

# Simplifying Neural Machine Translation with Addition-Subtraction Twin-Gated Recurrent Networks

Biao Zhang<sup>1</sup>, Deyi Xiong<sup>2</sup>, Jinsong Su<sup>1\*</sup>, Qian Lin<sup>1</sup> and Huiji Zhang<sup>3</sup>

Xiamen University, Xiamen, China<sup>1</sup>

Soochow University, Suzhou, China<sup>2</sup>

Xiamen Meiya Pico information Co.,Ltd. Xiamen, China<sup>3</sup>

{zb, qianl}@stu.xmu.edu.cn, dyxiong@suda.edu.cn, jssu@xmu.edu.cn  
zhanghj@300188.cn

## Abstract

In this paper, we propose an addition-subtraction twin-gated recurrent network (ATR) to simplify neural machine translation. The recurrent units of ATR are heavily simplified to have the smallest number of weight matrices among units of all existing gated RNNs. With the simple addition and subtraction operation, we introduce a twin-gated mechanism to build input and forget gates which are highly correlated. Despite this simplification, the essential non-linearities and capability of modeling long-distance dependencies are preserved. Additionally, the proposed ATR is more transparent than LSTM/GRU due to the simplification. Forward self-attention can be easily established in ATR, which makes the proposed network interpretable. Experiments on WMT14 translation tasks demonstrate that ATR-based neural machine translation can yield competitive performance on English-German and English-French language pairs in terms of both translation quality and speed. Further experiments on NIST Chinese-English translation, natural language inference and Chinese word segmentation verify the generality and applicability of ATR on different natural language processing tasks.

## 1 Introduction

Neural machine translation (NMT), typically with an attention-based encoder-decoder framework (Bahdanau et al., 2015), has recently become the dominant approach to machine translation and already been deployed for online translation services (Wu et al., 2016). Recurrent neural networks (RNN), e.g., LSTMs (Hochreiter and Schmidhuber, 1997) or GRUs (Chung et al., 2014), are widely used as the encoder and decoder for NMT. In order to alleviate the gradient

vanishing issue found in simple recurrent neural networks (SRNN) (Elman, 1990), recurrent units in LSTMs or GRUs normally introduce different gates to create shortcuts for gradient information to pass through.

Notwithstanding the capability of these gated recurrent networks in learning long-distance dependencies, they use remarkably more matrix transformations (i.e., more parameters) than SRNN. And with many non-linear functions modeling inputs, hidden states and outputs, they are also less transparent than SRNN. These make NMT which is based on these gated RNNs suffer from not only inefficiency in training and inference due to recurrency and heavy computation in recurrent units (Vaswani et al., 2017) but also difficulty in producing interpretable models (Lee et al., 2017). These also hinder the deployment of NMT models particularly on memory- and computation-limited devices.

In this paper, our key interest is to simplify recurrent units in RNN-based NMT. In doing so, we want to investigate how further we can advance RNN-based NMT in terms of the number of parameters (i.e., memory consumption), running speed and interpretability. This simplification shall preserve the capability of modeling long-distance dependencies in LSTMs/GRUs and the expressive power of recurrent non-linearities in SRNN. The simplification shall also reduce computation load and physical memory consumption in recurrent units on the one hand and allow us to take a good look into the inner workings of RNNs on the other hand.

In order to achieve this goal, we propose an addition-subtraction twin-gated recurrent network (ATR) for NMT. In the recurrent units of ATR, we only keep the very essential weight matrices: one over the input and the other over the history (similar to SRNN). Comparing with previous

\*Corresponding author.

RNN variants (e.g., LSTM or GRU), we have the smallest number of weight matrices. This will reduce the computation load of matrix multiplication. ATR also uses gates to bypass the vanishing gradient problem so as to capture long-range dependencies. Specifically, we use the addition and subtraction operations between the weighted history and input to estimate an input and forget gate respectively. These add-sub operations not only distinguish the two gates so that we do not need to have different weight matrices for them, but also make the two gates dynamically correlate to each other. Finally, we remove some non-linearities in recurrent units.

Due to these simplifications, we can easily show that each new state in ATR is an unnormalized weighted sum of previous inputs, similar to recurrent additive networks (Lee et al., 2017). This property not only allows us to trace each state back to those inputs which contribute more but also establishes unnormalized forward self-attention between the current state and all its previous inputs. The self-attention mechanism has already proved very useful in non-recurrent NMT (Vaswani et al., 2017).

We build our NMT systems on the proposed ATR with a single-layer encoder and decoder. Experiments on WMT14 English-German and English-French translation tasks show that our model yields competitive results compared with GRU/LSTM-based NMT. When we integrate an orthogonal context-aware encoder (still single layer) into ATR-based NMT, our model (yielding 24.97 and 39.06 BLEU on English-German and English-French translation respectively) is even comparable to deep RNN and non-RNN NMT models which are all with multiple encoder/decoder layers. In-depth analyses demonstrate that ATR is more efficient than LSTM/GRU in terms of NMT training and decoding speed.

We adapt our model to other language translation and natural language processing tasks, including NIST Chinese-English translation, natural language inference and Chinese word segmentation. Our conclusions still hold on all these tasks.

## 2 Related Work

The most widely used RNN models are LSTM (Hochreiter and Schmidhuber, 1997) and GRU (Chung et al., 2014), both of which are good at handling gradient vanishing problem, a

notorious bottleneck of the simple RNN (Elman, 1990). The design of gates in our model follows the gate philosophy in LSTM/GRU.

Our work is closely related to the recurrent additive network (RAN) proposed by Lee et al. (2017). They empirically demonstrate that many non-linearities commonly used in RNN transition dynamics can be removed, and that recurrent hidden states computed as purely the weighted sum of input vectors can be quite efficient in language modeling. Our work follows the same spirit of simplifying recurrent units as they do. But our proposed ATR is significantly different from RAN in three aspects. First, ATR is simpler than RAN with even fewer parameters. There are only two weight matrices in ATR while four different weight matrices in the simplest version of RAN (two for each gate in RAN). Second, since the only difference between the input and forget gate in ATR is the addition/subtraction operation between the history and input, the two gates can be learned to be highly correlated as shown in our analysis. Finally, although RAN is verified effective in language modeling, our experiments show that ATR is better than RAN in machine translation in terms of both speed and translation quality.

To speed up RNN models, a line of work has attempted to remove recurrent connections. For example, Bradbury et al. (2016) propose the quasi-recurrent neural network (QRNN) which uses convolutional layers and a minimalist recurrent pooling function to improve parallelism. Very recently, Lei and Zhang (2017) propose a simple recurrent unit (SRU). With the cuDNN optimization, their RNN model can be trained as fast as CNNs. However, to obtain promising results, QRNN and SRU have to use deep architectures. In practice, 4-layer QRNN encoder and decoder are used to gain translation quality that is comparable to that of single-layer LSTM/GRU NMT. In particular, our one-layer model achieves significantly higher performance than a 10-layer SRU system.

Finally, our work is also related to the efforts in developing alternative architectures for NMT models. Zhou et al. (2016) introduce fast-forward connections between adjacent stacked RNN layers to ease gradient propagation. Wang et al. (2017a) propose a linear associate unit to reduce the gradient propagation length along layers in deep NMT. Gehring et al. (2017b) and Vaswani et al. (2017) explore purely convolutional and attentional archi-




Figure 1: Architecture for LSTM, GRU and ATR.  $c_*$  indicates the memory cell specific to the LSTM network.  $x_*$  and  $h_*$  denote the input and output hidden states respectively.

tectures as alternatives to RNNs for neural translation. With careful configurations, their deep models achieve state-of-the-art performance on various datasets.

### 3 Addition-Subtraction Twin-Gated Recurrent Network

Given a sequence  $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ , RNN updates the hidden state  $\mathbf{h}_t$  recurrently as follows:

$$\mathbf{h}_t = \phi(\mathbf{h}_{t-1}, \mathbf{x}_t) \quad (1)$$

where  $\mathbf{h}_{t-1}$  is the previous hidden state, which is considered to store information from all previous inputs, and  $\mathbf{x}_t$  is the current input. The function  $\phi(\cdot)$  is a non-linear recurrent function, abstracting away from details in recurrent units.

GRU can be considered as a simplified version of LSTM. In this paper, theoretically, we use GRU as our benchmark and propose a new recurrent unit to further simplify it. The GRU function is defined as follows (see Figure 1b):

$$\mathbf{z}_t = \sigma(\mathbf{W}_z \mathbf{x}_t + \mathbf{U}_z \mathbf{h}_{t-1}) \quad (2)$$

$$\mathbf{r}_t = \sigma(\mathbf{W}_r \mathbf{x}_t + \mathbf{U}_r \mathbf{h}_{t-1}) \quad (3)$$

$$\tilde{\mathbf{h}}_t = \tanh(\mathbf{W}_h \mathbf{x}_t + \mathbf{U}_h (\mathbf{r}_t \odot \mathbf{h}_{t-1})) \quad (4)$$

$$\mathbf{h}_t = \mathbf{z}_t \odot \mathbf{h}_{t-1} + (1 - \mathbf{z}_t) \odot \tilde{\mathbf{h}}_t \quad (5)$$

where  $\odot$  denotes an element-wise multiplication. The *reset* gate  $\mathbf{r}_t$  and *update* gate  $\mathbf{z}_t$  enable manageable information flow from the history and the current input to the new state respectively. Despite the success of these two gates in handling gradient flow, they consume extensive matrix transformations and weight parameters.

We argue that many of these matrix transformations are not essential. We therefore propose an addition-subtraction twin-gated recurrent unit

(ATR), formulated as follows (see Figure 1c):

$$\mathbf{p}_t = \mathbf{W}_h \mathbf{h}_{t-1}, \quad \mathbf{q}_t = \mathbf{W}_x \mathbf{x}_t \quad (6)$$

$$\mathbf{i}_t = \sigma(\mathbf{p}_t + \mathbf{q}_t) \quad (7)$$

$$\mathbf{f}_t = \sigma(\mathbf{p}_t - \mathbf{q}_t) \quad (8)$$

$$\mathbf{h}_t = \mathbf{i}_t \odot \mathbf{q}_t + \mathbf{f}_t \odot \mathbf{h}_{t-1} \quad (9)$$

The hidden state  $\mathbf{h}_t$  in ATR is a weighted mixture of both the current input  $\mathbf{q}_t$  and the history  $\mathbf{h}_{t-1}$  controlled by an *input* gate  $\mathbf{i}_t$  and a *forget* gate  $\mathbf{f}_t$  respectively. Notice that we use the transformed representation  $\mathbf{q}_t$  for the current input rather than the raw vector  $\mathbf{x}_t$  due to the potential mismatch in dimensions between  $\mathbf{h}_t$  and  $\mathbf{x}_t$ .

Similar to GRU, we use gates, especially the forget gate, to control the back-propagated gradient flow to make sure gradients will neither vanish nor explode. We also preserve the non-linearities of SRNN in ATR but only in the two gates.

There are three significant differences of ATR from GRU. Some of these differences are due to the simplifications introduced in ATR. First, we squeeze the number of weight matrices in gate calculation from four to two (see Equation (2&3) and (7&8)). In all existing gated RNNs, the inputs to gates are weighted sum of the previous hidden state and input. In order to distinguish these gates, the weight matrices over the previous hidden state and the current input should be different for different gates. The number of different weight matrices in gates is therefore  $2|\#\text{gates}|$  in previous gated RNNs. Different from them, ATR introduces different operations (i.e., addition and subtraction) between the weighted history and input to distinguish the input and forget gate. Therefore, the weight matrices over the previous state/input in the two gates can be the same in ATR. Second, we keep the very essential non-linearities, only in the two gates. In ATR, the role of  $\mathbf{q}_t$  is similar to that of  $\tilde{\mathbf{h}}_t$  in GRU (see Equation (4)). However, we completely wipe out the recurrent non-linearity




Figure 2: Visualization of the difference between  $\sigma(x+y)$  (the input gate) and  $\sigma(x-y)$  (the forget gate). Here, we set  $x, y \in [-5, 5]$ .

of  $\tilde{\mathbf{h}}_t$  in  $\mathbf{q}_t$  (i.e.,  $\mathbf{q}_t = \mathbf{W}_x \mathbf{x}_t$ ). Lee et al. (2017) show that this non-linearity is not necessary in language modeling. We further empirically demonstrate that it is neither essential in machine translation. Third, in GRU the gates for  $\tilde{\mathbf{h}}_t$  and  $\mathbf{h}_{t-1}$  are coupled and normalized to 1 while we do not explicitly associate the two gates in ATR. Instead, they can be learned to be correlated in an implicit way, as shown in the next subsection and our empirical analysis in Section 5.1.

### 3.1 Twin-Gated Mechanism

Unlike GRU, we use an addition and subtraction operation over the transformed current input  $\mathbf{q}_t$  and history  $\mathbf{p}_t$  to differentiate the two gates in ATR. As the two gates have the same weights for their input components with only a single difference in the operation between the input components, they act like twins. We term the two gates in ATR as twin gates and the procedure, shown in Equation (7&8), as the *twin-gated mechanism*. This mechanism endows our model with the following two advantages: 1) Both addition and subtraction operations are completely linear so that fast computation can be expected; and 2) No other weight parameters are introduced for gates so that our model is more memory-compact.

A practical question for the twin-gated mechanism is whether twin gates are really capable of dynamically weighting the input and history information. To this end, we plot the surface of one-dimensional  $\sigma(x+y) - \sigma(x-y)$  in Figure 2. It is clear that both gates are highly non-linearly correlated, and that there are regions where  $\sigma(x+y)$  is equal to, greater or smaller than  $\sigma(x-y)$ . In other words, by adapting the distribution of input

Model	# WM	# MT
LSTM	8	8
GRU	6	6
RAN	4	4
ATR	2	2

Table 1: Comparison of LSTM, GRU, RAN and ATR in terms of the number of weight matrices (WM) and matrix transformations (MT).

and forget gates, the twin-gated mechanism has the potential to automatically seek suitable regions in Figure 2 to control its preference between the new and past information. We argue that the input and forget gates are negatively correlated after training, and empirically show their actual correlation in Section 5.1.

### 3.2 Computation Analysis

Here we provide a systematical comparison of computations in LSTM, GRU, RAN and our ATR with respect to the number of weight matrices and matrix transformations. Notice that all these units are building blocks of RNNs so that the total computational complexity and the minimum number of sequential operations required are unchanged, i.e.  $\mathcal{O}(n \cdot d^2)$  and  $\mathcal{O}(n)$  respectively where  $n$  is the sequence length and  $d$  is the dimensionality of hidden states. However, the actual number of matrix transformations in the unit indeed significantly affects the running speed of RNN in practice.

We summarize the results in Table 1. LSTM contains three different gates and a cell state, including 4 different neural layers with 8 weight matrices and transformations. GRU simplifies LSTM by removing a gate, but still involves two gates and a candidate hidden state. It includes 3 different neural layers with 6 weight matrices and transformations. RAN further simplifies GRU by removing the non-linearity in the state transition and therefore contains 4 weight matrices in its simplest version. Although our ATR also has two gates, however, there are only 2 weight matrices and transformations, accounting for only a third and a quarter of those in GRU and LSTM respectively. To the best of our knowledge, ATR has the smallest number of weight transformations in existing gated RNN units. We provide a detailed and empirical analysis on the speed in Section 5.2.



### 3.3 Interpretability Analysis of Hidden States

An appealing property of the proposed ATR is its interpretability. This can be demonstrated by rolling out Equation (9) as follows:

$$\begin{aligned} \mathbf{h}_t &= \mathbf{i}_t \odot \mathbf{q}_t + \mathbf{f}_t \odot \mathbf{h}_{t-1} \\ &= \mathbf{i}_t \odot \mathbf{W}_t \mathbf{x}_t + \sum_{k=1}^{t-1} \mathbf{i}_k \odot \left( \prod_{l=1}^{t-k} \mathbf{f}_{k+l} \right) \odot \mathbf{W}_x \mathbf{x}_k \\ &\approx \sum_{k=1}^t \mathbf{g}_k \odot \mathbf{W}_x \mathbf{x}_k \end{aligned} \quad (10)$$

where  $\mathbf{g}_k$  can be considered as an approximate weight assigned to the  $k$ -th input. Similar to the RAN model (Lee et al., 2017), the hidden state in ATR is a component-wise weighted sum of the inputs. This not only enables ATR to build up essential dependencies between preceding inputs and the current hidden state, but also allows us to easily detect which previous words have the promising impacts on the current state. This desirable property obviously makes ATR highly interpretable.

Additionally, this form of weighted sum is also related to self-attention (Vaswani et al., 2017). It can be considered as a forward unnormalized self-attention where each hidden state attends to all its previous positions. As the self-attention mechanism has proved very useful in NMT (Vaswani et al., 2017), we conjecture that such property of ATR partially contributes to its success in machine translation as shown in our experiments. We visualize the dependencies captured by Equation (10) in Section 5.3.

## 4 Experiments

### 4.1 Setup

We conducted our main experiments on WMT14 English-German and English-French translation tasks. Translation quality is measured by case-sensitive BLEU-4 metric (Papineni et al., 2002). Details about each dataset are as follows:

**English-German** To compare with previous reported results (Luong et al., 2015b; Jean et al., 2015; Zhou et al., 2016; Wang et al., 2017a), we used the same training data of WMT 2014, which consist of 4.5M sentence pairs. We used the newstest2013 as our dev set, and the newstest2014 as our test set.

**English-French** We used the WMT 2014 training data. This corpora contain 12M selected sentence pairs. We used the concatenation of newstest2012 and newstest2013 as our dev set, and the newstest2014 as our test set.

The used NMT system is an attention-based encoder-decoder system, which employs a bidirectional recurrent network as its encoder and a two-layer hierarchical unidirectional recurrent network as its decoder, companied with an additive attention mechanism (Bahdanau et al., 2015). We replaced the recurrent unit with our proposed ATR model. More details are given in Appendix A.1.

We also conducted experiments on Chinese-English translation, natural language inference and Chinese word segmentation. Details and experiment results are provided in Appendix A.2.

### 4.2 Training

We set the maximum length of training instances to 80 words for both English-German and English-French task. We used the *byte pair encoding* compression algorithm (Sennrich et al., 2016) to reduce the vocabulary size as well as to deal with the issue of rich morphology. We set the vocabulary size of both source and target languages to 40K for all translation tasks. All out-of-vocabulary words were replaced with a token “unk”.

We used 1000 hidden units for both encoder and decoder. All word embeddings had dimensionality 620. We initialized all model parameters randomly according to a uniform distribution ranging from -0.08 to 0.08. These tunable parameters were then optimized using Adam algorithm (Kingma and Ba, 2015) with the two momentum parameters set to 0.9 and 0.999 respectively. Gradient clipping 5.0 was applied to avoid the gradient explosion problem. We trained all models with a learning rate  $5e^{-4}$  and batch size 80. We decayed the learning rate with a factor of 0.5 between each training epoch. Translations were generated by a beam search algorithm that was based on log-likelihood scores normalized by sentence length. We used a beam size of 10 in all the experiments. We also applied dropout for English-German and English-French tasks on the output layer to avoid over-fitting, and the dropout rate was set to 0.2.

To train deep NMT models, we adopted the GNMT architecture (Wu et al., 2016). We kept all the above settings, except the dimensionality

System	Architecture	Vocab	tok BLEU	detok BLEU
Buck et al. (2014)	WMT14 winner system phrase-based + large LM	-	-	20.70
<i>Existing deep NMT systems (perhaps different tokenization)</i>				
Zhou et al. (2016)	LSTM with 16 layers + F-F connections	160K	20.60	-
Lei and Zhang (2017)	SRU with 10 layers	50K	20.70	-
Antonino and Federico (2018)	SR-NMT with 4 layers	32K	23.32	-
Wang et al. (2017a)	GRU with 4 layers + LAU + PosUnk	80K	23.80	-
Wang et al. (2017a)	GRU with 4 layers + LAU + PosUnk + ensemble	80K	26.10	-
Wu et al. (2016)	LSTM with 8 layers + RL-refined WPM	32K	24.61	-
Wu et al. (2016)	LSTM with 8 layers + RL-refined ensemble	80K	26.30	-
Vaswani et al. (2017)	Transformer with 6 layers + base model	37K	27.30	-
<i>Comparable NMT systems (the same tokenization)</i>				
Luong et al. (2015a)	LSTM with 4 layers + local att. + unk replace	50K	20.90	-
Zhang et al. (2017a)	GRU with gated attention + BPE	40K	23.84	-
Gehring et al. (2017b)	CNN with 15 layers + BPE	40K	25.16	-
Gehring et al. (2017b)	CNN with 15 layers + BPE + ensemble	40K	26.43	-
Zhang et al. (2018a)	Transformer with 6 layers + aan + base model	32K	26.31	-
<i>Our end-to-end NMT systems</i>				
<i>this work</i>	RNNSearch + GRU + BPE	40K	22.54	22.06
	RNNSearch + LSTM + BPE	40K	22.96	22.39
	RNNSearch + RAN + BPE	40K	22.14	21.40
	RNNSearch + ATR + BPE	40K	22.48	21.99
	RNNSearch + ATR + CA + BPE	40K	23.31	22.70
	GNMT + ATR + BPE	40K	24.16	23.59
	RNNSearch + ATR + CA + BPE + ensemble	40K	24.97	24.33

Table 2: Tokenized (*tok*) and detokenized (*detok*) case-sensitive BLEU scores on the WMT14 English-German translation task. “*unk replace*” and “*PosUnk*” denotes the approach of handling rare words in Jean et al. (2015) and Luong et al. (2015a) respectively. “*RL*” and “*WPM*” is the reinforcement learning optimization and word piece model used in Wu et al. (2016). “*CA*” is the context-aware recurrent encoder (Zhang et al., 2017b). “*LAU*” and “*F-F*” denote the linear associative unit and the fast-forward architecture proposed by Wang et al. (2017a) and Zhou et al. (2016) respectively. “*aan*” denotes the average attention network proposed by Zhang et al. (2018a).

of word embedding and hidden state which we set to be 512.

### 4.3 Results on English-German Translation

The translation results are shown in Table 2. We also provide results of several existing systems that are trained with comparable experimental settings to ours. In particular, our single model yields a detokenized BLEU score of 21.99. In order to show that the proposed model can be orthogonal to previous methods that improve LSTM/GRU-based NMT, we integrate a single-layer context-aware (CA) encoder (Zhang et al., 2017b) into our system. The ATR+CA system further reaches 22.7 BLEU, outperforming the winner system (Buck et al., 2014) by a substantial improvement of 2 BLEU points. Enhanced with the deep GNMT architecture, the GNMT+ATR system yields a gain of 0.89 BLEU points over the RNNSearch+ATR+CA and 1.6 BLEU points over the RNNSearch + ATR. Notice that different from our system which was trained on the parallel corpus alone, the winner system used a huge mono-

lingual text to enhance its language model.

Compared with the existing LSTM-based (Luong et al., 2015a) deep NMT system, our shallow/deep model achieves a gain of 2.41/3.26 tokenized BLEU points respectively. Under the same training condition, our ATR outperforms RAN by a margin of 0.34 tokenized BLEU points, and achieves competitive results against its GRU/LSTM counterpart. This suggests that although our ATR is much simpler than GRU, LSTM and RAN, it still possesses strong modeling capacity.

In comparison to several advanced deep NMT models, such as the Google NMT (8 layers, 24.61 tokenized BLEU) (Wu et al., 2016) and the LAU-connected NMT (4 layers, 23.80 tokenized BLEU) (Wang et al., 2017a), the performance of our shallow model (23.31) is competitive. Particularly, when replacing LSTM in the Google NMT with our ATR model, the GNMT+ATR system achieves a BLEU score of 24.16, merely 0.45 BLEU points lower. Notice that although all systems use the same training data of WMT14, the

System	Architecture	Vocab	tok BLEU	detok BLEU
<i>Existing end-to-end NMT systems</i>				
Jean et al. (2015)	RNNSearch (GRU) + unk replace + large vocab	500K	34.11	-
Luong et al. (2015b)	LSTM with 6 layers + PosUnk	40K	32.70	-
Sutskever et al. (2014)	LSTM with 4 layers	80K	30.59	-
Shen et al. (2016)	RNNSearch (GRU) + MRT + PosUnk	30K	34.23	-
Zhou et al. (2016)	LSTM with 16 layers + F-F connections + 36M data	80K	37.70	-
Wu et al. (2016)	LSTM with 8 layers + RL-refined WPM + 36M data	32K	38.95	-
Wang et al. (2017a)	RNNSearch (GRU) with 4 layers + LAU	30K	35.10	-
Gehring et al. (2017a)	Deep Convolutional Encoder 20 layers with kernel width 5	30K	35.70	-
Vaswani et al. (2017)	Transformer with 6 layers + 36M data + base model	32K	38.10	-
Gehring et al. (2017b)	ConvS2S with 15 layers + 36M data	40K	40.46	-
Vaswani et al. (2017)	Transformer with 6 layers + 36M data + big model	32K	41.80	-
Wu et al. (2016)	LSTM with 8 layers + RL WPM + 36M data + ensemble	32K	41.16	-
<i>Our end-to-end NMT systems</i>				
<i>this work</i>	RNNSearch + GRU + BPE	40K	35.89	33.41
	RNNSearch + LSTM + BPE	40K	36.95	34.15
	RNNSearch + ATR + BPE	40K	36.89	34.00
	RNNSearch + ATR + CA + BPE	40K	37.88	34.96
	GNMT + ATR + BPE	40K	38.59	35.67
	RNNSearch + ATR + CA + BPE + ensemble	40K	39.06	36.06

Table 3: Tokenized (*tok*) and detokenized (*detok*) case-sensitive BLEU scores on the WMT14 English-French translation task. “*12M data*” indicates the same training data as ours, while “*36M data*” is a significant larger dataset that contains the *12M data*.

tokenization of these work might be different from ours. However, the overall results can indicate the competitive strength of our model. In addition, SRU (Lei and Zhang, 2017), a recent proposed efficient recurrent unit, obtains a BLEU score of 20.70 with 10 layers, far more behind ATR’s.

We further ensemble eight likelihood-trained models with different random initializations for the ATR+CA system. The variance in the tokenized BLEU scores of these models is 0.07. As can be seen from Table 2, the ensemble system achieves a tokenized and detokenized BLEU score of 24.97 and 24.33 respectively, obtaining a gain of 1.66 and 1.63 BLEU points over the single model. The final result of the ensemble system, to the best of our knowledge, is a very promising result that can be reached by single-layer NMT systems on WMT14 English-German translation.

#### 4.4 Results on English-French Translation

Unlike the above translation task, the WMT14 English-French translation task provides a significant larger dataset. The full training data have approximately 36M sentence pairs, from which we only used 12M instances for experiments following previous work (Jean et al., 2015; Gehring et al., 2017a; Luong et al., 2015b; Wang et al., 2017a). We show the results in Table 3.

Our shallow model achieves a tokenized BLEU score of 36.89 and 37.88 when it is equipped

with the CA encoder, outperforming almost all the listed systems, except the Google NMT (Wu et al., 2016), the ConvS2S (Gehring et al., 2017b) and the Transformer (Vaswani et al., 2017). Enhanced with the deep GNMT architecture, the GNMT+ATR system reaches a BLEU score of 38.59, which beats the base model version of the Transformer by a margin of 0.49 BLEU points. When we use four ensemble models (the variance in the tokenized BLEU scores of these ensemble models is 0.16), the ATR+CA system obtains another gain of 0.47 BLEU points, reaching a tokenized BLEU score of 39.06, which is comparable with several state-of-the-art systems.

## 5 Analysis

### 5.1 Analysis on Twin-Gated Mechanism

We provide an illustration of the actual relation between the learned input and forget gate in Figure 3. Clearly, these two gates show strong negative correlation. When the input gate opens with high values, the forget gate prefer to be close. Quantitatively, on the whole test set, the Pearson’s *r* of the input and forget gate is -0.9819, indicating a high correlation.

### 5.2 Analysis on Speed and Model Parameters

As mentioned in Section 3.2, ATR has much fewer model parameters and matrix transformations. We

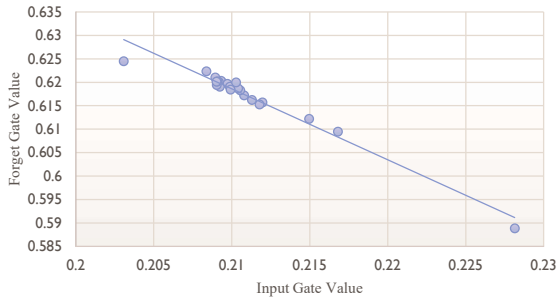


Figure 3: Visualization of correlation between the input and forget gate learned in Equation (13). For each gate, we record its mean value ( $\frac{1}{d_h} \sum_i \mathbf{i}_{t,i} / \frac{1}{d_h} \sum_i \mathbf{f}_{t,i}$ ) with its position on newstest2014, and then average all gate values for each position over all test sentences. This figure illustrates the position-wise input and forget gate values.

provide more details in this section by comparing against the following two NMT systems:

- *DeepRNNSearch (GRU)*: a deep GRU-equipped RNNSearch model (Wu et al., 2016) with 5 layers. We set the dimension of word embedding and hidden state to 620 and 1000 respectively.
- *Transformer*: a purely attentional translator (Vaswani et al., 2017). We set the dimension of word embedding and filter size to 512 and 2048 respectively. The model was trained with a minibatch size of 256.

We also compare with the GRU and LSTM-based RNNSearch. Without specific mention, all other experimental settings for all these models are the same as for our model. We implement all these models using the *Theano* library, and test the speed on one GeForce GTX TITAN X GPU card. We show the results on Table 4.

We observe that the Transformer achieves the best training speed, processing 4961 words per second. This is reasonable since the Transformer can be trained in full parallelization. On the contrary, DeepRNNSearch is the slowest system. As RNN performs sequentially, stacking more layers of RNNs inevitably reduces the training efficiency. However, this situation becomes the reverse when it comes to the decoding procedure. The Transformer merely generates 44 words per second while DeepRNNSearch reaches 70. This is because during decoding, all these beam search-

Model	#PMs	Train	Test
<i>RNNSearch+GRU</i>	83.5M	1996	168 (0.133)
<i>RNNSearch+LSTM</i>	93.3M	1919	167 (0.139)
<i>RNNSearch+RAN</i>	79.5M	2192	170 (0.129)
<i>DeepRNNSearch</i>	143.0M	894	70 (0.318)
<i>Transformer</i>	70.2M	4961	44 (0.485)
<i>RNNSearch+ATR</i>	67.8M	2518	177 (0.123)
<i>RNNSearch+ATR+CA</i>	63.1M	3993	186 (0.118)

Table 4: Comparison on the training and decoding speed and the number of model parameters of different NMT models on WMT14 English-German translation task with beam size 1. **#PMs**: the number of model parameters. **Train/Test**: the number of words in one second processed during training/testing. The number in bracket indicates the average decoding time per source sentence (in seconds).

based systems must generate translation one word after another. Therefore the parallelization advantage of the Transformer disappears. In comparison to DeepRNNSearch, the Transformer spends extra time on performing self-attention over all previous hidden states.

Our model with the CA structure, using only 63.1M parameters, processes 3993 words per second during training and generates 186 words per second during decoding, which yields substantial speed improvements over the GRU- and LSTM-equipped RNNSearch. This is due to the light matrix computation in recurrent units of ATR. Notice that the speed increase of ATR over GRU and LSTM does not reach 3x. This is because at each decoding step, there are mainly two types of computation: recurrent unit and softmax layer. The latter consumes the most calculation, which, however, is the same for different models (LSTM/GRU/ATR).

### 5.3 Analysis on Dependency Modeling

As shown in Section 3.3, a hidden state in our ATR can be formulated as a weighted sum of the previous inputs. In this section, we quantitatively analyze the weights  $\mathbf{g}_k$  in Equation (10) induced from Equation (13). Inspired by Lee et al. (2017), we visualize the captured dependencies of an example in Figure 4 where we connect each word to the corresponding previous word with the highest weight  $\mathbf{g}_k$ .

Obviously, our model can discover strong local dependencies. For example, the token “*unglück@@*” and “*lichen*” should be a






Figure 4: Visualization of dependencies on a target (German) sentence (selected from newstest2014). “@@” indicates a separator that splits one token into two pieces of sub-words.

single word. Our model successfully associates “*unglück@@*” closely to the generation of “*lichen*” during decoding. In addition, our model can also detect non-consecutive long-distance dependencies. Particularly, the prediction of “*Parteien*” relies heavily on the token “*unglücklichen*”, which actually entails an *amod* linguistic dependency relationship. These captured dependencies make our model more interpretable than LSTM/GRU.

## 6 Conclusion and Future Work

This paper has presented a twin-gated recurrent network (ATR) to simplify neural machine translation. There are only two weight matrices and matrix transformations in recurrent units of ATR, making it efficient in physical memory usage and running speed. To avoid the gradient vanishing problem, ATR introduces a twin-gated mechanism to generate an input gate and forget gate through linear addition and subtraction operation respectively, without introducing any additional parameters. The simplifications allow ATR to produce interpretable results.

Experiments on English-German and English-French translation tasks demonstrate the effectiveness of our model. They also show that ATR can be orthogonal to and applied with methods that improve LSTM/GRU-based NMT, indicated by the promising performance of the ATR+CA system. Further analyses reveal that ATR can be trained more efficiently than GRU. It is also able to transparently model long-distance dependencies.

We also adapt our ATR to other natural language processing tasks. Experiments show encouraging performance of our model on Chinese-English translation, natural language inference and Chinese word segmentation, demonstrating its generality and applicability on various NLP tasks.

In the future, we will continue to examine the effectiveness of ATR on different neural models for NMT, such as the hierarchical NMT model (Su et al., 2018b) as well as the generative NMT

model (Su et al., 2018a). We are also interested in adapting our ATR to summarization, semantic parsing etc.

## Acknowledgments

The authors were supported by National Natural Science Foundation of China (Grants No. 61672440, 61622209 and 61861130364), the Fundamental Research Funds for the Central Universities (Grant No. ZK1024), and Scientific Research Project of National Language Committee of China (Grant No. YB135-49). Biao Zhang greatly acknowledges the support of the Baidu Scholarship. We also thank the reviewers for their insightful comments.

## References

- M. Antonino and M. Federico. 2018. Deep Neural Machine Translation with Weakly-Recurrent Units. *ArXiv e-prints*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proc. of ICLR*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proc. of EMNLP*. Association for Computational Linguistics.
- Samuel R. Bowman, Jon Gauthier, Abhinav Rastogi, Raghav Gupta, Christopher D. Manning, and Christopher Potts. 2016. A fast unified model for parsing and sentence understanding. In *Proc. of ACL*, pages 1466–1477.
- James Bradbury, Stephen Merity, Caiming Xiong, and Richard Socher. 2016. Quasi-recurrent neural networks. *CoRR*, abs/1611.01576.
- Christian Buck, Kenneth Heafield, and Bas van Ooyen. 2014. N-gram counts and language models from the common crawl. In *Proc. of LREC*, pages 3579–3584, Reykjavik, Iceland.
- Xinchi Chen, Xipeng Qiu, Chenxi Zhu, Pengfei Liu, and Xuanjing Huang. 2015. Long short-term memory neural networks for chinese word segmentation. In *Proc. of EMNLP*, pages 1197–1206.

- Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. Long short-term memory-networks for machine reading. In *Proc. of EMNLP*, pages 551–561.
- Junyoung Chung, Çağlar Gülçehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR*.
- Jeffrey L Elman. 1990. Finding structure in time. *Cognitive science*, 14(2):179–211.
- Jonas Gehring, Michael Auli, David Grangier, and Yann N. Dauphin. 2017a. A convolutional encoder model for neural machine translation. In *Proc. of ACL*, pages 123–135.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017b. Convolutional sequence to sequence learning.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9:1735–1780.
- Sébastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On using very large target vocabulary for neural machine translation. In *Proc. of ACL-IJCNLP*, pages 1–10.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *Proc. of ICLR*.
- Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2017. Recurrent additive networks. *CoRR*, abs/1705.07393.
- T. Lei and Y. Zhang. 2017. Training RNNs as Fast as CNNs. *ArXiv e-prints*.
- Yang Liu, Chengjie Sun, Lei Lin, and Xiaolong Wang. 2016. Learning natural language inference using bidirectional LSTM model and inner-attention. *CoRR*, abs/1605.09090.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015a. Effective approaches to attention-based neural machine translation. In *Proc. of EMNLP*, pages 1412–1421.
- Thang Luong, Ilya Sutskever, Quoc Le, Oriol Vinyals, and Wojciech Zaremba. 2015b. Addressing the rare word problem in neural machine translation. In *Proc. of ACL-IJCNLP*, pages 11–19.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proc. of ACL*, pages 311–318.
- Wenzhe Pei, Tao Ge, and Baobao Chang. 2014. Max-margin tensor neural network for chinese word segmentation. In *Proc. of ACL*, pages 293–303, Baltimore, Maryland. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proc. of EMNLP*, pages 1532–1543.
- Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomas Kocisky, and Phil Blunsom. 2016. Reasoning about entailment with neural attention. In *Proc. of ICLR*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proc. of ACL*, pages 1715–1725.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Minimum risk training for neural machine translation. In *Proc. of ACL*, pages 1683–1692, Berlin, Germany. Association for Computational Linguistics.
- Richard Sproat and Thomas Emerson. 2003. The first international chinese word segmentation bake-off. In *Proceedings of the Second SIGHAN Workshop on Chinese Language Processing - Volume 17*, SIGHAN '03, pages 133–143.
- Jinsong Su, Shan Wu, Deyi Xiong, Yaojie Lu, Xianpei Han, and Biao Zhang. 2018a. Variational recurrent neural machine translation. *arXiv preprint arXiv:1801.05119*.
- Jinsong Su, Jiali Zeng, Deyi Xiong, Yang Liu, Mingxuan Wang, and Jun Xie. 2018b. A hierarchy-to-sequence attentional neural machine translation model. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 26(3):623–632.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2015. Order-embeddings of images and language. *CoRR*, abs/1511.06361.
- M. Wang, Z. Lu, J. Zhou, and Q. Liu. 2017. Deep Neural Machine Translation with Linear Associative Unit. *ArXiv e-prints*.
- Mingxuan Wang, Zhengdong Lu, Jie Zhou, and Qun Liu. 2017a. Deep neural machine translation with linear associative unit. In *Proc. of ACL*, pages 136–145, Vancouver, Canada.

- Shuohang Wang and Jing Jiang. 2016. Learning natural language inference with Istm. In *Proc. of NAACL*, pages 1442–1451.
- Zhiguo Wang, Wael Hamza, and Radu Florian. 2017b. Bilateral multi-perspective matching for natural language sentences. *CoRR*, abs/1702.03814.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144.
- Naiwen Xue, Fei Xia, Fu-dong Chiou, and Marta Palmer. 2005. The penn chinese treebank: Phrase structure annotation of a large corpus. *Nat. Lang. Eng.*, 11(2):207–238.
- Biao Zhang, Deyi Xiong, and Jinsong Su. 2017a. A gru-gated attention model for neural machine translation. *CoRR*, abs/1704.08430.
- Biao Zhang, Deyi Xiong, and Jinsong Su. 2018a. Accelerating neural transformer via an average attention network. In *Proc of ACL*, pages 1789–1798. Association for Computational Linguistics.
- Biao Zhang, Deyi Xiong, Jinsong Su, and Hong Duan. 2017b. A context-aware recurrent encoder for neural machine translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, PP(99):1–1.
- Jinchao Zhang, Mingxuan Wang, Qun Liu, and Jie Zhou. 2017c. Incorporating word reordering knowledge into attention-based neural machine translation. In *Proc. of ACL*, pages 1524–1534, Vancouver, Canada. Association for Computational Linguistics.
- Xiangwen Zhang, Jinsong Su, Yue Qin, Yang Liu, Rongrong Ji, and Hongji Wang. 2018b. Asynchronous bidirectional decoding for neural machine translation. *CoRR*, abs/1801.05122.
- Xiaoqing Zheng, Hanyang Chen, and Tianyu Xu. 2013. Deep learning for chinese word segmentation and pos tagging. In *Pro. of EMNLP*, pages 647–657.
- Zaixiang Zheng, Hao Zhou, Shujian Huang, Lili Mou, Xinyu Dai, Jiajun Chen, and Zhaopeng Tu. 2017. Modeling past and future for neural machine translation. *CoRR*, abs/1711.09502.
- Jie Zhou, Ying Cao, Xuguang Wang, Peng Li, and Wei Xu. 2016. Deep recurrent models with fast-forward connections for neural machine translation. *Transactions of the Association for Computational Linguistics*, 4:371–383.