

# Mask-CTC-based Encoder Pre-training for Streaming End-to-End Speech Recognition

Huaibo Zhao\*, Yosuke Higuchi\*, Yusuke Kida†, Tetsuji Ogawa\*, Tetsunori Kobayashi\*

\*Department of Communications and Computer Engineering, Waseda University, Tokyo, Japan

†LINE Corporation, Tokyo, Japan

zhao@pcl.cs.waseda.ac.jp

**Abstract**—Achieving high accuracy with low latency has always been a challenge in streaming end-to-end automatic speech recognition (ASR) systems. By attending to more future contexts, a streaming ASR model achieves higher accuracy but results in larger latency, which hurts the streaming performance. In the Mask-CTC framework, an encoder network is trained to learn the feature representation that anticipates long-term contexts, which is desirable for streaming ASR. Mask-CTC-based encoder pre-training has been shown beneficial in achieving low latency and high accuracy for triggered attention-based ASR. However, the effectiveness of this method has not been demonstrated for various model architectures, nor has it been verified that the encoder has the expected look-ahead capability to reduce latency. This study, therefore, examines the effectiveness of Mask-CTC-based pre-training for models with different architectures, such as Transformer-Transducer and contextual block streaming ASR. We also discuss the effect of the proposed pre-training method on obtaining accurate output spike timings, which contributes to the latency reduction in streaming ASR.

**Index Terms**—Streaming automatic speech recognition, latency reduction, Mask-CTC

## I. INTRODUCTION

In recent years, deep learning has become the core technology of automatic speech recognition (ASR) [1], [2]. End-to-end ASR further integrates the traditional separated components (i.e., acoustic, pronunciation, and language models) into a single deep neural network, significantly contributing to the simplicity of ASR developments [2]–[4]. End-to-end ASR models can be realized in various approaches, including connectionist temporal classification (CTC) [5], Transducer [6], and attention-based encoder-decoder [2], [4]. These end-to-end ASR approaches have greatly benefited from the adoption of Transformer [7]–[11], enabling a model to capture global contexts using the self-attention mechanism.

Streaming properties (i.e., real-time processing) are of vital importance in the applications of ASR systems. With the superior performance of Transformer, many efforts have been devoted to making Transformer-based ASR models streaming. Triggered attention-based ASR [12] obtains alignment information from CTC and realizes frame-synchronous decoding according to the CTC spike timings. Meanwhile, contextual block streaming ASR (CBS-ASR) [13] splits the input into blocks (chunks), and streaming encoder feature extraction is conducted on each block with the contexts inherited from the previous blocks. A block boundary detection algorithm is applied to detect the index boundary in each block, which

enables block-synchronous beam search decoding. Apart from the above streaming models based on attention-based encoder-decoder, the Transducer-based model can be naturally applied to streaming ASR. Transducer trains a model to align the output of the acoustic encoder with the output of the label encoder, enabling frame-synchronous decoding and making it a suitable framework for streaming ASR [8]. Transformer-Transducer (Transformer-T) [9] adopts Transformer for the acoustic encoder, where the chunk-wise attention mask limits the look-ahead range of the self-attention layer to ensure streaming properties.

For streaming ASR models in general, performance degradation occurs when the look-ahead range is limited from global to local, suppressing the advantage of the Transformer architecture (i.e., processing with long-range contexts). Consequently, longer look-ahead ranges are often required to provide adequate future contexts, leading to the growth of latency requirements. Therefore, capturing long-term contexts within short look-ahead ranges is essential for building successful streaming ASR systems.

One approach to realizing such a property is to utilize the Mask-CTC [14], [15] framework. With conditional masked language model (CMLM) [16], [17] and CTC multi-task training, Mask-CTC trains an encoder network to extract acoustic feature representation that contributes to capturing long-term output dependencies by anticipating future contexts. Its capability of learning context-rich bi-directional representations has been validated in the spoken language understanding task [18], [19]. In our previous work [20], we conducted supervised pre-training with the Mask-CTC objective on the encoder and CTC modules of the triggered attention-based model, which has shown effective results for improving accuracy while reducing latency. In this work, we aim to further examine the effectiveness of Mask-CTC-based pre-training for streaming models with various architectures.

In addition, a detailed perspective on the latency can be analyzed by focusing on the timing of output spikes. In streaming ASR, the timing of output spikes is generally delayed due to the lack of long-term context information. We expect that the proposed pre-training based on Mask-CTC will introduce the look-ahead capability to the encoder, thereby enabling accurate prediction of the timing of output spikes. This property contributes to the early determination of recognition results, which is essential for streaming applications (e.g., a system

that interacts with a user in real-time).

This study, therefore, attempts to demonstrate the effectiveness of Mask-CTC-based pre-training for streaming models with different architectures, including Transformer-T and CBS-ASR, and to discuss the contribution of the proposed pre-training to latency reduction.

The rest of this paper is organized as follows. Section II overviews the streaming ASR models and the Mask-CTC model. Section III describes the proposed pre-training approach for constructing low latency and high recognition accuracy streaming ASR models. In Section IV, we demonstrate the effectiveness of Mask-CTC-based pre-training through experiments and discuss the latency reduction effect with output spike timing measurements. Finally, Section V concludes this paper.

## II. BACKGROUND

In this section, we introduce two types of streaming ASR approaches that we study in this work: Transformer-Transducer [9] and contextual block streaming ASR [13]. We also describe Mask-CTC [14], which is the key to our proposed method.

### A. Transformer-Transducer

A Transducer-based ASR model contains three components: acoustic encoder, label encoder, and joint network. Given a streaming input to a current time index  $t$ , the output probability of the  $u$ -th token is calculated as follows:

$$\mathbf{h}_t^{\text{AE}} = \text{AcousticEncoder}(\mathbf{x}_{1:t}), \quad (1)$$

$$\mathbf{h}_{u-1}^{\text{LE}} = \text{LabelEncoder}(y_{1:u-1}), \quad (2)$$

$$\mathbf{h} = \text{Tanh}(\text{Linear}(\mathbf{h}_t^{\text{AE}}) + \text{Linear}(\mathbf{h}_{u-1}^{\text{LE}})), \quad (3)$$

$$P(y_u|y_{1:u-1}, \mathbf{x}_{1:t}) = \text{SoftMax}(\mathbf{h}). \quad (4)$$

First, the acoustic encoder embeds the input sequence  $\mathbf{x}_{1:t}$  into vector  $\mathbf{h}_t^{\text{AE}}$  (Eq. (1)). Meanwhile, the label encoder generates  $\mathbf{h}_{u-1}^{\text{LE}}$  from the previous output token sequence  $y_{1:u-1}$  (Eq. (2)). The two outputs are then sent to the joint network, projected to the same dimension, and added up (Eq. (3)). Finally, the output probabilities against tokens in a vocabulary  $\mathcal{V}$  are calculated based on the previous result (Eq. (4)). The Transducer framework predicts the current symbol for each input frame based on the past output tokens, which naturally introduces streaming fashion into decoding.

Various neural network types can be applied to implement the acoustic and label encoders [6], [21], [22]. In the work of [9], Transformer [7] is applied to the acoustic encoder to achieve high accuracy and LSTM [21] for the label encoder in consideration of the model size control. Chunk-wise attention masks are applied to the self-attention layers of the Transformer acoustic encoder to enable streaming feature extraction. This architecture is referred to as a Transformer-Transducer (Transformer-T).

### B. Contextual block streaming ASR

Contextual block streaming ASR (CBS-ASR) [13] introduces streaming properties to attention-based encoder-decoder models. For streaming feature extraction in the encoder, CBS-ASR utilizes block processing with a context inheritance mechanism proposed in [23]. The speech input is segmented into blocks containing past, central, and future frames with the numbers of  $N_l$ ,  $N_c$ , and  $N_r$ . The input blocks are passed on to the encoder, where the central frames are utilized for the output with local contexts provided by the past and future frames as well as the global contexts provided by a context embedding vector inherited from the previous block. Streaming decoding is achieved by a block boundary detection (BBD) algorithm [13], which takes end-of-sentence prediction or token repetition as stopping criteria from detecting the index boundaries on-the-fly and enables the beam search synchronous to the encoded blocks. The streaming processing in CBS-ASR is calculated as follows:

$$H_b, \mathbf{c}_b = \text{BlockEncoder}(Z_b, \mathbf{c}_{b-1}), \quad (5)$$

$$\alpha(y_{0:i}, H_{1:B}) \approx \sum_{b=1}^B \sum_{j=I_{b-1}+1}^{I_b} \log p(y_i|y_{0:j-1}, H_{1:b}). \quad (6)$$

Eq. (5) represents the streaming encoding of the  $b$ -th input sequence  $Z_b$ , where  $|Z_b| = N_l + N_c + N_r$ . The encoded acoustic features  $H_b$  is obtained from  $Z_b$  and the contextual vector from the previous block  $\mathbf{c}_{b-1}$ . Eq. (6) represents the score of the partial hypothesis  $y_{0:i}$  during streaming beam search decoding, where  $y_0$  is the start-of-sequence token.  $I_b$  denotes the index boundary of the  $b$ -th input block derived from the BBD algorithm.

### C. Mask-CTC

The Mask-CTC framework [14] aims to learn feature representations suitable for anticipation of future contexts. Mask-CTC trains an encoder-decoder model with the joint CMLM [16] and CTC objectives. During training, tokens in the ground truth are randomly masked, and the masked tokens are predicted based on contextual information captured by the encoder and other unmasked output tokens. For the input  $X$  and observed tokens  $Y_{\text{obs}}$ , the output probabilities of the masked tokens  $Y_{\text{mask}}$  are computed as follows:

$$P_{\text{cmlm}}(Y_{\text{mask}}|Y_{\text{obs}}, X) = \prod_{y \in Y_{\text{mask}}} P_{\text{cmlm}}(y|Y_{\text{obs}}, X), \quad (7)$$

where  $Y_{\text{obs}}$  is  $Y \setminus Y_{\text{mask}}$ . Based on the CMLM mask prediction of the decoder, the encoder network of Mask-CTC is trained to consider the long-term bidirectional dependencies between output tokens, which enables it to generate acoustic feature representations that anticipate future information.

Such properties are desirable in streaming ASR as it allows the model to capture more future contexts with a limited look-ahead range. In such a way, the Mask-CTC framework can be a potential solution for improving streaming ASR, enhancing a model to achieve high accuracy while keeping low latency.

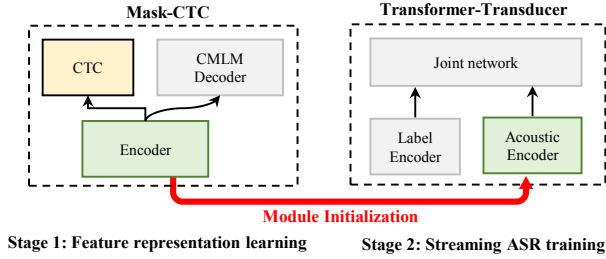


Fig. 1. Illustration of Mask-CTC-based pre-training using Transformer-Transducer model. In stage 1, encoder is trained with Mask-CTC framework. In stage 2, Transformer-Transducer model is initialized with pre-trained encoder and fine-tuned with streaming objective.

### III. MASK-CTC-BASED PRE-TRAINING METHOD

We present a simple and general Mask-CTC-based pre-training method for achieving high-accuracy and low-latency streaming ASR. Specifically, this paper aims to demonstrate the effectiveness of the Mask-CTC pre-training regardless of model architectures and discusses whether such pre-training can extract features suitable for anticipation as intended, focusing on the alignment of the output tokens.

As different end-to-end streaming ASR models, we focus on Transformer-T (see Section II-A) and CBS-ASR (see Section II-B), which cover both Transducer and encoder-decoder model architectures. For both models, the adoption of Transformer has realized high recognition accuracy in their non-streaming baselines. However, when applied to streaming scenarios, the look-ahead ranges of self-attention layers are limited from global to local. This leads to an inevitable performance drop by degrading the Transformer’s capability to capture long-range contexts, which limits applications where low latency is a top priority for recognition.

To remedy such an effect, we need the feature representation for the input sequence that considers long-term contextual dependencies and anticipates future information, which corresponds to the properties of the Mask-CTC encoder network as described in Section II-C. To introduce the desirable properties of the Mask-CTC model into the streaming ASR, we propose a simple two-step training method as follows, which is also described in Fig. 1:

- **Stage 1 (Feature representation learning):** The Mask-CTC model is pre-trained to obtain an encoder network that can consider long-term dependencies and anticipate future information.
- **Stage 2 (Streaming ASR training):** The pre-trained Mask-CTC model is exploited to initialize the streaming ASR models. For Transformer-T, the acoustic encoder with the chunk-wise attention is initialized with the Mask-CTC encoder. For CBS-ASR, both the Mask-CTC encoder and CTC networks are used to initialize the corresponding components.

With the two-step training method above, we expect to inherit the characteristics of Mask-CTC to a streaming ASR

model to capture long-term contextual information and reduce the latency dependency.

## IV. EXPERIMENTS

Speech recognition experiments were conducted to examine the effectiveness of the Mask-CTC-based pre-training method using ESPnet2 [24], [25]. We also investigated the essential effect of the proposed pre-training method by studying the output token alignments of the streaming ASR models.

### A. Datasets

The models were trained and evaluated using the Wall Street Journal (WSJ) [26] dataset, which contains 81h English utterances of read articles from the newspaper and the TED-LIUM2 (TED2) [27] dataset, which contains 207h English spontaneous speech. For the output tokens, we used SentencePiece [28] to construct a 80 subword vocabulary for WSJ and a 500 subword vocabulary for TED2, respectively. For robust model training, we applied SpecAugment [29] to the input data.

### B. Experimental setup

For the Transformer-T model, the acoustic encoder was implemented with 12 Transformer encoder layers and a single LSTM layer for the label encoder. For streaming feature extraction, a chunk-wise attention mask was implemented and applied to the encoder layers as in [9]. The latency value was calculated as the product of the maximum look-ahead range (i.e., chunk size  $- 1$ ) and a frame rate of 40ms.

For WSJ experiments, the CBS-ASR model consisted of 6 Conformer encoder layers [30] and 6 Transformer decoder layers. The input block settings followed  $N_l$  as eight,  $N_c$  as four, and  $N_r$  varying from 0 to 6. The latency for CBS-ASR was calculated as the product of the maximum look-ahead range in the block (i.e.,  $N_c + N_r - 1$ ) and a frame rate of 40ms. For TED2 experiments, the CBS-ASR model consisted of 12 Conformer encoder layers [30] and 6 Transformer decoder layers. The  $N_r$  was set to 6.

For the pre-trained Mask-CTC model, the encoder was constructed with the identical setting as the target streaming model. The CMLM decoder was built with six Transformer decoder layers. All the models were trained by 150 epochs, and the final models were obtained by averaging the snapshots of the ten epochs of the minimal loss for Transformer-T and the best accuracy for CBS-ASR. For decoding, a beam search was conducted with a beam size of ten for all. We used the word error rate (WER) for measuring the ASR performance.

### C. Experimental results

For both the Transformer-T and CBS-ASR systems, the performances of the following models are compared.

- **Baseline** [9], [13], [31]: Existing streaming ASR models, including Transformer-T and CBS-ASR. The parameters for all the components were randomly initialized.
- **Enhanced:** Streaming ASR models with Mask-CTC-based pre-training. Components of the streaming ASR were initialized with pre-trained Mask-CTC modules. For

TABLE I  
WORD ERROR RATES ON WSJ DATASET.

Model	Latency [ms]	Initialization	WER [%] ( $\downarrow$ )	
			eval92	dev93
<b>Baseline</b>				
Transformer-T	120	Random	19.5	23.3
	160		16.8	20.9
	200		<b>15.1</b>	<b>18.9</b>
	$\infty$	Random	14.7	17.3
CBS-ASR	200	Random	14.4	18.1
	280		13.2	16.2
	360		<b>12.9</b>	<b>16.1</b>
	1240	Random	11.2	14.2
<b>Enhanced</b>				
Transformer-T	120	Mask-CTC	16.6	20.8
	160		15.0	19.0
	200		<b>14.8</b>	<b>18.5</b>
CBS-ASR	200	Mask-CTC	13.5	17.2
	280		12.9	<b>16.0</b>
	360		<b>12.2</b>	16.1

TABLE II  
WORD ERROR RATE ON TED2 DATASET.

Model	Latency [ms]	Initialization	WER [%] ( $\downarrow$ )
<b>Baseline</b>			
CBS-ASR	280	Random	11.3
	1240	Random	9.8
<b>Enhanced</b>			
CBS-ASR	280	Mask-CTC	11.1

Transformer-T, the acoustic encoder was initialized with the Mask-CTC encoder. For CBS-ASR, both encoder and CTC modules were initialized with corresponding Mask-CTC modules.

The experimental results of Transformer-T and CBS-ASR are summarized in Table I and Table II. Non-streaming Transformer-T and CBS-ASR with 1240ms latency were used as lower bounds in the experiments.

The results on WSJ show that for both Transformer-T and CBS-ASR, the enhanced models outperformed the baseline models by achieving lower WERs under all latency settings, suggesting the accuracy enhancements introduced by the Mask-CTC-based pre-training method. For WSJ dataset, 40ms and 80ms latency reductions were reached for Transformer-T and CBS-ASR, respectively, while achieving better or equal recognition accuracy than the baseline models. For instance, the enhanced Transformer-T with 120ms latency achieved lower WERs (16.6% for eval92 and 20.8% for dev93) than the WERs of the baseline with 160ms latency (16.8% for eval92 and 20.9% for dev93). Such results demonstrated that our method contributed to the construction of streaming ASR

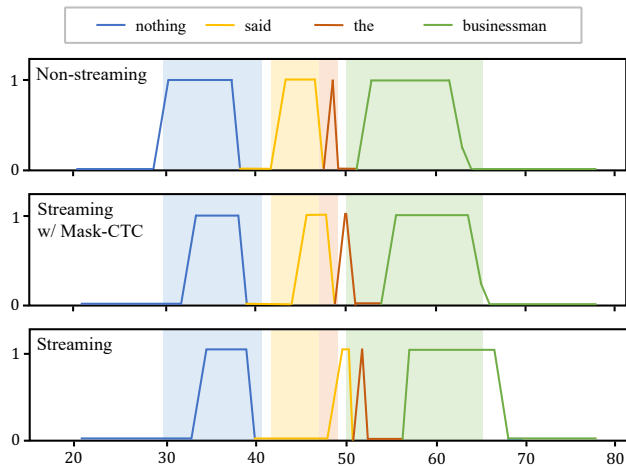


Fig. 2. Output token alignments of non-streaming and streaming Transformer-Transducer models.

models with low latency and high accuracy. For TED2 dataset, the enhanced CBS-ASR model also achieved 0.2 percentage point of WER reduction compared to the baseline model, which proves the general effectiveness of the proposed method regardless of the dataset. The results for systems with different architectures, such as Transducer and Encoder-Decoder, also demonstrated that the Mask-CTC-based pre-training was effective regardless of the model architecture.

#### D. Analysis of output token alignments

The work of [32] argued that the streaming model attends to shift the token boundaries to the future side to obtain more contextual information, which results in delay of the posterior probability spikes for the output tokens compared to non-streaming models. In contrast, if the encoder network learns the feature representations that anticipate future information, the output tokens can be confirmed earlier and the token boundary shifting issue should be remedied in some instances. Therefore, we measured the delay of the spike occurrences in streaming models by comparing them to the alignments obtained from a non-streaming model. The delay is expected to be reduced with the Mask-CTC-based pre-training method.

We conducted measurements on the dev93 validation set of WSJ. We used the baseline and enhanced models with 200ms latency settings for Transformer-T and compared their alignments with a non-streaming Transformer-T model. The alignments were obtained from the output of the joint network. For CBS-ASR, the latency was also set to 200ms, and we compared the output token boundaries between the baseline and enhanced models. The ASR alignments were obtained from the CTC predictions of CBS-ASR in the same manner as [32] and the reference alignments were obtained with the Montreal Forced Aligner [33]. Figure 2 illustrates one example of output token alignments given by Transformer-T. Here, the color in the background represents the reference alignment to the speech input. The non-streaming ASR (top) managed to predict accurate token alignments. However, the

baseline streaming ASR (bottom) showed a significant delay in the alignments, indicating token boundary shifting due to the lack of contexts. Meanwhile, our enhanced streaming ASR (middle), with a Mask-CTC-based pre-trained encoder network, largely improved the alignments of the streaming ASR. We calculated the average output delay reduction across the dev93 validation set for both Transformer-T and CBS-ASR. For Transformer-T, the spike output delay was reduced by 44ms, and for CBS-ASR, 46ms. Such results help us to understand the knowledge learned from the Mask-CTC-based pre-training method and the reason for the latency reduction capability.

## V. CONCLUSION

In this study, an attempt was made to demonstrate the effectiveness of Mask-CTC-based pre-training for achieving low latency and high accuracy in streaming speech recognition. Experimental results showed the effectiveness of the method on various model architectures, including Transformer-Transducer and contextual block streaming ASR. Furthermore, by studying the output spike timings of the streaming models, we discovered that more precise alignments of the input and output sequences are learnt by the pre-training, which contributes to the latency reduction in streaming ASR.

## REFERENCES

- [1] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al., “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [2] Jan Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio, “Attention-based models for speech recognition,” in *Proc. NeurIPS*, 2015.
- [3] Alex Graves and Navdeep Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *Proc. ICML*, 2014, pp. 1746–1772.
- [4] William Chan, Navdeep Jaitly, Quoc V. Le, and Oriol Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proc. ICASSP*, 2016, pp. 4960–4964.
- [5] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proc. ICML*, 2006, pp. 369–376.
- [6] Alex Graves, “Sequence transduction with recurrent neural networks,” *ArXiv*, vol. abs/1211.3711, 2012.
- [7] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Proc. NeurIPS*, 2017, pp. 5998–6008.
- [8] Qian Zhang, Han Lu, Hasim Sak, Anshuman Tripathi, Erik McDermott, Stephen Koo, and Shankar Kumar, “Transformer Transducer: A streamable speech recognition model with transformer encoders and RNN-T loss,” in *Proc. ICASSP*, 2020, pp. 7829–7833.
- [9] Xie Chen, Yuehua Wu, Zhenghao Wang, Shujie Liu, and Jinyu Li, “Developing real-time streaming Transformer Transducer for speech recognition on large-scale dataset,” in *Proc. ICASSP*, 2021, pp. 5904–5908.
- [10] Linhao Dong, Shuang Xu, and Bo Xu, “Speech-Transformer: A no-recurrence sequence-to-sequence model for speech recognition,” in *Proc. ICASSP*, 2018, pp. 5884–5888.
- [11] Shigeki Karita, Nelson Yalta, Shinji Watanabe, M. Delcroix, A. Ogawa, and T. Nakatani, “Improving Transformer-based end-to-end speech recognition with connectionist temporal classification and language model integration,” in *Proc. INTERSPEECH*, 2019.
- [12] Niko Moritz, Takaaki Hori, and Jonathan Le, “Streaming automatic speech recognition with the Transformer model,” in *Proc. ICASSP*, 2020, pp. 6074–6078.
- [13] Emiru Tsunoo, Yosuke Kashiwagi, and Shinji Watanabe, “Streaming Transformer ASR with blockwise synchronous beam search,” in *Proc. SLT*, 2021, pp. 22–29.
- [14] Yosuke Higuchi, Shinji Watanabe, Nanxin Chen, Tetsuji Ogawa, and Tetsunori Kobayashi, “Mask CTC: Non-autoregressive end-to-end ASR with CTC and mask predict,” in *Proc. INTERSPEECH*, 2020, pp. 3655–3659.
- [15] Yosuke Higuchi, Hirofumi Inaguma, Shinji Watanabe, Tetsuji Ogawa, and Tetsunori Kobayashi, “Improved Mask-CTC for non-autoregressive end-to-end ASR,” in *Proc. ICASSP*, 2021, pp. 8363–8367.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “BERT: Pre-training of deep bidirectional Transformers for language understanding,” in *Proc. NAACL-HLT*, 2019.
- [17] Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer, “Mask-predict: Parallel decoding of conditional masked language models,” in *Proc. EMNLP-IJCNLP*, 2019, pp. 6114–6123.
- [18] Mohan Li and Rama Doddipatla, “Non-autoregressive end-to-end approaches for joint automatic speech recognition and spoken language understanding,” *Proc. SLT*, pp. 390–397, 2023.
- [19] Quentin Meeus, Marie-Francine Moens, and Hugo Van hamme, “Bidirectional representations for low resource spoken language understanding,” *ArXiv*, vol. abs/2211.14320, 2022.
- [20] Huaibo Zhao, Yosuke Higuchi, Tetsuji Ogawa, and Tetsunori Kobayashi, “An investigation of enhancing CTC model for triggered attention-based streaming ASR,” in *Proc. APSIPA ASC*, 2021, pp. 477–483.
- [21] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, pp. 1735–1780, 1997.
- [22] Yanzhang He, Tara N. Sainath, Rohit Prabhavalkar, Ian McGraw, Raziel Álvarez, Ding Zhao, David Rybach, Anjali Kannan, Yonghui Wu, Ruoming Pang, Qiao Liang, Deepti Bhatia, Yuan Shangquan, Bo Li, Golan Pundak, Khe Chai Sim, Tom Bagby, Shuo yiin Chang, Kanishka Rao, and Alexander Gruenstein, “Streaming end-to-end speech recognition for mobile devices,” in *Proc. ICASSP*, 2019, pp. 6381–6385.
- [23] Emiru Tsunoo, Yosuke Kashiwagi, Toshiyuki Kumakura, and Shinji Watanabe, “Transformer ASR with contextual block processing,” in *Proc. ASRU*, 2019, pp. 427–433.
- [24] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai, “ESPnet: End-to-end speech processing toolkit,” in *Proc. INTERSPEECH*, 2018, pp. 2207–2211.
- [25] Florian Boyer, Yusuke Shinohara, Takaaki Ishii, Hirofumi Inaguma, and Shinji Watanabe, “A study of Transducer based end-to-end ASR with ESPnet: Architecture, auxiliary loss and decoding strategies,” *Proc. ASRU*, pp. 16–23, 2021.
- [26] Douglas B Paul and Janet Baker, “The design for the wall street journal-based CSR corpus,” in *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York*, 1992.
- [27] A. Rousseau, P. Deléglise, and Y. Estève, “Enhancing the TED-LIUM corpus with selected data for language modeling and more TED talks,” in *Proc. LREC*, 2014.
- [28] Taku Kudo and John Richardson, “Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” *arXiv preprint arXiv:1808.06226*, 2018.
- [29] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin Dogus Cubuk, and Quoc V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *ArXiv*, vol. abs/1904.08779, 2019.
- [30] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang, “Conformer: Convolution-augmented Transformer for speech recognition,” in *Proc. INTERSPEECH*, 2020, pp. 5036–5040.
- [31] Niko Moritz, Takaaki Hori, and Jonathan Le Roux, “Triggered attention for end-to-end speech recognition,” in *Proc. ICASSP*, 2019, pp. 5666–5670.
- [32] Hirofumi Inaguma and Tatsuya Kawahara, “Alignment knowledge distillation for online streaming attention-based speech recognition,” *ArXiv*, vol. abs/2103.00422, 2021.
- [33] Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using kaldif,” in *Interspeech*, 2017.