



# Improved Training for End-to-End Streaming Automatic Speech Recognition Model with Punctuation

*Hanbyul Kim, Seunghyun Seo, Lukas Lee, Seolki Baek*

NAVER Cloud, South Korea

{hanbyul.kim, real.seunghyun.seo, lukas.lee, seolki.baek}@navercorp.com

## Abstract

Punctuated text prediction is crucial for automatic speech recognition as it enhances readability and impacts downstream natural language processing tasks. In streaming scenarios, the ability to predict punctuation in real-time is particularly desirable but presents a difficult technical challenge. In this work, we propose a method for predicting punctuated text from input speech using a chunk-based Transformer encoder trained with Connectionist Temporal Classification (CTC) loss. The acoustic model trained with long sequences by concatenating the input and target sequences can learn punctuation marks attached to the end of sentences more effectively. Additionally, by combining CTC losses on the chunks and utterances, we achieved both the improved F1 score of punctuation prediction and Word Error Rate (WER).

**Index Terms:** speech recognition, streaming ASR, punctuation prediction

## 1. Introduction

Advancements in deep neural networks have led to substantial improvements in Automatic Speech Recognition (ASR) systems. These systems are extensively utilized in various applications, such as voice assistants and dictation software. In recent years, Transformer models [1], self-attention-based neural networks, have exhibited remarkable performance to extract high-level speech representations in ASR. Self-attention mechanism utilized in these models exploits contextual information for acoustic encoding by calculating attention weights for each pair of input features, allowing the modeling of temporal dependencies within a sequence. However, due to its requirement of having the entire input sequence available before processing, the limitation prevents it from being a suitable encoder for streaming ASR applications, which require the real-time recognition of spoken words shortly after they are uttered. This study focuses on Connectionist Temporal Classification (CTC) based streaming models, despite the existence of alternative methods such as RNN-T (Recurrent Neural Network Transducer) [2, 3] or encoder-decoder [4, 5] based models.

Several approaches can be employed to enable the utilization of transformer-based models for streaming scenarios. Time-restricted self-attention layer [6] was proposed to limit the number of look-ahead frames, alleviating the requirement of complete sequences. However, as the receptive field on the future frames scales linearly with the number of transformer layers, it results in significant latency for a deep architecture. An alternative approach is a chunk-based method, which divides the input signal into shorter segments and feeds them to the model sequentially. This method reduces the computational cost and latency compared to the full-sequence model while

maintaining comparable performance. Recent studies [7, 8, 9] have shown that the chunk-based method generally outperforms time-restricted self-attention regarding both accuracy and latency reduction, due to its ability to leverage contextual information within each segment better.

Many studies have attempted to enhance the accuracy of speech recognition systems, but relatively few have given sufficient attention to punctuation accuracy. Punctuation marks play a crucial role in enhancing the readability and comprehension of the transcriptions. Moreover, the accuracy of punctuation prediction can have significant implications for downstream natural language processing tasks. Previous studies [10, 11, 12] have assumed a cascaded system of two separate models for ASR and punctuation prediction. However, incorporating a punctuation model, such as BERT [13] into ASR systems poses several challenges. Firstly, pre-trained BERT can perform poorly because ASR output is fed to BERT, where ASR output has acoustic noises that BERT has never seen in training time. In addition, cascaded systems might introduce additional latency in streaming scenarios. Punctuation models generally cannot make predictions until the entire input sequence is received. Moreover, the incorporation of large neural networks in these systems may result in additional computational burdens. Due to these reasons, the end-to-end approach for formatted speech recognition is gaining attention from researchers [14, 15], but studies for streaming models still lack.

The task of end-to-end punctuation prediction in streaming ASR presents a unique challenge compared to non-streaming ASR. Most non-streaming ASR models are trained to accept full utterances with speech segments provided by additional modules, such as a Voice Activity Detection (VAD) model. Some key punctuation marks, such as period and question, are typically located at the end of sentences. Therefore, proper segmentation of the speech input makes the prediction of these key punctuation marks relatively straightforward, as reported in the previous study [15]. Nonetheless, recent studies about end-to-end VAD approach [16] or punctuation prediction with acoustic features [17, 18] proposed joint training approaches using both acoustic and semantic information.

In this paper, we propose a novel training method to enhance the accuracy of end-to-end streaming ASR models with punctuation prediction. Our approach is specifically designed for streaming ASR, where we utilize the transformer model trained with CTC loss. However, as we described above, predicting punctuation directly with a streaming ASR model poses challenges that are not present in non-streaming ASR. To address the issues, we present a technique that involves feeding concatenated speech mini-batches to a chunk-based transformer, which allows the model to learn the positions of punctuation marks in the middle of the sequence, rather than at the end.

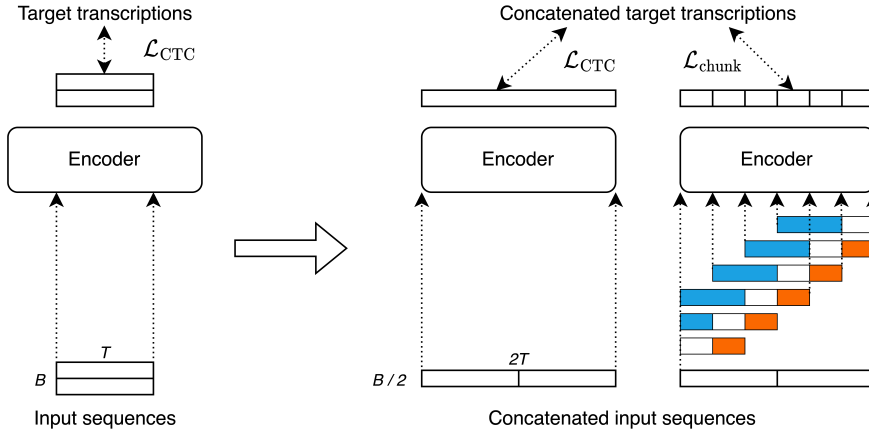


Figure 1: Overview of the proposed method for end-to-end streaming ASR with punctuation prediction. The approach concatenates input speech sequences and target text sequences for a CTC model to learn punctuation marks between the target labels. This doubles the number of frames while reducing the batch size by half. An additional CTC loss term is introduced, which is computed on the concatenated emissions of separated chunks. The blue and orange boxes in the figure represent past and future frames respectively, providing contextual information for a input sequence. The ASR model is trained by jointly optimizing CTC losses on full sequences and chunks.

To ensure stable speech recognition performance, we introduce an auxiliary loss that focuses on the chunk-hopping method. Our experiments demonstrate that this training method improves punctuation prediction in an end-to-end manner, while also enhancing speech recognition accuracy in streaming scenarios.

## 2. Related Work

### 2.1. Chunk-based streaming ASR

Chunk-based streaming ASR is an approach that aims to reduce latency in ASR systems. This method involves splitting the speech into chunks and processing each chunk using the ASR model, without waiting for the entire utterance to be completed [19, 20]. This allows the system to begin producing output in real time, providing users with faster and more efficient speech recognition capabilities. The chunk-based approach has been further developed to include context-sensitive chunks [21], where each chunk is encoded with additional left and right frames to provide the speech context. This approach has been explored in many studies [22, 7, 8], and found to be particularly effective when used with transformer-based ASR models.

### 2.2. Joint training in streaming ASR

ASR models with limited context can learn from a broader range of contextual information and improve their overall recognition performance by incorporating auxiliary loss functions in addition to the primary objective. This approach may be particularly effective for streaming ASR, as it only increases computational burden during the training phase and enables the model to maintain low latency during recognition.

For encoder-decoder models, the online CTC/attention mechanism [23, 24, 25] is a technique that combines both the attention loss from the decoder and the CTC loss from the encoder. Additionally, a “soft forgetting” approach [26, 8] aims to minimize the mean squared error between the full utterances and the chunk-based encoder output to improve the model’s per-

formance on streaming ASR tasks. The integration of streaming and non-streaming ASR systems by weight sharing and joint training a streaming model with a full-context model [27, 28] has emerged as a recent approach for improving the accuracy and low-latency performance of streaming ASR.

## 3. Proposed Method

In this paper, we propose an ASR encoder that can directly produce punctuated text from an acoustic input sequence. We utilized a Transformer encoder trained with CTC loss, which is suitable for streaming ASR. The CTC objective, first proposed by Graves, et al. [29] is a loss function that is commonly used in speech recognition tasks. It is designed to solve the problem of aligning input sequences and output sequences with different lengths in speech recognition due to the variability in speech duration and pronunciation.

Given a sequence of speech inputs  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$  and transcriptions  $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$ , the objective computes the likelihood of all possible alignments using representation vectors from ASR model and target sequence vectors.

$$P_{CTC}(\mathbf{y}_i|\mathbf{x}_i) = \sum_{\pi \in \beta^{-1}(\mathbf{y}_i)} P(\pi|\mathbf{x}_i)$$

$$\mathcal{L}_{CTC}(\mathbf{X}, \mathbf{Y}) = - \sum_{i=1}^N \log P_{CTC}(\mathbf{y}_i|\mathbf{x}_i)$$

where  $\beta^{-1}(y)$  is the set of all possible alignments between input sequences and target sequences.

The proposed method involves modifying the input sequence to enable the model to effectively learn to predict punctuation marks. First, given a current mini-batch, a new batch is created by concatenating two neighboring speech inputs and labels each from the current batch. In most publicly available speech recognition datasets, punctuation marks such as periods and question marks are conventionally placed at the end of sentences. This may lead to inconsistencies, as punctuation marks can appear at any position within a chunk of speech data.

It should be noted that the adjacent speech inputs being concatenated would be semantically unrelated because we use randomly shuffled mini-batches.

$$\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_{N/2}) = (\text{cat}[\mathbf{x}_1, \mathbf{x}_2], \dots, \text{cat}[\mathbf{x}_{N-1}, \mathbf{x}_N])$$

$$\tilde{\mathbf{Y}} = (\tilde{\mathbf{y}}_1, \dots, \tilde{\mathbf{y}}_{N/2}) = (\text{cat}[\mathbf{y}_1, \mathbf{y}_2], \dots, \text{cat}[\mathbf{y}_{N-1}, \mathbf{y}_N])$$

where  $\text{cat}[\cdot]$  is a function that concatenates the two inputs in the time direction. From this concatenated batch, the CTC posterior probability can be computed.

The new batch tensor is split into multiple chunks in the time dimension, which is then padded with the left and right chunks of 2s and 1s respectively. This is called the context-sensitive chunk technique to provide contextual information to chunks [26, 8, 28]. The ASR model takes each padded chunk as input separately and outputs representation vectors of each. Among CTC emissions from each chunk, past and future frames are discarded and the remaining emissions are merged to obtain the final posterior which has the same length as the concatenated sequences.

$$P_{\text{chunk}}(\tilde{\mathbf{y}}_i | \tilde{\mathbf{x}}_i) = \sum_{\pi \in \beta^{-1}(\mathbf{y}_i)} P(\pi | \tilde{\mathbf{x}}_{i1}, \dots, \tilde{\mathbf{x}}_{ic})$$

where  $c$  is the number of chunks and  $\tilde{\mathbf{x}}_{i1}, \dots, \tilde{\mathbf{x}}_{ic}$  are chunks of concatenated utterance. Note that, unlike local self-attention, this method does not suffer from the disadvantage of a larger receptive field as the layer gets deeper.

From each posterior probability,  $P_{\text{CTC}}$  and  $P_{\text{chunk}}$ , corresponding CTC losses  $\mathcal{L}_{\text{CTC}}$  and  $\mathcal{L}_{\text{chunk}}$  can be computed. Finally, the total loss is obtained by interpolating two losses.

$$\mathcal{L}_{\text{total}}(\mathbf{X}, \mathbf{Y}) = (1 - \lambda)\mathcal{L}_{\text{CTC}}(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}) + \lambda\mathcal{L}_{\text{chunk}}(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}})$$

The detailed training procedure is outlined in Algorithm 1.

---

**Algorithm 1** Pseudo-code of the proposed method

---

```

procedure TRAINSTEP(X, Y)
   $X_{\text{concat}} \leftarrow \text{CONCAT}(X)$     ▷ Concatenate input speech
   $Y_{\text{concat}} \leftarrow \text{CONCAT}(Y)$     ▷ Concatenate target labels

   $P_{\text{CTC}} \leftarrow \text{ENCODER}(X_{\text{concat}})$ 
   $L_{\text{CTC}} \leftarrow \text{CTC}(P_{\text{CTC}}, Y_{\text{concat}})$ 

   $[X_{\text{split}}]_1^N \leftarrow \text{SPLIT}(X_{\text{concat}})$     ▷ Split input speech into
  non-overlapping chunks
   $[X_{\text{pad}}]_1^N \leftarrow \text{PADCONTEXT}([X_{\text{split}}]_1^N)$     ▷ Pad the past
  and future frames for context
   $[P_{\text{pad}}]_1^N \leftarrow \text{ENCODER}([X_{\text{pad}}]_1^N)$ 
   $[P_{\text{chunk}}]_1^N \leftarrow \text{REMOVECONTEXT}([P_{\text{pad}}]_1^N)$     ▷ Remove
  the padded frames
   $P_{\text{merge}} \leftarrow \text{MERGE}([P_{\text{chunk}}]_1^N)$ 
   $L_{\text{chunk}} \leftarrow \text{CTC}(P_{\text{merge}}, Y_{\text{concat}})$ 
   $L_{\text{total}} \leftarrow (1 - \lambda) * L_{\text{CTC}} + \lambda * L_{\text{chunk}}$ 
  BACKWARD( $L_{\text{total}}$ )
end procedure

```

---

## 4. Experiments

### 4.1. Dataset

MuST-C dataset [30] was utilized for both training and evaluating the efficacy of our approach. MuST-C is a multilingual

speech translation corpus that includes at least 385 hours of audio recordings from English TED Talks for each of the eight target languages. We obtained data from the English-German speech translation dataset by extracting both English speech and its corresponding script, including punctuation.

The following processing steps were performed on the data labels. All labels were converted to lowercase to maintain consistency and reduce complexity. Punctuation marks, with the exception of apostrophes('), commas(,), periods(.), and question marks(?), were removed to simplify the text while preserving contextual information. Environmental labels enclosed in brackets (e.g. (Applause), (Laughter)) and speaker names followed by a colon (e.g. MJ:) were removed to prevent the ASR model from predicting them as output. The original train and dev sets were used in the same configuration for training and validation, respectively.

### 4.2. Training Setup

The training process for the proposed ASR model was conducted using the fairseq [31] framework. The ASR model architecture consisted of a Transformer with 12 encoder layers, 256 hidden dimensions, and 4 attention heads. We added subsampling layers to downsample the input features. The layers consisted of two convolutional layers with the first layer having an input channel size of 80 and an output channel size of 1024, and the second layer having an input channel size of 512 and an output channel size of 512. Both layers had a kernel size of 5 and a stride of 2. For streaming ASR, it is natural to use relative position encoding instead of absolute position encoding. Hence, we replaced the absolute position encoding layer with a convolutional positional encoding layer. The convolutional position encoding layer was implemented using a conv1d layer with a kernel size of 128, 256 output channels, and 16 groups. This implementation is consistent with the approach used in [32]. The total number of parameters in the model is approximately 18 million.

During training, the CTC loss with auxiliary chunk loss was used as the loss function, with interpolating factor  $\lambda$  set to 0.5. The input feature consisted of 80-dimensional log-mel filterbanks. The label data was encoded with 2000 wordpieces, which were trained using the label data of the training set with SentencePiece [33]. To optimize training speed and reduce memory consumption, the model was trained with half-precision floating-point (FP16). Adam [34] was chosen as the optimization algorithm for this study. The learning rate was set to 0.001, and the learning rate scheduler was the inverse square root. The warm-up steps were set to 10000, and the model was trained for 190 epochs, with approximately 100,000 training steps. The maximum token size was set to 40000 tokens, corresponding to a batch size of approximately 400. The training was conducted using NVIDIA A100 8 GPUs, each equipped with 80 GB of VRAM.

### 4.3. Evaluation

In this study, we used ‘‘tst-COMMON’’ subset of MuST-C dataset to evaluate our approach. The dataset is based on long English TED-talk recordings that consisted of audio transcriptions with corresponding timestamps. This enabled us to evaluate the proposed method on both long-form utterances using the original audio files, as well as short-form utterances by splitting the speech data according to the timestamps. The average duration of complete utterances was found to be 11 minutes, while the average duration of the individual segments obtained after

Model	WER (%)	Precision score (%)				Recall score (%)				F1 score (%)			
		,	.	?	avg.	,	.	?	avg.	,	.	?	avg.
Baseline	18.2	40.4	73.5	52.7	55.5	66.6	23.6	25.1	38.4	50.3	35.8	34.0	40.0
+ interCTC [15]	18.5	36.3	73.2	50.6	53.4	69.2	21.2	19.1	36.5	47.6	32.8	27.8	36.1
Proposed	15.8	39.0	76.0	53.5	56.2	72.8	33.6	29.4	45.3	50.8	46.6	37.9	45.1

Table 1: WER and punctuation prediction accuracy on long-form utterances in MuST-C testset

splitting them was 5.8 seconds.

We consisted of a chunk size of 1s with 2s of padding for past and 1s for future frames, consistent with the model training, during the chunk-encoding process. After computing the emission from the transformer encoder, we performed CTC beam search with a beam size of 5, without using any additional language model. For evaluating ASR accuracy, we measured the word error rate (WER), excluding any punctuation text. Additionally, we measured precision, recall and F1 score for each punctuation mark to evaluate the punctuation accuracy. Due to ASR errors, we first aligned the ASR prediction and ground truth and then calculated the F1 score.

#### 4.4. Experimental Results

Table 1 shows WER results for streaming speech recognition accuracy. The proposed method outperforms the baseline method by 13% relatively (18.2 vs. 15.8). In addition, we compared the punctuation recognition accuracy, and the proposed method demonstrated 12.5% higher average F1 scores across all punctuation marks. Further analysis of precision and recall reveals that the proposed method significantly improves the recall score. The recall of the period prediction shows a significant improvement with a nearly 50% increase (23.6 vs. 33.6). This suggests that our method better predicts the end of a sentence. We compared our proposed approach with the existing method [15] suggested previously. To reproduce the approach, we introduced an intermediate CTC loss with unpunctuated labels, which was jointly trained with the CTC loss. However, the proposed approach does not show any improvement in either WER or punctuation scores in streaming ASR.

Table 2 presents the evaluation results of previous studies and the proposed method on segmented speech with shorter lengths. Overall, the WER slightly degrades across all approaches. Despite the significant drop in F1 score on period mark compared to long-form decoding, the proposed approach demonstrates superior performance with the highest average punctuation prediction accuracy among all approaches.

#### 4.5. Analysis

In this section, we analyzed the factors that contributed to the effectiveness of our proposed approach. To determine the impact of specific factors, we re-trained our model with various modifications to our original configuration. Table 3 illustrates the effects of the modifications.

First, we removed the chunk CTC loss ( $\mathcal{L}_{\text{chunk}}$ ) and trained the model again with concatenated input sequences only. We found that WER increased significantly, indicating that proper regularization with the chunk CTC loss is crucial to prevent accuracy degradation. Additionally, we observed that our proposed method significantly boosted the accuracy of period marks, exceeding the improvement observed with the proposed approach. However, this has a negative effect on commas and question marks, leading to a lower average F1 score compared

to our method. From these observations, we conclude that chunk CTC loss not only regulates the CTC model to prevent WER degradation but also prevents period marks over-fitting while maintaining accuracy for other punctuation marks.

Model	WER (%)	Punctuation F1 score (%)			
		,	.	?	avg.
Baseline	18.5	47.7	16.0	39.4	34.4
+ interCTC [15]	18.9	44.6	12.5	29.1	28.7
Proposed	16.3	52.5	27.3	44.0	41.2

Table 2: Performance evaluation of WER and punctuation prediction accuracy on speech segments in MuST-C testset

Model	WER (%)	Punctuation F1 score (%)			
		,	.	?	avg.
Proposed	15.8	50.8	46.6	37.9	45.1
- $\mathcal{L}_{\text{chunk}}$	21.3	45.0	52.9	25.8	41.2
- concat.	16.0	50.1	41.0	34.0	41.7

Table 3: Ablation study of the effect of auxiliary loss term on chunk-based output and input sequence concatenation

We also conducted an experiment to assess the impact of excluding the input sequence concatenation in our training approach. The results show that the overall F1 scores for punctuation marks drop, which suggests that concatenating input sequences plays a crucial role in improving the punctuation recognition performance of our approach. Moreover, an interesting observation was that removing the concatenation also degraded the WER slightly. This could be due to the fact that the CTC loss computed on concatenated sequences also regularizes the joint training approach, indicating the importance of the regularization effect. These results highlight the effectiveness of our training method which improves punctuation recognition accuracy while maintaining the streaming ASR performance.

## 5. Conclusion

This work proposed a novel training method to improve both speech recognition and punctuation prediction performances in the streaming ASR task. The experimental results showed that the concatenation of input sequences played a significant role in improving the punctuation recognition performance of the proposed method. The chunk-based loss was crucial not only for preventing WER degradation by input concatenation but also for preventing period mark over-fitting while maintaining accuracy for other punctuation marks. While our approach does not introduce additional computation burden during inference, the training time and space complexity increased due to the chunk-wise forward steps. Further experimentation is required to optimize the training computation for our approach. We also would like to investigate the applicability of the proposed approach to other ASR models and datasets.

## 6. References

- [1] A. Vaswani, N. Shazeer, N. Parmar, and et al., “Attention is all you need,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [2] P. R. He Yanzhang, Sainath Tara N and et al., “Streaming end-to-end speech recognition for mobile devices,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6381–6385.
- [3] S. H. Zhang Qian, Lu Han and et al., “Transformer transducer: A streamable speech recognition model with transformer encoders and rnn-t loss,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [4] 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.
- [5] H. T. Moritz Niko and J. Le, “Streaming automatic speech recognition with the transformer model,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [6] D. Povey, H. Hadian, P. Ghahremani, and et al., “A time-restricted self-attention layer for asr,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [7] H. Miao, G. Cheng, C. Gao, and et al., “Transformer-based on-line ctc/attention end-to-end speech recognition architecture,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [8] K. An, H. Xiang, and Z. Ou, “CAT: A CTC-CRF Based ASR Toolkit Bridging the Hybrid and the End-to-End Approaches Towards Data Efficiency and Low Latency,” *Proc. Interspeech 2020*, 2020.
- [9] C. Wu, Y. Wang, Y. Shi, C.-F. Yeh, and F. Zhang, “Streaming Transformer-Based Acoustic Models Using Self-Attention with Augmented Memory,” *Proc. Interspeech 2020*, 2020.
- [10] O. Tilk and T. Alumäe, “Bidirectional recurrent neural network with attention mechanism for punctuation restoration,” *Proc. Interspeech 2016*, 2016.
- [11] M. Fang, H. Zhao, X. Song, and et al., “Using bidirectional lstm with bert for chinese punctuation prediction,” in *2019 IEEE International Conference on Signal, Information and Data Processing (ICSIDP)*, 2019.
- [12] P. V. Piyush Behre, Sharman Tan and S. Chang, “Streaming punctuation for long-form dictation with transformers,” in *Proceedings of the 8th International Conference on Signal, Image Processing and Embedded Systems*, 2022.
- [13] J. Devlin, M.-W. Chang, K. Lee, and et al., “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019.
- [14] M. Mimura, S. Sakai, and T. Kawahara, “An end-to-end model from speech to clean transcript for parliamentary meetings,” in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2021.
- [15] J. Nozaki, T. Kawahara, K. Ishizuka, and et al., “End-to-end Speech-to-Punctuated-Text Recognition,” *Proc. Interspeech 2022*, 2022.
- [16] W. R. Huang, S. yiin Chang, D. Rybach, and et al., “E2E Segmenter: Joint Segmenting and Decoding for Long-Form ASR,” *Proc. Interspeech 2022*, 2022.
- [17] H. Christensen, Y. Gotoh, and S. Renals, “Punctuation annotation using statistical prosody models,” in *ISCA Tutorial and Research Workshop (ITRW) on Prosody in Speech Recognition and Understanding*, 2001.
- [18] M. Sunkara, S. Ronanki, D. Bekal, and et al., “Multimodal semi-supervised learning framework for punctuation prediction in conversational speech,” *Proc. Interspeech 2020*, 2020.
- [19] X. Li and X. Wu, “Constructing long short-term memory based deep recurrent neural networks for large vocabulary speech recognition,” in *ICASSP 2015-2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [20] A. Zeyer, R. Schlüter, and H. Ney, “Towards online-recognition with deep bidirectional lstm acoustic models,” *Proc. Interspeech 2022*, 2016.
- [21] K. Chen and Q. Huo, “Training deep bidirectional lstm acoustic model for lvcsr by a context-sensitive-chunk bptt approach,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 7, 2016.
- [22] L. Dong, F. Wang, and B. Xu, “Self-attention aligner: A latency-control end-to-end model for asr using self-attention network and chunk-hopping,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019.
- [23] N. Moritz, T. Hori, and J. L. Roux, “Streaming end-to-end speech recognition with joint ctc-attention based models,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019.
- [24] H. Miao, G. Cheng, P. Zhang, and et al., “Online hybrid ctc/attention end-to-end automatic speech recognition architecture,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, 2020.
- [25] S. Zhang, Z. Gao, H. Luo, and et al., “Streaming Chunk-Aware Multihead Attention for Online End-to-End Speech Recognition,” *Proc. Interspeech 2020*, 2020.
- [26] K. Audhkhasi, G. Saon, Z. Tüske, and et al., “Forget a Bit to Learn Better: Soft Forgetting for CTC-Based Automatic Speech Recognition,” *Proc. Interspeech 2019*, 2019.
- [27] J. Yu, W. Han, A. Gulati, and et al., “Dual-mode ASR: unify and improve streaming ASR with full-context modeling,” in *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.
- [28] K. An, H. Zheng, Z. Ou, and et al., “CUSIDE: chunking, simulating future context and decoding for streaming ASR,” *Proc. Interspeech 2022*, 2022.
- [29] A. Graves, S. Fernández, F. Gomez, and et al., “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd international conference on Machine learning*, 2006.
- [30] M. A. Di Gangi, R. Cattoni, L. Bentivogli, and et al., “MuST-C: a Multilingual Speech Translation Corpus,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019.
- [31] M. Ott, S. Edunov, A. Baevski, and et al., “fairseq: A fast, extensible toolkit for sequence modeling,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019.
- [32] A. Baevski, Y. Zhou, A. Mohamed, and et al., “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [33] T. Kudo and J. Richardson, “Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018.
- [34] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations*, 2019.