

Multi-Head Attention with Disagreement Regularization

Jian Li^{1,2} Zhaopeng Tu^{3*} Baosong Yang⁴ Michael R. Lyu^{1,2} Tong Zhang³

¹Department of Computer Science and Engineering, The Chinese University of Hong Kong

²Shenzhen Research Institute, The Chinese University of Hong Kong

^{1,2}{jianli, lyu}@cse.cuhk.edu.hk

³Tencent AI Lab

⁴University of Macau

³{zptu, bradymzhang}@tencent.com

⁴nlp2ct.baosong@gmail.com

Abstract

Multi-head attention is appealing for the ability to jointly attend to information from different representation subspaces at different positions. In this work, we introduce a disagreement regularization to explicitly encourage the diversity among multiple attention heads. Specifically, we propose three types of disagreement regularization, which respectively encourage the subspace, the attended positions, and the output representation associated with each attention head to be different from other heads. Experimental results on widely-used WMT14 English \Rightarrow German and WMT17 Chinese \Rightarrow English translation tasks demonstrate the effectiveness and universality of the proposed approach.

1 Introduction

Attention model is now a standard component of the deep learning networks, contributing to impressive results in neural machine translation (Bahdanau et al., 2015; Luong et al., 2015), image captioning (Xu et al., 2015), speech recognition (Chorowski et al., 2015), among many other applications. Recently, Vaswani et al. (2017) introduced a multi-head attention mechanism to capture different context with multiple individual attention functions.

One strong point of multi-head attention is the ability to jointly attend to information from *different* representation subspaces at *different* positions. However, there is no mechanism to guarantee that different attention heads indeed capture distinct features. In response to this problem, we introduce a disagreement regularization term to explicitly encourage the diversity among multiple attention heads. The disagreement regularization

serves as an auxiliary objective to guide the training of the related attention component.

Specifically, we propose three types of disagreement regularization, which are applied to the three key components that refer to the calculation of feature vector using multi-head attention. Two regularization terms are respectively to maximize cosine distances of the input subspaces and output representations, while the last one is to disperse the positions attended by multiple heads with element-wise multiplication of the corresponding attention matrices. The three regularization terms can be either used individually or in combination.

We validate our approach on top of advanced TRANSFORMER model (Vaswani et al., 2017) for both English \Rightarrow German and Chinese \Rightarrow English translation tasks. Experimental results show that our approach consistently improves translation performance across language pairs. One encouraging finding is that TRANSFORMER-BASE with disagreement regularization achieves comparable performance with TRANSFORMER-BIG, while the training speed is nearly twice faster.

2 Background: Multi-Head Attention




Figure 1: Illustration of the multi-head attention, which jointly attends to different representation subspaces (colored boxes) at different positions (darker color denotes higher attention probability).

Attention mechanism aims at modeling the strength of relevance between representation pairs,

* Zhaopeng Tu is the corresponding author of the paper. This work was mainly conducted when Jian Li and Baosong Yang were interning at Tencent AI Lab.

such that a representation is allowed to build a direct relation with another representation. Instead of performing a single attention function, Vaswani et al. (2017) found it is beneficial to capture different context with multiple individual attention functions. Figure 1 shows an example of a two-head attention model. For the query word “Bush”, green and red head pay attention to different positions of “talk” and “Sharon” respectively.

Attention function softly maps a sequence of query $Q = \{Q_1, \dots, Q_N\}$ and a set of key-value pairs $\{K, V\} = \{(K_1, V_1), \dots, (K_M, V_M)\}$ to outputs. More specifically, multi-head attention model first transforms Q , K , and V into H subspaces, with different, learnable linear projections, namely:

$$Q^h, K^h, V^h = QW_h^Q, KW_h^K, VW_h^V,$$

where $\{Q^h, K^h, V^h\}$ are respective the query, key, and value representations of the h -th head. $\{W_h^Q, W_h^K, W_h^V\} \in \mathbb{R}^{d \times d_k}$ denote parameter matrices, d and d_k represent the dimensionality of the model and its subspace. Furthermore, H attention functions are applied in parallel to produce the output states $\{O^1, \dots, O^H\}$, among them:

$$O^h = A^h V^h \text{ with } A^h = \text{softmax}\left(\frac{Q^h K^{hT}}{\sqrt{d_k}}\right).$$

Here A^h is the attention distribution produced by the h -th attention head. Finally, the output states are concatenated to produce the final state.

3 Approach

Multi-head attention allows the model to jointly attend to information from *different* representation subspaces at *different* positions. To further guarantee the diversity, we enlarge the distances among multiple attention heads with disagreement regularization (Section 3.1). Specifically, we propose three types of disagreement regularization to encourage each head vector O^i to be different from other heads (Section 3.2).

3.1 Framework

In this work, we take the machine translation tasks as application. Given a source sentence \mathbf{x} and its translation \mathbf{y} , a neural machine translation model is trained to maximize the conditional translation probability over a parallel training corpus.

We introduce an auxiliary regularization term in order to encourage the diversity among multiple

attention heads. Formally, the training objective is revised as:

$$J(\theta) = \arg \max_{\theta} \left\{ \underbrace{L(\mathbf{y}|\mathbf{x}; \theta)}_{\text{likelihood}} + \lambda * \underbrace{D(\mathbf{a}|\mathbf{x}, \mathbf{y}; \theta)}_{\text{disagreement}} \right\},$$

where \mathbf{a} is the referred attention matrices, λ is a hyper-parameter and is empirically set to 1.0 in this paper. The auxiliary regularization term $D(\cdot)$ guides the related attention component to capture different features from the corresponding projected subspaces.

Note that the introduced regularization term works like $L1$ and $L2$ terms, which do not introduce any new parameters and only influence the training of the standard model parameters.

3.2 Disagreement Regularization

Three types of regularization term, which are applied to three parts of the original multi-head attention, are introduced in this section.

Disagreement on Subspaces (Sub.) This disagreement is designed to maximize the cosine distance between the projected values. Specifically, we first calculate the cosine similarity $\cos(\cdot)$ between the vector pair V^i and V^j in different value subspaces, through the dot product of the normalized vectors¹, which measures the cosine of the angle between V^i and V^j . Thus, the cosine distance is defined as negative similarity, i.e., $-\cos(\cdot)$. Our training objective is to enlarge the average cosine distance among all head pairs. The regularization term is formally expressed as:

$$D_{\text{subspace}} = -\frac{1}{H^2} \sum_{i=1}^H \sum_{j=1}^H \frac{V^i \cdot V^j}{\|V^i\| \|V^j\|}. \quad (1)$$

Disagreement on Attended Positions (Pos.)

Another strategy is to disperse the attended positions predicted by multiple heads. Inspired by the agreement regularization (Liang et al., 2006; Cheng et al., 2016) which encourages multiple alignments to be similar, in this work, we deploy a variant of the original term by introducing an alignment disagreement regularization. Formally, we employ the sum of element-wise multiplication of corresponding matrix cells², to measure the

¹We did not employ the Euler Distance between vectors since we do not care the absolute value in each vector.

²We also used the squared element-wise subtraction of two matrices in our preliminary experiments, and found it underperforms its multiplication counterpart, which is consistent with the results in (Cheng et al., 2016).

similarity between two matrices A^i and A^j of two heads:

$$D_{position} = -\frac{1}{H^2} \sum_{i=1}^H \sum_{j=1}^H |A^i \odot A^j|. \quad (2)$$

Disagreement on Outputs (Out.) This disagreement directly applies regularization on the outputs of each attention head, by maximizing the difference among them. Similar to the *subspace* strategy, we employ negative cosine similarity to measure the distance:

$$D_{output} = -\frac{1}{H^2} \sum_{i=1}^H \sum_{j=1}^H \frac{O^i \cdot O^j}{\|O^i\| \|O^j\|}. \quad (3)$$

4 Related Work

The regularization on attended positions is inspired by agreement learning in prior works, which encourages alignments or hidden variables of multiple models to be similar. Liang et al. (2006) first assigned agreement terms for jointly training word alignment in phrase-based statistic machine translation (Koehn et al., 2003). The idea was further extended into other natural language processing tasks such as grammar induction (Liang et al., 2008). Levinboim et al. (2015) extended the agreement for general bidirectional sequence alignment models with model inevitability regularization. Cheng et al. (2016) further explored the agreement on modeling the source-target and target-source alignments in neural machine translation model. In contrast to the mentioned approaches which assigned agreement terms into loss function, we deploy an alignment disagreement regularization by maximizing the distance among multiple attention heads.

As standard multi-head attention model lacks effective control on the influence of different attention heads, Ahmed et al. (2017) used a weighted mechanism to combine them rather than simple concatenation. As an alternative approach to multi-head attention, Shen et al. (2018a) and Shen et al. (2018b) extended the single relevance score to multi-dimensional attention weights, demonstrating the effectiveness of modeling multiple features for attention networks. Our approach is complementary to theirs: our model encourages the diversity among multiple heads, while theirs enhance the power of each head.

#	Regularization			Speed	BLEU
	Sub.	Pos.	Out.		
1	×	×	×	1.21	24.13
2	✓	×	×	1.15	24.64
3	×	✓	×	1.14	24.42
4	×	×	✓	1.15	24.78
5	✓	×	✓	1.12	24.73
6	✓	✓	×	1.11	24.38
7	✓	✓	✓	1.05	24.60

Table 1: Effect of regularization terms, which are applied to the encoder self-attention only. “Speed” denotes the training speed (steps/second).

5 Experiments

5.1 Setup

To compare with the results reported by previous work (Gehring et al., 2017; Vaswani et al., 2017; Hassan et al., 2018), we conduct experiments on both WMT2017 Chinese⇒English (Zh⇒En) and WMT2014 English⇒German (En⇒De) translation tasks. The Zh⇒En corpus consists of 20M sentence pairs, and the En⇒De corpus consists of 4M sentence pairs. We follow previous work to select the validation and test sets. Byte-pair encoding (BPE) is employed to alleviate the Out-of-Vocabulary problem (Sennrich et al., 2016) with 32K merge operations for both language pairs. We use the case-sensitive 4-gram NIST BLEU score (Papineni et al., 2002) as evaluation metric, and *sign-test* (Collins et al., 2005) for statistical significance test.

We evaluate the proposed approaches on the advanced TRANSFORMER model (Vaswani et al., 2017), and implement on top of an open-source toolkit – THUMT (Zhang et al., 2017). We follow Vaswani et al. (2017) to set the configurations and have reproduced their reported results on the En⇒De task. All the evaluations are conducted on the test sets. We have tested both *Base* and *Big* models, which differ at hidden size (512 vs. 1024) and number of attention heads (8 vs. 16). We study model variations with *Base* model on the Zh⇒En task (Section 5.2 and 5.3), and evaluate overall performance with *Big* model on both Zh⇒En and En⇒De tasks (Section 5.4).

5.2 Effect of Regularization Terms

In this section, we evaluate the impact of different regularization terms on the Zh⇒En task us-

System	Architecture	Zh⇒En		En⇒De	
		Speed	BLEU	Speed	BLEU
<i>Existing NMT systems</i>					
(Wu et al., 2016)	GNMT	n/a	n/a	n/a	26.30
(Gehring et al., 2017)	CONVS2S	n/a	n/a	n/a	26.36
(Vaswani et al., 2017)	TRANSFORMER-BASE	n/a	n/a	n/a	27.3
	TRANSFORMER-BIG	n/a	n/a	n/a	28.4
(Hassan et al., 2018)	TRANSFORMER-BIG	n/a	24.2	n/a	n/a
<i>Our NMT systems</i>					
<i>this work</i>	TRANSFORMER-BASE	1.21	24.13	1.28	27.64
	+ Disagreement	1.06	24.85 [↑]	1.10	28.51 [↑]
	TRANSFORMER-BIG	0.58	24.56	0.61	28.58
	+ Disagreement	0.47	25.08 [↑]	0.51	29.28 [↑]

Table 3: Comparing with existing NMT systems on WMT17 Chinese⇒English and WMT14 English⇒German translation tasks. “[↑]” indicates that the model is significantly better than its baseline counterpart ($p < 0.01$).

Applying to			Speed	BLEU
<i>Enc</i>	<i>E-D</i>	<i>Dec</i>		
×	×	×	1.21	24.13
✓	×	×	1.15	24.78
✓	✓	×	1.10	24.67
✓	×	✓	1.11	24.69
✓	✓	✓	1.06	24.85

Table 2: Effect of regularization on different attention networks, i.e., encoder self-attention (“*Enc*”), encoder-decoder attention (“*E-D*”), and decoder self-attention (“*Dec*”).

ing TRANSFORMER-BASE. For simplicity and efficiency, here we only apply regularizations on the encoder side. As shown in Table 1, all the models with the proposed disagreement regularizations (Rows 2-4) consistently outperform the vanilla TRANSFORMER (Row 1). Among them, the *Output* term performs best which is +0.65 BLEU score better than the baseline model, the *Position* term is less effective than the other two. In terms of training speed, we do not observe obvious decrease, which in turn demonstrates the advantage of our disagreement regularizations.

However, the combinations of different disagreement regularizations fail to further improve translation performance (Rows 5-7). One possible reason is that different regularization terms have overlapped guidance, and thus combining them does not introduce too much new information while makes training more difficult.

5.3 Effect on Different Attention Networks

The TRANSFORMER consists of three attention networks, including encoder self-attention, decoder self-attention, and encoder-decoder attention. In this experiment, we investigate how each attention network benefits from the disagreement regularization. As seen from Table 2, all models consistently improve upon the baseline model. When applying disagreement regularization to all three attention networks, we achieve the best performance, which is +0.72 BLEU score better than the baseline model. The training speed decreases by 12%, which is acceptable considering the performance improvement.

5.4 Main Results

Finally, we validate the proposed disagreement regularization on both WMT17 Chinese-to-English and WMT14 English-to-German translation tasks. Specifically, we adopt the *Output* disagreement regularization, which is applied to all three attention networks. The results are concluded in Table 3. We can see that our implementation of TRANSFORMER outperforms all existing NMT systems, and matches the results of TRANSFORMER reported in previous works. Incorporating disagreement regularization consistently improves translation performance for both base and big TRANSFORMER models across language pairs, demonstrating the effectiveness of the proposed approach. It is encouraging to see that TRANSFORMER-BASE with disagreement regu-

Regularization on	Disagreement on		
	<i>Sub.</i>	<i>Pos.</i>	<i>Out.</i>
n/a	0.882	0.007	0.881
Subspace	0.999	0.006	0.935
Position	0.882	0.219	0.882
Output	0.989	0.006	0.997

Table 4: Effect of different regularization terms on the three disagreement measurements. “n/a” denotes the baseline model without any regularization term. Larger value denotes more disagreement (at most 1.0).

Reg.	Layer					
	1	2	3	4	5	6
n/a	0.040	0.009	0.002	0.003	0.008	0.006
Sub.	0.039	0.009	0.001	0.003	0.006	0.005
Pos.	0.217	0.167	0.219	0.242	0.233	0.249
Out.	0.048	0.009	0.002	0.003	0.008	0.006

Table 5: Disagreement on attended positions with respect to the levels of the encoder layers.

larization achieves comparable performance with TRANSFORMER-BIG, while the training speed is nearly twice faster.

5.5 Quantitative Analysis of Regularization

In this section, we empirically investigate how the regularization terms affect the multi-head attention. To this end, we compare the disagreement scores on subspaces (“Sub.”), attended positions (“Pos.”), and outputs (“Out.”). Since the scores are negative values, we list $\exp(D)$ for readability, which has a maximum value of 1.0. Table 4 lists the results of encoder-side multi-head attention on the Zh \Rightarrow En validation set. As seen, the disagreement score on the individual component indeed increases with the corresponding regularization term. For example, the disagreement of outputs increases to almost 1.0 by using the *Output* regularization, which means that the output vectors are almost perpendicular to each other as we measure the cosine distance as the disagreement.

One interesting finding is that attending to different positions may not be the essential strength of multi-head attention on the translation task. As seen, the disagreement score on the attended positions for the standard multi-head attention is only 0.007, which indicates that almost all the heads attend to the same positions. Table 5 shows the disagreement scores on attended positions across en-

coder layers. Except for the 1st layer that attends to the input word embeddings, the disagreement scores on other layers (i.e. ranging from the 2nd to 6th layer) are very low, which confirms our above hypothesis.

Concerning the regularization terms, except that on position, the other two regularization terms (i.e. “Sub.” and “Out.”) do not increase the disagreement score on the attended positions. This can explain why positional regularization term does not work well with the other two terms, as shown in Table 1. This is also consistent with the finding in (Tu et al., 2016), which indicates that neural networks can model linguistic information in their own way. In contrast to attended positions, it seems that the multi-head attention prefer to encoding the differences among multiple heads in the learned representations.

6 Conclusion

In this work, we propose several disagreement regularizations to augment the multi-head attention model, which encourage the diversity among attention heads so that different head can learn distinct features. Experimental results across language pairs validate the effectiveness of the proposed approaches.

The models also suggest a wide range of potential advantages and extensions, from being able to improve the performance of multi-head attention in other tasks such as reading comprehension and language inference, to being able to combine with other techniques (Shaw et al., 2018; Shen et al., 2018a; Dou et al., 2018; Yang et al., 2018) to further improve performance.

Acknowledgments

The work was supported by the National Natural Science Foundation of China (Project No. 61332010 and No. 61472338), the Research Grants Council of the Hong Kong Special Administrative Region, China (No. CUHK 14234416 of the General Research Fund), and Microsoft Research Asia (2018 Microsoft Research Asia Collaborative Research Award). We thank the anonymous reviewers for their insightful comments.

References

Karim Ahmed, Nitish Shirish Keskar, and Richard Socher. 2017. Weighted transformer net-

- work for machine translation. *arXiv preprint arXiv:1711.02132*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*.
- Yong Cheng, Shiqi Shen, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Agreement-based joint training for bidirectional attention-based neural machine translation. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio. 2015. Attention-based Models for Speech Recognition. In *Advances in neural information processing systems (NIPS)*.
- M. Collins, P. Koehn, and I. Kučerová. 2005. Clause restructuring for statistical machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ziyi Dou, Zhaopeng Tu, Xing Wang, Shuming Shi, and Tong Zhang. 2018. Exploiting Deep Representations for Neural Machine Translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, et al. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv preprint arXiv:1803.05567*.
- Philipp Koehn, Franz Josef Och, and Daniel Marcu. 2003. Statistical phrase-based translation. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Tomer Levinboim, Ashish Vaswani, and David Chiang. 2015. Model invertibility regularization: Sequence alignment with or without parallel data. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Percy Liang, Dan Klein, and Michael I Jordan. 2008. Agreement-based Learning. In *Advances in neural information processing systems (NIPS)*.
- Percy Liang, Ben Taskar, and Dan Klein. 2006. Alignment by agreement. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-Attention with Relative Position Representations. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, Shirui Pan, and Chengqi Zhang. 2018a. DiSAN: Directional Self-Attention Network for RNN/CNN-free Language Understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, and Chengqi Zhang. 2018b. Bi-directional Block Self-Attention for Fast and Memory-Efficient Sequence Modeling. In *International Conference on Learning Representations (ICLR)*.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling coverage for neural machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems (NIPS)*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C Courville, Ruslan Salakhutdinov, Richard S Zemel, and Yoshua Bengio. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Baosong Yang, Zhaopeng Tu, Derek F. Wong, Fandong Meng, Lidia S. Chao, and Tong Zhang. 2018. Modeling Localness for Self-Attention Networks. In

Conference on Empirical Methods in Natural Language Processing (EMNLP).

Jiacheng Zhang, Yanzhuo Ding, Shiqi Shen, Yong Cheng, Maosong Sun, Huanbo Luan, and Yang Liu. 2017. THUMT: An Open Source Toolkit for Neural Machine Translation. *arXiv preprint arXiv:1706.06415*.