

# NORO: Noise-Robust One-shot Voice Conversion with Hidden Speaker Representation Learning

Haorui He<sup>†\*</sup>, Yuchen Song<sup>†\*</sup>, Yuancheng Wang<sup>†\*</sup>, Haoyang Li<sup>‡</sup>, Xueyao Zhang<sup>‡</sup>, Li Wang<sup>‡</sup>,  
Gongping Huang<sup>§</sup>, Eng Siong Chng<sup>‡</sup>, and Zhizheng Wu<sup>†</sup>

<sup>†</sup> School of Data Science, Chinese University of Hong Kong, Shenzhen, 518172, China

<sup>‡</sup> School of Computer Science and Engineering, Nanyang Technological University, Singapore 639798, Singapore

<sup>§</sup> School of Electronic Information, Wuhan University, Wuhan 430072, China

**Abstract**—The effectiveness of one-shot voice conversion (VC) decreases in real-world scenarios where reference speeches, which are often sourced from the internet, contain various disturbances like background noise. To address this issue, we introduce Noro, a noise-robust one-shot VC system. Noro features innovative components tailored for VC using noisy reference speeches, including a dual-branch reference encoding module and a noise-agnostic contrastive speaker loss. Experimental results demonstrate that Noro outperforms our baseline system in both clean and noisy scenarios, highlighting its efficacy for real-world applications. Additionally, we investigate the hidden speaker representation capabilities of our baseline system by repurposing its reference encoder as a speaker encoder. The results show that it is competitive with several advanced self-supervised learning models for speaker representation under the SUPERB settings, highlighting the potential for advancing speaker representation learning through one-shot VC tasks.

## I. INTRODUCTION

One-shot voice conversion (VC) alters the timbre of speech from a source speaker to that of a target speaker using just one reference speech sample from the target speaker, while maintaining the semantic content of the original speech. This task has been extensively studied in the deep learning era, employing various methods and producing promising results [1]–[5]. However, the strong performance of these one-shot VC methods is mainly demonstrated in controlled academic settings, where high-quality data is used for both training and evaluation. As noted in [6]–[9], their effectiveness significantly decreases in more challenging real-world scenarios, where reference speeches often collected from the internet are characterized by various interferences, such as background noise, leading to suboptimal converted speech.

Few works have addressed this challenge to achieve noise-robust VC. Some researchers [10], [11] attempt to integrate a pre-trained or jointly trained speech enhancement module into standard VC systems to improve robustness. However, these methods are not end-to-end and inevitably increase computational costs during inference. Other researchers [7], [9], [12], [13] suggest training strategies, such as data augmentation [7], to mitigate the impact of noise. However, there is still room for

improvement, especially in extremely noisy conditions where the Signal-to-Noise Ratio (SNR) is below 5 dB [12].

To address these challenges, we introduce a novel noise-robust one-shot VC system named Noro. Noro is built on a one-shot VC baseline system based on diffusion. This model is trained to generate speech from pitch and semantic representations obtained through a source encoder and speaker timbre representations via a reference encoder. To achieve noise-robustness with noisy reference speeches, we focus on learning speaker timbre representations that are agnostic to noise. We hypothesize that maintaining noise-agnostic input representations for the acoustic model will prevent its performance from deteriorating due to noise interference.

Therefore, we propose the following designs: Firstly, we introduce a dual-branch reference encoding module. This module consists of two transformer encoders with shared weights: one branch encodes clean reference speech, while the other encodes its noisy counterpart generated through data augmentation. Next, we devise a noise-agnostic contrastive speaker loss to maximize the similarity between samples (whether clean or noisy) from the same speaker while minimizing it for those from different speakers. This loss ensures that the dual-branch reference encoding module learns to represent speaker timbre independent of the acoustic environment of the reference speeches. Experimental results show that Noro significantly enhances the robustness of our baseline system in diverse noisy environments, with only a minor reduction in performance in clean settings, making it well-suited for real-world applications involving noisy reference speeches.

Additionally, inspired by [14], [15], we further investigate the speaker representation capabilities of our baseline system. The motivation stems from the fact that the reference encoder within the VC systems, trained to encode the vocal timbre of various speakers, including unseen ones, for one-shot VC, inherently functions as a speaker encoder. This encoder develops its capability to represent speakers in a self-supervised learning (SSL) manner without needing explicit speaker labels. Thus, we argue that the reference encoder might possess hidden abilities to function effectively as an SSL speaker encoder. To validate this hypothesis, we employ a pre-trained reference encoder from our baseline one-shot VC system as a speaker encoder, denoted as VC-SPK2VEC, and assess its

\*Equal contribution. Listed in alphabetical order.

<sup>†</sup>Corresponding author: wuzhizheng@cuhk.edu.cn

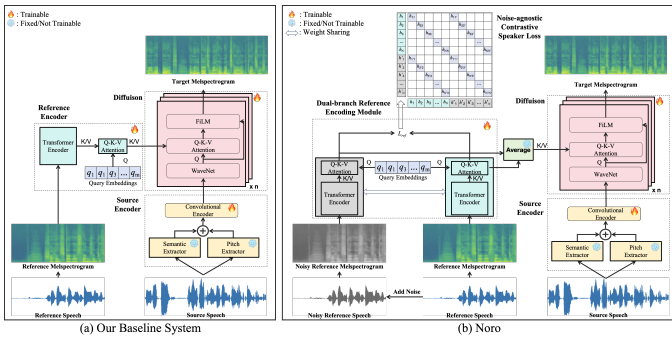


Fig. 1: Model architecture of our baseline system and Noro.

speaker representation ability through a speaker verification task under the SUPERB [16] setting. We compare its performance against several state-of-the-art (SOTA) SSL models, including Wav2vec [17], Wav2vec 2.0 [18], HuBERT [19], and WavLM [20]. Surprisingly, VC-SPK2VEC achieves a competitive Equal Error Rate (EER) of 5.32%, comparable to these SOTA SSL models. This result underscores the potential for leveraging one-shot VC tasks in advancing speaker representation learning.

## II. METHODOLOGY

In this section, we first introduce our baseline one-shot VC system and Noro. As depicted in Fig. 1, Noro builds upon the baseline system, but replaces its original reference encoders with a dual-branch reference encoding module and a noise-agnostic contrastive speaker loss.

### A. The Design of Our Baseline System

1) *Source Encoder*: To derive semantic representations from the source speech, our baseline system employs a pre-trained and frozen HuBERT model [19], [21] as the semantic extractor. This model encodes the source speech into continuous embeddings. We then apply K-Means quantization to these embeddings. Instead of using K-Means cluster IDs (discrete semantic tokens), we use the quantized continuous embeddings, which preserve more information, for semantic representation. The resulting semantic representation for the source speech is denoted as  $s_{\text{src}}$ . For pitch representations, we use the open-source software PyWorld to extract frame-level F0 values from the source speech. We apply min-max normalization to these values and use them for pitch representation, denoted as  $p_{\text{src}}$ . After extracting the semantic and pitch representations, we employ a convolutional neural network (CNN)-based encoder to simultaneously encode  $s_{\text{src}}$  and  $p_{\text{src}}$ , obtaining the final source representation  $h_{\text{src}}$ .

2) *Reference Encoder*: To encode the speaker timbre in the reference speech, we use a cross-attention scheme similar to the speech prompting mechanism in NaturalSpeech 2 [22]. Specifically, we first employ a Transformer-based reference encoder to convert the melspectrogram of the reference speech into a hidden sequence. Instead of allowing the diffusion model to directly attend to this hidden sequence, we employ

two attention blocks. In the first attention block, we use  $m$  randomly initialized query embeddings to attend to the hidden sequence, resulting in  $h_{\text{ref}}$ , a hidden sequence of length  $m$  that represents the reference speaker’s timbre at the utterance level. In the second attention block, the output hidden sequence of the WaveNet Layer serves as the query, while  $h_{\text{ref}}$  functions as both the key and value. The attention results of the second attention block are then used as conditional information for a FiLM layer in the final acoustic modeling process.

3) *Diffusion Model*: To enhance in-context learning for improved one-shot generation, each training sample is randomly segmented during the training phase. A segment comprising 25-45% of the total sample is designated as the reference speech, while the remaining portion functions as both the source and target speech. We utilize a WaveNet-based diffusion model [22], denoted as  $s_{\theta}$ , to predict the melspectrogram  $\mathbf{z}_{\text{tgt}}^t$  of the target speech. We describe both the diffusion (forward) process and the denoising (reverse) process using the stochastic differential equation. In the forward process, Gaussian noise is added to the melspectrogram  $\mathbf{z}_{\text{tgt}}^t$ :  $\mathbf{z}_{\text{tgt}}^t = e^{-\frac{1}{2} \int_0^t \beta_s ds} \mathbf{z}_{\text{tgt}}^t + (\mathbf{I} - e^{-\int_0^t \beta_s ds}) \epsilon$ , where  $\beta_t$  represents the noise schedule function,  $\epsilon$  denotes randomly sampled Gaussian noise, and  $t \in [0, 1]$ . The reverse process can be formulated as:  $d\mathbf{z}_{\text{tgt}}^t = -(\frac{1}{2} \mathbf{z}_{\text{tgt}}^t + \nabla \log p_t(\mathbf{z}_{\text{tgt}}^t)) \beta_t dt + \sqrt{\beta_t} \epsilon$ .

$s_{\theta}$  takes the noised melspectrogram  $\mathbf{z}_{\text{tgt}}^t$ , the time step  $t$ , and the conditions of the source representation  $h_{\text{src}}$ , containing both semantic and pitch information, and the reference representation  $h_{\text{ref}}$ , containing speaker timbre information, to estimate  $\nabla \log p_t(\mathbf{z}_{\text{tgt}}^t)$ . The loss function of the diffusion model can be written as:  $\mathcal{L}_{\text{diff}} = \|s_{\theta}(\mathbf{z}_{\text{tgt}}^t, t, h_{\text{src}}, h_{\text{ref}}) - \nabla \log p_t(\mathbf{z}_{\text{tgt}}^t)\|_1$ . During the inference stage, we gradually denoise  $\mathbf{z}_{\text{tgt}}^t$  utilizing the estimated score  $s_{\theta}(\mathbf{z}_{\text{tgt}}^t, t, s_{\text{src}}, p_{\text{src}}, h_{\text{ref}})$ . This process initiates with Gaussian noise, initially sampled as  $\mathbf{z}_{\text{tgt}}^1$ .

### B. The Design of Noro

1) *Dual-branch Reference Encoding Module*: To improve the noise robustness of the baseline system, we replaced its original reference encoder with a dual-branch reference encoding module. This module consists of two reference encoders sharing identical model weights as a dual-branch structure. For each clean training sample, we randomly mix clean reference speech with eight types of noise from the DEMAND database [23] at SNRs following a normal distribution (0,20) dB to create noisy reference speeches. Subsequently, the clean and noisy reference speeches are processed by separate reference encoders using the same query embeddings. This process generates clean and noisy reference representations  $h_{\text{ref}}$  and  $h'_{\text{ref}}$ , respectively. Considering the weight-sharing scheme used in this dual-branch reference encoding module, the method described here qualifies as a “training strategy” for enhancing our baseline system. Initially, we load the Noro system with weights from a pre-trained baseline system. During the training phase, the average of  $h_{\text{ref}}$  and  $h'_{\text{ref}}$  is fed into the diffusion model for acoustic modeling and acts as the keys and values in the second attention block. During the inference phase,

only one reference encoder is utilized, and the Noro system’s structure mirrors that of the baseline system.

2) *Noise-agnostic Contrastive Speaker Loss*: After obtaining the representations of the clean speeches  $h_{\text{ref}}$  and their noisy counterparts  $h'_{\text{ref}}$ , we introduce a noise-agnostic contrastive speaker loss. This loss function aims to maximize the similarity between samples (clean or noisy) from the same speaker while minimizing it for samples from different speakers. First, as  $h_{\text{ref}}$  and  $h'_{\text{ref}}$  are hidden sequences of length  $m$ , we perform average pooling over the length dimension to form a comprehensive reference representation that captures the reference speaker’s timbre at the utterance level. We then concatenate  $h_{\text{ref}}$  and  $h'_{\text{ref}}$  along the batch size dimension:  $h_{\text{all}} = [h_{\text{ref}}; h'_{\text{ref}}]$ . Given that the noisy speeches share the same speakers as the clean ones,  $y_{\text{all}} = [y_{\text{spk}}; y_{\text{spk}}]$ . Finally, our proposed noise-agnostic contrastive speaker loss is formulated as:  $\mathcal{L}_{\text{ref}} = \frac{1}{2N} \sum_{i=1}^{2N} \text{CrossEntropy} \left( \frac{h_i \cdot h_j^T}{\tau}, \mathbf{M}_{i,j} \right)$ , where  $h_i$  and  $h_j$  are the  $i$ -th and  $j$ -th reference representations in  $h_{\text{all}}$ ,  $\tau$  is the temperature parameter adjusting the scaling of the logits,  $N$  is the batch size, and the mask matrix  $\mathbf{M}_{i,j}$  is defined as follows:  $\mathbf{M}_{i,j} = \begin{cases} 1 & \text{if } y_i = y_j, \\ 0 & \text{otherwise,} \end{cases}$  where  $y_i$  and  $y_j$  are the speaker labels for  $h_i$  and  $h_j$  respectively. This loss ensures that the reference encoders represent the vocal timbre of different speakers regardless of noise interference, thus enhancing the robustness of the system. The total loss function for the Noro system is defined as follows:  $\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{diff}} + \beta \mathcal{L}_{\text{ref}}$ , where  $\alpha$  and  $\beta$  are the weights assigned to each loss component.

### III. EXPERIMENTS

#### A. One-shot VC Under Different Acoustic Environments

1) *Experimental Setups*: To evaluate Noro, we use two acoustic settings: clean and noisy. In the clean setting, the high-quality VCTK Corpus, recorded in a studio, serves as the evaluation dataset. VCTK contains recordings from 110 English speakers with diverse accents, characterized by minimal noise interference. The test set comprises 150 randomly selected pairs of source and reference speech from VCTK. For the noisy setting, we introduce noise of unseen types from the training stage to the reference speeches in the test set at different SNR conditions to create challenging conditions to evaluate the robustness of VC systems. Under these two settings, we compare Noro with our baseline system and three state-of-the-art one-shot VC models: FaCodec-VC [24], FreeVC [25], and DiffVC [1]. For both our systems, we employ the LibriLight [26] dataset for training, which consists of 60k hours of 16 kHz speech data from over 8k distinct speakers. We leverage the official segmentation scripts to segment the speech data into approximately 15-second sequences by concatenating consecutive chunks with voice activity. Utterances longer than 30 seconds or shorter than 3 seconds are discarded. Detailed hyper-parameters is publicly available at Amphion [3].

We use both objective and subjective metrics to evaluate the systems. For objective evaluation, we follow [2] to employ

speaker embedding cosine similarity (SECS) and character error rate (CER). SECS is determined using a cutting-edge speaker representation model,<sup>1</sup> with the results reflecting similarity to the original voice prompt. CER is calculated using an ASR model<sup>2</sup> to assess the intelligibility of the generated speech. For subjective evaluation, we follow [22], [24] to use the Comparative Mean Opinion Score (CMOS) and the Similarity Mean Opinion Score (SMOS) to evaluate naturalness and similarity, respectively. We use ten randomly selected pairs from the clean test set and another ten from the noisy test set. Twelve proficient English speakers conducted the assessments. The naturalness scores range from one (“Bad”) to five (“Excellent”), and the similarity scores range from one (“Different speaker, sure”) to four (“Same speaker, sure”).

TABLE I: The experimental results for objective evaluation.

Model	Clean		Noisy (0–5 dB)		Noisy (5–10 dB)	
	CER ↓	SECS ↑	CER ↓	SECS ↑	CER ↓	SECS ↑
Ground Truth	1.37	90.4	1.37	90.4	1.37	90.4
FaCodec-VC	3.66	82.33	3.56	74.78	3.26	77.88
FreeVC	3.31	80.07	3.49	73.83	3.13	76.14
DiffVC	16.91	85.32	15.72	82.00	13.91	82.33
Our Baseline	4.71	82.35	7.26	77.28	6.43	78.89
Noro	4.74	82.38	4.66	80.09	5.06	80.71

TABLE II: The experimental results for subjective evaluation.

Model	Clean		Noisy (0–5 dB)	
	CMOS ↑	SMOS ↑	CMOS ↑	SMOS ↑
Ground Truth	4.04	3.90	4.04	3.90
Our Baseline	<b>3.29</b>	<b>3.02</b>	2.09	2.75
Noro	3.16	2.90	<b>2.95</b>	<b>2.97</b>

2) *Results and Analysis*: Table I presents the objective evaluation results for the Noro and baseline systems in both clean and noisy environments. In clean conditions, the baseline system achieved a CER of 4.71 and a SECS of 82.35, while Noro had a comparable CER of 4.74 and a SECS of 82.38. Both systems demonstrated effective one-shot voice conversion, accurately preserving semantic content and adapting to speaker timbre. Under noisy conditions, the baseline system’s performance deteriorated significantly, with the CER increasing to 7.26 and the SECS dropping to 77.28 at 0–5 dB, reflecting its poor noise-handling capabilities. In contrast, Noro exhibited significantly better noise robustness, with a CER of 4.66 and a SECS of 80.09, indicating its reliability in real-world, noisy environments. The subjective evaluation results in Table II are consistent with the objective findings. Both systems performed well in clean conditions, maintaining speech quality and speaker identity. However, under noisy conditions, the baseline system’s performance declined sharply, with a CMOS of 2.09 and an SMOS of

<sup>1</sup><https://huggingface.co/microsoft/wavlm-base-plus-sv>

<sup>2</sup><https://huggingface.co/facebook/hubert-large-ls960-ft>




Fig. 2: t-SNE visualization of reference representations.

2.75, while Noro maintained significantly higher scores, with a CMOS of 2.95 and an SMOS of 2.97. These results underscore Noro’s superior noise robustness and its practicality for use in diverse noisy environments.

We further visualize the reference representation  $h_{\text{ref}}$  produced by the baseline system and Noro using t-SNE, as shown in Fig. 2. In the baseline system, the reference representations for clean and noisy speech show a clear separation, which contributes to the significant drop in performance when the reference speech is noisy. In contrast, Noro’s reference representations for both clean and noisy speech are well-mixed. This is due to Noro’s use of the proposed dual-branch reference encoding module and noise-agnostic contrastive speaker loss during training, which enables it to learn a noise-agnostic reference representation. This capability accounts for Noro’s superior performance in diverse noisy environments.

#### B. One-shot VC as SSL for Speaker Representation

1) *Experimental Setups*: The reference encoder effectively captures speaker timbre characteristics. Here, we adapt the reference encoder—originally designed for one-shot VC—into a self-supervised learning (SSL) speaker encoder, named VC-SPK2VEC. We evaluate its performance in speaker representation by comparing it with several widely used pre-trained SSL models [17]–[20], [27], [28] under the SUPERB framework. Following the SUPERB [16] setup, we employ SSL features for speaker verification. Input waveforms are processed through a frozen pre-trained SSL model to extract frame-level features, and a simple x-vector model serves as the downstream model, designed to classify whether two utterances belong to the same speaker, framing this as a binary classification task. All models are fine-tuned using the VoxCeleb2 dataset, augmented with noise, and tested on VoxCeleb1 [29]. For traditional SSL models, the final speaker representation is produced by computing the weighted sum of hidden states from each layer:  $\hat{o}_t = \sum_{l=1}^L w_l \cdot h_l^t$ , where  $h_l^t$  represents the hidden state from layer  $l$  at time  $t$ , and  $w_l$  denotes the normalized weight for each layer. In contrast, for VC-SPK2VEC, the final representation is taken from the hidden states of the reference encoder’s transformer encoder’s last layer. The x-vector downstream model utilizes an additive-margin softmax loss, maintaining the same hyperparameters as outlined in [29]. Cosine similarity is then used to compute the matching scores between speaker embeddings for each enrollment-test utterance pair.

TABLE III: The SV performance comparison. The results for these baseline models are sourced from [16] and [20].

Model	Params.	Data	EER (%)
FBANK	-	-	9.56
Modified CPC [27]	1.84M	60k hr	12.86
wav2vec [17]	32.54M	960 hr	7.99
vq-wav2vec [28]	34.15M	960 hr	10.38
wav2vec 2.0 Base [18]	95.04M	960 hr	6.02
HuBERT Base [19]	94.68M	960 hr	5.11
WavLM Base [20]	94.70M	960 hr	4.69
WavLM Base+ [20]	94.70M	94k hr	4.07
wav2vec 2.0 Large [18]	317.38M	60k hr	5.65
HuBERT Large [19]	316.61M	60k hr	5.98
WavLM Large [20]	316.62M	94k hr	<b>3.77</b>
VC-SPK2VEC	72.4M	60k hr	5.32

2) *Results and Analysis*: Table III compares VC-SPK2VEC with various baseline models on the speaker verification (SV) task. VC-SPK2VEC achieves a competitive equal error rate (EER) of 5.32%. While it trails slightly behind HuBERT Base (5.11%) and the three WavLM model variants (4.69%, 4.07%, and 3.77% for Base, Base+, and Large, respectively), it outperforms other baseline models such as wav2vec 2.0 Base, wav2vec 2.0 Large, and HuBERT Large. Notably, VC-SPK2VEC uses a comparatively smaller model size of 72.4M parameters. These results underscore VC-SPK2VEC’s effectiveness as a self-supervised speaker encoder and highlight its comparable performance to SOTA SSL models. These findings highlight the great potential of leveraging speaker representation learning in one-shot speech generation tasks, as initially explored in [14], [15]. With the advent of speech generation models trained on extensive datasets exceeding 100k hours of speech data, unifying speaker representation learning with generative speech tasks like voice conversion (VC) may lead to further advancements in both fields. Future research could investigate the relationship between the capability of speaker representation and factors such as the number of speakers included in the training set of one-shot VC systems or the model size of the reference encoder in these systems.

#### IV. CONCLUSIONS

In this paper, we introduced Noro, a noise-robust one-shot voice conversion system with a dual-branch reference encoding module and noise-agnostic contrastive speaker loss to improve performance in diverse noisy environments. Our experiments show that Noro significantly outperforms the baseline in noisy settings with minimal performance loss in clean conditions, making it a practical solution for real-world use. We also repurposed the reference encoder as an SSL speaker encoder, VC-SPK2VEC, and demonstrated its competitive performance on the SUPERB benchmark compared to advanced SSL models. This suggests a promising direction for unifying speaker representation learning and voice conversion, with potential for leveraging large-scale voice conversion systems and unlabeled data for further advancements in both fields.

#### ACKNOWLEDGMENT

The work is supported by the National Natural Science Foundation of China (No. 62376237), the 2023 Shenzhen Stability Science Program, the Shenzhen Science and Technology Program (No. ZDSYS20230626091302006), and the Program for Guangdong Introducing Innovative and Entrepreneurial Teams (No. 2023ZT10X044).

#### REFERENCES

- [1] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, M. S. Kudinov, and J. Wei, "Diffusion-Based Voice Conversion with Fast Maximum Likelihood Sampling Scheme," in *ICLR*, 2022.
- [2] J. Li, Y. Guo, X. Chen, and K. Yu, "SEF-VC: Speaker Embedding Free Zero-Shot Voice Conversion with Cross Attention," in *ICASSP*, 2024.
- [3] X. Zhang, L. Xue, Y. Gu, *et al.*, "Amphion: An open-source audio, music and speech generation toolkit," in *SLT*, 2024.
- [4] H. He, Z. Shang, C. Wang, *et al.*, "Emilia: An Extensive, Multilingual, and Diverse Speech Dataset for Large-Scale Speech Generation," in *SLT*, 2024.
- [5] H. He, Z. Shang, C. Wang, *et al.*, "Emilia: A Large-Scale, Extensive, Multilingual, and Diverse Dataset for Speech Generation," in *arXiv:2501.15907*, 2025.
- [6] T.-H. Huang, J.-H. Lin, and H.-Y. Lee, "How Far are We from Robust Voice Conversion: A Survey," in *SLT*, 2021.
- [7] C.-Y. Huang, K.-W. Chang, and H.-Y. Lee, "Toward Degradation-Robust Voice Conversion," in *ICASSP*, 2022.
- [8] C. Xie and T. Toda, "Noisy-to-Noisy Voice Conversion Under Variations of Noisy Condition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 3871–3882, 2023.
- [9] L. Xue, S. Yang, N. Hu, D. Su, and L. Xie, "Learning Noise-Independent Speech Representation for High-Quality Voice Conversion for Noisy Target Speakers," in *INTERSPEECH*, 2021.
- [10] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Investigating RNN-Based Speech Enhancement Methods for Noise-Robust Text-to-Speech," in *SSW*, 2016.
- [11] Y.-J. Chan, C.-J. Peng, S.-S. Wang, H.-M. Wang, Y. Tsao, and T.-S. Chi, "Speech Enhancement-Assisted Stargan Voice Conversion in Noisy Environments," *arXiv preprint arXiv:2110.09923*, 2021.
- [12] H. Du, L. Xie, and H. Li, "Noise-Robust Voice Conversion with Domain Adversarial Training," *Neural Networks*, vol. 148, pp. 74–84, 2022.
- [13] T. Igarashi, Y. Saito, K. Seki, *et al.*, "Noise-Robust Voice Conversion by Conditional Denoising Training Using Latent Variables of Recording Quality and Environment," in *INTERSPEECH*, 2024.
- [14] J. Cho, P. Zelasko, J. Villalba, S. Watanabe, and N. Dehak, "Learning Speaker Embedding from Text-to-Speech," in *INTERSPEECH*, 2020.
- [15] J. Cho, P. Żelasko, J. Villalba, and N. Dehak, "Improving Reconstruction Loss Based Speaker Embedding in Unsupervised and Semi-Supervised Scenarios," in *ICASSP*, 2021.
- [16] S.-w. Yang, P.-H. Chi, Y.-S. Chuang, *et al.*, "SUPERB: Speech Processing Universal Performance Benchmark," in *INTERSPEECH*, 2021.
- [17] S. Schneider, A. Baeovski, R. Collobert, and M. Auli, "Wav2vec: Unsupervised Pre-Training for Speech Recognition," in *INTERSPEECH*, 2019.
- [18] A. Baeovski, Y. Zhou, A. Mohamed, and M. Auli, "Wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations," in *NeurIPS*, 2020.
- [19] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [20] S. Chen, C. Wang, Z. Chen, *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.
- [21] A. Lee, H. Gong, P.-A. Duquenne, *et al.*, "Textless Speech-to-Speech Translation on Real Data," in *NACCL*, 2022.
- [22] K. Shen, Z. Ju, X. Tan, *et al.*, "Naturalspeech 2: Latent Diffusion Models are Natural and Zero-Shot Speech and Singing Synthesizers," in *ICLR*, 2023.
- [23] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Investigating RNN-Based Speech Enhancement Methods for Noise-Robust Text-to-Speech," in *SSW*, 2016.
- [24] Z. Ju, Y. Wang, K. Shen, *et al.*, "Naturalspeech 3: Zero-Shot Speech Synthesis with Factorized Codec and Diffusion Models," *ICML*, 2024.
- [25] J. Li, W. Tu, and L. Xiao, "FreeVC: Towards High-Quality Text-Free One-Shot Voice Conversion," in *ICASSP*, 2023.
- [26] J. Kahn, M. Riviere, W. Zheng, *et al.*, "Libri-Light: A Benchmark for ASR with Limited or No Supervision," in *ICASSP*, 2020.
- [27] M. Rivière, A. Joulin, P.-E. Mazaré, and E. Dupoux, "Unsupervised Pretraining Transfers Well Across Languages," in *ICASSP*, 2020.
- [28] A. Baeovski, S. Schneider, and M. Auli, "Vq-Wav2vec: Self-Supervised Learning of Discrete Speech Representations," in *ICLR*, 2019.
- [29] A. Nagrani, J. S. Chung, W. Xie, and A. Zisserman, "VoxCeleb: Large-Scale Speaker Verification in the Wild," *Computer Speech & Language*, vol. 60, p. 101027, 2020.