

Exploiting Sentential Context for Neural Machine Translation

Xing Wang

Tencent AI Lab

brightxwang@tencent.com

Zhaopeng Tu

Tencent AI Lab

zptu@tencent.com

Longyue Wang

Tencent AI Lab

vinnylywang@tencent.com

Shuming Shi

Tencent AI Lab

shumingshi@tencent.com

Abstract

In this work, we present novel approaches to exploit sentential context for neural machine translation (NMT). Specifically, we first show that a *shallow sentential context* extracted from the top encoder layer only, can improve translation performance via contextualizing the encoding representations of individual words. Next, we introduce a *deep sentential context*, which aggregates the sentential context representations from all the internal layers of the encoder to form a more comprehensive context representation. Experimental results on the WMT14 English \Rightarrow German and English \Rightarrow French benchmarks show that our model consistently improves performance over the strong TRANSFORMER model (Vaswani et al., 2017), demonstrating the necessity and effectiveness of exploiting sentential context for NMT.

1 Introduction

Sentential context, which involves deep syntactic and semantic structure of the source and target languages (Nida, 1969), is crucial for machine translation. In statistical machine translation (SMT), the sentential context has proven beneficial for predicting local translations (Meng et al., 2015; Zhang et al., 2015). The exploitation of sentential context in neural machine translation (NMT, Bahdanau et al., 2015), however, is not well studied. Recently, Lin et al. (2018) showed that the translation at each time step should be conditioned on the whole *target-side* context. They introduced a deconvolution-based decoder to provide the global information from the target-side context for guidance of decoding.

In this work, we propose simple yet effective approaches to exploiting *source-side* global sentence-level context for NMT models. We use encoder representations to represent the source-

side context, which are summarized into a *sentential context* vector. The source-side context vector is fed to the decoder, so that translation at each step is conditioned on the whole *source-side* context. Specifically, we propose two types of sentential context: 1) the *shallow* one that only exploits the top encoder layer, and 2) the *deep* one that aggregates the sentence representations of all the encoder layers. The deep sentential context can be viewed as a more comprehensive global sentence representation, since different types of syntax and semantic information are encoded in different encoder layers (Shi et al., 2016; Peters et al., 2018; Raganato and Tiedemann, 2018).

We validate our approaches on top of the state-of-the-art TRANSFORMER model (Vaswani et al., 2017). Experimental results on the benchmarks WMT14 English \Rightarrow German and English \Rightarrow French translation tasks show that exploiting sentential context consistently improves translation performance across language pairs. Among the model variations, the deep strategies consistently outperform their shallow counterparts, which confirms our claim. Linguistic analyses (Conneau et al., 2018) on the learned representations reveal that the proposed approach indeed provides richer linguistic information.

The contributions of this paper are:

- Our study demonstrates the necessity and effectiveness of exploiting source-side sentential context for NMT, which benefits from fusing useful contextual information across encoder layers.
- We propose several strategies to better capture useful sentential context for neural machine translation. Experimental results empirically show that the proposed approaches achieve improvement over the strong baseline model TRANSFORMER.




Figure 1: Illustration of the proposed approach. As on a 3-layer encoder: (a) vanilla model without sentential context, (b) *shallow* sentential context representation (i.e. **blue square**) by exploiting the top encoder layer only; and (c) *deep* sentential context representation (i.e. **brown square**) by exploiting all encoder layers. The circles denote hidden states of individual tokens in the input sentence, and the squares denote the sentential context representations. The **red** up arrows denote that the representations are fed to the subsequent decoder. This figure is best viewed in color.

2 Approach

Like a human translator, the encoding process is analogous to reading a sentence in the source language and summarizing its meaning (i.e. sentential context) for generating the equivalents in the target language. When humans translate a source sentence, they generally scan the sentence to create a whole understanding, with which in mind they incrementally generate the target sentence by selecting parts of the source sentence to translate at each decoding step. In current NMT models, the attention model plays the role of selecting parts of the source sentence, but lacking a mechanism to guarantee that the decoder is aware of the whole meaning of the sentence. In response to this problem, we propose to augment NMT models with sentential context, which represents the whole meaning of the source sentence.

2.1 Framework

Figure 1 illustrates the framework of the proposed approach. Let $\mathbf{g} = g(\mathbf{X})$ be the sentential context vector, and $g(\cdot)$ denotes the function to summarize the source sentence \mathbf{X} , which we will discuss in the next sections. There are many possible ways to integrate the sentential context into the decoder. The target of this paper is not to explore this whole space but simply to show that one fairly straightforward implementation works well and that sentential context helps. In this work, we incorporate the sentential context into decoder as

$$\mathbf{d}_i^l = f(\text{LAYER}_{dec}(\widehat{\mathbf{D}}^{l-1}), \mathbf{c}_i^l), \quad (1)$$

$$\widehat{\mathbf{D}}^{l-1} = \text{FFN}_l(\mathbf{D}^{l-1}, \mathbf{g}), \quad (2)$$

where \mathbf{d}_i^l is the l -th layer decoder state at decoding step i , \mathbf{c}_i^l is a dynamic vector that selects certain parts of the encoder output, $\text{FFN}_l(\cdot)$ is a distinct feed-forward network associated with the l -th layer of the decoder, which reads the $l - 1$ -th layer output \mathbf{D}^{l-1} and the sentential context \mathbf{g} . In this way, at each decoding step i , the decoder is aware of the sentential context \mathbf{g} embedded in $\widehat{\mathbf{D}}^{l-1}$.

In the following sections, we discuss the choice of $g(\cdot)$, namely *shallow sentential context* (Figure 1b) and *deep sentential context* (Figure 1c), which differ at the encoder layers to be exploited. It should be pointed out that the new parameters introduced in the proposed approach are jointly updated with NMT model parameters in an end-to-end manner.

2.2 Shallow Sentential Context

Shallow sentential context is a function of the top encoder layer output \mathbf{H}^L :

$$\mathbf{g} = g(\mathbf{H}^L) = \text{GLOBAL}(\mathbf{H}^L), \quad (3)$$

where $\text{GLOBAL}(\cdot)$ is the composition function.

Choices of $\text{GLOBAL}(\cdot)$ Two intuitive choices are mean pooling (Iyyer et al., 2015) and max pooling (Kalchbrenner et al., 2014):

$$\text{GLOBAL}_{\text{MEAN}} = \text{MEAN}(\mathbf{H}^L), \quad (4)$$

$$\text{GLOBAL}_{\text{MAX}} = \text{MAX}(\mathbf{H}^L). \quad (5)$$

Recently, Lin et al. (2017) proposed a self-attention mechanism to form sentence representation, which is appealing for its flexibility on extracting implicit global features. Inspired by this,

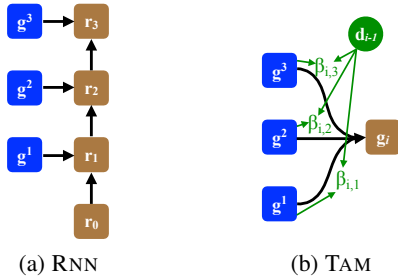


Figure 2: Illustration of the deep functions. “TAM” model dynamically aggregates sentence representations at each decoding step with state \mathbf{d}_{i-1} .

we propose an attentive mechanism to learn sentence representation:

$$\text{GLOBAL}_{\text{ATT}} = \text{ATT}(\mathbf{g}^0, \mathbf{H}^L), \quad (6)$$

$$\mathbf{g}^0 = \text{MAX}(\mathbf{H}^0), \quad (7)$$

where \mathbf{H}^0 is the word embedding layer, and its max pooling vector \mathbf{g}^0 serves as the query to extract features to form the final sentential context representation.

2.3 Deep Sentential Context

Deep sentential context is a function of all encoder layers outputs $\{\mathbf{H}^1, \dots, \mathbf{H}^L\}$:

$$\mathbf{g} = g(\mathbf{H}^1, \dots, \mathbf{H}^L) = \text{DEEP}(\mathbf{g}^1, \dots, \mathbf{g}^L), \quad (8)$$

where \mathbf{g}^l is the sentence representation of the l -th layer \mathbf{H}^l , which is calculated by Equation 3. The motivation for this mechanism is that recent studies reveal that different encoder layers capture linguistic properties of the input sentence at different levels (Peters et al., 2018), and aggregating layers to better fuse semantic information has proven to be of profound value (Shen et al., 2018; Dou et al., 2018; Wang et al., 2018; Dou et al., 2019). In this work, we propose to fuse the global information across layers.

Choices of DEEP(\cdot) In this work, we investigate two representative functions to aggregate information across layers, which differ at whether the decoding information is taken into account.

RNN Intuitively, we can treat $\mathbf{G} = \{\mathbf{g}^1, \dots, \mathbf{g}^L\}$ as a sequence of representations, and recurring all the representations with an RNN:

$$\text{DEEP}_{\text{RNN}} = \text{RNN}(\mathbf{G}). \quad (9)$$

We use the last RNN state as the sentence representation: $\mathbf{g} = \mathbf{r}_L$. As seen, the RNN-based

aggregation repeatedly revises the sentence representations of the sequence with each recurrent step. As a side effect coming together with the proposed approach, the added recurrent inductive bias of RNNs has proven beneficial for many sequence-to-sequence learning tasks such as machine translation (Dehghani et al., 2018).

TAM Recently, Bapna et al. (2018) proposed a novel *transparent attention model* (TAM) to train very deep NMT models. In this work, we apply TAM to aggregate sentence representations:

$$\text{DEEP}_{\text{TAM}} = \sum_{l=1}^L \beta_{i,l} \mathbf{g}^l, \quad (10)$$

$$\beta_i = \text{ATT}_g(\mathbf{d}_{i-1}^l, \mathbf{G}), \quad (11)$$

where $\text{ATT}_g(\cdot)$ is an attention model with its own parameters, that specifies which context representations is relevant for each decoding output. Again, \mathbf{d}_{i-1}^l is the decoder state in the l -th layer.

Comparing with its RNN counterpart, the TAM mechanism has three appealing strengths. First, TAM dynamically generates the weights β_i based on the decoding information at every decoding step \mathbf{d}_{i-1}^l , while RNN is unaware of the decoder states and the associated parameters are fixed after training. Second, TAM allows the model to adjust the gradient flow to different layers in the encoder depending on its training phase.

3 Experiment

We conducted experiments on WMT14 En \Rightarrow De and En \Rightarrow Fr benchmarks, which contain 4.5M and 35.5M sentence pairs respectively. We reported experimental results with case-sensitive 4-gram BLEU score. We used byte-pair encoding (BPE) (Sennrich et al., 2016) with 32K merge operations to alleviate the out-of-vocabulary problem. We implemented the proposed approaches on top of TRANSFORMER model (Vaswani et al., 2017). We followed Vaswani et al. (2017) to set the model configurations, and reproduced their reported results. We tested both *Base* and *Big* models, which differ at the layer size (512 vs. 1024) and the number of attention heads (8 vs. 16).

3.1 Ablation Study

We first investigated the effect of components in the proposed approaches, as listed in Table 1.

Shallow Sentential Context (Rows 3-5) All the shallow strategies achieve improvement over the

#	Model	GLOBAL(\cdot)	DEEP(\cdot)	# Para.	Train	Decode	BLEU
1	BASE	n/a	n/a	88.0M	1.39	3.85	27.31
2	MEDIUM	n/a	n/a	+25.2M	1.08	3.09	27.81
3	SHALLOW	Mean Pooling	n/a	+18.9M	1.35	3.45	27.58
4		Max Pooling		+18.9M	1.34	3.43	27.81 \uparrow
5		Attention		+19.9M	1.22	3.23	28.04 \uparrow
6	DEEP	Attention	RNN	+26.8M	1.03	3.14	28.38 \uparrow
7			TAM	+26.4M	1.07	3.03	28.33 \uparrow

Table 1: Impact of components on WMT14 En \Rightarrow De translation task. BLEU scores in the table are case sensitive. “Train” denotes the training speed (steps/second), and “Decode” denotes the decoding speed (sentences/second) on a Tesla P40. “TAM” denotes the transparent attention model to implement the function DEEP(\cdot). “ \uparrow / \uparrow ”: significant over TRANSFORMER counterpart ($p < 0.05/0.01$), tested by bootstrap resampling (Koehn, 2004).

baseline *Base* model, validating the importance of sentential context in NMT. Among them, attentive mechanism (Row 5) obtains the best performance in terms of BLEU score, while maintains the training and decoding speeds. Therefore, we used the attentive mechanism to implement the function GLOBAL(\cdot) as the default setting in the following experiments.

Deep Sentential Context (Rows 6-7) As seen, both RNN and TAM consistently outperform their shallow counterparts, proving the effectiveness of deep sentential context. Introducing deep context significantly improves translation performance by over 1.0 BLEU point, while only marginally decreases the training and decoding speeds.

Compared to Strong Base Model (Row 2) As our model has more parameters than the *Base* model, we build a new baseline model (MEDIUM in Table 1) which has a similar model size as the proposed deep sentential context model. We change the filter size from 1024 to 3072 in the decoder’s feed-forward network (Eq.2). As seen, the proposed deep sentential context models also outperform the MEDIUM model over 0.5 BLEU point.

3.2 Main Result

Experimental results on both WMT14 En \Rightarrow De and En \Rightarrow Fr translation tasks are listed in Table 2. As seen, exploiting deep sentential context representation consistently improves translation performance across language pairs and model architectures, demonstrating the necessity and effectiveness of modeling sentential context for NMT. Among them, TRANSFORMER-BASE with deep sentential context achieves comparable performance with the vanilla TRANSFORMER-BIG, with only less than half of the parameters (114.4M

Model	En \Rightarrow De	En \Rightarrow Fr
TRANSFORMER-BASE	27.31	39.32
+ DEEP (RNN)	28.38\uparrow	40.15 \uparrow
+ DEEP (TAM)	28.33 \uparrow	40.27\uparrow
TRANSFORMER-BIG	28.58	41.41
+ DEEP (RNN)	29.04 \uparrow	41.87
+ DEEP (TAM)	29.19\uparrow	42.04\uparrow

Table 2: Case-sensitive BLEU scores on WMT14 En \Rightarrow De and En \Rightarrow Fr test sets. “ \uparrow / \uparrow ”: significant over TRANSFORMER counterpart ($p < 0.05/0.01$), tested by bootstrap resampling.

vs. 264.1M, not shown in the table). Furthermore, DEEP (TAM) consistently outperforms DEEP (RNN) in the TRANSFORMER-BIG configuration. One possible reason is that the big models benefit more from the improved gradient flow with the transparent attention (Bapna et al., 2018).

3.3 Linguistic Analysis

To gain linguistic insights into the global and deep sentence representation, we conducted probing tasks¹ (Conneau et al., 2018) to evaluate linguistics knowledge embedded in the encoder output and the sentence representation in the variations of the *Base* model that are trained on En \Rightarrow De translation task. The probing tasks are classification problems that focus on simple linguistic properties of sentences. The 10 probing tasks are categories into three groups: (1) Surface information. (2) Syntactic information. (3) Semantic information. For each task, we trained the classifier on the train set, and validated the classifier on the validation set. We followed Hao et al. (2019) and Li

¹<https://github.com/facebookresearch/SentEval/tree/master/data/probing>

Model	Surface			Syntactic				Semantic					
	SeLen	WC	Avg	TrDep	ToCo	BShif	Avg	Tense	SubN	ObjN	SoMo	CoIn	Avg
L4 IN BASE	94.18	66.24	80.21	43.91	77.36	69.25	63.51	88.03	83.77	83.68	52.22	60.57	73.65
L5 IN BASE	93.40	63.95	78.68	44.36	78.26	71.36	64.66	88.84	84.05	84.56	52.58	61.56	74.32
L6 IN BASE	92.20	63.00	77.60	44.74	79.02	71.24	65.00	89.24	84.69	84.53	52.13	62.47	74.61
+ SSR	92.09	62.54	77.32	44.94	78.39	71.31	64.88	89.17	85.79	85.21	53.14	63.32	75.33
+ DSR	91.86	65.61	78.74	45.52	78.77	71.62	65.30	89.08	85.89	84.91	53.40	63.33	75.32

Table 3: Performance on the linguistic probing tasks of evaluating linguistics embedded in the encoder outputs. “BASE” denotes the representations from TRANSFORMER-BASED encoder. “SSR” denotes shallow sentence representation. “DSR” denotes deep sentence representation. “AVG” denotes the average accuracy of each category.

et al. (2019) to set the model configurations. We also listed the results of lower layer representations ($L = 4, 5$) in TRANSFORMER-BASE to conduct better comparison.

The accuracy results on the different test sets are shown in Table 3. From the table, we can see that

- For different encoder layers in the baseline model (see “L4 in BASE”, “L5 in BASE” and “L6 in BASE”), lower layers embed more about surface information while higher layers encode more semantics, which are consistent with previous findings in (Raganato and Tiedemann, 2018).
- Integrating the shallow sentence representation (“+ SSR”) obtains improvement over the baseline on semantic tasks (75.33 vs. 74.61), while fails to improve on the surface (77.32 vs. 77.60) and syntactic tasks (64.88 vs. 65.00). This may indicate that the shallow representations that exploits only the top encoder layer (“L6 in BASE”) encodes more semantic information.
- Introducing deep sentence representation (“+ DSR”) brings more improvements. The reason is that our deep sentence representation is induced from the sentence representations of all the encoder layers, and lower layers that contain about surface and syntactic information are exploited.

Along with the above translation experiments, we believe that the sentential context is necessary for NMT by enriching the source sentence representation. The *deep sentential context* which is induced from all encoder layers can improve translation performance by offering different types of syntax and semantic information.

4 Related Work

Sentential context has been successfully applied in SMT (Meng et al., 2015; Zhang et al., 2015). In these works, sentential context representation which is generated by the CNNs is exploited to guided the target sentence generation. In broad terms, sentential context can be viewed as a sentence abstraction from a specific aspect. From this point of view, domain information (Foster and Kuhn, 2007; Hasler et al., 2014; Wang et al., 2017b) and topic information (Xiao et al., 2012; Xiong et al., 2015; Zhang et al., 2016) can also be treated as the sentential context, the exploitation of which we leave for future work.

In the context of NMT, several researchers leverage document-level context for NMT (Wang et al., 2017a; Choi et al., 2017; Tu et al., 2018), while we opt for sentential context. In addition, contextual information are used to improve the encoder representations (Yang et al., 2018, 2019; Lin et al., 2018). Our approach is complementary to theirs by better exploiting the encoder representations for the subsequent decoder. Concerning guiding the NMT generation with source-side context, Zheng et al. (2018) split the source content into translated and untranslated parts, while we focus on exploiting global sentence-level context.

5 Conclusion

In this work, we propose to exploit sentential context for neural machine translation. Specifically, the *shallow* and the *deep* strategies exploit the top encoder layer and all the encoder layers, respectively. Experimental results on WMT14 benchmarks show that exploiting sentential context improves performances over the state-of-the-art TRANSFORMER model. Linguistic analyses reveal that the proposed approach indeed captures more linguistic information as expected.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *ICLR*.
- Ankur Bapna, Mia Chen, Orhan Firat, Yuan Cao, and Yonghui Wu. 2018. Training deeper neural machine translation models with transparent attention. In *EMNLP*.
- Heeyoul Choi, Kyunghyun Cho, and Yoshua Bengio. 2017. Context-dependent word representation for neural machine translation. *Computer Speech & Language*, 45:149–160.
- Alexis Conneau, Germán Kruszewski, Guillaume Lample, Łoïc Barrault, and Marco Baroni. 2018. What you can cram into a single \mathbb{R}^d vector: Probing sentence embeddings for linguistic properties. In *ACL*.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. 2018. Universal transformers. *arXiv preprint arXiv:1807.03819*.
- Zi-Yi Dou, Zhaopeng Tu, Xing Wang, Shuming Shi, and Tong Zhang. 2018. Exploiting deep representations for neural machine translation. In *EMNLP*.
- Zi-Yi Dou, Zhaopeng Tu, Xing Wang, Longyue Wang, Shuming Shi, and Tong Zhang. 2019. Dynamic layer aggregation for neural machine translation with routing-by-agreement. In *AAAI*.
- George Foster and Roland Kuhn. 2007. Mixture-model adaptation for smt. In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 128–135. Association for Computational Linguistics.
- Jie Hao, Xing Wang, Baosong Yang, Longyue Wang, Jinfeng Zhang, and Zhaopeng Tu. 2019. Modeling recurrence for transformer. In *NAACL*.
- Eva Hasler, Barry Haddow, and Philipp Koehn. 2014. Combining domain and topic adaptation for smt. In *Proceedings of AMTA*, volume 1, pages 139–151.
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. Deep unordered composition rivals syntactic methods for text classification. In *ACL*.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *ACL*.
- Philipp Koehn. 2004. Statistical Significance Tests for Machine Translation Evaluation. 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*.
- Junyang Lin, Xu Sun, Xuancheng Ren, Shuming Ma, Jinsong Su, and Qi Su. 2018. Deconvolution-based global decoding for neural machine translation. In *COLING*.
- 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*.
- Fandong Meng, Zhengdong Lu, Mingxuan Wang, Hang Li, Wenbin Jiang, and Qun Liu. 2015. Encoding source language with convolutional neural network for machine translation. In *ACL*.
- Eugene A Nida. 1969. Science of translation. *Language*, pages 483–498.
- 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 2018 workshop BlackboxNLP*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In *ACL*.
- Yanyao Shen, Xu Tan, Di He, Tao Qin, and Tie-Yan Liu. 2018. Dense information flow for neural machine translation. In *NAACL*.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. Does string-based neural mt learn source syntax? In *EMNLP*.
- Zhaopeng Tu, Yang Liu, Shuming Shi, and Tong Zhang. 2018. Learning to remember translation history with a continuous cache. *TACL*.
- 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*.
- Longyue Wang, Zhaopeng Tu, Andy Way, and Qun Liu. 2017a. Exploiting cross-sentence context for neural machine translation. In *EMNLP*.
- Qiang Wang, Fuxue Li, Tong Xiao, Yanyang Li, Yinqiao Li, and Jingbo Zhu. 2018. Multi-layer representation fusion for neural machine translation. In *COLING*.
- Rui Wang, Andrew Finch, Masao Utiyama, and Ei-ichiro Sumita. 2017b. Sentence embedding for neural machine translation domain adaptation. In *ACL*.
- Xinyan Xiao, Deyi Xiong, Min Zhang, Qun Liu, and Shouxun Lin. 2012. A topic similarity model for hierarchical phrase-based translation.

- Deyi Xiong, Min Zhang, and Xing Wang. 2015. Topic-based coherence modeling for statistical machine translation. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 23(3):483–493.
- 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*.
- Jiajun Zhang, Dakun Zhang, and Jie Hao. 2015. Local translation prediction with global sentence representation. In *IJCAI*.
- Jian Zhang, Liangyou Li, Andy Way, and Qun Liu. 2016. Topic-informed neural machine translation. In *COLING*.
- Zaixiang Zheng, Hao Zhou, Shujian Huang, Lili Mou, Xinyu Dai, Jiajun Chen, and Zhaopeng Tu. 2018. Modeling past and future for neural machine translation. *TACL*.