

Memory Attentive Fusion: External Language Model Integration for Transformer-based Sequence-to-Sequence Model

Mana Ihuri, Ryo Masumura, Naoki Makishima,
Tomohiro Tanaka, Akihiko Takashima, Shota Orihashi
NTT Media Intelligence Laboratories, NTT Corporation
1-1 Hikarinooka, Yokosuka-Shi, Kanagawa 239-0847, Japan
mana.ihori.kx@hco.ntt.co.jp

Abstract

This paper presents a novel fusion method for integrating an external language model (LM) into the Transformer based sequence-to-sequence (seq2seq) model. While paired data are basically required to train the seq2seq model, the external LM can be trained with only unpaired data. Thus, it is important to leverage memorized knowledge in the external LM for building the seq2seq model, since it is hard to prepare a large amount of paired data. However, the existing fusion methods assume that the LM is integrated with recurrent neural network-based seq2seq models instead of the Transformer. Therefore, this paper proposes a fusion method that can explicitly utilize network structures in the Transformer. The proposed method, called **memory attentive fusion**, leverages the Transformer-style attention mechanism that repeats source-target attention in a multi-hop manner for reading the memorized knowledge in the LM. Our experiments on two text-style conversion tasks demonstrate that the proposed method performs better than conventional fusion methods.

1 Introduction

In recent studies, the Transformer sequence-to-sequence (seq2seq) model (Vaswani et al., 2017) has successfully performed in various natural language generation tasks, such as machine translation (Wang et al., 2019; Barrault et al., 2019), image captioning (Li et al., 2019b; Yu et al., 2019; Li et al., 2019a), and automatic speech recognition (Dong et al., 2018; Karita et al., 2019; Salazar et al., 2019). Although the Transformer training needs paired data, a large amount of paired data cannot often be prepared. Moreover, unpaired data cannot be used for training the Transformer even though such data can be collected on a large scale.

To utilize a large amount of unpaired data, methods of integrating an external language model (LM) trained with these data into seq2seq mod-

els have been proposed (Kannan et al., 2018; Gulcehre et al., 2015; Sriram et al., 2018). These methods can improve the fluency of sentences that are generated by seq2seq models; however, they integrate the LM into recurrent neural network (RNN) based seq2seq models rather than the Transformer. In other words, LM fusion methods specific to the Transformer have not been considered yet.

Here, the Transformer employs the multi-hop attention mechanism (Sukhbaatar et al., 2015) that repeats the source-target attention mechanism in each Transformer decoder block. Thus, it is supposed that the source-target attention mechanism promotes to extract effective source information for target tasks more exactly than RNN based seq2seq models. Therefore, we assume that the Transformer can utilize memorized knowledge in the external LM more effectively by using the multi-hop attention mechanism for the LM fusion.

In this paper, we propose a novel fusion method, called **memory attentive fusion**, to integrate an external LM into the Transformer. This fusion method utilizes a multi-hop source-target attention mechanism for combining the Transformer decoder with the external LM. We performed experiments with two text-style conversion tasks: spoken-to-written style conversion and dialect conversion. Our experiments demonstrate that the proposed method performs better than conventional fusion methods.

2 Related work

The simplest fusion method is to train the seq2seq model and the LM independently and then combine their outputs (Kannan et al., 2018; Chorowski and Jaitly, 2017; Sutskever et al., 2014). These methods are called shallow fusion. Moreover, methods that integrate an external LM into seq2seq models during training have been proposed: deep fusion (Gulcehre et al., 2015) and

cold fusion (Sriram et al., 2018). These methods utilize the information of not only paired data but also unpaired data in training. Figure 1 shows a Transformer with cold fusion. These methods assume that the LM is integrated into RNN-based seq2seq models instead of the Transformer.

3 Memory attentive fusion

This section details memory attentive fusion for the Transformer seq2seq model. In fact, memory attentive fusion is an extended method of the cold fusion (Sriram et al., 2018). While the cold fusion uses memorized knowledge in the LM at an output layer only once, the memory attentive fusion repeatedly uses the knowledge at Transformer decoder blocks based on a source-target attention mechanism.

We define an input sequence as $\mathbf{X} = \{x_1, \dots, x_M\}$ and an output sequence as $\mathbf{Y} = \{y_1, \dots, y_N\}$, where x_m and y_n are tokens in the input and output sequence. In text-style conversion, the model predicts the generation probabilities of the output sequence given the input sequence. The generation probability of \mathbf{Y} is defined as

$$P(\mathbf{Y}|\mathbf{X}; \Theta) = \prod_{n=1}^N P(y_n|y_{1:n-1}, \mathbf{X}; \Theta), \quad (1)$$

where $\Theta = \{\theta_{\text{enc}}, \theta_{\text{dec}}, \theta_{\text{lm}}\}$ represents model parameter sets. θ_{enc} and θ_{dec} are trainable parameter sets with encoder and decoder, respectively. θ_{lm} is parameter set for the external LM. $P(y_n|y_{1:n-1}, \mathbf{X}; \Theta)$ can be computed using an encoder and a decoder with memory attentive fusion in the Transformer. Figure 2 shows the Transformer with memory attentive fusion.

Encoder: The encoder converts an input sequence \mathbf{X} into the hidden representations $\mathbf{S}^{(K)}$ using K Transformer encoder blocks. First, the input hidden representation of the Transformer encoder block $\mathbf{S}^{(0)} = \{s_{1:M}^{(0)}\}$ is produced by

$$s_m^{(0)} = \text{Embedding}(x_m; \theta_{\text{enc}}), \quad (2)$$

where $\text{Embedding}(\cdot)$ consists of positional encoding and a linear layer. Next, the k -th Transformer encoder block composes the k -th hidden representations $\mathbf{S}^{(k)}$ from the lower inputs $\mathbf{S}^{(k-1)}$ as

$$\mathbf{S}^{(k)} = \text{TransformerEncBlock}(\mathbf{S}^{(k-1)}; \theta_{\text{enc}}), \quad (3)$$

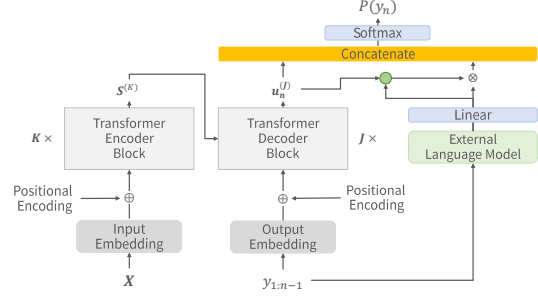


Figure 1: Transformer with cold fusion.

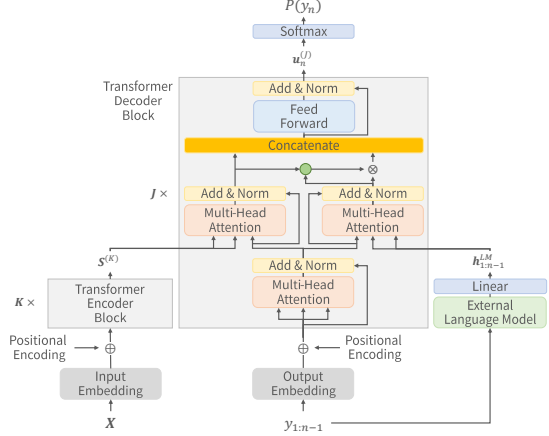


Figure 2: Transformer with memory attentive fusion.

where $\text{TransformerEncBlock}(\cdot)$ is the Transformer encoder block that consists of a scaled dot product multi-head self-attention layer and a position-wise feed-forward network (Vaswani et al., 2017).

Decoder with memory attentive fusion: The decoder with memory attentive fusion computes the generation probability of a token from the preceding tokens and hidden representations of the input sequence and the LM information. The predicted probabilities of the n -th token y_n are calculated as

$$P(y_n|y_{1:n-1}, \mathbf{X}) = \text{softmax}(\mathbf{u}_n^{(J)}; \theta_{\text{dec}}), \quad (4)$$

where $\text{softmax}(\cdot)$ is a softmax layer with a linear transformation. The input hidden vector $\mathbf{u}_n^{(J)}$ is computed from $\mathbf{S}^{(K)}$ and $y_{1:n-1}$ using J Transformer decoder blocks with an external LM. First, the input hidden representation of the Transformer decoder block $\mathbf{u}_{n-1}^{(0)}$ and $\mathbf{h}_{n-1}^{\text{LM}}$ are produced by

$$\mathbf{u}_{n-1}^{(0)} = \text{Embedding}(y_{n-1}; \theta_{\text{dec}}), \quad (5)$$

$$\mathbf{l}_{n-1}^{\text{LM}} = \text{LanguageModel}(y_{1:n-1}; \theta_{\text{lm}}), \quad (6)$$

$$\mathbf{h}_{n-1}^{\text{LM}} = \text{linear}(\mathbf{l}_{n-1}^{\text{LM}}; \theta_{\text{dec}}), \quad (7)$$

where $\text{LanguageModel}(\cdot)$ is the trained external LM, and l_{n-1}^{LM} is the logit output. Next, we convert hidden representations in the lower layer $\mathbf{u}_{1:n-1}^{(j-1)}$ and the encoder output $\mathbf{S}^{(k)}$ into a hidden vector $\mathbf{c}_n^{(j)}$. The hidden vector is computed as

$$\bar{\mathbf{v}}_n^{(j)} = \text{SourceTarget}(\mathbf{u}_{1:n-1}^{(j-1)}, \mathbf{u}_{n-1}^{(j-1)}; \theta_{\text{dec}}), \quad (8)$$

$$\mathbf{v}_n^{(j)} = \text{LayerNorm}(\mathbf{u}_{n-1}^{(j-1)} + \bar{\mathbf{v}}_n^{(j)}), \quad (9)$$

$$\bar{\mathbf{c}}_n^{(j)} = \text{SourceTarget}(\mathbf{S}^{(K)}, \mathbf{v}_n^{(j)}; \theta_{\text{dec}}), \quad (10)$$

$$\mathbf{c}_n^{(j)} = \text{LayerNorm}(\mathbf{v}_n^{(j)} + \bar{\mathbf{c}}_n^{(j)}), \quad (11)$$

where $\text{SourceTarget}(\cdot)$ is a scaled dot product multi-head source target attention layer and $\text{LayerNorm}(\cdot)$ is layer normalization (Ba et al., 2016). In memory attentive fusion, we also convert the LM output $\mathbf{h}_{1:n-1}^{\text{LM}}$ and the hidden vector $\mathbf{v}_n^{(j)}$ into a hidden vector $\mathbf{b}_n^{(j)}$ with attention mechanism. The hidden vector is computed as

$$\bar{\mathbf{b}}_n^{(j)} = \text{SourceTarget}(\mathbf{h}_{1:n-1}^{\text{LM}}, \mathbf{v}_n^{(j)}; \theta_{\text{dec}}), \quad (12)$$

$$\mathbf{b}_n^{(j)} = \text{LayerNorm}(\mathbf{v}_n^{(j)} + \bar{\mathbf{b}}_n^{(j)}). \quad (13)$$

This attention mechanism is repeated with Transformer decoder block in the multi-hop manner. Therefore, we expect to read the memorized memory in the LM effectively. Next, we concatenate the hidden vector that have target and source information, and that have target and the LM information with gating mechanism by

$$\mathbf{g}_n^{(j)} = \text{sigmoid}([\mathbf{c}_n^{(j)\text{T}}, \mathbf{b}_n^{(j)\text{T}}]^\text{T}; \theta_{\text{dec}}), \quad (14)$$

$$\mathbf{q}_n^{(j)} = [\mathbf{c}_n^{(j)\text{T}}, \mathbf{g}_n^{(j)} \odot \mathbf{b}_n^{(j)\text{T}}]^\text{T}, \quad (15)$$

where $\text{sigmoid}(\cdot)$ is a sigmoid layer with a linear transformation. Next, the hidden vector $\mathbf{q}_n^{(j)}$ is converted into the j -th hidden representation $\mathbf{u}_n^{(j)}$. The hidden representation is computed as

$$\bar{\mathbf{u}}_n^{(j)} = \text{FeedForward}(\mathbf{q}_n^{(j)}; \theta_{\text{dec}}), \quad (16)$$

$$\mathbf{u}_n^{(j)} = \text{LayerNorm}(\mathbf{q}_n^{(j)} + \bar{\mathbf{u}}_n^{(j)}), \quad (17)$$

where $\text{FeedForwrd}(\cdot)$ is a position-wise feed-forward network.

Training: In the Transformer, the model parameter set can be optimized from training dataset $\mathcal{D} = \{(\mathbf{X}^1, \mathbf{Y}^1), \dots, (\mathbf{X}^{|\mathcal{D}|}, \mathbf{Y}^{|\mathcal{D}|})\}$. The objective function for optimizing the model parameter

is defined as

$$\mathcal{L} = -\frac{1}{|\mathcal{D}|} \sum_{d=1}^{|\mathcal{D}|} \log P(\mathbf{Y}^d | \mathbf{X}^d; \Theta). \quad (18)$$

Note that the external LM uses the freezing parameter θ_{lm} .

4 Experiments

We evaluated our method on text-style conversion tasks. In particular, we chose *spoken-to-written style conversion task* and *dialect conversion task* in Japanese. In the spoken-to-written style conversion task, spoken-style text produced by an automatic speech recognition system is converted into written-style text that has correct punctuation and no disfluency (Ihori et al., 2020). In the dialect conversion task, Japanese dialects are converted into standard Japanese.

4.1 Datasets

Spoken-to-written style conversion: We used the Corpus of Spontaneous Japanese (CSJ) (Maekawa et al., 2000) and the parallel corpus for Japanese spoken-to-written style conversion (CJSW) (Ihori et al., 2020). We divided the CSJ into a training set, validation set, and test set. The training set, validation set, and test set have 46,847, 13,510, and 3,949 sentences, respectively. The CJSW has four domains, and we divided it up following (Ihori et al., 2020). We used all of the training and validation sets for training and each test set (CJSW-1, 2, 3, 4) for the evaluation. All of these datasets are paired data of spoken-style text (manual transcriptions of speech) and written-style text (created with crowd-sourcing).

Dialect conversion: We prepared three paired data of dialects (Tohoku-ben, Osaka-ben, Kyushu-ben) and standard Japanese with crowd-sourcing. We divided these data into a training set, validation set, and test set for each dialect. We used all of the training and validation sets for training and three test sets for the evaluation. The training set, validation set and test set have 15,506, 3,924 and 2,160 sentences, respectively. Moreover, the test set consists of 700 Tohoku-ben, 862 Osaka-ben, and 598 Kyushu-ben sentences, respectively.

External text: We prepared a large-scale Japanese web text as the unpaired written-style text data. The Web text was downloaded from

various topic Web pages using our home-made crawler. The downloaded pages were filtered in such a way that HTML tags, Javascript codes and other parts that were not useful for these tasks were excluded. Finally, we prepare *one million sentences* for training the external LM. Moreover, we divided this data into a training set, validation set. The training set and validation set have 800,000 and 200,000 sentences, respectively.

4.2 Setups

Transformer: We constructed the Transformer with shallow fusion (Kannan et al., 2018), cold fusion (Sriram et al., 2018) and memory attentive fusion methods. In addition, we constructed the Transformer without fusion methods as a baseline. We used the following configurations. The dimensions of the output continuous representations and the inner outputs in the position-wise feed-forward network were set to 256, and the number of heads in the multi-head attentions was set to 8. ReLU activation was used in nonlinear transformation function. We stacked 4-layer Transformer encoder blocks, and 2-layer Transformer decoder blocks. We set the output unit size (which corresponded to the amount of tokens in the training set for language model) to 5,640. To train these models, we used the adam optimizer and label smoothing with a smoothing parameter of 0.1. The training steps were stopped based on early stopping using the part of the training data. We set the mini-batch size to 64 sentences and the dropout rate in the Transformer blocks to 0.2. For the mini-batch training, we truncated each sentence to 200 tokens. We used characters as tokens. All trainable parameters were randomly initialized. For the decoding, we used a beam search algorithm in which the beam size was set to 4.

External LM: We utilized Open AI GPT (Radford et al., 2019) for the LM fusion, although any LM can potentially be used. We used the following configurations. The number of heads in the multi-head attentions was set to 4. We stacked 8-layer Transformer blocks. The training steps were stopped based on early stopping using the part of the training data. We set the dropout rate in the Transformer blocks to 0.1. The other settings were the same as the Transformer settings. After training, perplexity of this LM was 11.8. Note that this LM was used in both two tasks and the Transformer and the external LM were not pre-trained.

4.3 Results

Table 1 shows the experimental results in the spoken-to-written style conversion task. Also, Table 2 shows the experimental results in the dialect conversion task. We calculated automatic evaluation scores in three metrics: BLEU-3 (B-3) (Papineni et al., 2002), ROUGE-L (R-L) (Lin and Och, 2004), and METEOR (Banerjee and Lavie, 2005). *Baseline* in the tables mean the results of the Transformer without the external LM.

Table 1 shows that shallow fusion and cold fusion performed worse than the baseline on the CSJ dataset. On the other hand, memory attentive fusion outperformed the baseline. Moreover, memory attentive fusion outperformed the baseline and shallow fusion on the CJSW dataset. In addition, cold fusion outperformed the baseline on CJSW-1, -3 and -4. As in the spoken-to-written style conversion task, Table 2 shows that memory attentive fusion outperformed the other methods.

The above results show that shallow fusion is not suitable for the Transformer because it degraded performance in all cases. Moreover, when the LM was integrated with cold fusion, the performance was better than baseline in some domains. Thus, we consider that cold fusion can be used with the Transformer in limited cases.

Test set		B-3	R-L	METEOR
CSJ	a).	0.667	0.855	0.853
	b).	0.667	0.850	0.853
	c).	0.657	0.852	0.847
	d).	0.669	0.860	0.856
CJSW-1	a).	0.723	0.785	0.881
	b).	0.705	0.775	0.870
	c).	0.734	0.791	0.889
	d).	0.735	0.792	0.887
CJSW-2	a).	0.657	0.718	0.847
	b).	0.630	0.702	0.830
	c).	0.655	0.714	0.840
	d).	0.671	0.726	0.859
CJSW-3	a).	0.671	0.732	0.839
	b).	0.663	0.726	0.836
	c).	0.672	0.73	0.842
	d).	0.686	0.737	0.85
CJSW-4	a).	0.772	0.818	0.898
	b).	0.752	0.806	0.884
	c).	0.775	0.819	0.897
	d).	0.779	0.821	0.900

a). Baseline b). Shallow fusion
c). Cold fusion d). Memory attentive fusion

Table 1: Results on spoken-to-written style conversion tasks.

Test set		B-3	R-L	METEOR
Osaka-ben	a).	0.649	0.784	0.790
	b).	0.638	0.774	0.780
	c).	0.648	0.784	0.787
	d).	0.663	0.795	0.802
Kyushu-ben	a).	0.741	0.857	0.872
	b).	0.729	0.849	0.859
	c).	0.738	0.855	0.867
	d).	0.752	0.864	0.880
Tohoku-ben	a).	0.619	0.767	0.742
	b).	0.603	0.755	0.721
	c).	0.610	0.761	0.730
	d).	0.630	0.772	0.752

a). Baseline b). Shallow fusion
c). Cold fusion d). Memory attentive fusion

Table 2: Results on dialect conversion tasks.

Input	新鮮で海のものがおいしいです < Fresh and delicious sea food >
Reference	新鮮で海のものがおいしいです。
Baseline	新朝海で海のものがおいしいです。
Shallow fusion	新海で海のものがおいしいです。
Cold fusion	新朝で海のものがおいしいです。
Memory attentive fusion	新鮮で海のものがおいしいです。

Figure 3: Example of spoken-to-written style conversion in CSJ dataset with each fusion method.

On the other hand, memory attentive fusion outperformed the other fusion methods in almost all of the domains and tasks. Therefore, we consider that memory attentive fusion is suitable for integration of the external LM into the Transformer. In addition, memory attentive fusion worked well especially in the dialect conversion task. Thus, we can infer that the fusion method for the Transformer is more effective when there is small training data.

We show the converted example of spoken-to-written style conversion in CSJ dataset with each fusion method in Figure 3. Figure 3 shows that the word “新鮮” (fresh) was output correctly with memory attentive fusion, but other methods were not output the word correctly. The word “新鮮” was not included in training data for the Transformer, but it was included in training data for the external LM. This indicates that only memory attentive fusion was successful in extracting knowledge of the external LM.

5 Conclusion

We proposed memory attentive fusion, a novel method to integrate an external LM into the Transformer. Conventional fusion methods assume that the LM is integrated into the RNN-based seq2seq. On the other hand, the proposed method

employs a Transformer-specific fusion method which repeats the attention mechanism for the LM many times. Experiments demonstrated that the proposed method outperformed the conventional methods in two tasks. We conclude that the proposed method is suitable for integrating the LM into the Transformer. In the future work, we will try using the proposed method in other natural language generation tasks such as automatic speech recognition.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proc. the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. Findings of the 2019 conference on machine translation (WMT19). In *Proc. Conference on Machine Translation (WMT)*, pages 1–61.
- Jan Chorowski and Navdeep Jaitly. 2017. Towards better decoding and language model integration in sequence to sequence models. In *Proc. International Speech Communication Association (INTER-SPEECH)*, pages 523–527.
- Linhao Dong, Shuang Xu, and Bo Xu. 2018. Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition. In *Proc. International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5884–5888.
- Caglar Gulcehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, Loic Barrault, Huei-Chi Lin, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2015. On using monolingual corpora in neural machine translation. *arXiv preprint arXiv:1503.03535*.
- Mana Ihuri, Akihiko Takashima, and Ryo Masumura. 2020. Parallel corpus for Japanese spoken-to-written style conversion. In *Proc. Language Resources and Evaluation Conference (LREC)*, pages 6346–6353.
- Anjali Kannan, Yonghui Wu, Patrick Nguyen, Tara N Sainath, Zhijeng Chen, and Rohit Prabhavalkar. 2018. An analysis of incorporating an external language model into a sequence-to-sequence

- model. In *Proc. International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5824–5828.
- Shigeki Karita, Nanxin Chen, Tomoki Hayashi, Takaaki Hori, Hirofumi Inaguma, Ziyang Jiang, Masao Someki, Nelson Enrique Yalta Soplín, Ryuichi Yamamoto, Xiaofei Wang, et al. 2019. A comparative study on transformer vs RNN in speech applications. In *Proc. Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 449–456.
- Guang Li, Linchao Zhu, Ping Liu, and Yi Yang. 2019a. Entangled transformer for image captioning. In *Proc. International Conference on Computer Vision (ICCV)*, pages 8928–8937.
- Jie Li, Xiaorui Wang, Yan Li, et al. 2019b. The speech transformer for large-scale Mandarin Chinese speech recognition. In *Proc. International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7095–7099.
- Chin-Yew Lin and Franz Josef Och. 2004. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proc. Annual Meeting on Association for Computational Linguistics (ACL)*, pages 605–612.
- Kikuo Maekawa, Hanae Koiso, Sadaoki Furui, and Hitoshi Isahara. 2000. Spontaneous speech corpus of Japanese. In *Proc. International Conference on Language Resources and Evaluation (LREC)*, pages 947–952.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proc. Annual Meeting on Association for Computational Linguistics (ACL)*, pages 311–318.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, page 9.
- Julian Salazar, Katrin Kirchhoff, and Zhiheng Huang. 2019. Self-attention networks for connectionist temporal classification in speech recognition. In *Proc. International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7115–7119.
- Anuroop Sriram, Heewoo Jun, Sanjeev Satheesh, and Adam Coates. 2018. Cold fusion: Training seq2seq models together with language models. In *Proc. International Speech Communication Association (INTERSPEECH)*, pages 387–391.
- Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Proc. Advances in neural information processing systems (NIPS)*, pages 2440–2448.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Proc. Advances in neural information processing systems (NIPS)*, pages 3104–3112.
- 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 *Proc. Advances in neural information processing systems (NIPS)*, pages 5998–6008.
- Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. 2019. Learning deep transformer models for machine translation. In *Proc. Association for Computational Linguistics (ACL)*, pages 1810–1822.
- Jun Yu, Jing Li, Zhou Yu, and Qingming Huang. 2019. Multimodal transformer with multi-view visual representation for image captioning. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1.