

# Speech Enhancement Using Self-Supervised Pre-Trained Model and Vector Quantization

Xiao-Ying Zhao\*, Qiu-Shi Zhu<sup>†</sup> and Jie Zhang<sup>†</sup>

\*University of Science and Technology of China (USTC), Hefei, Anhui, China

<sup>†</sup>National Engineering Research Center for Speech and Language Information Processing,

University of Science and Technology of China (USTC), Hefei, Anhui, China

E-mail: {xyzhao1123, qszhu}@mail.ustc.edu.cn, jzhang6@ustc.edu.cn

**Abstract**—With the development of deep learning, neural network-based speech enhancement (SE) models have shown excellent performance. Meanwhile, it was shown that the development of self-supervised pre-trained models can be applied to various downstream tasks. In this paper, we will consider the application of the pre-trained model to the real-time SE problem. Specifically, the encoder and bottleneck layer of the DEMUCS model are initialized using the self-supervised pre-trained WavLM model, the convolution in the encoder is replaced by causal convolution, and the transformer encoder in the bottleneck layer is based on causal attention mask. In addition, as discretizing the noisy speech representations is more beneficial for denoising, we utilize a quantization module to discretize the representation output from the bottleneck layer, which is then fed into the decoder to reconstruct the clean speech waveform. Experimental results on the Valentini dataset and an internal dataset show that the pre-trained model based initialization can improve the SE performance and the discretization operation suppresses the noise component in the representations to some extent, which can further improve the performance.

## I. INTRODUCTION

Speech enhancement (SE) aims to improve the perceptual speech quality and intelligibility by removing the background noises contained in the noisy input signal, which is usually exploited as a front-end pre-processing module in many applications, such as automatic speech recognition (ASR), speaker diarization, hearing aids [1], [2]. With the development of deep learning, neural network-based SE models [3]–[5] can even outperform the traditional counterparts [6], which can be divided into two categories: time-frequency domain [7]–[10] and time-domain methods [3]–[5]. In the time-frequency domain, a masking matrix is usually first estimated via supervised training and then multiplied with the noisy spectrum to estimate the clean spectrum, which is finally transformed into the time domain to recover the clean speech. For the time-domain SE models, a convolutional encoder-decoder or Unet [11] framework can be utilized to predict the clean speech waveform directly from noisy speech waveforms, where long short-term memory (LSTM) [5] or self-attention [12] is adopted to model the temporal information. It has been experimentally shown that these methods can improve the speech quality/intelligibility to some extent.

In the speech community, self-supervised pre-training models have been developed rapidly recently, which are pre-trained using large amounts of unlabeled data and then transferred to downstream tasks, e.g., ASR, speaker recognition. For example, contrastive predictive coding (CPC) [13] was proposed to predict the future frames using a contrastive loss. Wav2vec2.0 [14] leverages contextual information to predict the information of the masked frames using a contrastive loss function. HuBERT [15] performs offline clustering for the representation output from the middle layer of the model, which enables to directly predict the clustering labels at the masked positions. On the basis of HuBERT, WavLM [16] can achieve the state-of-the-art performance on SUPERB benchmark [17] by using an utterance mixing training and more unlabeled data. It was shown that there are some other pre-trained models, e.g., [18]–[20] that are beneficial for downstream tasks in both clean and noisy scenes.

Further, self-supervised pre-trained models can also be used to improve the SE performance in literature. In [21], thirteen pre-trained models are applied to the SE task, which are taken as feature extractors to generate spectral masks to reconstruct the clean speech waveform. It was shown that the high-level features extracted by self-supervised pre-trained models are also applicable to SE compared to traditional acoustic features. In [22], the SE performance is improved by using a combination of features extracted by a self-supervised model and traditional spectral features. However, currently these methods can only be applied to offline SE tasks and few of them consider the application of self-supervised pre-trained models to the real-time case. In addition, as it was shown in [23], [24] that representation clustering can improve the noise robustness of the ASR, inspired by the clustering representation approach in HuBERT, it might be the case that discretizing the noisy speech representations is more beneficial for denoising as well as the reconstruction of clean speech waveforms.

In this paper, we therefore consider the application of a self-supervised pre-trained model to the real-time SE task. The overall configuration of the proposed model is shown in Fig. 1. The basis SE model adopts the Unet-based framework by initializing the encoder and bottleneck layers of the model using the pre-trained WavLM model and replacing the convolution of the encoder in the WavLM with the causal convolution. The causal attention mask is adopted for the Transformer

<sup>0</sup>Jie Zhang is the corresponding author.

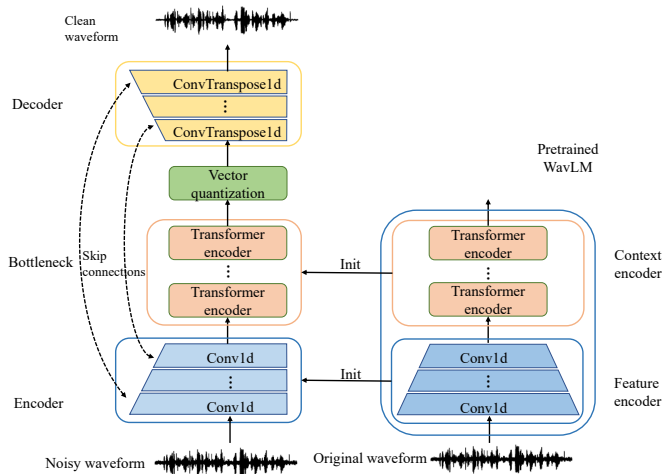


Fig. 1. An illustration of the proposed speech enhancement model.

encoder in the bottleneck layer. In addition, We utilize a vector quantization (VQ) module to discretize the representation output from the bottleneck layer, which is then fed into the decoder for reconstructing clean speech waveforms. Experimental results on the Valentini dataset [25] and an internal dataset show that using the pre-trained model can improve the SE performance, and the proposed VQ module can further improve the performance. Some enhanced speech samples and reference code can be found at: <https://zxy0001.github.io>.

## II. METHOD

### A. The self-supervised pre-trained WavLM model

WavLM [16] is a pre-training model based on HuBERT, which contains a convolutional encoder and a Transformer context encoder, where the convolutional encoder has several layers (each contains a time-domain one dimensional (1D) convolutional layer), a normalization layer, and a Gaussian error linear unit (GELU) activation layer. In the Transformer context encoder, relative position information is introduced for the attention network by gated relative position encoding to better model the local information. In the pre-training stage, WavLM randomly transforms the input speech waveform, e.g., by mixing two waveforms or adding background noise. Afterwards, about 49% of the speech signal is masked randomly in a sentence, and the discrete labels corresponding to the masked positions are predicted at the Transformer output, where the discrete labels are generated by discretizing the continuous signal via K-means clustering. It was shown that WavLM can improve the performance of SUPERB benchmarks for various speech tasks with more unlabeled data and model parameters. For more details on WavLM, please refer to [16].

### B. The model architecture

The proposed SE model is built based on the U-net structure, which contains an  $encoder : X \mapsto Z$ , a bottleneck  $f : Z \mapsto C$ , a  $VQ : C \mapsto Q$  and a  $decoder : Q \mapsto Y$ , and the corresponding model structure is shown in Fig. 1. The encoder

has  $D$  layers, and each contains a 1D causal convolution layer, a normalization layer and a GELU activation layer, where the convolution kernel size is  $K$ , the stride size is  $S$  and the number of channels is  $H$ . The decoder also has  $D$  layers, and each contains a 1D causal transpose convolutional layer, a normalization layer and a GELU activation layer. A skip connection is employed between the output of the  $i$ -th encoder and the input of the  $i$ -th decoder. The bottleneck contains  $N$  Transformer encoder layers, each of which contains a multi-head self-attention layer and a position-wise fully connected feed-forward layer. Skip connection and layer normalization are also utilized for each layer.

Specifically, given a noisy speech waveform  $\mathbf{x}$ , the feature  $\mathbf{z} = encoder(\mathbf{x})$  is obtained by the encoder, which is then input to the bottleneck layer to obtain the contextual representation  $\mathbf{c} = f(\mathbf{z})$ . Given the context representation  $\mathbf{c}$ , the discrete representation is calculated by the VQ module as  $\mathbf{q} = VQ(\mathbf{c})$ , which is finally exploited by the decoder to reconstruct the enhanced speech waveform, i.e.,  $\hat{\mathbf{y}} = decoder(\mathbf{q})$ . As shown in Fig. 1, in order to ensure the causality of the proposed SE model (i.e., at time-frame  $t$  only the information of previous frames can be used), we adopt a causal attention masking matrix to mask out the information after time  $t$ . Since the convolutional receptive fields of the WavLM model and the SE model are different, we only utilize a few layers of the WavLM-base<sup>1</sup> model to initialize the encoder and bottleneck of the SE model.

We obtain a finite set of representations by discretizing the bottleneck layer output via production quantization [26], which is involved similarly to [14]. Given  $G$  codebooks, each with  $V$  learnable  $d$ -dimensional codewords, we first map the bottleneck representations  $\mathbf{c}$  to  $\mathbf{l} \in \mathbb{R}^{G \times V}$  logits, and then select the discrete vectors by the Gumbel-softmax [27] operation in a differential way. The probability of choosing the  $v$ -th codewords from the  $g$ -th codebook is given by

$$\bar{p}_{g,v} = \frac{\exp(\bar{l}_{g,v} + n_v)/\tau}{\sum_{k=1}^V \exp(\bar{l}_{g,k} + n_k)/\tau}, \quad (1)$$

where  $\tau$  is a non-negative temperature coefficient,  $n_v = -\log(-\log(u))$  with  $u$  uniformly distributed over  $[0, 1]$ . In the forward stage, the vector  $i$  is selected as  $i = \operatorname{argmax}_v p_{g,v}$ , and in the backward propagation stage the true gradient of the Gumbel-softmax output is utilized. The quantization expects to utilize more codewords by maximizing the softmax distribution  $\mathbf{l}$ , so we use the diversity loss function  $L_d$ , which is given by

$$L_d = \frac{1}{GV} \sum_{g=1}^G \sum_{v=1}^V \bar{p}_{g,v} \log \bar{p}_{g,v}. \quad (2)$$

### C. The total loss function

The total loss function  $L_{total}$  consists of an SE loss  $L_{se}$  and the diversity loss  $L_d$ , which reads

$$L_{total} = L_{se} + \lambda L_d, \quad (3)$$

<sup>1</sup><https://github.com/microsoft/unilm/tree/master/wavlm>

where  $\lambda$  is the balancing hyper-parameter. Note that the SE loss function  $L_{se}$  contains both the time-domain loss and the frequency-domain loss, which is similarly derived as in [5], [12], where the time-domain component adopts  $\ell_1$  loss and the frequency-domain component utilizes the multi-resolution STFT loss. As a result,  $L_{se}$  can be formulated as

$$L_{se} = \frac{1}{T} \left( \|\mathbf{y} - \hat{\mathbf{y}}\|_1 + \sum_{i=1}^M L_{stft}^{(i)}(\mathbf{y}, \hat{\mathbf{y}}) \right), \quad (4)$$

where

$$L_{stft}(\mathbf{y}, \hat{\mathbf{y}}) = L_{sc}(\mathbf{y}, \hat{\mathbf{y}}) + L_{mag}(\mathbf{y}, \hat{\mathbf{y}}), \quad (5)$$

$$L_{sc}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{\| |\text{STFT}(\mathbf{y})| - |\text{STFT}(\hat{\mathbf{y}})| \|_F}{\| \text{STFT}(\mathbf{y}) \|_F}, \quad (6)$$

$$L_{mag}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{T} \|\log |\text{STFT}(\mathbf{y})| - \log |\text{STFT}(\hat{\mathbf{y}})|\|_1. \quad (7)$$

where  $M$  denotes the number of STFT loss functions and  $\|\cdot\|$  the Frobenius norm. The multi-resolution loss  $L_{stft}^{(i)}$  utilizes the STFT loss with the number of FFT bins ranging from  $\{512, 1024, 2048\}$ , hop size from  $\{50, 120, 240\}$  and window length from  $\in \{240, 600, 1200\}$ , respectively.

### III. EXPERIMENTAL SETUP

#### A. Datasets and evaluation metrics

The Valentini [25] dataset contains 28.4 hours of clean and noisy speech pairs at a sampling rate of 48 kHz. These data are collected from 84 speakers at 4 signal-to-noise ratios (SNRs) (i.e., 0, 5, 10 and 15 dB) in the training set and 4 SNRs (2.5, 7.5, 12.5 and 17.5 dB) in the test set. We downsampled the raw speech waveforms to 16kHz and then applied the same remix augmentation and bandmask augmentation methods as in [5]. In addition, we also validate the proposed model using the internal challenging 100-hour noisy data.

In addition, in order to show the generalizability of the proposed model, we also evaluate the performance on a self-built dataset, which is more challenging than the Valentini dataset in principal. The training set of the internal dataset involves 100 hours data from Librispeech clean and noisy speech [28], where the noisy mixtures are generated by mixing the clean speech and the noise sources from the Freesound [29] noise dataset at a SNR randomly chosen from  $[0, 10]$  dB. The testing set has 15 hours speech consisting of 5000 utterances from 10 speakers recorded in the real traffic and in-car environments. The sampling frequency is 16 kHz.

The performance of the proposed method is evaluated using objective metrics, including 1) perceptual evaluation of speech quality (PESQ) [30], 2) short-time objective intelligibility (STOI) [31], 3) mean opinion score (MOS) prediction of distortion of speech signal (SIG) [32], 4) MOS prediction of intrusiveness of background noise (BAK) [32], 5) MOS prediction of overall quality (OVRL) [32]. Some enhanced speech samples and reference code can be downloaded from <https://zxy0001.github.io>.

#### B. Model configuration

In order to enable the initialization of the proposed SE model using the pre-trained WavLM, the entire SE model utilizes a 3-layer ( $D = 3$ ) 1D convolution for the encoder, a 2-layer ( $N = 2$ ) Transformer encoder for the bottleneck and a 3-layer 1D transposed convolution for the decoder. Both the encoder and decoder have a dimension of 512 ( $H = 512$ ). The Transformer encoder in the bottleneck layer has a dimension of 768, the feedforward neural network has a dimension of 2048, and the self-attention module has 12 heads. For the 1D convolution in encoder, the convolution kernel size and stride size are (10, 3, 3) and (5, 2, 2), respectively. For the VQ module, we adopt  $G = 1$  codebook with  $V = 320$  learnable 128-dimensional vectors in each codebook, and  $\lambda$  in (3) is set to be 0.01 due to the range of the diversity loss. We use the Adam optimizer to train the SE model for 1M iterations, where the batch size is 64 and the maximum learning rate is  $2 \times 10^{-4}$ . All models are trained on 4 Tesla-V100-32G GPUs.

### IV. EXPERIMENTAL RESULTS

**Comparison methods:** We first measure objective evaluation metrics on the noisy test set as a baseline. The time-domain SE comparison methods include SEGAN [33], SEGAN-D [35] and MetricGAN [36], which are based on the generative adversarial networks (GAN). The time-frequency domain SE comparison approaches include MMSE-GAN [8], which is also relied on the GAN model, and DeepMMSE [37], which estimates the noise spectral density using the minimum mean square error (MMSE) criterion. Wave U-Net [34], DEMUCS [5] and CleanUNet [12] are also compared, which utilize the Unet network for the time-domain SE.

Table I shows the experimental results of the comparison methods on the Valentini dataset. It is clear that the proposed model achieves a PESQ of 3.13 and an STOI of 96.1%, respectively using the pre-trained WavLM for initialization and the extra quantization module, which outperforms DEMUCS and CleanUNet. For the real-time SE case, our model achieves a PESQ of 3.02 and an STOI of 95.8%, respectively, which are also better than the best off-the-shelf models, e.g., DEMUCS and CleanUNet. As the average input SNR of the Valentini dataset is moderately high, in order to show the generality of the proposed method we employed a more challenging internal dataset for testing, and the obtained results are shown in Table II. As CleanUNet and DEMUCS obtain the best performance among all chosen comparison methods, we will only compare them with the proposed method in the sequel. We can see that the performance of the proposed method with random initialization is comparable to that of the CleanUNet model, since both employ self-attention modules as bottleneck and the amounts of model parameters are roughly equal. In case WavLM is adopted to initialize our model, the performance is clearly improved compared to the random initialization. Note that using the VQ module in addition to the WavLM-based initialization does not always mean a positive performance gain, which should be dependent on the number of learnable

TABLE I  
PERFORMANCE COMPARISON OF OBJECTIVELY EVALUATION METRICS ON THE VALENTINI TEST SET.

Model	Domain	PESQ (WB)	STOI (%)	pred. CSIG	pred. CBAK	pred. COVL	Causality
Noisy	-	1.97	92.1	3.35	2.44	2.63	-
SEGAN [33]	waveform	2.16	-	3.48	2.94	2.80	No
Wave U-Net [34]	waveform	2.40	-	3.52	3.24	2.96	No
SEGAN-D [35]	waveform	2.39	-	3.46	3.11	3.50	No
MMSE-GAN [8]	time-frequency	2.53	93.0	3.80	3.12	3.14	No
MetricGAN [36]	waveform	2.86	-	3.99	3.18	3.42	No
DeepMMSE [37]	waveform	2.95	94.0	4.28	3.46	3.64	No
DEMUCS [5]	waveform	3.07	95.0	4.31	3.40	3.63	No
CleanUNet [12]	waveform	3.09	95.8	4.38	3.47	3.69	No
Ours (D=3, N=2)	waveform	<b>3.13</b>	<b>96.1</b>	<b>4.46</b>	<b>3.56</b>	<b>3.82</b>	No
Wiener	-	2.22	93.0	3.23	2.68	2.67	Yes
DeepMMSE [37]	waveform	2.77	93.0	4.14	3.32	3.46	Yes
DEMUCS [5]	waveform	2.93	95.0	4.22	3.25	3.52	Yes
CleanUNet [12]	waveform	2.91	95.6	4.34	3.42	3.65	Yes
Ours (D=3, N=2)	waveform	<b>3.02</b>	<b>95.8</b>	<b>4.40</b>	<b>3.49</b>	<b>3.72</b>	Yes

TABLE II  
PERFORMANCE COMPARISON OF OBJECTIVE EVALUATION METRICS FOR REAL-TIME SPEECH ENHANCEMENT ON THE INTERNAL TEST SET.

Model	PESQ (WB)	STOI (%)	pred. CSIG	pred. CBAK	pred. COVL
Noisy	1.32	80.8	2.54	1.87	1.86
DEMUCS [5]	1.89	86.2	3.25	2.37	2.50
CleanUNet [12]	2.17	87.1	3.61	2.54	2.62
Ours (random init)	2.04	86.5	3.53	2.49	2.55
Ours (Init)	2.29	88.1	3.70	2.61	2.72
Ours (Init, V=160)	2.20	87.3	3.65	2.59	2.65
Ours (Init, V=320)	2.35	88.6	3.79	2.66	2.78
Ours (Init, V=480)	<b>2.37</b>	<b>88.9</b>	<b>3.80</b>	<b>2.67</b>	<b>2.81</b>

codewords. When the context representation is quantized with sufficient codewords, the proposed VQ module can further improve the SE performance in terms of all metrics (e.g., Init V = 480 vs. Init). This is due to the fact that in case the number of codewords is small, the quantization noise will heavily affect the context representation as well as the diversity loss function and a certain loss in speech context information becomes inevitable. In general, the less the codewords, the more the information loss and the higher the quantization noise variance.

This can also be seen from different choices of V in Table II, as the SE performance of the proposed model decreases in case the number of learnable vectors decreases. Finally, we visualize the filterbank features of speech signals in Fig. 2, from which it can be clearly seen that SE can remove the noise component from the noisy speech to some extent. In addition, for the causal models we evaluate the real-time factor (RTF), which is defined as the required processing time of unit-second speech. The RTF is computed on a 12-core Intel E5-2680 v3 CPU. We find that DEMUCS, CleanUet and the proposed method obtain a comparable RTF of 0.66, which tells that the proposed model does not increase the time complexity.

### V. CONCLUSIONS

In this paper, we investigated the effect of using pre-trained models for initialization and vector quantization for context representations on the real-time SE problem. It was shown

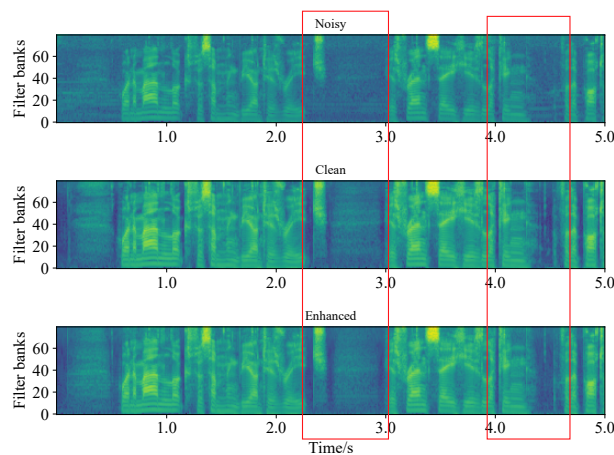


Fig. 2. A filterbank feature of an example speech.

that the pre-trained WavLM based initialization for the encoder and bottleneck layers of the SE model can always lead to a performance improvement compared to the random initialization, while the utilization of the VQ module should be careful. Only when the size of representation codebook is large enough (as the quatization noise does not dominate), it can further improve the SE performance on the basis of the WavLM-based initialization. In principal, applying the quantization module to discretize the representation of the bottleneck layer output functions as a preliminary noise reduction processor on the speech representation.

### ACKNOWLEDGMENT

This work was supported by the National Natural Science Foundation of China (62101523), Hefei Municipal Natural Science Foundation (2022012) and Fundamental Research Funds for the Central Universities.

### REFERENCES

[1] Lei Sun, Jun Du, Xueyang Zhang, Tian Gao, Xin Fang, and Chin-Hui Lee, "Progressive multi-target network based speech enhancement with snr-preselection for robust speaker diarization," in *IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 7099–7103.

- [2] Jie Zhang and Guanghui Zhang, "A parametric unconstrained beamformer based binaural noise reduction for assistive hearing," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 292–304, 2022.
- [3] Yi Luo and Nima Mesgarani, "Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [4] Ashutosh Pandey and DeLiang Wang, "A new framework for cnn-based speech enhancement in the time domain," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 7, pp. 1179–1188, 2019.
- [5] Alexandre Défossez, Gabriel Synnaeve, and Yossi Adi, "Real Time Speech Enhancement in the Waveform Domain," in *ISCA Interspeech*, 2020, pp. 3291–3295.
- [6] S. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 2, pp. 113–120, 1979.
- [7] Yong Xu, Jun Du, Li-Rong Dai, and Chin-Hui Lee, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 23, no. 1, pp. 7–19, 2015.
- [8] Meet H. Soni, Neil Shah, and Hemant A. Patil, "Time-frequency masking-based speech enhancement using generative adversarial network," in *IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 5039–5043.
- [9] Zhong-Qiu Wang, Peidong Wang, and DeLiang Wang, "Complex spectral mapping for single- and multi-channel speech enhancement and robust asr," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 1778–1787, 2020.
- [10] Yanxin Hu, Yun Liu, Shubo Lv, Mengtao Xing, Shimin Zhang, Yihui Fu, Jian Wu, Bihong Zhang, and Lei Xie, "DCCRN: Deep Complex Convolution Recurrent Network for Phase-Aware Speech Enhancement," in *ISCA Interspeech*, 2020, pp. 2472–2476.
- [11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, "U-net: Convolutional networks for biomedical image segmentation," in *MICCAI*. Springer, 2015, pp. 234–241.
- [12] Zhifeng Kong, Wei Ping, Amrith Dantrey, and Bryan Catanzaro, "Speech denoising in the waveform domain with self-attention," in *IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 7867–7871.
- [13] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [14] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proc. of NIPS*, 2020, pp. 12449–12460.
- [15] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
- [16] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, Jian Wu, Long Zhou, Shuo Ren, Yanmin Qian, Yao Qian, Jian Wu, Michael Zeng, Xiangzhan Yu, and Furu Wei, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE J. of Selected Topics in Signal Process.*, pp. 1–14, 2022.
- [17] Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhotia, Yist Y Lin, Andy T Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, et al., "Superb: Speech processing universal performance benchmark," *arXiv preprint arXiv:2105.01051*, 2021.
- [18] Andy T. Liu, Shang-Wen Li, and Hung-yi Lee, "Tera: Self-supervised learning of transformer encoder representation for speech," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 2351–2366, 2021.
- [19] W. Hsu, A. Sriram, A. Baevski, T. Likhomanenko, Q. Xu, V. Pratap, J. Kahn, A. Lee, R. Collobert, G. Synnaeve, et al., "Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training," *arXiv preprint arXiv:2104.01027*, 2021.
- [20] Qiu-Shi Zhu, Jie Zhang, Zi-Qiang Zhang, Ming-Hui Wu, Xin Fang, and Li-Rong Dai, "A noise-robust self-supervised pre-training model based speech representation learning for automatic speech recognition," in *IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 3174–3178.
- [21] Zili Huang, Shinji Watanabe, Shu-wen Yang, Paola García, and Sanjeev Khudanpur, "Investigating self-supervised learning for speech enhancement and separation," in *IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 6837–6841.
- [22] Kuo-Hsuan Hung, Szu-wei Fu, Huan-Hsin Tseng, Hsin-Tien Chiang, Yu Tsao, and Chii-Wann Lin, "Boosting self-supervised embeddings for speech enhancement," *arXiv preprint arXiv:2204.03339*, 2022.
- [23] Yu-An Chung, Yu Zhang, Wei Han, Chung-Cheng Chiu, James Qin, Ruoming Pang, and Yonghui Wu, "w2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training," in *ASRU*, 2021, pp. 244–250.
- [24] Qiu-Shi Zhu, Jie Zhang, Zi-Qiang Zhang, and Li-Rong Dai, "Joint training of speech enhancement and self-supervised model for noise-robust asr," *arXiv preprint arXiv:2205.13293*, 2022.
- [25] C.Valentini-Botinhao, "Noisy speech database for training speech enhancement algorithms and tts models," 2017.
- [26] Herve Jégou, Matthijs Douze, and Cordelia Schmid, "Product quantization for nearest neighbor search," *IEEE TPAMI*, vol. 33, no. 1, pp. 117–128, 2011.
- [27] Eric Jang, Shixiang Gu, and Ben Poole, "Categorical reparameterization with gumbel-softmax," *arXiv preprint arXiv:1611.01144*, 2016.
- [28] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2015, pp. 5206–5210.
- [29] Frederic Font, Gerard Roma, and Xavier Serra, "Freesound technical demo," in *Proc. 21st ACM Int. Conf. Multimedia*, 2013, pp. 411–412.
- [30] I.-T. Recommendation, "Perceptual evaluation of speech quality (pesq): An objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codecs," *Rec. ITU-T P. 862*, 2001.
- [31] Cees H. Taal, Richard C. Hendriks, Richard Heusdens, and Jesper Jensen, "An algorithm for intelligibility prediction of time-frequency weighted noisy speech," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 7, pp. 2125–2136, 2011.
- [32] Yi Hu and Philippos C. Loizou, "Evaluation of objective quality measures for speech enhancement," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 16, no. 1, pp. 229–238, 2008.
- [33] Santiago Pascual, Antonio Bonafonte, and Joan Serra, "Segan: Speech enhancement generative adversarial network," in *ISCA Interspeech*, 2017, pp. 3642–3646.
- [34] Craig Macartney and Tillman Weyde, "Improved speech enhancement with the wave-u-net," *arXiv preprint arXiv:1811.11307*, 2018.
- [35] Huy Phan, Ian V. McLoughlin, Lam Pham, Oliver Y. Chén, Philipp Koch, Maarten De Vos, and Alfred Mertins, "Improving gans for speech enhancement," *IEEE Signal Process. Lett.*, vol. 27, pp. 1700–1704, 2020.
- [36] Szu-Wei Fu, Chien-Feng Liao, Yu Tsao, and Shou-De Lin, "MetricGAN: Generative adversarial networks based black-box metric scores optimization for speech enhancement," in *ICML*. 2019, vol. 97, pp. 2031–2041, PMLR.
- [37] Qiquan Zhang, Aaron Nicolson, Mingjiang Wang, Kuldip K. Paliwal, and Chenxu Wang, "Deepmmse: A deep learning approach to mmse-based noise power spectral density estimation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 1404–1415, 2020.