same time. Given the training data consisting of T sentence pairs $(x^{(t)}, y^{(t)})_{t=1}^T$, the objective is to maximize the log-likelihood over the training data.

$$J(\theta) = \sum_{t=1}^T (\log p(\vec{y}^{(t)} | x^{(t)}) + \log p(\overleftarrow{y}^{(t)} | x^{(t)})) \quad (29)$$

When calculating $p(\vec{y}_i^{(t)})$, L2R decoder usually employs the gold reference $\vec{y}_{<i}^{(t)}$ as the condition $p(\vec{y}_i^{(t)} | \vec{y}_{<i}^{(t)}, x^{(t)})$. In synchronous bidirectional inference, a problem will arise if we directly utilize the gold reference $\overleftarrow{y}_{<i}^{(t)}$ from the other side to compute $p(\vec{y}_i^{(t)} | \vec{y}_{<i}^{(t)}, \overleftarrow{y}_{<i}^{(t)}, x^{(t)})$. For example, in the calculation of $p(\vec{y}_{n-1}^{(t)} | \vec{y}_{<n-1}^{(t)}, \overleftarrow{y}_{<n-1}^{(t)}, x^{(t)})$, $\overleftarrow{y}_{<n-1}^{(t)}$ includes $\overleftarrow{y}_0^{(t)} = \vec{y}_{n-1}^{(t)}$. It indicates that $\vec{y}_{n-1}^{(t)}$ is used to predict itself. Obviously, it is not reasonable. To address this issue during training, we propose two optimization strategies to learn network parameters: the two-pass training strategy [12] and the fine-tuning algorithm.

5.1. Two-pass training

In the first training pass, we learn independent L2R and R2L inference models on the training data. Then, L2R and R2L models are employed to decode the input sentences of the training data, resulting in $(x^{(t)}, \vec{y}^{*(t)})_{t=1}^T$ and $(x^{(t)}, \overleftarrow{y}^{*(t)})_{t=1}^T$ respectively. During the second training pass, $p(\vec{y}_i^{(t)})$ is calculated using $p(\vec{y}_i^{(t)} | \vec{y}_{<i}^{(t)}, \overleftarrow{y}^{*(t)}_{<i}, x^{(t)})$, indicating that the

future context is the model predictions $\hat{y}_{<i}^{(t)}$ rather than gold reference $\bar{y}_{<i}^{(t)}$. Similarly, we calculate $p(\hat{y}_i^{(t)})$ using $p(\hat{y}_i^{(t)} | \bar{y}_{<i}^{(t)}, \hat{y}_{<i}^{(t)}, x^{(t)})$.

5.2. Fine-tuning algorithm

The newly proposed fine-tuning algorithm decouples our synchronous bidirectional inference model into two parts. One performs the synchronous inference function, and the other one performs bidirectional interaction function. We first train the parallel inference for L2R and R2L without interaction just as Equation (30) shows. Each training instance for this step is a triple (x, \bar{y}, \hat{y}) .

$$P(y|x) = \begin{cases} \prod_{i=0}^{n-1} p(\bar{y}_i | \bar{y}_0 \cdots \bar{y}_{i-1}, x) & \text{if L2R} \\ \prod_{i=0}^{n-1} p(\hat{y}_i | \hat{y}_0 \cdots \hat{y}_{i-1}, x) & \text{if R2L} \end{cases} \quad (30)$$

After this simple training procedure converges, we use this model to decode a small subset of the source sentences in the original training data (e.g. 10% of the dataset) and get the new triple $(x, \bar{y}^*, \hat{y}^*)$. Then, we can fine-tune our synchronous bidirectional inference model to enhance the bidirectional interaction function similar to the second pass of the two-pass training strategy.

Compared to the two-pass training strategy, the fine-tuning algorithm is much cheaper to implement since there is no need to train two separate models, to decode the entire training set and to do the second training over the whole dataset. In the experiments, we mainly employ the two-pass strategy and compare these two strategies in the experimental analysis part.

6. Experimental setup

In our experiments, two typical Seq2Seq tasks of machine translation and abstractive summarization are employed to test the effectiveness of our synchronous bidirectional inference model.

6.1. Machine translation

6.1.1. Dataset

We evaluate the proposed synchronous bidirectional inference model on both Chinese-to-English and English-to-German translation tasks. For the Chinese-to-English task, the training data consists of about 2.1M sentence pairs extracted from LDC corpora.⁶ We choose NIST 2002 (MT02) dataset for validation. For testing, we employ NIST 2003-2006 (MT03-06) datasets. We apply Byte-Pair Encoding (BPE)[15] with 30K merge operations and maintain the source and target vocabularies to the most frequent 30K tokens.

For the English-to-German task, we utilize the same subset of the WMT 2014 training corpus employed by [4,14,16,17]. It contains 4.5M sentence pairs.⁷ The concatenation of news-test 2012 and news-test 2013 is used as the validation set. The news-test 2014 is employed as the test set.

6.1.2. Training and evaluation details

For our synchronous bidirectional inference model with **LSTM-based architecture**, we implement the system by reusing and modifying the open source toolkit Zoph_RNN⁸ which is written in C++/CUDA and provides efficient training across multiple GPUs. The encoder includes two stacked LSTM layers and the first layer employs the bidirectional LSTMs. The decoder also contains two stacked LSTM layers followed by the softmax layer. The dimension of word embedding and the size of hidden layers are all set to 1000. The dropout rate is set to 0.2. At test time, we employ beam search with beam size $k = 4$.

For the synchronous bidirectional inference model with **Transformer**, we modify the tensor2tensor⁹ toolkit for training and evaluation. We employ the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.998$, and $\epsilon = 10^{-9}$. The warmup and decay strategy for learning rate are the same as [4], with 16,000 warmup steps. During training, we employ label smoothing of value $\epsilon_{ls} = 0.1$. For evaluation, we use beam search with a beam size of $k = 4$ and length penalty $\alpha = 0.6$. Additionally, we use 6 encoder and decoder layers. In each layer, we employ $d_{model} = 1024$ hidden size, 16 attention-heads, 4096 feed forward inner-layer dimensions, and $P_{dropout} = 0.1$. Our settings are close to *transformer_big* setting as defined in [4]. We employ three Titan Xp GPUs to train English-to-German translation and one GPU for Chinese-to-English translation pairs. In addition,

⁶ LDC2000T50, LDC2002L27, LDC2002T01, LDC2002E18, LDC2003E07, LDC2003E14, LDC2003T17, LDC2004T07.

⁷ All preprocessed dataset and vocabulary can be directly download in tensor2tensor website https://drive.google.com/open?id=0B_bZck-ksdkpM25jRUN2X2UxMm8.

⁸ https://github.com/isi-nlp/Zoph_RNN.

⁹ <https://github.com/tensorflow/tensor2tensor>.

we average the last 20 checkpoints to get the final model for English-to-German but do not perform checkpoint averaging for Chinese-to-English.

We evaluate the final translation quality with case-insensitive BLEU [18] for Chinese-to-English and with case-sensitive BLEU for English-to-German. Significance test is performed using the pairwise re-sampling approach [19].

6.1.3. Translation systems

We use **BI-RNMT** to denote our proposed synchronous bidirectional inference model implemented in LSTM-based recurrent neural machine translation. The proposed synchronous bidirectional inference model for Transformer is named **BIFT**.¹⁰

We compare the proposed models against the following state-of-the-art NMT systems:

- **RNMT** [20]: it is a state-of-the-art LSTM-based NMT system with the same setting as **BI-RNMT**.
- **RNMT (R2L)**: it is a variant of RNMT and produces translations from right to left.
- **Transformer**: it is the state-of-the-art machine translation system with self-attention mechanism using the default left-to-right generation [4].
- **Transformer (R2L)**: it is a variant of Transformer which performs translation in a right-to-left manner.
- **Rerank-NMT**: following [9], we first run beam search for L2R and R2L inference models independently to obtain two k-best lists, and then re-score the union of these two k-best lists. This method assumes that some source sentences are appropriate to translate from left to right, while others are better to translate from right to left. It is denoted as **Rerank-RNMT** for RNMT architecture and **Rerank-Transformer** for Transformer framework.
- **ABD-Transformer**: it is an asynchronous bidirectional inference model under the Transformer framework that performs L2R inference with the results generated by R2L inference model [11]. During inference, two-pass decoding scheme is employed. First, the R2L inference model generates the backward hidden states and corresponding translation results. Then, ABD-Transformer optimizes the L2R inference with the help of the backward hidden states.

6.2. Abstractive summarization

6.2.1. Dataset

Abstractive sentence summarization is a task that generates a title-like summary for a long sentence. Our **training data** is a (text, summary) parallel corpus from the Annotated English Gigaword dataset [21,22]. It contains about 3.8M text-summary pairs for training and 189 K pairs for validation. The encoder and decoder share the same vocabulary of about 90K word types.

For the **test set**, we use both DUC 2004 and the English Gigaword. In the test set of DUC 2004, there are 500 examples and each example pairs a document with four different human-written reference summaries. For the test set of the English Gigaword, we employ the same randomly selected subset of 2000 text-summary pairs as [22,23].

6.2.2. Training and evaluation details

For both LSTM-based architecture and Transformer framework, we use the same model settings as neural machine translation.

For evaluation, we use ROUGE [24] as the metric. ROUGE measures the quality of summary by computing overlapping lexical units, such as unigram, bigram, trigram, and longest common subsequence (LCS). It becomes the standard evaluation metric for DUC shared tasks and popular for summarization evaluation. Following previous work, we use ROUGE-1 (unigram), ROUGE-2 (bigram) and ROUGE-L (LCS) as our evaluation metrics in the reported experimental results.

6.2.3. Summarization systems

We compare our proposed model with the following strong baselines.

ABS: [22] first proposed the abstractive summarization task and used an attentive CNN encoder and NNLM (neural network language model) decoder to perform this task.

LSTM-Sum: it is the abstractive summarization system with the same architecture as RNMT in which the encoder and decoder are both LSTM-based recurrent neural networks.

Feats2S: it is also a RNN encoder-decoder model using gated recurrent unit (GRU) [2] and provides more features (e.g. POS and NER) to enrich the encoder [25].

Selective-Enc: [23] proposed a selective mechanism to select important information from encoder before generating summary.

Transformer: it is a Transformer model which is applied to the abstractive sentence summarization task.

¹⁰ Our code is freely available in github <https://github.com/ZNLP/sb-nmt>.

Table 2

Translation quality for Chinese-to-English tasks using case-insensitive BLEU scores. Our model **BI-RNMT** and **BIFT**($\lambda = 0.1$) are respectively significantly better than corresponding baselines ($p < 0.01$).

Model	MT03	MT04	M05	MT06	AVE	Δ
RNMT	42.07	43.40	40.73	41.11	41.83	–
RNMT (R2L)	41.47	43.13	40.62	40.94	41.54	–0.29
Rerank-RNMT	42.66	43.72	41.18	41.71	42.32	+0.49
BI-RNMT	43.50	43.98	41.37	42.48	42.83	+1.00
Transformer	47.63	48.32	47.51	45.31	47.19	–
Transformer (R2L)	46.79	47.01	46.50	44.13	46.11	–1.08
Rerank-Transformer	48.23	48.91	48.73	46.51	48.10	+0.91
ABD-Transformer	49.47	48.01	48.19	47.09	48.19	+1.00
BIFT	51.87	51.50	51.23	49.83	51.11	+3.92

7. Results and analysis

7.1. Machine translation

7.1.1. Effect of hyper-parameter λ

Since BIFT employs a hyper-parameter λ to weigh the contribution of history and future contexts (see equation (26)), we need to determine the best setting for this hyper-parameter. We empirically set $\lambda = 0.1$, $\lambda = 0.5$ and $\lambda = 1.0$ to train BIFT respectively and get three BIFT models. Finally, we choose the setting that achieves the best performance on the development set.

We conduct this experiment on Chinese-to-English translation and the BLEU scores on the development set (NIST02) are 50.99, 50.72 and 50.96 for $\lambda = 0.1$, $\lambda = 0.5$ and $\lambda = 1.0$ respectively. Interestingly, the translation quality is not very sensitive to the hyper-parameter. Nevertheless, we choose $\lambda = 0.1$ in the following experiments due to the best translation performance.

7.1.2. Overall translation quality

Table 2 reports the translation performance of different systems on the Chinese-English task. The results are mainly divided into two parts. The first part in this table shows the BLEU scores of the systems based on LSTM framework while the second part gives the results of various systems based on the Transformer architecture.

Comparing the baselines using different architectures, we can easily see that the self-attention based Transformer remarkably outperforms the LSTM-based RNMT, with the average improvement of 5.36 BLEU points (47.19 vs. 41.83), suggesting the superiority of the Transformer.

As conventional decoding performs left to right, a question may arise that which inference direction is better. It is easy to find from Table 2 that the right-to-left decoding performs worse than the left-to-right style in Chinese-to-English translation no matter which neural network architecture is adopted (RNMT (R2L) vs. RNMT, Transformer (R2L) vs. Transformer). Specifically, the gap under the Transformer architecture is much bigger (1.08 vs. 0.29), indicating that Transformer is more sensitive to the inference direction.

We also investigate previous methods that take advantage of two inference directions. The results in Table 2 shows that the reranking approach [9] (Rerank-RNMT and Rerank-Transformer) performs much better than the baseline for both RNMT and Transformer architecture. We also notice that the asynchronous bidirectional decoding method ABD-Transformer [11] can slightly outperform the reranking approach. The average gains can be up to 1.0 BLEU point (48.19 vs. 47.19), indicating that L2R decoding and R2L decoding can be complementary to each other.

We go step further to exploit synchronous bidirectional inference that makes full use of L2R and R2L decoding. It is obvious to see from Table 2 that our proposed method performs best. The first part in Table 2 says that the synchronous bidirectional inference model under LSTM framework **BI-RNMT** can obtain an average improvement of 1.0 BLEU point over RNMT. The second part demonstrates that our synchronous bidirectional inference model under the Transformer architecture **BIFT** achieves promising BLEU gains and the gap can be as large as 3.92 BLEU points on average. The remarkable improvements suggest that compared to asynchronous bidirectional decoding, our synchronous bidirectional inference can better explore the history and future contexts on the target side. In addition to the evaluation on the NIST data, we further test our **BIFT** model on WMT Chinese-to-English dataset in which we utilize the CWMT subset including about 9.03 million sentence pairs as the training data. Newstest2017, newstest2018 and newstest2019 are employed as the test sets. Table 3 reports the translation results. We can observe that on the large-scale Chinese-to-English dataset, the synchronous bidirectional inference model BIFT can also substantially outperform the baselines with average gains of 2.3 BLEU points,¹¹ demonstrating the superiority of the proposed synchronous bidirectional inference model.

¹¹ Note that the performance is not on par with the top one achieved in the evaluation campaign. We believe there are two reasons. From the data perspective, we do not employ large-scale monolingual data. From the model perspective, we use the Transformer-base model and do not employ ensemble, system combination and other sophisticated techniques.

Table 3

Translation quality for different systems on WMT Chinese-English machine translation tasks.

Model	test2017	test2018	test2019	AVE	Δ
Transformer	22.92	23.57	24.61	23.70	–
Transformer (R2L)	22.70	22.14	24.21	23.02	–0.68
BIFT	25.69	26.35	25.97	26.00	+2.30

Table 4

Translation results on WMT14 English-to-German task using case-sensitive BLEU.

Model	TEST (WMT 14)
GNMT [20]	24.61
Conv [3]	25.16
AttIsAll [4]	28.40
RNMT	22.85
RNMT (R2L)	22.17
Rerank-RNMT	23.06
BI-RNMT	23.97
Transformer ^a	27.72
Transformer (R2L)	27.13
Rerank-Transformer	27.81
ABD-Transformer	28.22
BIFT	29.21

^a The BLEU score of the Transformer model is reproduced in our hardware environment and is slightly lower than AttIsAll [4]. [26,27] also reported that their reproduction is lower than their original result in [4]. For example, [27] reported 27.9 BLEU points. In our experiments, we use only 3 GPUs for English-to-German, whereas [4,26] adopted TPUs for model training.

Similar phenomena can be observed from the English-German translation results as shown in Table 4. The finding is that **BI-RNMT** outperforms RNMT with 1.12 BLEU points. **BIFT** performs best¹² among all the systems including GNMT [20], Conv [3] and AttIsAll [4]. The improvements of **BIFT** over the baseline Transformer can be up to 1.4 absolute BLEU points. Considering that only one reference is available for English-to-German translation, the improvements are very promising.

It should be noted that the reranking approach performs only slightly better than the baseline on English-to-German translation, while the improvement on Chinese-to-English translation is significant. We believe that there are three reasons behind. First, the beam size is small (e.g. 4, the same as baselines) and the candidate translation hypotheses lack enough diversity. Second, we only employ two features (L2R and R2L translation scores) to rerank the candidates, while others adopt many features such as large-scale language model, word alignment and sentence similarity. Third, there is only one reference for English-to-German translation but there four references for Chinese-to-English NIST translation. It is well known that more references can lead to bigger improvements.

7.1.3. Model size and efficiency

Our synchronous bidirectional inference model is slightly complicated than conventional L2R or R2L inference model. It is interesting to figure out the model size and efficiency of our system compared to other baseline systems. Table 5 reports the corresponding statistics of different NMT models. The model size denotes the total number of network parameters. Since the synchronous bidirectional inference model only introduces one parameter λ in Equation (26) for BIFT, the model size is the same as that of the Transformer. In contrast, Rerank-Transformer has double the number of parameters compared to the Transformer because it requires two individual encoder-decoder models for L2R and R2L decoding respectively. As for the asynchronous bidirectional decoding model, ABD-Transformer shares one encoder and has two decoders, and thus contains more than a half parameters against the Transformer baseline.

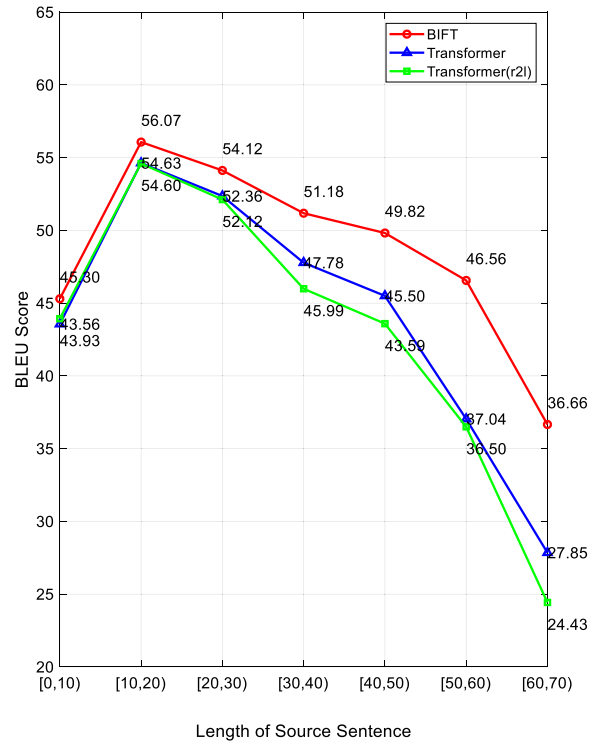

The *Train* column shows the number of global training steps per second. Because the training procedure of BIFT needs to match both of the L2R and R2L references, it takes more time to converge. It is about 40% slower than the baseline Transformer. We further calculated the total training time for different systems. **Transformer** took about 53.7 hours to converge. Our model **BIFT** consumed about 87.5 hours for convergence. Since **BIFT** requires to decode the source sentences of bilingual training data before parameter optimization, we also collected the decoding time together. Greedy search was employed when decoding the training data and it took about 10.3 hours to decode all the source sentences of the training data.

¹² In practice, it is not the new state-of-the-art on English-to-German translation. For example, [28] reported 29.7 BLEU points on the same dataset.

Table 5

Comparison results of model size, training and testing efficiency. *Train* denotes the number of global training steps processed per second at the same batch-sized sentences; *Test* indicates the amount of translated sentences in one second.

Model	Model size	Efficiency	
		<i>Train</i>	<i>Test</i>
Transformer	207.8M	2.07	19.97
Transformer (R2L)	207.8M	2.07	19.81
Rerank-Transformer	415.6M	1.03	6.51
ABD-Transformer	333.8M	1.18	7.20
BIFT	207.8M	1.26	17.87

**Fig. 4.** BLEU score trends for different intervals of sentence length.

Therefore, **BIFT** spent a total of 97.8 hours for model training. However, regarding the decoding efficiency, our synchronous bidirectional inference model performs on par with the Transformer baseline and is only 10% slowdown (17.87 vs. 19.97 sentences per second), whereas Rerank-Transformer and ABD-Transformer are much slower. The statistics suggest that BIFT is acceptable regarding the decoding efficiency.

7.1.4. Performance trends on sentence length


In previous sections, we argued that BI-RNMT and BIFT can utilize both of the history and future contexts during translation. A natural question may arise that whether our proposed model would perform much better on the long sentences. To answer this question, we group the source sentences of similar lengths in the test set and calculate the corresponding BLEU scores for each length interval.

Fig. 4 displays the statistics. The left part of Fig. 4 represents the results of the systems using LSTM-based framework. The right one shows the statistics of Transformer-based systems. Overall, no matter which architecture is adopted, our proposed models (BI-RNMT and BIFT) are superior to baselines over sentences with all different lengths. Generally, the gap becomes bigger and bigger when the length grows. Comparing the two architectures, we find that the Transformer-based BIFT excels in long sentence translation. The results indicate that our proposed synchronous bidirectional inference model are better at translating long sentences with the help of both the history and future contexts during decoding.

Table 6

Matching accuracy of the first and last four tokens between model predictions and references in NIST Chinese-English machine translation tasks for different NMT systems.

Model	First four	Last four
RNMT	36.35%	31.64%
RNMT (R2L)	31.22%	34.01%
BI-RNMT	36.88%	34.65%
Transformer	40.21%	35.10%
Transformer (R2L)	35.67%	39.47%
Rerank-Transformer	38.98%	38.91%
ABD-Transformer	38.36%	38.11%
BIFT	40.89%	40.08%

**Fig. 5.** Match precision of translation tokens for different positions.

7.1.5. Translation precision over different positions

We mentioned in introduction that L2R inference is good at predicting prefix while R2L inference is adept at suffix prediction. We may wonder that how does our synchronous bidirectional inference model perform on prefix and suffix prediction. Furthermore, which part of the translation sentence will be improved most, the prefix, middle part or the suffix?

To figure out these questions, we first analyze the prediction precision of the first-four and last-four tokens of the translation compared to the references for different NMT systems. Table 6 reports the comparison results. Obviously, the proposed synchronous bidirectional inference model performs best for matching precision of both the first-four and last-four tokens, showing the superiority of our methods.

Then, we go step further and conduct a deep analysis. We divide each translation hypothesis and its reference into 10 equal parts and calculate the average word prediction accuracy for each part. In this way, we attempt to investigate the contribution of the synchronous bidirectional inference model over different positions. Fig. 5 illustrates the comparison results. It is interesting to see from this figure that both ends of the translation hypothesis are much easier to predict. In contrast, the prediction accuracy of the middle part is much lower, suggesting more demands of both the history and future contexts. Our proposed BI-RNMT and BIFT facilitate the usage of the left and right predictions, leading to large improvements over the middle part (40-80% in Fig. 5).

Table 7

Translation results of different training strategies for our synchronous bidirectional inference model on Chinese-to-English tasks.

Model	MT03	MT04	M05	MT06	AVE	Δ
Transformer	47.63	48.32	47.51	45.31	47.19	–
BIFT (original data)	45.24	45.03	43.35	43.35	44.24	–2.95
BIFT (two-pass)	51.87	51.50	51.23	49.83	51.11	+3.92
BIFT (fine-tuning)	50.76	49.72	48.32	47.91	49.18	+1.99

Table 8

Abstractive summarization quality on DUC 2004 and English Gigaword for different methods.

Model	DUC-2004			English Gigaword		
	R1	R2	R-L	R1	R2	R-L
ABS	26.55	7.06	22.05	29.55	11.32	26.42
Feats2s	28.35	9.46	24.59	32.67	15.59	30.64
Selective-Env	29.21	9.56	25.51	36.15	17.54	33.63
RNMT	28.22	10.21	25.14	34.54	16.85	32.32
BI-RNMT	29.05	10.90	26.05	35.47	17.62	32.90
Transformer	28.09	9.52	24.91	34.12	16.04	31.46
BIFT	29.17	10.30	26.05	35.68	17.39	32.89

7.1.6. Two-pass training vs. fine-tuning algorithm

In this subsection, we attempt to investigate the effects of different parameter optimization strategies for our synchronous bidirectional inference model. We employ BIFT and Chinese-to-English translation task to compare between the two-pass training strategy and the fine-tuning algorithm. In the fine-tuning step, we randomly choose 10% source sentences of the training data. In addition to these two training algorithms, we also test the performance if we leave the gap between training and inference as is, so as to figure out whether the two-pass training or fine-tuning algorithm is necessary or not. Table 7 reports the comparison results.

As shown in the table, if we do not tackle the training-inference inconsistency and utilize the original training data directly, namely using the triple (x, \bar{y}, \bar{y}) as training instance, the translation quality is much decreased with a drop of 2.95 BLEU points in average compared to the baseline (BIFT (original data) vs. Transformer). Obviously, the synchronous bidirectional inference model cannot be optimized in such a way that target words are predicted conditioned on themselves.

Fortunately, we observe from the last two rows in Table 7 that both of the proposed training algorithms can remarkably improve the translation performance compared to the Transformer baseline. Although the fine-tuning strategy is not as powerful as the two-pass training strategy, it can still achieve a remarkable improvement of 2 BLEU points in average over the strong baseline. Considering that the fine-tuning strategy is much easier and cheaper for system deployment, we believe this strategy will be more popular than the two-pass strategy.

7.2. Abstractive summarization

Abstractive sentence summarization is another well-known testbed for sequence to sequence learning. We then apply our synchronous bidirectional inference model into this task. Table 8 presents the results of different systems over two test sets DUC-2004 and English Gigaword.

The first three rows show the performances of strong baselines. Among these three baselines, both Feats2s and Selective-Env aim at improving the summarization quality by enhancing the representation learning of the encoder. Selective-Env augments the encoder with key information selection and performs best.

In contrast, our method attempts to improve the decoder (inference module) by enabling synchronous bidirectional decoding. The last four rows in Table 8 demonstrate that the LSTM-based RNMT stably outperforms the self-attention based Transformer, which is quite different from that for neural machine translation in which Transformer is the better one. After applying our synchronous bidirectional inference model, BI-RNMT and BIFT respectively achieve significantly better results on two test sets over RNMT and Transformer, despite that these two models cannot outperform Selective-Env on the Gigaword test set. The reason behind may be that we just apply BI-RNMT and BIFT into abstractive summarization without any special adaptation processing. From another perspective, the two kinds of the models handle encoder and decoder respectively, and can be complementary to each other. Nevertheless, the statistics given in the table further show the effectiveness of our synchronous bidirectional inference model beyond machine translation.

7.3. Some examples

To better understand the models, we further investigate some specific examples for machine translation and text summarization respectively.

Input	中国有13亿人口,不管多么小的问题,只要乘以13亿,那就成为很大很大的问题;不管多么可观的财力、物力,只要除以13亿,那就成为很低很低的人均水平。
Ref	China has a population of 1.3 billion. No matter how small a problem is, if it times 1.3 billion, it will become a very big problem; no matter how significant the financial and material resources are, if divided by 1.3 billion, it will be at a very low per capita level.
Transformer	China has 1.3 billion people, no matter how small the population, so long as it means 1.3 billion, it will become a big problem.
Transformer (R2L)	Regardless of the enormous amount of financial and material resources , so long as it is divided by nt \$1.3 billion, the per capita level will be very low.
BIFT	China has a population of 1.3 billion, no matter how small, which means taking advantage of 1.3 billion. No matter how considerable financial and material resources are, it will become a very low per capita level.

Input	中方相信,各方都能意识到肩负的和平使命,以建设性的态度,合作与谅解的精神,相互尊重,显示灵活,缩小差距。
Ref:	The Chinese party believes that all the parties are aware of their peace missions, will adopt constructive attitudes, show the spirit of cooperation and understanding, respect each other, display flexibility and narrow down the differences.
Transformer	The Chinese side believes that all sides will be able to realize their peaceful mission, and show mutual respect, flexibility, and gaps in a constructive manner in the spirit of cooperation and understanding.
Transformer (R2L)	The Chinese side believes that all sides respect each other, show flexibility, and narrow the gap with a constructive attitude, cooperation, and understanding.
BIFT	The Chinese side believes that all sides are aware of the peace mission on their shoulders, with a constructive attitude, and the spirit of cooperation and understanding, respect each other, show flexibility, and narrow the gap .

Fig. 6. Two examples of Chinese-to-English translation.

Input	resident nelson mandela acknowledged saturday the african national congress violated human rights during apartheid , setting him at odds with his deputy president over a report that has divided much of south africa.
Ref	mandela acknowledges human rights violations by african national congress
Transformer	mandela acknowledges human rights violation at odds with deputy president
Transformer (R2L)	mandela says south africa violated human rights
BIFT	mandela says south african national congress violated human rights

Input	the new york times said in an editorial on monday, oct. UNK : since the deadly bombing of two american embassies in africa in august , there has been a troubling accumulation of evidence that the state department inexplicably ignored warnings of possible terrorist attacks against the installations .
Ref	editorial claims state department ignored warnings of terrorist
Transformer	new york times says evidence of terror attacks is ignored
Transformer (R2L)	new york times warns of possible terrorist attacks on u.s. embassies in africa
BIFT	new york times says state department ignored terror warnings

Fig. 7. Two examples of abstractive summarization.

Fig. 6 shows two examples for Chinese-to-English translation. In the first example, the left-to-right Transformer omits a sub-sentence in the right part and the right-to-left Transformer neglects to translate the left sub-sentence. In contrast, with the help of both history and future contexts, our synchronous bidirectional inference model BIFT obtains a much better translation which covers all the contents of the source input. The phenomenon in the second example is similar. The left-to-right baseline forgets to translate *display flexibility and narrow down the differences* at the end, while the right-to-left baseline fails to generate *are aware of their peace missions* at the beginning part. However, the proposed BIFT model could produce all the contents correctly, suggesting the effectiveness of synchronous bidirectional inference.

The summarization examples are displayed in Fig. 7. In each example, the input is a long sentence and the output is a title-like summary.

In the first example, *human rights* and *african national congress* are two key contents. The baseline Transformer fails to generate *african national congress* in the tail part and Transformer (R2L) neglects *natrail congress* in the head part. In contrast, our model BIFT renders all of the key points. As for the second example, the baselines have made similar mistakes. Transformer omits *by state department* at the end, and the summary generated by Transformer (R2L) expresses the wrong meaning due to the absence of *says state department fails to* after *new york times*. However, our model BIFT can generate the summary with correct and complete contents. The examples demonstrate the superiority of our synchronous bidirectional inference model on the modeling of both history and future contexts.

8. Related work

This work addresses the generalized synchronous bidirectional inference for sequence to sequence learning tasks, aiming to take full advantage of history and future predictions on the output. Generally, the related work can be divided into two categories, namely bidirectional inference and future context usage.

Bidirectional inference is well studied in sequential labeling tasks [29–31], in which each input token corresponds to an output label and the output shares the same length with the input. In general, there are 2^{n-1} decomposition ways of the conditional probability $p(y_0^{n-1}|x_0^{n-1})$ for an n -token input sequence, since each token is predicted after the left one (from left to right) or the right one (from right to left). In this way, bidirectional inference is not difficult for sequential labeling. However, it is not trivial to leverage bidirectional inference for sequence generation problems mainly due to the length nondeterminacy of the output. [9,7,11] added agreement constraints to enforce L2R inference output to be consistent with R2L inference output for sequence generation tasks. [32] proposed the twin network that encourages the target hidden states of the L2R and R2L inferences at the same position to be as close as possible to predict the same token during training. Recently, [11] introduced an asynchronous bidirectional inference model for neural machine translation. They first obtained the translation hypothesis using R2L inference and then optimized the L2R inference model with the help of the R2L inference result. Despite of performance improvement, all these studies require two individual inference models, making the architecture more complicated. Furthermore, the interactions between L2R and R2L inferences are not adequate. Taking the asynchronous bidirectional inference model for example, L2R model can utilize the information of R2L model but R2L inference cannot use the L2R predictions. In contrast, our synchronous bidirectional inference model has only a single decoder in which L2R and R2L inferences interact with each at each decoding time step.

Using future contexts has drawn more and more attention in sequence prediction tasks. Intuitively, the R2L inference model can be employed to re-rank the n -best hypotheses of the L2R inference model, preferring balanced output [33–38]. To use the future context which is unavailable in conventional inference model, [39–41] proposed the reinforcement learning methods to estimate the possible future information. To mimic the human cognitive behaviors, [42] presented a deliberation network, which leverages the global information with the help of both forward and backward predictions in sequence generation through a deliberation process. [43] introduced two additional recurrent layers to model translated past contents and untranslated future contents. They much improved the sequence generation quality with the cost of model complexity. They either employed two-pass decoding strategy or added more layers to the original network. Compared to this kind of work, our proposed model uses a smart way to exploit both history and future predictions by allowing L2R and R2L inferences to perform in parallel but interactively. In our previous work [12], we address the bidirectional decoding for neural machine translation. In this current work, the new contributions are three-fold. First, we generalize the decoding model into synchronous bidirectional inference for the general sequence-to-sequence models (e.g. LSTM and Transformer) and general sequence generation tasks (e.g. translation and summarization). Second, in addition to the two-pass training method, we further propose and investigate a new fine-tuning based algorithm for efficient parameter optimization. Third, we conduct deep analysis about the characteristics of the proposed synchronous bidirectional inference model.

9. Conclusion and future work

This work proposes a synchronous bidirectional inference model for sequential generation tasks. We first presented a synchronous bidirectional beam search algorithm for sequence generation, in which left-to-right and right-to-left decoding perform in parallel but interactively. We have exploited the usage of synchronous bidirectional inference model on both LSTM-based and Transformer-based seq2seq architectures. We have also proposed and investigated two parameter optimization strategies. The comprehensive experiments on machine translation and abstractive summarization have demonstrated that our proposed synchronous bidirectional inference model remarkably outperforms the strong baselines. The deep analysis further shows that our model can indeed take full advantage of both history and future predictions during inference.

In the future, we plan to explore three directions. The first one is to further improve the training algorithm. Currently, we need to produce some pseudo training corpus to mimic the test environment and this requirement adds the complexity to the training procedure. Accordingly, we attempt to design a more efficient training algorithm to avoid the need of pseudo data construction. For example, we will try the method proposed by [44] to bridge the gap between training and inference.

The second direction is to explore more applications of the synchronous bidirectional inference model. For example, we will apply the proposed model into other sequential generation tasks, such as question answering, chatbot and image caption.

The third one is to generalize the synchronous bidirectional inference model. The current model handles the synchronous interaction between two inference directions. We believe that the synchronous interaction can be leveraged in many multi-task generation tasks. For instance, speech recognition and speech-to-text translation share the same input and the two inference processes can interact with other other to make use of complementary information.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The research work described in this paper has been supported by the Natural Science Foundation of China under Grant No. U1836221 and 61673380. The research work in this paper has also been supported by the Beijing Municipal Science and Technology Project No. Z181100008918017.

References

- [1] I. Sutskever, O. Vinyals, Q.V. Le, Sequence to sequence learning with neural networks, in: Proceedings of NIPS, 2014.
- [2] D. Bahdanau, K. Cho, Y. Bengio, Neural machine translation by jointly learning to align and translate, in: Proceedings of ICLR, 2015.
- [3] J. Gehring, M. Auli, D. Grangier, D. Yarats, Y.N. Dauphin, Convolutional sequence to sequence learning, in: Proceedings of ICML, 2017.
- [4] A. Vawani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, arXiv preprint arXiv:1706.03762.
- [5] T. Watanabe, E. Sumita, Bidirectional decoding for statistical machine translation, in: Proceedings of the 19th International Conference on Computational Linguistics, vol. 1, 2002, pp. 1–7.
- [6] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (8) (1997) 1735–1780.
- [7] L. Liu, A.M. Finch, M. Utiyama, E. Sumita, Agreement on target-bidirectional LSTMs for sequence-to-sequence learning, in: AAAI, 2016, pp. 2630–2637.
- [8] Z. Zhang, S. Wu, S. Liu, M. Li, M. Zhou, E. Chen, Regularizing neural machine translation by target-bidirectional agreement, preprint arXiv:1808.04064.
- [9] L. Liu, M. Utiyama, A. Finch, E. Sumita, Agreement on target-bidirectional neural machine translation, in: Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2016, pp. 411–416.
- [10] Y. Wang, S. Cheng, L. Jiang, J. Yang, W. Chen, M. Li, L. Shi, Y. Wang, H. Yang, Sogou neural machine translation systems for WMT17, in: Proceedings of the Second Conference on Machine Translation, 2017, pp. 410–415.
- [11] X. Zhang, J. Su, Y. Qin, Y. Liu, R. Ji, H. Wang, Asynchronous bidirectional decoding for neural machine translation, in: Proceedings of AAAI, 2018.
- [12] L. Zhou, J. Zhang, C. Zong, Synchronous bidirectional neural machine translation, in: TACL, 2019.
- [13] M. Elbayad, L. Besacier, J. Verbeek, Pervasive attention: 2d convolutional neural networks for sequence-to-sequence prediction, in: Proceedings of the 22nd Conference on Computational Natural Language Learning, 2018, pp. 97–107.
- [14] M.-T. Luong, H. Pham, C.D. Manning, Effective approaches to attention-based neural machine translation, in: Proceedings of EMNLP, 2015.
- [15] R. Sennrich, B. Haddow, A. Birch, Neural machine translation of rare words with subword units, in: Proceedings of ACL, 2016.
- [16] S. Shen, Y. Cheng, Z. He, W. He, H. Wu, M. Sun, Y. Liu, Minimum risk training for neural machine translation, in: Proceedings of ACL, 2016.
- [17] J. Zhou, Y. Cao, X. Wang, P. Li, W. Xu, Deep recurrent models with fast-forward connections for neural machine translation, arXiv preprint arXiv:1606.04199.
- [18] K. Papineni, S. Roukos, T. Ward, W.-J. Zhu, Bleu: a method for automatic evaluation of machine translation, in: Proceedings of ACL, 2002, pp. 311–318.
- [19] P. Koehn, Statistical significance tests for machine translation evaluation, in: Proceedings of EMNLP, 2004, pp. 388–395.
- [20] Y. Wu, M. Schuster, Z. Chen, Q.V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, et al., Google's neural machine translation system: bridging the gap between human and machine translation, arXiv preprint arXiv:1609.08144.
- [21] C. Napoles, M. Gormley, B. Van Durme, Annotated gigaword, in: Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction, 2012, pp. 95–100.
- [22] A.M. Rush, S. Chopra, J. Weston, A neural attention model for abstractive sentence summarization, in: Proceedings of EMNLP, 2015.
- [23] Q. Zhou, N. Yang, F. Wei, M. Zhou, Selective encoding for abstractive sentence summarization, in: Proceedings of ACL, 2017.
- [24] C.-Y. Lin, Rouge: a package for automatic evaluation of summaries, in: Text Summarization Branches Out: Proceedings of the ACL-04 Workshop, vol. 8, 2004.
- [25] R. Nallapati, B. Zhou, C. Gulcehre, B. Xiang, Abstractive text summarization using sequence-to-sequence RNNs and beyond, in: Computational Natural Language Learning, 2016.
- [26] M.X. Chen, O. Firat, A. Bapna, M. Johnson, W. Macherey, G. Foster, L. Jones, N. Parmar, M. Schuster, Z. Chen, et al., The best of both worlds: combining recent advances in neural machine translation, in: Proceedings of ACL, 2018, pp. 76–86.
- [27] P. Shaw, J. Uszkoreit, A. Vaswani, Self-attention with relative position representations, in: Proceedings of NAACL, 2018, pp. 464–468.
- [28] F. Wu, A. Fan, A. Baevski, Y.N. Dauphin, M. Auli, Pay less attention with lightweight and dynamic convolutions, arXiv preprint arXiv:1901.10430.
- [29] K. Toutanova, D. Klein, C.D. Manning, Y. Singer, Feature-rich part-of-speech tagging with a cyclic dependency network, in: Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology, vol. 1, 2003, pp. 173–180.
- [30] Y. Tsuruoka, J. Tsujii, Bidirectional inference with the easiest-first strategy for tagging sequence data, in: Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing, 2005, pp. 467–474.
- [31] L. Shen, G. Satta, A. Joshi, Guided learning for bidirectional sequence classification, in: Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, 2007, pp. 760–767.
- [32] D. Serdyuk, N.R. Ke, A. Sordani, A. Trischler, C. Pal, Y. Bengio, Twin networks: matching the future for sequence generation, in: International Conference on Learning Representations, 2018.
- [33] R. Sennrich, B. Haddow, A. Birch, Edinburgh neural machine translation systems for WMT 16, in: Proceedings of the First Conference on Machine Translation: Vol. 2, Shared Task Papers, 2016, pp. 371–376.

- [34] R. Sennrich, A. Birch, A. Currey, U. Germann, B. Haddow, K. Heafield, A.V.M. Barone, P. Williams, The university of Edinburgh's neural mt systems for WMT17, in: *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, 2017, pp. 389–399.
- [35] C.D.V. Hoang, G. Haffari, T. Cohn, Towards decoding as continuous optimisation in neural machine translation, in: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 146–156.
- [36] Z. Tan, B. Wang, J. Hu, Y. Chen, X. Shi, XMU neural machine translation systems for WMT 17, in: *Proceedings of the Second Conference on Machine Translation*, 2017, pp. 400–404.
- [37] Y. Deng, S. Cheng, J. Lu, K. Song, J. Wang, S. Wu, L. Yao, G. Zhang, H. Zhang, P. Zhang, et al., Alibaba's neural machine translation systems for WMT18, in: *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, 2018, pp. 368–376.
- [38] Y. Liu, L. Zhou, Y. Wang, Y. Zhao, J. Zhang, C. Zong, A comparable study on model averaging, ensembling and reranking in NMT, in: *CCF International Conference on Natural Language Processing and Chinese Computing*, Springer, 2018, pp. 299–308.
- [39] D. Bahdanau, P. Brakel, K. Xu, A. Goyal, R. Lowe, J. Pineau, A. Courville, Y. Bengio, An actor-critic algorithm for sequence prediction, in: *Proceedings of ICLR*, 2017.
- [40] D. He, H. Lu, Y. Xia, T. Qin, L. Wang, T.-Y. Liu, Decoding with value networks for neural machine translation, in: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 178–187.
- [41] J. Li, W. Monroe, D. Jurafsky, Learning to decode for future success, preprint arXiv:1701.06549.
- [42] Y. Xia, F. Tian, L. Wu, J. Lin, T. Qin, N. Yu, T.-Y. Liu, Deliberation networks: sequence generation beyond one-pass decoding, in: *Advances in Neural Information Processing Systems*, 2017, pp. 1784–1794.
- [43] Z. Zheng, H. Zhou, S. Huang, L. Mou, X. Dai, J. Chen, Z. Tu, Modeling past and future for neural machine translation, in: *TACL*, vol. 6, 2018, pp. 145–157.
- [44] W. Zhang, Y. Feng, F. Meng, D. You, Q. Liu, Bridging the gap between training and inference for neural machine translation, in: *Proceedings of ACL*, 2019, pp. 4334–4343.