



Streaming Transformer-based Acoustic Models Using Self-attention with Augmented Memory

Chunyang Wu, Yongqiang Wang, Yangyang Shi, Ching-Feng Yeh, Frank Zhang

Facebook AI

{chunyang, yqw, yyshi, cfyeh, frankz}@fb.com

Abstract

Transformer-based acoustic modeling has achieved great success for both hybrid and sequence-to-sequence speech recognition. However, it requires access to the full sequence, and the computational cost grows quadratically with respect to the input sequence length. These factors limit its adoption for streaming applications. In this work, we proposed a novel augmented memory self-attention, which attends on a short segment of the input sequence and a bank of memories. The memory bank stores the embedding information for all the processed segments. On the librispeech benchmark, our proposed method outperforms all the existing streamable transformer methods by a large margin and achieved over 15% relative error reduction, compared with the widely used LC-BLSTM baseline. Our findings are also confirmed on some large internal datasets.

Index Terms: streaming speech recognition, transformer, acoustic modeling

1. Introduction

Sequence modeling is an important problem in speech recognition. In both conventional hybrid [1] and end-to-end style (e.g., attention-based encoder-decoder [2, 3] or neural transducer[4]) architectures, a neural encoder is used to extract a sequence of high-level embeddings from an input feature vector sequence. A feed-forward neural network extracts embeddings from a fixed window of local features [5]. Recurrent neural networks (RNNs), especially the long short-term memory (LSTM) [6], improve the embedding extraction by exploiting both long-term and short-term temporal patterns [7]. Recently, attention (or self-attention if there is only one input sequence) has emerged as an alternative technique for sequence modeling [8]. Different from RNNs, attention connects arbitrary pairs of positions in the input sequences directly. To forward (or backward) signals between two positions that are n steps away in the input, it only needs one step to traverse the network, compared with $O(n)$ steps in RNNs. Built on top of the attention operation, the transformer model [8] leverages multi-head attention and interleaves with feed-forward layers. It has achieved great success in both natural language processing [9, 10] and speech applications [11, 12].

However, two significant issues make transformer-based models impractical for online speech recognition applications. First, it requires access to the entire utterance before it can start producing output; second, the computational cost and memory usage grow quadratically with respect to the input sequence length if an infinite left context is used. There are a few methods that can partially solve these issues. First, *time-restricted self-attention* [13] can be used in which the computation of attention only uses the past input vectors and a limited length of future inputs (e.g., [14, 15]). However, since the reception field

is linearly growing for the number of transformer layers, it usually generates a significant delay; it does not address the issue of quadratically growing cost either. Second, *block processing* is used in [16], which chunks the input utterances into segments, and self-attention performs on each segment. In this way, the computation cost and memory usage don't grow quadratically. It is similar to context-sensitive-chunk BPTT in [17] and truncated BLSTM in [18], which was successfully deployed to build online speech recognition system based on BLSTM models. However, since the transformer cannot attend beyond the current segment, it is observed that this method yields significant accuracy degradation [19, 20]. Third, *recurrent connection*, in which embeddings from the previous segment are carried over to the current one, can be combined with the block processing. This approach is similar to the idea proposed in latency controlled BLSTM (LC-BLSTM) [21]. An example of this approach is transformer-XL [20], in which it can model a very long dependency on text data for language modeling. The works in [19, 22] have explored similar ideas for acoustic modeling.

Carrying over segment level information enables attention to access information beyond the current segment. A recurrent connection compresses the segment level information into a single memory slot. For a segment that is k steps away, it takes $O(k)$ steps to retrieve the embedding extracted from that segment. Inspired by the neural Turing machine[23], we propose a novel *augmented memory* transformer, which accumulates the segment level information into a memory bank with multiple memory slots. Attention is then performed over the memory bank, together with the embeddings from the current segment. In this way, all the information, regardless of whether it is in the current segment or k segments away, can be equally accessible. We applied this *augmented memory* transformer to hybrid speech recognition architecture and performed an in-depth comparison with other methods on a widely used LibriSpeech benchmark [24]. Experimental results demonstrate that the proposed augmented memory transformer outperforms all the other methods by a large margin. Using our proposed method, we show that with similar look-ahead sizes, augmented memory transformer improves over the widely used LC-BLSTM model by over 15% relatively. Though we only evaluate the proposed method in a hybrid speech recognition scenario, it is equally applicable to end-to-end style architectures.

The rest of this paper is organized as follows. In Section 2, we briefly review the self-attention and transformer-based acoustic model. We present the augmented memory transformer in Section 3. Section 4 demonstrates and analyzes the experimental results, followed by a summary in Section 5.

2. Transformer-based acoustic models

We first give a brief introduction of self-attention that is the core of the transformer-based model. Then we describe the architec-

The authors would like to thank Christian Fuegen and Michael L. Seltzer for their support and discussion.

ture of the transformer-based acoustic model from [12]. The model in this paper extends its model architecture for online streaming speech recognition.

2.1. Self-attention

Given an input embedding sequence $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ where $\mathbf{x}_t \in \mathbb{R}^D$, self-attention projects the input to *query*, *key* and *value* space using \mathbf{W}_q , \mathbf{W}_k and \mathbf{W}_v , respectively,

$$\mathbf{Q} = \mathbf{W}_q \mathbf{X}, \quad \mathbf{K} = \mathbf{W}_k \mathbf{X}, \quad \mathbf{V} = \mathbf{W}_v \mathbf{X} \quad (1)$$

where \mathbf{W}_q , \mathbf{W}_k , \mathbf{W}_v are learnable parameters. Self-attention uses dot-product to get the attention distribution over *query* and *key*, i.e., for position t in *query*, a distribution α_t is obtained by:

$$\alpha_{t\tau} = \frac{\exp(\beta \cdot \mathbf{Q}_t^\top \mathbf{K}_\tau)}{\sum_{\tau'} \exp(\beta \cdot \mathbf{Q}_t^\top \mathbf{K}_{\tau'})} \quad (2)$$

where $\beta = \frac{1}{\sqrt{D}}$ is a scaling factor. Given α_t , the output embedding of self-attention is obtained via:

$$\mathbf{z}_t = \sum_{\tau} \text{Dropout}(\alpha_{t\tau}) \cdot \mathbf{V}_\tau. \quad (3)$$

In [8], multiple head attentions are introduced. Each of the attention heads is applied individually on the input sequences. The output of each head is concatenated and linearly transformed into the final output.

2.2. Transformer-based acoustic model

The transformer-based acoustic model [12] is a deep stack of transformer layers on top of VGG blocks [25]. Each transformer layer consists of a multi-head self-attention followed by a position-wise feed-forward layer. Rather than using Sinusoid positional embedding [8], the transformer-based acoustic model [12] uses VGG blocks to implicitly encode the relative positional information [26]. The layer normalization [27], the iterated loss [28], residual connections, and dropout are applied to train the deep stack transformer layers effectively. More model details can be found from [12].

3. Augmented Memory Transformer

The original transformer model generates the outputs according to the attention on the whole input sequence, which is not suitable for streaming speech recognition. The proposed augmented memory transformer addresses this issue by the combination of two mechanisms. First, similar to *block processing* [16], the whole utterance is segmented into segments padding with left context and right context. The size of each segment limits the computation and memory consumption in each transformer layer. Second, to carry over information across segments, an *augmented memory bank* is used. Each slot in the *augmented memory bank* is the embedding representation of an observed segment.

Figure 1 illustrates one forward step on the n -th segment using augmented memory transformer. An augmented memory bank (red) is introduced to the self-attention function. The input sequence is first segmented into segments. Each segment $\mathbf{C}_n = (\mathbf{x}_{nB+1}, \dots, \mathbf{x}_{(n+1)B})$ contains B input embedding vectors, where B is referred to as the *segment length*. The n -th segment is formed by patching the current segment with left context \mathbf{L}_n (length L) and right context \mathbf{R}_n (length R). An

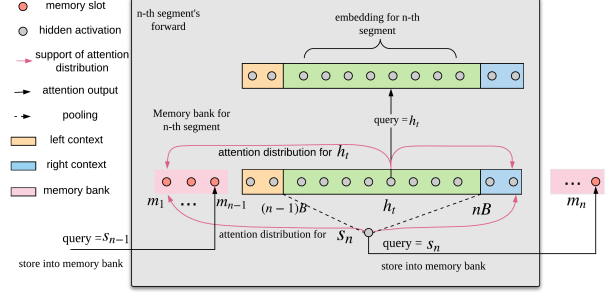


Figure 1: Illustration of one forward step for the augmented memory transformer on the n -th segment.

embedding vector \mathbf{s}_n , referred to as the *summarization embedding* is then computed by pooling over \mathbf{C}_n . Different pooling methods, e.g., average pooling, max pooling, and the linear combination, can be used. This paper focuses on the average pooling. In the self-attention with augmented memory, the *query* is the projection from the concatenation of current segment with context frames and the summarization embedding. The *key* and the *value* are the projections from the concatenation of the augmented memory bank and the current segment with context frames. They are formalized as

$$\mathbf{Q} = \mathbf{W}_q[\mathbf{L}_n, \mathbf{C}_n, \mathbf{R}_n], \quad \mathbf{q}_{\text{sum}} = \mathbf{W}_q \mathbf{s}_n, \quad (4)$$

$$\mathbf{K} = \mathbf{W}_k[\mathbf{M}_n, \mathbf{L}_n, \mathbf{C}_n, \mathbf{R}_n], \quad (5)$$

$$\mathbf{V} = \mathbf{W}_v[\mathbf{M}_n, \mathbf{L}_n, \mathbf{C}_n, \mathbf{R}_n] \quad (6)$$

where $\mathbf{M}_n = (\mathbf{m}_1, \dots, \mathbf{m}_{n-1})$ is the augmented memory bank. The attention output from \mathbf{Q} is feed to the next layer, except for the last transformer layer, only the center B vectors, i.e., $(\mathbf{q}_{L+1}, \dots, \mathbf{q}_{L+B})$, are used. The query \mathbf{q}_{sum} is the projection from the summarization embedding \mathbf{s}_n ; the associated attention output is stored into augmented memory bank as \mathbf{m}_n for future forward steps,

$$\mathbf{m}_n = \sum_{\tau} \text{Dropout}(\alpha_{\text{sum},\tau}) \cdot \mathbf{V}_\tau \quad (7)$$

where $\alpha_{\text{sum},\tau}$ is the attention weight for \mathbf{q}_{sum} . The output for the whole utterance is the concatenation of outputs from all the segments.

The proposed method is different from existing models in a variety of aspects. Transformer-XL [20] incorporates history information only from previous segment \mathbf{C}_{n-1} via

$$\begin{aligned} \mathbf{Q} &= \mathbf{W}_q \mathbf{C}_n, \\ \mathbf{K} &= \mathbf{W}_k[\mathbf{C}_{n-1}, \mathbf{C}_n], \quad \mathbf{V} = \mathbf{W}_v[\mathbf{C}_{n-1}, \mathbf{C}_n]. \end{aligned} \quad (8)$$

Also note that, in transformer-XL, \mathbf{C}_{n-1} is from the lower layer. This makes the upper layers have an increasing long reception field. Our proposed augmented memory transformer explicitly holds the information from all the previous segments (Eq. 5 and 6) and all the layers have the same reception field. Using a bank of memories to represent past segments is also explored in [29], primarily in language modeling tasks. In [13], the *time-restricted transformer* restricts the attention to a context window in each transformer layer. This means the look-ahead length is linearly growing by the number of transformer layers. Our proposed method has a fixed look-ahead window, thus enabling us to use many transformer layers without increasing look-ahead window size.

4. Experiments

The proposed model was evaluated on the LibriSpeech ASR task, and two of our internal video ASR tasks, German and Russian. Neural network models were trained and evaluated using an in-house extension of the PyTorch-based *fairseq* [30] toolkit. In this paper, different streamable models were compared with similar look-ahead windows, for a fair comparison.

4.1. LibriSpeech

We first performed experiments on the LibriSpeech task [24]. This dataset contains about 960 hours of read speech data for training, and 4 development and test sets (`{dev, test} - {clean, other}`) for evaluation, where `other` sets are more acoustically challenging. The standard 4-gram language model (LM) with a 200K vocabulary was used for all first-pass decoding. In all experiments, 80 dimensional log Mel-filter bank features with a 10ms frame-shift were used as input features. The context- and position-dependent graphemes, i.e., *chenones* [31], were used as output labels.

4.1.1. Experiment Setups

A GMM-HMM system was first trained following the standard Kaldi [32] Librispeech recipe. To speed up the training of neural networks, the training data were segmented into utterances that were up to 10 seconds¹; speed perturbation [33] and *SpecAugment* [34] were performed on the training data. In evaluation, no segmentation was performed on the test data. This paper focuses on cross-entropy (CE) training for neural network models. The proposed augmented memory transformer (AMTrf) was compared with streamable baselines including LC-BLSTM [21], transformer-XL (Trf-XL) [20] and time-restricted transformer (TRTrf) [13]. Also, the non-streamable original transformer (Trf) was included to indicate potential performance lower-bound.

We started with investigating models of a small configuration with approximately 40M parameters. The LC-BLSTM baseline consists of 5 layers with 400 nodes in each layer each direction. Mixed frame rate [35], i.e., the output frames of the first layer are sub-sampled by factor of 2 before propagated to the second layer, is used. The look-ahead window is set to 0.4 second, i.e., 40 frames; the chunk size in LC-BLSTM is 1.5 seconds. For transformers, the topology is 12 transformer layers with 512 input embedding dimensions, 8 multi-heads, and 2048 feed-forward network (FFN) dimensions in each layer. Following [12], two VGG blocks [25] are introduced as lower layers before the stacked transformer layers². Each VGG block consists of two consecutive 3-by-3 convolution layers, a ReLU activation function, and a max-pooling layer; the first VGG block includes 32 channels, and the second VGG block has 64; 2-by-2 max-pooling is used in each block with stride 2 in the first and stride 1 in the second. The VGG blocks generate a 2560-D feature sequence at a 20ms frame rate. In training, the Adam optimizer [36] was used for all the models. Dropout was used: 0.5 for LC-BLSTMs and 0.1 for transformers. The LC-BLSTM baseline was optimized for at most 30 epochs on 16

¹The training-data segmentation was obtained from the alignments of an initial LC-BLSTM model. According to our studies, shorter segments in training can both improve the training throughput and decoding performance.

²As studied in [12], the VGG blocks are a best practice of input positional embedding for transformers. In experiments using VGG blocks on LC-BLSTM, insignificant gains obtained.

Nvidia V100 GPUs. The learning rate was initially 10^{-4} and reduced 50% after each epoch with accuracy degradation on the cross-validation data. Transformer models were optimized using a tri-stage learning-rate strategy: 8K updates with a learning rate increased linearly from 10^{-5} to a holding learning rate 3×10^{-4} , 100K updates with the holding learning rate, and further updates with the learning rate decreased exponentially. 32 GPUs were used in training one transformer model. Transformer models were updated up to 70 epochs.

The large configuration, i.e., approximately 80M parameters, was then investigated. The large LC-BLSTM baseline consists of 5 layers with 800 nodes in each layer each direction. The large transformer consists of 24 layers. The layer setting is identical to that of the small configuration; also, the same VGG blocks are used. The training schedule of LC-BLSTM and transformers followed a similar fashion as that in the small configuration. For large transformers, to alleviate the gradient vanishing issue, iterated loss [28] is applied. The outputs of the 6/12/18-th transformer layers are non-linearly transformed (projected to a 256-dimensional space with a linear transformation followed by a ReLU activation function), and auxiliary CE losses are calculated separately. These additional losses are interpolated with the original loss with an 0.3 weight.

In evaluation, a fully-optimized, static 4-gram decoding graph built by Kaldi was used. The results on test sets were obtained using the best epoch on the development set³. Following [37], the best checkpoints for `test-clean` and `test-other` are selected the respective development sets.

4.1.2. Segment and Context Length

We investigate the effect of segment length and context size first. A key issue on the proposed model is how to compromise between the accuracy and the look-ahead window size.

Table 1: *Effect of segment, left and right context length on LibriSpeech. Length is measured by number of frames, where frames are shifted in a 10ms frame rate.*

Left	Segment	Right	test-clean	test-other
0	64	0	10.7	13.9
0	96	0	9.8	13.0
0	128	0	7.7	10.4
16	128	0	5.2	9.5
32	128	0	3.6	8.5
64	128	0	3.5	8.5
0	128	16	5.5	9.3
0	128	32	3.8	8.1
32	128	32	3.3	8.0
64	128	32	3.3	7.6
-	∞	-	3.1	7.1

The decoding performance is reported in Table 1. The first block shows the results without context. By increasing the segment length from 64 to 128 frames, the word error rate (WER) decreased. Next, various context settings were investigated with the segment length fixed to 128 frames. The second and third blocks illustrate the effect of left and right contexts, respectively. Either left or right contexts contributed to alleviating the boundary effect. A more extended context was shown to improve the recognition accuracy. Finally, the effect of using both contexts was shown in the fourth block. The left and right con-

³A model that averaged the last 10 epochs was included as a model candidate.

texts showed some level of complementarity; thus, the performance further improved. The ∞ system refers to a transformer-based acoustic model presented in [12], indicating the performance lower-bound. The setting of 128 segment length, 64 left, and 32 right contexts was investigated in the following experiments. It yields a look-ahead window of 32 frames, i.e., 0.32 seconds, which is comparable to that of the LC-BLSTM baseline.

4.1.3. Limited Memory

The second set of experiments investigated the effect of limited memory size. Instead of the complete observation of augmented memory bank, models in this section were trained and tested by observing a fixed number of the most recent memory vectors. Note, when memory size equals to 1, our method becomes almost the same as the encoder used in [19]. These experiments were performed to investigate how much long-term history contributed to the final performance. Table 2 reports the results

Table 2: Effect of limited memory size on LibriSpeech.

MemSize	test-clean	test-other
0	3.2	8.1
1	3.3	8.0
3	3.2	7.9
5	3.3	7.9
∞	3.3	7.6

using different memory sizes. On the noisy set `test-other`, the performance was consistently improved from no memory (0) to unlimited memory (∞)⁴. However on the clean data `test-clean`, little improvement was obtained. This observation indicates that the global information in long-term memory can alleviate more challenging acoustic conditions.

4.1.4. Comparison with Other Streamable Models

Table 3 compares the WERs of different models. For a fair comparison, corresponding models of similar look-ahead window are compared. The first block compares models with about 40M parameters. The transformer-XL baseline used a segment length of 128, which is identical to that of the proposed model. The “+look-ahead” reports the extension of transformer-XL with right context⁵. The TRTrf baseline used a context of 3 in each layer, resulting in a look-ahead window of 0.72 second. The proposed augmented memory transformer outperformed all the streamable baselines. Larger models with about 80M parameters are compared in the second block. The augmented memory transformer shows consistent gains as the small-size one. For further improvement, the weak-attention suppression (WAS) [38] was applied on top of the proposed model, denoted by “+WAS”. Compared with the LC-BLSTM baseline, the augmented memory transformer (with WAS) achieved 15%-18% relative error reduction on the two test sets. At the time of writing, this is the best number that we acknowledge on LibriSpeech for streamable models.

4.2. Video ASR

To evaluate the model in more challenging acoustic conditions, our in-house Russian and German video ASR datasets were

⁴The longest utterance in the LibriSpeech test sets is about 35 seconds. Thus, the ∞ system used maximum 28 memory slots.

⁵There is no context in the original design of transformer-XL. We applied a similar idea of right context (32 frames) on transformer-XL as the proposed model. Thus, both model has a look-ahead window of 0.32 second.

Table 3: Performance of different models on LibriSpeech. “Str” stands for streamable, specifying if a model is a streamable one.

#Param	Str	Model	test-clean	test-other
$\approx 40M$	✓	LC-BLSTM	3.8	9.9
		Trf-XL	4.2	10.7
		+look-ahead	3.9	10.1
		TRTrf	4.1	9.0
		AMTrf	3.3	7.6
	✗	Trf	3.1	7.1
$\approx 80M$	✓	LC-BLSTM	3.3	8.2
		Trf-XL	3.5	8.3
		+ look-ahead	3.2	7.7
		AMTrf	3.1	7.1
		+WAS	2.8	6.7
	✗	Trf	2.6	5.6

used. The videos in this dataset are originally shared publicly by users; only the audio part of those videos is used in our experiments. These data are completely de-identified; both transcribers and researchers do not have access to any user-identifiable information. For the Russian task, the training data consisted of 1.8K hours from 100K video clips. 14.6 hours of audio (790 video clips) were used as validation data. Two test sets were used in evaluation: the 11-hour `clean` (466 videos), and 24-hour `noisy` (1.3K videos) sets. For the German task, the training data consisted of 3K hour audios (135K videos). The validation data was 14.5 hours (632 videos). The test data were the 25-hour `clean` (989 videos) and 24-hour `noisy` (1K videos) sets.

Table 4: Experiment results on our internal video ASR tasks.

Language	Model	clean	noisy
Russian	LC-BLSTM	19.8	24.4
	AMTrf	18.0	23.3
	Trf	16.6	21.1
German	LC-BLSTM	19.6	19.5
	AMTrf	17.4	17.1
	Trf	16.2	15.6

The large network configuration, i.e., 80M-parameter models, was examined. The training of all the models was performed in a similar fashion as presented in Section 4.1.1 (large configuration). Table 4 summarizes the decoding results. On both languages, the proposed model consistently outperformed the LC-BLSTM baseline by 9-11% on clean test sets and 5-12% on noisy test sets. There are still some accuracy gaps compared with the transformer which has the access to the whole utterance.

5. Conclusions

In this work, we proposed the augmented memory transformer for streaming transformer-based models for speech recognition. It processes sequence data incrementally using short segments and an augmented memory, thus has the potential for latency-constrained tasks. On LibriSpeech, the proposed model outperformed LC-BLSTM and all the existing streamable transformer baselines. Initial study on more challenging Russian and German video datasets also illustrated similar conclusions.

In the future, we will investigate the real-time latency and measure the throughput of this model. The proposed method can be also applied to transformer transducer [14, 39] or transformer-based sequence-to-sequence models (e.g., [26, 11]).

6. References

- [1] H. A. Bourlard and N. Morgan, *Connectionist speech recognition: a hybrid approach*. Springer Science & Business Media, vol. 247.
- [2] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *Proc. ICASSP*, 2016, pp. 4945–4949.
- [3] C.-C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen, Z. Chen, A. Kannan, R. J. Weiss, K. Rao, E. Gonina *et al.*, "State-of-the-art speech recognition with sequence-to-sequence models," in *Proc. ICASSP*, 2018.
- [4] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang *et al.*, "Streaming end-to-end speech recognition for mobile devices," in *Proc. ICASSP*, 2019, pp. 6381–6385.
- [5] G. Hinton, L. Deng, D. Yu *et al.*, "Deep neural networks for acoustic modeling in speech recognition," *IEEE Signal processing magazine*, vol. 29, 2012.
- [6] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [7] H. Sak, A. Senior, and F. Beaufays, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling," in *Proc. Interspeech*, 2014.
- [8] A. Vaswani, N. Shazeer, N. Parmar *et al.*, "Attention is all you need," in *Proc. NIPS*, 2017, pp. 5998–6008.
- [9] J. Devlin, M.-W. Chang, K. Lee *et al.*, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [10] A. Radford, K. Narasimhan, T. S. *et al.*, "Improving language understanding by generative pre-training," 2018.
- [11] S. Karita, N. Chen, T. Hayashi *et al.*, "A Comparative Study on Transformer vs RNN in Speech Applications," *arXiv preprint arXiv:1909.06317*, 2019.
- [12] Y. Wang, A. Mohamed, D. Le, C. Liu *et al.*, "Transformer-based acoustic modeling for hybrid speech recognition," *arXiv preprint arXiv:1910.09799*, 2019.
- [13] D. Povey, H. Hadian, P. Ghahremani *et al.*, "A time-restricted self-attention layer for asr," in *Proc. ICASSP*, 2018, pp. 5874–5878.
- [14] Q. Zhang, H. Lu, H. Sak, A. Tripathi, E. McDermott, S. Koo, and S. Kumar, "Transformer transducer: A streamable speech recognition model with transformer encoders and RNN-T loss," in *Proc. ICASSP*, 2020, pp. 7829–7833.
- [15] N. Moritz, T. Hori, and J. L. Roux, "Streaming automatic speech recognition with the transformer model," *arXiv preprint arXiv:2001.02674*, 2020.
- [16] L. Dong, F. Wang, and B. Xu, "Self-attention aligner: A latency-control end-to-end model for asr using self-attention network and chunk-hopping," in *Proc. ICASSP*, 2019, pp. 5656–5660.
- [17] K. Chen and Q. Huo, "Training deep bidirectional LSTM acoustic model for LVCSR by a Context-Sensitive-Chunk BPTT approach," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, 07 2016.
- [18] A.-r. Mohamed, F. Seide, D. Yu, J. Droppo *et al.*, "Deep bidirectional recurrent networks over spectral windows," in *Proc. Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2015, pp. 78–83.
- [19] E. Tsunoo, Y. Kashiwagi, T. Kumakura, and S. Watanabe, "Transformer ASR with Contextual Block Processing," in *Proc. Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 427–433.
- [20] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-XL: Attentive language models beyond a fixed-length context," *arXiv preprint arXiv:1901.02860*, 2019.
- [21] Y. Zhang, G. Chen, D. Yu, K. Yao *et al.*, "Highway long short-term memory RNNs for distant speech recognition," in *Proc. ICASSP*, 2016, pp. 5755–5759.
- [22] Z. Tian, J. Yi, Y. Bai *et al.*, "Synchronous transformers for end-to-end speech recognition," *arXiv preprint arXiv:1912.02958*, 2019.
- [23] A. Graves, G. Wayne, and I. Danihelka, "Neural Turing machines," *arXiv preprint arXiv:1410.5401*, 2014.
- [24] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *Proc. ICASSP*, 2015, pp. 5206–5210.
- [25] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [26] A. Mohamed, D. Okhonko, and L. Zettlemoyer, "Transformers with convolutional context for asr," *arXiv preprint arXiv:1904.11660*, 2019.
- [27] J. Lei Ba, J. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [28] A. Tjandra, C. Liu, F. Zhang, X. Zhang, Y. Wang, G. Synnaeve, S. Nakamura, and G. Zweig, "Deja-vu: Double feature presentation and iterated loss in deep transformer networks," in *Proc. ICASSP*, 2020, pp. 6899–6903.
- [29] J. W. Rae, A. Potapenko, S. M. Jayakumar, and T. P. Lillicrap, "Compressive transformers for long-range sequence modelling," *arXiv preprint arXiv:1911.05507*, 2019.
- [30] O. Myle, E. Sergey, B. Alexei, F. Angela *et al.*, "fairseq: A Fast, Extensible Toolkit for Sequence Modeling," in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- [31] D. Le, X. Zhang, W. Zheng *et al.*, "From senones to chenones: Tied context-dependent graphemes for hybrid speech recognition," *arXiv preprint arXiv:1910.01493*, 2019.
- [32] D. Povey, A. Ghoshal, G. Boulianne *et al.*, "The kaldia speech recognition toolkit," in *IEEE Workshop on Automatic Speech Recognition and Understanding*, 2011.
- [33] T. Ko, V. Peddinti, D. Povey *et al.*, "Audio augmentation for speech recognition," in *Proc. Interspeech*, 2015.
- [34] D. S. Park, W. Chan, Y. Zhang *et al.*, "SpecAugment: A simple data augmentation method for automatic speech recognition," *arXiv preprint arXiv:1904.08779*, 2019.
- [35] V. Peddinti, Y. Wang, D. Povey, and S. Khudanpur, "Low latency acoustic modeling using temporal convolution and lstms," *IEEE Signal Processing Letters*, vol. 25, no. 3, pp. 373–377, 2017.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [37] C. Lüscher, E. Beck, K. Irie *et al.*, "RWTH ASR Systems for LibriSpeech: Hybrid vs Attention-w/o Data Augmentation," *arXiv preprint arXiv:1905.03072*, 2019.
- [38] Y. Shi, Y. Wang, C. Wu, C. Fuegen, F. Zhang, D. Le, C.-F. Yeh, and M. Seltzer, "Weak-Attention Suppression For Transformer Based Speech Recognition," in *Proc. Interspeech*, 2020.
- [39] C.-F. Yeh, J. Mahadeokar, K. Kalgaonkar, Y. Wang, D. Le, M. Jain, K. Schubert, C. Fuegen, and M. L. Seltzer, "Transformer-Transducer: End-to-End Speech Recognition with Self-Attention," *arXiv preprint arXiv:1910.12977*, 2019.