



Self-regularised Minimum Latency Training for Streaming Transformer-based Speech Recognition

Mohan Li, Rama Doddipatla and Cătălin Zorilă

Cambridge Research Laboratory, Toshiba Europe Ltd, Cambridge, UK

{mohan.li, rama.doddipatla, catalin.zorila}@crl.toshiba.co.uk

Abstract

This paper proposes a self-regularised minimum latency training (SR-MLT) method for streaming Transformer-based automatic speech recognition (ASR) systems. In previous works, latency was optimised by truncating the online attention weights based on the hard alignments obtained from conventional ASR models, without taking into account the potential loss of ASR accuracy. On the contrary, here we present a strategy to obtain the alignments as a part of the model training without external supervision. The alignments produced by the proposed method are dynamically regularised on the training data, such that the latency reduction does not result in the loss of ASR accuracy. SR-MLT is applied as a fine-tuning step on the pre-trained Transformer models that are based on either monotonic chunkwise attention (MoChA) or cumulative attention (CA) algorithms for online decoding. ASR experiments on the AIShell-1 and LibriSpeech datasets show that when applied on a decent pre-trained MoChA or CA baseline model, SR-MLT can effectively reduce the latency with the relative gains ranging from 11.8% to 39.5%. Furthermore, we also demonstrate that under certain accuracy levels, the models trained with SR-MLT can achieve lower latency when compared to those supervised using external hard alignments.

Index Terms: streaming Transformer-based ASR, self-regularised minimum latency training

1. Introduction

The Transformer architecture [1], as an epitome of the attention-based encoder-decoder framework [2, 3], has become one of the dominant end-to-end (E2E) ASR techniques. Compared with other modelling strategies such as connectionist temporal classification (CTC) [4, 5] and recurrent neural network (RNN) transducer [6], Transformer sustains severe latency issues at inference time, since it requires access to the full speech utterance for decoding. To facilitate streaming ASR, a number of online attention mechanisms have been proposed in literature, including monotonic chunkwise attention (MoChA) [7, 8, 9], CTC-triggered attention [10], decoder-end adaptive computation steps (DACS) based algorithms [11, 12], as well as the recent cumulative attention (CA) [13]. Although the aforementioned algorithms enable Transformer to emit ASR outputs in real time, the emission can be delayed from the actual acoustic boundary, because the system always tends to involve some future information to enhance the prediction confidence.

There have been attempts to reduce the latency level of attention-based online decoding, and the most popular method is minimum latency training (MLT) [14]. MLT introduces the "ground-truth" token boundaries into the attention calculation, which can be obtained from the forced alignments produced by conventional hidden Markov model (HMM) based ASR models [15]. When training the E2E system, any attending probabili-

ties beyond the boundary will be ignored, so as to attenuate the chance of triggering later than the real acoustic border.

Besides, the CTC posterior spikes corresponding to the non-blank labels are also regarded as effective attention boundaries used in MLT, given that they sparsely appear around the endpoint of acoustic events [10]. Similar to the HMM-based method, a Viterbi decoding can be performed with the CTC model to generate the alignments that locate the output tokens. As a result, the triggering points produced by the attention model are calibrated as close to the CTC spikes as possible [16].

In MLT, both HMM and CTC alignments impose arbitrary restrictions on the attention computation, where the effect of latency reduction is largely subject to the quality of the alignments themselves, which determine the upper bound of MLT performance. Furthermore, the sacrifice of ASR accuracy in exchange of lower latency is not well considered by MLT, making the trade-off easily biased to the latency side.

Deep reinforcement learning boosted head-synchronous DACS (DRL-HS-DACS) [17], on the contrary, saves acquiring prior knowledge by self-exploring dynamic attention boundaries through its own learning experience. During E2E training, the system keeps track of the halting position along with the ASR prediction correctness observed on the training data, based on which the halting agent learns to seek the earliest triggering point with sufficient acoustic information. As opposed to MLT methods, DRL-HS-DACS could not only breakthrough the lower limit of latency posed by external alignments, but also maintain a balance between the accuracy and latency levels.

Even if the halting positions produced by the DRL-HS-DACS model are generally optimised by the DRL algorithm, the absence of strong constraints (e.g. hard acoustic boundaries) upon the agent training can still lead to the failure of halting for certain decoding steps. Thus, other auxiliary halting strategies such as the maximum look-ahead steps are also adopted to help the system trigger ASR outputs in time.

To overcome the issues encountered by MLT and DRL-HS-DACS while take advantage of their strengths, we propose the self-regularised minimum latency training (SR-MLT) algorithm. On the one hand, hard acoustic boundaries are utilised to cut off the redundant attention weights that cause delayed triggering, eliminating the need of DRL. On the other hand, these boundaries are dynamically obtained from the course of E2E training itself, rather than from external resources like forced alignments, which notably cuts down the training complexity when compared with the MLT methods.

We applied SR-MLT on two streaming Transformer ASR systems that are based on MoChA [7] and CA [13] algorithms respectively, in order to show it compatible to various online attention mechanisms. Our experiments demonstrate that the proposed method can achieve significant latency reduction against the baseline systems, including the vanilla MLT systems that leverage HMM produced acoustic boundaries.

2. Streaming Transformer-based ASR

2.1. Model architecture

The Transformer ASR system is composed of three parts, as front-end, encoder and decoder. The front-end is usually implemented by convolutional neural networks (CNNs), which enhance the feature extraction and conduct frame-rate reduction. A stack of identical layers construct the Transformer encoder, with each consisting of a self-attention module and a pointwise feed-forward network (FFN). The decoder takes a similar architecture, but having an additional cross-attention module at each layer, where the speech-to-text alignment is generated.

Transformer adopts the scaled dot-product attention mechanism to capture the dependency between sequence elements:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V}, \quad (1)$$

where $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ denote the projected feature representations within either the encoder or decoder attention modules, given d_k as the dimension of projection. Moreover, multiple attention heads are utilised by Transformer, aiming to obtain information from distinct feature spaces:

$$\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Concat}(\text{head}_1, \dots, \text{head}_H)\mathbf{W}^O, \quad (2)$$

$$\text{and head}_h = \text{Attention}(\mathbf{Q}\mathbf{W}_h^Q, \mathbf{K}\mathbf{W}_h^K, \mathbf{V}\mathbf{W}_h^V), \quad (3)$$

where $\mathbf{W}_h^{Q,K,V}$ and \mathbf{W}^O denote the linear projection matrices of the individual heads and the output layer.

The standard Transformer system performs global attention on the entire input sequence, which precludes the streaming inference of ASR. At the encoder side, the solution can be segmenting the speech utterance into a stream of chunks with fixed [18] or adaptive [19, 20] size, so that the attention computation is restricted to local contexts. As for the decoder side is concerned, a few online decoding algorithms have been proposed, and here we present two of them as follows.

2.2. Monotonic chunkwise attention (MoChA)

Rather than working on the full speech utterance, MoChA first estimates the probability of stopping the decoding at each frame, which is monotonically computed along the timestep j , as:

$$p_{i,j} = \text{Sigmoid}(e_{i,j}), \quad (4)$$

$$e_{i,j} = \frac{\mathbf{q}_i(\mathbf{k}_{i,j})^T}{\sqrt{d_k}}, \quad (5)$$

where \mathbf{q} and \mathbf{k} denote the decoder and encoder states seen at the decoding step i , respectively. The sigmoid function here serves as a substitute to the softmax function in eq. (1), performing instant normalisation. To facilitate the system training, an expectation form of the triggering probability is given as:

$$\alpha_{i,j} = p_{i,j} \left((1 - p_{i,j-1}) \frac{\alpha_{i,j-1}}{p_{i,j-1}} + \alpha_{i-1,j} \right), \quad (6)$$

which can be recursively calculated over the decoding steps.

As the name suggests, MoChA also conducts a second pass soft attention upon the chunks of encoder states, in order to add flexibility to the speech-to-text alignment:

$$\beta_{i,j} = \sum_{k=j}^{j+w-1} \left(\alpha_{i,k} \exp(u_{i,j}) / \sum_{l=k-w+1}^k \exp(u_{i,l}) \right), \quad (7)$$

given w as the width of chunks, and $u_{i,j}$ the chunkwise attention energy that is computed as in eq. (5).

Finally, a context vector is generated as the weighted sum of the encoder states \mathbf{v} :

$$\mathbf{c}_i = \sum_{j=1}^T \beta_{i,j} \cdot \mathbf{v}_j, \quad (8)$$

where T denotes the length of speech.

During test time, the earliest timestep that has $p_{i,j} > 0.5$ will be chosen as the decoding endpoint, and the chunk ending at the corresponding j is used to produce the context vector.

2.3. Cumulative attention (CA)

The CA algorithm has been proposed to mitigate the problem of MoChA that the decoding is halted only based on the acoustic feature of single frames. Instead, the halting decisions in CA are made upon all history encoder states.

Contrary to MoChA, CA works out the interim context vectors in the first place, following the autoregressive manner:

$$\mathbf{c}_{i,j} = \mathbf{c}_{i,j-1} + a_{i,j} \cdot \mathbf{v}_{i,j}, \quad (9)$$

where $a_{i,j}$ is computed in the same way as eq. (4), but rather interpreted as the monotonic attention weight. It is understood that $\mathbf{c}_{i,j}$ contains all the relevant acoustics information accumulated at each timestep. Next, a trainable device called halting selector is introduced to take $\mathbf{c}_{i,j}$ as input, from which a halting probability is calculated:

$$p_{i,j} = \text{Sigmoid}(\text{HaltSelect}(\mathbf{c}_{i,j}) + r + \epsilon), \quad (10)$$

where r is a bias term and ϵ denotes the additive Gaussian noise applied during training. The halting selector is usually implemented as a simple one-layer deep neural network (DNN) to speed up the inference. Similar to MoChA, an expected halting probability $\alpha_{i,j}$ is computed through eq. (6) to enable the update of halting selector parameters.

As a result, the expected context vector is produced to cover all the alignment paths as:

$$\mathbf{c}_i = \sum_{j=1}^T \alpha_{i,j} \cdot \mathbf{c}_{i,j}. \quad (11)$$

For inference, the interim context vectors are dynamically updated along the timesteps as in training, and likewise the first one with $p_{i,j} > 0.5$ will be directly picked up as the output of the attention module, which concludes the decoding.

3. Proposed self-regularised minimum latency training

The goal of SR-MLT is to achieve the maximum latency reduction during inference at the minimum cost of ASR accuracy. Thus, we must carefully balance both metrics when designing the method. To accomplish this, SR-MLT is carried out in two stages: pre-training and fine-tuning.

3.1. Pre-training

As no external alignments are leveraged, the SR-MLT method has to obtain the appropriate acoustic boundaries completely on its own. We assume that the triggering points (halting positions) that correspond to accurate ASR predictions can be used

Table 1: Update policy of SR-MLT, with regard to the prediction correctness and triggering point on training set.

acc \ cov	equal	up	down
	equal	False	False
up	True	True	True
down	False	False	False

as good initial boundaries for the later fine-tuning stage. Thus, in the course of pre-training, we train the online Transformer system as normal with MoChA or CA applied, and keep recording (a) the triggering points and (b) the ASR prediction correctness of each decoding step within the training set. Based on these factors, the accuracy and coverage ratio (which reflects latency) of the data are worked out:

$$acc = \frac{\text{number of correctly predicted tokens}}{\text{total number of tokens}}, \quad (12)$$

$$cov = \frac{\text{triggering point index}}{\text{total length of utterance}}. \quad (13)$$

The recorded factors will be updated as soon as a higher accuracy is seen at the current data iteration when compared with the record, regardless of the coverage ratio variation, or otherwise remain unchanged to the previous figures.

Specifically, compared to [17] where the metrics in eq. (12) and (13) are calculated with respect to individual tokens, SR-MLT calculates the metrics upon the whole mini-batch of training data. We believe that the statistics measured on larger granularity, such as mini-batch, can provide a more reliable estimate toward the system performance, so as to help make better update decisions during training. A comparison of ASR performance using different levels of granularity including token, utterance and mini-batch will be given in section 4.3.

Note that the number of pre-training epochs can be neither too small, where the ASR model hasn't matured to reach an adequate accuracy level and provide stable triggering points, nor too large such that the model is over-fitted to the training data. In our experiments, the models are pre-trained up to around 60% of the full training time. In the end, the pre-training stage evaluates the upper bound of the ASR accuracy that is achievable by the streaming system, which is exactly what we aim to hold to when reducing the latency at the next stage.

3.2. Fine-tuning

At the fine-tuning stage, the initial boundaries obtained from the pre-training are applied to restricting the attention computation, similar to [14]:

$$\alpha_{i,j} = \begin{cases} p_{i,j} \left((1 - p_{i,j-1}) \frac{\alpha_{i,j-1}}{p_{i,j-1}} + \alpha_{i-1,j} \right), & j \leq b_i + \delta \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

where b_i is the triggering point in the record, and δ denotes an offset appending to the acoustic boundary which allows some extra frames to be included to the attention, easing the restriction imposed by hard alignments.

Meanwhile, we continue to record the prediction correctness and triggering point on the training set, but the update of these two factors follows the policy presented in Table. 1. It can be seen that the accuracy still dictates the update decision. The records are made up-to-date whenever the accuracy goes up on the same mini-batch of data, or left untouched if it declines.

Table 2: Character error rates (CERs %) on AIShell-1.

Model	dev	test
Offline [23]	-	6.7
MMA-MoChA [9]	-	7.5
BS-DEC [24]	6.4	7.3
MoChA [13]	6.4	7.2
SR-MLT MoChA	6.3	7.2
CA [13]	6.3	7.0
SR-MLT CA	6.4	7.1

Table 3: Word error rates (WERs %) on Librispeech.

Model	dev		test	
	clean	other	clean	other
Offline [13]	2.4	6.0	2.6	6.1
Triggered attention [10]	-	-	2.8	7.2
BS-DEC [24]	2.5	6.8	2.7	7.1
MoChA [13]	2.8	7.1	3.1	7.4
SR-MLT MoChA	2.8	7.4	3.1	7.5
CA [13]	2.5	6.7	2.7	6.8
SR-MLT CA	2.6	6.8	2.7	7.1

As for the equal accuracy is concerned, the update is only performed when the coverage ratio shows a downward tendency, which indicates the reduction of latency label.

Once being updated, The new triggering point will be immediately applied to eq. (14) as b_i for the following training epochs. By doing this, the acoustic boundaries could be gradually pushed to the lower limit as the training proceeds. However, in normal MLT, the streaming system may have rare opportunities to produce triggering points that are substantially earlier than the arbitrary HMM alignments.

At the same time, SR-MLT doesn't lose the potential of achieving equally good ASR accuracy as before the fine-tuning, because the inappropriate triggering points that lead to decayed accuracy are precluded into becoming the valid acoustic boundaries. Hence, we could say that the ASR performance is self-regularised through such a balancing strategy.

4. Experiments

4.1. Experimental setup

The proposed method has been evaluated on two open-sourced datasets: AIShell-1 [21] and Librispeech [22]. We adopt precisely the same data preparation, model architecture, and external language model as in [13] for fair comparisons. The number of pre-training and fine-tuning epochs, as well as the offset δ (after frame-rate reduction by front-end) are set to $\{30, 20, 2\}$ for AIShell-1, and $\{70, 50, 8\}$ for Librispeech. When recording the triggering point for each decoding step, since each timestep is assigned with a triggering (halting) probability $\alpha_{i,j}$ (see eq. (6)) rather than a binary decision. we put the timestep with the largest $\alpha_{i,j}$ as the one into practice. Besides, to enhance the update rate in the fine-tuning stage, any variations of accuracy and coverage ratio that are less than 0.1% are deemed as 'equal', given that the exact equality of the metrics can be very seldom seen on the larger mini-batch of speech utterances.

4.2. Experimental results

Table. 2 and 3 present the character-error-rate (CER) and the word-error-rate (WER) achieved by the proposed method on AIShell-1 and Librispeech datasets, respectively. From the re-

Table 4: Latency (frames) on AIShell-1.

Model	dev	test
Offline [13]	232.8	257.0
MoChA [13]	90.1	92.5
SR-MLT MoChA	56.3	56.0
CA [13]	52.8	51.8
SR-MLT CA	47.9	45.7

Table 5: Latency (frames) on Librispeech.

Model	dev		test	
	clean	other	clean	other
Offline [13]	497.9	440.7	528.5	463.8
MoChA [13]	295.2	256.4	303.1	271.2
SR-MLT MoChA	295.0	256.5	303.3	270.5
CA [13]	68.1	63.5	68.5	65.5
SR-MLT CA	51.3	50.0	50.6	50.6

sults, we could find that SR-MLT has successfully maintained the accuracy level for the streaming ASR systems. On AIShell-1, the CERs of both SR-MLT MoChA and SR-MLT CA models have been preserved very closely to the baseline numbers. As for Librispeech is concerned, even if the accuracy of SR-MLT based models drops on the test-other set when compared to the vanilla MoChA and CA models, they are still kept in the similar range with the reference systems in literature.

The latency level during the inference of SR-MLT models are presented in Table. 4 and 5 for the two datasets, in comparison with the vanilla MoChA and CA based counterparts. The evaluation of latency is adopted as the same in [13], which measures the difference between the triggering point of the E2E model \hat{b}_i^k and the ground-truth token boundary b_i^k :

$$\Delta_{corpus} = \frac{1}{\sum_{k=1}^N |\mathbf{y}^k|} \sum_{k=1}^N \sum_{i=1}^{|\mathbf{y}^k|} (\hat{b}_i^k - b_i^k), \quad (15)$$

where N denotes the total number of utterances in the dataset, and $|\mathbf{y}^k|$ is the number of tokens in each utterance. On the test set of AIShell-1, One can observe that SR-MLT has effectively reduced the latency of both streaming systems, with the relative gains of 39.5% and 11.8% for MoChA and CA respectively. The significant reduction of latency has also been seen by SR-MLT CA on Librispeech, where the relative gains of 26.1% and 22.7% are achieved on the test-clean and test-other sets. However, in contrast with the success on AIShell-1, SR-MLT does not bring obvious latency reductions to the MoChA model on Librispeech. This is due to that some attention heads of MoChA-based Transformer can behave very poorly [13], where SR-MLT is not able to optimise their attention boundaries. Thus, we always need a decent baseline model to perform SR-MLT well at the stage of fine-tuning.

4.3. Analysis

Figure. 1 illustrates the effect of using different granularity to calculate the accuracy and latency statistics on the training set as described in section 3.1. We experiment with SR-MLT CA on AIShell-1 at three levels, as token, utterance and mini-batch. Meanwhile, a range of offset values are manipulated to demonstrate the trade-off between the accuracy and latency metrics. From the figure, it is noticed that in general, using larger granularity could lead to improved accuracy and correspondingly compromised latency performance. Despite the harsh restrictions posed by smaller δ ($\delta = 1$), the mini-batch setting is still




Figure 1: ASR performance of using different update granularity in SR-MLT CA on AIShell-1.




Figure 2: ASR performance of MLT CA and SR-MLT CA on AIShell-1.

able to contain the decay of accuracy under a reasonable level, as opposed to the token and utterance settings.

We further compare the performance of the proposed SR-MLT with the original MLT on top of the CA-base Transformer system. The results achieved on AIShell-1 are presented in Figure. 2. It can be observed that MLT is less sensitive to the change of δ , with regard to both accuracy and latency levels. On the contrary, SR-MLT manages to acquire a better balance of the metrics as δ moderates. Specifically, when $\delta = 2$ and 3, the SR-MLT model obtains lower latency with the same accuracy compared to the MLT model, showing the advantage of dynamic acoustic boundaries over the fixed HMM alignments.

5. Conclusions

In this work, we present a novel minimum latency training method, called SR-MLT, for streaming Transformer-based ASR. Compared to the original MLT method that leverages external forced alignments, SR-MLT produces dynamic acoustic boundaries fully from the model training, in order to restrict the frames that are involved to the online attention computation. The proposed method is conducted in two stages. At the pre-training stage, the optimal triggering points along with the ASR prediction correctness are recorded for all the training data. Then at the fine-tuning stage, the triggering points in the records are utilised as initial acoustic boundaries for the decoding steps. Meanwhile, as the training proceeds, the records are updated once reduced latency is observed on the training set with maintained accuracy. We applied SR-MLT onto two streaming ASR algorithms as MoChA and CA. The experimental results show that the proposed method can reduce the latency level for both MoChA and CA based systems, without degrading the ASR accuracy. Besides, we also show that SR-MLT could achieve a better balance between latency and accuracy when compared with normal MLT.

6. References

- [1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, "Attention is all you need," *arXiv preprint arXiv:1706.03762*, 2017.
- [2] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 4960–4964.
- [3] Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio, "Attention-based models for speech recognition," *Advances in neural information processing systems*, vol. 28, 2015.
- [4] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [5] Alex Graves and Navdeep Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *International conference on machine learning*. PMLR, 2014, pp. 1764–1772.
- [6] Alex Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [7] Chung-Cheng Chiu and Colin Raffel, "Monotonic chunkwise attention," *arXiv preprint arXiv:1712.05382*, 2017.
- [8] Emiru Tsunoo, Yosuke Kashiwagi, Toshiyuki Kumakura, and Shinji Watanabe, "Towards online end-to-end Transformer automatic speech recognition," *arXiv preprint arXiv:1910.11871*, 2019.
- [9] Hirofumi Inaguma, Masato Mimura, and Tatsuya Kawahara, "Enhancing monotonic multihead attention for streaming ASR," *arXiv preprint arXiv:2005.09394*, 2020.
- [10] Niko Moritz, Takaaki Hori, and Jonathan Le, "Streaming automatic speech recognition with the Transformer model," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6074–6078.
- [11] Mohan Li, Cătălin Zorilă, and Rama Doddipatla, "Transformer-based online speech recognition with decoder-end adaptive computation steps," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 1–7.
- [12] Mohan Li, Cătălin Zorilă, and Rama Doddipatla, "Head-synchronous decoding for Transformer-based streaming ASR," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5909–5913.
- [13] Mohan Li, Shucong Zhang, Catalin Zorila, and Rama Doddipatla, "Transformer-based streaming ASR with cumulative attention," *arXiv preprint arXiv:2203.05736*, 2022.
- [14] Hirofumi Inaguma, Yashesh Gaur, Liang Lu, Jinyu Li, and Yifan Gong, "Minimum latency training strategies for streaming sequence-to-sequence ASR," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6064–6068.
- [15] Chung-Cheng Chiu, Tara N Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjuli Kannan, Ron J Weiss, Kanishka Rao, Ekaterina Gonina, et al., "State-of-the-art speech recognition with sequence-to-sequence models," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4774–4778.
- [16] Haoran Miao, Gaofeng Cheng, and Pengyuan Zhang, "Low-latency Transformer model for streaming automatic speech recognition," *Electronics Letters*, vol. 58, no. 1, pp. 44–46, 2022.
- [17] Mohan Li and Rama Doddipatla, "Improving HS-DACS based streaming Transformer ASR with deep reinforcement learning," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 154–161.
- [18] Emiru Tsunoo, Yosuke Kashiwagi, Toshiyuki Kumakura, and Shinji Watanabe, "Transformer ASR with contextual block processing," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 427–433.
- [19] Mohan Li, Min Liu, and Hattori Masanori, "End-to-end speech recognition with adaptive computation steps," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6246–6250.
- [20] Linhao Dong and Bo Xu, "Cif: Continuous integrate-and-fire for end-to-end speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6079–6083.
- [21] Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao Zheng, "Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline," in *2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*. IEEE, 2017, pp. 1–5.
- [22] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [23] Shigeki Karita, Nanxin Chen, Tomoki Hayashi, Takaaki Hori, Hirofumi Inaguma, Ziyang Jiang, Masao Someki, Nelson Enrique Yalta Soplín, Ryuichi Yamamoto, Xiaofei Wang, et al., "A comparative study on Transformer vs RNN in speech applications," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 449–456.
- [24] Emiru Tsunoo, Yosuke Kashiwagi, and Shinji Watanabe, "Streaming Transformer ASR with blockwise synchronous beam search," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 22–29.