

Tied Transformers: Neural Machine Translation with Shared Encoder and Decoder

Yingce Xia,¹ Tianyu He,² Xu Tan,¹ Fei Tian,¹ Di He,³ Tao Qin¹

¹Microsoft Research, ²University of Science and Technology of China

³Key Laboratory of Machine Perception, MOE, School of EECS, Peking University
{yinxia,xuta,fetia,taoqin}@microsoft.com, hetianyu@mail.ustc.edu.cn, di_he@pku.edu.cn

Abstract

Sharing source and target side vocabularies and word embeddings has been a popular practice in neural machine translation (briefly, NMT) for similar languages (e.g., English to French or German translation). The success of such word-level sharing motivates us to move one step further: we consider model-level sharing and tie the whole parts of the encoder and decoder of an NMT model. We share the encoder and decoder of *Transformer* (Vaswani et al. 2017), the state-of-the-art NMT model, and obtain a compact model named *Tied Transformer*. Experimental results demonstrate that such a simple method works well for both similar and dissimilar language pairs. We empirically verify our framework for both supervised NMT and unsupervised NMT: we achieve a 35.52 BLEU score on IWSLT 2014 German to English translation, 28.98/29.89 BLEU scores on WMT 2014 English to German translation without/with monolingual data, and a 22.05 BLEU score on WMT 2016 unsupervised German to English translation.

1 Introduction

Neural machine translation (briefly, NMT), based on the encoder-decoder framework with an attention module, has made significant progress in recent years (Vaswani et al. 2017; Hassan et al. 2018; He et al. 2016; Xia et al. 2017b). A common practice for training an NMT system is to share source and target side vocabularies, especially when the two languages are similar or in the same language family, like English↔French translation (Gehring et al. 2017; Vaswani et al. 2017). Such a shared vocabulary is built upon subword units like BPE (Sennrich, Haddow, and Birch 2016b) or word pieces (Wu et al. 2016), and both source and target words are split into smaller segmentations in this shared vocabulary. Acting in this way, the words in source sentences and target sentences are mapped to the same embedding space, which strengthens the semantic correlation between the two languages, and regularizes a neural network model with high capacity.

Parameter sharing has a long history in machine learning and has been applied in different learning settings such as multi-task learning (Ruder 2017; Zhang and Yang 2017), transfer learning (Pan and Yang 2010), meta-learning (Pham

et al. 2018), etc. Parameter sharing means that multiple models are bridged by tying parts of or all the model parameters. It has multiple potential benefits, such as reducing the number of model parameters so as to control model complexity (Firat et al. 2016), introducing prior knowledge to regularize models (Xia et al. 2018), and saving the storage space or memory size (Johnson et al. 2017). Parameter sharing is also adopted within network structure, including spatial sharing for the convolutional kernels in convolutional neural networks and temporal sharing for the transition functions in recurrent neural networks (Goodfellow et al. 2016). Inspired by the success of sharing vocabularies (and so word embeddings) of source and target languages, a question naturally comes out: Can we go one step forward to further share the whole parts of the encoder and decoder of an NMT model?

We make an initial attempt to answer this question and then propose *tied transformer*. We cast the typical encoder-decoder based sequence-to-sequence model into a more compact one in that there is only one copy of parameter set, which is applicable to both the encoder and decoder. In that way we force the sharing among the weights of the encoder and the decoder, rather than only among the source-side and target-side word embeddings.

We conduct extensive experiments to test the effectiveness of our proposed model. (1) For the translation between similar languages (i.e., languages within the similar language family), like {German, Spanish, Romanian}-from/to-English, our model achieves promising even state-of-the-art results: We achieve 35.52 for IWSLT German to English translation (see Figure 2), 28.98/29.89 for WMT 2014 English to German translation without/with monolingual data (see Table 4), and 34.67 for WMT 2016 English to Romanian translation (see Table 5). (2) For the translation of dissimilar languages (e.g., languages in different language families that cannot share vocabularies such as {Russian, Hebrew, Chinese}-from/to-English), our model also obtains good improvements (See Table 7&8). (3) For unsupervised NMT where no bilingual data is available, on WMT 2016 German→English translation, we improve the BLEU score from 21.0 to 22.05 (see Table 9).

2 Background

In this section, we introduce the background of a general NMT model and the basic structure of Transformer.

2.1 NMT Architecture

The task of NMT is to translate a sentence $x \in \mathcal{X}$ in the source language to the one $y \in \mathcal{Y}$ in the target language, where \mathcal{X} and \mathcal{Y} are two language spaces. An NMT model consists of an encoder f_E , a decoder f_D and a source-to-target attention model f_A . A source sentence x is first encoded into a series of hidden representations by using f_E , that is, $(h_1, h_2, \dots, h_{T_x}) = f_E(x)$, where T_x is the length of x , h_i is the i 'th hidden representation $\forall i \in [T_x]$. Then, when translating the j 'th word y_j in the target side, the attention model f_A is used to calculate the context vector of the encoder's outputs, that is,

$$c_j = \sum_{i=1}^{T_x} \alpha_{i,j} h_i, \quad \alpha_{i,j} = f_A(h_i, s_{j-1}), \quad (1)$$

where $\alpha_{i,j}$'s have to satisfy $\sum_{i=1}^{T_x} \alpha_{i,j} = 1$ and $\alpha_{i,j} \geq 0$. Next, based on the context vector c_j , we calculate the j 's hidden representation at target side by $s_j = f_D(s_{j-1}, y_{j-1}, c_j)$, where y_{j-1} is the $(j-1)$ 'th target word. Finally, s_j is mapped to $y_j \in \mathcal{Y}$ by a simple network. The f_E, f_D can be specialized by LSTM (Wu et al. 2016), GRU (Bahdanau, Cho, and Bengio 2015), CNN (Gehring et al. 2017), Transformer (Vaswani et al. 2017) and so on. f_A can be implemented by a bilinear function like $\alpha_{i,j} \propto \exp(h_i W_A s_{j-1}) \forall i \in [T_x]$, a sum of two affine transformations like $\alpha_{i,j} \propto \exp(W_{A,h} h_i + W_{A,s} s_{j-1}) \forall i \in [T_x]$ where the W 's are the parameters to be learned, and so on. Luong, Pham, and Manning (2015) summarize different attention models.

2.2 Parameter Sharing in NMT

Several parameter sharing mechanisms have been explored for multilingual NMT. Dong et al. (2015) leveraged a network with one encoder and multiple decoders for one-to-many language translations. Luong et al. (2016) proposed a structure with multiple encoders and one decoder for many-to-one language translation. Zoph and Knight (2016) targeted at a multi-source translation problem, where the decoder is shared. Firat, Cho, and Bengio (2016) designed a network with multiple encoders and decoders plus a shared attention mechanism across different language pairs for many-to-many language translation. (Ha, Niehues, and Waibel 2016; Johnson et al. 2017) used a single encoder-decoder model to work on many-to-many languages translation. While there exist quite a few papers studying parameter sharing for multiple language translation, there are few work studying parameter sharing in one task and we attack this problem.

2.3 Brief Introduction of Transformer

Transformer (Vaswani et al. 2017) is a general framework built upon self-attention that achieves state-of-the-art results on many translation tasks. Transformer is a stacked network with several blocks. Each block consists of two or three basic modules:

(1) A self-attention module φ_S , used to output a weighted version of its inputs. That is, given an m -element set $Z = \{z_1, z_2, \dots, z_m\}$, for any $i \in [m]$, $z_i^S = \varphi_S(z_i, Z) =$

$\sum_{j=1}^m \alpha_{i,j} z_j$, where $\alpha_{i,j}$ is calculated by the multi-head attention as that proposed in Vaswani et al. (2017) with z_i and Z as inputs. Z can be a collection of either source side hidden states or target side hidden states. This layer is associated with a residual connection and layer normalization LN (Ba, Kiros, and Hinton 2016), whose eventual output is $\text{LN}(z_i + z_i^S)$ for any $i \in [m]$.

(2) An optional cross-lingual (i.e., source-to-target) attention model φ_C that will appear in the decoder only. Similar to φ_S , set $H = \{h_1, h_2, \dots, h_{T_x}\}$, T_x is the length of source sentence x , H is the output of the last layer in the encoder, s_i is the i 'th hidden state in the decoder, $s_i^C = \varphi_C(s_i, H)$, where $\varphi_C(s_i, H)$ is executed in a similar way as the φ_S in step (1) with s_i and H as inputs. This layer will eventually outputs $\text{LN}(s_i + s_i^C)$. Both φ_S and φ_C are implemented as multi-head attention models and details can be found in Vaswani et al. (2017).

(3) A feed-forward network $\varphi_F(z) = W_2 \max(W_1 z + b_1, 0) + b_2$, where the W 's and b 's are the parameters to be learned. Again, this layer is followed with a residual connection and layer normalization.

3 Our Model

In this section, we introduce the architecture of our proposed tied transformer. We describe the model architecture in Section 3.1 and give a theoretical discussion in Section 3.2.

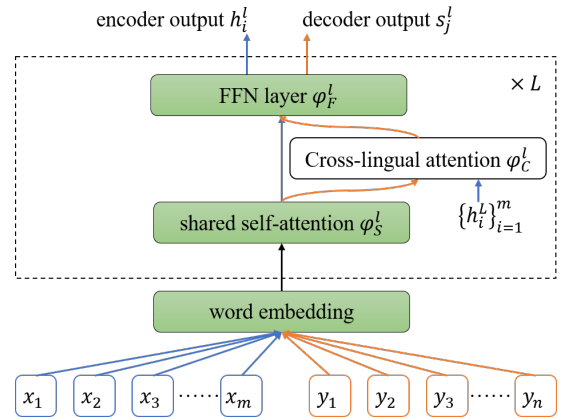


Figure 1: The architecture of tied transformer

3.1 Model Architecture

The architecture of tied transformer is shown in Figure 1. We follow the notations defined in Section 2.3 to mathematically describe our model. Tied transformer is a stacked model with L blocks, where the l 'th block consists of a self-attention module φ_S^l , a cross-lingual attention φ_C^l and a non-linear function φ_F^l , where superscript l represents the layer id. We show how the encoder and decoder are shared.

(1) *Encoding* For any input x , we initialize $H^0 = \{h_1^0, h_2^0, \dots, h_{T_x}^0\}$, where h_i^0 is the word embedding of the i 'th word. The encoding process can be formulated as fol-

lows: For any $l \in [L]$,

$$\begin{aligned} h_i^l &= \varphi_F^l(\varphi_S^l(h_i^{l-1}, H^{l-1})); \\ H^l &= \{h_1^l, h_2^l, \dots, h_{T_X}^l\}. \end{aligned} \quad (2)$$

Eventually, we obtain H^L , a set containing all the outputs of the last layer in the encoder.

(2) *Decoding* At the $(j + 1)$ 'th decoding step, denote the generated hidden states at the target side as $S_j^l = \{s_1^l, s_2^l, \dots, s_j^l\}$, $l \in \{0, 1, \dots, L\}$ where S_j^0 represents the collection of word embeddings at target side. We first set $s_{j+1}^0 = y_j$ and initialize $S_{j+1}^0 = S_j^0 \cup \{y_j\}$, where y_j is the word embedding of the predicted j 'th word y_j . Then, for any $l \in [L]$,

$$\begin{aligned} s_{j+1}^l &= \varphi_F^l\left(\varphi_S^l(\varphi_S^l(s_{j+1}^{l-1}, S_{j+1}^{l-1}), H^L)\right), \\ S_{j+1}^l &= S_j^l \cup \{s_{j+1}^l\}. \end{aligned} \quad (3)$$

Once we obtain s_{j+1}^L , we could use a simple network with a softmax layer to predict the corresponding word y_{j+1} . We repeat decoding until meeting the end-of-sentence token.

Note that given any $l \in [L]$, the φ_S^l and φ_F^l in Eqn.(2) and Eqn.(3) are shared; however, in a standard encoder-decoder based network, such two components are not tied. This shows that our proposed model is strongly regularized during training.

Note that such a parameter sharing scheme could also be applied to LSTM based encoder-decoder frameworks. An empirical analysis is illustrated in Section 4.4.

3.2 Theoretical Discussion

We provide a brief theoretical discussion about the proposed tied transformer. Denote the parameters of the encoder, decoder and source-to-target attention modules as θ_e, θ_d and θ_a respectively. For a tied transformer, $\theta_e = \theta_d$. In an NMT task, what we learn is a mapping $f : \mathcal{X} \mapsto \mathcal{Y}$, that can translate sentences from source language space \mathcal{X} to target language space \mathcal{Y} . Our objective is to minimize the (expected) risk of a model f , which is defined as follows:

$$R(f) = \mathbb{E}_{(x,y) \sim \mathcal{P}} [\ell(f(x), y)]. \quad (4)$$

where \mathcal{P} represents the underlying distribution of a source and target language sentence pair over $\mathcal{X} \times \mathcal{Y}$, ℓ is the loss function that can achieve the mapping $\mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$. Denote Θ_e, Θ_d and Θ_a as the parameter spaces of the encoder, decoder and source-to-target attention model. For the conventional transformer, $f \in \mathcal{F}_c$ where $\mathcal{F}_c = \{f(x; \theta_e, \theta_d, \theta_a) | \theta_e \in \Theta_e, \theta_d \in \Theta_d, \theta_a \in \Theta_a\}$. For the tied transformer, $f \in \mathcal{F}_t$ where $\mathcal{F}_t = \{f(x; \theta_e, \theta_d, \theta_a) | \theta_e \in \Theta_e, \theta_d \in \Theta_d, \theta_a \in \Theta_a, \theta_e = \theta_d\}$. Obviously, $\mathcal{F}_t \subseteq \mathcal{F}_c$.

When training an NMT model, we minimize the empirical risk on the n training samples $\{(x_i, y_i)\}_{i=1}^n$, $x_i \in \mathcal{X}$, $y_i \in \mathcal{Y}$, which is defined as follows: for any $f \in \mathcal{F}_c$ or $f \in \mathcal{F}_t$,

$$R_n(f) = \frac{1}{n} \sum_{i=1}^n \ell(f(x_i), y_i).$$

We introduce Rademacher complexity (Mohri et al., 2012) for our proposed method, a measure for the complexity of the function space.

Definition 1 Denote the Rademacher complexity of the standard transformer and the tied transformer as \mathfrak{R}_n^c and \mathfrak{R}_n^t respectively. We have

$$\begin{aligned} \mathfrak{R}_n^c &= \mathbb{E}_{z, \sigma} \left[\sup_{f \in \mathcal{F}_c} \frac{1}{n} \sum_{i=1}^n \sigma_i \ell(f(x_i), y_i) \right], \\ \mathfrak{R}_n^t &= \mathbb{E}_{z, \sigma} \left[\sup_{f \in \mathcal{F}_t} \frac{1}{n} \sum_{i=1}^n \sigma_i \ell(f(x_i), y_i) \right], \end{aligned} \quad (5)$$

where $z = \{z_1, z_2, \dots, z_n\} \sim \mathcal{P}^n$, $z_i = (x_i, y_i)$ in which $x_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$, $\sigma = \{\sigma_1, \dots, \sigma_n\}$ are i.i.d sampled with $P(\sigma_i = 1) = P(\sigma_i = -1) = 0.5$.

The following theorem holds for our proposed method:

Theorem 1 (Theorem 3.1, (Mohri et al., 2012)) Let $\ell(f(x), y)$ be a mapping from $\mathcal{X} \times \mathcal{Y}$ to $[0, 1]$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the following inequalities holds:

$$\begin{aligned} R(f) &\leq R_n(f) + 2\mathfrak{R}_n^c + \sqrt{\frac{1}{2n} \ln\left(\frac{1}{\delta}\right)}, \quad \forall f \in \mathcal{F}_c, \\ R(f) &\leq R_n(f) + 2\mathfrak{R}_n^t + \sqrt{\frac{1}{2n} \ln\left(\frac{1}{\delta}\right)}, \quad \forall f \in \mathcal{F}_t. \end{aligned}$$

Since $\mathcal{F}_t \subseteq \mathcal{F}_c$, we have that $\mathfrak{R}_n^t \leq \mathfrak{R}_n^c$, which shows that tied transformers have better generalization abilities than standard transformer.

4 Experiments on Similar Languages

In this section, we verify our proposed method on languages in similar language families. We work on various amounts of training data, including IWSLT translation tasks (less than 200k sentences pairs), WMT translation tasks (more than 2M sentence pairs) and the usage of monolingual data (additional 8M unlabeled sentences).

4.1 Settings

We choose IWSLT 2014 {German, Spanish, Romanian}-from/to-English translation and WMT 2014 English-to-{German, Romanian} to verify our proposed algorithm. We briefly denote German, Romanian, Spanish and English as ‘‘De’’, ‘‘Ro’’, ‘‘Es’’ and ‘‘En’’ respectively.

Datasets For IWSLT 2014 De \leftrightarrow En, Es \leftrightarrow En and Ro \leftrightarrow En translation tasks, there are respectively 153k, 181k, 182k sentence pairs in each datasets¹. For De \rightarrow En translation, we follow the common practice to lowercase all words; and we use the same validation and test sets as those used in (Edunov et al. 2018), which consists of 7k and 7k sentences respectively. For the other translation tasks, we do not lowercase the words; we use IWSLT14.TED.tst2013 as the validation sets and IWSLT14.TED.tst2014 as test sets.

For WMT 2014 En \rightarrow De and WMT 2016 En \rightarrow Ro translation, there are 4.5M and 2.8M bilingual sentence pairs respectively. For En \rightarrow De translation, we concatenate newstest2012 and newstest2013 as the validation set and use

¹All IWSLT 2014 training data can be found at <https://wit3.fbk.eu/archive/2014-01/texts>

newstest2014 as the test set. For En→Ro, we use news-dev2016 as the validation set and use newstest2016 as the test. For both IWSLT and WMT translation tasks, all the datasets are preprocessed into workpieces following (Wu et al. 2016).

Model Configurations For the IWSLT 2014 translation tasks, we choose the *transformer_small* setting with 8 blocks, where each block contains a self-attention layer, an optional encoder-to-decoder attention layer and a feed-forward layer. The word embedding dimension, hidden state dimension, non-linear layer dimension and the number of head are 256, 256 and 1024 and 4 respectively. For WMT 2014 En→De translation and WMT 2016 En→Ro task, we choose the *transformer_big* setting, where the four numbers are 1024, 1024, 4096 and 16 respectively. The dropout rate for the two settings are 0.1 and 0.2. For the three IWSLT 2014 translation tasks, we train each model on two V100 GPUs for up to three days until convergence and the minibatch size is fixed as 4096 tokens per GPU, including both the source to target and target to source data. For the two WMT translation tasks, we train each model on four P40 GPUs for ten days until convergence. We try to use the largest possible minibatch sizes for the WMT settings to fulfill GPU memory limitation.

Evaluation. Following (Wu et al. 2016; Vaswani et al. 2017), the eventually output is picked up by $\operatorname{argmax}_{y' \in \mathcal{C}} \log P(y'|x)/|y'|^\alpha$, where x denotes the source sentence, \mathcal{C} represents the candidates carried out by beam search, $|y'|$ denotes the number of words in y' and α is the length penalty tuned by the validation set. For IWSLT related tasks, we use beam search with beam width 6 and $\alpha = 1.1$ to generate candidates. For WMT related tasks, the beam size is changed to 4 and the α is chosen as 0.6 for a fair comparison with (Vaswani et al. 2017). We use the BLEU scores as the evaluation metrics, which is the geometric mean of four n -gram precisions between the translation results and reference sentences, $n = 1, 2, 3, 4$. To be more specific, (1) For IWSLT De→En translation and WMT En→De translations, we calculate the tokenized BLEU by `multi-bleu.perl`² for fair comparisons with previous methods. (2) For other IWSLT translation tasks, we use `sacreBLEU`³. (3) For En→Ro, we use detokenized BLEU for fair comparison with previous methods⁴.

4.2 Results on IWSLT Related Tasks

The experimental results of IWSLT 2014 translation tasks are shown in Table 1. The three rows in Table 1 represent the basic Transformer algorithm, our proposed algorithm and the improvements brought by our algorithm. We also list the IWSLT 2014 De→En results reported by previous literature.

The tied transformer significantly outperforms the standard transformer. The improvements are two-folded: first,

²<https://github.com/moses-smt/mosesdecoder/blob/master/scripts/generic/multi-bleu.perl>

³<https://github.com/aws-labs/sockeye/tree/master/contrib/sacrebleu>

⁴<https://github.com/moses-smt/mosesdecoder/blob/master/scripts/tokenizer/detokenizer.perl>

	De→En	Es→En	Ro→En
<i>Transformer</i>	33.27	38.67	29.64
<i>Tied Transformer</i>	35.10	40.51	30.99
<i>Improvements</i>	1.83	1.84	1.35
Existing Results on IWSLT De→En			
GRU + Dual Transfer Learning (Wang et al. 2018)			32.35
CNN + Reinforcement learning (Edunov et al. 2018)			32.85
LSTM + Variational Attention (Deng et al. 2018)			33.10
Transformer + Model-level dual (Xia et al. 2018)			34.71

Table 1: Results of IWSLT 2014 {De, Es, Ro}-to-En translation tasks.

we could see *Transformer* outperforms non-transformer systems, which is helpful to get a great score; second, on top of such a strong baseline, we are still capable to further improve the performances, which demonstrate the effectiveness of our method. On the three tasks De→En, Es→En and Ro→En, our proposed method could outperform the Transformer baseline by 1.83, 1.84 and 1.35 points, which indeed verifies our motivation introduced in Section 1 that such a framework is helpful to improve the translations within a same language family. As far as we could survey, we achieve the best result on IWSLT 2014 De→En, whose previously best result in 33.81 provided by (Elbayad et al. 2018).

	En→De	En→Es	En→Ro
<i>Transformer</i>	27.72	37.51	23.80
<i>Tied Transformer</i>	29.07	38.81	24.31
<i>Improvements</i>	1.35	1.30	0.51

Table 2: Results of IWSLT 2014 En-to-{De, Es, Ro} translation tasks.

We also check the ability of translating from English to other languages. The results are shown in Table 2. Again, tying the weights of encoder and decoder outperforms the baseline. The results in Table 1 and Table 2 demonstrate that our proposed method works pretty good on the translation tasks within the same language family again.

Then, we compare our proposed tied transformer with two well-known baselines, which are both based on 8-block transformers:

1. Cross-tasking sharing (Johnson et al. 2017; Ha, Niehues, and Waibel 2016), briefly denoted as CTS, which means to use one encoder-decoder based network to solve two tasks together. In this framework, the encoder is used to encode two languages and the decoder is used to decode two languages.
2. Dual supervised learning (Xia et al. 2017a), briefly denoted as DSL, in which the training objective of De→En and En→De are jointly constrained by $P(x)P(y|x; \theta_{xy}) = P(y)P(x|y; \theta_{yx})$, where $x, y, \theta_{xy}, \theta_{yx}$ represent a German sentence, an English sentence, the parameters of De→En translation model and the parameters of En→De translation model respectively. We

list it as a baseline here considering we report the results of two dual tasks and they could be considered together.


The results are shown in Table 3. We have several observations: (1) Our method outperforms both the standard baseline and DSL, which shows that parameter sharing is more powerful in these two tasks. (2) Compared with the cross-task sharing, our method outperforms this baseline by 0.77 and 0.74 point, which shows that sharing parameter like ours is more useful than the others. Besides, it demonstrates that parameter sharing is indeed a useful technique in NMT for both cross-task sharing and our proposed sharing scheme, which can be seen as a kind of intra-task sharing, especially when the training data is limited.

	<i>Baseline</i>	<i>DSL</i>	<i>CTS</i>	<i>Ours</i>
De→En	33.27	33.60	34.33	35.10
En→De	27.72	27.88	28.33	29.07

Table 3: Comparison with several other baselines on IWSLT 2014 De↔En translation tasks.

Next, we explore how the test performances of our proposed framework change with different number of blocks and different hidden dimensions.

In Figure 2(a), we fix the hidden dimension as 256 and increase the block number from 6 to 14. We have the following observations: (1) our proposed algorithm always outperforms the baseline with different number of blocks, demonstrating our algorithm has better generalization ability. (2) Both our algorithm and the baseline achieve the best results at the 8-block setting. As we increase more blocks, the performances will drop, indicating that deeper models need stronger regularization.



(a) Different block numbers (b) Different hidden dimensions

Figure 2: IWSLT 2014 De→En BLEU with different model parameters.

In Figure 2(b), we fix the block number as 8 and vary the hidden dimensions. For the baseline, the BLEU starts to drop when the hidden dimension is larger than 256. In comparison, the BLEU of our algorithm starts to drop when the dimension is larger than 384, which shows that our algorithm has better generalization ability. Another observation is that our shared model cannot be too small: when the hidden dimension is 128, the BLEU score of our algorithm is 0.82 point lower than the baseline, which shows that the model does not have enough capacity to handle the task. Finally, by

setting the dimension as 384, we improve the BLEU score from 35.10 to 35.52, setting a new record on this task.

4.3 Results on WMT Related Tasks

To verify the performances of our algorithm on larger datasets, we move to WMT 2014 En→De and WMT 2016 En→Ro translation tasks.

In Table 4, we show the BLEU scores of En→De translation. In this case, we can see that a 6-block tied transformer (with BLEU 28.35) can almost catch up with the standard 6-block transformer (with BLEU 28.4) but using only around half of the parameters. This shows that even though the training corpus is large, our method is able to catch up with the baselines.

<i>8-layer LSTM + RL</i> (Wu et al. 2016)	24.60
<i>15-layer CNN network</i> (Gehring et al. 2017)	25.16
<i>6-block transformer</i> (Vaswani et al. 2017)	28.40
<i>6-block tied transformer</i>	28.35
<i>8-block tied transformer</i>	28.88
<i>12-block tied transformer</i>	28.98
<i>8-block tied transformer + 8M unlabeled data</i>	29.89

Table 4: WMT 2014 En→De BLEU.

Then we increase the number of blocks. When the network has 8 blocks, the BLEU score is 28.88, 0.48 point ahead of the baseline. When we increase it to 12 blocks, the BLEU score is 28.98, which achieves 0.58 point improvement over the baseline. Note that our model with 12 blocks has the approximate number of parameter with the baseline with 6 blocks.

We also check the ability of our model with the support of monolingual data. We use the back-translation technique to leverage monolingual data (Sennrich, Haddow, and Birch 2016a). To achieve that, we first use the 4.5M bilingual data to train an initial De→En translation model with 32.12 BLEU score. Then, we choose 8M monolingual German data from *newscrawl 2013*⁵. Next, we use the obtained De→En model to translate the 8M sentences with beam-size 4 and α 1.0. Finally, we merge the original bilingual dataset and the generated English-German pairs from the monolingual data and train another En→De model based on tied transformer. In this way, we obtain 29.89 BLEU score, which shows that our proposed tied model can handle large amount of training data.

The En→Ro translation results are summarized in Table 5. We implement the standard transformer as baseline and get a 33.05 BLEU score. Our method with 6 blocks can achieve 34.67 BLEU, which is a state-of-the-art result. We also tried the 8-block network but obtained a lower score. This is because the En→Ro task does not have as much training data as En→De and thus could not apply too large models.

⁵<http://www.statmt.org/wmt14/training-monolingual-news-crawl/news.2013.de.shuffled.gz>

20-layer CNN network (Gehring et al. 2017)	30.02
6-block transformer (Vaswani et al. 2017)	33.05
6-block tied transformer	34.67
8-block tied transformer	34.55

Table 5: WMT 2016 En→Ro BLEU.

4.4 Tied LSTM Models for NMT

To verify the generality of sharing the encoder and decoder, we adapt our framework into the LSTM and obtained a tied LSTM model. We follow (Wu et al. 2016) to build a deep LSTM NMT model but sharing the parameters of the encoder and the decoder (excluding the attention model).

We work on IWSLT 2014 De→En translation to verify our idea. For the baseline, we choose a 4-layer single-direction LSTM with vocabulary size $24k$, word embedding dimension 256 and hidden dimension 512. For the tied LSTM, we also fix the word embedding dimension as 256 and the number of layers as 4. We set the hidden dimensions as 512 and 1024, corresponding to almost 50% and 100% parameters of the baseline LSTM. We share the source embeddings, target embeddings and the softmax matrix. We use Adadelta (Zeiler 2012) to optimize the network. The results are shown in Table 6.

Baseline	hidden 512	hidden 1024
29.41	29.57	30.61

Table 6: IWSLT 2014 De→En with tied LSTMs

We can see the baseline achieves 29.41 BLEU score. After sharing the encoder and decoder, we obtain 0.16 gain. The improvement is not very significant due to the limited model capacity. When we increase the hidden dimension to 1024, we can achieve 30.61 BLEU score, which is 1.2 point improvement over the baseline. This shows that our proposed method can be generalized to more structures like LSTM. We leave how to apply our framework to more complex LSTMs as future work.

5 Experiments of Dissimilar Languages

In Section 4, we have verified the effectiveness of our proposed method for translating between languages within the same family. In this section, we analyze how our method performs to translating between languages in different families.

Settings. We choose IWSLT 2014 {Russian, Hebrew, Chinese}-from/to-English translation tasks to verify our proposed algorithm, where there are $160k$, $151k$, $182k$ sentence pairs in each datasets. Russian, Hebrew and Chinese are briefly denoted as “Ru, He, Zh” respectively. We use IWSLT14.TED.tst2013 as the validation sets and IWSLT14.TED.tst2014 as test sets. Again, all sentences are split into wordpieces by (Wu et al. 2016). We use the same model structure and the evaluation metrics as that used in Section 4 for IWSLT 2014 tasks.

Results. The experimental results are shown in Table 7. Our proposed method outperforms the Transformer baseline by 1.16, 0.97 and 0.87 points on the three tasks.

	Ru→En	He→En	Zh→En
<i>Transformer</i>	17.89	31.96	16.45
<i>Tied Transformer</i>	19.05	32.93	17.32
<i>Improvements</i>	1.16	0.97	0.87

Table 7: Experimental results on IWSLT 2014 {Ru, He, Zh}-to-English translation tasks.

The English to Russian, Hebrew and Chinese translation results are shown in Table 8. We could still make improvements on these tasks.

	En→Ru	En→He	En→Zh
<i>Transformer</i>	14.59	20.84	11.88
<i>Tied Transformer</i>	14.90	21.47	12.38
<i>Improvements</i>	0.31	0.63	0.50

Table 8: Experimental results on IWSLT 2014 English-to-{Ru, He, Zh} translation tasks.

As shown in Figure 3, an interesting observation is that, no matter for similar languages or different languages, overall, the improvement of $X \rightarrow \text{En}$ translation quality brought by our method is larger than the corresponding $\text{En} \rightarrow X$, where X is the language we used in the IWSLT datasets. We conjecture that English is an easier language to decode, which results in the aforementioned observation.

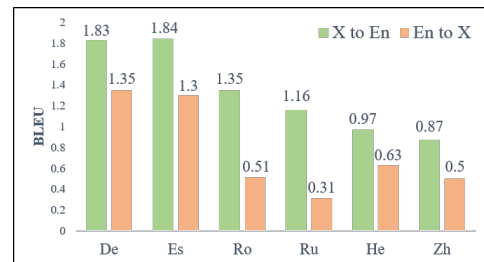


Figure 3: Improvements of our method compared with the Transformer baseline.

We also find that the improvements of translating within same language family are larger than that within different language families, no matter for X -to-English translation or English-to- X translation. Our explanation is that the languages within the same family are easily mapped into the same language space, considering their vocabulary are shared. Working on the same semantic space will result in better improvements than those in different spaces.

6 Experiments on Unsupervised NMT

Unsupervised NMT targets at making translation between two languages possibly without bilingual data. Several literature leverages parameter sharing and dual learning to

solve this problem (Lample, Denoyer, and Ranzato 2018; Lample et al. 2018), which motivates us to solve unsupervised NMT with a tied transformer.

6.1 Model Adaption

Background We briefly introduce (Lample et al. 2018), a recent state-of-the-art unsupervised NMT algorithm and then adapt it with our tied transformer. Like the standard NMT system, an unsupervised NMT model also consists of an encoder, a source-to-target attention module, and a decoder as well as a shared embedding for similar languages like German and English. The source sentences and target sentences are mapped into a same vocabulary using BPE techniques. The shared embedding is pretrained with fast-Text (Bojanowski et al. 2017) to get good initial values. The proposed model will handle both source-to-target translation and target-to-source translation like Johnson et al. (2017). The reason to use one model for two translation tasks is that a single model could better map two different languages into the same representation space, which is easier for decoding.

The training loss of an unsupervised NMT model usually consists of two parts, a language model loss and a back-translation loss. Let $P_{\mathcal{X} \rightarrow \mathcal{Y}}$ denote the translation from space \mathcal{X} to space \mathcal{Y} , and so for the other similar notations.

(1) Language model loss, implemented by a denoising autoencoder. Mathematically,

$$\mathcal{L}^{\text{lm}} = \mathbb{E}_{x \sim \mathcal{X}} [-\log P_{\mathcal{X} \rightarrow \mathcal{X}}(x|\sigma(x))] + \mathbb{E}_{y \sim \mathcal{Y}} [-\log P_{\mathcal{Y} \rightarrow \mathcal{Y}}(y|\sigma(y))], \quad (6)$$

where $\sigma(\cdot)$ is a noise model with randomly dropping several words, swapping words, etc.

(2) Back translation loss, implemented by back-translating the monolingual data and feeding into the reversed models. Mathematically,

$$\mathcal{L}^{\text{back}} = \mathbb{E}_{x \sim \mathcal{X}} [-\log P_{\mathcal{Y} \rightarrow \mathcal{X}}(x|\hat{y}(x))] + \mathbb{E}_{y \sim \mathcal{Y}} [-\log P_{\mathcal{X} \rightarrow \mathcal{Y}}(y|\hat{x}(y))], \quad (7)$$

in which $\hat{y}(x) = \arg \max_{u \in \mathcal{Y}} P_{\mathcal{X} \rightarrow \mathcal{Y}}(u|x)$ and $\hat{x}(y) = \arg \max_{w \in \mathcal{X}} P_{\mathcal{Y} \rightarrow \mathcal{X}}(w|y)$. Note that the four P_{\dots} 's are implemented in a single encoder-decoder based model, where each translation task has a different "task embedding", i.e., a learnable vector indicating the translating directions.

Adaption To achieve unsupervised NMT, the aforementioned four models P_{\dots} 's in Eqn. (6) and Eqn. (7) are implemented in a single tied transformer. The model architecture is the same as that introduced in Section 3, where each task has a task embedding as that used in Lample et al. (2018). We also use $\mathcal{L}^{\text{lm}} + \mathcal{L}^{\text{back}}$ as the training loss.

6.2 Settings

We work on German \leftrightarrow English translation. Following Lample et al. (2018), we collect 50M WMT monolingual data from newscrawl 2014 to newscrawl 2017. The BPE (Sennrich, Haddow, and Birch 2016b) with 60k merge operations is applied to pre-process the data. Newstest2014 and newstest2016 German \leftrightarrow English translation datasets are used as the validation set and test set respectively.

We follow the *transformer_base* configuration (Vaswani et al. 2017) to set the hyper-parameter of our model. The word embedding dimension, hidden state dimension, non-linear layer dimension and the number of heads of the multi-head attention are 512, 512 and 2048 and 8 respectively. We try different settings: 4 blocks and 8 blocks for the tied transformer. In the training phase, we fix the dropout rate as 0.2, apply Adam optimizer with learning rate 0.0002 to optimize the network. The model is implemented in TensorFlow and trained on 8 M40 GPUs. In the inference phase, we use beam search with beam width 4 and set length penalty α as 0.6.

6.3 Results

The results of unsupervised NMT are shown in Table 9. To speed up the training, we first train an initial unsupervised NMT model, back translate the monolingual data, train a warm-start tied transformer and then keep tuning the obtained model by online sampling. We can see that our proposed model performs best among all NMT based translation systems.

	BLEU
LSTM model (Lample, Denoyer, and Ranzato 2018)	13.3
4-block Transformer (Lample et al. 2018)	21.0
<i>Our method with 4 blocks</i>	21.61
<i>Our method with 8 blocks</i>	22.05

Table 9: Experimental results of unsupervised neural machine translation on WMT 2016 De \rightarrow En.

Specifically, when using the 4-block model, which accounts only half of the parameters used in (Lample et al. 2018), on De \rightarrow En translation, we can achieve 21.61 BLEU score compared to the 21.0 obtained by the baseline algorithm. When we increase the number of blocks from 4 to 8, the BLEU score increases to 22.05. These results demonstrate the effectiveness of our method.

As for En \rightarrow De translation, we do not observe significant improvement. Specially, we improve the En \rightarrow De from 17.2 (Lample et al. 2018) to 17.31/17.42 on the 4/8-block model. Such a phenomenon is consistent with the results shown in Figure 3. We will leave a more detailed study in the future work.

7 Conclusion and Future Work

In this paper, we studied a new parameter sharing mechanism in NMT by tying the encoder and decoder and proposed the tied transformer. Experimental results on several different tasks demonstrated the effectiveness of our proposed model. For future work, there are several interesting directions to explore in the future: First, we studied hard parameter sharing in this work. How to constrain the parameters in a soft way like regularizing the distances of parameters is an interesting topic. Second, how to apply this idea to multi-lingual translation is worthy of investigation. Third, we will study whether this idea works for other applications such text summarization and question answering.

References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2017. Enriching word vectors with subword information. *TACL* 135–146.
- Deng, Y.; Kim, Y.; Chiu, J.; Guo, D.; and Rush, A. M. 2018. Latent alignment and variational attention. *arXiv preprint arXiv:1807.03756*.
- Dong, D.; Wu, H.; He, W.; Yu, D.; and Wang, H. 2015. Multi-task learning for multiple language translation. In *ACL*, volume 1, 1723–1732.
- Edunov, S.; Ott, M.; Auli, M.; Grangier, D.; and Ranzato, M. 2018. Classical structured prediction losses for sequence to sequence learning. In *NAACL*.
- Elbayad, M.; Besacier, L.; and Verbeek, J. 2018. Pervasive attention: 2d convolutional neural networks for sequence-to-sequence prediction. In *The SIGNLL Conference on Computational Natural Language Learning*.
- Firat, O.; Sankaran, B.; Al-Onaizan, Y.; Vural, F. T. Y.; and Cho, K. 2016. Zero-resource translation with multi-lingual neural machine translation. *EMNLP*.
- Firat, O.; Cho, K.; and Bengio, Y. 2016. Multi-way, multilingual neural machine translation with a shared attention mechanism. *arXiv preprint arXiv:1601.01073*.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. N. 2017. Convolutional sequence to sequence learning. In *ICML*.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MIT Press.
- Ha, T.-L.; Niehues, J.; and Waibel, A. 2016. Toward multilingual neural machine translation with universal encoder and decoder. *arXiv preprint arXiv:1611.04798*.
- Hassan, H.; Aue, A.; Chen, C.; Chowdhary, V.; Clark, J.; Federmann, C.; Huang, X.; Junczys-Dowmunt, M.; Lewis, W.; Li, M.; et al. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv preprint arXiv:1803.05567*.
- He, D.; Xia, Y.; Qin, T.; Wang, L.; Yu, N.; Liu, T.; and Ma, W.-Y. 2016. Dual learning for machine translation. In *Advances in Neural Information Processing Systems*, 820–828.
- Johnson, M.; Schuster, M.; Le, Q. V.; Krikun, M.; Wu, Y.; Chen, Z.; Thorat, N.; Viégas, F.; Wattenberg, M.; Corrado, G.; et al. 2017. Google’s multilingual neural machine translation system: enabling zero-shot translation. *TACL*.
- Lample, G.; Ott, M.; Conneau, A.; Denoyer, L.; and Ranzato, M. 2018. Phrase-based & neural unsupervised machine translation. In *EMNLP*.
- Lample, G.; Denoyer, L.; and Ranzato, M. 2018. Unsupervised machine translation using monolingual corpora only. In *ICLR*.
- Luong, M.-T.; Le, Q. V.; Sutskever, I.; Vinyals, O.; and Kaiser, L. 2016. Multi-task sequence to sequence learning. *ICLR*.
- Luong, M.-T.; Pham, H.; and Manning, C. D. 2015. Effective approaches to attention-based neural machine translation. In *EMNLP*.
- Mohri, M.; Rostamizadeh, A.; and Talwalkar, A. 2012. *Foundations of machine learning*. MIT press.
- Pan, S. J., and Yang, Q. 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* 22(10):1345–1359.
- Pham, H.; Guan, M. Y.; Zoph, B.; Le, Q. V.; and Dean, J. 2018. Efficient neural architecture search via parameter sharing. *arXiv preprint arXiv:1802.03268*.
- Ruder, S. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016a. Improving neural machine translation models with monolingual data. In *ACL*.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016b. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 1715–1725.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 6000–6010.
- Wang, Y.; Xia, Y.; Zhao, L.; Bian, J.; Qin, T.; Liu, G.; and Liu, T. 2018. Dual transfer learning for neural machine translation with marginal distribution regularization. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*.
- Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR* abs/1609.08144.
- Xia, Y.; Qin, T.; Chen, W.; Bian, J.; Yu, N.; and Liu, T.-Y. 2017a. Dual supervised learning. In *ICML2017*.
- Xia, Y.; Tian, F.; Qin, T.; Yu, N.; and Liu, T.-Y. 2017b. Sequence generation with target attention. In *ECMLPKDD*, 816–831. Springer.
- Xia, Y.; Tan, X.; Tian, F.; Qin, T.; Yu, N.; and Liu, T.-Y. 2018. Model-level dual learning. In *International Conference on Machine Learning*, 5379–5388.
- Zeiler, M. D. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Zhang, Y., and Yang, Q. 2017. A survey on multi-task learning. *arXiv preprint arXiv:1707.08114*.
- Zoph, B., and Knight, K. 2016. Multi-source neural translation. *NAACL*.