



CTC-synchronous Training for Monotonic Attention Model

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

Monotonic chunkwise attention (MoChA) has been studied for the online streaming automatic speech recognition (ASR) based on a sequence-to-sequence framework. In contrast to connectionist temporal classification (CTC), backward probabilities cannot be leveraged in the alignment marginalization process during training due to left-to-right dependency in the decoder. This results in the error propagation of alignments to subsequent token generation. To address this problem, we propose CTC-synchronous training (CTC-ST), in which MoChA uses CTC alignments to learn optimal monotonic alignments. Reference CTC alignments are extracted from a CTC branch sharing the same encoder with the decoder. The entire model is jointly optimized so that the expected boundaries from MoChA are synchronized with the alignments. Experimental evaluations of the TEDLIUM release-2 and Librispeech corpora show that the proposed method significantly improves recognition, especially for long utterances. We also show that CTC-ST can bring out the full potential of SpecAugment for MoChA.

Index Terms: Streaming sequence-to-sequence ASR, connectionist temporal classification, monotonic chunkwise attention

1. Introduction

Streaming automatic speech recognition (ASR) is a core technology used in simultaneous interpretation such as live captioning, simultaneous translation, and dialogue systems. Recently, the performance of end-to-end (E2E) ASR systems has been nearing that of hybrid systems with much more simplified architectures [1, 2]. Therefore, building effective streaming E2E-ASR systems is an important step towards making E2E models applicable in the real world.

The streaming E2E models proposed thus far can be categorized as *time-synchronous* or *label-synchronous* models. Time-synchronous models include connectionist temporal classification (CTC) [3], recurrent neural network transducer (RNN-T) [4], and recurrent neural aligner (RNA) [5]. These models can make predictions from the input stream frame-by-frame, so they can easily satisfy the demands of streaming ASR. Label-synchronous models based on an attention-based sequence-to-sequence (S2S) framework [6] outperform time-synchronous models on several benchmarks in offline scenarios [7–9]. However, label-synchronous models are not suitable for streaming ASR because the initial token cannot be generated until all of the speech frames in an utterance are encoded. To make label-synchronous models streamable, several variants have been studied, such as hard monotonic attention [10], monotonic chunkwise attention (MoChA) [11], triggered attention [12], adaptive computation steps [13], and continuous integrate-and-fire [14]. In this work, we focus on MoChA as a streaming attention-based S2S model since it can be trained efficiently and shows promising results for various ASR tasks [15–19].

Hard monotonic attention [10] enables online linear-time decoding during inference by introducing a discrete binary vari-




Figure 1: Visualization of decision boundaries (yellow dots) of LC-BLSTM-40+40 - MoChA (top, T5) and proposed model w/ CTC-ST (bottom, T6)

able. To train the model using standard backpropagation, the alignment probabilities are marginalized over all possible paths during training. MoChA is an extension of hard monotonic attention and introduces additional soft attention over a small chunk. This lessens the strict monotonic constraints of input-output alignments and leads to more accurate recognition [11]. However, the results of training MoChA have often been reported to be unstable [16–19]. This is because the attention scores are not globally normalized over all encoder memories, which results in poorly scaled attention scores in the early training stage. Consequently, the gap in attention behaviors widens between training and test time. A possible solution to this problem is to regularize the model so that the total attention scores over all output timesteps sum up to the output sequence length. This was originally proposed in [14, 18] and referred to as *quantity loss*. Although [18] reported that quantity loss is not effective for MoChA for large-scale data (3.4k hours), we have found that it is very effective when the training data size is small (< 1k hours). However, apart from the scaling issue, alignments in MoChA are susceptible to noise from previous timesteps because they are calculated only with the forward algorithm due to left-to-right dependency in the autoregressive decoder. This is because the backward algorithm cannot be used to estimate accurate alignments, unlike in CTC. This leads to significant error propagation of the predicted token boundary positions to subsequent token generation, especially for long utterances.

In this work, we propose *CTC-synchronous training* (CTC-ST) to learn reliable monotonic alignments by using CTC alignments. Since CTC is optimized with the forward-backward algorithm, the resulting alignments are more reliable than those from MoChA. Moreover, the CTC posterior probabilities tend to peak in spikes [3], so they can be regarded as decision boundaries for token generation and a good reference for MoChA to learn the appropriate timing to generate output tokens. We extract the reference boundaries from the CTC model with forced alignment, and the expected decision boundaries from MoChA

are optimized to be close to the corresponding CTC boundaries. Since we jointly optimize the MoChA and CTC branches by sharing the same encoder, the entire model can be trained in an end-to-end fashion. We also propose a curriculum training strategy to stabilize the training.

We experimentally evaluated the TEDLIUM release-2 and Librispeech corpora to demonstrate that the proposed CTC-ST significantly improves recognition. We also investigate reliability of CTC alignments in the proposed CTC-ST by combining it with SpecAugment [20].

2. Related work

CTC alignments are used in triggered attention for streaming E2E-ASR [12,21,22], in which encoder memories are truncated with every CTC spike and the global attention mechanism is applied subsequently. Our work differs in that (1) we do not use CTC spikes in the test phase, and (2) the decoding complexity of MoChA is linear while that of the triggered attention is quadratic but streamable.

Regarding manipulating boundary positions in hard monotonic attention, [18] uses the external framewise alignments extracted from the hybrid ASR system to reduce latency for token generation in MoChA. In contrast, our proposed CTC-ST does not require any external framewise alignments and is designed to improve recognition.

3. Streaming attention-based S2S

3.1. Latency-controlled BLSTM (LC-BLSTM)

In order to enable low-latency feature encoding while maintaining bidirectionality, latency-controlled bidirectional long short-term memory (LC-BLSTM) encoder has been introduced to restrict future frames to small chunks (N_r frames) [16,17,23,24]. LC-BLSTM consists of forward and backward LSTMs similar to the BLSTM but processes a small chunk by sliding a window of size N_c without overlap between adjacent chunks. The backward LSTM processes later N_r frames as future contexts in addition to N_c frames in the current chunk. Thus, the total latency at each chunk is $(N_c + N_r)$ frames. The backward LSTM state is reset at every chunk while the previous state of the forward LSTM is carried over to the next chunk as the initial state. We refer to this as LC-BLSTM- N_c+N_r in the later experiments.

3.2. Monotonic chunkwise attention (MoChA)

To enable online linear-time decoding for attention-based S2S models during inference, hard monotonic attention was proposed by introducing discrete binary decision processes [10]. Unlike the global attention mechanism [6], all attention scores are assigned to a single memory at each output timestep i without global score normalization. As hard attention is not differentiable, the expected alignment probabilities are marginalized over all possible paths during training by introducing a selection probability $p_{i,j} \in [0, 1]$. $p_{i,j}$ is a function of the monotonic energy activation $e_{i,j}$, which is parameterized with the i -th decoder state s_i and the j -th encoder state h_j as follows:

$$\begin{aligned} \alpha_{i,j} &= p_{i,j} \sum_{k=1}^j \left(\alpha_{i-1,k} \prod_{l=k}^{j-1} (1 - p_{i,l}) \right) \\ &= p_{i,j} \left((1 - p_{i,j-1}) \frac{\alpha_{i,j-1}}{p_{i,j-1}} + \alpha_{i-1,j} \right) \end{aligned} \quad (1)$$

$$e_{i,j} = g \frac{v^\top}{\|v\|} \text{ReLU}(\mathbf{W}_h h_j + \mathbf{W}_s s_i + b) + r \quad (2)$$

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

where $g, v, \mathbf{W}_h, \mathbf{W}_s, b$, and r are learnable parameters.

Monotonic chunkwise attention (MoChA) [11] is an extension of the above method which introduces additional soft chunkwise attention to loosen the strict input-output alignment in hard attention. The chunkwise attention scores $\beta_{i,j}$ over the small w frames are calculated from each expected boundary j :

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

where $u_{i,j}$ is the chunk energy activation formulated similar to $e_{i,j}$ in Eq. (2) without weight normalization and the offset r . The i -th context vector is calculated as a weighted sum of encoder memories by $\beta_{i,j}$ and the subsequent token generation processes are the same as in the global attention. Both $\alpha_{i,j}$ and $\beta_{i,j}$ can be calculated in parallel with the cumulative sum/product and the moving sum operations.

At the time of the test, each token is generated once $p_{i,j}$ exceeds a threshold of 0.5. The next token boundary is determined further to the right than the current boundary. For more details on the decoding algorithm, refer to [10, 11].

3.3. Quantity regularization

Since MoChA does not normalize monotonic attention scores $\alpha_{i,j}$ across all encoder memories $\{h_j\}_{j=1}^T$, there is no guarantee that $\sum_{j=1}^T \alpha_{i,j} = 1$ is satisfied during training. This results in poorly scaled selection probabilities $p_{i,j}$ and widens the gap in behaviors between training and testing because $\alpha_{i,j}$ can attenuate quickly during marginalization. To address this drawback, we use a simple regularization term \mathcal{L}_{qua} , i.e., *quantity loss* [14, 18], to encourage the total number of the expected boundaries to be close to the reference sequence length U : $\mathcal{L}_{\text{qua}} = |U - \sum_{i=1}^U \sum_{j=1}^T \alpha_{i,j}|$. Scaling $\alpha_{i,j}$ properly is expected to encourage the discreteness of $p_{i,j}$ during training, which also leads to better estimation of $\beta_{i,j}$ in Eq. (3). The objective function is designed with the interpolation of the negative log-likelihood $\mathcal{L}_{\text{mocha}}$, CTC loss \mathcal{L}_{ctc} , and quantity loss \mathcal{L}_{qua} with tunable interpolation weights λ_{ctc} ($0 \leq \lambda_{\text{ctc}} \leq 1$) and λ_{qua} (≥ 0) as follows:

$$\mathcal{L}_{\text{total}} = (1 - \lambda_{\text{ctc}}) \mathcal{L}_{\text{mocha}} + \lambda_{\text{ctc}} \mathcal{L}_{\text{ctc}} + \lambda_{\text{qua}} \mathcal{L}_{\text{qua}} \quad (4)$$

We perform joint optimization with the CTC objective by sharing the encoder sub-network to encourage monotonicity of the input-output alignment [25].

4. CTC-synchronous training (CTC-ST)

Hard monotonic attention in MoChA depends entirely on past alignments because of the left-to-right dependency in Eq. (1). Although the scaling of the monotonic attention scores $\alpha_{i,j}$ can be facilitated with quantity regularization described in Section 3.3, the alignment errors in the middle and latter steps cannot be recovered, which leads to the significant error propagation of $\alpha_{i,j}$ to the latter tokens as the output sequence becomes longer. Unlike CTC, the decoder is autoregressive, so it is difficult to apply the backward algorithm during the alignment marginalization process. Consequently, the decision boundaries tend to shift to the right side (future) from the corresponding reference acoustic boundaries [15, 17, 18].

On the other hand, since CTC marginalizes all possible alignments with the forward-backward algorithm during training, its decision boundaries are more reliable and tend to shift to the left compared to those of MoChA. We indeed observed a delay of a few frames between the decision boundaries from the initial MoChA model and the corresponding CTC spikes when

sharing the same encoder (see Figure 1). This also arises a problem that hard monotonic attention activates at the different position over the shared encoder outputs from the corresponding CTC spike. Meanwhile, the well-trained CTC posterior probability distributions tend to peak in sharp spikes [3]. Therefore, the decision boundaries from CTC are expected to serve as an effective guide for MoChA to learn accurate alignments. In other words, MoChA can correct error propagation from past decision boundaries with the help of the CTC alignments.

In this work, we propose *CTC-synchronous training* (CTC-ST) to provide reliable CTC alignments as references to MoChA for learning robust monotonic alignments. The MoChA model is trained to mimic the CTC model to generate the similar decision boundaries. Both the MoChA and CTC branches are jointly optimized by sharing the same encoder, and the decision boundaries as the reference are extracted from the CTC branch. Thus, we can train the model in an end-to-end fashion without the use of external alignments, unlike in [18]. Synchronizing both decision boundaries can be regarded as explicit interaction between MoChA and CTC on the decoder side.

We use the most probable CTC path $\bar{\pi}$ of length T after forced alignment with the forward-backward algorithm [12] and use time indices of non-blank labels in $\bar{\pi}$ as the reference boundary positions $\mathbf{b}^{\text{ctc}} = (b_1^{\text{ctc}}, \dots, b_U^{\text{ctc}})$. Note that the leftmost index is used when non-blank labels are repeated without interleaving any blank labels. For the end-of-sentence token, the last input index (i.e., T) is used. \mathbf{b}^{ctc} is generated with model parameters at each training step on-the-fly and updates as the training continues. The objective function of the CTC-ST $\mathcal{L}_{\text{sync}}$ is defined as follows:

$$\mathcal{L}_{\text{sync}} = \frac{1}{U} \sum_{i=1}^U |b_i^{\text{ctc}} - b_i^{\text{mocha}}|$$

where $b_i^{\text{mocha}} = \sum_{j=1}^T j \alpha_{i,j}$ is the expected decision boundary of MoChA for the i -th token during training. The total objective function in Eq. (4) is modified accordingly as follows:

$$\mathcal{L}_{\text{total}} = (1 - \lambda_{\text{ctc}}) \mathcal{L}_{\text{mocha}} + \lambda_{\text{ctc}} \mathcal{L}_{\text{ctc}} + \lambda_{\text{qua}} \mathcal{L}_{\text{qua}} + \lambda_{\text{sync}} \mathcal{L}_{\text{sync}} \quad (5)$$

where $\lambda_{\text{sync}} (\geq 0)$ is a tunable parameter, set to 1.0 in this work. Unless otherwise noted, λ_{qua} is set to 0 when using CTC-ST.

4.1. Curriculum learning strategy

Since CTC-ST is designed to bring decision boundaries from both MoChA and CTC closer, their alignment probabilities must be peaky to minimize $\mathcal{L}_{\text{sync}}$ in Eq. (5). In the early training stage, however, monotonic attention scores $\alpha_{i,j}$ tend to be diffused over several frames and are not normalized to sum up to one. Applying CTC-ST from scratch is ineffective, so we use a curriculum learning strategy instead. First, the MoChA model with a standard BLSTM encoder is trained with random initialization together with Eq. (4) until convergence (**stage-1**). Next, after loading model parameters in stage-1, future contexts are restricted for the LC-BLSTM encoder and the parameters are optimized with CTC-ST by Eq. (5) (**stage-2**). Although we mainly use the LC-BLSTM encoder, this two-staged training can also be applied to the MoChA model with the unidirectional LSTM encoder by using the same encoder in both stages.

4.2. Combination with SpecAugment

SpecAugment has been shown to greatly enhance the decoder in the global attention model by performing on-the-fly data augmentation [20]. However, since it introduces time and frequency masks into the input log-mel spectrogram, the recur-

Table 1: *TEDLIUM2 results*. *Quantity regularization is used.

	Model	%WER
Offline	LSTM - Global attention	11.9
	BLSTM - Global attention (T1)	9.5
	+ LC-BLSTM-40+20 (seed: T1)	10.1
	+ LC-BLSTM-40+40 (seed: T1)	9.7
	BLSTM - MoChA	12.6
	+ Quantity regularization (T2) + CTC-ST	9.8 10.2
Streaming	LSTM - MoChA* (T3) + CTC-ST (T4, seed: T3)	15.0 13.2
	LC-BLSTM-40+20 - MoChA* (seed: T2) + CTC-ST	12.2 10.5
	LC-BLSTM-40+40 - MoChA* (T5, seed: T2) + CTC-ST (T6)	11.3 9.9

ency in Eq. (1) can be easily collapsed after the masked region. In our experiments, MoChA was not shown to improve as the global attention model with SpecAugment. In contrast, CTC-ST can recover $\alpha_{i,j}$ after the masked region with CTC spikes since CTC is formulated on the conditional independence assumption per frame. Therefore, CTC-ST is beneficial for MoChA to learn the monotonic alignments that withstand noisy inputs. We will also analyze the impact of mask size.

5. Experiments

5.1. Experimental setup

We used the TEDLIUM release-2 (210 hours, lecture) [26] and Librispeech (960 hours, reading) [27] corpora for experimental evaluations. We extracted 80-channel log-mel filterbank coefficients computed with a 25-ms window size shifted every 10 ms using Kaldi [28]. We performed 3-fold speed perturbation [29] on the TEDLIUM2 corpus with factors of 0.9, 1.0, and 1.1.

The encoders were composed of two CNN blocks followed by five layers of (LC-)BLSTM. Each CNN block was composed of two layers of CNN having a 3×3 filter followed by a max-pooling layer with a stride of 2×2 , which resulted in 4-fold frame rate reduction. We set the number of cells in each (LC-)BLSTM layer to 512 per direction. We summed up the LSTM outputs in both directions at each layer to reduce the input dimension of the subsequent (LC-)BLSTM layer [30]. The memory cells were doubled when using the unidirectional LSTM encoder. The decoder was a single layer of unidirectional LSTM with 1024-dimensional memory cells. For offline models, we used the location-based attention [6]. We set the chunk size w of MoChA to 4, which was tuned in our preliminary experiments. r in Eq. (2) was initialized with -4 . We used 10k vocabularies based on the Byte Pair Encoding (BPE) algorithm [31].

Optimization was performed using Adam [32] with learning rate $1e-3$ and it was exponentially decayed. We used dropout and label smoothing [33] with probabilities 0.4 and 0.1, respectively. λ_{ctc} was set to 0.3. We set λ_{qua} to 2.0 and 0.01 on the TEDLIUM2 and Librispeech corpora, respectively. We used a 4-layer LSTM language model (LM) with 1024 memory cells for decoding with a beam width of 10 [34]. Scores were normalized by the number of tokens for MoChA. CTC scores were not leveraged during inference.¹

5.2. Results

Table 1 shows the results for the TEDLIUM2. For offline models, MoChA (T2) approached the performance of the global at-

¹We implemented models with Pytorch [35]. Detailed hyperparameter settings during training and decoding are available at https://github.com/hirofumi0810/neural_sp.

Table 2: Results of curriculum learning on the TEDLIUM2

Model		Quantity regularization	CTC-ST	%WER
Streaming	LC-BLSTM-40+40	✓	–	12.3
	+ Curriculum learning	–	–	16.9
	(seed: T2)	✓	–	11.3
		–	✓	9.9
		✓	✓	10.1

Table 3: Results with SpecAugment on the TEDLIUM2

Model		F	T	%WER
Offline	Transformer [2]	30	40	8.1
	BLSTM - Global attention [36]	N/A	N/A	8.8
	BLSTM - Global attention	–	–	9.5
		27	100	8.1
Streaming	LC-BLSTM-40+40 - MoChA (seed: T2)	–	–	11.3
		27	100	12.8
		13	50	11.2
	+ CTC-ST	–	–	9.9
		27	100	9.0
		27	50	8.6
		13	50	9.0




Figure 2: WER performance bucketed by different input length on the TEDLIUM2

tention model (T1). Quantity regularization was essential for attaining suitable performance for the baseline MoChA. CTC-ST also improved the performance by a large margin, but it was less effective than quantity regularization when applying from scratch. This is because the scale of $\alpha_{i,j}$ was not adequate in the early training stage. For streaming models, our proposed CTC-ST with the curriculum learning strategy significantly improved the performances of the unidirectional LSTM, LC-BLSTM-40+20, and LC-BLSTM-40+40 MoChA models by 12.0, 13.9, and 12.3%, respectively. Note that CTC-ST was not combined with quantity regularization in stage-2. Although a larger future context was helpful for boosting performance, the effectiveness of CTC-ST was orthogonal.

Next, we investigated the effectiveness of regularization terms and the curriculum learning strategy for the LC-BLSTM-MoChA, shown in Table 2. We used BLSTM-MoChA (T2) as a seed model except for the first row and optimized the model with either CTC-ST, quantity regularization, or both. Curriculum learning was highly effective and CTC-ST (T6) significantly outperformed the case using quantity regularization (T5). Combining the two did not lead to any further improvements although it was more effective than the model with quantity regularization only. This is likely because CTC-ST has an effect to encourage the proper scale of $\alpha_{i,j}$ in hard monotonic attention similarly to quantity regularization.

The combination of CTC-ST and SpecAugment is shown in Table 3. We used two time masks with time mask parameter T and two frequency masks with frequency mask parameter F in [20]. We applied SpecAugment to MoChA in stage-2 only because applying SpecAugment from scratch did not converge. SpecAugment did not improve the performance of the MoChA models without the guide from CTC alignments because the attention scores in Eq. (1) can be easily collapsed as mentioned in Section 4.2. CTC-ST solved this issue and led to an additional

Table 4: Librispeech results. *Quantity regularization is used. \diamond SpecAugment is used.

Model		%WER	
		clean	other
Offline	BLSTM - Global attention \diamond [2]	3.3	10.8
	BLSTM - Global attention (L1)	3.1	9.5
	+ SpecAugment (seed: L1)	2.8	7.6
	BLSTM - MoChA	3.6	10.5
	+ Quantity regularization (L2)	3.3	10.0
Streaming	Transformer - CIF \diamond [14]	3.3	9.7
	Transformer - Triggered attention \diamond [22]	2.8	7.2
	PTDLSTM - Triggered attention [21]	5.9	16.8
	LSTM - MoChA + MWER \diamond [15]	5.6	15.6
	LSTM - MoChA + {char, BPE}-CTC [37]	4.4	15.2
	LC-BLSTM - sMoChA [16]	6.0	16.7
	LC-BLSTM - MTA [19]	4.2	12.3
	LSTM - MoChA * (L3)	5.3	14.5
	+ CTC-ST (seed: L3)	4.7	13.6
	++ SpecAugment ($F=13, T=50$)	4.2	11.2
	LC-BLSTM-40+40 - MoChA * (seed: L2)	4.1	11.2
	+ SpecAugment * ($F=27, T=100$)	5.0	9.7
	+ SpecAugment * ($F=13, T=50$)	4.0	9.5
	+ CTC-ST	3.9	11.2
	++ SpecAugment ($F=27, T=100$)	3.6	9.2
++ SpecAugment ($F=27, T=50$)	3.5	9.1	
++ SpecAugment ($F=13, T=50$)	3.6	9.4	

13.1% relative improvement. Moreover, CTC-ST was robust to input mask size.

We plotted the WERs bucketed by input length in Figure 2. The largest CTC-ST gains came from long utterances. The offline global attention model (T1) did not have difficulty in recognizing long utterances, whereas the initial LC-BLSTM MoChA model (T5) did. The proposed CTC-ST mitigated this problem (T6). The decision boundaries from the MoChA and CTC branches extracted from T3 and T4 are visualized in Figure 1. We found that the gap between the two boundaries was reduced, and the CTC spikes slightly shifted to the left. This is also beneficial for reducing user perceived latency [18].

Finally, the results for Librispeech are shown in Table 4. When using the unidirectional LSTM encoder for MoChA, we obtained 11.3% and 6.2% relative improvements with CTC-ST on the test-clean and test-other sets, respectively. SpecAugment further improved the performance. As the training data size is much larger, the gains from quantity regularization for the BLSTM encoder and CTC-ST for the LC-BLSTM encoder were smaller than those in TEDLIUM2, although both were still beneficial. CTC-ST was effective for obtaining gains from SpecAugment and led to additional improvements of 10.2% and 18.7% on the test-clean and test-other sets, respectively. The model parameter size of LC-BLSTM-MoChA was 53.77M and fixed through all experiments. Our optimal model requires only 860ms (= 400ms (N_c) + 400ms (N_r) + 60ms (CNN)) lookahead frames and the decoding complexity is linear.

6. Conclusion

We proposed CTC-synchronous training (CTC-ST) to provide hard monotonic attention in MoChA with reliable alignments extracted from CTC. By jointly training both sub-networks with the shared encoder and generating CTC alignments simultaneously, we enabled effective interaction between MoChA and CTC. Experimental evaluations revealed that CTC-ST significantly improved the performance of MoChA and greatly reduced the gap from the offline models. Further gains with SpecAugment were obtained when CTC-ST was applied, thus verifying its robustness to noisy alignments.

7. 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 *Proceedings of ICASSP*. IEEE, 2018, pp. 4774–4778.
- [2] 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 *Proceedings of ASRU*. IEEE, 2019, pp. 499–456.
- [3] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of ICML*, 2006, pp. 369–376.
- [4] A. Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [5] H. Sak, M. Shannon, K. Rao, and F. Beaufays, “Recurrent Neural Aligner: An encoder-decoder neural network model for sequence to sequence mapping,” in *Interspeech*, 2017, pp. 1298–1302.
- [6] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *Proceedings of NeurIPS*, 2015, pp. 577–585.
- [7] 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 *Proceedings of Interspeech*, 2017, pp. 939–943.
- [8] 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 *Proceedings of ASRU*. IEEE, 2017, pp. 206–213.
- [9] M. Huang, Y. Lu, L. Wang, Y. Qian, and K. Yu, “Exploring model units and training strategies for end-to-end speech recognition,” in *Proceedings of ASRU*. IEEE, 2019, pp. 524–531.
- [10] C. Raffel, M.-T. Luong, P. J. Liu, R. J. Weiss, and D. Eck, “Online and linear-time attention by enforcing monotonic alignments,” in *Proceedings of ICML*, 2017, pp. 2837–2846.
- [11] C.-C. Chiu and C. Raffel, “Monotonic chunkwise attention,” in *Proceedings of ICLR*, 2018.
- [12] N. Moritz, T. Hori, and J. Le Roux, “Triggered attention for end-to-end speech recognition,” in *Proceedings of ICASSP*. IEEE, 2019, pp. 5666–5670.
- [13] M. Li, M. Liu, and H. Masanori, “End-to-end speech recognition with adaptive computation steps,” in *Proceedings of ICASSP*. IEEE, 2019, pp. 6246–6250.
- [14] L. Dong and B. Xu, “CIF: Continuous integrate-and-fire for end-to-end speech recognition,” in *Proceedings of ICASSP*. IEEE, 2020, pp. 6079–6083.
- [15] K. Kim, K. Lee, D. Gowda, J. Park, S. Kim, S. Jin, Y.-Y. Lee, J. Yeo, D. Kim, S. Jung *et al.*, “Attention based on-device streaming speech recognition with large speech corpus,” in *Proceedings of ASRU*. IEEE, 2019, pp. 956–963.
- [16] H. Miao, G. Cheng, P. Zhang, T. Li, and Y. Yan, “Online hybrid CTC/attention architecture for end-to-end speech recognition,” in *Proceedings of Interspeech*, 2019, pp. 2623–2627.
- [17] R. Fan, P. Zhou, W. Chen, J. Jia, and G. Liu, “An online attention-based model for speech recognition,” in *Proceedings of Interspeech*, 2019, pp. 4390–4394.
- [18] H. Inaguma, Y. Gaur, L. Lu, J. Li, , and Y. Gong, “Minimum latency training strategies for streaming sequence-to-sequence ASR,” in *Proceedings of ICASSP*. IEEE, 2020, pp. 6064–6068.
- [19] H. Miao, G. Cheng, P. Zhang, T. Li, and Y. Yan, “Online hybrid CTC/attention end-to-end automatic speech recognition architecture,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1452–1465, 2020.
- [20] 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 *Proceedings of Interspeech*, 2019, pp. 2613–2617.
- [21] N. Moritz, T. Hori, and J. Le Roux, “Streaming end-to-end speech recognition with joint CTC-attention based models,” in *Proceedings of ASRU*. IEEE, 2019, pp. 936–943.
- [22] N. Moritz, T. Hori, and J. L. Roux, “Streaming automatic speech recognition with the Transformer model,” *arXiv preprint arXiv:2001.02674*, 2020.
- [23] Y. Zhang, G. Chen, D. Yu, K. Yaco, S. Khudanpur, and J. Glass, “Highway long short-term memory RNNs for distant speech recognition,” in *Proceedings of ICASSP*. IEEE, 2016, pp. 5755–5759.
- [24] S. Xue and Z. Yan, “Improving latency-controlled BLSTM acoustic models for online speech recognition,” in *Proceedings of ICASSP*. IEEE, 2017, pp. 5340–5344.
- [25] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, “Hybrid CTC/attention architecture for end-to-end speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [26] A. Rousseau, P. Deléglise, and Y. Estève, “TED-LIUM: An automatic speech recognition dedicated corpus,” in *Proceedings of LREC*, 2012, pp. 125–129.
- [27] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *Proceedings of ICASSP*. IEEE, 2015, pp. 5206–5210.
- [28] 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 *Proceedings of ASRU*. IEEE, 2011.
- [29] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, “Audio augmentation for speech recognition,” in *Proceedings of Interspeech*, 2015, pp. 3586–3589.
- [30] Z. Tüske, K. Audhkhasi, and G. Saon, “Advancing sequence-to-sequence based speech recognition,” in *Proceedings of Interspeech*, 2019, pp. 3780–3784.
- [31] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” in *Proceedings of ACL*, 2016, pp. 1715–1725.
- [32] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [33] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of CVPR*, 2016, pp. 2818–2826.
- [34] A. Kannan, Y. Wu, P. Nguyen, T. N. Sainath, Z. Chen, and R. Prabhavalkar, “An analysis of incorporating an external language model into a sequence-to-sequence model,” in *Proceedings of ICASSP*. IEEE, 2018, pp. 5824–5828.
- [35] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, “PyTorch: An imperative style, high-performance deep learning library,” in *Proceedings of NeurIPS*, 2019, pp. 8024–8035.
- [36] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, “A comparison of Transformer and LSTM encoder decoder models for ASR,” in *Proceedings of ASRU*. IEEE, 2019.
- [37] A. Garg, D. Gowda, A. Kumar, K. Kim, M. Kumar, and C. Kim, “Improved multi-stage training of online attention-based encoder-decoder models,” in *Proceedings of ASRU*. IEEE, 2019, pp. 70–77.