

Generating Diverse Translation by Manipulating Multi-Head Attention

Zewei Sun, Shujian Huang, Hao-Ran Wei, Xin-yu Dai, Jiajun Chen

State Key Laboratory for Novel Software Technology
Nanjing University, Nanjing 210023, China
sunzw@smail.nju.edu.cn, whr94621@foxmail.com
{huangsj, daixinyu, chenjj}@nju.edu.cn

Abstract

Transformer model (Vaswani et al. 2017) has been widely used in machine translation tasks and obtained state-of-the-art results. In this paper, we report an interesting phenomenon in its encoder-decoder multi-head attention: different attention heads of the final decoder layer align to different word translation candidates. We empirically verify this discovery and propose a method to generate diverse translations by manipulating heads. Furthermore, we make use of these diverse translations with the back-translation technique for better data augmentation. Experiment results show that our method generates diverse translations without a severe drop in translation quality. Experiments also show that back-translation with these diverse translations could bring a significant improvement in performance on translation tasks. An auxiliary experiment of conversation response generation task proves the effect of diversity as well.

Introduction

In recent years, neural machine translation (NMT) has shown its ability to produce precise and fluent translations (Sutskever, Vinyals, and Le 2014; Bahdanau, Cho, and Bengio 2015; Luong, Pham, and Manning 2015). More and more novel network structures has been proposed (Barone et al. 2017; Gehring et al. 2017; Vaswani et al. 2017), among which Transformer (Vaswani et al. 2017) achieves the best results. The main differences between Transformer and other translation models are: i) self-attention architecture, ii) multi-head attention mechanism. We focus on the second one in this paper.

Intuitively, the attention mechanism in traditional attention-based sequence-to-sequence models plays the role of choosing the next source word to be translated, which could be seen as an alignment between source and target words (Bahdanau, Cho, and Bengio 2015; Luong, Pham, and Manning 2015). However, how multi-head attention works seems unclear.

In this paper, we report an interesting phenomenon in Transformer: in the final layer of its decoder, each individual encoder-decoder attention head dispersedly aligns to a specific source word which is highly likely to be translated next.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In other words, multi-head attention actually learns multiple alignment choices. Further, by means of picking different attention heads, we can precisely control the following word generation. We verify this characteristic by a series of statistic study afterwards.

Straightway, we consider taking advantage of this intrinsic characteristic to generate diverse translations due to the multiple generation candidates. Natural language can be diversely translated through different syntax structures or word orders. However, it has been well recognized that NMT system severely lacks translation diversity, distinguished from human beings (He, Haffari, and Norouzi 2018; Ott et al. 2018; Edunov et al. 2018). We try to tackle this issue with a new method based on our observation.

There have been a few works attempting to generate more diverse translation, which can be roughly divided into two categories. The first category tries to encourage diversity during beam search by adding some regularization items (Li, Monroe, and Jurafsky 2016; Vijayakumar et al. 2018). However, these methods actually fail to produce satisfied diversity in our re-implementation experiments. The other category tends to augment diversity by introducing latent variables (He, Haffari, and Norouzi 2018; Shen et al. 2019). However, they heavily increase training barrier and lack interpretability.

Different from them, we make use of the diverse factors we observe inside the model structure, which is more lightweight and interpretable. Since multi-head attention has the potential to identify different translation candidates, we propose a method to manipulate it to generate diverse translations. Our method works simply but more effectively than previous works, bringing in no extra parameter or regularization item. Furthermore, we propose to combine diverse translations with the back-translation technique for better data augmentation.

Experiment results show that the proposed method could generate diverse translations without a severe drop in translation performance. Besides, improvements could be achieved by employing our more diverse back-translation results for machine translation. An auxiliary experiment of conversation response generation proves the effect of diversity as well.

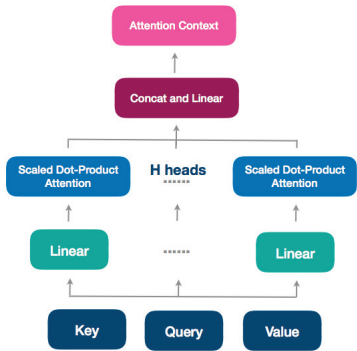


Figure 1: Multi-Head Attention consists of several attention heads running in parallel.

Background

Transformer (Vaswani et al. 2017) architecture adopts the encoder-decoder structure, utilizing self-attention instead of recurrent or convolutional networks. The encoder iteratively processes its hidden representation through 6 layers of self-attention and feed-forward network, coupled with layer normalization and residual connection. Afterwards, the decoder takes a similar circuit with injected by an encoder-decoder attention layer between self-attention and feed-forward components. An important difference from previous models is that Transformer turns all the attention mechanism into a multi-head version.

Instead of performing a single attention function with d -dimensional keys, values and queries, multi-head attention projects them into H different sub-components. After calculating attention for every sub-component, each yielding a d/H -dimensional output context, these context vectors are concatenated and projected, resulting in the final context, as depicted in Figure 1. Specifically:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^o \quad (1)$$

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V) \quad (2)$$

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (3)$$

To complete the decoding part, the model uses learned linear transformation and softmax function to convert the decoder output to next-token probabilities. The embedding and final transformation parameters are shared mutually.

Analysis of Multi-Head Attention

Each Head Indicates an Alignment

Previous works show that multi-head attention plays a key role in the significant improvement of translation performance (Vaswani et al. 2017; Chen et al. 2018). However, not much observation was made on its inside pattern. We visualize the multi-head attention to see whether different heads play different roles. After observing plenty of examples, we find that at every decoding step, all the source words that are identified by heads are highly likely to be translated next. In other words, each head aligns to a source word candidate.

	NLL		NLL
Rank 1	0.59	Head 0	5.05
Rank 3	4.12	Head 1	5.40
Rank 5	5.19	Head 2	4.69
Head-Average	5.27	Head 3	5.04
Rank 6	5.53	Head 4	5.33
Rank 10	6.40	Head 5	5.72
Rank 100	9.47	Head 6	5.68
Rank 1000	11.54	Head 7	5.32
Random	13.57	Head-Average	5.27

Table 1: Negative log-likelihood of attention heads and words ranked R^{th} on average.

We do the statistic study to verify this observation on NIST MT03 dataset (see Datasets in Experiment). At each timestep, we pick H referred source words (may overlap) that H heads correspond to. “Referred source word” means the source word with max attention for each head. Then we translate these referred source words into target language with baseline model and name these translations “referred target words”. We count the number of times these “referred target words” appear at different rankings of the softmax target probability and plot them in Figure 2. We can see that the vast majority of heads align to the most possible words.

Also, we collect the negative log-likelihood (NLL) of these referred target words to see whether they really have high generation probability. To make a comparison, we list the average NLL of words ranked R^{th} as well. The results in Table 1 verify our assumption. The chosen words are ranked around 5^{th} on average, which implies they are indeed quite possible to be selected at each decoding step.

Each Head Determines a Generation Word

Furthermore, we can control the next word generation by choosing the corresponding source word by choosing different heads. As presented in Table 2, a Zh2En model has translated “he said :” and waits for the following context. At this step, different heads refer to several source words. From Figure 3, we can see that head 4,5,6 refer to YiLai (since), XiaJiang (decline), ChuKou (exports), respectively. We con-

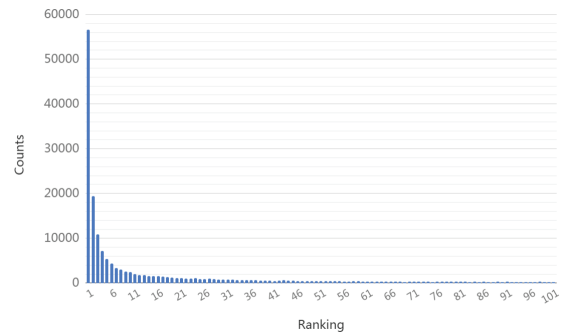


Figure 2: “Referred target words” ranking counts of top 100. The vast majority of heads align to the most possible words.




Figure 3: Different heads have different attention, referring to different words. For example, head 4,5,6 refer to YiLai (since), XiaJiang (decline), ChuKou (exports), respectively.

Source	Ta Shuo , QuNian JiuYue YiLai , ChuKou XiaJiang DaoZhi YinDu JingJi EHua .
Reference	he said : the drop in exports has caused india 's economy deterioration since september last year .
Translated	he said :
YiLai (since)	he said : since september last year , the decline in exports has led to a deterioration in india 's economy .
XiaJiang (decline)	he said : the decline in exports has led to a deterioration in india 's economy since september last year .
Chukou (exports)	he said : exports have declined since september last year , causing india 's economy to deteriorate .

Table 2: A Zh2En model has translated “he said :” and waits for the following context. Different heads (referred to different candidates) determine different following generation.

Control the model to generate the specific word by selecting the corresponding head and copying its attention weights to the other $H - 1$ heads. In this way, we indeed obtain different translation results with expected translation candidates (see three translation outputs in Table 2).

Intuitively, we can utilize these characteristics to generate diverse translations by picking different candidates to change the word choices or the sentence structure. More importantly, the diversity is from an interpretable mechanism rather than an abstract latent variable like previous works.

Diversity-Encouraged Generation

Since we have confirmed multiple head alignments can be utilized, it is natural for us to sample different heads at every timestep, so that diverse word candidates can be generated. However, we found that it will badly harm the translation quality if sampling everywhere. So we propose a sample policy to balance quality and diversity.

As stated in Algorithm 1, at every decoding step t , we denote att_{it}^h as the attention of $head_h$ from target side hidden state s_t to source side word src_i . The most possible candi-

Algorithm 1 Sample Policy

Input: The source sentence length T , a hyper parameters K , the head number H , a counting array $[n_0, \dots, n_i, \dots, n_{T-1}]$

Output: Adjusted attention

```

1: for  $t$  in decoding timesteps do
2:   for  $i$  in range( $T$ ) do
3:      $n_i = 0$ 
4:   end for
5:   calculate  $att_{it}^h$ ,  $i \in [0, T), h \in [0, H)$ 
6:   for  $h$  in range( $H$ ) do
7:      $candidate_t^h = \arg \max_i att_{it}^h$ 
8:      $n_{candidate_t^h} + 1$ 
9:   end for
10:  if  $\max(n) \leq K$  then
11:    head = sample( $0, H$ )
12:    for all  $h$  do
13:       $att_{it}^h = att_{it}^{head}$ 
14:    end for
15:  end if
16: end for

```

date for $head_h$ next step is :

$$candidate_t^h = \arg \max_i att_{it}^h \quad (4)$$

We denote an array of $[n_0, \dots, n_i, \dots, n_{T-1}]$ as the number of times that $word_i$ is chosen from equation 4, where T is the length of source sentence. Obviously:

$$\sum_{i=0}^{T-1} n_i = H \quad (5)$$

where H is the number of heads. Diverse translations are generated when multiple candidates are offered. In other words, not all heads focus on the same source word. Therefore, we define a *confusing* condition when:

$$\max(n_i) \leq K \quad (6)$$

where K is a hyper-parameter. *Confusing* condition means “referred words” are disperse and multiple candidates can be accepted. Under *confusing* condition, we sample one of the heads as attention and force other heads to be the same.

Otherwise, the decoding step remains unchanged. If $K = 0$, the model is the same as the original version. If $K = H$, the model samples at every step. For the sake of balance of quality and diversity, we may choose different K in different conditions. We do the whole decoding for M times and pick the most possible output in the beam every time. In this paper, we let $M = 5$.

Our another contribution is adopting the method with back-translation technique as data augmentation. Back-translation has been proved helpful for neural machine translation (Sennrich, Haddow, and Birch 2016a; Poncelas et al. 2018). However, the lack of diversity restricts its effect (Edunov et al. 2018). We provide a new scheme for back-translation with diverse corpus generated by our method and gain improvement.

Experiment

Our translation experiments include two parts: diverse translation and diverse back-translation. In addition, a conversation response generation experiment is also performed as auxiliary evidence.

Setup

Datasets We choose five datasets as our experiment corpus.

- NIST Chinese-to-English (NIST Zh-En). The training data consists of 1.34 million sentence pairs extracted from LDC corpus. We use MT03 as the development set, MT04, MT05, MT06 as the test sets.
- WMT14 English-to-German (WMT En-De). The training data consists of 4.5 million sentence pairs from WMT14 news translation task. We use newstest2013 as the development set and newstest2014 as the test set.
- WMT16 English-to-Romanian (WMT En-Ro). The training data consists of 0.6 million sentence pairs from WMT16 news translation task. We use newstest2015 as the development set and newstest2016 as the test set.
- Monolingual English corpus (for back-translation) from IWSLT17 Chinese-to-English (IWSLT Zh-En). The training data consists of 0.2 million sentences from IWSIT17 spoken language translation task. We used dev2010 and tst2010 as the development set and tst2011 as the test set.
- Short Text Conversation (STC) (Shang, Lu, and Li 2015). The corpus contains about 4.4 million Chinese post-response sentence pairs crawled from Weibo, built for single turn conversation tasks. We remove sentence pairs that are exactly the same between two sides. For the test set, we extract 3000 post sentences that have 10 responses in the corpus, forming 10 references. The develop set is made up similarly.

For NIST Zh-En, we use BPE (Sennrich, Haddow, and Birch 2016b) with 30K merge operations on both sides. For En-De and En-Ro, we also apply BPE to segment sentences and limit the vocabulary size to 32K. We filter out sentence pairs whose source or target side contains more than 100 words for Zh-En and En-Ro sets. For STC corpus, we also apply BPE and keep a vocabulary size to 36K. All the out-of-vocabulary words are mapped to a distinct token $\langle \text{UNK} \rangle$.

Experiment Settings Without extra statement, we follow the Transformer base v1 settings¹, with 6 layers in encoder and 2 layers in decoder², 512 hidden units, 8 heads in multi-head attention and 2048 hidden units in feed-forward layers. Parameters are optimized using Adam optimizer (Kingma and Ba 2015), with $\beta_1 = 0.9$, $\beta_2 = 0.98$, and $\epsilon = 10^{-9}$. The learning rate is scheduled according to the method proposed in Vaswani et al. (2017), with $warmup_steps = 8000$. Label smoothing (Szegedy et al. 2016) of $value = 0.1$ is also adopted. For K , we do not observe diversity enhancement when K is too small like $K = 1, 2$. And conditions of $K = 6, 7$ are very similar to $K = 8$. Hence we use $K = 3, 4, 5, 8$ as comparisons.

Diverse Translation

Comparing Objects We compare our models with original beam search (as Baseline) and sampling from the probability distribution (as Multinomial Sampling). Besides, we compare our methods with a few previous works:

- Li, Monroe, and Jurafsky (2016) propose a decoding trick to penalize hypotheses that are siblings (expansions of the same parent node) in the beam search to increase the translation diversity.
- Vijayakumar et al. (2018) add a regularization item in beam search to penalize the same word generation.
- Shen et al. (2019) and He, Haffari, and Norouzi (2018) use multiple decoders as mixture of experts to increase diversity by manipulating latent variables. Considering their similarity, we choose Shen et al. (2019) since they report better results. We re-implement the model with Transformer architecture and choose the hMup (online-shared) version since the authors recommend it.

Metrics We evaluate our method from both diversity and quality. For diversity, we adopt average pair-wise BLEU of outputs (denoted as pwb) to measure the difference among translations like previous work. For quality, we use BLEU with the references (denoted as rfb). Lower pwb and higher rfb mean better results. In this paper, the reference BLEU of Baseline is the highest score in the beam while the other groups take the average reference BLEU of M outputs. And to synthetically evaluate the performance, we propose an overall index: Diversity Enhancement per Quality (denoted as DEQ). Specifically:

$$DEQ = \frac{pwb^* - pwb}{rfb^* - rfb} \quad (7)$$

where pwb and pwb^* refer to pair-wise BLEU of the evaluated system and baseline respectively, rfb and rfb^* refer to

¹<https://github.com/tensorflow/tensor2tensor/blob/v1.3.0/tensor2tensor/models/transformer.py>

²We check different decoder layer numbers settings and find less-decoder-layers Transformer shows comparable performance with original six-layers-decoder Transformer while it is much easier to manipulate and faster to decode. The diversity enhancement is also more significant.

Model	MT03 (dev)		MT04		MT05		MT06		Average		DEQ \uparrow
	rfb \uparrow	pwb \downarrow	rfb \uparrow	pwb \downarrow	rfb \uparrow	pwb \downarrow	rfb \uparrow	pwb \downarrow	rfb \uparrow	pwb \downarrow	
Baseline	45.64	84.63	47.25	84.62	43.45	84.78	42.26	82.46	44.32	83.95	-
Multinomial Sampling	21.75	11.29	22.19	11.42	20.54	11.08	19.12	9.67	20.62	10.72	3.09
(Li, Monroe, and Jurafsky 2016)	44.63	80.92	45.81	81.33	42.86	81.28	40.87	78.11	43.18	80.24	3.25
(Vijayakumar et al. 2018)	40.38	59.55	41.99	60.11	39.46	59.56	37.28	54.54	39.58	58.07	5.46
(Shen et al. 2019)	40.59	62.24	41.55	62.68	38.51	61.37	35.57	58.04	38.54	60.70	4.02
Sample $K = 3$	43.73	66.48	45.38	67.82	42.43	65.80	40.18	64.93	42.66	66.18	10.70
Sample $K = 4$	40.88	51.26	42.50	53.63	39.18	51.07	37.73	50.28	39.80	51.66	7.14
Sample $K = 5$	38.60	43.64	40.21	45.69	37.05	43.14	35.45	42.38	37.57	43.74	5.96
Sample $K = 8$	36.68	38.29	38.03	40.02	34.65	37.30	32.93	36.15	35.20	37.82	5.06

Table 3: Pair-wise BLEU and Reference BLEU in Zh2En experiments.

Model	rfb \uparrow	pwb \downarrow	DEQ \uparrow
Baseline	26.31	80.41	-
Multinomial Sampling	11.99	12.84	4.72
(Li, Monroe, and Jurafsky 2016)	25.27	78.57	1.77
(Vijayakumar et al. 2018)	23.27	66.13	4.70
(Shen et al. 2019)	23.22	68.03	4.01
Sample $K = 3$	25.62	78.96	2.10
Sample $K = 4$	24.26	62.04	8.96
Sample $K = 5$	22.62	50.14	8.20
Sample $K = 8$	19.76	38.36	6.42

Table 4: Pair-wise BLEU and Reference BLEU in En2De experiments.

Model	rfb \uparrow	pwb \downarrow	DEQ \uparrow
Baseline	31.76	81.29	-
Multinomial Sampling	18.85	20.82	4.68
(Li, Monroe, and Jurafsky 2016)	31.02	78.42	3.88
(Vijayakumar et al. 2018)	28.91	69.67	4.08
(Shen et al. 2019)	31.07	85.71	-6.04
Sample $K = 3$	31.33	82.41	-2.60
Sample $K = 4$	30.06	71.12	5.98
Sample $K = 5$	27.89	59.42	5.65
Sample $K = 8$	26.43	50.56	5.77

Table 5: Pair-wise BLEU and Reference BLEU in En2Ro experiments.

reference BLEU of the evaluated system and baseline respectively. It measures how much diversity can be produced per quality drop.

Results From Table 3, in Zh2En experiment, we can see that traditional beam search translations severely lack diversity while multinomial sampling extremely harms translation quality. Li, Monroe, and Jurafsky (2016) bring very limited enhancement in diversity, failing to achieve the goal. Vijayakumar et al. (2018) and Shen et al. (2019) show the ability to produce diversity, but our method ($K = 4$) attains more significant diversity as well as better quality comparing with them. And $K = 3$ achieve the highest DEQ , gaining the most satisfactory result. Also, unlike Shen et al. (2019), our work needs no extra training or extra parameters. Furthermore, the diversity can be well interpreted and does not

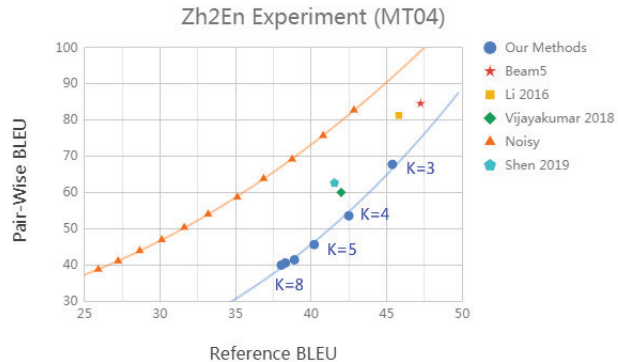


Figure 4: Pair-wise BLEU with reference BLEU in Zh2En Experiments (MT04). The bottom right corner means the best result. All previous work including noisy sets lie on the top left of the curve of K .

rely on an abstract latent variable. See Table 6 for a case. We can reach the similar conclusion in En2De and En2Ro experiments from Table 4, 5.

Besides, to exclude the possibility that randomly interfering causes the effect, we compare with the sets with noise. We add noise to the translations of baseline model to generate different outputs. Specifically, for each sentence, we replace one of its words with ‘<UNK>’ with probability p and randomly swap two words with probability p as well. And we make multiple experiment sets by controlling p . See Figure 4, at the same level of pair-wise BLEU, our method maintains much higher reference BLEU, which means our method improves diversity through seeking diverse translations rather than just generating randomly.

For K (see Figure 4 again), as expected, as K grows (sample more), the diversity increases (pair-wise BLEU decreases) while the quality decreases (reference BLEU decreases). And previous works all lie on the top left of the curve of K . What’s more, we can choose different K to diversely balance the diversity and the quality depending on our needs. We make the trade-off more continuous.

Some may not be satisfied with the sacrifice of the reference BLEU. But considering the calculation of BLEU is based on n-gram rather than semantic similarity, we regard it as a normal phenomenon. After all, if we want to obtain

Input	Liang Ge ZhuJue – ChaoXian He MeiGuo Dou MeiYou BiaoXian Chu RangBu , ShuangFang De JiBen LiChang Ye Dou MeiYou SongDong .
Beam	1. the two leading characters – the dprk and the united states – did not make any concessions , and the basic positions of both sides were not relaxed . 2. the two leading characters – the dprk and the united states – did not make any concessions , and the basic positions of both sides were not loosened . 3. the two leading characters – the dprk and the united states – did not make any concessions , and both sides did not relax their basic positions . 4. the two leading characters – the dprk and the united states – did not make any concessions . both sides did not relax their basic positions . 5. the two leading characters – the dprk and the united states – did not make any concessions , and both sides ' basic positions were not relaxed .
K=4	1. the two leading characters – the dprk and the united states – did not make any concessions . both sides ' basic positions were not relaxed . 2. neither the dprk nor the united states has made any concessions . both sides have not relaxed their basic positions . 3. the two leading roles – the dprk and the united states – have made no concessions , and neither have they relaxed their basic positions . 4. neither the dprk nor the united states – the two leading characters – did make any concessions , and the basic positions of both sides were not relaxed . 5. neither the democratic people 's republic of korea and the united states have made any concessions , and the basic positions of both sides have not been relaxed .

Table 6: One Zh2En case. Our method shows obviously more diversity compared with beam search.

sentences with different grammar structure or word order, the overlap of n-gram will inevitably decrease to some extent even the meaning remains the same. Meanwhile, we empirically prove our method maintains a relatively high quality comparing to noisy sets as well as previous works.

We also investigate the effect of sentence length. Theoretically, longer sentences shall have more diversity due to their broader searching space. However, beam search with MAP prefers to abandon different but slightly less possible candidates, making hypotheses lack diversity and are all close to specific translations. Conversely, our method increases diversity as the sentences getting longer (see figure 5), which conforms to the statistical law.

Diverse Back-Translation

Back-translation has been proved helpful for neural machine translation (Sennrich, Haddow, and Birch 2016a; Poncelas et al. 2018). However, the lack of diversity restricts its effect (Edunov et al. 2018). We try to utilize our methods to enhance the translation performance by improving back-translation. According to Edunov et al. (2018), unrestricted sampling from the model distribution yields the best performance. Therefore, we compare with 1) baseline without utilizing back-translation, denoted as *Baseline*, 2) beam search as back-translation, denoted as *Beam-5*, 3) unrestricted sampling as back-translation, denoted as *Sampling*.

We do experiments under conditions with and without additional monolingual data.

Self Back-Translation Firstly, We focus on the condition where original training data is repeatedly used by back-translation. When translating language pair f to e , for each target sentence e , we get M translations with a reverse translation model. We combine those translations with e as synthetic sentence pairs and add them to the training data. As previously stated, we let $M = 5$. Experiments are conducted on Zh-En NIST dataset.

See Table 7 and 8, all of our experiment sets report better results, among which, the best set of $K = 3$ in Zh2En experiments yields 1.82 improvement and the best set of $K = 4$ in En2Zh experiments yields 0.82 improvement.

Utilizing Additional Monolingual Data Secondly, we evaluate our method with additional monolingual data. We select one side of parallel data from IWSLT17 as monolingual data. We use the same method to generate synthetic sentence pairs. Then we train our model on the mixture of original NIST dataset and the synthetic dataset.

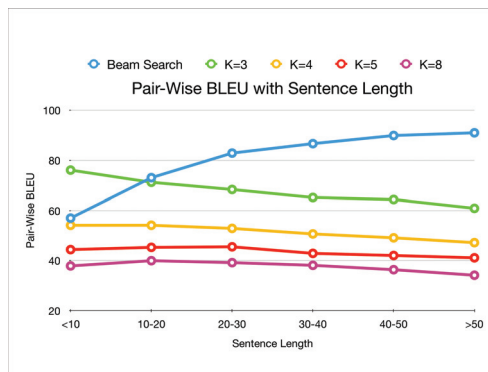


Figure 5: Our method increases diversity (pair-wise BLEU decreases) as the sentences getting longer.

See Table 9, experiment results show that our algorithm brings the most significant improvement for translation performance as our work adds generation diversity and maintains the quality simultaneously.

Model	MT03	MT04	MT05	MT06	Average
Baseline	45.64	47.25	43.45	42.26	44.32
Beam-5	46.31	47.26	44.87	43.43	45.19
Sampling	47.03	47.96	45.72	44.06	45.91
$K = 3$	47.24	48.24	45.70	44.48	46.14
$K = 4$	47.39	47.93	45.38	43.98	45.76
$K = 5$	47.31	48.31	45.34	43.95	45.87
$K = 8$	47.15	48.15	45.69	43.95	45.93

Table 7: Zh2En translation experiments with back-translation of original training data.

Conversation Response Generation

Responses generated by neural conversational models tend to lack informativeness and diversity (Li et al. 2016; Shao et al. 2017; Baheti et al. 2018; Zhang et al. 2018). Therefore, we try to ease this issue by utilizing our method on Conversation Response Generation tasks. Still, we perform decoding for M times and pick the N^{th} output in the beam for the N^{th} group ($N \in [1, M]$).

Metrics Since the responses of human conversation can be pretty subjective, which is hard to evaluate automatically.

Model	MT03	MT04	MT05	MT06	Average
Baseline	22.75	22.33	20.35	21.35	21.34
Beam-5	23.73	21.69	20.61	22.33	21.54
Sampling	23.69	22.78	20.85	22.34	21.99
$K = 3$	24.21	22.23	20.65	22.52	21.80
$K = 4$	24.01	23.15	21.04	22.30	22.16
$K = 5$	23.76	21.93	20.57	22.50	21.67
$K = 8$	23.93	21.66	20.72	22.23	21.54

Table 8: En2Zh translation experiments with back-translation of original training data.

Model	Zh2En
Baseline	9.18
Beam-5	13.06
Sampling	13.38
$K = 3$	14.03
$K = 4$	13.76
$K = 5$	13.66
$K = 8$	13.76

Table 9: Zh2En translation experiments with back-translation with additional monolingual data.

Hence, except for reference BLEU, we also measure the response quality by human evaluation through three indexes: relevance, fluency and informativeness. Relevance reveals how much the responses match the expectation of the question. Fluency means to what extent the translation is well-formed grammatically. Both of them are scored from 1 to 5. Informativeness measures the degree of meaningfulness. We classify responses into two groups, informative and uninformative. Uninformative means the safe answer like “I don’t know” or simply copying from the original post. We then calculate the proportion of the informative groups. For diversity, pair-wise BLEU maintains used.

Results In Table 10, we compare our method with basic Seq2Seq model and Li et al. (2016), which use Maximum Mutual Information (MMI) as the objective function (MMI-antiLM version). On one hand, our method achieves significant improvement in generation diversity. On the other hand, the quality including relevance, fluency and informativeness all rise to some degree. After looking into cases, we suppose it is because original Seq2Seq model tends to generate safe outputs like “I don’t know” or simply copying from the source side. In contrast, our method brings in randomness, reaching broader generation space.

	BLEU \uparrow	Rel \uparrow	Flu \uparrow	Inf \uparrow	Div \downarrow
Baseline	13.06	2.45	4.72	0.604	52.94
MMI	13.39	2.58	4.45	0.639	45.42
$K = 3$	13.48	2.63	4.76	0.678	39.82
$K = 8$	12.73	2.53	4.67	0.652	27.73

Table 10: Conversation Response Generation experiment results on STC dataset.

Related Work

Lack of diversity has been a disturbing problem for neural machine translation. In recent years, a few works put forward some related methods. Li, Monroe, and Jurafsky (2016) proposes a decoding trick to penalize hypotheses that are siblings (expansions of the same parent node) in the beam search to increase the translation diversity. Vijayakumar et al. (2018) adds a regularization item in beam search to penalize the same word generation. He, Haffari, and Norouzi (2018) and Shen et al. (2019) use multiple decoders as different latent components, trying to control the generation by different latent variables. Basically, there are two categories: either to add diverse regularization in beam search or to utilize latent variables. Our method achieves better results than both two categories. And specifically, compared with the latter class, our work needs no extra training or extra parameters. Besides, it is hard to tell what the latent variables exactly represent and why they differ while our method shows a clear explanation that heads align to word candidates.

Apart from machine translation, there are also other works concerning generation diversity, including Visual Question Generation (Jain, Zhang, and Schwing 2017), Conversational Response Generation (Li et al. 2016; Shao et al. 2017; Baheti et al. 2018; Zhang et al. 2018), Paraphrase (Gupta et al. 2018; Xu et al. 2018b), Summarization (Nema et al. 2017) and Text Generation (Guu et al. 2018; Xu et al. 2018a).

As for multi-head attention, Strubell et al. (2018) employ different heads to capture different linguistic features. Tu et al. (2018) introduce disagreement regularization to encourage diversity among attention heads. Li et al. (2019) propose to aggregate information captured by different heads. Yang et al. (2019) model the interactions among attention heads. Raganato and Tiedemann (2018) do an analysis of encoder representation and find there exists dependency relations, syntactic and semantic connections across layers.

Conclusion

In this paper, we discover an internal characteristic of Transformer encoder-decoder multi-head attention that each head aligns to a source word which is a possible candidate to be translated. We take advantage of this phenomenon to generate diverse translations by manipulating heads in particular conditions. Experiments show that our algorithm outperforms previous work and obtain the most satisfactory result of quality and diversity. Besides, the multiple trade-off setting can be adopted diversely depending on different needs. Finally, applications on back-translation as data augmentation and conversation response significantly improve the performance, proving our method effective.

Acknowledgement

Shujian Huang is the corresponding author. This work is supported by the National Key R&D Program of China (No. U1836221, 61772261), the Jiangsu Provincial Research Foundation for Basic Research (No. BK20170074).

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Baheti, A.; Ritter, A.; Li, J.; and Dolan, W. B. 2018. Generating more interesting responses in neural conversation models with distributional constraints. In *EMNLP*.
- Barone, A. V. M.; Helcl, J.; Sennrich, R.; Haddow, B.; and Birch, A. 2017. Deep architectures for neural machine translation. In *WMT*.
- Chen, M. X.; Firat, O.; Bapna, A.; Johnson, M.; Macherey, W.; Foster, G.; Jones, L.; Parmar, N.; Schuster, M.; Chen, Z.; Wu, Y.; and Hughes, M. 2018. The best of both worlds: Combining recent advances in neural machine translation. In *ACL*.
- Edunov, S.; Ott, M.; Auli, M.; and Grangier, D. 2018. Understanding back-translation at scale. In *EMNLP*.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. 2017. Convolutional sequence to sequence learning. In *ICML*.
- Gupta, A.; Agarwal, A.; Singh, P.; and Rai, P. 2018. A deep generative framework for paraphrase generation. In *AAAI*.
- Guu, K.; Hashimoto, T. B.; Oren, Y.; and Liang, P. S. 2018. Generating sentences by editing prototypes. *Transactions of the Association for Computational Linguistics* 6:437–450.
- He, X.; Haffari, G.; and Norouzi, M. 2018. Sequence to sequence mixture model for diverse machine translation. In *CoNLL*.
- Jain, U.; Zhang, Z.; and Schwing, A. G. 2017. Creativity: Generating diverse questions using variational autoencoders. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 5415–5424.
- Kingma, D. P., and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Li, J.; Galley, M.; Brockett, C.; Gao, J.; and Dolan, W. B. 2016. A diversity-promoting objective function for neural conversation models. In *HLT-NAACL*.
- Li, J.; Yang, B.; Dou, Z.-Y.; Wang, X.; Lyu, M. R.; and Tu, Z. 2019. Information aggregation for multi-head attention with routing-by-agreement. In *NAACL*.
- Li, J.; Monroe, W.; and Jurafsky, D. 2016. A simple, fast diverse decoding algorithm for neural generation. *CoRR* abs/1611.08562.
- Luong, T.; Pham, H. Q.; and Manning, C. D. 2015. Effective approaches to attention-based neural machine translation. In *EMNLP*.
- Nema, P.; Khapra, M. M.; Laha, A.; and Ravindran, B. 2017. Diversity driven attention model for query-based abstractive summarization. In *ACL*.
- Ott, M.; Auli, M.; Grangier, D.; and Ranzato, M. 2018. Analyzing uncertainty in neural machine translation. In *ICML*.
- Poncelas, A.; Shterionov, D.; Way, A.; de Buy Wenniger, G. M.; and Passban, P. 2018. Investigating backtranslation in neural machine translation. *CoRR* abs/1804.06189.
- Raganato, A., and Tiedemann, J. 2018. An analysis of encoder representations in transformer-based machine translation. In *BlackboxNLP@EMNLP*.
- 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 *ACL*.
- Shang, L.; Lu, Z.; and Li, H. 2015. Neural responding machine for short-text conversation. In *ACL*.
- Shao, Y.; Gouws, S.; Britz, D.; Goldie, A.; Strophe, B.; and Kurzweil, R. 2017. Generating high-quality and informative conversation responses with sequence-to-sequence models. In *EMNLP*.
- Shen, T.; Ott, M.; Auli, M.; and Ranzato, M. 2019. Mixture models for diverse machine translation: Tricks of the trade. In *ICML*.
- Strubell, E.; Verga, P.; Andor, D.; Weiss, D. I.; and McCallum, A. 2018. Linguistically-informed self-attention for semantic role labeling. In *EMNLP*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 2818–2826.
- Tu, Z.; Yang, B.; Lyu, M. R.; and Zhang, T. 2018. Multi-head attention with disagreement regularization. In *EMNLP*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*.
- Vijayakumar, A. K.; Cogswell, M.; Selvaraju, R. R.; Sun, Q.; Lee, S.; Crandall, D. J.; and Batra, D. 2018. Diverse beam search for improved description of complex scenes. In *AAAI*.
- Xu, J.; Ren, X.; Lin, J.; and Sun, X. 2018a. Diversity-promoting gan: A cross-entropy based generative adversarial network for diversified text generation. In *EMNLP*.
- Xu, Q.; Zhang, J.; Qu, L.; Xie, L.; and Nock, R. 2018b. D-page: Diverse paraphrase generation. *CoRR* abs/1808.04364.
- Yang, B.; Wang, L.; Wong, D.; Chao, L. S.; and Tu, Z. 2019. Convolutional self-attention networks. In *NAACL*.
- Zhang, Y.; Galley, M.; Gao, J.; Gan, Z.; Li, X.; Brockett, C.; and Dolan, W. B. 2018. Generating informative and diverse conversational responses via adversarial information maximization. In *NeurIPS*.