Improving Speech Representation Learning via Speech-level and Phoneme-level Masking Approach

Xulong Zhang[1], Jianzong Wang[1*], Ning Cheng[1], Kexin Zhu[1,2], Jing Xiao[1]

[1]*Ping An Technology (Shenzhen) Co., Ltd.*
[2]*School of Software, Fudan University, Shanghai, China*

Abstract—Recovering the masked speech frames is widely applied in speech representation learning. However, most of these models use random masking in the pre-training. In this work, we proposed two kinds of masking approaches: (1) speech-level masking, making the model to mask more speech segments than silence segments, (2) phoneme-level masking, forcing the model to mask the whole frames of the phoneme, instead of phoneme pieces. We pre-trained the model via these two approaches, and evaluated on two downstream tasks, phoneme classification and speaker recognition. The experiments demonstrated that the proposed masking approaches are beneficial to improve the performance of speech representation.

Index Terms—speech representation learning, masking approach, phoneme classification, speaker recognition, TEAR

I. Introduction

Speech representation learning has proved to be an effective method of extracting high-level speech information [1]. The speech representation learning is usually based on pre-training on a large scale unlabeled speech data. It is well known that labeled speech data requires huge amounts of labour costs, while the unlabeled corpus is relatively easy to obtain [2]. The extracted speech representation could be used to improve the downstream speech and language processing (SLP) tasks.

Several Self-supervised learning (SSL) methods have been proposed for speech representation learning [3], [4]. Autoregressive predictive coding (APC) [5] is designed to learn an auto-regressive model by predicting future speech frames. Contrastive predictive coding (CPC) [6] uses a contrastive loss that maximizing the mutual information between present representations and future signals. Wav2vec [7] also learns a contrastive objective that distinguishs the true future audio from negative samples. All of the above SSL models could extract speech representation through pre-training on large amounts of unlabeled data.

Masked language model (MLM) [8] is another popular SSL architecture of speech representation learning. The MLM model often uses a Transformer-based network to reconstruct the masked or altered speech frames. Masked predictive coding (MPC) [9] used a Transformer encoder to predict the masked filter bank (Fbank) features, and fine-tuned the Transformer decoder for transcript prediction. Mockingjay [10] proposed to use a BERT-style masking strategy with random selected frames for pre-training. Transformer encoder representations

*Corresponding author: Jianzong Wang, jzwang@188.com.

from alteration (TERA) [11] extended the work of Mockingjay, introducing three auxiliary multi-task learning objectives (temporal, channel, and magnitude) to the self-supervised learning for speech. Wav2vec2.0 [12] masks the speech inputs in the latent space, and pre-trains the model by a contrastive loss of the quantized latent speech representation.

Despite the impressive performance of these MLM models, some pre-training strategies still could be improved. One is that most of the masking approaches is random masking for these models. They do not consider any prior knowledge in the speech. In natural language processing (NLP), some previous works were proposed to use knowledge based masking strategies, instead of random masking. Enhanced Representation through kNowledge IntEgration (ERNIE) [13] is designed to learn language representation by entity-level masking and phrase-level masking. SpanBERT [14] proposed to mask contiguous spans and designed a span boundary objective loss relying on the relative position within the masked span. Cui et al. [15] presented whole word masking to BERT [16], providing a more challenging pre-training task of predicting all the characters in a complete Chinese word. RoBERTa [17] found that dynamic masking is beneficial to pre-training on large datasets, which generates the masking pattern dynamically when feeding it to the model.

For SLP tasks [18], random masking approach is likely to select the non-speech segments. Some speech data may contain lots of silence like telephone conversational corpus. These less informative segments make the pre-training task too easy to recover these non-speech frames. Inspired by the above works in NLP, we proposed two levels of masking approaches to the speech representation learning: (1) speech-level masking, (2) phoneme-level masking. The speech-level approach will mask more speech segments than silence segments at the pre-training stage. We think the speech frames may contain more useful acoustic information than non-speech frames. The voice activity detection (VAD) algorithm [19] is applied to split the original speech into speech segments or not. Furthermore, the phoneme-level approach will force the model to reconstruct the whole frames of the masked phoneme. This is a more challenging task than just reconstructing some pieces of the phoneme. A pre-trained automatic speech recognition (ASR) model [20] is used to force align each speech frame to a corresponding phoneme. We also combined the speech-level and phoneme-level masking together, to obtain better speech representation in the pre-training and better performance in

downstream tasks.

II. METHODOLOGY

A. Model Architecture

This paper exploits the transformer-based masked language model (MLM) as the overall model architecture. Our works mainly focus on the masking approach in temporal channel at the pre-training stage. A masking approach alters or masks a number of speech frames of the original input $X = (x_1, x_2, ..., x_T)$. The masking sequence is $M = (m_1, m_2, ..., m_T)$, where T is the length of acoustic sequence. We denote the masking process as conducting element-wise product $\odot$ between X and M. Then, the MLM P_{MLM} predicts an output $\widetilde{X} = (\widetilde{x}_1, \widetilde{x}_2, ..., \widetilde{x}_T)$ based on the masked input.

The objective of pre-training is to minimize the error between the predicted output $\widetilde{X}$ and the original input X. As in Mockingjay [10] and TERA [11], we also used the $\mathcal{L}_1$ loss as follows:

$$\mathcal{L}_1 = |X - \widetilde{X}| \tag{1}$$

The most common masking approach in temporal channel is random masking, which selects the masked frames randomly in time domain. As in [10], the approach masks successive C frames once they randomly pick a frame as the starting point. This is to avoid the model utilizing local smoothness of the acoustic frames.
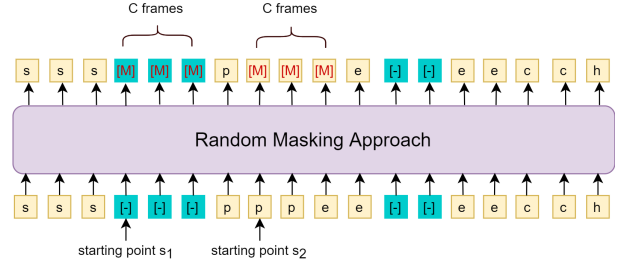
B. Speech-Level Masking Approach

We firstly proposed a speech-level masking approach based on the VAD algorithm [19]. The VAD algorithm is a binary classification problem determining whether an input signal contains speech or not. Certain features (e.g. energy, cepstral coefficient, etc.) are extracted from a segment of the input audio signals. Then, a threshold θ is set to classify the segment as the speech if the value of extracted features exceeds the threshold θ.
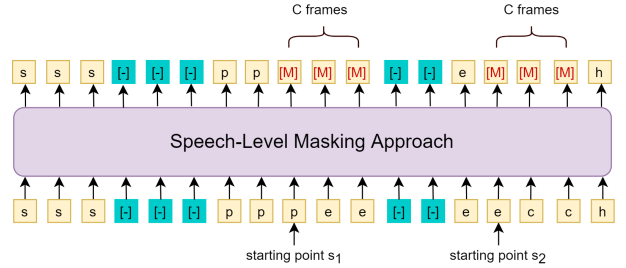
We think masking more speech segments could help the model to learn more useful acoustic information. A ratio parameter ρ is set to control the proportion of masked speech and non-speech segments. The reason why a small portion of non-speech segments are still masked is that the silence may sometimes contain high semantic knowledge. We firstly used VAD algorithm to classify whether each frame is speech or not. Then, the starting points of masking are selected from speech list A or non-speech list B, according to the ratio ρ. After that, C successive frames after each starting point are also masked, to generate the speech-level masking sequence M_S.
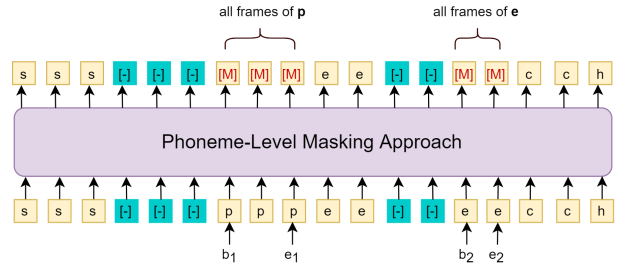
C. Phoneme-Level Masking Approach

In this section, we proposed a more challenging task by masking the whole frames of phoneme. We firstly used a pre-trained ASR model [20] P_{ASR} to predict the text Y of acoustic features X. In real application, the P_{ASR} is usually trained by a small amount of labeled data. Then, we applied force alignment to map each speech frame to one phoneme,

(a) Random Masking Speech and Silence Segments

(b) Speech-Level Masking More Speech Segments

(c) Phoneme-Level Masking All Frames of Phoneme **p** and **e**

Fig. 1. Visualization of Different Masking Approaches

generating the aligned phoneme sequence Y'. Force alignment [21] is a task of determining the time boundaries between phonemes of acoustic features. After that, N phonemes are selected randomly. All the frames between the begin index b_j and end index e_j of each selected phoneme y''_j are masked. Especially, we masked the whole frames of each phoneme in this approach, instead of C successive frames.

D. Visualization

We illustrated the masking process of above three masking approaches in Figure 1. The input signal is the speech frames of a word *speech* (in light yellow boxes). The non-speech frames are denoted as a symbol $[-]$ (in light green boxes). The masked frames after masking approach are denoted as a red symbol $[M]$.

For random masking approach (in Figure 1(a)), two starting points s_1 and s_2 are randomly selected and C successive frames are masked. The random masking is likely to select the silence segment (s_1 in Figure 1(a)). For speech-level masking approach (in Figure 1(b)), speech and non-speech frames are classified by VAD algorithm firstly. Then, most of the masked

TABLE I
Compared with Different Masking Approaches, Results on Librispeech, Accuracy(%)

Pre-training	Masking Approach	train-clean-100				train-clean-360			
		Phoneme-L	Phoneme-1H	Speaker-F	Speaker-U	Phoneme-L	Phoneme-1H	Speaker-F	Speaker-U
Mockingjay [10]	Random	69.6	78.8	68.4	96.1	67.5	78.2	86.9	97.3
	Speech-Level	70.2	79.3	97.6	97.2	68.0	78.1	97.8	98.3
	Phoneme-Level	70.2	79.7	97.9	**98.5**	67.8	78.8	**98.1**	**98.9**
	Speech&Phoneme-Level	**70.3**	**79.9**	**98.2**	98.2	**68.5**	**78.9**	97.2	98.3
TERA [11]	Random	71.3	79.1	98.9	99.2	70.8	79.6	99.1	99.3
	Speech-Level	71.5	**80.3**	99.6	99.3	71.3	**80.7**	99.1	99.2
	Phoneme-Level	71.4	79.5	**99.7**	**99.7**	71.7	80.4	99.3	**99.5**
	Speech&Phoneme-Level	**71.8**	80.1	99.5	99.4	**71.8**	80.5	**99.3**	99.4

frames are selected as the informative speech frames (s_1 and s_2 in Figure 1(b)). The speech-level approach will also mask C successive frames. For phoneme-level masking approach (in Figure 1(c)), phoneme boundaries are detected by force alignment algorithm. In which, b_1 and e_1 denote the begin and end frame of phoneme p, and b_2 and e_2 cover the frames of phoneme e. All frames of each phoneme are masked for a more challenging model pre-training.

In addition, speech-level and phoneme-level masking approach could also be combined together. Firstly, we apply the VAD algorithm to distinguish the speech and non-speech frames. Secondly, all frames of a detected phoneme are masked, when we choose a starting point in speech segment with probability ρ. Otherwise, fixed C frames from a starting point in silence segment are masked. The experimental results of this combined masking approach will be shown in Section III-B.

III. EXPERIMENTS

A. Experimental Setup

We used two subsets of Librispeech [22] corpus for pre-training: the *train-clean-100* and the *train-clean-360*. To improve the speech representation learning with well-designed masking strategies. Therefore, the accuracy of downstream tasks are compared using different masking approaches, without changing the network architecture. Two MLM models P_{MLM}: Mockingjay [10] and TERA [11], are used to extract the speech representations from the last layer of the model after pre-training.

For phoneme classification tasks, we utilized linear classifier (denoted as *Phoneme-L*) and classifier with one single hidden layer (denoted as *Phoneme-1H*). For speaker recognition tasks, we performed frame-wise (denoted as *Speaker-F*) and utterance-wise (denoted as *Speaker-U*) classification to predict the speaker identity.

The input acoustic features are 80-dimensional Fbank. We set the time masking width C to 7 frames for random approach. Following the previous works in [10] and [11], 3-layers Transformer encoder network is used. The multi-head self-attention layer can extract feature information from multiple dimensions [23]. Each layer produces an output of the same dimension. The hidden size of intermediate feed-forward layer is 3072 with dropout rate 0.1. In addition,

the VAD algorithm is implemented by Google WebRTC framework [24]. The force-alignment results are obtained by Kaldi recipes [25]. We conducted all the pre-training and downstream experiments on S3PRL toolkit [26].

B. Results

As depicted in Table I, we illustrated the accuracy results ($20k$ pre-training steps, $20k$ downstream steps) of phoneme classification and speaker recognition with different masking approaches on Librispeech dataset. For Mockingjay model, the proposed three masking approaches (*Speech-Level*, *Phoneme-Level*, and *Speech&Phoneme-Level*) have higher accuracy than random masking approach, in both two downstream tasks. Especially in the phoneme classification task, the combined masking approach achieves the highest accuracy on both datasets. It achieves accuracy rates of 70.3% and 79.9% respectively when using the *train-clean-100* dataset for the *Phoneme-L* and *Phoneme-1H*, while its effect on the *train-clean-360* dataset is 68.5% and 78.9%, respectively. Compared to random masking approach, our three masking approaches have achieved very significant improvements in the accuracy of the frame-wise speaker recognition task. In this task, our three approaches achieved accuracies of 97.6%, 97.9% and 98.2% respectively when using the *train-clean-100* dataset, compared to 68.4% for the random approach. Our approaches also achieved 97.8%, 98.1% and 97.2% accuracies on the *train-clean-360* dataset, respectively, as compared to the random approach accuracy of 86.9%. For TERA model, all of our three masking approaches could outperform random masking approach in phoneme classification task. The proposed approaches also have comparable performance with random masking approach in speaker recognition tasks, despite the results are very close with each other. The *Speech&Phoneme-Level* approach performs best on the *Phoneme-L*, achieving 71.8% accuracy on both datasets. The *Speech-Level* approach performed best on the *Phoneme-1H*, achieving 80.3% and 80.7% accuracy on the *train-clean-100* and *train-clean-360* datasets, respectively. The *Phoneme-Level* approach performs best on speaker recognition tasks, where it achieves 99.7% accuracy on both *Speaker-F* and *Speaker-U* when using the *train-clean-100* dataset. When using the *train-clean-360* dataset, the *Phoneme-Level* approach also achieves 99.5%

TABLE II
Quick Tests with Different Proportion Ratios ρ, Results on Librispeech, Accuracy(%)

ρ	Speech-Level		Speech&Phoneme-Level	
	Phoneme-L	Phoneme-1H	Phoneme-L	Phoneme-1H
0.80	59.3	65.8	61.2	66.4
0.85	60.2	67.0	62.1	66.6
0.90	**61.0**	**68.0**	**62.3**	**68.0**
0.95	60.0	65.9	62.2	67.2
1.00	59.6	60.0	65.9	66.2

accuracy on the *Speaker-U*, slightly higher than other approaches, and 99.3% accuracy on the *Speaker-F* which is on par with the *Speech&Phoneme-Level* approach.

We also explored different values of proportion ratio ρ. The results of phoneme classification task are shown in Table II. We made quick tests ($20k$ pre-training steps, $5k$ downstream steps) on *train-clean-100* dataset. Two masking approaches (*Speech-Level* and *Speech&Phoneme-Level*) were investigated. The results indicated that $\rho = 0.9$ is the best choice. It means 90% of the masked segments are speech, and 10% are non-speech. When we set ρ to 0.9, both approaches showed better results than when set to several other values, in which the *Speech-Level* approach achieved the accuracy of *Phoneme-L* and *Phoneme-1H* the highest, at 61.0% and 68.0%, respectively. Similarly, when the parameter is set to 0.9, the *Speech-Level* approach also shows the highest accuracy of 68.0% on the *Phoneme-1H*. It also proves that some silence segments may contain high-level semantic knowledge, and they should not be discarded at all in the pre-training.

C. Spectrogram Analysis

In this section, we plotted the masking parts of spectrogram, and reconstructed spectrogram after pre-training by TERA model. As depicted in Figure 2, we made a comparison between random and *Speech&Phoneme-Level* masking approach, which are operated on one audio sample. The masking parts are highlighted in yellow lines.

For random masking, a lot of silence frames are likely masked. In addition, the masking areas have the same width in temporal dimension because the random approach masks fixed-length C successive frames (in Figure 2(a)). While for *Speech&Phoneme-Level* masking approach, the masking widths are variable, which are determined by the time duration of selected masking phonemes (in Figure 2(c)).

After pre-training, the spectrogram is predicted, and the masking parts are supplemented. We found that the reconstructed spectrogram is over smooth for random approach (in Figure 2(b)). It might be attributed to the local smoothness problem, which averages the surrounding signals when reconstructing the masking frames. On the contrary, our *Speech&Phoneme-Level* approach leads to a more sharp spectrogram in the masking areas (in Figure 2(d)). It proves that the proposed methods could alleviate the over smoothness prob-

(a) Masking Parts
(Random Approach)

(b) Reconstructed Spectrogram
(Random Approach)

(c) Masking Parts
(Speech&Phoneme-Level)

(d) Reconstructed Spectrogram
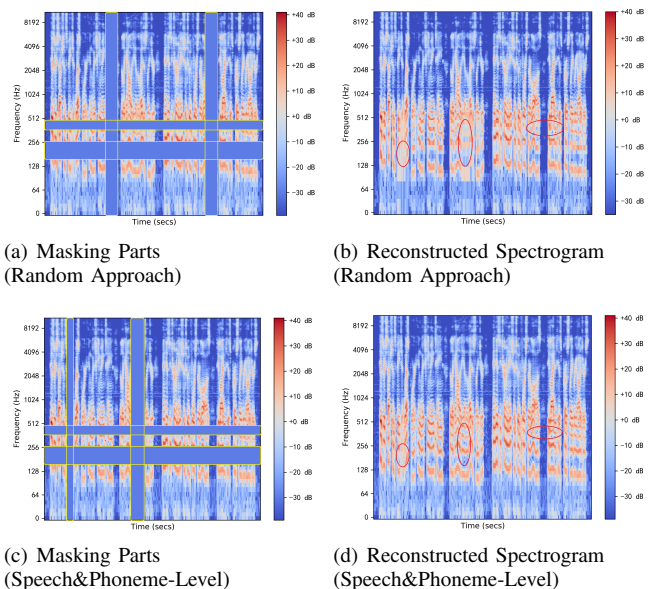(Speech&Phoneme-Level)

Fig. 2. Spectrogram Comparison Between Random and Speech&Phoneme-Level Masking Approach.

lem, and thus extract more meaningful speech representation than random approach.

IV. CONCLUSIONS

Random masking is widely used in existing speech representation learning models. However, previous random masking method masks non-speech segments from which useful acoustic information is difficult to obtain. This work proposed two well-designed strategies, *speech-level* and *phoneme-level* masking approaches. The experiments show that the proposed approaches have better results on downstream tasks, than random masking. We also found that combining two masking approaches could further improve the performance. In addition, some non-speech masked segments should be reserved to provide high-level information. We could set different ratios to control the proportion of silence and speech segments. Spectrogram analysis indicated that the proposed methods could alleviate the over smoothness problem, resulting in a more sharp reconstructed spectrogram. In future works, we will investigate unsupervised method of obtaining the phoneme boundaries, instead of force-alignment, such as gate activation signal method or phoneme clustering algorithm.

V. ACKNOWLEDGEMENT

REFERENCES

[1] Y.-A. Chung, C.-C. Wu, C.-H. Shen, and et al., "Audio word2vec: Unsupervised learning of audio segment representations using sequence-to-sequence autoencoder," in *IEEE International Speech Communication Association (INTERSPEECH)*, 2016, pp. 765–769.

[2] X. Zhang, J. Wang, N. Cheng, and J. Xiao, "Tdass: Target domain adaptation speech synthesis framework for multi-speaker low-resource tts," in *2022 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2022, pp. 1–7.

[3] S. Sadhu, D. He, C. W. Huang, and et al., "Wav2vec-c: A self-supervised model for speech representation learning," in *IEEE International Speech Communication Association (INTERSPEECH)*, 2021, pp. 711–715.

[4] H. Tang, X. Zhang, J. Wang, N. Cheng, and J. Xiao, "Avqvc: One-shot voice conversion by vector quantization with applying contrastive learning," in *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP2022)*. IEEE, 2022, pp. 4613–4617.

[5] Y.-A. Chung, W.-N. Hsu, H. Tang, and et al., "An unsupervised autoregressive model for speech representation learning," in *INTERSPEECH 2019*, 2019, pp. 146–150.

[6] A. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," in *IEEE Neural Information Processing Systems*, 2018.

[7] S. Schneider, A. Baevski, R. Collobert, and et al., "wav2vec: Unsupervised pre-training for speech recognition," in *INTERSPEECH 2019*, 2019, pp. 3465–3469.

[8] W. Wang, Q. Tang, and K. Livescu, "Unsupervised pre-training of bidirectional speech encoders via masked reconstruction," in *ICASSP 2020*, 2020, pp. 6889–6893.

[9] D. Jiang, X. Lei, W. Li, and et al., "Improving transformer-based speech recognition using unsupervised pre-training," in *arXiv preprint arXiv:1910.09932*, 2019.

[10] A. T. Liu, S. Yang, P. Chi, and et al., "Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders," in *ICASSP 2020*, 2020, pp. 6419–6423.

[11] A. Liu, S.-W. Li, and H.-y. Lee, "Tera: Self-supervised learning of transformer encoder representation for speech," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 29, pp. 2351–2366, 2020.

[12] A. Baevski, H. Zhou, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *IEEE Neural Information Processing Systems (NIPS)*, 2020, pp. 12 449–12 460.

[13] Z. Zhang, X. Han, Z. Liu, and et al., "Ernie: Enhanced language representation with informative entities," in *the Annual Meeting of The Association for Computational Linguistics*, 2019, pp. 1441–1451.

[14] M. Joshi, D. Chen, Y. Liu, and et al., "Spanbert: Improving pre-training by representing and predicting spans," in *IEEE Transactions of the Association for Computational Linguistics*, vol. 8, 2020, pp. 64–77.

[15] Y. Cui, W. Che, T. Liu, and et al., "Pre-training with whole word masking for chinese bert," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 29, pp. 3504–3514, 2019.

[16] J. Devlin, M.-W. Chang, K. Lee, and et al., "Bert: Pre-training of deep bidirectional transformers for language understanding," in *IEEE North American Association for Computational Linguistics*, 2018, pp. 4171–4186.

[17] Y. Liu, M. Ott, N. Goyal, and et al., "Roberta: A robustly optimized bert pretraining approach," in *arXiv preprint arXiv:1907.11692*, 2019.

[18] X. Jia, J. Wang, Z. Zhang, and et al., "Large-scale transfer learning for low-resource spoken language understanding," in *IEEE International Speech Communication Association (INTERSPEECH)*, 2020, pp. 1555–1559.

[19] S. Salishev, A. Barabanov, D. Kocharov, and et al., "Voice activity detector (vad) based on long-term mel frequency band features," in *Text, Speech, and Dialogue (TSD)*, 2016, pp. 352–358.

[20] S. Kim, T. Hori, and S. Watanabe, "Joint ctc-attention based end-to-end speech recognition using multi-task learning," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 4835–4839.

[21] M. McAuliffe, M. Socolof, S. Mihuc, and er al., "Montreal forced aligner: Trainable text-speech alignment using kaldi," in *IEEE International Speech Communication Association (INTERSPEECH)*, 2017, pp. 498–502.

[22] V. Panayotov, G. Chen, D. Povey, and et al., "Librispeech: An asr corpus based on public domain audio books," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.

[23] A. Vaswani, N. Shazeer, N. Parmar, and et al., "Attention is all you need," in *IEEE Neural Information Processing Systems (NIPS)*, 2017, pp. 5998–6008.

[24] J. Wiseman, "Google webrtc," in *online GitHub repository: https://github.com/wiseman/py-webrtcvad*.

[25] D. Povey, "Kaldi," in *online GitHub repository: https://github.com/kaldi-asr/kaldi*.

[26] A. T. Liu and Y. Shu-wen, "S3prl: The self-supervised speech pre-training and representation learning toolkit," in *online GitHub repository: https://github.com/s3prl/s3prl*, 2020.