



TOWARDS HIGH-FIDELITY SINGING VOICE CONVERSION WITH ACOUSTIC REFERENCE AND CONTRASTIVE PREDICTIVE CODING

Chao Wang^{1,2}, Zhonghao Li², Benlai Tang², Xiang Yin², Yuan Wan², Yibiao Yu¹, Zejun Ma²,

¹Soochow University

²ByteDance AI Lab

georgefly1@163.com, {lizhonghao.01,tangbenlai}@bytedance.com

ABSTRACT

Recently, phonetic posteriorgrams (PPGs) based methods have been quite popular in non-parallel singing voice conversion systems. However, due to the lack of acoustic information in PPGs, style and naturalness of the converted singing voices are still limited. To solve these problems, in this paper, we utilize an acoustic reference encoder to implicitly model singing characteristics. We experiment with different auxiliary features, including mel spectrograms, HuBERT, and the middle hidden feature (PPG-Mid) of pretrained automatic speech recognition (ASR) model, as the input of the reference encoder, and finally find the HuBERT feature is the best choice. In addition, we use contrastive predictive coding (CPC) module to further smooth the voices by predicting future observations in latent space. Experiments show that, compared with the baseline models, our proposed model can significantly improve the naturalness of converted singing voices and the similarity with the target singer. Moreover, our proposed model can also make the speakers with just speech data sing.

Index Terms— Singing voice conversion, phonetic posteriorgrams, acoustic reference, contrastive predictive coding

1. INTRODUCTION

Singing voice conversion is intended to change the timbre of singing voices while keeping the musical contents like melody and lyrics maintained. It has achieved more and more attention in recent years, since it can be used in many aspects, like improving the singer's vocal quality and creating a virtual singer.

The early studies are mainly focused on parallel singing voice conversion, which means there are parallel training data between the source singer and the target singer. These traditional methods[1][2][3] use statistical models to build the mapping between the parallel samples. However, it's quite expensive to collect parallel samples of different singers in real application scenarios. So recently, the studies of non-parallel voice conversion have been the mainstream. The key idea of non-parallel singing voice conversion is to extract the singer-independent content representation first, and then generate the converted singing voices using the target singer embedding. According to whether the singer-independent content representation is extracted in an unsupervised manner, non-parallel

singing voice conversion can be categorized into auto-encoder based and PPGs based methods.

Auto-encoder based methods disentangle the content feature from speech in an unsupervised manner by imposing constraints on encoder output. Domain confusion [4] module is used in [5][6] to remove speaker-dependent information like timbre from speech. In [7][8], feature disentanglement is realized using variational auto-encoder (VAE). [9] utilizes vector quantized variational auto-encoder (VQ-VAE) as an information bottleneck to get the content information. Although auto-encoder based methods can achieve acceptable results in non-parallel singing voice conversion, its ability of feature disentanglement is not robust enough for various singing input.

In PPGs based methods[10][11] [12][13][14], they use the last layer output of the pretrained ASR model, namely PPGs, to represent content information. Since the ASR model is pretrained with huge amount of training data and the training targets are phoneme labels, its ability of eliminating acoustic information is quite robust. However, compared with speech, singing voice contains much more characteristics like prosody and rhythm, which are difficult to model explicitly only using PPGs. Several works have tried to supplement acoustic features in singing voice conversion. In [13], they use pitch to represent melody. However, pitch itself is still not enough to model singing characteristics. In [14], apart from pitch, they feed mel spectrograms into reference encoder to model singing characteristics implicitly. However, as mentioned in auto-encoder based methods, there exists too much singer-dependent information in mel spectrograms, which is difficult to remove and will finally degrade the performance.

In this work, our contribution is two-fold: (1) We try to find an appropriate auxiliary feature that contains enough singer-independent acoustic and musical information to serve as the reference encoder input. Apart from mel spectrograms, We select HuBERT[15], which is a kind of feature originally proposed in self-supervised speech representation learning, and middle hidden feature of ASR model, denoted as PPG-Mid. Both HuBERT feature and PPG-Mid feature contain linguistic and acoustic information. Meanwhile, they contain less singer information than mel spectrograms. In the experiments, we compare these three features when they serve as the reference encoder inputs. To further eliminate the remaining singer-dependent information like timbre in the ref-

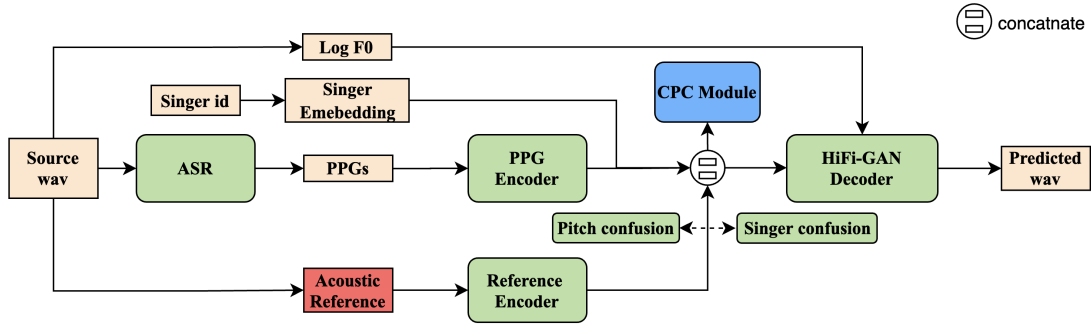


Fig. 1. Overall architecture of our proposed method. Singer confusion, pitch confusion, and CPC module are not used in the inference phase.

reference encoder output, both singer confusion and pitch confusion module are also used. (2) To alleviate local errors and enhance the smoothness of voices, inspired by recent studies proposed in speech self-supervised representation learning[16][17][18], we utilize contrastive predictive coding (CPC) module which predicts future representations in the encoder output feature space. In [19], they use the CPC module to eliminate the local errors in the vector quantization output. In our paper, we are the first one to utilize the CPC module in the PPGs based singing voice conversion system.

The rest of the paper is organized as follows: Section 2 will present the proposed method. In Section 3, Experiments results are presented. Section 4 concludes this paper.

2. METHODS

The overall architecture of our proposed method is shown in Fig.1. It is mainly composed of pretrained ASR model, PPG encoder, acoustic reference encoder, CPC module, and HiFi-GAN[20] decoder.

During training, the PPG encoder is to acquire singing content from PPGs feature. Meanwhile, the auxiliary acoustic feature of source voices is fed into the reference encoder. The outputs of PPG encoder and reference encoder are concatenated to form the final encoder output. Then encoder output, singer embedding, and pitch are fed into the HiFi-GAN decoder to reconstruct the speech waveform, where pitch is fused into the decoder in the same way as [12]. Meanwhile, the singer and pitch confusion module are used to remove the singer information in the reference encoder output. During conversion, the singer and pitch confusion module are removed. The ASR model is based on Deep-FSMN[21], which is a hierarchical network with 38 layers. The model architectures of the PPG encoder and acoustic reference encoder are CBHG network proposed in [22]. We use the HiFi-GAN vocoder proposed in [20] as the decoder to reconstruct waveform directly. The reconstruction loss used for training the network follows [20].

In the following part, we will introduce the acoustic reference features, and the CPC module in detail.

2.1. Acoustic Reference

Singing voice contains rich acoustic and musical information. However, PPGs do not contain enough acoustic and musical information, which makes it difficult to generate high-fidelity singing voices. To model the acoustic and musical information for singing voice conversion, we incorporate a reference encoder. As for the reference encoder input, we select mel spectrograms, PPG-Mid, and HuBERT feature as the candidates:

Mel spectrograms: Although mel spectrograms contains rich acoustic information, too much singer-dependent information like timbre also exists in mel spectrograms, which will finally degrade the quality of the converted singing voices and reduce the similarity with the target speaker.

PPG-Mid: PPG-Mid is the output feature of 30th hidden layer in DFSMN[21] based ASR model. On one hand, Compared with PPGs, the middle hidden layer output PPG-Mid contains more acoustic information because it is the shallower layer. On the other hand, The singer information in PPG-Mid is much less than that in mel spectrograms, because ASR model has already eliminated most of the singer information by linguistic label training.

HuBERT: HuBERT feature[15] is a kind of self-supervised speech representation, which achieves the state-of-the-art results in the ASR task. As shown in Fig.2, it utilizes an offline clustering step to provide aligned pseudo target labels for prediction. Because audio data of different speakers are trained towards the same pseudo target labels, most of the speaker information in the HuBERT feature are removed. The prediction loss is over the masked regions only, forcing the model to extract high-level representations from the unmasked frames to infer the masked targets correctly. Therefore, the HuBERT feature contains a combined acoustic and linguistic information.

However, the above three acoustic reference features still keep some singer information like pitch and timbre, so we still need the singer confusion and pitch confusion module like in [6] to eliminate these information. The reference encoder output is fed into the singer-identity classifier and pitch predictor. Both the confusion modules consist of four 1-D convolution layers activated by ReLU layer followed by one final linear projection layer. The singer clas-

sifier is trained using cross entropy loss L_s to correctly classify the speaker label, and pitch predictor is optimized using MSE loss L_f to predict the pitch. The confusion loss $L_{confusion}$ to update the reference encoder is defined as below:

$$L_{confusion} = -\lambda * L_s - \omega * L_f \quad (1)$$

where scalar λ and ω are the weight of the singer confusion loss and pitch confusion loss, respectively.

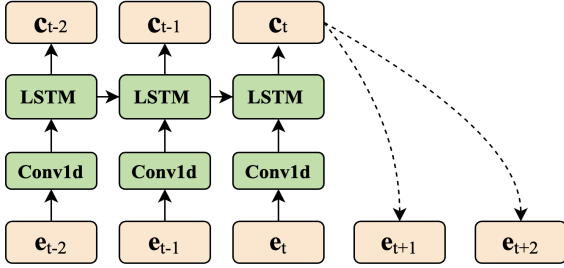


Fig. 2. The architecture of the CPC module

2.2. CPC module

The quality of the converted singing voices depends on the accuracy of the pretrained ASR model. The local recognition errors in the PPGs feature will bring noise in the converted singing voices and degrade the quality of singing voice conversion. In addition, although the singer and pitch confusion module can get rid of the remaining singer information like timbre in the reference encoder output, it inevitably harms the phonetic and vocal information at the same time, which will cause local instability in the converted singing voices.

To overcome these problems, we employ contrastive predictive coding (CPC) module as the regularization loss to predict future representations in the encoder output feature space. The CPC module is originally proposed in the self-supervised learning[16][17][18] to learn meaningful speech representation, which can be used in various downstream tasks. In this paper, we utilize CPC module to discard local noise and errors in the encoder output and enhance the voice smoothness.

As shown in Fig.2, the encoder output $E = \{e_1, \dots, e_T\}$ of length T is fed to the context network, which in our paper is composed of long short-term memory network (LSTM) and 1-D convolution layer, to get the context output $C = \{c_1, \dots, c_T\}$. Given current timestep context feature c_t , the CPC module needs to distinguish the positive future k -step feature e_{t+k} from the remaining negative feature samples. If e_t contains vocal errors, then c_t also contains errors. Given e_{t+k} contains precise vocal information, the network learns to discard the local errors and noise in e_t and c_t so that e_{t+k} can be predicted correctly. The CPC loss term is as follow:

$$L_{CPC} = -\beta * \sum_{t=0}^{T-t-1} \sum_{k=1}^K \log \left[\frac{\exp(e_{t+k}^T \mathbf{W}_k c_t)}{\sum_{\tilde{e} \in \Omega} \exp(\tilde{e}^T \mathbf{W}_k c_t)} \right] \quad (2)$$

where T is the length of the encoder output, K is the largest timestep we predict, and β is the CPC loss weight. $\mathbf{W}_k (k = 1, 2, \dots, K)$ is a trainable projection matrix and Ω is the set containing the negative samples. In our paper, the negative samples set is formed by randomly selected from the same utterance as the positive sample. The K is set to 12.

Let the HiFi-GAN[20] reconstruction loss be the L_{hifi} , which consists of Mel spectrograms loss, deep feature matching loss, and adversarial loss. The overall generation loss of our proposed model is shown as below:

$$L_G = L_{hifi} + L_{confusion} + L_{CPC} \quad (3)$$

3. EXPERIMENTS

3.1. Experiment setup

We evaluate our proposed model on a mixed training dataset which contains VCTK speech corpus[23] and NUS-48E singing corpus[24]. 90% of the audios in datasets are selected for training, 5% for validation. The test set consists of 20 segments from 10 singers. For evaluation, we select a female singer and a male singer as the target timbres. Except for section 3.2.4, the test source audios and target singers for evaluation are selected from NUS-48E corpus, namely in-domain singing voice conversion. All songs are sampled at 24kHz. 80-dimensional mel spectrograms are extracted with a Hanning window of 40 ms and 10 ms frame shift.

We use 20k singing data to pretrain the ASR model. The HuBERT features are extracted using the pretrained model provided in [15]. We train our proposed model using Adam [25] optimizer for 400K steps. The learning rate is 0.0004 and the batch size is 16. Weighting factors λ , ω , and β are all set to 0.1. We use pyreaper¹ to extract the pitch from the audios. Samples can be found online².

3.2. Evaluation

We invite 15 music professionals to evaluate the results. For subjective evaluation, we select 1-5 mean opinion score (MOS) (1-bad, 2-poor, 3-fair, 4-good, 5-excellent) for both singing naturalness and timbre-style similarity with the target singer. For objective evaluation, we select normalized cross-correlation (NCC) to measure the pitch matching degree between the source and converted singing audios. The higher the NCC score is, the more accurate the pitch is.

Table 1. The evaluation results of different acoustic reference.

Method	Naturalness	Simiarity	NCC
Ground_truth	4.76 ± 0.07	–	–
Prop_mel	3.56 ± 0.16	3.48 ± 0.16	0.976
Prop_ppg-mid	3.49 ± 0.18	3.58 ± 0.18	0.982
Prop_hub	3.62 ± 0.17	3.63 ± 0.16	0.985

¹<https://github.com/tr9y9/pyreaper>

²<https://georgehappy1.github.io/svcdemo/>

Table 2. The evaluation results of the ablation study for the CPC module

Method	Naturalness	Simiarity	NCC
Prop_hub	3.62 ± 0.17	3.63 ± 0.16	0.985
Prop_hub_cpc	3.74 ± 0.19	3.87 ± 0.17	0.987

Table 3. The evaluation results of baseline models and our proposed model

Method	Naturalness	Simiarity	NCC
Baseline1	2.49 ± 0.19	2.66 ± 0.19	0.969
Baseline2	1.85 ± 0.13	1.94 ± 0.13	0.949
Prop_hub_cpc	3.74 ± 0.19	3.87 ± 0.17	0.987

3.2.1. Comparison among different acoustic reference

In this part, we compare three acoustic reference features. Systems using HuBERT feature, middle hidden feature PPG-Mid, and mel spectrograms are called Prop_hub, Prop_ppg-mid, Prop_mel, respectively.

The naturalness, similarity MOS results, and NCC are shown in Table.1. The Prop_hub model all achieves the best results in terms of naturalness, similarity, and NCC scores. Although Prop_mel model achieves higher naturalness MOS scores than Prop_ppg-mid, its similarity and NCC score is lower. This is because the mel spectrograms contains too much singer-dependent information to remove. Now we can conclude that the HuBERT feature is the best choice as the acoustic reference. So in the following experiment, we all use the HuBERT feature as the reference encoder input.

3.2.2. Ablation study for the CPC module

To show the effectiveness of the CPC module, we compare the Prop_hub model with the Prop_hub_cpc model, which incorporates the CPC module. As shown by the results in Table.2, incorporating the CPC module in Prop_hub improves both singing naturalness and similarity. It reveals that CPC module, which serves as the regularization loss applied to the encoder output, can discard local noise and errors and enhance the voice smoothness.

3.2.3. Comparison with the baseline model

In this part, we compare our best configuration Prop_hub_cpc with the following two state-of-the-art baseline models. Baseline1 is proposed in [13]. It consists of PPG encoder, pitch encoder, speaker encoder, and decoder. Similar to our proposed model, baseline1 is also an end-to-end model, which maps PPGs to waveform directly. Baseline2 is proposed in [14]. Apart from modules in Baseline1, Baseline2 incorporates a mel reference encoder. Different from our proposed model, Baseline2 is not an end-to-end model. It maps PPGs to mel spectrograms and then we use a pretrained HiFi-GAN vocoder to generate the waveform.

Table 4. The evaluation results of cross-domain singing voice conversion

Method	NCC	COS-SIM
Baseline1 (cross-domain)	0.926	0.64
Baseline2 (cross-domain)	0.902	0.57
Prop_hub_cpc (cross-domain)	0.983	0.67
Prop_hub_cpc (in-domain)	0.987	0.77

From the evaluation results shown in Table.3, we can see our proposed model significantly improves the naturalness and similarity compared with the two baselines, which reveals the the superiority of our proposed method. It is noted that the MOS scores of two baseline systems in our paper are lower than those in the original paper[13][14]. This is because the training singing data used in our paper is only a 3-hour open-source singing corpus, which is much less than the training data used in [13][14].

3.2.4. Cross-domain singing voice conversion

Our work focuses on in-domain singing voice conversion. However, with the VCTK speech dataset involved in training, our model can also make the speaker in the VCTK dataset sing, which is called cross-domain singing voice conversion. We compare our proposed model with the Baseline1 and Baseline2 model. We just evaluate the performance via objective metrics normalized cross-correlation (NCC) and Cosine similarity (COS-SIM) between the d-vectors of the converted and the real target audio. The d-vectors are extracted from the pretrained speaker verification model³. Higher the COS-SIM score is, more similar the timbres are. The evaluation results are presented in Tab.4. Our proposed method still achieves significantly better results than the two baseline models, which is consistent with in-domain singing voice conversion. Meanwhile, the results of cross domain singing voice conversion are inferior to these of in domain singing voice conversion since in domain task is easier.

4. CONCLUSION AND FUTURE WORK

In this paper, we proposed an end-to-end high-quality singing voice conversion system. It employs a reference encoder to implicitly model the prosody and rhythm of the singing voices. We compare three different features: mel spectrograms, PPG-Mid, and HuBERT feature as the acoustic reference. Experiments show that the HuBERT feature is the best choice. Meanwhile, we incorporate a CPC module to further enhance the smoothness of singing voices. Comprehensive experiments have been performed to verify the superiority of our proposed model. As for future work, we will further improve the performance of cross-domain singing voice conversion.

³<https://github.com/resemble-ai/Resemblyzer>

5. REFERENCES

- [1] Kazuhiro Kobayashi, Tomoki Toda, Graham Neubig, Sakriani Sakti, and Satoshi Nakamura, "Statistical singing voice conversion based on direct waveform modification with global variance," *INTERSPEECH*, pp. 2754–2758, 2015.
- [2] Kazuhiro Kobayashi, Tomoki Toda, Graham Neubig, Sakriani Sakti, and Satoshi Nakamura, "Statistical singing voice conversion with direct waveform modification based on the spectrum differential," *INTERSPEECH*, pp. 2514–2518, 2014.
- [3] T. Toda, Y. Ohtani, and K. Shikano, "One-to-many and many-to-one voice conversion based on eigenvoices," *ICASSP*, pp. 1249–1252, 2007.
- [4] Ganin, Yaroslav, et al., "Domain-adversarial training of neural networks," *The journal of machine learning research*, vol. 17, pp. 2096–2030, 2016.
- [5] E. Nachmani and L. Wolf, "Unsupervised singing voice conversion," *INTERSPEECH*, pp. 2583–2587, 2019.
- [6] C. Deng, C. Yu, et al., "Pitchnet: Unsupervised singing voice conversion with pitch adversarial network," *ICASSP*, pp. 7749–7753, 2020.
- [7] Y. Luo, C. Hsu, et al., "Singing voice conversion with disentangled representations of singer and vocal technique using variational autoencoders," *ICASSP*, pp. 3277–3281, 2020.
- [8] J. Lu, K. Zhou, et al., "Vaw-gan for singing voice conversion with non-parallel training data," *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, pp. 514–519, 2020.
- [9] Naoya Takahashi, Mayank Kumar Singh, et al., "Hierarchical disentangled representation learning for singing voice conversion," arXiv preprint, 2021.
- [10] X. Chen, W. Chu, J. Guo, and N. Xu, "Singing voice conversion with non-parallel data," *MIPR*, p. 292–296, 2019.
- [11] Adam Polyak, Lior Wolf, Yossi Adi, and Yaniv Taigman, "Unsupervised cross-domain singing voice conversion," *INTERSPEECH*, pp. 801–805, 2020.
- [12] Songxiang Liu, Yüewen Cao, Na Hu, Dan Su, and Helen Meng, "Fastsvc: Fast cross-domain singing voice conversion with feature-wise linear modulation," *ICME*, pp. 1–6, 2021.
- [13] Haohan Guo, Heng Lu, Na Hu, Chunlei Zhang, Shan Yang, Lei Xie, Dan Su, and Dong Yu, "Phonetic posteriorgrams based many-to-many singing voice conversion via adversarial training," arXiv preprint, 2020.
- [14] Zhonghao Li, Benlai Tang, Xiang Yin, Yuan Wan, Ling Xu, Chen Shen, and Zejun Ma, "Ppg-based singing voice conversion with adversarial representation learning," *ICASSP*, pp. 7073–7077, 2021.
- [15] W. N. Hsu, B. Bolte, and Tsai Y H H, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," preprint, 2021.
- [16] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," preprint, 2019.
- [17] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," preprint, 2020.
- [18] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," preprint, 2019.
- [19] Disong Wang et al., "Vqmivc: Vector quantization and mutual information-based unsupervised speech representation disentanglement for one-shot voice conversion," *INTERSPEECH*, pp. 1344–1348, 2021.
- [20] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," *Advances in Neural Information Processing Systems*, pp. 17022–17033, 2020.
- [21] S. Zhang, M. Lei, Z. Yan, and L. Dai, "Deep-fsmn for large vocabulary continuous speech recognition," *ICASSP*, pp. 5869–5873, 2018.
- [22] Yuxuan Wang et al., "Tacotron: Towards end-to-end speech synthesis," *INTERSPEECH*, pp. 4006–4010, 2017.
- [23] C. Veaux, J. Yamagishi, and Kirsten MacDonald, *Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit*, 2019.
- [24] Z. Duan, H. Fang, B. Li, K. C. Sim, and Y. Wang, "The nus sung and spoken lyrics corpus: A quantitative comparison of singing and speech," *APSIPA ASC*, pp. 1–9, 2013.
- [25] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, vol. 2015.