



# Noise-Robust Voice Conversion by Conditional Denoising Training Using Latent Variables of Recording Quality and Environment

Takuto Igarashi<sup>1</sup>, Yuki Saito<sup>1</sup>, Kentaro Seki<sup>1</sup>, Shinnosuke Takamichi<sup>1,3</sup>, Ryuichi Yamamoto<sup>2</sup>, Kentaro Tachibana<sup>2</sup>, Hiroshi Saruwatari<sup>1</sup>

<sup>1</sup>The University of Tokyo, Japan, <sup>2</sup>LY Corp., Japan, <sup>3</sup>Keio University, Japan

takuto0228@g.ecc.u-tokyo.ac.jp

## Abstract

We propose noise-robust voice conversion (VC) which takes into account the recording quality and environment of noisy source speech. Conventional denoising training improves the noise robustness of a VC model by learning noisy-to-clean VC process. However, the naturalness of the converted speech is limited when the noise of the source speech is unseen during the training. To this end, our proposed training conditions a VC model on two latent variables representing the recording quality and environment of the source speech. These latent variables are derived from deep neural networks pre-trained on recording quality assessment and acoustic scene classification and calculated in an utterance-wise or frame-wise manner. As a result, the trained VC model can explicitly learn information about speech degradation during the training. Objective and subjective evaluations show that our training improves the quality of the converted speech compared to the conventional training.

**Index Terms:** voice conversion, noise-robust voice conversion, denoising training, any-to-any voice conversion, latent variables

## 1. Introduction

Voice conversion (VC) is a technology that converts a source speaker's timbre to that of a target speaker while preserving the linguistic content. VC can be widely applied in the real world, such as in movie dubbing [1], personalized text-to-speech [2], and speaking assistance [3]. As a result of advancements in deep learning, deep neural network (DNN)-based VC methods have significantly improved the quality of converted speech [4].

Typical DNN-based VC methods train a VC model with a large multi-speaker corpus containing high-quality speech from a variety of speakers. However, actual speech samples recorded in the real world are often degraded by various factors, such as background noise and recording channels. Huang et al. [5] empirically demonstrated that the recording-quality mismatch of input speech between the training and inference (i.e., clean and noisy) significantly deteriorates VC performance.

Denoising training (DT) [6] is one promising approach to achieving noise-robust DNN-based VC. In DT, a VC model is trained using pseudo-noisy (i.e., artificially degraded) speech, with the aim of implicitly denoising input speech during the VC process. Although this training mitigates the distortion of converted speech caused by noisy input speech in inference, the naturalness of the converted speech is still limited when the degradation factor of input speech is unseen during the training. One primary reason is that the trained VC model does not explicitly learn information about speech degradation, such as noise characteristics (e.g., stationary or non-stationary) and noise levels, and does not guarantee generalization performance to various degradation factors.

In this study, we propose *conditional DT (CDT)*, an improved version of conventional (i.e., unconditional) DT, to improve the noise robustness of VC towards unseen degradation. CDT conditions a DNN-based VC model on two latent variables regarding the degradation of input speech: recording quality and environment. These latent variables are derived from deep neural networks pre-trained on recording quality assessment and acoustic scene classification and calculated in an utterance-wise or frame-wise manner. As a result, the trained VC model can explicitly learn information about speech degradation during the training. We present the CDT framework using NISQA [7] and PaSST [8] as representative models to extract latent variables of the recording quality and environment, respectively. In the experimental evaluation, we validate the effectiveness of CDT using S2VC [9] as the baseline VC model following the conventional DT framework [6]. Our contributions are summarized as follows.

- We propose *CDT* that can improve the noise-robustness of DNN-based VC models by explicitly conditioning the model on speech degradation information.
- We present two conditioning strategies with an utterance-wise and frame-wise manner so that the trained VC model can take into account not only global but also time-variant characteristics of degradation in input speech.
- From the evaluation results, we show that conditioning the VC model on frame-wise latent variables of recording environment is essential for improving the naturalness of the converted speech in the noisy-to-clean VC scenario.

## 2. Conventional DT-based noise-robust VC

### 2.1. DT algorithm

DT is an end-to-end learning method for noisy-to-clean VC, in which data augmentations are utilized to train a noise-robust VC model. In the learning process, pseudo-noisy speech is artificially generated from clean speech by mixing various environmental noises with random signal-to-noise ratios (SNRs) and fed into the VC model as input. The training objective function is computed by comparing the ground-truth clean speech and converted speech, i.e., the VC model's output. Typically, the L1 or L2 loss between mel-spectrograms of ground-truth clean speech and converted speech is used as the objective function. Thus, DT can be interpreted as a learning method in which the VC model acts as a denoising autoencoder [6].

Huang et al. [6] demonstrated that the noise robustness of VC models trained by DT improved when the VC models were based on an autoencoding process such as AdaIN-VC [10] and S2VC [9], and S2VC with DT was the most effective for VC.

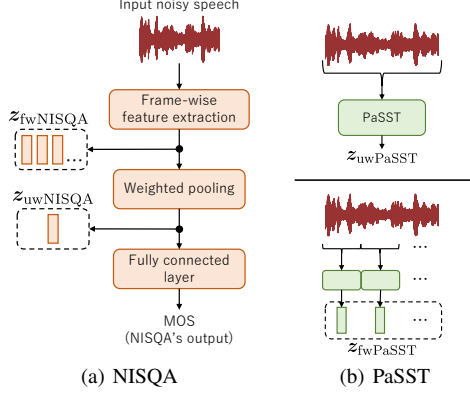


Figure 1: Latent variables extracted by NISQA and PaSST in an utterance-wise or frame-wise manner.

## 2.2. Limitations

Another approach to achieving noise-robust VC without modifying the model structure can be the concatenation of a pre-trained speech enhancement model with an existing any-to-any VC model. However, this approach tends to limit the VC performance compared to an end-to-end approach such as DT. This is because the artifacts caused by speech enhancement to suppress unseen noises negatively affects the downstream VC task [11].

Although DT is an end-to-end approach, the naturalness of the converted speech is still limited when the noise of the input speech is unseen during the training. One of the main reasons is that the trained VC model does not explicitly learn information about speech degradation, such as noise characteristics and noise levels, and does not guarantee generalization to various degradation factors.

## 3. Proposed CDT-based noise-robust VC

### 3.1. Motivation

One possible solution to the problem presented in Section 2.2 is to enable the DNN-based VC model to explicitly learn information about speech degradation. To this end, we propose CDT, which conditions the model on two latent variables regarding the degradation of input speech: recording quality and environment. For practicality, our CDT algorithms assume that only the source speech samples are degraded by noise. This is because it is more challenging for VC users, who are not always experts in speech technology or high-quality speech recording, to record their own clean source speech than it is to collect target speakers' clean speech. Nevertheless, we believe that our algorithm can be extended to VC where both the source and target speech are degraded.

### 3.2. CDT algorithm

Let  $\mathbf{x}^c$  be any clean speech in the training dataset. The loss function  $\mathcal{L}$  of our CDT is defined as follows:

$$\mathbf{x}^s = \mathbf{x}^c + \mathbf{n}, \quad (1)$$

$$\mathbf{x}^t = \mathbf{x}^c, \quad (2)$$

$$\mathbf{z}_{\text{SSL}}^s = f_{\text{SSL}}(\mathbf{x}^s), \quad (3)$$

$$\mathbf{z}_{\text{rqa}}^s = f_{\text{rqa}}(\mathbf{x}^s), \quad (4)$$

$$\mathbf{z}_{\text{asc}}^s = f_{\text{asc}}(\mathbf{x}^s), \quad (5)$$

$$\mathbf{z}_{\text{SSL}}^t = f_{\text{SSL}}(\mathbf{x}^t), \quad (6)$$

$$\mathcal{L} = |f_{\theta}(\mathbf{z}_{\text{SSL}}^s, \mathbf{z}_{\text{rqa}}^s, \mathbf{z}_{\text{asc}}^s, \mathbf{z}_{\text{SSL}}^t) - g_{\text{mel}}(\mathbf{x}^c)|, \quad (7)$$

where  $\mathbf{x}^s$  and  $\mathbf{x}^t$  are the source and target speech, respectively. The pseudo-noisy source speech  $\mathbf{x}^s$  is calculated by adding noise  $\mathbf{n}$  to the clean speech  $\mathbf{x}^c$ .  $g_{\text{mel}}(\cdot)$  is a function which calculates the log mel-spectrogram from input speech.  $f_{\text{SSL}}(\cdot)$  is a pre-trained SSL model that extracts intermediate feature representations from the source and target speech used for the VC process, i.e.,  $\mathbf{z}_{\text{SSL}}^s$  and  $\mathbf{z}_{\text{SSL}}^t$ . The source speaker's conditional latent variables for recording quality and environment, i.e.,  $\mathbf{z}_{\text{rqa}}^s$  and  $\mathbf{z}_{\text{asc}}^s$ , are extracted by pre-trained DNNs for recording quality assessment  $f_{\text{rqa}}(\cdot)$  and acoustic scene classification  $f_{\text{asc}}(\cdot)$ , respectively.  $f_{\theta}(\cdot)$  is a DNN-based model parameterized by  $\theta$ , which predicts the target speaker's log mel-spectrogram from the input features. The model parameter  $\theta$  is updated to minimize the log mel-spectrogram prediction error shown in Eq. (7).

We present the CDT framework using NISQA [7] and PaSST [8] as representative models to extract latent variables of the recording quality and environment:  $\mathbf{z}_{\text{rqa}}^s$  and  $\mathbf{z}_{\text{asc}}^s$ , respectively.

**NISQA:** NISQA is a method for automatically estimating recording quality scores without reference speech. The open-source model is trained on pseudo-noisy and real-noisy speech, allowing it to robustly predict recording quality values against a variety of noises.

**PaSST:** PaSST is a method for obtaining audio tags from input sounds. The model trained on AudioSet [12], which contains 527 types of tags, achieves significantly higher predictive performance compared to other models [8].

### 3.3. Frame-wise conditioning

NISQA and PaSST were originally designed to output an utterance-wise prediction: the mean opinion score (MOS) and audio tag, respectively. However, frame-wise features can be obtained as described in the next paragraph. The frame-wise features can be expected to represent the non-stationary characteristics of noise in the noisy source speech.

Let  $(\mathbf{z}_{\text{uwnISQA}}, \mathbf{z}_{\text{uwnPaSST}})$  and  $(\mathbf{z}_{\text{fwnISQA}}, \mathbf{z}_{\text{fwnPaSST}})$  be latent variables extracted by (NISQA, PaSST) in an utterance-wise and frame-wise manner, respectively. Figure 1 shows the feature extraction methods. As shown in Figure 1(a), NISQA's model segments an input speech, estimates frame-wise features  $\mathbf{z}_{\text{fwnISQA}}$ , computes their weighted average  $\mathbf{z}_{\text{uwnISQA}}$  along with the frame axis, and outputs the final prediction (MOS of input speech). In contrast, PaSST's model outputs audio tags with no dimensions in the frame direction, as shown in the upper part of Figure 1(b). However, we can also use the model to extract frame-wise audio tags by segmenting the input speech in advance [13]. The former extracts  $\mathbf{z}_{\text{uwnPaSST}}$  and the latter extracts  $\mathbf{z}_{\text{fwnPaSST}}$ .

## 4. Experimental evaluation

### 4.1. Experimental conditions

We used the parallel100 subset from Japanese Versatile Speech (JVS) [14] and downsampled all speech data to 16 kHz. The subset contains 22 hours of speech data for 100 Japanese speakers (100 utterances per speaker). The numbers of speakers included in the training, validation, and evaluation data were 90 (“jvs001” to “jvs086”), 4 (“jvs087” to “jvs090”), and 10 (“jvs091” to “jvs100”), respectively. The 10 test speakers excluded from the training data were used as unseen speakers for the any-to-any VC evaluation.




Figure 2: Conditioning the VC model on latent variables

We created the pseudo-noisy speech dataset by adding noise  $\mathbf{n}$  to the clean speech  $\mathbf{x}^c$  with SNR randomly sampled from a uniform distribution  $U(0, 20)$  [dB]. Here,  $\mathbf{n}$  was taken from the DEMAND [15] noise and the WHAM!48kHz [16] noise for the training and evaluation datasets, respectively. Thus, the noise of speech from the evaluation dataset was unseen in the training process. We downsampled the WHAM!48kHz to 16 kHz. Both contained various types of background noises.

Referring to Huang et al.’s study [6], we adopted S2VC [9] as a backbone VC model with the publicly available implementation on GitHub (robust-vc)<sup>1</sup>. In CDT, we conditioned the VC model on ( $\mathbf{z}_{\text{uwNISQA}}$  or  $\mathbf{z}_{\text{fwNISQA}}$ ) and ( $\mathbf{z}_{\text{uwPaSST}}$  or  $\mathbf{z}_{\text{fwPaSST}}$ ) as shown in Figure 2. We unified the size of these latent variables to that of  $\mathbf{z}_{\text{SSL}}$ , i.e.,  $256 \times$  frame size in this study. We concatenated the latent variables along with the feature axis and input them to the source encoder. First, to unify the frame dimension of ( $\mathbf{z}_{\text{uwNISQA}}$ ,  $\mathbf{z}_{\text{uwPaSST}}$ ) to that of  $\mathbf{z}_{\text{SSL}}$ , we replicated them by a factor of the frame size of  $\mathbf{z}_{\text{SSL}}$  along with the frame axis. Meanwhile, we upsampled ( $\mathbf{z}_{\text{fwNISQA}}$ ,  $\mathbf{z}_{\text{fwPaSST}}$ ) to align the frame lengths. Second, to unify the feature dimensions of ( $\mathbf{z}_{\text{uwNISQA}}$ ,  $\mathbf{z}_{\text{fwNISQA}}$ ) and ( $\mathbf{z}_{\text{uwPaSST}}$ ,  $\mathbf{z}_{\text{fwPaSST}}$ ), we inserted the linear projection layers from 64 to 256 dimensions and from 768 to 256 dimensions, respectively.

In CDT, the S2VC model was trained based on the loss function shown in Eq. (7) while the parameters of the pre-trained models,  $f_{\text{SSL}}(\cdot)$ ,  $f_{\text{rqa}}(\cdot)$ , and  $f_{\text{asc}}(\cdot)$ , were fixed. Specifically,  $f_{\text{SSL}}(\cdot)$  is a feature extractor using contrastive predictive coding [17] as is the case of the study of Huang et al. [6]. The latent variable extractors,  $f_{\text{rqa}}(\cdot)$  and  $f_{\text{asc}}(\cdot)$ , were based on the official implementations of NISQA<sup>2</sup> and PaSST<sup>3</sup>, respectively. In the implementation of NISQA, the frame length was 150 ms and the frame shift was 40 ms. In the implementation of PaSST, the former was 160 ms and the latter was 50 ms. To generate the converted speech waveform from the predicted log mel-spectrogram, we used the HiFi-GAN vocoder [18] trained on the JVS training data with a batch size of 16, while all the VC models were trained with a batch size of 6. The optimizer was AdamW [19] with a learning rate of  $5 \times 10^{-5}$ , and  $\beta_1 = 0.9, \beta_2 = 0.999$ . Each of the VC models had approximately 33 million parameters which were randomly initialized. The training was stopped when the validation loss converged completely and the training time was approximately 60 hours. We trained the VC model using both the conventional DT [6] and our CDT frameworks and evaluated their effectiveness.

## 4.2. Objective evaluation

We randomly selected 250 test pairs of source and target speech samples with different speakers taken from the evaluation dataset (10 JVS speakers). Then, we performed VC on

<sup>1</sup><https://github.com/cyhuang-tw/robust-vc>

<sup>2</sup><https://github.com/gabrielmittag/NISQA>

<sup>3</sup>[https://github.com/kkoutini/passt\\_hear21](https://github.com/kkoutini/passt_hear21)

Table 1: Average CER and SECS of the 250 samples converted by the VC models trained with conditional DT and CDT. The method without any conditioning corresponds to conventional DT. “uw” means utterance-wise and “fw” means frame-wise. Values in the parentheses indicate standard deviations.

Method		CER [%]	SECS
NISQA	PaSST		
-	-	26.3	0.935 ( $\pm 0.034$ )
uw	uw	27.2	<b>0.938</b> ( $\pm 0.036$ )
uw	fw	24.6	0.935 ( $\pm 0.034$ )
fw	uw	24.7	0.931 ( $\pm 0.036$ )
fw	fw	<b>23.3</b>	0.934 ( $\pm 0.033$ )

each pair with the VC models trained by the conventional DT and CDT, which took less than 80 ms. Our CDT methods are denoted as “uwNISQA-uwPaSST”, “uwNISQA-fwPaSST”, “fwNISQA-uwPaSST”, and “fwNISQA-fwPaSST”, depending on whether the model was conditioned on utterance-wise or frame-wise features.

To evaluate the intelligibility of the converted speech, we used the character error rate (CER) estimated by the automatic speech recognition (ASR) system. The ASR model for calculating CER is a pre-trained ReasonSpeech model available on HuggingFace<sup>4</sup>. The smaller the CER, the more the converted speech preserves the linguistic content of the source speech, and the larger the CER, the more severely distorted the converted speech. Thus, CER is regarded as a measure reflecting the intelligibility of the converted speech. In contrast, to measure the speaker similarity between the converted speech and the target speaker, we used the speaker embedding cosine similarity (SECS). To compute SECS, we generated two fixed-dimensional embedding vectors representing the speaker identity of the converted and target speech and computed their cosine similarity. The higher the SECS, the more similar the converted and target speech are in terms of the speaker identity. We adopted x-vector [20] extracted by using a pre-trained model of WavLM [21] as the embedding vectors, which is available on HuggingFace<sup>5</sup>.

Table 1 shows the average CER and SECS of the converted speech samples corresponding to the 250 test pairs. As shown, SECS is almost constant across methods, which may be because the VC models trained by CDT received the same information on target speech as the conventional DT when only the source speech was noisy. In contrast, CER differed between the methods. Although NISQA and PaSST were originally designed to output utterance-wise prediction, we can see from Table 1 that conditioning the VC model on the utterance-wise features is not very effective for improving CER. In particular, the CER of

<sup>4</sup><https://huggingface.co/reason-research/reasonspeech-espnet-next>

<sup>5</sup><https://huggingface.co/microsoft/wavlm-base-sv>

Table 2: *Naturalness and speaker similarity of the speech samples converted by conventional DT and CDT. Values in the parentheses indicate 95% confidential intervals of the scores. Bold scores are the highest among the five compared methods.*

Method		Naturalness	Speaker similarity
NISQA	PaSST		
-	-	2.75 ( $\pm 0.085$ )	2.43 ( $\pm 0.082$ )
uw	uw	2.67 ( $\pm 0.084$ )	2.39 ( $\pm 0.082$ )
uw	fw	2.84 ( $\pm 0.087$ )	<b>2.50</b> ( $\pm 0.082$ )
fw	uw	2.74 ( $\pm 0.086$ )	2.43 ( $\pm 0.083$ )
fw	fw	<b>2.85</b> ( $\pm 0.086$ )	2.47 ( $\pm 0.082$ )

Table 3: *Preference AB test results for naturalness of converted speech for any pair of (Conventional DT, uwNISQA-fwPaSST, fwNISQA-fwPaSST). “Con. DT” means conventional DT, “(uw, fw)” means uwNISQA-fwPaSST, and “(fw, fw)” means fwNISQA-fwPaSST.*

A vs B	Naturalness	p-value
Con. DT vs (uw, fw)	0.414 vs <b>0.586</b>	$< 10^{-6}$
Con. DT vs (fw, fw)	0.442 vs <b>0.558</b>	$< 10^{-3}$
(uw, fw) vs (fw, fw)	<b>0.514</b> vs 0.486	0.38

“uwNISQA-uwPaSST” is lower than that of the conventional DT. This may be because the utterance-wise features consist of the exact same matrices along the frame axis. Although they are quite large, they do not contain much useful information for VC. Thus, they can prevent the VC model from receiving information that is essential for VC. In contrast, conditioning the model on frame-wise features improved CER. This implies that the frame-wise features represent the non-stationary characteristics of noise in the noisy source speech and that the VC model leverages useful information from the conditional latent variables to improve VC performance.

### 4.3. Subjective evaluation

We conducted subjective evaluations using crowdsourcing on Lancers<sup>6</sup> regarding the naturalness and speaker similarity of converted speech. We combined every 250 converted speech samples generated in Section 4.2 into a single dataset and used it as the dataset for subjective evaluation. Thus, the evaluation dataset contained  $250 \times 5 = 1250$  converted speech samples. When evaluating their naturalness, evaluators were given a converted utterance randomly sampled from the subjective evaluation dataset. Then, they rated the perceptual quality on a 5-point MOS scale from 1 (very bad) to 5 (very good). When evaluating speaker similarity, evaluators were given a ground-truth target utterance and the converted utterance sampled from the JVS corpus and the subjective evaluation dataset, respectively. Then, they answered how similar the speakers were who produced the two utterances on a MOS score ranging from 1 (completely different) to 5 (completely same). In both cases, the listening experiment was repeated 20 times per evaluator. There were 100 evaluators for each subjective evaluation on naturalness and speaker similarity.

As the subjective evaluation results in Table 2 show, “uwNISQA-uwPaSST” demonstrated the lowest VC performance in terms of both naturalness and speaker similarity as is the case of the CER results shown in Table 1. On the other

hand, “uwNISQA-fwPaSST” and “fwNISQA-fwPaSST” outperformed the conventional DT, although there is no statistical significance between the scores.

To investigate the differences between “Conventional DT”, “uwNISQA-fwPaSST”, and “fwNISQA-fwPaSST”, we further conducted preference AB tests to compare the naturalness of the converted speech for any pair of these three methods. Fifty listeners took part in each test, and each listener evaluated 10 pairs of converted speech samples using our crowdsourcing-based evaluation platform. The results are shown in Table 3. The proposed methods (i.e. “uwNISQA-fwPaSST” and “fwNISQA-fwPaSST”) significantly outperformed the conventional DT ( $p < 0.05$ ), but there was no significant difference between “uwNISQA-fwPaSST” and “fwNISQA-fwPaSST”. This indicates that conditioning the VC model on  $z_{fwPaSST}$  is effective for improving the naturalness of the converted speech in the noisy-to-clean VC scenario.

## 5. Discussion

We investigated noise-robust VC from noisy source speech to clean noisy speech in this paper and demonstrated the effectiveness of CDT. We anticipate that CDT should be applicable to VC where the target speech is also noisy. Noisy-to-noisy VC [22, 23], which aims to preserve background noise of source speech during the VC process, is another possible situation in which CDT can be applied.

We used two latent variables, recording quality and environment, as the conditional features on the DT algorithm. Other variables that characterize input speech, such as speech naturalness (e.g., UTMOS) [24] and reverberation (e.g., T60 estimator [25]), can be also introduced to our CDT.

Regarding the training objective function, we only considered the simple L1 loss between the ground-truth and generated mel-spectrograms. We can introduce other machine learning techniques to improve the noise robustness of the trained VC model, such as adversarial training with feature decoupling [26].

## 6. Conclusion

We proposed conditional denoising training (CDT), which conditions a VC model on two latent variables regarding the recording quality and acoustic environment of noisy source speech. We verified the effectiveness of four CDT methods in the cases where these two latent variables were utterance-wise or frame-wise. The objective and subjective evaluations showed that conditioning the VC model on frame-wise features can effectively improve VC performance, while conditioning it on utterance-wise features does not necessarily improve VC performance.

In the future, we will consider VC models that take into account various degradations of input speech, including reverberation and bandwidth rejection, as well as the additive noise considered in this study. In addition, towards real-world applications, we will extend our CDT to address noisy speech recorded by actual devices such as smartphones.

<sup>6</sup><https://www.lancers.jp/>

**Acknowledgements:** This research was conducted as joint research between LY Corporation and Saruwatari-Takamichi Laboratory of The University of Tokyo, Japan. This work was supported by Research Grant S of the Tateishi Science and Technology Foundation.

## 7. References

- [1] F. M. Mukhneri, I. Wijayanto, and S. Hadiyoso, "Voice conversion for dubbing using linear predictive coding and hidden Markov model," *Journal of Southwest Jiaotong University*, vol. 55, no. 4, 2020.
- [2] A. Kain and M. W. Macon, "Spectral voice conversion for text-to-speech synthesis," in *Proc. ICASSP*, Seattle, U.S.A., May 1998, pp. 285–288.
- [3] K. Nakamura, T. Toda, H. Saruwatari, and K. Shikano, "Speaking-aid systems using gmm-based voice conversion for electrolaryngeal speech," *Speech communication*, vol. 54, no. 1, pp. 134–146, Jan. 2012.
- [4] Y. Zhao, W.-C. Huang, X. Tian, J. Yamagishi, R. K. Das, T. Kinunen, Z. Ling, and T. Toda, "Voice Conversion Challenge 2020 — intra-lingual semi-parallel and cross-lingual voice conversion —," in *Proc. Joint Workshop for the Blizzard Challenge and Voice Conversion Challenge 2020 (VCCBC)*, Shanghai, China, Oct. 2020, pp. 80–98.
- [5] T.-H. Huang, J.-H. Lin, and H.-Y. Lee, "How far are we from robust voice conversion: a survey," in *Proc. SLT*, Virtual Conference, Jan 2021, pp. 514–521.
- [6] C.-Y. Huang, K.-W. Chang, and H.-Y. Lee, "Toward degradation-robust voice conversion," in *Proc. ICASSP*, Singapore, May 2022, pp. 6777–6781.
- [7] G. Mittag, B. Naderi, A. Chehadi, and S. Möller, "NISQA: A deep cnn-self-attention model for multidimensional speech quality prediction with crowdsourced datasets," in *Proc. INTERSPEECH*, Brno, Czech Republic, Sep. 2021, pp. 2127–2131.
- [8] K. Koutini, J. Schlüter, H. Eghbal-zadeh, and G. Widmer, "Efficient training of audio Transformers with patchout," in *Proc. INTERSPEECH*, Incheon, South Korea, Sep. 2022, pp. 2753–2757.
- [9] J. Lin, Y. Y. Lin, C.-M. Chien, and H.-Y. Lee, "S2VC: A framework for any-to-any voice conversion with self-supervised pre-trained representations," in *Proc. INTERSPEECH*, Brno, Czech Republic, Sep. 2021, pp. 836–840.
- [10] J.-C. Chou and H.-Y. Lee, "One-shot voice conversion by separating speaker and content representations with instance normalization," in *Proc. INTERSPEECH*, Graz, Austria, Sep. 2019, pp. 664–668.
- [11] C. Zhang, Y. Ren, X. Tan, J. Liu, K. Zhang, T. Qin, S. Zhao, and T.-Y. Liu, "Denoispeech: Denoising text to speech with frame-level noise modeling," in *Proc. ICASSP*, Toronto, Canada, Jun. 2021, pp. 7063–7067.
- [12] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio Set: An ontology and human-labeled dataset for audio events," in *Proc. ICASSP*, Brighton, U.K., Mar. 2017, pp. 776–780.
- [13] J. Turian, J. Shier, H. R. Khan, B. Raj, B. W. Schuller, C. J. Steinmetz, C. Malloy, G. Tzanetakis, G. Velarde, K. McNally *et al.*, "HEAR: Holistic evaluation of audio representations," in *NeurIPS 2021 Competitions and Demonstrations Track*, Dec. 2022, pp. 125–145.
- [14] S. Takamichi, R. Sonobe, K. Mitsui, Y. Saito, T. Koriyama, N. Tanji, and H. Saruwatari, "JSUT and JVS: Free Japanese voice corpora for accelerating speech synthesis research," *Acoustical Science and Technology*, vol. 41, no. 5, pp. 761–768, 2020.
- [15] J. Thiemann, N. Ito, and E. Vincent, "The Diverse Environments Multi-channel Acoustic Noise Database (DEMAND): A database of multichannel environmental noise recordings," in *Proc. ICA*, Montreal, Canada, Jun. 2013.
- [16] G. Wichern, J. Antognini, M. Flynn, L. R. Zhu, E. McQuinn, D. Crow, E. Manilow, and J. Le Roux, "WHAM!: Extending speech separation to noisy environments," in *Proc. INTERSPEECH*, Graz, Austria, Sep. 2019, pp. 1368–1372.
- [17] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [18] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," *Proc. NeurIPS*, vol. 33, pp. 17 022–17 033, Dec. 2020.
- [19] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. ICLR*, New Orleans, U.S.A., May 2019.
- [20] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *Proc. ICASSP*, Calgary, Canada, Apr. 2018, pp. 5329–5333.
- [21] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [22] C. Xie, Y.-C. Wu, P. L. Tobing, W.-C. Huang, and T. Toda, "Direct noisy speech modeling for noisy-to-noisy voice conversion," in *Proc. ICASSP*, Singapore, May 2022, pp. 6787–6791.
- [23] J. Yao, Y. Lei, Q. Wang, P. Guo, Z. Ning, L. Xie, H. Li, J. Liu, and D. Xie, "Preserving background sound in noise-robust voice conversion via multi-task learning," in *Proc. ICASSP*, Rhodes, Greece, Jun. 2023.
- [24] T. Saeki, D. Xin, W. Nakata, T. Koriyama, S. Takamichi, and H. Saruwatari, "UTMOS: UTokyo-SaruLab system for Voice-MOS Challenge 2022," in *Proc. INTERSPEECH*, Incheon, South Korea, Sep. 2022, pp. 4521–4525.
- [25] Y. Choi, C. Xie, and T. Toda, "Reverberation-controllable voice conversion using reverberation time estimator," in *Proc. INTERSPEECH*, Dublin, Ireland, Aug. 2023, pp. 2103–2107.
- [26] L. Chen, X. Zhang, Y. Li, and M. Sun, "Noise-robust voice conversion using adversarial training with multi-feature decoupling," *Engineering Applications of Artificial Intelligence*, vol. 131, 2024.