

EMFORMER: EFFICIENT MEMORY TRANSFORMER BASED ACOUSTIC MODEL FOR LOW LATENCY STREAMING SPEECH RECOGNITION

Yangyang Shi, Yongqiang Wang, Chunyang Wu, Ching-Feng Yeh, Julian Chan, Frank Zhang, Duc Le, Mike Seltzer

Facebook AI

ABSTRACT

This paper proposes an efficient memory transformer *Emformer* for low latency streaming speech recognition. In *Emformer*, the long-range history context is distilled into an augmented memory bank to reduce self-attention’s computation complexity. A cache mechanism saves the computation for the key and value in self-attention for the left context. *Emformer* applies a parallelized block processing in training to support low latency models. We carry out experiments on benchmark LibriSpeech data. Under average latency of 960 ms, *Emformer* gets WER 2.50% on test-clean and 5.62% on test-other. Comparing with a strong baseline augmented memory transformer (AM-TRF), *Emformer* gets 4.6 folds training speedup and 18% relative real-time factor (RTF) reduction in decoding with relative WER reduction 17% on test-clean and 9% on test-other. For a low latency scenario with an average latency of 80 ms, *Emformer* achieves WER 3.01% on test-clean and 7.09% on test-other. Comparing with the LSTM baseline with the same latency and model size, *Emformer* gets relative WER reduction 9% and 16% on test-clean and test-other, respectively.

Index Terms— Low Latency, Transformer, Emformer

1. INTRODUCTION

Transformers [1] have achieved dominated performance for various tasks in natural language processing area [2, 3, 4]. Rather than using memory state to capture long-range dependencies in recurrent neural networks, the multi-head self-attention method connects arbitrary positions in the whole sequence directly in parallel.

Recently, transformer-based model architectures have also been successfully applied to automatic speech recognition (ASR) area across various modeling paradigms, including sequence-to-sequence [5, 6, 7, 8, 9], neural transducer [10, 11, 12], Connectionist temporal classification (CTC) [13, 14] and traditional hybrid [15, 16] systems.

Unlike most natural language processing tasks, many ASR applications deal with streaming scenarios challenging for vanilla transformers. The streaming recognizer needs to produce output given partially available speech utterance rather than entire utterance. Several methods advance the transformer for streaming speech recognition. The work [15, 10, 17] proposed to constrain the attention computation with a limited length of look-ahead inputs. However, these methods have a significant delay due to the look-ahead context leaking issue where essential look-ahead context grows linearly with the number of transformer layers stacking on top of one another. A scout network is proposed in [9] to detect the word boundary. In scout networks, only the context information before the word boundary is used by the transformer to make predictions. However, the scout network does not address the heavy self-attention computation that grows quadratically with the left context length. A

streaming transformer with augmented memory (AM-TRF) is proposed in [18] to reduce latency and the self-attention computation.

AM-TRF uses a similar block processing method as [19]. The block processing chunks the whole utterance into multiple segments. To reduce the computation in capturing the long-range left context, AM-TRF introduces a memory bank. Each vector in the memory bank is an abstract embedding from the previous one segment. The direct left context block from the current segment and look-ahead context block provides context information for current segment recognition in addition to the memory bank. However, AM-TRF has duplicated computations for the direct left context block in both training and decoding. The memory bank carries over the context information from previous segments in a similar auto-regression way as recurrent neural networks. The inherent auto-regression characteristic makes AM-TRF challenging to parallelize the block processing in training.

In this paper, we propose the *Emformer* that improves the AM-TRF from the following aspects. First, *Emformer* removes the duplicated computation from the left context block by caching the key and value in previous segments’ self-attention. Second, rather than passing the memory bank within the current layer in AM-TRF, inspired by transformer-xl [2] and its application in speech recognition [20], *Emformer* carries over the memory bank from the lower layer. Third, *Emformer* disables the summary vector’s attention with memory bank to avoid overweighting the most left part of context information. Finally, *Emformer* applies a parallelized block processing training method, which is important to train *Emformer* for low latency speech recognition.

To verify the performance of the proposed method, we carry out experiments on LibriSpeech [21]. More experiments using industry dataset with variant scenarios are in [22]. Under the average latency of 640 ms constraint, comparing with AM-TRF, *Emformer* gets relative WER reduction 17% on test-clean and 9% on test-other. Meanwhile, *Emformer* reduces the training time by almost 80% and decoding RTF by 18%. For a low latency scenario with an average latency of 80 ms, *Emformer* saves more than 91% computation from AM-TRF and obtain WER 3.01% on test-clean and 7.09% on test-other. According to our knowledge, this is the first work to give streaming transformer results on LibriSpeech with such low latency. Under the average latency of 960 ms and 640 ms constraint, *Emformer* also gives the best result on LibriSpeech so far.

2. EMFORMER

Emformer improves over the AM-TRF. The following subsection gives a short introduction to AM-TRF.

2.1. AM-TRF

Figure (1a) illustrates the operations in one AM-TRF layer. A sequence of input feature vectors are chunked into multiple non-




Fig. 1: Comparison of AM-TRF with Emformer

overlapping segments C_1^n, \dots, C_{i-1}^n , where the i denotes the index of segment, and n the layer's index. In order to reduce boundary effect, left and right contextual blocks, L_i^n and R_i^n , are concatenated with C_i^n to form a contextual segment $\mathbf{X}_i^n = [L_i^n, C_i^n, R_i^n]$. At the i -th segment, the n -th AM-TRF layer accepts \mathbf{X}_i^n and a bank of memory vector $\mathbf{M}_i^n = [m_1^n, \dots, m_{i-1}^n]$ as the input, and produces $\mathbf{X}_i^{n+1} = [L_i^{n+1}, C_i^{n+1}, R_i^{n+1}]$ and m_i^n as the output, whereas \mathbf{X}_i^{n+1} is feed to the next layer and m_i^n is inserted into the memory bank to generate \mathbf{M}_{i+1}^n and carried over to the next segment. After all the AM-TRF layers, the center blocks $\{C_i^{N-1}\}_{i=0}^{I-1}$ are concatenated as the encoder output sequence; the contextual blocks $\{L_i^{N-1}\}_{i=0}^{I-1}$ and $\{R_i^{N-1}\}_{i=0}^{I-1}$ are discarded.

At the core of each AM-TRF layer, there is a modified attention mechanism which attends to the memory bank and yields a new memory vector at each segment:

$$\hat{\mathbf{X}}_i^n = \text{LayerNorm}(\mathbf{X}_i^n) \quad (1)$$

$$\mathbf{K}_i^n = \mathbf{W}_k [M_i^n, \hat{\mathbf{X}}_i^n], \quad (2)$$

$$\mathbf{V}_i^n = \mathbf{W}_v [M_i^n, \hat{\mathbf{X}}_i^n], \quad (3)$$

$$[\mathbf{Z}_{L,i}^n, \mathbf{Z}_{C,i}^n, \mathbf{Z}_{R,i}^n] = \text{Attn}(\mathbf{W}_q \hat{\mathbf{X}}_i^n, \mathbf{K}_i^n, \mathbf{V}_i^n) + \mathbf{X}_i^n \quad (4)$$

$$\mathbf{m}_i^n = \text{Attn}(\mathbf{W}_q s_i^n, \mathbf{K}_i^n, \mathbf{V}_i^n) \quad (5)$$

whereas $\mathbf{Z}_{L,i}^n$, $\mathbf{Z}_{C,i}^n$ and $\mathbf{Z}_{R,i}^n$ are the attention output for L_i^n , C_i^n and R_i^n respectively; s_i^n is the mean of center block C_i^n ; $\text{Attn}(\mathbf{q}; \mathbf{k}, \mathbf{v})$ is the attention operation defined in [1] with \mathbf{q} , \mathbf{k} and \mathbf{v} being the query, key and value, respectively.

$\mathbf{Z}_{L,i}^n$, $\mathbf{Z}_{C,i}^n$, $\mathbf{Z}_{R,i}^n$ are passed to a point-wise feed-forward network (FFN) with layer normalization and residual connection to generate the output of this AM-TRF layer, i.e.,

$$\hat{\mathbf{X}}_{i+1}^n = \text{FFN}(\text{LayerNorm}([\mathbf{Z}_{L,i}^n, \mathbf{Z}_{C,i}^n, \mathbf{Z}_{R,i}^n])) \quad (6)$$

$$\mathbf{X}_i^{n+1} = \text{LayerNorm}(\hat{\mathbf{X}}_{i+1}^n + [\mathbf{Z}_{L,i}^n, \mathbf{Z}_{C,i}^n, \mathbf{Z}_{R,i}^n]) \quad (7)$$

where FFN is a two-layer feed-forward network with Relu non-linearity. The last layer normalization in Eq. (7) is used to prevent a path to bypass all the AM-TRF layers.

2.2. Emformer

As shown in [18], given the similar latency constraint, AM-TRF has outperformed previous streaming transformer models. However, there are several issues with AM-TRF. The usage of the left context is not efficient. AM-TRF training relies on the sequential block

processing that is not suitable for low latency model training. Having observed these limitations, we proposed a new streamable transformer architecture, namely, Emformer. One layer of Emformer is demonstrated in Figure (1b). The following subsections describe the important improvements made in Emformer.

2.2.1. Cache key and value from previous segments

As illustrated in Figure (1a), for the i -th segment, the embedding of the left context L_i^n needs to be re-computed for every step, even though L_i^n is overlapped with C_{i-1}^n (or possibly even more previous center blocks). Thus we only need to cache the projections from the previous segments. As shown in Figure (1b), Emformer only computes the *key*, *value* projections for the memory bank, center, and right context; Emformer saves the computation of *query* projection of left context, as it does not need to give output from the left context block for the next layer. Compared with AM-TRF, the attention part in Emformer operates in the following sequence:

$$[\hat{C}_i^n, \hat{R}_i^n] = \text{LayerNorm}([C_i^n, R_i^n]) \quad (8)$$

$$\mathbf{K}_i^n = [\mathbf{W}_k M_i^n, \mathbf{K}_{L,i}^n, \mathbf{W}_k C_i^n, \mathbf{W}_k R_i^n], \quad (9)$$

$$\mathbf{V}_i^n = [\mathbf{W}_v M_i^n, \mathbf{V}_{L,i}^n, \mathbf{W}_v C_i^n, \mathbf{W}_v R_i^n], \quad (10)$$

$$\mathbf{Z}_{C,i}^n = \text{Attn}(\mathbf{W}_q \hat{C}_i^n, \mathbf{K}_i^n, \mathbf{V}_i^n) + C_i^n \quad (11)$$

$$\mathbf{Z}_{R,i}^n = \text{Attn}(\mathbf{W}_q \hat{R}_i^n, \mathbf{K}_i^n, \mathbf{V}_i^n) + R_i^n \quad (12)$$

$$\mathbf{m}_i^n = \text{Attn}(s_i^n; \mathbf{K}_i^n, \mathbf{V}_i^n) \quad (13)$$

where $\mathbf{K}_{L,i}^n$ and $\mathbf{V}_{L,i}^n$ are the *key* and *value* copies from previous segments with no additional computations.

Let's assume L , C , R , and M are the lengths for the left context block, the center context, the right context, and the memory bank. The number of heads in the multi-head self-attention is h and per head dimension is d . Note the summary vector is the mean of the center segment, of which length is always 1. In practice, the memory bank is implemented in ring buffer way with small length, and the model dimension, dh , is much larger than any of L , C , R , and M . Emformer saves approximately $\frac{L}{L+C+R}$ of AM-TRF computation. For low latency scenario with center context length 80 ms, right context length 40 ms, and left context length 1280 ms, Emformer reduces more than 91% computation from AM-TRF.

2.2.2. Carryover memory vector from previous segments in the lower layer

The attention output from the summary vector s_i^n is a memory vector in the memory bank. The memory bank carries all the previous

context information for future segments. As we can see from Figure (1a), the memory vector \mathbf{m}_i^n from the i -th segment in the n -th layer is a prerequisite for the $(i+1)$ -th segment from the same layer. In training, the auto-regression characteristic of AM-TRF forces the block processing to be in a sequential way that is not suitable for GPU computing. Especially for low latency model training, where the center segment is small, sequential block processing chunks the whole utterance computation into a small computation loop, which renders extremely low GPU usage.

To support parallelization for block processing training, Emformer takes the memory bank input from previous segments in the lower layer rather than the same layer. In this way, for each Emformer layer, the whole sequence is trained in parallel, fully taking advantage of the GPU computing resources.

2.2.3. Disallow attention between the summary vector with the memory bank

According to Eq. (5), the memory vector is a weighted interpolation of *values* projected from the memory bank, the left context block, the center block, and the right context block. For both AM-TRF and Emformer, assigning the attention weight between the summary vector and the memory bank to zero stabilizes the training and improves recognition accuracy for long-form speech. There are two issues by including the memory bank information in the current memory vector. The most left context information is over-weighted because it is contained in every memory vector in the memory bank. The other one is similar to a recurrent neural network, where a recurrent connection of the memory bank with a memory vector could cause gradient vanishing or explosion. For AM-TRF, the usage of the weak-attention suppression method [18, 23] partially addresses the problem by setting weak-attention weights to zero.

2.2.4. Deal with look-ahead context leaking

The sequential block processing in AM-TRF training chunks the input sequence physically. The right context size bounds the look-ahead reception field. However, sequentially processing blocks significantly slows the training. Now Emformer processes the in-




Fig. 2: Illustration of avoiding look-ahead context leaking

put sequence in a fully parallel manner in the training stage. Like transformer-XL [2], Emformer applies attention masks to limit the reception field in each layer without physically chunking the input sequence. However, this method has the risk of a look-ahead of context leaking. As shown in the left part of Figure (2), even the look-ahead context size is limited to 1 in each layer; the total look-ahead context size is 4 when four layers of Emformer are used. To deal with the look-ahead context leaking issue in training, Emformer makes a hard copy of each segment’s look-ahead context and puts the look-ahead context copy at the input sequence’s beginning as illustrated in the right part of Figure (2).

3. EXPERIMENTS

3.1. Data and Setup

We verify the proposed method on the LibriSpeech corpus [21]. LibriSpeech has 1000 hours of book reading utterances derived from the LibriVox project. There are two subsets of development data and evaluation data in LibriSpeech. The “clean” subsets contain simple and clean utterances. The “other” subset contains complex and noisy utterances. Based on the WER on the dev data, we select the best model and report its WER on test data. In the experiments, Emformer is used as an encoder for both the hybrid [16, 18, 14] and transducer [10, 11, 12] models.

3.1.1. Hybrid model

The context and positional dependent graphemes are used as output units [24]. We use the standard Kaldi [25] LibriSpeech recipe to build bootstrap the HMM-GMM system. The 80-dimensional log Mel filter bank features at a 10 ms frame rate are used. We also apply speed perturbation [26] and *SpecAugment* [27] without time warping to stabilize the training.

A linear layer maps the 80-dimensional features to 128 dimension vectors. Four continuous 128-dimensional vectors are concatenated with stride 4 to form a 512 vector that is the input to Emformer. In Emformer, each layer has eight heads of self-attention. The input and output for each layer have 512 nodes. The inner-layer of FFN has dimensionality 2048. Dropout is 0.1 for all layers across all experiments. For medium latency, memory bank length is 4. For low latency experiments where the segment size is small, memory bank information largely overlaps with direct left context. Therefore, we set the memory bank length to 0. An auxiliary incremental loss [28] with weight 0.3 is used to overcome the training divergence issue for deep transformer models. All hybrid models are trained with the adam optimizer [29] using 180 epochs. The learning rate increases to 1e-3 in 20K warming-up updates. Then it is fixed until 100 epochs. From then on, the learning rate shrinks every epoch with factor 0.95. All the models are trained using 32 Nvidia V100 GPUs with fp16 precision. We use hosts with Intel Xeon D-2191A 18-core CPUs to measure real time factors (RTFs). In measuring RTFs, 10 utterances are concurrently decoded.

3.1.2. Transducer model

The output units are 1024 sentence pieces [30] with byte pair encoding (BPE) [31] as the segmentation algorithm. In the predictor, the tokens are first represented by 256-dimensional embeddings before going through two LSTM layers with 512 hidden nodes, followed by a linear projection to 640-dimensional features before the joiner. For the joiner, the combined embeddings from the encoder and the predictor first go through a Tanh activation and then another linear projection to the target number of sentence pieces. Both the LCBLSTM and Emformer encoders are pre-trained from the hybrid systems. Similar to [12], we use a neural network language model (NNLM) for shallow fusion during beam search where the weight for NNLM probabilities was 0.3 across experiments. The training data for NNLM is the combined transcripts of the train set and the 800M text-only set.

3.2. Results

3.2.1. Algorithmic latency induced by the encoder (EIL)

In block processing based decoding, the latency comes from the center block size and the look-ahead context size. For the most left frame in the center block, the latency is the center block size plus look-ahead context size. The latency for the most right frame in the center block is look-ahead context size. Therefore, we use algorithmic latency induced by the encoder (EIL), an average latency of all

the frames in the center block, which equals to the look-ahead context latency plus center block latency discounted by 0.5.

3.2.2. From AM-TRF to Emformer

Table 1 gives a performance comparison of AM-TRF with Emformer with a latency of 960 ms. Caching the key and value computation speeds up the training from 1.14 hours per epoch to 0.5 hours per epoch and decoding from RTF (real-time factor) 0.19 to 0.17. The left context caching also reduces the redundant gradient in training that results in some WER reduction¹. Finally, using all improvements, comparing with AM-TRF, Emformer speeds up the training by 4.6 folds. Emformer also gets relative WER reduction 17% on test-clean, 9% on test-other and 18% relative RTF reduction in decoding. For a low latency scenario, Emformer saves up to 91% of computations from AM-TRF without considering parallel block processing. It is impractical to train AM-TRF for a low latency scenario. Therefore we ignore the detailed comparison.

Model	RTF	test		train hours per epoch
		clean	other	
AM-TRF-24L	0.16	3.27	6.66	1.14h
+ left context caching	0.13	2.88	6.44	0.50h
EM-24L	0.13	2.72	6.01	0.25h

Table 1: From AM-TRF to Emformer based on hybrid systems. All models have 80M parameters. Left context size, center block size and right context size are 640 ms, 1280 ms and 320 ms, respectively.

3.2.3. Results from hybrid systems

Model	LC size	Center Size	test		RTF
			clean	other	
LCBLSTM	–	1280	2.90	6.76	0.25
	–	640	2.96	6.97	0.27
EM-24L	320	1280	2.75	6.08	0.13
	640		2.72	6.01	0.13
	1280		2.59	5.90	0.13
	320	640	2.80	6.47	0.13
	640		2.78	6.46	0.13
	1280		2.76	6.59	0.15
EM-36L +smbr	1280	1280	2.58	5.75	0.17
			2.50	5.62	0.17
EM-36L +smbr	1280	640	2.69	6.14	0.20
			2.62	5.97	0.19

Table 2: Impact of left context (LC) size (in millisecond) on WER and RTF under medium latency constraints for hybrid models. Look-ahead size is 320 ms, the EIL is 640 ms or 960 ms when center size is 640 ms and 1280 ms, respectively. Both LCBLSTM and EM-24L have the similar 80M parameters. EM-36L has 120M parameters.

Table 2 and Table 3 presents the performance of the Emformer based hybrid systems for medium latency and low latency, respectively. For both tables, larger left context size gives better WER and slightly worse decoding RTF. In Table 2, LCBLSTM consists of 5 layers with 800 nodes in each layer each direction. Using a similar model size and latency constraint, Emformer gets a relative 48% RTF deduction. Under EIL 1280 ms, Emformer obtained over relative 12% WER reduction over LCBLSTM on both test-clean and test-other datasets. Together with sMBR training [32], the Emformer with 120M parameters achieves WER 2.50% on test-clean

¹For large datasets, the caching strategy does not give WER reduction.

and 5.62% on test-other under EIL 960 ms, and 2.62% on test-clean and 5.97% on test-other under EIL 640 ms.

In Table 3, the LSTM consists of 7 layers with 1200 nodes in each layer. The input to LSTM is a concatenation of the current frame with eight look-ahead context frames. Low latency speech recognition gives higher RTF than medium latency speech recognition. Because medium latency speech recognition chunks an utterance into fewer larger segments, it speeds up the neural network’s computation. Using a similar model size and latency constraint, Emformer gets relative WER reduction 9% and 15% on test-clean and test-other, respectively. Together with sMBR training [32], the 36 layer Emformer achieves WER 3.01% on test-clean and 7.09% on test-other. According to our knowledge, for low latency 80 ms, Emformer gives the best WER on LibriSpeech data.

Model	LC size in milliseconds	latency	test		RTF
			clean	other	
LSTM	–	80	3.75	9.18	0.25
EM-24L	320	80	3.44	8.37	0.30
	640		3.37	8.05	0.31
	1280		3.41	7.75	0.33
EM-36L +smbr	1280	80	3.32	7.56	0.49
			3.01	7.09	0.49

Table 3: Impact of left context (LC) size (in millisecond) on word error rate and RTF under a low latency constraint for hybrid models. The look-ahead and center context size are 40 ms and 80 ms, respectively. Latency is defined by encoder induced latency (EIL). Both LSTM and EM-24L have the similar 80M parameters. EM-36L has 120M parameters.

3.2.4. Results from transducer systems

Table 4 summarizes the comparison between LCBLSTM and Emformer as encoders in the transducer system. Similar to the previous observations with hybrid systems, we see that given the same EIL (640 ms), Emformer consistently outperforms LCBLSTM on WER. With the external NNLM, the transducer systems achieved similar WER to those from hybrid systems.

Model	NNLM	test	
		clean	other
LCBLSTM	✗	3.04	8.25
	✓	2.65	7.26
EM-24L	✗	2.78	6.92
	✓	2.37	6.07

Table 4: WER of Emformer with the neural transducers. Both models use an EIL 640 ms with center context 640 ms and look-ahead context 320 ms. Left context size is 1280 ms.

4. CONCLUSIONS

The proposed Emformer applied a cache strategy to remove the duplicated computation in augmented memory transformer (AM-TRF) for the left context. Emformer disabled the summary vector attention with a memory bank to stabilize the training. By redefining the memory carryover procedure and avoiding the right context leaking, Emformer supported parallelized block processing in training. Comparing with AM-TRF, Emformer got 4.6 folds of training speedup and 18% decoding RTF reduction. Experiments on LibriSpeech showed that Emformer outperformed the baselines in both hybrid and transducer systems. Under average latency EIL 960 ms, Emformer achieved WER 2.50% on test-clean and 5.62% on test-other with decoding RTF 0.13. Under low latency 80 ms constraint, Emformer achieved WER 3.01% on test-clean and 7.09% on test-other.

5. REFERENCES

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017.
- [2] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Quoc V. Le, and R. Salakhutdinov, "Transformer-XL: Attentive language models beyond a fixed-length context," *ACL*, pp. 2978–2988, 2019.
- [3] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *NAACL*, vol. 1, pp. 4171–4186, 2019.
- [4] C. Raffel, N. Shazeer, A. Roberts, and Others, "Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer," *arXiv preprint arXiv:1910.10683*, 2019.
- [5] L. Dong, S. Xu, and B. Xu, "Speech-transformer: a non-recurrence sequence-to-sequence model for speech recognition," in *ICASSP*, 2018.
- [6] S. Karita, N. Chen, T. Hayashi, and Others, "A Comparative Study on Transformer vs RNN in Speech Applications," *arXiv preprint arXiv:1909.06317*, 2019.
- [7] M. Sperber, J. Niehues, G. Neubig, and Others, "Self-attentional acoustic models," *arXiv preprint arXiv:1803.09519*, 2018.
- [8] S. Zhou, L. Dong, S. Xu, and B. Xu, "Syllable-based sequence-to-sequence speech recognition with the transformer in mandarin Chinese," *arXiv preprint arXiv:1804.10752*, 2018.
- [9] C. Wang, Y. Wu, S. Liu, J. Li, et al., "Low Latency End-to-End Streaming Speech Recognition with a Scout Network," *arXiv preprint arXiv:12003.10369*, 2020.
- [10] 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," *ICASSP*, vol. 2020-May, pp. 7829–7833, 2020.
- [11] C.-F. Yeh, J. Mahadeokar, and Others, "Transformer-Transducer: End-to-End Speech Recognition with Self-Attention," *arXiv preprint arXiv:11910.12977*, 2019.
- [12] A. Gulati, J. Qin, C.-C. Chiu, et al., "Conformer: Convolution-augmented Transformer for Speech Recognition," *arXiv preprint arXiv:12005.08100*, 2020.
- [13] J. Salazar, K. Kirchhoff, and Z. Huang, "Self-Attention Networks for Connectionist Temporal Classification in Speech Recognition," in *Proceedings of ICASSP*, 2019.
- [14] F. Zhang, Y. Wang, X. Zhang, C. Liu, et al., "Fast, Simpler and More Accurate Hybrid ASR Systems Using Wordpieces," *InterSpeech*, 2020.
- [15] D. Povey, H. Hadian, P. Ghahremani, and Others, "A time-restricted self-attention layer for asr," in *Proc. ICASSP*, 2018, pp. 5874–5878.
- [16] Y. Wang, A. Mohamed, D. Le, and Others, "Transformer-Based Acoustic Modeling for Hybrid Speech Recognition," *ICASSP*, 2019.
- [17] N. Moritz, T. Hori, and J. L. Roux, "Streaming automatic speech recognition with the transformer model," *arXiv preprint arXiv:2001.02674*, 2020.
- [18] C. Wu, Y. Shi, Y. Wang, and C.-F. Yeh, "Streaming Transformer-based Acoustic Modeling Using Self-attention with Augmented Memory," in *InterSpeech*, 2020.
- [19] 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," *ICASSP*, pp. 5656–5660, 2019.
- [20] L. Lu, C. Liu, J. Li, and Y. Gong, "Exploring Transformers for Large-Scale Speech Recognition," in *InterSpeech*, 2020.
- [21] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *ICASSP*, 2015.
- [22] Y. Wang, Y. Shi, F. Zhang, C. Wu, and Others, "Transformer in action: a comparative study of transformer-based acoustic model for large scale speech recognition applications," <https://arxiv.org/abs/2010.14665>, 2020.
- [23] Y. Shi, Y. Wang, C. Wu, C. Fuegen, et al., "Weak-Attention Suppression For Transformer Based Speech Recognition," in *InterSpeech*, 2020.
- [24] D. Le, X. Zhang, W. Zheng, and Others, "From Senones to Chenones: Tied Context-Dependent Graphemes for Hybrid Speech Recognition," *arXiv preprint arXiv:1910.01493*, 2019.
- [25] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, "The Kaldi speech recognition toolkit," *ASRU*, pp. 1–4, 2011.
- [26] T. Ko, V. Peddinti, D. Povey, and Others, "Audio augmentation for speech recognition," in *InterSpeech*, 2015.
- [27] D S Park, W Chan, Y Zhang, and Others, "SpecAugment: A simple data augmentation method for automatic speech recognition," *arXiv preprint arXiv:1904.08779*, 2019.
- [28] A. Tjandra, C. Liu, F. Zhang, and Others, "Deja-vu: Double Feature Presentation and Iterated loss in Deep Transformer Networks," *ICASSP*, 2020.
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [30] T. Kudo and J. Richardson, "SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," *EMNLP*, pp. 66–71, 2018.
- [31] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," *ACL*, 2016.
- [32] K. Vesely, A. Ghoshal, L. Burget, and D. Povey, "Sequence-discriminative training of deep neural networks.," in *InterSpeech*, 2013, vol. 2013, pp. 2345–2349.