

Convolutional Self-Attention Networks

Baosong Yang[†] Longyue Wang[‡] Derek F. Wong[†] Lidia S. Chao[†] Zhaopeng Tu^{‡*}

[†]NLP²CT Lab, Department of Computer and Information Science, University of Macau

[‡]Tencent AI Lab

[†]nlp2ct.baosong@gmail.com, {derekfw, lidiasc}@umac.mo

[‡]{vinnylywang, zptu}@tencent.com

Abstract

Self-attention networks (SANs) have drawn increasing interest due to their high parallelization in computation and flexibility in modeling dependencies. SANs can be further enhanced with multi-head attention by allowing the model to attend to information from different representation subspaces. In this work, we propose novel *convolutional self-attention networks*, which offer SANs the abilities to 1) strengthen dependencies among neighboring elements, and 2) model the interaction between features extracted by multiple attention heads. Experimental results of machine translation on different language pairs and model settings show that our approach outperforms both the strong Transformer baseline and other existing models on enhancing the locality of SANs. Comparing with prior studies, the proposed model is parameter free in terms of introducing no more parameters.

1 Introduction

Self-attention networks (SANs) (Parikh et al., 2016; Lin et al., 2017) have shown promising empirical results in various natural language processing (NLP) tasks, such as machine translation (Vaswani et al., 2017), natural language inference (Shen et al., 2018a), and acoustic modeling (Sperber et al., 2018). One appealing strength of SANs lies in their ability to capture dependencies regardless of distance by explicitly attending to all the elements. In addition, the performance of SANs can be improved by multi-head attention (Vaswani et al., 2017), which projects the input sequence into multiple subspaces and applies attention to the representation in each subspace.

Despite their success, SANs have two major limitations. First, the model fully take into ac-

count all the elements, which disperses the attention distribution and thus overlooks the relation of neighboring elements and phrasal patterns (Yang et al., 2018; Wu et al., 2018; Guo et al., 2019). Second, multi-head attention extracts distinct linguistic properties from each subspace in a parallel fashion (Raganato and Tiedemann, 2018), which fails to exploit useful interactions across different heads. Recent work shows that better features can be learned if different sets of representations are present at feature learning time (Ngiam et al., 2011; Lin et al., 2014).

To this end, we propose novel *convolutional self-attention networks* (CSANs), which model locality for self-attention model and interactions between features learned by different attention heads in a unified framework. Specifically, in order to pay more attention to a local part of the input sequence, we restrict the attention scope to a window of neighboring elements. The localness is therefore enhanced via a parameter-free 1-dimensional convolution. Moreover, we extend the convolution to a 2-dimensional area with the axis of attention head. Thus, the proposed model allows each head to interact local features with its adjacent subspaces at attention time. We expect that the interaction across different subspaces can further improve the performance of SANs.

We evaluate the effectiveness of the proposed model on three widely-used translation tasks: WMT14 English-to-German, WMT17 Chinese-to-English, and WAT17 Japanese-to-English. Experimental results demonstrate that our approach consistently improves performance over the strong TRANSFORMER model (Vaswani et al., 2017) across language pairs. Comparing with previous work on modeling locality for SANs (e.g. Shaw et al., 2018; Yang et al., 2018; Sperber et al., 2018), our model boosts performance on both translation quality and training efficiency.

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




Figure 1: Illustration of (a) vanilla SANs; (b) 1-dimensional convolution with the window size being 3; and (c) 2-dimensional convolution with the area being 3×3 . Different colors and patterns represent different subspaces modeled by multi-head attention, and transparent colors denote masked tokens that are invisible to SANs.

2 Multi-Head Self-Attention Networks

SANs produce representations by applying attention to each pair of tokens from the input sequence, regardless of their distance. Vaswani et al. (2017) found it is beneficial to capture different contextual features with multiple individual attention functions. Given an input sequence $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_I\} \in \mathbb{R}^{I \times d}$, the model first transforms it into queries \mathbf{Q} , keys \mathbf{K} , and values \mathbf{V} :

$$\mathbf{Q}, \mathbf{K}, \mathbf{V} = \mathbf{X}\mathbf{W}_Q, \mathbf{X}\mathbf{W}_K, \mathbf{X}\mathbf{W}_V \in \mathbb{R}^{I \times d} \quad (1)$$

where $\{\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V\} \in \mathbb{R}^{d \times d}$ are trainable parameters and d indicates the hidden size. The three types of representations are split into H different subspaces, e.g., $[\mathbf{Q}^1, \dots, \mathbf{Q}^H] = \mathbf{Q}$ with $\mathbf{Q}^h \in \mathbb{R}^{I \times \frac{d}{H}}$. In each subspace h , the element \mathbf{o}_i^h in the output sequence $\mathbf{O}^h = \{\mathbf{o}_1^h, \dots, \mathbf{o}_I^h\}$ is calculated by

$$\mathbf{o}_i^h = \text{ATT}(\mathbf{q}_i^h, \mathbf{K}^h)\mathbf{V}^h \in \mathbb{R}^{\frac{d}{H}} \quad (2)$$

where $\text{ATT}(\cdot)$ is an attention model (Bahdanau et al., 2015; Vaswani et al., 2017) that retrieves the keys \mathbf{K}^h with the query \mathbf{q}_i^h . The final output representation \mathbf{O} is the concatenation of outputs generated by multiple attention models:

$$\mathbf{O} = [\mathbf{O}^1, \dots, \mathbf{O}^H] \in \mathbb{R}^{I \times d} \quad (3)$$

3 Approach

As shown in Figure 1(a), the vanilla SANs use the query \mathbf{q}_i^h to compute a categorical distribution over all elements from \mathbf{K}^h (Equation 2). It may inherit the attention to neighboring information (Yu et al., 2018; Yang et al., 2018; Guo et al., 2019). In this work, we propose to model locality for SANs by restricting the model to attend to a local region via convolution operations (1D-CSANs, Figure 1(b)). Accordingly, it provides distance-aware

information (e.g. phrasal patterns), which is complementary to the distance-agnostic dependencies modeled by the standard SANs (Section 3.1).

Moreover, the calculation of output \mathbf{o}^h are restricted to the a single individual subspace, overlooking the richness of contexts and the dependencies among groups of features, which have proven beneficial to the feature learning (Ngiam et al., 2011; Wu and He, 2018). We thus propose to convolute the items in adjacent heads (2D-CSANs, Figure 1(c)). The proposed model is expected to improve performance through interacting linguistic properties across heads (Section 3.2).

3.1 Locality Modeling via 1D Convolution

For each query \mathbf{q}_i^h , we restrict its attention region (e.g., $\mathbf{K}^h = \{\mathbf{k}_1^h, \dots, \mathbf{k}_i^h, \dots, \mathbf{k}_I^h\}$) to a local scope with a fixed size $M + 1$ ($M \leq I$) centered at the position i :

$$\widehat{\mathbf{K}}^h = \{\mathbf{k}_{i-\frac{M}{2}}^h, \dots, \mathbf{k}_i^h, \dots, \mathbf{k}_{i+\frac{M}{2}}^h\} \quad (4)$$

$$\widehat{\mathbf{V}}^h = \{\mathbf{v}_{i-\frac{M}{2}}^h, \dots, \mathbf{v}_i^h, \dots, \mathbf{v}_{i+\frac{M}{2}}^h\} \quad (5)$$

Accordingly, the calculation of corresponding output in Equation (2) is modified as:

$$\mathbf{o}_i^h = \text{ATT}(\mathbf{q}_i^h, \widehat{\mathbf{K}}^h)\widehat{\mathbf{V}}^h \quad (6)$$

As seen, SANs are only allowed to attend to the neighboring tokens (e.g., $\widehat{\mathbf{K}}^h, \widehat{\mathbf{V}}^h$), instead of all the tokens in the sequence (e.g., $\mathbf{K}^h, \mathbf{V}^h$).

The SAN-based models are generally implemented as multiple layers, in which higher layers tend to learn semantic information while lower layers capture surface and lexical information (Peters et al., 2018; Raganato and Tiedemann, 2018). Therefore, we merely apply locality modeling to the lower layers, which same to the configuration in Yu et al. (2018) and Yang et al. (2018). In

this way, the representations are learned in a hierarchical fashion (Yang et al., 2017). That is, the distance-aware and local information extracted by the lower SAN layers, is expected to complement distance-agnostic and global information captured by the higher SAN layers.

3.2 Attention Interaction via 2D Convolution

Mutli-head mechanism allows different heads to capture distinct linguistic properties (Raganato and Tiedemann, 2018; Li et al., 2018), especially in diverse local contexts (Yang et al., 2018). We hypothesize that exploiting local properties across heads can further improve the performance of SANs. To this end, we expand the 1-dimensional window to a 2-dimensional area with the new dimension being the index of attention head. Suppose that the area size is $(N + 1) \times (M + 1)$ ($N \leq H$), the keys and values in the area are:

$$\tilde{\mathbf{K}}^h = \bigcup[\hat{\mathbf{K}}^{h-\frac{N}{2}}, \dots, \hat{\mathbf{K}}^h, \dots, \hat{\mathbf{K}}^{h+\frac{N}{2}}] \quad (7)$$

$$\tilde{\mathbf{V}}^h = \bigcup[\hat{\mathbf{V}}^{h-\frac{N}{2}}, \dots, \hat{\mathbf{V}}^h, \dots, \hat{\mathbf{V}}^{h+\frac{N}{2}}] \quad (8)$$

where $\hat{\mathbf{K}}^h, \hat{\mathbf{V}}^h$ are elements in the h -th subspace, which are calculated by Equations 4 and 5 respectively. The union operation \bigcup means combining the keys and values in different subspaces. The corresponding output is calculated as:

$$\mathbf{o}_i^h = \text{ATT}(\mathbf{q}_i^h, \tilde{\mathbf{K}}^h) \tilde{\mathbf{V}}^h \quad (9)$$

The 2D convolution allows SANs to build relevance between elements across adjacent heads, thus flexibly extract local features from different subspaces rather than merely from an unique head.

The vanilla SAN models linearly aggregate features from different heads, and this procedure limits the extent of abstraction (Fukui et al., 2016; Li et al., 2019). Multiple sets of representations presented at feature learning time can further improve the expressivity of the learned features (Ngiam et al., 2011; Wu and He, 2018).

4 Related Work

Self-Attention Networks Recent studies have shown that SANs can be further improved by capturing complementary information. For example, Hao et al. (2019) complemented SANs with recurrence modeling, while Yang et al. (2019) modeled contextual information for SANs.

Concerning modeling locality for SANs, Yu et al. (2018) injected several CNN layers (Kim,

2014) to fuse local information, the output of which is fed to the subsequent SAN layer. Several researches proposed to revise the attention distribution with a parametric localness bias, and succeed on machine translation (Yang et al., 2018) and natural language inference (Guo et al., 2019). While both models introduce additional parameters, our approach is a more lightweight solution without introducing any new parameters. Closely related to this work, Shen et al. (2018a) applied a positional mask to encode temporal order, which only allows SANs to attend to the previous or following tokens in the sequence. In contrast, we employ a positional mask (i.e. the tokens outside the local window is masked as 0) to encode the distance-aware local information.

In the context of distance-aware SANs, Shaw et al. (2018) introduced relative position encoding to consider the relative distances between sequence elements. While they modeled locality from position embedding, we improve locality modeling from revising attention scope. To make a fair comparison, we re-implemented the above approaches under a same framework. Empirical results on machine translation tasks show the superiority of our approach in both translation quality and training efficiency.

Multi-Head Attention Multi-head attention mechanism (Vaswani et al., 2017) employs different attention heads to capture distinct features (Raganato and Tiedemann, 2018). Along this direction, Shen et al. (2018a) explicitly used multiple attention heads to model different dependencies of the same word pair, and Strubell et al. (2018) employed different attention heads to capture different linguistic features. Li et al. (2018) introduced disagreement regularizations to encourage the diversity among attention heads. Inspired by recent successes on fusing information across layers (Dou et al., 2018, 2019), Li et al. (2019) proposed to aggregate information captured by different attention heads. Based on these findings, we model interactions among attention heads to exploit the richness of local properties distributed in different heads.

5 Experiments

We conducted experiments with the Transformer model (Vaswani et al., 2017) on English \Rightarrow German (En \Rightarrow De), Chinese \Rightarrow English (Zh \Rightarrow En) and Japanese \Rightarrow English (Ja \Rightarrow En) translation tasks.


For the $\text{En} \Rightarrow \text{De}$ and $\text{Zh} \Rightarrow \text{En}$ tasks, the models were trained on widely-used WMT14 and WMT17 corpora, consisting of around 4.5 and 20.62 million sentence pairs, respectively. Concerning $\text{Ja} \Rightarrow \text{En}$, we used the first two sections of WAT17 corpus as the training data, which consists of 2M sentence pairs. To reduce the vocabulary size, all the data were tokenized and segmented into subword symbols using byte-pair encoding (Sennrich et al., 2016) with 32K merge operations. Following Shaw et al. (2018), we incorporated the proposed model into the encoder, which is a stack of 6 SAN layers. Prior studies revealed that modeling locality in lower layers can achieve better performance (Shen et al., 2018b; Yu et al., 2018; Yang et al., 2018), we applied our approach to the lowest three layers of the encoder. About configurations of NMT models, we used the *Base* and *Big* settings same as Vaswani et al. (2017), and all models were trained on 8 NVIDIA P40 GPUs with a batch of 4096 tokens.

5.1 Effects of Window/Area Size


We first investigated the effects of window size (1D-CSANs) and area size (2D-CSANs) on $\text{En} \Rightarrow \text{De}$ validation set, as plotted in Figure 2. For 1D-CSANs, the local size with 11 is superior to other settings. This is consistent with Luong et al. (2015) who found that 10 is the best window size in their local attention experiments. Then, we fixed the number of neighboring tokens being 11 and varied the number of heads. As seen, by considering the features across heads (i.e. > 1), 2D-CSANs further improve the translation quality. However, when the number of heads in attention goes up, the translation quality inversely drops. One possible reason is that the model still has the flexibility of learning a different distribution for each head with few interactions, while a large amount of interactions assumes more heads make “similar contributions” (Wu and He, 2018).

5.2 Accuracy of Phrase Translation

One intuition of our approach is to capture useful phrasal patterns via modeling locality. To evaluate the accuracy of phrase translations, we calculate the improvement of the proposed approaches over multiple granularities of n-grams, as shown in Figure 3. Both the two model variations consistently outperform the baseline on larger granularities, indicating that modeling locality can raise the ability of self-attention model on capturing the



(a) 1D-CSANs



(b) 2D-CSANs

Figure 2: Effects of (a) window size on 1D-CSANs, and (b) attended head numbers on 2D-CSANs. For 2D-CSANs, the window size dimension is fixed to be 11.




Figure 3: Performance improvement on different n-grams. “Gap of BLEU” denotes the improvement achieved by the proposed models over the baseline.

phrasal information. Furthermore, the dependencies among heads can be complementary to the localness modeling, which reveals the necessity of the interaction of features in different subspaces.

5.3 Comparison to Related Work

We re-implemented and compared several exiting works (Section 4) upon the same framework. Table 1 lists the results on the $\text{En} \Rightarrow \text{De}$ translation task. As seen, all the models improve translation quality, reconfirming the necessity of modeling locality and distance information. Besides, our models outperform all the existing works, indicating the superiority of the proposed approaches. In particular, CSANs achieve better performance than

Model	Parameter	Speed	BLEU	Δ
TRANSFORMER-BASE (Vaswani et al., 2017)	88.0M	1.28	27.31	-
+ BI_DIRECT (Shen et al., 2018a)	+0.0M	-0.00	27.58	+0.27
+ REL_POS (Shaw et al., 2018)	+0.1M	-0.11	27.63	+0.32
+ NEIGHBOR (Sperber et al., 2018)	+0.4M	-0.06	27.60	+0.29
+ LOCAL_HARD (Luong et al., 2015)	+0.4M	-0.06	27.73	+0.42
+ LOCAL_SOFT (Yang et al., 2018)	+0.8M	-0.09	27.81	+0.50
+ BLOCK (Shen et al., 2018b)	+6.0M	-0.33	27.59	+0.28
+ CNNs (Yu et al., 2018)	+42.6M	-0.54	27.70	+0.39
+ 1D-CSANS	+0.0M	-0.00	27.86	+0.55
+ 2D-CSANS	+0.0M	-0.06	28.18	+0.87

Table 1: Comparing with the existing approaches on WMT14 En \Rightarrow De translation task. For a fair comparison, we re-implemented the existing locality approaches under the same framework. ‘‘Parameter’’ denotes the number of model parameters (M = million) and ‘‘Speed’’ denotes the training speed (steps/second). ‘‘ Δ ’’ column denotes performance improvements over the Transformer baseline.

Model	WMT14 En \Rightarrow De		WMT17 Zh \Rightarrow En		WAT17 Ja \Rightarrow En	
	Speed	BLEU	Speed	BLEU	Speed	BLEU
TRANSFORMER-BASE	1.28	27.31	1.21	24.13	1.33	28.10
+ CSANS	1.22	28.18 \uparrow	1.16	24.80 \uparrow	1.28	28.50 \uparrow
TRANSFORMER-BIG	0.61	28.58	0.58	24.56	0.65	28.41
+ CSANS	0.50	28.74	0.48	25.01 \uparrow	0.55	28.73 \uparrow

Table 2: Experimental results on WMT14 En \Rightarrow De, WMT17 Zh \Rightarrow En and WAT17 Ja \Rightarrow En test sets. ‘‘Speed’’ denotes the training speed (steps/second). ‘‘ \uparrow / \uparrow ’’ indicates statistically significant difference from the vanilla self-attention counterpart ($p < 0.05/0.01$), tested by bootstrap resampling (Koehn, 2004).

CNNs, revealing that extracting local features with dynamic weights is superior to assigning fixed parameters. Moreover, while most of the existing approaches (except for Shen et al. (2018a)) introduce new parameters, our methods are parameter-free and thus only marginally affect training efficiency.

5.4 Universality of The Proposed Model

To validate the universality of our approach on MT tasks, we evaluated the proposed approach on different language pairs and model settings. Table 2 lists the results on En \Rightarrow De, Zh \Rightarrow En and Ja \Rightarrow En translation tasks. As seen, our model consistently improves translation quality across language pairs, which demonstrates the effectiveness and universality of the proposed approach. It is encouraging to see that CSANS with *base* setting yields comparable performance with TRANSFORMER-BIG.

6 Conclusion

In this paper, we propose a parameter-free convolutional self-attention model to enhance the feature extraction of neighboring elements across

multiple heads. Empirical results of machine translation task on a variety of language pairs demonstrate the effectiveness and universality of the proposed methods. The extensive analyses suggest that: 1) modeling locality is beneficial to SANS; 2) interacting features across multiple heads at attention time can further improve the performance; and 3) to some extent, the dynamic weights are superior to their fixed counterpart (i.e. CSANS vs. CNNs) on local feature extraction. Moreover, it is interesting to validate the proposed model in other sequence modeling tasks.

Acknowledgments

The work was partly supported by the National Natural Science Foundation of China (Grant No. 61672555), the Joint Project of Macao Science and Technology Development Fund and National Natural Science Foundation of China (Grant No. 045/2017/AFJ) and the Multiyear Research Grant from the University of Macau (Grant No. MYRG2017-00087-FST). We thank the anonymous reviewers for their insightful comments.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *ICLR*.
- Ziyi Dou, Zhaopeng Tu, Xing Wang, Shuming Shi, and Tong Zhang. 2018. Exploiting Deep Representations for Neural Machine Translation. In *EMNLP*.
- Ziyi Dou, Zhaopeng Tu, Xing Wang, Longyue Wang, Shuming Shi, and Tong Zhang. 2019. Dynamic Layer Aggregation for Neural Machine Translation. In *AAAI*.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. Multimodal Compact Bilinear Pooling for Visual Question Answering and Visual Grounding. In *EMNLP*.
- Maosheng Guo, Yu Zhang, and Ting Liu. 2019. Gaussian Transformer: A Lightweight Approach for Natural Language Inference. In *AAAI*.
- Jie Hao, Xing Wang, Baosong Yang, Longyue Wang, Jinfeng Zhang, and Zhaopeng Tu. 2019. Modeling Recurrence for Transformer. In *NAACL*.
- Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *EMNLP*.
- Philipp Koehn. 2004. Statistical Significance Tests for Machine Translation Evaluation. In *EMNLP*.
- Jian Li, Zhaopeng Tu, Baosong Yang, Michael R. Lyu, and Tong Zhang. 2018. Multi-Head Attention with Disagreement Regularization. In *EMNLP*.
- Jian Li, Baosong Yang, Zi-Yi Dou, Xing Wang, Michael R. Lyu, and Zhaopeng Tu. 2019. Information Aggregation for Multi-Head Attention with Routing-by-Agreement. In *NAACL*.
- Min Lin, Qiang Chen, and Shuicheng Yan. 2014. Network in Network. In *ICLR*.
- Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A Structured Self-Attentive Sentence Embedding. In *ICLR*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective Approaches to Attention-based Neural Machine Translation. In *EMNLP*.
- Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. 2011. Multimodal Deep Learning. In *ICML*.
- Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A Decomposable Attention Model for Natural Language Inference. In *EMNLP*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *NAACL*.
- Alessandro Raganato and Jörg Tiedemann. 2018. An Analysis of Encoder Representations in Transformer-Based Machine Translation. In *EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In *ACL*.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-Attention with Relative Position Representations. In *NAACL*.
- 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 *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 *ICLR*.
- Matthias Sperber, Jan Niehues, Graham Neubig, Sebastian Stüker, and Alex Waibel. 2018. Self-Attentional Acoustic Models. *Interspeech*.
- Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. Linguistically-Informed Self-Attention for Semantic Role Labeling. In *EMNLP*.
- 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 *NIPS*.
- Wei Wu, Houfeng Wang, Tianyu Liu, and Shuming Ma. 2018. Phrase-level Self-Attention Networks for Universal Sentence Encoding. In *EMNLP*.
- Yuxin Wu and Kaiming He. 2018. Group Normalization. *arXiv:1803.08494*.
- Baosong Yang, Jian Li, Derek F. Wong, Lidia S. Chao, Xing Wang, and Zhaopeng Tu. 2019. Context-Aware Self-Attention Networks. In *AAAI*.
- Baosong Yang, Zhaopeng Tu, Derek F. Wong, Fandong Meng, Lidia S. Chao, and Tong Zhang. 2018. Modeling Localness for Self-Attention Networks. In *EMNLP*.
- Baosong Yang, Derek F Wong, Tong Xiao, Lidia S Chao, and Jingbo Zhu. 2017. Towards Bidirectional Hierarchical Representations for Attention-based Neural Machine Translation. In *EMNLP*.
- Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. QANet: Combining Local Convolution with Global Self-attention for Reading Comprehension. In *ICLR*.