

One In A Hundred: Selecting the Best Predicted Sequence from Numerous Candidates for Speech Recognition

Zhengkun Tian^{1,2} and Jiangyan Yi¹ and Ye Bai^{1,2} and Jianhua Tao^{1,2,3} and Shuai Zhang^{1,2} and Zhengqi Wen¹

¹National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing, China

²School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

³CAS Center for Excellence in Brain Science and Intelligence Technology, Beijing, China

{zhengkun.tian, jiangyan.yi, jhtao, ye.bai, shuai.zhang, zqwen}@nlpr.ia.ac.cn

Abstract—The RNN-Transducers and improved attention-based encoder-decoder models are widely applied to streaming speech recognition. Compared with these two end-to-end models, the CTC model is more efficient in training and inference. However, it cannot capture the linguistic dependencies between the output tokens. Inspired by the success of two-pass end-to-end models, we introduce a transformer decoder and the two-stage inference method into the streaming CTC model. During inference, the CTC decoder first generates many candidates in a streaming fashion. Then the transformer decoder selects the best candidate based on the corresponding acoustic encoded states. The second-stage transformer decoder can be regarded as a conditional language model. We assume that a large enough number and enough diversity of candidates generated in the first stage can compensate the CTC model for the lack of language modeling ability. All the experiments are conducted on a Chinese Mandarin dataset AISHELL-1. The results show that our proposed model can implement streaming decoding in a fast and straightforward way. Our model can achieve up to a 20% reduction in the character error rate than the baseline CTC model. In addition, our model can also perform non-streaming inference with only a little performance degradation.

I. INTRODUCTION

Streaming speech recognition has been applied to many real scenarios, like meeting real-time transcription and keyboard dictation systems on mobile phones. The mainstream models for streaming speech recognition include RNN-Transducer (RNN-T) models and the improved attention-based encoder-decoder models [1], [2], [3], [4], [5], [6], [7].

The RNN-T model, which utilizes unidirectional recurrent neural networks as the encoder, can be directly applied to streaming speech recognition [3], [6], [8], [9]. However, the RNN-T model suffers from the inefficiency of training and inference [10], [11]. The traditional attention-based encoder-decoder model decodes the output sequences based on the previously predicted tokens and the entire acoustic encoded states, which prevents it from decoding the output sequence in a streaming way [7]. The improved attention-based encoder-decoder models mainly include the following versions: the monotonic chunk-wise attention (MoChA) [12], the triggered attention [4], the continuous integrate-and-fire (CIF) [13], the

synchronous transformer [7], and so on. These models utilize many complicated methods to segment the acoustic encoded states and then compute the attention weights on the segment states. As a kind of end-to-end model, connectionist temporal classification (CTC) models were first proposed to transcribe the acoustic feature sequences into the corresponding text sequence [14], [15], [16]. The CTC models cannot capture the linguistic dependencies between the output tokens [8]. To address this problem, CTC models tend to depend on an external language model or a WFST-based graph to improve the performance [16], [17]. Recently, some very deep convolution CTC models, like Jasper [18] and ContextNet [19], have achieved competitive performance with other end-to-end models. Although very deep convolution layers make the model able to model long-range contexts, it will also result in a large latency and prevent the model from being applied to streaming speech recognition.

The CTC models are more efficient in inference than the other two end-to-end models. Inspired by the success of two-pass end-to-end models [20], we introduce a transformer decoder and the two-stage inference method into the CTC model to improve the model performance on streaming speech recognition. Now that the CTC model cannot model the linguistic dependencies between the output tokens, the transformer decoder can be regarded as a conditional language model, which can pick out the best sequence from the top-N candidates based on the linguistic knowledge and the corresponding acoustic information. We assume that a large enough number and diversity of candidates generated by the CTC model can compensate for the lack of language modeling ability. We rename the process that selecting the best sequence from numerous candidates as One-In-A-Hundred(OAH). Our model consists of three components, a latency-controlled streaming transformer encoder, a CTC decoder, and a transformer decoder. We first force the self-attention mechanism in the transformer encoder to focus on the local context to model the input sequence in a streaming fashion. Then we introduce a latency-controlled context layer at the top of the transformer encoder to model the future context of limited range. During

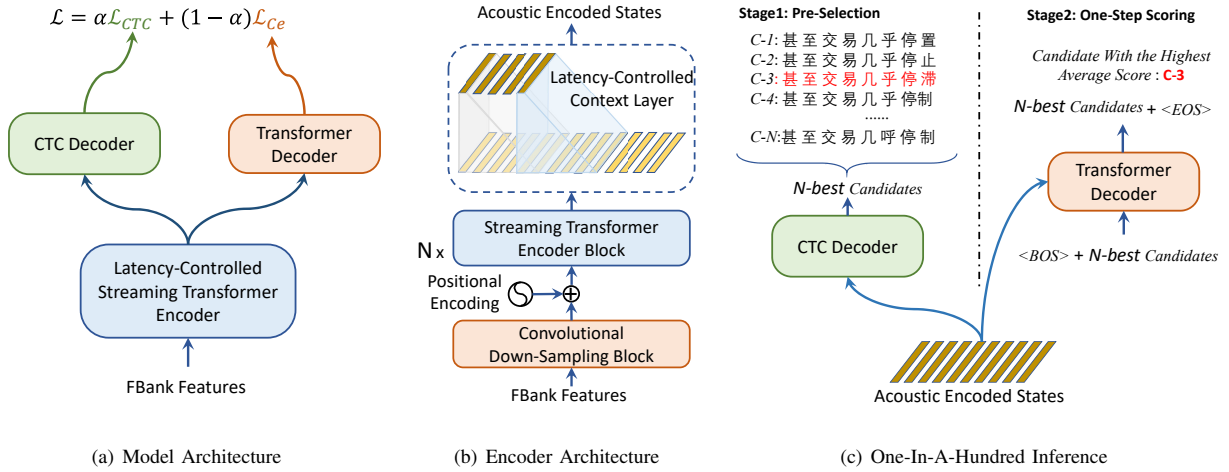


Fig. 1: (a) illustrates the overall structure of our proposed model. The model consists of a latency-controlled streaming transformer encoder (LC-STE), a CTC decoder, and a transformer decoder. (b) illustrates the details of the encoder. The LC-STE consists of a convolutional down-sampling block, N streaming transformer blocks, and a latency-controlled context layer. (c) illustrates the one-in-a-hundred inference process. The inference process can be split into two stages: pre-selection and one-step scoring. The CTC decoder can generate many possible candidates at the pre-selection stage. Then, the transformer decoder selects the best predicted sequence from the candidates in one step. The best candidate C-3 means that “the deal is almost stalled”.

training, these three components are optimized jointly. The one-in-a-hundred inference process can be divided into two stages: pre-selection and one-step scoring. At the pre-selection stage, the CTC decoder generates many possible sequences as candidates. At the one-step scoring stage, the transformer decoder scores all these candidates based on the corresponding acoustic encoded states and then selects the sequence with the highest average score as the final predict sequence.

The contributions of this paper can be summarized in the following two aspects. Firstly, we introduce the two-pass inference method into the streaming CTC model and propose an assumption that a large number of candidates generated in the first inference stage can greatly improve the model performance. Secondly, we propose a dual inference mode for the hybrid CTC and transformer model, which can be applied to the streaming and non-streaming tasks directly. All experiments are conducted on a public Chinese Mandarin dataset AISHELL-1. The results show that our proposed model can achieve up to a 20% reduction in the character error rate compared to the baseline CTC model. Furthermore, our model can also perform non-streaming decoding.

The remainder of this paper is organized as follows. We introduce the details of the model and related works in Section 2 and Section 3 respectively. Our experimental setup and results will be presented in Section 4. The conclusions will be given in Section 5.

II. OUR PROPOSED METHOD

Our model consists of three components, a latency-controlled streaming transformer encoder, a CTC decoder, and

a transformer decoder, as shown in Fig.1(a).

A. Model Architecture

1) *The Latency-Controlled Streaming Encoder*: The transformer encoder generally consists of a convolutional down-sampling block, a positional embedding, and N transformer encoder blocks. The transformer encoder block is composed of a multi-head self-attention block and a feed-forward network layer [22]. The original transformer encoder depends on all context to compute the self-attention mechanism, which makes it powerful to model long-range temporal information, and prevents it from being applied to streaming speech recognition [23]. As depicted in Fig.1(b), we proposed a new streaming encoder named latency-controlled streaming transformer encoder (LC-STE), which improves the transformer encoder in two aspects.

On the one hand, we force the self-attention of LC-STE to focus on the fixed-range previous context and ignore the future context completely. It is well known that extending the range of future context will improve the performance of the model. However, with the increase of network depth, it will lead to excessive dependence on future information, making the latency of inference increase sharply. The streaming self-attention (SSA) can be expressed as follows.

$$SSA(Q_t, K_{t-\tau:t}, V_{t-\tau:t}) = softmax\left(\frac{Q_t K_{t-\tau:t}^T}{\sqrt{d_k}}\right) V_{t-\tau:t} \quad (1)$$

where τ means the range of the previous context. Q_t is the query vector at the time t . K and V indicate the context. d_k is the dimension of K .

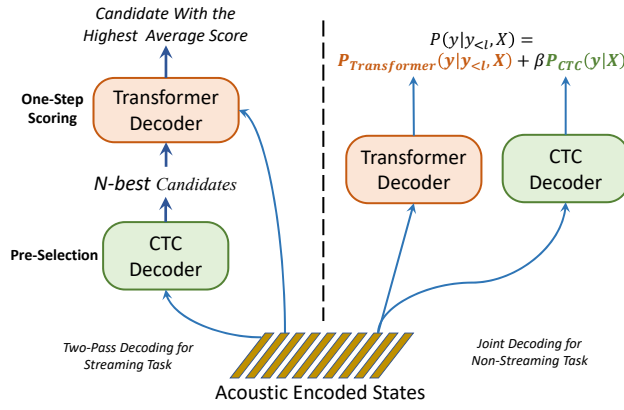


Fig. 2: The Dual Inference Mode for Streaming and Non-Streaming Tasks.

On the other hand, inspired by DeepSpeech2 [16], we put a latency-controlled context layer at the top of the encoder to model the fixed-range future context. The context layer is composed of a 1D convolution layer with kernel size $\varepsilon + 1$. The ε means the range of future context. The acoustic encoded states can be expressed as \hat{h}_t .

$$\hat{h}_t = \sum_{i=0}^{\varepsilon} w_i h_{t+i} + b_i \quad (2)$$

Where h is the output of the last streaming transformer encoder block. w and b indicate the weight and bias respectively. The latency can be expressed as $40 \times (\varepsilon + 1)$ ms, where 40ms means the convolutional down-sampling block can reduce the length of input sequences with frame shift 10ms by 4 times. With the latency-controlled context layer, the computation of ideal latency does not rely on the depth of the encoder.

2) *The CTC Decoder and Transformer Decoder:* The CTC decoder contains only a linear project layer, which is utilized to compute the label posterior probabilities of CTC. The transformer decoder is the same as it in speech transformer [23], which is composed of masked multi-head self-attention, multi-head cross-attention, feed-forward network and positional embedding. During training, the decoder adopts a triangle-like mask to mask force the decoder to focus on the previous tokens [22]. We compute the CTC loss \mathcal{L}_{CTC} and the cross entropy loss \mathcal{L}_{CE} with label smooth for the CTC decoder and transformer decoder respectively. These three components are optimized jointly. The joint loss is expressed as \mathcal{L}_{Joint} .

$$\mathcal{L}_{Joint} = \alpha \mathcal{L}_{CTC} + (1 - \alpha) \mathcal{L}_{CE} \quad (3)$$

where α is the weight of \mathcal{L}_{CTC} .

B. Inference

We introduce the two-pass inference method for streaming inference, as shown in Fig.1(c). We assume that a large enough number of candidates generated by the CTC model at the first-stage inference can compensate for the lack of language modeling ability. Therefore, we figuratively rename

the process that selecting the best sequence from numerous candidates as One-In-A-Hundred(OAH). The inference process can be split into two stages: pre-selection and one-step scoring. At the pre-selection stage, the CTC decoder can generate N candidate sequences by prefix beam search [24]¹. The number of candidates is equal to the width of beams. At the one-step scoring stage, the transformer decoder scores all the candidates. Different from the language model rescoring, the one-step scoring relies on the corresponding acoustic encoded states to match the selected candidate with the same pronunciation. Due to the transformer decoder can perform parallel computation, the scoring can be finished in one step, which greatly improve the efficiency of inference. As long as the candidates are diverse enough, we might be able to pick out the best one. The score of a sentence of length L can be expressed as

$$S_{y_{1:L}} = OneStepScoring(y_{1:L}, \hat{h}_{enc}) * 1/L \quad (4)$$

where \hat{h}_{enc} means the corresponding acoustic encoded states. We apply length normalization to the scoring process to prevent the model from tending to select the short candidate as the predicted sequence.

In addition, our model support two inference mode, streaming and non-streaming, as depicted in Fig.2. When we adopt the CTC decoder as the leading role and the transformer decoder to re-score the candidates generated by the CTC decoder, the model can perform the streaming inference. The process is named one-in-a-hundred. From another viewpoint, when we regard the transformer decoder as the leading role and the CTC decoder as the assistant, the transformer decoder can still model the whole context and decode in a non-streaming fashion [21]. The non-streaming decoding process starts with the beginning token $\langle S/E \rangle$. At every step, we interpolate the scores of two decoders. The model will repeat the above process until the end-of-sentence token is predicted.

III. RELATED WORKS

Our work can be regarded as a good extension and supplement of the previous two-pass end-to-end model [20]. The inference of these two methods can be divided into two stages, generating N-best candidates and selecting the best candidates by the extra decoder. However, there are at least three significant differences between these two works.

Firstly, the previous two-pass method is applied to the hybrid RNNT and LAS model. It utilizes the RNN-based LAS decoder to scoring the candidates, which is hard to perform in parallel. However, we introduce the two-pass method into the classical hybrid CTC and attention model and utilize a transformer decoder to rescore the candidate in one step, which is faster and more effective.

Secondly, it's well known that the RNNT model has a strong ability to model language. For the hybrid RNNT and LAS model, the two-stage can be regarded as the fusion of two

¹The code of prefix beam search is available at <https://github.com/PaddlePaddle/DeepSpeech/tree/develop/decoders>

conditional language models. By contrast, the CTC model lacks the ability of modeling language. The second stage can be considered as a process that fuses linguistic information effectively.

Finally, it's very challenging for RNNT to generate many candidate sequences by beam search during inference, which will cost plenty of time and memory. However, the CTC model performs more effectively during inference and training. Therefore, we try to generate up to one hundred candidates in the first stage of inference, which far exceeds the candidates from RNNT models. We assume that a large enough number and enough diversity of candidates can compensate the CTC model for the lack of language modeling ability.

IV. EXPERIMENTS AND RESULTS

A. Dataset

In this work, all experiments are conducted on a public Mandarin speech corpus AISHELL-1². The training set contains about 150 hours of speech (120,098 utterances) recorded by 340 speakers. The development set contains about 20 hours (14,326 utterances) recorded by 40 speakers. And about 10 hours (7,176 utterances / 36109 seconds) of speech is used as the test set. The speakers of different sets are not overlapped.

B. Experimental Setup

For all experiments, we use 40-dimensional FBANK features computed on a 25ms window with a 10ms shift. We choose 4233 characters (including a padding symbol <PAD>, an unknown symbol <UNK>, and an start-or-end-of-sentence symbol <S/E>) as modeling units. <S/E> is also utilized as a blank symbol for CTC Decoder.

Our model ³ consists of 12 encoder blocks and 6 decoder blocks. There are 4 heads in multi-head attention. The 2D convolution front end utilizes two-layer time-axis CNN with ReLU activation, stride size 2, channels 320, and kernel size 3. Both the output size of the multi-head attention and the feed-forward layers are 512. The hidden size of feed-forward layers is 768. The range of the self-attention in streaming transformer encoder blocks is limited from 10 frames on the left to the current position ($\tau = 10$). We adopt an Adam optimizer with warmup steps 12000 and the learning rate scheduler reported in [22]. After 100 epochs, we average the parameters saved in the last 30 epochs. We also use the time mask and frequency mask method proposed in [25] instead of speed perturbation.

We use the character error rate (CER) to evaluate the performance of different models. For evaluating the inference speed of different models, we decode utterances one by one to compute real-time factor (RTF) on the test set. The RTF is the time taken to decode one second of speech. All experiments are conducted on a GeForce GTX TITAN X 12G GPU.

²<https://openslr.org/33/>

³Our model is built on the open-source code <https://github.com/ZhengkunTian/OpenTransformer>

C. Results

1) *Comparison of models with different CTC weights:* We first compare the models with different CTC weights. We set the range of future context in latency-controlled context layer to 10 and adopt beam search with width 10 during inference. We evaluate the CER on the development and test set in two ways, which are one-pass beam-search (marked as OPS) and two-pass one-in-a-hundred (marked as OAH). This naming setting is still adopted in subsequent experiments.

As shown in Table.1, the model with CTC weight 0.1 can achieve the best performance on development and test test. Furthermore, it's obvious that applying our OAH strategy can generally improve the performance of the model. However, with the increase of CTC weight, the performance of the model gradually deteriorated. We assume that the weight α can balance the importance of the CTC model in the training process. An inappropriate weight will make the model unbalanced in the training process, leading to the decline of model performance. When the weight α is 1.0, the transformer decoder is discarded. Under this condition, the model can be regarded as a CTC model. By contrast, we find that the joint training with an appropriate weight can improve performance.

TABLE I: Comparison of models with different CTC weights (CER %).

Weight α	0.05	0.1	0.2	0.3	0.5	0.7	1.0
Dev	OPS	8.46	8.22	8.26	8.31	8.89	8.88
	OAH	7.20	6.92	7.04	7.14	7.55	8.03
Test	OPS	9.37	9.25	9.36	9.67	10.21	11.03
	OAH	8.15	7.76	8.02	8.18	8.69	9.26

2) *Comparison of models with different ranges of future context:* We evaluate the models with different range ε of future context in Table 2. We set the CTC weight of all models to 0.1 and adopt the beam search with width 10 during inference. The latency is proportional to the range of future context. It can be calculated by $40 \times (\varepsilon + 1)$ ms. It appears that the more future information the model focuses on, the better performance it achieves. Meanwhile, it will increase the computation and the real-time rate. When the range ε is set to 10 and 20, there is no significant difference in the performance. Considering that the model that focuses on the next 10 frames can achieve lower latency and faster inference speed, we set the range of future context to 10 in subsequent experiments.

TABLE II: Comparison of models with different ranges of future context (CER %).

ε	0	1	5	10	20	
Latency(ms)	40	80	240	440	840	
Dev	OPS	11.70	10.79	8.31	8.21	8.07
	OAH	8.75	8.02	6.97	6.91	6.82
Test	OPS	13.15	12.18	9.43	9.25	9.26
	OAH	10.18	9.26	7.82	7.76	7.73
RTF	OPS	0.0175	0.0209	0.0254	0.0263	0.0274
	OAH	0.0217	0.0223	0.0286	0.0281	0.0295

3) *Comparison of the model with different beam widths:*

We choose the model with the CTC weight 0.1 and the future context range 10 to conduct these experiments in Table 3. The results show that the beam widths have little effect on the performance of the CTC decode results (without OAH). We assume that the CTC part of our model has learned a very sharp posterior probability distribution, which leads to a numerous difference between the probabilities of search paths. However, the beam width plays a vital role in our two-stage inference method. The model with OAH scoring can achieve up to 20% reduction in CER compared to the performance of the base CTC model (without OAH). The beam widths are equal to the number of candidate sequences. The larger the beam width, the better performance the model achieves. The model with beam width 50 achieves a CER of 7.41% and an RTF of 0.0380. When the width is large than 50, the performance will decline. We assume that generating more hypotheses with larger beam sizes will also produce more wrong sequences with better score, which may have a negative impact on the final performance of the model [26].

TABLE III: Comparison of the model with different beam widths (CER %).

Beam Width		1	5	10	50	100
Dev	OPS	8.22	8.22	8.22	8.21	8.21
	OAH	-	7.15	6.92	6.73	6.76
Test	OPS	9.25	9.25	9.25	9.25	9.25
	OAH	-	8.09	7.76	7.41	7.49
RTF	OPS	0.0202	0.0216	0.0263	0.0262	0.0421
	OAH	-	0.0212	0.0281	0.0380	0.0487

TABLE IV: Comparison with other models (CER %).

Model	Is Streaming	Test	RTF
TDNN-Chain(Kaldi) [27]	N	7.45	-
Speech-Transformer* [23]	N	6.47	0.0564
SA-Transducer [6]	N	9.30	0.1536
ST-NAT [28]	N	7.67	0.0056
Sync-Transformer [7]	Y	8.91	0.1183
SAN-CTC* [29]	N	7.80	0.0269
SAN-CTC*	Y	10.27	0.0248
+ RNNLM Rescoring	Y	9.02	0.0385
+ TransLM Rescoring	Y	8.93	0.0384
Our Model (OAH)	Y	7.41	0.0380
Our Model (Non-Streaming)	N	7.05	0.0812

* These models are re-implemented by ourselves according to the papers, which has the same parameters configuration as our model.

4) *Comparison with other models:* We also compare our model with other models. As shown in Table 4, our model with OAH can achieve a comparable performance with TDNN-Chain model [27] and ST-NAT [28]. What's more, our model has a real-time factor of 0.0380, which exceeds SA-Transducer [6], Sync-Transformer[7], and speech-transformer [23].

Our model with OAH achieves better performance compared to the CTC with language rescoring. We assume that the model can't distinguish between two grammatical sentences that have similar pronunciation only depending on linguistic

information. Therefore, the acoustic information still plays an import role in the second rescoring stage.

In addition, our model is able to decoding the output sequence in a non-streaming fashion (NS). Under this condition, our model achieves a CER of 7.05%. Compared with the speech-transformer with the same parameters configuration, our model has a little performance degradation.

V. CONCLUSIONS

In this paper, we improve the hybrid CTC and attention model and introduce a two-stage inference method named one-in-a-hundred (OAH). Our proposed model consists of three components, a latency-controlled streaming transformer encoder, a CTC decoder, and a transformer decoder. The latency-controlled streaming transformer encoder can model the streaming input feature sequence in very low latency. The streaming inference process can be split into two stages: sampling and one-step scoring. At the first stage, the CTC decoder can generate up to a hundred candidate sequences quickly. At the second stage, the transformer decoder score all the candidate sequences based on the corresponding acoustic encoded states in one step. We conduct the experiments on a public Chinese mandarin dataset AISHELL-1. The results show that our proposed method can achieve up to 20% reduction in CER compared to the baseline CTC model, which also proves our assumption that the two-stage OAH inference can compensate the CTC model for the lack of language modeling ability. What's more, our model also can perform non-streaming decoding with a little performance degradation.

VI. ACKNOWLEDGE

This work is supported by the National Key Research & Development Plan of China (No. 2017YFC0820602), the National Natural Science Foundation of China (No.61831022, No.61901473, No.61771472, No.61773379), and Inria-CAS Joint Research Project (No.173211KYSB20190049).

REFERENCES

- [1] N. Jaitly, D. Sussillo, Q. V. Le, O. Vinyals, I. Sutskever, and S. Bengio, "A neural transducer," *arXiv preprint arXiv:1511.04868*, 2015.
- [2] T. N. Sainath, C.-C. Chiu, R. Prabhavalkar, A. Kannan, Y. Wu, P. Nguyen, and Z. Chen, "Improving the performance of online neural transducer models," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5864–5868.
- [3] 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," *arXiv: Computation and Language*, 2018.
- [4] N. Moritz, T. Hori, and J. Le Roux, "Triggered attention for end-to-end speech recognition," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5666–5670.
- [5] L. Dong and B. 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.
- [6] Z. Tian, J. Yi, J. Tao, Y. Bai, and Z. Wen, "Self-Attention Transducers for End-to-End Speech Recognition," in *Proc. Interspeech 2019*, 2019, pp. 4395–4399. [Online]. Available:
- [7] Z. Tian, J. Yi, Y. Bai, J. Tao, S. Zhang, and Z. Wen, "Synchronous transformers for end-to-end speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7884–7888.

- [8] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [9] A. Graves, A.-r. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *2013 IEEE international conference on acoustics, speech and signal processing*. IEEE, 2013, pp. 6645–6649.
- [10] T. Bagby, K. Rao, and K. C. Sim, "Efficient implementation of recurrent neural network transducer in tensorflow," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 506–512.
- [11] J. Li, R. Zhao, H. Hu, and Y. Gong, "Improving rnn transducer modeling for end-to-end speech recognition," *arXiv: Computation and Language*, 2019.
- [12] C.-C. Chiu and C. Raffel, "Monotonic chunkwise attention," *arXiv preprint arXiv:1712.05382*, 2017.
- [13] L. Dong and B. Xu, "Cif: Continuous integrate-and-fire for end-to-end speech recognition," *arXiv preprint arXiv:1905.11235*, 2019.
- [14] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 369–376.
- [15] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates *et al.*, "Deep speech: Scaling up end-to-end speech recognition," *arXiv preprint arXiv:1412.5567*, 2014.
- [16] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *International conference on machine learning*, 2016, pp. 173–182.
- [17] Y. Miao, M. Gowayyed, and F. Metze, "Eesen: End-to-end speech recognition using deep rnn models and wfst-based decoding," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 167–174.
- [18] J. Li, V. Lavrukhin, B. Ginsburg, R. Leary, O. Kuchaiev, J. M. Cohen, H. Nguyen, and R. T. Gadde, "Jasper: An end-to-end convolutional neural acoustic model," *arXiv preprint arXiv:1904.03288*, 2019.
- [19] W. Han, Z. Zhang, Y. Zhang, J. Yu, C.-C. Chiu, J. Qin, A. Gulati, R. Pang, and Y. Wu, "Contextnet: Improving convolutional neural networks for automatic speech recognition with global context," *arXiv preprint arXiv:2005.03191*, 2020.
- [20] T. N. Sainath, R. Pang, D. Rybach, Y. He, R. Prabhavalkar, W. Li, M. Visontai, Q. Liang, T. Strohmaier, Y. Wu *et al.*, "Two-pass end-to-end speech recognition," *arXiv preprint arXiv:1908.10992*, 2019.
- [21] 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.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [23] L. Dong, S. Xu, and B. Xu, "Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5884–5888.
- [24] A. Y. Hannun, A. L. Maas, D. Jurafsky, and A. Y. Ng, "First-pass large vocabulary continuous speech recognition using bi-directional recurrent dnns," *arXiv preprint arXiv:1408.2873*, 2014.
- [25] 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," *arXiv preprint arXiv:1904.08779*, 2019.
- [26] W. Zhou, S. Ralf, and N. Hermann, "Robust beam search for encoder-decoder attention based speech recognition without length bias," in *Proc. Interspeech 2020*, 2020, pp. 1768–1772.
- [27] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi," in *Interspeech*, 2016, pp. 2751–2755.
- [28] Z. Tian, J. Yi, J. Tao, Y. Bai, S. Zhang, and Z. Wen, "Spike-triggered non-autoregressive transformer for end-to-end speech recognition," *arXiv preprint arXiv:2005.07903*, 2020.
- [29] J. Salazar, K. Kirchhoff, and Z. Huang, "Self-attention networks for connectionist temporal classification in speech recognition," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7115–7119.