

Deconvolution-Based Global Decoding for Neural Machine Translation

Junyang Lin^{1,2}, Xu Sun², Xuancheng Ren², Shuming Ma², Jinsong Su³, Qi Su¹

¹School of Foreign Languages, Peking University

²MOE Key Lab of Computational Linguistics, School of EECS, Peking University

³School of Software, Xiamen University

{linjunyang, xusun, renxc, shumingma, sukia}@pku.edu.cn
jssu@xmu.edu.cn

Abstract

A great proportion of sequence-to-sequence (Seq2Seq) models for Neural Machine Translation (NMT) adopt Recurrent Neural Network (RNN) to generate translation word by word following a sequential order. As the studies of linguistics have proved that language is not linear word sequence but sequence of complex structure, translation at each step should be conditioned on the whole target-side context. To tackle the problem, we propose a new NMT model that decodes the sequence with the guidance of its structural prediction of the context of the target sequence. Our model generates translation based on the structural prediction of the target-side context so that the translation can be freed from the bind of sequential order. Experimental results demonstrate that our model is more competitive compared with the state-of-the-art methods, and the analysis reflects that our model is also robust to translating sentences of different lengths and it also reduces repetition with the instruction from the target-side context for decoding.

1 Introduction

Deep learning has achieved tremendous success in machine translation, outperforming the traditional linguistic-rule-based and statistical methods. In recent studies of Neural Machine Translation (NMT), most models are based on the sequence-to-sequence (Seq2Seq) model based on the encoder-decoder framework (Kalchbrenner and Blunsom, 2013; Sutskever et al., 2014; Cho et al., 2014) with the attention mechanism (Bahdanau et al., 2014; Luong et al., 2015). While traditional linguistic-rule-based and statistical methods of machine translation require much work of feature engineering, NMT can be trained in the end-to-end fashion. Besides, the attention mechanism can model the alignment relationship between the source text and translation (Bahdanau et al., 2014; Luong et al., 2015), and some recent improved versions of attention have proved successful in this task (Tu et al., 2016; Mi et al., 2016; Meng et al., 2016; Xiong et al., 2017; Vaswani et al., 2017).

However, the decoding pattern of the recent Seq2Seq models is inconsistent with the linguistic analysis. As the conventional decoder translates words in a sequential order, the current generation is highly dependent on the previous generation and it is short of the knowledge about future generation. Nida (1969) pointed out that translation goes through a process of analysis, transfer and reconstruction, involving the deep syntactic and semantic structure of the source and target languages. Language generation involves complex syntactic analysis and semantic integration, instead of a step-by-step word generation (Frazier, 1987). Moreover, from the perspective of semantics and pragmatics, the syntactic analysis of utterance can be guided by the global lexical-semantic and discourse information (Altmann and Steedman, 1988; Trueswell et al., 1994, 1993; Tyler and Marslen-Wilson, 1977). In brief, the process of translation is in need of the global information from the target-side context, but the decoding pattern of the conventional Seq2Seq model in NMT does not meet the requirement.

Recent researches in NMT have taken this issue into consideration by the implementation of bidirectional decoding. Some methods of bidirectional decoding (Liu et al., 2016; Cong et al., 2017) rerank the candidate translations with the scores from the bidirectional decoding. However, these bidirectional

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

decoding methods cannot provide effective complementary information due to the limited search space of beam search.

In this article, we extend the conventional attention-based Seq2Seq model by introducing the deconvolution-based decoder, which is a Convolutional Neural Network (CNN) to perform deconvolution. Recently, deconvolution has been applied to the studies of natural language (Zhang et al., 2017; Shen et al., 2017), which can be regarded as the transposition of the convolution (Long et al., 2015; Noh et al., 2015). Zhang et al. (2017) applied this method to natural language by modeling sentences with a convolution-deconvolution autoencoder. The study of Zhang et al. (2017) showed that the deconvolution solves the problems by reconstructing a representation of high quality irrespective to the order or length. It can be found that deconvolution owns the potential to provide global information for guidance of decoding. Therefore, we follow this idea and propose a new model with deconvolution for NMT.

To be specific, the conventional RNN encoder encodes the source sentences to new representations and sends the final state to the decoder, and the conventional RNN decoder decodes it to the target sentences with the attention to the encoder outputs. In our model, our designed deconvolution-based decoder decodes the final state of the encoder to a matrix representing the global information of the target-side contexts. Each column of the matrix is learned to be close to the word embedding of the target words. The conventional RNN decoder can attend to the columns for the information of the target-side context to perform global decoding in the translation.

Our contributions in this study are illustrated in the following:

- We propose a new model for NMT, which contains a deconvolution-based decoder to provide global information of the target-side contexts to the RNN decoder, so that the model is able to perform global decoding¹.
- Experimental results demonstrate that our model outperforms the baseline models in both the Chinese-to-English translation and the English-to-Vietnamese translation, outperforming the Seq2Seq model in the BLEU score evaluation with the advantages of BLEU score 2.82 and 1.54 respectively.
- The analysis shows that our model that performs global decoding is more capable of reducing repetition and more robust to the translation of sentences of different lengths, and the case study reflects that it is able to capture the syntactic structure for the translation and has a better reflection of the semantic meaning of the source text.

2 Model

In the following, we introduce the details of our model, including the encoder, the deconvolution-based decoder and the conventional RNN-based decoder. The functions of each decoder are illustrated below to show how they collaborate to improve the quality of the translation.

2.1 Encoder

In our model, the encoder reads the embeddings of the input text sequence $x = \{x_1, \dots, x_n\}$ and encodes a sequence of encoder outputs $h = \{h_1, \dots, h_n\}$. The final hidden state h_n is sent to the decoder as the initial state for it to decode a sequence of output text. The encoder outputs provide the information of the source-side contexts to our RNN-Based decoder through the attention mechanism.

The encoder in our model is a bidirectional LSTM (Hochreiter and Schmidhuber, 1997), which reads the input in two directions to generate two sequences of hidden states $\vec{h} = \{\vec{h}_1, \vec{h}_2, \vec{h}_3, \dots, \vec{h}_n\}$ and $\overleftarrow{h} = \{\overleftarrow{h}_1, \overleftarrow{h}_2, \overleftarrow{h}_3, \dots, \overleftarrow{h}_n\}$, where:

$$\vec{h}_i = LSTM(x_i, \overrightarrow{h_{i-1}}, C_{i-1}) \quad (1)$$

$$\overleftarrow{h}_i = LSTM(x_i, \overleftarrow{h_{i-1}}, C_{i-1}) \quad (2)$$

The encoder outputs corresponding to each time step are concatenated as $h_i = [\vec{h}_i; \overleftarrow{h}_i]$.

¹The code is released in <https://github.com/lancopku/DeconvDec>




Figure 1: Model architecture. There are three components in the proposed model, i.e., the LSTM encoder, the deconvolution-based decoder, and the conventional LSTM decoder. The encoder distills the input sentence into a state h_n , which is then used in the deconvolution-based decoder to obtain the global information of the target-side contexts. Based on the target-side contexts and the input-side contexts, the conventional LSTM decoder generates the output from the state h_n .

2.2 Deconvolution-Based Decoder

In our model, there are two decoders, which perform different tasks for the whole decoding process. While the RNN-based decoder is similar to the conventional decoder, which decodes the output text sequence in a sequential order and attends to the annotations of the encoder via the attention mechanism, our proposed deconvolution-based decoder does not decode the text but provide global information of the target-side contexts to the RNN-based decoder so that it can decode structurally instead of sequentially. To be specific, the deconvolution-based decoder learns to generate the word embedding matrix of the target text sequence.

In order to provide global information of the target-side contexts to the RNN decoder, we implement a multilayer CNN as the deconvolution-based decoder to perform deconvolution. With deconvolution, it is available for a vector or a small matrix to be transformed to a large matrix. In our model, the deconvolution is implemented on the final states in both directions from the encoder. As words in our model are represented with word embedding vectors, sentences can be formed as word embedding matrices. In our model, the deconvolution-based decoder is designed to learn word embedding matrices of the target sequences with the representation matrix from the encoder. As the conjugate operation of convolution, deconvolution expands the dimension of the input representation to a matrix of our designed size. There are L layers in the deconvolution, each of which has f_l filters of kernel size k_l . The i^{th} filter $W_l^i \in R^{k \times dim}$ (dim refers to the size of the input representation vector) with stride s_l^i and padding p_l^i performs deconvolution on the input representation matrix $I \in R^{m \times dim}$ ($m \times dim$ refers to the size of the input representation matrix), the final hidden state of the encoder. The computation of convolution is illustrated as below:

$$c_l^i = g(W_l^i * X + b) \quad (3)$$

where X refers to the convolved matrix and g refers to non-linear activation function, which is ReLU (Nair and Hinton, 2010) following Zhang et al. (2017), and deconvolution is its transposed operation. With the input I , our objective of the deconvolution operation is to generate a word embedding matrix $E \in R^{T \times dim}$ where T refers to the sentence length designed for the output text sequence, which is a hyper-parameter. At the l^{th} layer, deconvolution generates a matrix $E_l \in R^{T_l \times dim}$ where $T_l = T_{l-1} \times s_l + k_l - 2 \times p_l$. With the control of stride and kernel size, the height of the matrix can be assigned, and with the control of the number of filters, the width of the matrix can also be assigned. In our model, they are the length of the output sequence and the dimension of the word embedding respectively.

The deconvolution-based decoder can generate meaningful representation with information different from that in the conventional RNN decoder. The conventional RNN decoder generates sequence in a way similar to Markov Decision Process, which is highly dependent on the previous generation and follows

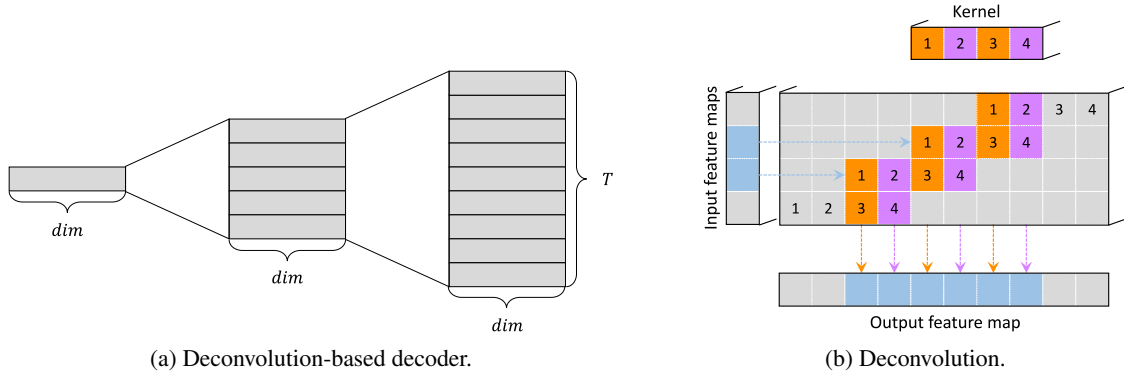


Figure 2: Deconvolution-based decoder. On the right shows an example of a 1d deconvolution on a input of size 2 with a kernel of size 4, a padding of 1 and, a stride of 2. The depth means the dimension of the channel, which is dim in our case.

a strict sequential order, so it contains high sequential dependency. On the contrary, the deconvolution-based decoder generates a word embedding matrix depending on the representation from the encoder without considering the order, which does not have the problem of long-term dependency. Moreover, although it is not capable as the RNN encoder to generate coherent text, it reveals the information of the text from a global perspective, including syntactic and semantic features.

2.3 RNN-Based Decoder

Different from the deconvolution-based decoder, the RNN-based decoder is responsible for decoding the representation h_n to generate the translation $y = \{y_1, \dots, y_m\}$. With the final encoder state as the initial state, the decoder is initialized to decode in sequential order, until it generates the token representing the end of sentence. During decoding, the attention mechanism is applied for the decoder to extract the information from the source-side contexts, which are the annotations of the encoder, as well as the information from the target-side contexts, which are the outputs of the deconvolution-based decoder.

For the RNN-based decoder, we implement a unidirectional LSTM. The output of the RNN-based decoder at each time step is sent into a feed-forward neural network to be projected into the space of the target vocabulary $Y \in R^{|Y| \times dim}$ for the prediction of the translated word. At each time step, the decoder generates a word y_t by sampling from a conditional probability distribution of the target vocabulary P_{vocab} , where:

$$P_{vocab} = softmax(W_o v_t) \quad (4)$$

$$v_t = g(s_t, c_t, \tilde{c}_t) \quad (5)$$

$$s_t = LSTM(y_{t-1}, s_{t-1}, C_{t-1}) \quad (6)$$

where $g(\cdot)$ refers to non-linear activation function, and c_t and \tilde{c}_t are the outputs of the attention mechanism, which are illustrated in the following.

The attention mechanism in our model is the global attention mechanism (Luong et al., 2015). Different from the conventional attention mechanism, which only computes the attention scores on the source-side contexts, the attention mechanism in our model consists of two parts. The first one is similar to the conventional one, attending to the source-side contexts from the encoder, but the second one is original, which attends to the target-side contexts, which is the word embedding matrix generated by the deconvolution-based decoder. By attending to the encoder annotations, the model computes the attention $\alpha_{t,i}$ of the RNN-based decoder output s_t on the annotations of the encoder h_i and generates the context vector c_t . Similarly, by attending to the outputs of the deconvolution-based decoder, the RNN-based decoder computes the attentions of s_t on each column E_i of its matrix E and generates the context vector

\tilde{c}_t :

$$c_t = \sum_{i=1}^n \alpha_{t,i} h_i \quad (7)$$

$$\tilde{c}_t = \sum_{i=1}^n \tilde{\alpha}_{t,i} E_i \quad (8)$$

where $\alpha_{t,i}$ and $\tilde{\alpha}_{t,i}$ are defined as below (as they are computed in the same way, they are both represented by $\alpha_{t,i}$ and the annotations are represented by x_i):

$$\alpha_{t,i} = \frac{\exp(e_{t,i})}{\sum_{j=1}^n \exp(e_{t,j})} \quad (9)$$

$$e_{t,i} = s_{t-1}^\top W_a x_i \quad (10)$$

2.4 Training

The training for the Seq2Seq model is usually based on maximum likelihood estimation. Given the parameters θ and source text x , the model generates a sequence \tilde{y} . The learning process is to minimize the negative log-likelihood between the generated text \tilde{y} and reference y , which in our context is the sequence in target language for machine translation:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \log P(y_t^{(i)} | \tilde{y}_{<t}^{(i)}, x^{(i)}, \theta) \quad (11)$$

where the loss function is equivalent to maximizing the conditional probability of sequence y given parameters θ and source sequence x .

However, as there are two decoders in our model, the loss function should also be designed for the deconvolution-based decoder. In our model, we compute the smooth L1 loss between the generated matrix of the deconvolution-based decoder E and the word embedding matrix \tilde{E} (which both contain M elements), which is more robust to outliers (Girshick, 2015), as well as the cross-entropy loss between the prediction of the deconvolution-based decoder \hat{y} and reference y given the parameters of the encoder and the deconvolution-based decoder θ' . Therefore, the generated matrix E can be closer to the word embedding matrix \tilde{E} , and it contains information beneficial to the prediction of the target words. Moreover, for the cross entropy loss of the deconvolution-based decoder, we apply the method of Ma et al. (2018a) as it increases no parameter for the prediction by computing the cosine similarity between the output and the word embeddings. To sum up, the loss function is defined as below:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \left(\sum_{t=1}^T \log P(y_t^{(i)} | \tilde{y}_{<t}^{(i)}, x^{(i)}, \theta) + \sum_{m=1}^M \text{smooth}_{L1}(E_m - \tilde{E}_m) + \sum_{t=1}^T \log P(y_t^{(i)} | x^{(i)}, \theta') \right) \quad (12)$$

where smooth L1 loss is defined below:

$$\text{smooth}_{L1}(x, y) = \begin{cases} 0.5 \|x - y\|_2^2 & \text{if } \|x - y\| < 1 \\ \|x - y\|_1 - 0.5 & \text{if } \|x - y\| \geq 1 \end{cases} \quad (13)$$

We have tested L1 loss, L2 loss as well as smooth L1 loss in our experiments and found that smooth L1 loss encourages the model to reach the best performance.

3 Experiment

3.1 Datasets

We evaluate our proposed model on the NIST translation task for the Chinese-to-English translation and provide the analysis on the same task. Moreover, in order to evaluate the performance of our model on

the low-resource translation, we also evaluate our model on the IWLST 2015 (Cettolo et al., 2015) for the English-to-Vietnamese translation task.

Chinese-to-English Translation For the NIST translation task, we train our model on 1.25M sentence pairs extracted from LDC2002E18, LDC2003E07, LDC2003E14, Hansards portion of LDC2004T07, LDC2004T08 and LDC2005T06, with 27.9M Chinese words and 34.5M English words. Following Wang et al. (2016), we validate our model on the dataset for the NIST 2002 translation task and tested our model on that for the NIST 2003, 2004, 2005, 2006 translation tasks. We use the most frequent 30K words for both the Chinese vocabulary and the English vocabulary, which includes around 97.4% and 99.5% of the Chinese and English words in the training data. The sentence pairs longer than 50 words are filtered. The evaluation metric is BLEU (Papineni et al., 2002).

English-to-Vietnamese Translation The data is from the translated TED talks, containing around 133K training sentence pairs provided by the IWSLT 2015 Evaluation Campaign (Cettolo et al., 2015). We follow the studies of Huang et al. (2017), and use the same preprocessing methods as well as the same validation and the test set. The validation set is the TED tst2012 with 1553 sentences and the test set is the TED tst2013 with 1268 sentences. The English vocabulary is 17.7K words and the Vietnamese vocabulary is 7K words. The evaluation metric is also BLEU score.

3.2 Setting

We implement the models on PyTorch², and the experiments are conducted on an NVIDIA 1080Ti GPU. Both the size of word embedding and the number of units in the hidden layers are 512, and the batch size is 64. We use Adam optimizer (Kingma and Ba, 2014) to train the model with the setting $\beta_1 = 0.9$, $\beta_2 = 0.98$ and $\epsilon = 1 \times 10^{-9}$ following Vaswani et al. (2017), and we initialize the learning rate to 0.0003.

Gradient clipping is applied so that the norm of the gradients cannot be larger than a constant, which is 10 in our experiments. Dropout is used with the dropout rate set to 0.3 for the Chinese-to-English translation and 0.4 for the English-to-Vietnamese translation, based on the performance on the validation set.

Based on the performance on the validation set, we use beam search with a beam width of 10 to generate translation for the evaluation and test, and we normalize the log-likelihood scores by sentence length.

3.3 Baselines

For the Chinese-to-English translation, we compare our model with the state-of-the-art NMT systems for the task.

- **Moses** An open source phrase-based translation system with default configurations and a 4-gram language model trained on the training data for the target language;
- **RNNsearch** An attention-based Seq2Seq with fine-tuned hyperparameters (Bahdanau et al., 2014);
- **Coverage** The method extends RNNSearch with a coverage model for the attention mechanism that tackles the problem of over-translation and under-translation (Tu et al., 2016);
- **Lattice** The Seq2Seq model with a word-lattice-based RNN encoder that tackles the problem of tokenization in NMT (Su et al., 2016);
- **InterAtten** The Seq2Seq model that records the interactive history of decoding (Meng et al., 2016);
- **MemDec** Based on the RNNSearch, it is equipped with external memory that the model reads and writes during decoding (Wang et al., 2016).

For the English-to-Vietnamese translation, we compare our model with the recent NMT models for this task, and we present the results of the baselines reported in their articles.

²<http://pytorch.org>

Model	MT-03	MT-04	MT-05	MT-06	Ave.
Moses	32.43	34.14	31.47	30.81	32.21
RNNSearch	33.08	35.32	31.42	31.61	32.86
Lattice	34.32	36.50	32.40	32.77	34.00
Coverage	34.49	38.34	34.91	34.25	35.49
InterAtten	35.09	37.73	35.53	34.32	35.67
MemDec	36.16	39.81	35.91	35.98	36.97
Seq2Seq+Attention	35.32	37.25	33.52	33.54	34.91
+DeconvDec	38.04	39.75	36.77	36.32	37.73

Table 1: Results of our model and the baselines (the results are those reported in the referred articles, and the models are trained on the identical training data or larger training data) on the Chinese-to-English translation, tested on the NIST Machine Translation tasks in 2003, 2004, 2005, 2006 by BLEU score evaluation.

Model	BLEU
RNNSearch-1	23.30
RNNSearch-2	26.10
LabelEmb	26.80
NPMT	27.69
Seq2Seq+Attention	26.93
+DeconvDec	28.47

Table 2: Results of our model and the baselines (directly reported in the referred articles) on the English-to-Vietnamese translation, tested on the TED tst2013 with the BLEU score evaluation.

- **RNNsearch-1** The attention-based Seq2Seq model by Luong and Manning (2015);
- **RNNsearch-2** The implementation of the attention-based Seq2Seq by Huang et al. (2017);
- **LabelEmb** Extending RNNSearch with soft target representation (Sun et al., 2017);
- **NPMT** The Neural Phrased-based Machine Translation model by Huang et al. (2017);

4 Results and Analysis

4.1 Results

Table 1 shows the overall results of the models on the Chinese-to-English translation task. Beside our reimplementation of the attention-based Seq2Seq model, we report the results of the recent NMT models, which are results in their original articles or improved results of the reimplementation. To facilitate fair comparison, we compare with the baselines that are trained on the same training data. The results have shown that for the NIST 2003, 2004, 2005 and 2006 translation tasks, our model with the deconvolution-based decoder outperforms the baselines, and the advantage of BLEU score over the attention-based Seq2Seq model is 2.82 on average compared with our reimplementation of the attention-based Seq2Seq model. From the results mentioned above, it can be inferred that the global information of the target-side contexts retrieved from the deconvolution-based decoder is contributive to the translation. Our analysis and case study in the following can further demonstrate how the deconvolution-based decoder improves the attention-based Seq2Seq model.

Table 2 presents the results of the models on the English-to-Vietnamese translation. Compared with the attention-based Seq2Seq model, including the implementation with the strongest performance, our model with the deconvolution-based decoder can outperform it with the advantage of BLEU score 1.54. We also display the most recent model NPMT (Huang et al., 2017) trained and tested on the dataset.




Figure 3: **Percentage of the duplicates at sentence level and the BLEU scores on sentences of different lengths** Tested on the NIST 2003 dataset. The red bar and line indicate the performance of our model, and the blue bar and line indicate that of the attention-based SeqSeq model.

Compared with NPMT, our model has an advantage of BLEU score 0.78. It can be indicated that for the low-resource translation, the information from the deconvolution-based decoder is important, which brings significant improvement to the conventional attention-based Seq2Seq model.

4.2 Analysis

As our model generates translation with global information from the deconvolution-based decoder, it should learn to reduce repetition as it can learn to avoid generating same contents according to the conjecture by the deconvolution-based decoder about the target-side contexts. In order to test whether our model can mitigate the problem of repetition in translation, we test the repetition on the NIST 2003 dataset, following See et al. (2017). The proportions of the duplicates of 1-gram, 2-gram, 3-gram and 4-gram in each sentence are calculated. Results on Figure 3(a) show that our model generates less repetitive translation. In particular, the proportion of duplicates of our model is less than half of that of the conventional Seq2Seq model.

Moreover, to validate its robustness on different sentence-length levels, we test the BLEU scores on sentences of length no shorter than 10 to 60 of the NIST 2003 dataset. According to the results on Figure 3(b), though with the increase in length, the performance of our model is always stronger than the Seq2Seq model. However, with the increase of length, the advantage of our model becomes smaller. This is consistent with our hypothesis. Since the length of generation of the deconvolution-based decoder is assigned a particular value (30 words in Chinese-to-English translation) due to the limited computation resources, there is not enough global target-side information for translating long sentences (say, longer than 30 words). In our future work, we will delve into this problem and conduct further research to reduce computation cost.

Figure 4 presents the attention heatmaps of the RNN-based decoder on the generated matrix of the deconvolution-based decoder in the English-to-Vietnamese translation. They reflect that the RNN-based decoder has diverse local focuses on the self-contained target-side contexts at different time steps. Contrary to the conventional attention on the source-side contexts which captures the corresponding annotations, it focuses on groups of the columns of the generated matrix from the deconvolution-based decoder. With the guidance of the information of global decoding, the model generates translation of higher accuracy and higher coherence. However, as the deconvolution-based decoder is not responsible for generating translation, it is hard to interpret what each column of the generated matrix represents. Moreover, as it does not capture alignment relationship as the conventional attention mechanism does, it is our future work to improve the attention on the outputs of the deconvolution-based decoder and explain the group focuses as shown in the heatmaps.




Figure 4: **Attention heatmaps of the RNN-based decoder on the deconvolution-based decoder**
 Words on the left refer to the translation of the RNN-based decoder. The heatmaps show that the RNN-based decoder can focus on certain parts of the outputs from the deconvolution-based decoder.

4.3 Case Study

Table 3 demonstrates three examples of the translation of our model in comparison with the Seq2Seq model and the golden translation. In Table 3(a), while the Seq2Seq model cannot recognize the objects of the main clause and the infinitive, causing inaccuracy and repetition, our model better captures the syntactic structure of the sentence and translates the main idea of the source text, though leaving the information “that causes disease”. In Table 3(b), the source sentence is more complicated. With a temporal adverbial clause, its syntactic structure is more complex than simple sentence. The translation of the conventional Seq2Seq model cannot capture the syntactic information in the source text and regards the “parliament members” as the argument of the predicate “talk”. Moreover, it is confused by the word “中期”, meaning “middle”, and translates “mid - autumn festival”. In comparison, our model can recognize the adverbial clause and the main clause as well as their syntactic structures. In Table 3(c), it can be shown that when translating the long and relatively complex text, the baseline model makes a series of mistake of repetition. In contrast, the translation generated by our model though repeats the word “disaster”, it is much more coherent and more semantically consistent with the source text as it successfully presents “sent 250,000 yuan” corresponding to the source text “调拨25万元人民币”, while the baseline cannot translate the content.

5 Related Work

In the following, we review the studies in NMT and the application of deconvolution in NLP.

Kalchbrenner and Blunsom (2013); Cho et al. (2014); Sutskever et al. (2014) studied the application of the encoder-decoder framework on the machine translation task, which launched the development of NMT. Another significant innovation in this field is the attention mechanism, which builds connection between the translated contents and the source text (Bahdanau et al., 2014; Luong et al., 2015). To improve the quality of NMT, researchers have focused on improving the attention mechanism. Tu et al. (2016) and Mi et al. (2016) modeled coverage in the NMT, Meng et al. (2016) and Xiong et al. (2017) incorporated the external memory to the attention, and Xia et al. (2017) as well as Lin et al. (2018a) utilized the information from the previous generation by target-side attention and memory for attention history respectively. For more target-side information, Ma et al. (2018b) incorporated bag of words as target. A breakthrough of NMT in recent years is that Vaswani et al. (2017) invented a model only with the attention mechanism that reached the state-of-the-art performance.

Although many researches in NLP focused on the application of RNN, CNN is also an important type of network for the study of language (Kim, 2014; Kalchbrenner et al., 2014; Zhang et al., 2015; Lin et al., 2018b). Also, its application in NMT has been successful (Gehring et al., 2017). Recently, deconvolution was applied to modeling text (Zhang et al., 2017; Shen et al., 2017), which is able to construct a representation of high quality with the self-contained information.

Text: 基因科学家的目标是,提供诊断工具以发现致病的缺陷基因
Gold: the goal of geneticists is to provide diagnostic tools to identify defective genes that cause diseases
Seq2Seq: the objective of genetic scientists is to provide genes to detect genetic genetic genes
DeconvDec: the objective of the gene scientist is to provide diagnostic tools to detect the defective genes
(a)
Text: 叛军暗杀两位菲国国会议员后,菲律宾总统雅罗育在二零零一年中期停止与共党谈判。
Gold: after the rebels assassinated two philippine legislators , philippine president arroyo ceased negotiations with the communist party in mid 2001 .
Seq2Seq: philippine president gloria arroyo stopped the two philippine parliament members in the mid - autumn festival .
DeconvDec: philippine president gloria arroyo stopped holding talks with the communist party after the rebels assassinated two philippine parliament members .
(b)
Text: 中国红十字会已在24日地震发生后紧急向新疆灾区调拨25万元人民币,又于25日向灾区派出救灾工作组。
Gold: china red cross has released 250 thousand renminbi for the xinjiang disaster area immediately after the earthquake on the 24th . a disaster relief team was dispatched to the area on the 25th .
Seq2Seq: the red cross society of china (red cross) , the china red cross society (red cross) , emergency relief team sent an emergency team to xinjiang for disaster relief in the disaster areas after the earthquake on 24 june .
DeconvDec: the china red cross society has sent 250,000 yuan to the disaster areas in xinjiang after the earthquake occurred on the 24 th, and sent a relief team to disaster disaster areas on the 25 th.
(c)

Table 3: Two examples of the translation of our model in comparison with the conventional attention-based Seq2Seq model on the NIST 2003 Chinese-to-English translation task. The errors in the translation are colored in red and the successful translation of some particular contents are colored in yellow (e.g., the contents that the model successfully translates but the other does not).

6 Conclusion and Future Work

In this paper, we propose a new model with the global decoding mechanism. With our deconvolution-based decoder, which provides global information of the target-side contexts, the model can effectively exploit the information for the inference of syntactic structure and semantic meaning in the translation. Experimental results on the Chinese-to-English translation and English-to-Vietnamese translation demonstrate that our model outperforms the baseline models, and the analysis shows that our model generates less repetitive translation and demonstrates higher robustness to the sentences of different lengths. Furthermore, the case study shows that the translation of our model better observes the syntactic and semantic requirements for the translation and generates coherent and accurate translation with fewer irrelevant contents.

In the future, we will further develop analysis of the mechanism of deconvolution in NMT and try to figure out its generalized patterns for the construction of the target-side contexts.

Acknowledgements

This work was supported in part by National Natural Science Foundation of China (No. 61673028) and the National Thousand Young Talents Program. Xu Sun is the corresponding author of this paper.

References

- G Altmann and M Steedman. 1988. Interaction with context during human sentence processing. *Cognition*, 30(3):191.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, Roldano Cattoni, and Marcello Federico. 2015. The iwslt 2015 evaluation campaign. *Proc. of IWSLT, Da Nang, Vietnam*.
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP 2014*, pages 1724–1734.
- Duy Vu Hoang Cong, Gholamreza Haffari, and Trevor Cohn. 2017. Towards decoding as continuous optimisation in neural machine translation. In *EMNLP 2017*, pages 146–156.
- Lyn Frazier. 1987. *Sentence Processing: A Tutorial Review*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. In *ICML 2017*, pages 1243–1252.
- Ross Girshick. 2015. Fast r-cnn. *Computer Science*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Po-Sen Huang, Chong Wang, Dengyong Zhou, and Li Deng. 2017. Neural phrase-based machine translation. *CoRR*, abs/1706.05565.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent continuous translation models. In *EMNLP 2013*, pages 1700–1709.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *ACL 2014*, pages 655–665.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *EMNLP 2014*, pages 1746–1751.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Junyang Lin, Shuming Ma, Qi Su, and Xu Sun. 2018a. Decoding-history-based adaptive control of attention for neural machine translation. *CoRR*, abs/1802.01812.
- Junyang Lin, Xu Sun, Shuming Ma, and Qi Su. 2018b. Global encoding for abstractive summarization. *CoRR*, abs/1805.03989.
- Lemao Liu, Masao Utiyama, Andrew Finch, and Eiichiro Sumita. 2016. Agreement on target-bidirectional neural machine translation. In *NAACL HLT 2016*, pages 411–416.
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *CVPR 2015*, pages 3431–3440.
- Minh-Thang Luong and Christopher D Manning. 2015. Stanford neural machine translation systems for spoken language domains. In *Proceedings of the International Workshop on Spoken Language Translation*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *EMNLP 2015*, pages 1412–1421.
- Shuming Ma, Xu Sun, Wei Li, Sujian Li, Wenjie Li, and Xuancheng Ren. 2018a. Query and output: Generating words by querying distributed word representations for paraphrase generation. In *NAACL-HLT 2018*, pages 196–206.

- Shuming Ma, Xu Sun, Yizhong Wang, and Junyang Lin. 2018b. Bag-of-words as target for neural machine translation. *CoRR*, abs/1805.04871.
- Fandong Meng, Zhengdong Lu, Hang Li, and Qun Liu. 2016. Interactive attention for neural machine translation. In *COLING 2016*, pages 2174–2185.
- Haitao Mi, Baskaran Sankaran, Zhiguo Wang, and Abe Ittycheriah. 2016. Coverage embedding models for neural machine translation. In *EMNLP 2016*, pages 955–960.
- Vinod Nair and Geoffrey E. Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *ICML 2010*, pages 807–814.
- Eugene A Nida. 1969. Science of translation. *Language*, pages 483–498.
- Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. 2015. Learning deconvolution network for semantic segmentation. In *ICCV 2015*, pages 1520–1528.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL, 2002*, pages 311–318.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *ACL 2017*, pages 1073–1083.
- Dinghan Shen, Yizhe Zhang, Ricardo Henao, Qinliang Su, and Lawrence Carin. 2017. Deconvolutional latent-variable model for text sequence matching. *CoRR*, abs/1709.07109.
- Jinsong Su, Zhixing Tan, Deyi Xiong, and Yang Liu. 2016. Lattice-based recurrent neural network encoders for neural machine translation. *CoRR*, abs/1609.07730.
- Xu Sun, Bingzhen Wei, Xuancheng Ren, and Shuming Ma. 2017. Label embedding network: Learning label representation for soft training of deep networks. *CoRR*, abs/1710.10393.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *NIPS, 2014*, pages 3104–3112.
- J. C. Trueswell, M. K. Tanenhaus, and S. M. Garnsey. 1994. Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language*, 33(3):285–318.
- J. C. Trueswell, M. K. Tanenhaus, and C Kello. 1993. Verb-specific constraints in sentence processing: separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology Learning Memory and Cognition*, 19(3):528–553.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling coverage for neural machine translation. In *ACL 2016*.
- Lorraine K. Tyler and William D. Marslen-Wilson. 1977. The on-line effects of semantic context on syntactic processing. *Journal of Verbal Learning and Verbal Behavior*, 16(6):683–692.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *CoRR*, abs/1706.03762.
- Mingxuan Wang, Zhengdong Lu, Hang Li, and Qun Liu. 2016. Memory-enhanced decoder for neural machine translation. In *EMNLP 2016*, pages 278–286.
- Yingce Xia, Fei Tian, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2017. Sequence generation with target attention. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2017, Skopje, Macedonia, September 18-22, 2017, Proceedings, Part I*, pages 816–831.
- Hao Xiong, Zhongjun He, Xiaoguang Hu, and Hua Wu. 2017. Multi-channel encoder for neural machine translation. *CoRR*, abs/1712.02109.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *NIPS, 2015*, pages 649–657.
- Yizhe Zhang, Dinghan Shen, Guoyin Wang, Zhe Gan, Ricardo Henao, and Lawrence Carin. 2017. Deconvolutional paragraph representation learning. In *NIPS 2017*, pages 4172–4182.