



Semantic VAD: Low-Latency Voice Activity Detection for Speech Interaction

Mohan Shi¹, Yuchun Shu², Lingyun Zuo³, Qian Chen³, Shiliang Zhang³, Jie Zhang¹, Li-Rong Dai¹

¹NERC-SLIP, University of Science and Technology of China (USTC), China

²Tianjin University, China

³Speech Lab of DAMO Academy, Alibaba Group, China

smohan@mail.ustc.edu.cn, {jzhang6; lrdai}@ustc.edu.cn, yuchunshu@tju.edu.cn,
{ailsa.zly; tanqing.cq; sly.zsl}@alibaba-inc.com

Abstract

For speech interaction, voice activity detection (VAD) is often used as a front-end. However, traditional VAD algorithms usually need to wait for a continuous tail silence to reach a preset maximum duration before segmentation, resulting in a large latency that affects user experience. In this paper, we propose a novel semantic VAD for low-latency segmentation. Different from existing methods, a frame-level punctuation prediction task is added to the semantic VAD, and the artificial endpoint is included in the classification category in addition to the often-used speech presence and absence. To enhance the semantic information of the model, we also incorporate an automatic speech recognition (ASR) related semantic loss. Evaluations on an internal dataset show that the proposed method can reduce the average latency by 53.3% without significant deterioration of character error rate in the back-end ASR compared to the traditional VAD approach.

Index Terms: Voice activity detection, semantic loss, low latency, speech interaction, character error rate.

1. Introduction

Speech interaction [1] is one of the most representative applications of artificial intelligence, where long audio recordings need to be cut into short audio segments for downstream tasks, e.g., automatic speech recognition (ASR) [2], speech translation [3], speech emotion recognition [4], etc. For this purpose, a front-end voice activity detection (VAD) [5–7] is often built to distinguish between speech and non-speech periods, as the former is more related to the interaction goals.

Over the last few decades, many VAD approaches have been proposed. Early VAD methods classify speech and non-speech regions based on energy [8–12], zero-crossing rate [13], and periodicity measure [14]. More recently, with the advance in deep learning, the application of deep neural networks (DNNs) has become dominant in this field [15–18], where most aim to improve the classification capability. Also, there are some works on improving noise robustness, e.g., [19–21].

Directly applying these traditional VAD approaches may cause a new problem in realistic speech interaction scenarios. Specifically, traditional VAD algorithms usually need to wait for a continuous tail silence to reach a preset maximum tail silence duration (e.g., 700ms) before deciding whether to perform




Figure 1: Schematic of the tail silence.

tail segmentation, e.g., see the example in Figure 1, which results in a relatively high tail segmentation latency and seriously affects the user experience [22,23]. In order to address the high-latency issue of tail segmentation, we consider a novel semantic VAD approach in this work.

Firstly, as traditional VAD models only distinguish speech or silence and neglect whether each part of silence is a complete semantic breakpoint, they have to wait for a long continuous tail silence (e.g., 700ms) before making a tail segmentation judgment to avoid splitting the whole sentences incorrectly. It follows that if a semantic breakpoint is contained within the silence segment, the tail segmentation decision can then be simply made by waiting for a relatively short tail silence. In order to obtain semantic breakpoint information, we thus consider adding a frame-level punctuation prediction task to the basic VAD model. In case an ending punctuation is detected (e.g., period, question mark), indicating that there is a full semantic breakpoint, it is reasonable to wait for a short tail silence (e.g., 300ms). In the case of detecting a non-ending punctuation (e.g., comma), the tail silence required for segmentation needs to be slightly longer than that of ending punctuation (e.g., 400ms). Only when punctuation can not be predicted, the preset maximum tail silence of the traditional VAD (e.g., 700ms) will be used to determine the segmentation point.

Second, in order to further increase the proportion of early judgments and reduce the tail segmentation latency, we add predictions for artificially defined endpoints. Specifically, the endpoint label is introduced in addition to conventional binary speech-presence and speech-absence categories, forming a three-class classification task. The endpoint label is generated by a combination of punctuation and silence duration, similar to how the punctuation is utilized to determine the tail segmentation. The difference lies in that combining punctuation to determine the tail segmentation is implemented in the inference stage, while combining punctuation to generate endpoint labels artificially changes the training target of VAD.

Finally, since the prediction of punctuation and endpoint requires the model to be semantically strong, we introduce a semantic loss as an auxiliary task to enhance the semantic modeling ability. It should be noted that, different from existing joint VAD and ASR multi-task training methods [24,25], where VAD is cascaded with ASR so that the speech is processed as a whole, resulting in a relatively large parameter amount, while

This work was done while the first author was an intern at Speech Lab of DAMO Academy, Alibaba Group, and supported by the National Natural Science Foundation of China (62101523), Hefei Municipal Natural Science Foundation (2022012) and USTC Research Funds of the Double First-Class Initiative (YD2100002008). (Correspondence: jzhang6@ustc.edu.cn)

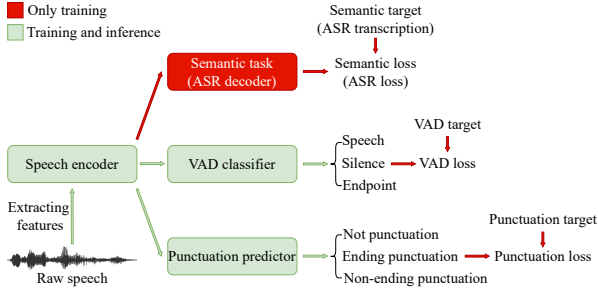


Figure 2: The schematic of the proposed semantic VAD, where the semantic ASR decoder only appears for training and the rest components are required for both training and inference.

the proposed semantic loss is only used to assist model training and not present at inference time. As such, the independence and lightweight of our semantic VAD model can be guaranteed, meaning that the proposed VAD model can be separately applied or flexibly combined with other downstream tasks.

Evaluations on an internal dataset show that the proposed method can reduce the average tail segmentation latency by 53.3% without significant deterioration of character error rate in the back-end ASR compared to the traditional VAD approach. The rest of this paper is organized as follows. Section 2 presents the proposed semantic VAD. Section 3 shows the experimental setup, followed by results in Section 4. Finally, Section 5 concludes this work.

2. Semantic VAD

The schematic of semantic VAD is shown in Figure 2, which consists of the speech encoder, punctuation prediction, VAD classifier, and semantic task.

2.1. Speech encoder

To facilitate the subsequent migration to the streaming platform, we use contextual block conformer [26, 27] as our speech encoder, which is a block processing conformer structure designed for streaming ASR. In the speech encoder, a context embedding vector is introduced to convey contextual information between blocks. The acoustic features are passed through the speech encoder to generate a frame-level speech representation.

2.2. Punctuation prediction

In traditional VAD, there are no decisions on semantic breakpoints, so it is necessary to wait for a long tail silence before performing tail segmentation. In order to reduce the tail segmentation latency, it might be helpful to obtain semantic segmentation via punctuation prediction. This can be accomplished by directly incorporating the frame-level speech representation obtained by the speech encoder into the classifier to obtain the frame-level punctuation prediction. In this work, we consider three types of punctuations:

$$P = \begin{cases} \text{Not punctuation} \\ \text{Ending punctuation (E-punc)} \\ \text{Non-ending punctuation (NE-punc)} \end{cases} \quad (1)$$

where the ending punctuation means complete semantic breakpoints (e.g., periods, question marks) and the non-ending punctuation implies incomplete semantic breakpoints (e.g., comma).

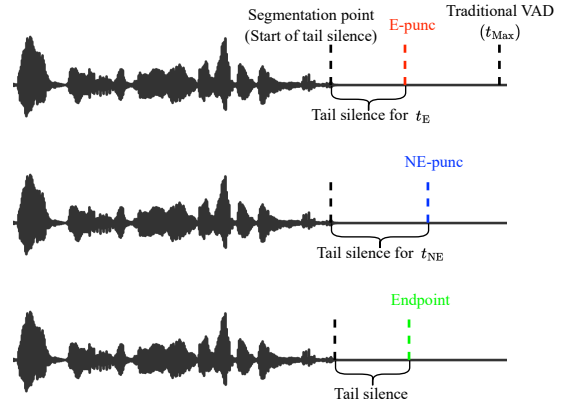


Figure 3: Schematic diagram of several judgment methods during inference.

2.3. VAD with endpoint

Conventionally, each speech frame is only simply labeled with speech-presence or speech-absence. In order to further improve the proportion of tail segmentation points judged in advance and reduce the tail segmentation latency, the artificially defined endpoint label is introduced as the third category, resulting in a three-class classification problem:

$$S \in \{\text{Speech, Silence, Endpoint}\}, \quad (2)$$

where the endpoint is the newly defined label in the training set. On the basis of obtaining the labels of speech and silence, in case a sentence ends with an E-punc, the frame after a continuous t_E silence until the next speech begins will be set with an endpoint; in case of a NE-punc, the duration of the tail silence is then set to be t_{NE} . Similarly to punctuation prediction, the output of the speech encoder has to pass through a classification layer to obtain the frame-level VAD decision.

2.4. Semantic task

Since both punctuation classification and the prediction of the artificially specified endpoint frames in the considered VAD, the representation obtained by the speech encoder should contain strong semantic information. Therefore, it is necessary to introduce a semantic loss as an auxiliary task to preserve the semantic information in speech representation. In this work, we use the ASR loss as the auxiliary semantic loss. Specifically, the speech encoder is connected to the decoder for pre-training on the ASR task, and the punctuation prediction, VAD and ASR are then jointly trained. The overall loss function for joint training can be represented as

$$L = \mu L^{\text{punc}}(\hat{p}, p) + \lambda L^{\text{asr}}(\hat{y}, y) + (1 - \mu - \lambda) L^{\text{vad}}(\hat{s}, s) \quad (3)$$

where p , y , and s represent the true punctuation, transcription, and VAD label, respectively, while \hat{p} , \hat{y} , and \hat{s} represent the corresponding hypothesis. Notice that unlike previous multi-task works, ASR is only used here as an auxiliary task to enrich the semantic information during training, and does not appear during inference. This guarantees the lightweight and independence of the front-end VAD module at the inference phase.

2.5. Inference

For traditional VAD methods, after the tail silence reaches t_{Max} continuously, it will be assigned a break point, and segmentation is performed at the beginning of the tail silence. This leads

Table 2: Performance comparison of traditional and semantic VAD methods on the three test sets in terms of the tail segmentation latency (ms), CER (%) for back-end ASR and DCF (%) for front-end segmentation.

Approach	Average			Test1			Test2			Test3		
	Latency	CER	DCF	Latency	CER	DCF	Latency	CER	DCF	Latency	CER	DCF
Traditional VAD	700	16.75	10.78	700	19.95	13.51	700	27.41	9.94	700	6.78	7.60
	600	16.80	10.78	600	19.96	13.50	600	27.74	9.95	600	6.77	7.61
	500	16.98	10.70	500	20.10	13.42	500	28.08	9.82	500	6.93	7.55
	400	17.42	10.71	400	20.62	13.42	400	28.56	9.84	400	7.23	7.56
	330	17.94	10.70	330	21.17	13.41	330	29.04	9.81	330	7.73	7.56
Semantic VAD	326.92	16.88	9.60	326.45	19.91	11.52	338.99	27.84	9.60	326.32	7.04	7.10

Table 1: The label distributions (%) in the training set.

punctuation labels			VAD labels		
Not punc	E-punc	NE-punc	Speech	Silence	Endpoint
73.79	22.05	4.16	58.51	18.80	22.69

to a high tail segmentation latency and thus an annoying user experience. In order to avoid this, we consider four cases to perform tail segmentation:

- An endpoint is detected in the continuous tail silence;
- An E-punc is detected and the continuous tail silence reaches t_E ;
- A NE-punc is detected and the continuous tail silence reaches t_{NE} ;
- The continuous tail silence reaches t_{Max} .

These are graphically shown in Figure 3. It will be experimentally shown that combining the prediction of punctuation and endpoint can greatly reduce the tail segmentation latency.

3. Experimental Setup

3.1. Dataset

In this work, we use 325 hours of internal Mandarin speech data, obtained from various websites, such as YouTube, as training and validation sets. We use an internal Mandarin long speech test set of three different scenarios for the final evaluation, including a meeting (Test1, for 3.3 hours), a recording pen (Test2, for 1 hour), and a live webcast (Test3, for 2.1 hours). Punctuation labels per frame are obtained by forcing alignment. Table 1 shows the distribution of various labels in the training set.

3.2. Evaluation metrics

Latency: Since the focus of this work is on the reduction of the tail latency of VAD segmentation during speech interaction, the latency time is taken as one of the evaluation metrics. Specifically, the latency is measured by the average continuous tail silence duration during VAD segmentation. Since the same VAD model is used in experiments, the latency of the model itself is ignored for convenience.

Character error rate (CER): Because the downstream task of the back-end best reflects the performance of VAD segmentation, in this work, we choose ASR as the downstream task of the back-end. After VAD segmentation, the segmented short audio will be sent to the pre-trained ASR system to obtain the corresponding transcriptions. We use pre-trained Universal ASR [28] as our back-end ASR model. Since the Mandarin-speaking cor-

pus is used for both training and evaluation, we exploit the CER as the performance measure of ASR.

Detection cost function (DCF): The DCF is a public evaluation metric of NIST sound activity detection (SAD), which is used to measure the signal-level VAD performance in this work. In VAD segmentation, four types of decisions might appear: True Negative (TN), True Positive (TP), False Negative (FN), and False Positive (FP). The DCF can thus be calculated as:

$$P_{miss} = \frac{FN \text{ time}}{TP \text{ time} + FN \text{ time}}, \quad (4)$$

$$P_{fa} = \frac{FP \text{ time}}{TN \text{ time} + FP \text{ time}}, \quad (5)$$

$$DCF = 0.75 * P_{miss} + 0.25 * P_{fa}, \quad (6)$$

where P_{miss} and P_{fa} stand for the miss rate and false alarm rate, respectively. The weights in (6) are empirically chosen. Note that a smaller DCF means a more accurate VAD segmentation.

3.3. Model configuration

In this work, we use the 80-dimensional log-mel filterbank (Fbank) as the input feature. The window size is 25 ms, and the window shift is 10 ms. Based on experience as well as tuning results on the development set, t_E is set to 300ms, t_{NE} is set to 400ms, and t_{Max} is set to 700ms. A 6-layer contextual block conformer with 4-head of multi-head attention (MHA) is used as the speech encoder. The dimensions of MHA and feed-forward network (FFN) are set to be 256 and 512, respectively, the convolution kernel size is set to be 15, the block size is 64 (corresponding to 640 ms), and the hop size is 16 (i.e., 160 ms). In each block, the first 160 ms is the future information. For the auxiliary ASR task, a mixture of 4-layer transformer decoder and CTC [29] is used for training. The decoder configuration keeps the same as the speech encoder with a CTC loss weight of 0.3. The parameter amount of the speech encoder is around 5.98M. We first pre-train the ASR task for 50 epochs, and then jointly train the punctuation prediction, ASR and VAD tasks for 50 epochs, where both μ and λ are set to be 0.2. We use the ESPnet toolkit [30] for our experiments.

4. Experimental Results

First of all, we show in Table 2 the comparison of the performance of traditional and semantic VAD methods on the three test sets in terms of the tail segmentation latency, back-end ASR and front-end segmentation. For a fair comparison, the same speech encoder structure as the semantic VAD is used in the traditional VAD, but the punctuation prediction is excluded. That is, only the binary classification is performed by the traditional VAD approach to distinguish speech from silence, and the semantic loss is not considered during training. It can be seen

Table 3: Performance comparison of the performance of semantic VAD under different strategies on the three test sets.

Approach	Average			Test1			Test2			Test3		
	Latency	CER	DCF	Latency	CER	DCF	Latency	CER	DCF	Latency	CER	DCF
Semantic VAD	362.95	16.92	10.91	355.29	20.08	13.47	359.07	27.32	9.43	377.20	7.13	8.23
+Semantic loss	344.63	16.75	9.70	342.14	19.75	11.57	355.19	27.75	9.68	347.86	6.93	7.27
+Endpoint	326.92	16.88	9.60	326.45	19.91	11.52	338.99	27.84	9.60	326.32	7.04	7.10




Figure 4: The visual comparison of the tail judgment points of semantic and traditional VAD methods.




Figure 5: The judgment proportions of the semantic VAD.

that in traditional VAD, the performance of back-end ASR becomes worse (with an average CER increase from 16.75% to 17.94%) as the preset tail segmentation latency decreases from 700 ms to 330 ms. The proposed semantic VAD approach can significantly reduce the average tail segmentation latency (from 700 ms to 326.92 ms) without a serious sacrifice in the back-end ASR performance (with an average CER increase from 16.75% to 16.88%). More importantly, the proposed method can improve the front-end segmentation performance in DCF from 10.78% to 9.60% due to more semantic segmentation. It is worth mentioning that in case the tail segmentation latency of traditional VAD is set to be 330 ms, it becomes equivalent to the proposed semantic VAD, while the corresponding performance of back-end ASR is much worse than that of semantic VAD. These clearly show that the proposed method can not only reduce the latency and improve the segmentation accuracy, but also preserve the back-end ASR performance.

Table 3 presents the comparison of the proposed semantic VAD with different strategies on the three test sets in terms of the tail segmentation latency, back-end ASR and front-end segmentation. It is clear that adding the semantic loss as an auxiliary task in training leads to better performance in all three metrics (e.g., latency from 362.95ms to 344.63ms, CER from 16.92% to 16.75%, DCF from 10.91% to 9.70% on average). This indicates that with the assistance of semantic loss, the pro-

portion of segmentation through punctuation and endpoint becomes larger, and the semantic of the segmented sentence is more complete, which is informative to the back-end ASR module. Involving endpoint judgment in inference leads to a further reduction in tail segmentation latency and DCF, but slightly degrades the back-end ASR performance. This is due to the fact that the endpoint is a kind of artificially defined label, which somehow increases the error in inference.

Finally, we conduct a visual analysis of the proposed semantic VAD. In Figure 4, we compare the tail judgment points of traditional VAD and semantic VAD for an audio segment of Test 1. The upward and downward stems denote the judgment points of the traditional and proposed VAD methods, and the red, blue, and green stems represent the judgments of E-punc, NE-punc, and Endpoint, respectively. It is clear that the judgment points of the semantic VAD are all before those of the traditional VAD. This verifies the efficacy of the proposed method in reducing latency. Figure 5 illustrates the proportion of various ways of obtaining tail judgment points for the semantic VAD on all test sets. It can be clearly seen that the maximum silence judgment method (i.e., traditional VAD) accounts for only 5.41%, which indicates that our semantic VAD method can make most tail segmentation points be judged in advance. As the endpoint accounts for a relatively small proportion (14.42%), the artificially defined label is semantically abstract, and it is thus necessary to make a pre-judgment through punctuation prediction (e.g., with high proportions).

5. Conclusion

In this paper, we proposed a novelty semantic VAD method, which can synthetically determine whether tail segmentation should be performed by introducing the punctuation prediction, artificially defined endpoints, and auxiliary semantic task. Experiments on the internal dataset show that the proposed semantic approach significantly reduces the tail segmentation latency without an obvious sacrifice in the back-end ASR performance compared to the traditional VAD approach. In the future, we will investigate the application to other semantic-dependent speech tasks (e.g., speech translation, dialogue understanding) as well as the generalization to multi-task cases.

6. References

- [1] G. Philip and E. S. Young, "Man machine interaction by voice: Developments in speech technology: Part I: the state-of-the-art," *J. Inf. Sci.*, vol. 13, no. 1, pp. 3–14, 1987.
- [2] J. Li, "Recent advances in end-to-end automatic speech recognition," *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.
- [3] J. Niehues, "Survey talk: A survey on speech translation," in *INTERSPEECH*. ISCA, 2019.
- [4] G. Drakopoulos, G. Pikramenos, E. D. Spyrou, and S. J. Perantonis, "Emotion recognition from speech: A survey," in *WEBIST*. ScitePress, 2019, pp. 432–439.
- [5] B. Atal and L. Rabiner, "A pattern recognition approach to voiced-unvoiced-silence classification with applications to speech recognition," *IEEE Trans. Audio Speech Lang. Process.*, vol. 24, no. 3, pp. 201–212, 1976.
- [6] S. I. Salishev, A. Barabanov, D. Kocharov, P. A. Skrelin, and M. J. Moiseev, "Voice activity detector (VAD) based on long-term mel frequency band features," in *TSD*, ser. Lecture Notes in Computer Science, vol. 9924. Springer, 2016, pp. 352–358.
- [7] I. Kraljevski, Z. Tan, and M. P. Bissiri, "Comparison of forced-alignment speech recognition and humans for generating reference VAD," in *INTERSPEECH*. ISCA, 2015, pp. 2937–2941.
- [8] S. G. Tanyer and H. Özer, "Voice activity detection in nonstationary noise," *IEEE Trans. Speech Audio Process.*, vol. 8, no. 4, pp. 478–482, 2000.
- [9] A. Davis, S. Nordholm, and R. Togneri, "Statistical voice activity detection using low-variance spectrum estimation and an adaptive threshold," *IEEE Trans. Speech Audio Process.*, vol. 14, no. 2, pp. 412–424, 2006.
- [10] K.-H. Woo, T.-Y. Yang, K.-J. Park, and C. Lee, "Robust voice activity detection algorithm for estimating noise spectrum," *Electronics Letters*, vol. 36, no. 2, pp. 180–181, 2000.
- [11] I. Yoo, H. Lim, and D. Yook, "Formant-based robust voice activity detection," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 23, no. 12, pp. 2238–2245, 2015.
- [12] M. Hadi, M. R. Pakravan, and M. M. Razavi, "An efficient real-time voice activity detection algorithm using teager energy to energy ratio," in *2019 27th Iranian Conference on Electrical Engineering (ICEE)*. IEEE, 2019, pp. 1420–1424.
- [13] J. Junqua, B. Reaves, and B. Mak, "A study of endpoint detection algorithms in adverse conditions: incidence on a DTW and HMM recognizer," in *EUROSPEECH*. ISCA, 1991.
- [14] R. Tucker, "Voice activity detection using a periodicity measure," *IEE Proceedings I (Communications, Speech and Vision)*, vol. 139, no. 4, pp. 377–380, 1992.
- [15] T. Hughes and K. Mierle, "Recurrent neural networks for voice activity detection," in *ICASSP*. IEEE, 2013, pp. 7378–7382.
- [16] J. Kim, J. Kim, S. Lee, J. Park, and M. Hahn, "Vowel based voice activity detection with LSTM recurrent neural network," in *ICSPS*. ACM, 2016, pp. 134–137.
- [17] Y. Lee, J. Min, D. K. Han, and H. Ko, "Spectro-temporal attention-based voice activity detection," *IEEE Signal Process. Lett.*, vol. 27, pp. 131–135, 2020.
- [18] T. G. Kang and N. S. Kim, "Dnn-based voice activity detection with multi-task learning," *IEICE Trans. Inf. Syst.*, vol. 99-D, no. 2, pp. 550–553, 2016.
- [19] X. Zhang and D. Wang, "Boosting contextual information for deep neural network based voice activity detection," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 24, no. 2, pp. 252–264, 2016.
- [20] R. Masumura, K. Matsui, Y. Koizumi, T. Fukutomi, T. Oba, and Y. Aono, "Context-aware neural voice activity detection using auxiliary networks for phoneme recognition, speech enhancement and acoustic scene classification," in *EUSIPCO*. IEEE, 2019, pp. 1–5.
- [21] J. Kim and M. Hahn, "Voice activity detection using an adaptive context attention model," *IEEE Signal Process. Lett.*, vol. 25, no. 8, pp. 1181–1185, 2018.
- [22] Y. Shangguan, R. Prabhavalkar, H. Su, J. Mahadeokar, Y. Shi, J. Zhou, C. Wu, D. Le, O. Kalinli, C. Fuegen, and M. L. Seltzer, "Dissecting user-perceived latency of on-device E2E speech recognition," in *Interspeech*. ISCA, 2021, pp. 4553–4557.
- [23] J. Hou, W. Guo, Y. Song, and L. Dai, "Segment boundary detection directed attention for online end-to-end speech recognition," *EURASIP J. Audio Speech Music. Process.*, vol. 2020, no. 1, p. 3, 2020.
- [24] W. R. Huang, S. Chang, D. Rybach, T. N. Sainath, R. Prabhavalkar, C. Peyser, Z. Lu, and C. Allauzen, "E2E segmenter: Joint segmenting and decoding for long-form ASR," in *INTERSPEECH*. ISCA, 2022, pp. 4995–4999.
- [25] S. Bijwadia, S. Chang, B. Li, T. N. Sainath, C. Zhang, and Y. He, "Unified end-to-end speech recognition and endpointing for fast and efficient speech systems," in *SLT*. IEEE, 2022, pp. 310–316.
- [26] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *INTERSPEECH*. ISCA, 2020, pp. 5036–5040.
- [27] E. Tsunoo, Y. Kashiwagi, T. Kumakura, and S. Watanabe, "Transformer ASR with contextual block processing," in *ASRU*. IEEE, 2019, pp. 427–433.
- [28] Z. Gao, S. Zhang, M. Lei, and I. McLoughlin, "Universal ASR: unifying streaming and non-streaming ASR using a single encoder-decoder model," *CoRR*, vol. abs/2010.14099, 2020.
- [29] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-end speech recognition," *IEEE J. Sel. Top. Signal Process.*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [30] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba *et al.*, "Espnet: End-to-end speech processing toolkit," in *INTERSPEECH*. ISCA, 2018, pp. 2207–2211.