



Enhancing Monotonic Multihead Attention for Streaming ASR

Hirofumi Inaguma¹, Masato Mimura¹, Tatsuya Kawahara¹

¹Graduate School of Informatics, Kyoto University, Kyoto, Japan

{inaguma, mimura, kawahara}@sap.ist.i.kyoto-u.ac.jp

Abstract

We investigate a monotonic multihead attention (MMA) by extending hard monotonic attention to Transformer-based automatic speech recognition (ASR) for online streaming applications. For streaming inference, all monotonic attention (MA) heads should learn proper alignments because the next token is not generated until all heads detect the corresponding token boundaries. However, we found not all MA heads learn alignments with a naïve implementation. To encourage every head to learn alignments properly, we propose *HeadDrop* regularization by masking out a part of heads stochastically during training. Furthermore, we propose to prune redundant heads to improve consensus among heads for boundary detection and prevent delayed token generation caused by such heads. Chunkwise attention on each MA head is extended to the multihead counterpart. Finally, we propose *head-synchronous* beam search decoding to guarantee stable streaming inference.

Index Terms: Transformer, streaming end-to-end ASR, monotonic multihead attention, beam search decoding

1. Introduction

Recent progress of end-to-end (E2E) automatic speech recognition (ASR) models bridges the gap from the state-of-the-art hybrid systems [1]. To make E2E models applicable to simultaneous interpretations in lecture and meeting domains, online streaming processing is necessary. For E2E models, connectionist temporal classification (CTC) [2] and recurrent neural network transducer (RNN-T) [3] have been dominant approaches and reached a level of real applications [4, 5]. Meanwhile, attention-based encoder-decoder (AED) models [6, 7] have demonstrated the powerful modeling capability in offline tasks [8–10] and a number of streaming models have been investigated for RNN-based models [11–17].

Recently, the Transformer architecture [18], based on self-attention and multihead attention, has shown to outperform the RNN counterparts in various domains [19, 20], and several streaming models have been proposed such as triggered attention [21], continuous-integrate-and-fire (CIF) [22], hard monotonic attention (HMA) [23, 24], and other variants [25]. Triggered attention truncates encoder outputs by using CTC spikes and performs an attention mechanism over all past frames. CIF learns acoustic boundaries explicitly and extracts context vectors from the segmented region. Therefore, these models have adaptive segmentation policies relying on acoustic cues only.

On the other hand, HMA detects token boundaries on the decoder side by using lexical information as well. Thus, it is more flexible for modeling non-monotonic alignments and has been investigated in simultaneous machine translation (MT) [26]. Recently, HMA was extended to the Transformer architecture, named *monotonic multihead attention* (MMA), by replacing each encoder-decoder attention head in the decoder with a monotonic attention (MA) head [27]. Unlike a single MA head used in RNN-based models, each MA head can extract source

contexts with different pace and learn complex alignments between input and output sequences. Concurrently, similar methods have been investigated for Transformer-based streaming ASR [23, 24]. Miao *et al.* [24] simplified the MMA framework by equipping a single MA head with each decoder layer to truncate encoder outputs as in triggered attention and perform attention over all past frames. Tsunoo *et al.* [23] also investigated the MMA framework but resorted to using all past frames to obtain a decent performance. However, looking back to the beginning of input frames lessens the advantage of linear-time decoding with HMA as the input length gets longer.

In this work, we investigate the MMA framework using restricted input context for the streaming ASR task. To perform streaming recognition with the MMA framework, it is necessary for every MA head to learn alignments properly. This is because the next token is not generated until all heads detect the corresponding token boundaries. If some heads fail to detect the boundaries until seeing the encoder output of the final frame, the next token generation is delayed accordingly. However, with a naïve implementation, we found that proper monotonic alignments are learnt in dominant MA heads only. To prevent this, we propose *HeadDrop*, in which a part of heads are entirely masked out at random as a regularization during training to encourage the rest non-masked heads to learn alignments properly. Moreover, we propose to prune redundant MA heads in lower decoder layers to further improve consensus among heads on token boundary detection. Chunkwise attention [17] on top of each MA head is further extended to the multihead counterpart to extract useful representations and compensate the limited context size. Finally, we propose *head-synchronous* beam search decoding to guarantee streamable inference.

Experimental evaluations on Librispeech corpus show that our proposed methods effectively encourage MA heads to learn alignments properly, which leads to improvement of ASR performance. Our optimal model enables stable streaming inference on other corpora as well without architecture modification.

2. Transformer ASR architecture

In this section, we detail the Transformer base architecture used in this paper. Our Transformer architecture consists of stacked E encoder layers followed by C front-end CNN blocks, and D decoder layers [19]. A CNN block has two CNN layers with a 3×3 filter followed by a ReLU activation with a channel size 32. The frame rate is reduced by a factor of 2^C by a max-pooling layer with a stride of 2×2 after every block. Each encoder layer is composed of a self-attention (SAN) sub-layer followed by a position-wise feed-forward network (FFN) sub-layer, wrapped by residual connections and layer normalization [28]. A key component of SAN sub-layers is a multihead attention (MHA) mechanism, in which key, value, and query matrices are split into H portions with a dimension $d_k = d_{\text{model}}/H$ after linear transformations and each head performs a scaled-dot attention mechanism: $\text{Attention}(Q, K, V) = \text{softmax}(QK^T/\sqrt{d_k})V$, where K , Q , and V represent key,

query, and value matrices on each head, respectively. The outputs from all heads are concatenated along the feature dimension followed by a linear transformation. A FFN sub-layer is composed of two linear layers with the inner dimension d_{ff} , interleaved with a ReLU activation between them.

In each decoder layer, unlike the encoder layer, additional encoder-decoder attention sub-layer is inserted between SAN and FFN sub-layers, and causal masking is performed to prevent the decoder from seeing the future tokens. We adopt three 1D-convolutional layers for positional embeddings [29]. The entire network is optimized by minimizing the negative log-likelihood and CTC loss with an interpolation weight $\lambda_{ctc} = 0.3$ [19].

3. Monotonic multihead attention (MMA)

In this section, we review hard monotonic attention (HMA) [16], monotonic chunkwise attention (MoChA) [17], and monotonic multihead attention (MMA) [27] as an extension of them.

3.1. Hard monotonic attention (HMA)

HMA was originally proposed for online linear-time decoding with RNN-based AED models. At output step i , the decoder scans encoder outputs from left to right and stops at an index $j = t_i$ (token boundary) to attend the corresponding single encoder output h_j . The decoder has options to stop at the current index or move forward to the next index. The next boundary is determined by resuming scanning from the previous boundary. As hard attention is not differentiable, the expected alignments $\alpha_{i,j}$ are calculated by marginalizing over all possible paths during training as follows:

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

$$p_{i,j} = \text{Sigmoid}(\text{MonotonicEnergy}(h_j, s_i)) \quad (2)$$

where $p_{i,j}$ is a selection probability and a monotonic energy function $\text{MonotonicEnergy}(\cdot)$ takes the i -th decoder state s_i and j -th encoder output h_j as inputs. Whenever $p_{i,j} \geq 0.5$ is satisfied at test time, $\alpha_{i,j}$ is activated (i.e., set to 1.0).

3.2. Monotonic chunkwise attention (MoChA)

To relax strict input-output alignment by using the surrounding contexts, MoChA introduces additional soft attention mechanism on top of HMA. Given the boundary j , chunkwise attention is performed over a fixed window of width w from there:

$$\beta_{i,j} = \sum_{k=j}^{j+w-1} \left(\alpha_{i,k} \exp(u_{i,j}) / \sum_{l=k-w+1}^k \exp(u_{i,l}) \right) \quad (3)$$

where $u_{i,j}$ is a chunk energy parameterized similar to the monotonic energy in Eq. (2) using separate parameters. $\alpha_{i,j}$ in Eq. (3) is a continuous value during training, but is a binary value according to $p_{i,j}$ at test time.

3.3. Monotonic multihead attention (MMA)

To keep the expressive power of Transformer with the multihead attention mechanism while enabling online linear-time decoding, MMA was proposed as an extension of HMA [27]. Each encoder-decoder attention head in the decoder is replaced with a monotonic attention (MA) head in Eq. (1) by defining the monotonic energy function in Eq. (2) as follows:

$$\text{MonotonicEnergy}(h_j, s_i) = \frac{W_s s_i (W_h h_j)^T}{\sqrt{d_k}} + r \quad (4)$$

where W_s and W_h are parameter matrices, and r is a learnable offset parameter (initialized with -2 in this work). Unlike a case of a single MA head in Section 3.1, each MA head can attend to input speech with different pace because its decision process regarding timing to activate $\alpha_{i,j}$ does not influence each

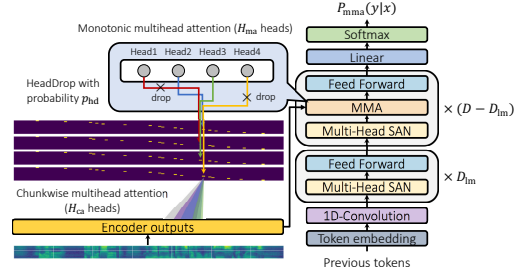


Figure 1: System overview. Residual connections and layer normalization are omitted.

other at each output step. The side effect is that all heads must be activated to generate a token. Otherwise, some MA heads continue to scan encoder outputs until the last time index, which leads to significant increase of latency during inference.

Furthermore, unlike previous works [23, 27] having a single chunkwise attention (CA) head on each MA head, we extend it to the multi-head version having H_{ca} heads per MA head to extract useful representations with multiple views from each boundary (*chunkwise multihead attention*). Assuming each decoder layer has H_{ma} MA heads, the total number of CA heads is $H_{ma} \cdot H_{ca}$ at the layer. However, we found that sharing parameters of CA heads among MA heads in the same layer is effective in our pilot experiments and adopt this strategy in all experiments. The chunk energy $u_{i,j}$ for each CA head is designed as in Eq. (4) without r .

4. Enhancing monotonic alignments

In the Transformer models, there exist many attention heads and residual connections, so it is unclear that all heads contributes to the final predictions. Michel *et al.* [30] reported that most heads can be pruned at test time without significant performance degradation in standard MT and BERT [31] architectures. They also revealed that important heads are determined in early training stage. Concurrently, Voita *et al.* [32] also reported the similar observations by automatically pruning a part of heads with L_0 penalty [33]. In our preliminary experiments, we also observed that not all MA heads learn alignments properly in the MMA-based ASR models and monotonic alignments are learnt only by dominant heads in upper decoder layers. Since $\alpha_{i,j}$ in Eq. (1) are not normalized over inputs so as to sum up to one during training, context vectors from heads which do not learn alignments are more likely to become zero vectors at test time. This is a severe problem because (1) it leads to mismatch between training and testing conditions, and (2) the subsequent tokens cannot be generated until all heads are activated. To alleviate this problem, we propose the following methods.

4.1. HeadDrop

We first propose a regularization method to encourage each MA head to equally contribute to the target task. During training, we stochastically zero out all elements in each MA head (i.e., $\alpha_{i,j} = 0, 1 \leq j \leq |h_i|$) with a fixed probability p_{hd} to force the other heads to learn alignments. The decision of dropping each head is independent of other heads regardless of the depth of the decoder. The output of a MMA function is normalized by dividing it by H_{ma}^+ / H_{ma} , where H_{ma}^+ is the number of non-masked MA heads. We name this *HeadDrop*, inspired by dropout [34] and DropConnect [35].

4.2. Pruning monotonic attention heads in lower layers

Although HeadDrop is effective for improving the contribution of each MA head, we found that some MA heads in the lower decoder layers still do not learn alignments properly. Therefore, we propose to prune such redundant heads because they are

harmful for streaming decoding. We remove the MMA function in the first D_{lm} decoder layers from the bottom ($0 \leq D_{lm} \leq D - 1$) during both training and test time (see Figure 1). These layers have SAN and FFN sub-layers only and serve as a pseudo language model (LM). The total number of effective MA heads H_{ma}^{total} is $(D - D_{lm}) \cdot H_{ma}$. This method is also based on findings in [32] that the lower layers of the Transformer decoder are mostly responsible for language modeling. Another advantage of pruning redundant heads is that inference speed is improved, which is effective for streaming ASR.

5. Head-synchronous beam search decoding

During beam search in the MMA framework, failure of boundary detection in some MA heads in some beam candidates at an output step easily prevents the decoder from continuing streaming inference. This is because other candidates must wait for the hypothesis pruning until all heads in all candidates are activated at each output step. To continue streaming inference, we propose *head-synchronous* beam search decoding (Algorithm 1). The idea is to force non-activated heads to activate after a small fixed delay. If a head at the l -th layer cannot detect a boundary for ϵ_{wait} frames after the leftmost boundary detected by other heads in the same layer, the boundary of such a head is set to the rightmost boundary t_{tail}^l among already detected boundaries at the current output step (line 14). Therefore, latency between the fastest (rightmost) and slowest (leftmost) boundary positions in the same layer is less than ϵ_{wait} frames. We note that the decisions of boundary detection at the l -th layer are dependent on outputs from the $(l - 1)$ -th layer, and at least one head must be activated at each layer to generate a token. For the efficient implementation, we search boundaries of all heads in a layer in parallel, thus the loop in line 7 can be ignored. Moreover, we perform batch processing over multiple hypotheses in the beam and cache previous decoder states for efficiency. Note that head synchronization is not performed during training to maintain the divergence of boundary positions. Thus, synchronization can have the ensemble effect for boundary detection.

6. Experimental evaluations

6.1. Experimental setup

We used the 960-hour Librispeech dataset [36] for experimental evaluations. We extracted 80-channel log-mel filterbank coefficients computed with a 25-ms window size shifted every 10 ms using Kaldi [37]. We used a 10k vocabulary based on the Byte Pair Encoding (BPE) algorithm [38]. For the Transformer model configurations, we used $(d_{model}, d_{ff}, H, H_{ma}, H_{ca}, w, C, E, D) = (256, 2048, 4, 4, 1, 4, 3, 12, 6)$ for the baseline MMA models. Adam [39] optimizer was used for training with Noam learning rate schedule [18]. Warmup steps and a learning rate constant were set to 25000 and 5.0, respectively. We averaged model parameters at the last 10 epochs for evaluation. Both dropout and label smoothing [40] were used with a probability 0.1. We set p_{hd} to 0.5. We used a 4-layer LSTM LM with 1024 memory cells for decoding with a beam width of 10. We used decoding hyperparameters $(\alpha, \beta) = (0.5, 2.0)$.¹

6.2. Evaluation measure of boundary detection

To assess consensus among heads for token boundary detection, we propose metrics to evaluate (1) how well each MA head learns alignments (*boundary coverage*) and (2) how often the model satisfies the streamable condition (*streamability*). This is because even if better word error rate (WER) is obtained, the model cannot continue streaming inference if some heads do not learn alignments well, whose evaluation is missing in [23, 24].

¹Code: https://github.com/hirofumi0810/neural_sp.

Algorithm 1 Head-synchronous beam search decoding

Input: h : encoder outputs, ϵ_{wait} : wait time threshold, B : beam width
Output: Ω_{end} : top-k hypotheses

```

1: Initialize:  $t_0^{l,h} \leftarrow 1, \Omega \leftarrow \{\}, \Omega_{end} \leftarrow \{\}, L_{max} \leftarrow 200$ 
2: for  $i \leftarrow 1$  to  $L_{max}$  do
3:    $\Omega_{new} \leftarrow \{\}$ 
4:   for  $k \leftarrow 1$  to  $|\Omega|$  do ▷ Batchfy
5:     for  $l \leftarrow D_{lm}+1$  to  $D$  do
6:        $t_{tail}^l \leftarrow 1$ 
7:       for  $h \leftarrow 1$  to  $H_{ma}$  in parallel do
8:          $p_{i,j}^{l,h} = \text{Sigmoid}(\text{MonotonicEnergy}(h_j, s_{i-1}^{l-1}))$ 
9:         for  $j \leftarrow t_{i-1}^{l,h}$  to  $|h|$  do
10:          if  $p_{i,j}^{l,h} \geq 0.5$  then
11:             $t_i^{l,h} \leftarrow j; t_{tail}^l \leftarrow \max(t_{tail}^l, j);$  break;
12:          else
13:            if  $j \geq t_{tail}^l + \epsilon_{wait}$  then ▷ Forced activation
14:               $t_i^{l,h} \leftarrow t_{tail}^l$ 
15:              break
16:            end if
17:          end if
18:        end for
19:      end for
20:    end for
21:    Append topK( $\log P_{mma} + \alpha \log P_{lm} + \beta i$ ) to  $\Omega_{new}$ 
22:  end for
23:   $\Omega, \Omega_{end} \leftarrow \text{Prune}(\Omega_{new}, \Omega_{end})$ 
24:  if  $i = L_{max}$  or  $|\Omega_{end}| = B$  then break; end if
25: end for
26: return  $\Omega_{end}$ 

```

6.2.1. Boundary coverage

During beam search, we count the total number of boundaries (j such that $\alpha_{i,j} = 1$) up to the i -th output step averaged over all MA heads, $Q_i^{n,k}$, for every candidate $\tilde{y}^{n,k}$ in the n -th utterance:

$$Q_i^{n,k} = \frac{1}{H_{ma}^{total}} \sum_{h=1}^{H_{ma}^{total}} \sum_{i'=1}^i \sum_{j=1}^{|h|} \alpha_{i',j}^h$$

The boundary coverage R_{cov} is defined as the ratio of $Q_{|\tilde{y}^{n,1}|}^{n,1}$ to the corresponding hypothesis length $|\tilde{y}^{n,1}|$ of the best candidate and averaged over N utterances in the evaluation set as follows:

$$R_{cov} [\%] = \frac{1}{N} \sum_{n=1}^N \frac{Q_{|\tilde{y}^{n,1}|}^{n,1}}{|\tilde{y}^{n,1}|} \times 100$$

6.2.2. Streamability

The streamability R_{str} is defined as the ratio of utterances satisfying a condition where $Q_i^{n,k} = |\tilde{y}^{n,k}|$ over all candidates up to the $|\tilde{y}^{n,1}|$ -th output step (i.e. until generation of the best hypothesis is completed) as follows:

$$R_{str} [\%] = \frac{1}{N} \sum_{n=1}^N \delta_n \times 100$$

$$\delta_n = \begin{cases} 1 & (Q_i^{n,k} = |\tilde{y}^{n,k}|, 1 \leq i^\forall \leq |\tilde{y}^{n,1}|, 1 \leq k^\forall \leq |\Omega_i^n|) \\ 0 & (\text{otherwise}) \end{cases}$$

where δ_n is the delta function and Ω_i^n is a hypothesis set at i -th output step of the n -th utterance. $\delta_n = 0$ indicates that the model failed streaming recognition somewhere in the n -th utterance, i.e., continued scanning the encoder outputs until the last frame. However, we note that it does not mean the model leverages additional context.

6.3. Offline ASR results

HeadDrop and pruning MA heads in lower layers improve WER and streamability Table 1 shows the results for offline MMA models on the Librispeech dev-clean/other sets. "Offline" means the encoder is the offline architecture. Boundary coverage and streamability were averaged over two sets. A naive implementation A1 showed a very poor performance. By

Table 1: Results of offline MMA models on the dev sets with standard beam search decoding. HD: HeadDrop.

ID	D_{1m}	H_{ma}	H_{ma}^{total}	HD	dev-clean / dev-other		
					%WER	R_{cov}	R_{str}
A1	0		24		8.6 / 16.5	67.40	0.0
A2	1		20		7.3 / 16.3	79.02	0.0
A3	2	4	16		4.7 / 12.6	86.07	0.0
A4	3		12		4.5 / 12.8	83.87	0.0
A5	4		8		3.6 / 10.8	93.80	0.9
B1	0		24		3.7 / 11.4	60.59	0.0
B2	1		20		4.0 / 11.9	73.73	0.0
B3	2	4	16	✓	3.9 / 10.8	98.85	3.7
B4	3		12		4.1 / 11.0	99.36	6.4
B5	4		8		4.1 / 11.3	99.50	15.8
C1		1	6	-	4.9 / 11.7	99.38	15.7
C2	0	1	6	✓	3.7 / 10.4	99.86	35.9
C3		2	12	✓	3.5 / 10.7	72.08	0.0

Table 2: Results of offline MMA models on the dev sets with head-synchronous beam search decoding and HeadDrop.

ID	D_{1m}	H_{ma}	w	H_{ca}	dev-clean / dev-other		
					%WER	R_{cov}	R_{str}
B3	2				3.9 / 10.7	99.74	21.6
B4	3	4	4	1	3.9 / 10.6	99.76	25.1
B5	4				3.8 / 11.1	99.84	40.5
D1	2				3.3 / 9.9	99.78	37.4
D2	3	4	16	1	3.7 / 10.8	99.83	36.5
D3	4				3.5 / 10.4	99.93	60.4
E1	2				3.3 / 10.2	99.78	40.6
E2	3			2	3.6 / 10.3	99.87	51.2
E3	4	4	16		3.5 / 10.7	99.92	50.0
E4	2				3.3 / 9.8	99.91	77.9
E5	3			4	3.4 / 9.9	99.91	84.5
E6	4				3.6 / 10.4	99.92	63.2
F1	0	1	16	4	3.5 / 10.5	96.23	40.6

pruning MA heads in lower layers with increasing D_{1m} , WER was significantly reduced, but the boundary coverage was not so high (A5). The proposed HeadDrop also significantly improved WER, and with the increase of D_{1m} , the boundary coverage was drastically improved to almost 100% (B3-B5). We can conclude that MA heads in lower layers are not necessarily important. This is probably because (1) the modalities between input and output sequences are different in the ASR task and (2) hidden states in lower decoder layers tend to represent the previous tokens, thus are not suitable for alignment.

Head-synchronous beam search decoding improves WER and streamability Next, the results with head-synchronous beam search decoding are shown in Table 2. ϵ_{wait} is set to 8 in all models. Head-synchronous decoding improved both boundary coverage and streamability. We found that if a head cannot detect the boundary around the corresponding actual acoustic boundary, it tends to stop around the next acoustic boundary twice to compensate the failure when using a standard beam search. Head-synchronous decoding alleviated this mismatch of boundary positions and led to small WER improvement.

Chunkwise multihead attention is effective Furthermore, we increased the window size w and number of heads H_{ca} in chunkwise attention, both of which further improved WER. With $D_{1m} = 4$, E3 and E6 did not obtain benefits from larger H_{ca} . Increasing w to longer than 16 was not effective.

Multiple MA heads in each layer are necessary We also examined the effect of the number of MA heads H_{ma} in each layer (C1-C3 in Table 1). C1 with only one MA head per layer showed a high boundary coverage and was further improved with HeadDrop (C2). C3 with two heads per layer degraded streamability very much. Although C2 showed better performances than B*, it did not obtain much gains from larger w and H_{ma} (F1 in Table 2). This confirms that having multiple MA heads in upper layers is more effective than simply reducing the number of MA heads per layer. In other words, the place of MA

Table 3: Results of streaming MMA models on the Librispeech, TEDLIUM2, and AISHELL-1 test sets

	Model	%WER		%CER
		Librispeech clean	TED LIUM2 other	
Offline	Transformer (ours)	3.3	9.1	10.1
	+ data augmentation	2.8	7.0	-
	++ large model	2.5	6.1	-
	MMA (E5)	3.4	9.9	10.5
Streaming	Triggerred attention [41]	2.8	7.2	-
	CIF [22]	3.3	9.7	-
	MoChA [42]	4.0	9.5	11.3
	MMA ($w = \infty$) [23]	-	-	-
	MMA (narrow chunk)	3.5	11.1	11.0
	MMA (wide chunk)	3.3	10.5	10.2
	+ data augmentation	3.0	8.5	-
	++ large model	2.7	7.1	-

heads is more important than the total number of them.

Here, what does the rest 15.5% for streamability in E5 account for? We found that the last few tokens corresponding to the tail part of input frames were predicted after head pointers on upper layers reached the last encoder output. For these 15.5% utterances, E5 was able to continue streaming decoding until 76.9% of the input frames on average. Since the tail part is mostly silence, this does not affect streaming recognition. In our manual analysis, we observed that MA heads in the same layer move forward with the similar pace, and the pace gets faster in upper layers.² This is because decoder states are dependent on the output from the lower layer. Considering the balance between WER and streamability performance, we will use the E5 setting for streaming experiments in the next section.

6.4. Streaming ASR results

Finally, we present the results of streaming MMA models for the Librispeech test sets in Table 3. We also included results on TEDLIUM2 [43] and AISHELL-1 [44] to confirm whether the optimal configuration tuned on Librispeech can work in other corpora as well. We adopted the chunk hopping mechanism [45] for the online encoder. Following [22, 24], we set the left/current (hop)/right chunk sizes to 960/640/320 (narrow chunk) and 640/1280/640 (wide chunk) [ms]. We used speed perturbation [46] and SpecAugment [47] for data augmentation, but speed perturbation was applied by default for TEDLIUM2 and AISHELL-1. For large models, we used (d_{model} , d_{ff} , H) = (768, 3072, 8) and other hyperparameters were kept. Head-synchronous decoding was used for all MMA models. We used CTC scores during inference only for standard Transformer models. Our streaming MMA models achieved comparable results to the state-of-the-art Transformer-based streaming ASR model [41] without looking back to the first input frame. Moreover, our model outperformed the MMA model with $w = \infty$ [23] by a large margin. Increasing the model size was also effective. The streamabilities of the streaming MMA models on TEDLIUM2 and AISHELL-1 with the wide chunk were 80.0%, and 81.5%, respectively. This confirms that the E5 setting generalizes to other corpora without architecture modification.

7. Conclusion

We tackled the alignment issue in monotonic multihead attention (MMA) for online streaming ASR with *HeadDrop* regularization and head pruning in lower decoder layers. We also stabilized streamable inference by *head-synchronous* decoding. Our future work includes investigation of adaptive policies for head pruning and regularization methods to make the most of the MA heads instead of discarding them. Minimum latency training as done in MoChA [48] is another interesting direction.

²Examples available at https://github.com/hirofumi0810.github.io/demo/enhancing_mma_asr.

8. References

- [1] C.-C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen, Z. Chen, A. Kannan, R. J. Weiss, K. Rao, K. Gonina *et al.*, “State-of-the-art speech recognition with sequence-to-sequence models,” in *Proc. of ICASSP*, 2018, pp. 4774–4778.
- [2] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proc. of ICML*, 2006, pp. 369–376.
- [3] A. Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [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. of ICASSP*, 2019, pp. 6381–6385.
- [5] T. N. Sainath, Y. He, B. Li, A. Narayanan, R. Pang, A. Bruguier, S.-y. Chang, W. Li, R. Alvarez, Z. Chen *et al.*, “A streaming on-device end-to-end model surpassing server-side conventional model quality and latency,” in *Proc. of ICASSP*, 2020, pp. 6059–6063.
- [6] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *Proc. of NeurIPS*, 2015, pp. 577–585.
- [7] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proc. of ICASSP*, 2016, pp. 4960–4964.
- [8] R. Prabhavalkar, K. Rao, T. N. Sainath, B. Li, L. Johnson, and N. Jaitly, “A comparison of sequence-to-sequence models for speech recognition,” in *Proc. of Interspeech*, 2017, pp. 939–943.
- [9] E. Battenberg, J. Chen, R. Child, A. Coates, Y. G. Y. Li, H. Liu, S. Satheesh, A. Sriram, and Z. Zhu, “Exploring neural transducers for end-to-end speech recognition,” in *Proc. of ASRU*, 2017, pp. 206–213.
- [10] A. Zeyer, K. Irie, R. Schlüter, and H. Ney, “Improved training of end-to-end attention models for speech recognition,” in *Proc. of Interspeech*, 2018, pp. 7–11.
- [11] J. Hou, S. Zhang, and L.-R. Dai, “Gaussian prediction based attention for online end-to-end speech recognition,” in *Proc. of Interspeech*, 2017, pp. 3692–3696.
- [12] A. Tjandra, S. Sakti, and S. Nakamura, “Local monotonic attention mechanism for end-to-end speech and language processing,” in *Proc. of IJCNLP*, 2017, pp. 431–440.
- [13] D. Lawson, C.-C. Chiu, G. Tucker, C. Raffel, K. Swersky, and N. Jaitly, “Learning hard alignments with variational inference,” in *Proc. of ICASSP*, 2018, pp. 5799–5803.
- [14] M. Li, M. Liu, and H. Masanori, “End-to-end speech recognition with adaptive computation steps,” in *Proc. of ICASSP*, 2019, pp. 6246–6250.
- [15] N. Moritz, T. Hori, and J. Le Roux, “Triggered attention for end-to-end speech recognition,” in *Proc. of ICASSP*, 2019, pp. 5666–5670.
- [16] C. Raffel, M.-T. Luong, P. J. Liu, R. J. Weiss, and D. Eck, “Online and linear-time attention by enforcing monotonic alignments,” in *Proc. of ICML*, 2017, pp. 2837–2846.
- [17] C.-C. Chiu and C. Raffel, “Monotonic chunkwise attention,” in *Proc. of ICLR*, 2018.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proc. of NeurIPS*, 2017, pp. 5998–6008.
- [19] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang *et al.*, “A comparative study on Transformer vs RNN in speech applications,” in *Proc. of ASRU*, 2019, pp. 499–456.
- [20] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, “A comparison of Transformer and LSTM encoder decoder models for ASR,” in *Proc. of ASRU*, 2019, pp. 8–15.
- [21] N. Moritz, T. Hori, and J. L. Roux, “Streaming automatic speech recognition with the Transformer model,” in *Proc. of ICASSP*, 2020, pp. 6074–6078.
- [22] L. Dong and B. Xu, “CIF: Continuous integrate-and-fire for end-to-end speech recognition,” in *Proc. of ICASSP*, 2020, pp. 6079–6083.
- [23] E. Tsunoo, Y. Kashiwagi, T. Kumakura, and S. Watanabe, “Towards online end-to-end Transformer automatic speech recognition,” *arXiv preprint arXiv:1910.11871*, 2019.
- [24] H. Miao, G. Cheng, C. Gao, P. Zhang, and Y. Yan, “Transformer-based online CTC/attention end-to-end speech recognition architecture,” in *Proc. of ICASSP*, 2020, pp. 6084–6088.
- [25] Z. Tian, J. Yi, Y. Bai, J. Tao, S. Zhang, and Z. Wen, “Synchronous Transformers for end-to-end speech recognition,” in *Proc. of ICASSP*, 2020, pp. 7884–7888.
- [26] N. Arivazhagan, C. Cherry, W. Macherey, C.-C. Chiu, S. Yavuz, R. Pang, W. Li, and C. Raffel, “Monotonic infinite lookback attention for simultaneous machine translation,” in *Proc. of ACL*, 2019, pp. 1313–1323.
- [27] X. Ma, J. Pino, J. Cross, L. Puzon, and J. Gu, “Monotonic multi-head attention,” in *Proc. of ICLR*, 2020.
- [28] J. L. Ba, J. R. Kiros, and G. E. Hinton, “Layer normalization,” *arXiv preprint arXiv:1607.06450*, 2016.
- [29] A. Mohamed, D. Okhonko, and L. Zettlemoyer, “Transformers with convolutional context for ASR,” *arXiv preprint arXiv:1904.11660*, 2019.
- [30] P. Michel, O. Levy, and G. Neubig, “Are sixteen heads really better than one?” in *Proc. of NeurIPS*, 2019, pp. 14 014–14 024.
- [31] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional Transformers for language understanding,” in *Proc. of NAACL-HLT*, 2019, pp. 4171–4186.
- [32] E. Voita, D. Talbot, F. Moiseev, R. Sennrich, and I. Titov, “Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned,” in *Proc. of ACL*, 2019, pp. 5797–5808.
- [33] C. Louizos, M. Welling, and D. P. Kingma, “Learning sparse neural networks through L_0 regularization,” in *Proc. of ICLR*, 2018.
- [34] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 56, pp. 1929–1958, 2014.
- [35] L. Wan, M. Zeiler, S. Zhang, Y. Le Cun, and R. Fergus, “Regularization of neural networks using dropconnect,” in *Proc. of ICML*, 2013, pp. 1058–1066.
- [36] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *Proc. of ICASSP*, 2015, pp. 5206–5210.
- [37] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, “The Kaldi speech recognition toolkit,” in *Proc. of ASRU*, 2011.
- [38] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” in *Proc. of ACL*, 2016, pp. 1715–1725.
- [39] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [40] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proc. of CVPR*, 2016, pp. 2818–2826.
- [41] N. Moritz, T. Hori, and J. L. Roux, “Streaming automatic speech recognition with the Transformer model,” *arXiv preprint arXiv:2001.02674*, 2020.
- [42] H. Inaguma, M. Mimura, and T. Kawahara, “CTC-synchronous training for monotonic attention model,” *arXiv preprint arXiv:2005.04712*, 2020.
- [43] A. Rousseau, P. Deléglise, and Y. Estève, “TED-LIUM: An automatic speech recognition dedicated corpus,” in *Proc. of LREC*, 2012, pp. 125–129.
- [44] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, “AISHELL-1: An open-source mandarin speech corpus and a speech recognition baseline,” in *Proc. of O-COCOSDA*, 2017, pp. 1–5.
- [45] 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. of ICASSP*, 2019, pp. 5656–5660.
- [46] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, “Audio augmentation for speech recognition,” in *Proc. of Interspeech*, 2015, pp. 3586–3589.
- [47] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proc. of Interspeech*, 2019.
- [48] H. Inaguma, Y. Gaur, L. Lu, J. Li, , and Y. Gong, “Minimum latency training strategies for streaming sequence-to-sequence ASR,” in *Proc. of ICASSP*, 2020, pp. 6064–6068.