

HQ-SVC: Towards High-Quality Zero-Shot Singing Voice Conversion in Low-Resource Scenarios

Bingsong Bai, Yizhong Geng, Fengping Wang, Cong Wang, Puyuan Guo,
Yingming Gao, Ya Li*

Beijing University of Posts and Telecommunications, China
{bingsongbai, yzgend, wfp, congwang, guopy, yingming.gao, yli01}@bupt.edu.cn

Abstract

Zero-shot singing voice conversion (SVC) transforms a source singer’s timbre to an unseen target speaker’s voice while preserving melodic content without fine-tuning. Existing methods model speaker timbre and vocal content separately, losing essential acoustic information that degrades output quality while requiring significant computational resources. To overcome these limitations, we propose HQ-SVC, an efficient framework for high-quality zero-shot SVC. HQ-SVC first extracts jointly content and speaker features using a decoupled codec. It then enhances fidelity through pitch and volume modeling, preserving critical acoustic information typically lost in separate modeling approaches, and progressively refines outputs via differentiable signal processing and diffusion techniques. Evaluations confirm HQ-SVC significantly outperforms state-of-the-art zero-shot SVC methods in conversion quality and efficiency. Beyond voice conversion, HQ-SVC achieves superior voice naturalness compared to specialized audio super-resolution methods while natively supporting voice super-resolution tasks.

Code — <https://github.com/ShawnPi233/HQ-SVC>

Introduction

Zero-shot singing voice conversion (SVC) aims to transform a singer’s timbre to that of another unseen target speaker while preserving the melody and content, without requiring prior data or fine-tuning for the target speaker. It can be widely applied in music production and virtual singers. However, zero-shot singing voice conversion faces challenges due to the high sampling rates and broader fundamental frequency (F_0) ranges in singing voices, making singing voice synthesis more difficult. Additionally, the intricate coupling between timbre and pitch further complicates the task. Moreover, the limited availability of high-quality singing datasets and computational resources highlights the need of studying zero-shot singing voice conversion in low-resource scenarios.

Existing SVC methods (Liu et al. 2021b; SVC-Develop-Team 2023; Zhou et al. 2023; Bai et al. 2024) have demonstrated high-quality singing voice conversion by leveraging generative models like VITS (Kim, Kong, and Son

2021) and diffusion (Ho, Jain, and Abbeel 2020) models. However, these models typically rely on explicit speaker IDs and separate content encoders (Hsu et al. 2021; Qian et al. 2022; Chen et al. 2022), which limit their generalization ability to unseen speakers, making them unsuitable for zero-shot singing voice conversion. FastSVC (Liu et al. 2021a) achieves zero-shot singing voice conversion with a lightweight model but struggles with poor synthesis quality. LDM-SVC (Chen et al. 2024) provides high-quality zero-shot singing voice conversion on small datasets but demands significant computational resources and training time due to its two-stage and adversarial training. SaMoye-SVC (Wang et al. 2024) also delivers high-quality zero-shot singing voice conversion but requires large datasets and extensive adversarial training, making it highly resource-intensive. Moreover, these methods often rely on separate content and speaker encoders, leading to information loss and hindering effective feature fusion and reconstruction of natural and expressive voice.

Given that both speech and singing belong to voice, the inherent similarities between them allow large-scale speech data to provide valuable priors, enabling transfer learning methods to effectively improve performance in zero-shot singing voice conversion tasks. Recent advancements in vector quantization (VQ) for zero-shot text-to-speech models like NaturalSpeech3 (Ju et al. 2024) and UniCATS (Du et al. 2024) have achieved significant success. NaturalSpeech3 introduces a VQ-based audio codec, FACodec, which turn audio into distinct latent spaces for content and speaker features simultaneously. This architecture allows zero-shot disentanglement of speaker and content features for unseen speakers, while also reducing information loss and making feature alignment easier compared to separate modeling. Although FACodec enables effective disentanglement of speaker and content features, it does not fully capture the complex acoustic variations necessary for high-quality synthesis, especially in challenging zero-shot singing voice conversion tasks.

Therefore, based on unified decoupled audio codec, we propose HQ-SVC, a novel framework designed to achieve high-quality zero-shot singing voice conversion in low-resource scenarios. Additionally, to improve high-quality synthesis modeling, we first leverage decoupled codec’s zero-shot disentanglement capability to extract content and

*Corresponding author.

speaker features from the decoder’s intermediate layer outputs. We then introduce supplementary features to improve the performance of SVC. To further enhance feature disentanglement and fusion, we propose the Enhanced Voice Adaptation (EVA) module, which incorporates additional acoustic features for multi-feature fusion, introduces a speaker loss, and includes a Speaker-F0 Predictor to enhance synthesis quality. Following the approach in DDSP-SVC (Yxl1lc 2023), the synthesis process is progressively optimized using Differentiable Digital Signal Processing (DDSP) (Engel et al. 2020) and diffusion (Ho, Jain, and Abbeel 2020) models to generate Mel spectrograms. Finally, a vocoder is employed to produce singing voices.

By avoiding unnecessary retraining and adversarial training, HQ-SVC is lightweight and time-efficient, enabling zero-shot singing voice conversion for unseen speakers with low computational resources and small datasets. Additionally, since the input features are extracted at a low sampling rate and aligned with high-sampling-rate Mel spectrograms, HQ-SVC also supports zero-shot voice super-resolution (SR). Therefore, the main contributions of this work are summarized as follows:

- We propose HQ-SVC, which achieves high-quality zero-shot singing voice conversion under low-resource scenarios, such as training on a single consumer-grade GPU and small-scale datasets. HQ-SVC outperforms the current SOTA method, across most objective and subjective evaluation metrics.
- We introduce the EVA module to improve multi-feature fusion and enhance synthesis quality. It incorporates speaker loss and a Speaker-F0 Predictor, enabling more natural singing voice conversion.
- HQ-SVC extends its capabilities to zero-shot singing and speech SR tasks. Compared to the current SOTA audio SR method, HQ-SVC generates voice with higher naturalness and speaker similarity. Our work may inspire further research on using zero-shot singing voice conversion methods to achieve higher-quality voice SR tasks.

Related Works

Singing Voice Conversion

Singing Voice Conversion (SVC) aims to transform a singer’s timbre into that of another target singer while preserving the melody and content. Early SVC methods (Doi et al. 2012; Kobayashi et al. 2014, 2015) were primarily based on Gaussian Mixture Models (GMM), which relied on parallel data. These methods require both source and target singers to perform the same song. This dependency not only demanded a large amount of parallel data but also limited the generation quality of the models. With the rise of Generative Adversarial Networks (GANs) (Goodfellow et al. 2020), several GAN-based SVC approaches (Sisman et al. 2019; Sisman and Li 2020; Liu et al. 2021a) emerged that operate without parallel data. These methods employ adversarial training, where a generator and a discriminator engage in a game to improve the audio quality. However, GAN-based methods typically take longer to train and can suffer from

issues like mode collapse, making the overall training process more challenging. SoVITS-SVC (SVC-Develop-Team 2023) builds on adversarial training by adding a flow-based model to prevent mode collapse during training. It is based on VITS (Kim, Kong, and Son 2021), which uses an end-to-end flow-based architecture with adversarial training to generate high-quality singing voices. While effective, it still requires considerable computational resources for training.

High-Quality SVC in Low-resource Scenarios

Open-source singing voice datasets available for academic research remain relatively scarce. At the same time, the growing practical demand for efficient training and inference on consumer-grade GPUs has led to increasing research attention on achieving high-quality SVC under low-resource conditions. Low-resource SVC means using only a small amount of singing voice data and limited GPU power during training, while still producing high-quality audio. Several methods have been proposed to address this challenge. These include DDSP-SVC (Yxl1lc 2023) based on Differentiable Digital Signal Processing (DDSP) models (Engel et al. 2020) and diffusion-based SVC methods (Liu et al. 2021b; Lu et al. 2024; Huang et al. 2025). These methods not only demonstrate superior generation quality, but also exhibit strong training stability, enabling high-quality singing voice conversion even in low-resource scenarios.

Zero-shot SVC

Nevertheless, while the aforementioned SVC methods have made significant strides, their reliance on explicit speaker identities restricts their ability to generalize to unseen speakers, making them unsuitable for zero-shot SVC scenarios. In zero-shot SVC, the objective is to achieve high-quality conversion without requiring target singer data or fine-tuning. Traditional SVC approaches typically depend on parallel data and explicit speaker labels, and in practical applications, issues such as high sampling rates, broader pitch ranges, and the intricate coupling between timbre and pitch present significant challenges in modeling and training. FastSVC (Liu et al. 2021a) is one of the earlier zero-shot singing voice conversion methods; it leverages GANs combined with FiLM (Perez et al. 2018) to enable a lightweight zero-shot conversion. However, its generation quality leaves room for improvement. In contrast, methods like LDM-SVC (Chen et al. 2024), FreeSVC (Ferreira et al. 2025), and SaMoye-SVC (Wang et al. 2024), which are based on VITS, can synthesize high-quality singing voices in a zero-shot manner, although at the cost of requiring considerable computational power during training.

Methods

As shown in Fig. 1 (a), HQ-SVC is built upon decoupled audio codec to achieve disentangled content features x_{con} and speaker embeddings x_{spk} . EVA module is proposed to integrate other acoustic features such as pitch and energy, enhancing the capture of vocal details. Furthermore, HQ-SVC aligns with high-sampling-rate Mel spectrograms by optimizing the reconstruction loss within both DDSP and diffusion models. This enables high-quality zero-shot singing

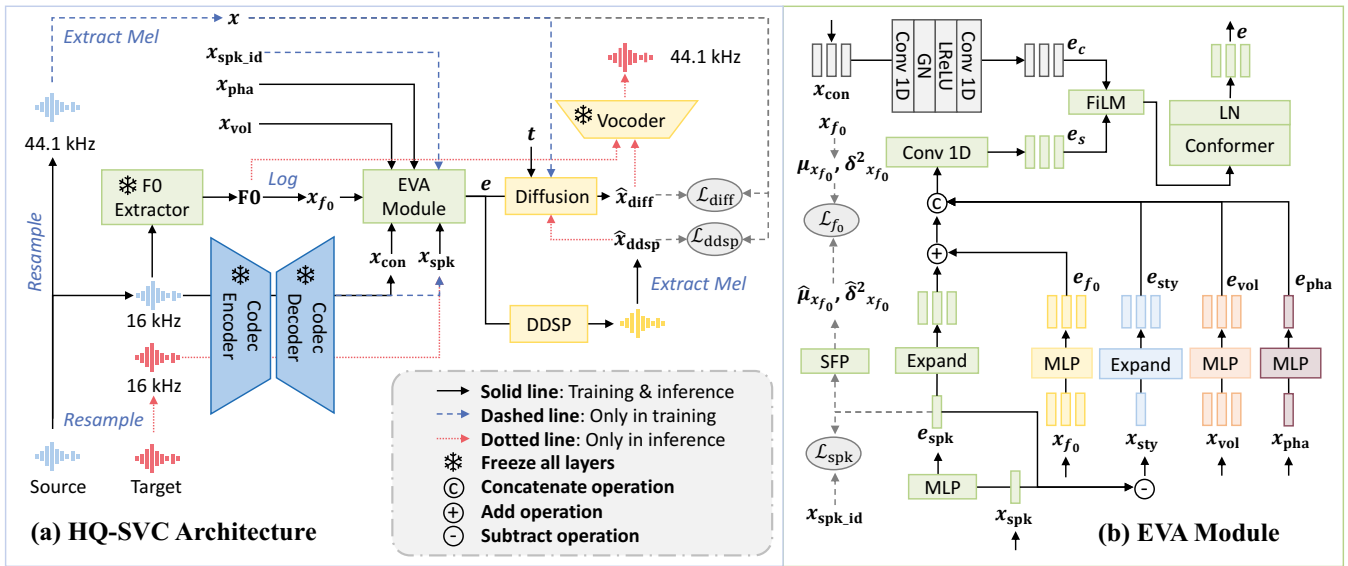


Figure 1: Overall architecture of the proposed HQ-SVC and the architecture of the proposed EVA module.

voice conversion and even voice super-resolution. In the following sections, we will provide a detailed explanation of the model architecture and the modules employed.

Decoupled Codec

We utilize FAcCodec (Ju et al. 2024) as the unified disentangler, due to its high generation quality and strong disentanglement capability. Specifically, we freeze all of the layers of the encoder and decoder of the FAcCodec and get the x_{con} and x_{spk} as the content and speaker features respectively. The basic architecture of FAcCodec follows Kumar et al. (Kumar et al. 2024) and employs the SnakeBeta activation function (gil Lee et al. 2023). The speaker extractor consists of several Conformer (Gulati et al. 2020) blocks. For the three factorized vector quantization (FVQ) components Q^c , Q^p , Q^d , we use $N_{qc} = 2$, $N_{qp} = 1$, $N_{qd} = 3$ quantizers respectively, with all quantizers having a codebook size of 1024.

DDSP Model

DDSP (Engel et al. 2020) is a method that integrates classical digital signal processing (DSP) techniques with deep learning. It introduces strong inductive biases while preserving the expressive power of neural networks, enabling high-fidelity audio synthesis without relying on large autoregressive models or adversarial losses. Its modular architecture allows for fine-grained control over attributes such as timbre and loudness. Specifically, the DDSP synthesizer consists of a harmonic synthesizer and a noise synthesizer, which generate periodic and aperiodic components respectively and combine them to produce the final waveform. Incorporating DDSP into the SVC task enhances control over audio attributes, thereby improving conversion performance.

Diffusion Model

Moreover, to improve the quality of vocal synthesis, the diffusion (Ho, Jain, and Abbeel 2020) model is utilized to pro-

cess the output of DDSP. Integrating the diffusion model with DDSP enriches the description of sound characteristics by filling in any gaps, thereby boosting the performance of SVC. The strong performance of diffusion model in SVC tasks has been demonstrated in DiffSVC (Liu et al. 2021b) and CoMoSVC (Lu et al. 2024).

Enhanced Voice Adaption Module

Singing voices are often rich in melody and energy dynamics, making it challenging to ensure high-quality generation using only content and speaker features. To address this, we propose the Enhanced Voice Adaption (EVA) module, which introduces RMVPE (Wei et al. 2023) as an F0 encoder to extract pitch features and incorporates energy features along with phase features required by the DDSP (Engel et al. 2020) model for multi-feature fusion.

Acoustic Feature Extraction and Mapping As shown in Fig. 1 (b), since FAcCodec’s speaker features x_{spk} may include unrelated characteristics like style, they are mapped to speaker embeddings e_{spk} using an MLP. And the residual between the speaker features x_{spk} and speaker embeddings e_{spk} is computed to obtain residual style features x_{sty} , which is expanded to yield e_{sty} that matches the sequence length dimension. Following the approach in DDSP-SVC (Yxlilc 2023), we perform the transformation on the F0 to obtain $x_{f0} = \log_e(\frac{f_0}{700} + 1)$ to help the model better capture the melody. And we map it to pitch embeddings e_{f0} through an MLP. Similarly, volume feature x_{vol} and phase x_{pha} feature are fed into MLPs to generate their respective embeddings e_{vol} and e_{pha} . All the above MLPs consist of two linear layers, with a SiLU (Elfwing, Uchibe, and Doya 2018) activation function applied between them. The output dimension of MLPs is 256.

Fusion of Style and Content Embeddings Given the close relationship between pitch and timbre, while other fea-

tures remain relatively independent, we combine the speaker and f0 embeddings into a unified representation. This combined representation is concatenated with the remaining embeddings to form a 1024-dimensional style embedding

$$e_s = \text{Concat}(e_{\text{spk}} + e_{f_0}, e_{\text{sty}}, e_{\text{vol}}, e_{\text{pha}}), \quad (1)$$

which is subsequently compressed to 256 dimensions via a 1D convolution to match the content. The content features x_{con} , after convolution and normalization, are fused with the aforementioned style embeddings using FiLM (Perez et al. 2018), as described by

$$\text{FiLM}(e_c, e_s) = f_\alpha(e_s) \cdot e_c + f_\beta(e_s), \quad (2)$$

where e_c denotes the content embeddings, e_s represents the style embeddings, and $f_\alpha(\cdot)$ and $f_\beta(\cdot)$ correspond to two independent linear layers. To achieve more effective fusion of features, the FiLM output is passed through a Conformer (Gulati et al. 2020) with 8-head self-attention (Vaswani 2017) and layer normalization, producing e , which is then used as input for the DDSP and diffusion models.

In addition, to enhance the model’s performance in zero-shot singing voice conversion for unseen speakers, we introduce the speaker loss \mathcal{L}_{spk} based on the InfoNCE loss (Oord, Li, and Vinyals 2018), a widely-used contrastive learning objective. It helps the model distinguish speakers by pulling positive samples closer and pushing negative samples farther within a batch:

$$\mathcal{L}_{\text{spk}} = -\frac{1}{N} \sum_{i=1}^N \log \left(\frac{\exp \left(\frac{e_{\text{spk},i} \cdot e_{\text{spk},j^+}}{\tau} \right)}{\sum_{j=1}^N \exp \left(\frac{e_{\text{spk},i} \cdot e_{\text{spk},j^-}}{\tau} \right)} \right), \quad (3)$$

where $e_{\text{spk},i}$ represents embedding of the i -th speaker. The positive sample e_{spk,j^+} matches the same speaker ID $x_{\text{spk_id}}$, while negative sample e_{spk,j^-} comes from different speaker IDs. τ is the temperature parameter that controls the sharpness of the similarity distribution. N is the batch size.

The strong coupling between pitch range and timbre makes it essential to obtain statistical characteristics of pitch to generate singing voices that closely match the timbre of the target speaker. However, in zero-shot singing voice conversion tasks, it is typically infeasible to acquire sufficient pitch data from the target speaker to derive them. To address this, we designed a simple yet effective Speaker-F0 Predictor (SFP) $f_{\text{sfp}}(\cdot)$ based on an MLP. It uses shared layers extract features from speaker embeddings e_{spk} , followed by separate branches to predict mean and variance

$$\hat{\mu}_{x_{f_0}}, \hat{\sigma}_{x_{f_0}}^2 = f_{\text{sfp}}(e_{\text{spk}}). \quad (4)$$

Specifically, it takes e_{spk} as input to predict the mean and variance of the x_{f_0} , and the loss is calculated by minimizing the L1 loss between the predicted values and the actual mean and variance of the x_{f_0} sequence of the input. The calculation formula is

$$\mathcal{L}_{f_0} = \mathbb{E}[\|\mu_{x_{f_0}} - \hat{\mu}_{x_{f_0}}\|_1 + \|\sigma_{x_{f_0}}^2 - \hat{\sigma}_{x_{f_0}}^2\|_1]. \quad (5)$$

Singing Voice Reconstruction

We first pass the embedding e from EVA module through DDSP model to generate audio. We then turn this audio into

a Mel spectrogram \hat{x}_{ddsp} , and minimize the Mean Squared Error (MSE) loss between \hat{x}_{ddsp} and the real Mel spectrogram x :

$$\mathcal{L}_{\text{ddsp}} = \mathbb{E}[\|x - \hat{x}_{\text{ddsp}}\|_2^2]. \quad (6)$$

Additionally, we pass e to the diffusion model for further optimization. We use the WaveNet (Oord et al. 2016) from FastSpeech2 (Ren et al. 2020) as the diffusion denoiser, configured with 128-dimensional input features, 20 residual block layers, 512 convolutional output channels, and a 256-dimensional encoder hidden layer. The formula for the diffusion loss is given by:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{t,x_0,\epsilon} [\|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, t, e)\|_2^2], \quad (7)$$

where θ denotes the model parameters, ϵ denotes random noise, x_0 denotes the clean Mel spectrogram, $\bar{\alpha}_t$ is a parameter in the model, and t is the time step. The expectation expression encapsulates the expectation value over all possible t , x_0 , and ϵ . Thus, the total loss of HQ-SVC is

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{ddsp}} + \mathcal{L}_{\text{diff}} + \mathcal{L}_{\text{spk}} + \mathcal{L}_{f_0}. \quad (8)$$

During inference, DPM-Solver++ (Lu et al. 2022) is used for sampling in the diffusion model, with chunk-wise processing adopted to reduce inference latency, 100 diffusion steps and a speed-up rate of 10 to balance quality and speed. To improve synthesis quality, HQ-SVC uses NSF-HiFiGAN (OpenVPI 2022), an enhanced version of HiFi-GAN (Kong, Kim, and Bae 2020) with a Neural Source-Filter (NSF) (Wang, Takaki, and Yamagishi 2019) module, to convert Mel spectrogram with F0 into audio.

Experiments

Datasets

Training Datasets In this study, we utilized the high-quality, open-source Mandarin singing datasets Opensinger (Huang et al. 2021) and M4Singer (Zhang et al. 2022) for model training. From Opensinger, 2 male and 2 female singers were randomly selected as unseen singers. These singers were excluded from both training and validation and were only used for model testing. Additionally, 1% of the remaining singing data from M4Singer and Opensinger was randomly selected as a validation set to monitor model convergence. To avoid potential noise and performance degradation caused by excessively short audio samples, we excluded audio clips shorter than 2.1 seconds.

Data Preprocessing To enable zero-shot singing voice conversion in low-resource scenarios, we utilized frozen pre-trained models for feature extraction, effectively reducing memory consumption during training and inference. Since the singing audio files have two different rates: 44.1 kHz and 48 kHz, we resampled them to 44.1 kHz. Then, we used a hop size of 512 and obtained 128-dimensional Mel spectrograms and energy features from the audios. Next, we down-sampled the audios to 16 kHz, used RMVPE to extract pitch features, and employed the FACodec (Ju et al. 2024; Zhang et al. 2024) to extract 256-dimensional content and speaker features. Furthermore, we constructed a lookup table to map all speakers across datasets to unique speaker IDs in a unified table for speaker loss calculation.

Method	Train Config.	Dataset	STOI (↑)	SECS (↑)	F0 RMSE (↓)	FPC (↑)	NISQA (↑)	NMOS (↑)	SMOS (↑)
FACodec-SVC	RTX 3090 (1 h)	<80 h	0.533	0.074	77.798	0.601	1.791	2.391 ± 0.201	2.740 ± 0.192
SaMoye-SVC	A100 (7 days)	1700 h	0.724	0.647	17.418	0.617	3.528	3.958 ± 0.154	3.569 ± 0.147
HQ-SVC (Ours)	RTX 3090 (11 h)	<80 h	<u>0.799</u>	<u>0.627</u>	8.681	0.891	3.841	4.215 ± 0.124	3.578 ± 0.192
HQ-SVC w/o \mathcal{L}_{spk}	RTX 3090 (11 h)	<80 h	0.800	0.622	<u>8.828</u>	0.889	3.734	4.174 ± 0.135	<u>3.688 ± 0.183</u>
HQ-SVC w/o \mathcal{L}_{f_0}	RTX 3090 (11 h)	<80 h	0.800	0.622	8.855	<u>0.889</u>	<u>3.764</u>	<u>4.195 ± 0.138</u>	3.729 ± 0.168

Table 1: Objective and subjective evaluation results of zero-shot singing voice conversion. Bold and underline values indicate the best and second best results, respectively. Confidence interval for MOS scores is 95%.




Figure 2: Zero-shot singing voice conversion spectral comparison. The red curve represents the F0 contour, with red boxes highlighting differences in the F0 regions across the spectrograms. Yellow boxes indicate high-frequency harmonics.

Experimental Setups

Model Training During training, we employed the AdamW (Loshchilov and Hutter 2017) optimizer with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, an initial learning rate of 1.5×10^{-4} , a batch size of $N = 64$. For the speaker loss, the temperature parameter was set to $\tau = 0.1$. The model was trained for 250k steps on a single NVIDIA RTX 3090 GPU. Notably, HQ-SVC occupies less than 6 GB of GPU memory during training and can be trained within about 11 hours, making it well-suited for zero-shot singing voice conversion in low-resource scenarios.

Model Evaluation on Zero-shot Singing Voice Conversion We conducted a comparative evaluation of the proposed HQ-SVC, against the SOTA zero-shot singing voice conversion model. Since most zero-shot SVC methods are not open-source, the VITS-based SaMoye-SVC (Wang et al. 2024) model pretrained on 1700 hours of clean singing dataset, was used as a strong baseline for comparison, whereas HQ-SVC was trained on less than 80 hours of singing data. To further validate the effectiveness of our improvements to the FACodec, we fine-tuned FACodec. For fairness, we modified FACodec’s output structure to produce Mel spectrograms, which were combined with F0 and synthesized into singing audio using NSF-HiFiGAN (OpenVPI 2022), like HQ-SVC. In addition, all generated audios were resampled to 32 kHz for zero-shot singing voice conversion evaluation, since HQ-SVC generates audios at 44.1 kHz, while SaMoye outputs at 32 kHz.

We assessed the models on unseen speakers using objective measures—intelligibility (STOI), timbre similarity (SECS), pitch accuracy (F0 RMSE and FPC), overall quality (NISQA) (Mittag et al. 2021) and subjective 5-point ratings of naturalness (NMOS) and speaker similarity (SMOS).

For MOS ratings, 12 volunteers participated in the evaluation. To evaluate zero-shot singing voice conversion methods in unseen scenarios, we randomly selected 4 audio samples from 2 male and 2 female singers in the Opensinger test set, giving us 8 source and 8 target audio samples. By performing cross-conversion between all source and target samples, we generated a total of 128 audio samples for objective evaluation. For subjective evaluation, 8 samples were randomly selected, totaling 96 per model. Fewer samples were used compared to the objective tests to maintain evaluator focus and ensure more consistent and reliable scoring.

Model Evaluation on Voice Super-resolution HQ-SVC utilizes 16 kHz data for feature extraction during training and employs authentic 44.1 kHz Mel spectrograms for model optimization. This design allows HQ-SVC to achieve zero-shot voice super-resolution (SR) using the zero-shot singing voice conversion method while preserving the speaker’s identity. To demonstrate this, we compared HQ-SVC with the current SOTA method AudioSR (Liu et al. 2024). AudioSR uses a Latent Diffusion Model (LDM) trained on over 7,000 hours of sound effects, speech, and music, allowing it to upscale audio to 48 kHz.

Since HQ-SVC uniformly resamples ground-truth audio to 44.1 kHz during training, we set its output sampling rate to 44.1 kHz for a fair comparison. We ensure fairness and demonstrate the generalization of HQ-SVC on Voice Super-resolution tasks for both singing and speech using unseen test datasets. These include Mandarin Opensinger and English NHSS-Song (Sharma et al. 2021) (singing datasets), as well as English LibriTTS-P (Kawamura et al. 2024) and NHSS-Speech (speech datasets).

For evaluation on voice super-resolution task, we randomly selected 2 male and 2 female speakers from each

of the four datasets, totaling 8 male and 8 female speakers. From each speaker, we randomly chose 8 audio samples, giving us 128 samples for objective evaluation. For subjective evaluation, we randomly selected 1 audio sample per speaker, resulting in 8 singing and 8 speech samples. This yielded a total of 96 singing and 96 speech samples per model. The objective metrics include STOI, SECS, FPC, NISQA, and log-spectral distance (LSD) (Gray and Markel 1976) metric adopted by AudioSR, while the subjective metrics include NMOS and SMOS.

Experimental Results and Analysis

Zero-shot Singing Voice Conversion Comparison The experimental results for zero-shot singing voice conversion are presented in Table 1. First, a comparison between HQ-SVC and FACodec-SVC reveals that HQ-SVC significantly outperforms FACodec-SVC across all metrics. This demonstrates the effectiveness of integrating the EVA module for multi-feature fusion, as well as leveraging DDSP and diffusion models for performance optimization. Furthermore, despite being trained in low-resource scenarios, HQ-SVC outperforms SaMoye, which was trained on large-scale datasets for an extended period. HQ-SVC achieves superior results in objective metrics such as STOI and demonstrates notable advantages in F0 RMSE, FPC, and NISQA. Additionally, it maintains an advantage in subjective metrics like NMOS and SMOS, indicating that HQ-SVC achieves higher performance in preserving melodic consistency, as well as enhancing the naturalness and similarity of singing voices.

However, we noticed that while SaMoye performs slightly worse in SMOS, it performs better in the SECS metric. This difference may be due to how SECS and human perception respond to pitch. SECS is calculated using cosine similarity from speaker embeddings extracted by the CAM++ (Wang et al. 2023) speaker verification model. These embeddings implicitly contain statistical information about pitch, allowing model to effectively identify the speaker in the converted audio, even when the pitch remains unchanged. In contrast, human perception of speaker identity is more sensitive to variations in pitch. Thus, even if the speaker characteristics are altered, a consistent pitch can make it difficult for humans to identify the speaker.

The results of the ablation experiments further validate the above observations. Removing either the \mathcal{L}_{f_0} or the \mathcal{L}_{spk} leads to a varying degrees of decline in SECS, F0 RMSE, NISQA, and NMOS. This indicates that these two losses contribute to generating singing voices with more accurate pitch, while ensuring that the synthesized voices remain natural and are more easily distinguishable by the speaker verification model. We also observed a slight improvement in SMOS when \mathcal{L}_{f_0} or \mathcal{L}_{spk} is removed. This further supports the earlier point: adding \mathcal{L}_{f_0} and \mathcal{L}_{spk} enhances pitch-related information in the speaker embeddings, making them more identifiable by CAM++ (Wang et al. 2023) and better at maintaining a consistent F0. This consistency makes it difficult for human listeners to perceive changes in speaker characteristics, particularly when the target speaker has a significantly different vocal range, thus leading to a decrease in SMOS. Nevertheless, the timbre similarity of HQ-SVC




Figure 3: Voice super-resolution spectral comparison.

Method	STOI(↑)	SECS(↑)	FPC(↑)	LSD(↓)	NISQA(↑)
Ground Truth	-	-	-	-	4.234
Downsampled	-	-	-	-	4.181
AudioSR	0.986	0.759	0.998	2.087	4.094
HQ-SVC (Ours)	0.841	0.766	0.868	1.842	4.193

Table 2: Objective evaluation results of voice super-resolution.

remains superior to other baseline models.

To provide a more intuitive demonstration of HQ-SVC’s performance in zero-shot singing voice conversion, we visualized and analyzed the spectrograms, as shown in Fig. 2. Comparing the F0 contours of the source audio and those generated by different models, we found that HQ-SVC’s F0 is closest to the source, while SaMoye has overly smooth F0 contours and FACodec-SVC shows large fluctuations, both deviating from the source. These findings align with the objective experimental results, where HQ-SVC outperformed SaMoye and FACodec-SVC in F0 RMSE and FPC metrics. For high-frequency harmonics, SaMoye also shows overly smoothing, creating an unnatural sound, while FACodec-SVC adds significant noise, making the singing voice less clear. These spectrogram findings align with the objective and subjective results mentioned above.

Voice Super-resolution Comparison The objective experimental results for voice super-resolution are shown in Table 2. HQ-SVC shows clear advantages in metrics like LSD and NISQA, producing higher-quality voices closer to the ground truth. These results highlight HQ-SVC’s ability to generate high-fidelity audio and effectively reconstruct details in voice super-resolution tasks, even in low-resource situations. However, HQ-SVC performs poorly in STOI and FPC compared to AudioSR, indicating that AudioSR produces audio with better intelligibility in terms of F0 and content. This difference is mainly due to the distinct training methods of the models. AudioSR, a specialized audio super-resolution model, is trained on downsampled and high-sample-rate Mel spectrogram pairs, helping

Method	Dataset	NMOS(\uparrow)	SMOS(\uparrow)
Ground Truth	-	4.667 \pm 0.089	-
Downsampled	-	4.083 \pm 0.100	-
AudioSR	7000 h	4.188 \pm 0.103	4.235 \pm 0.096
HQ-SVC (Ours)	<80 h	4.332 \pm 0.088	4.479 \pm 0.087

Table 3: Subjective evaluation results of voice super-resolution with the confidence interval 95%.

Method	Speed	RTF (\downarrow)	STOI (\uparrow)	SECS (\uparrow)	F0 RMSE (\downarrow)	FPC (\uparrow)	NISQA (\uparrow)
Dpm-solver++	1 \times	0.295	0.801	0.620	<u>8.578</u>	0.888	3.803
	5 \times	0.090	0.799	<u>0.624</u>	8.575	0.888	<u>3.825</u>
	10 \times	0.065	0.799	0.627	8.681	0.891	3.841
	20 \times	0.054	0.808	0.570	8.695	0.889	3.684
UniPC	1 \times	0.275	0.801	0.619	8.611	0.888	3.798
	5 \times	0.085	0.802	0.612	<u>8.578</u>	0.889	3.775
	10 \times	0.061	<u>0.804</u>	0.598	8.699	<u>0.890</u>	3.737
	20 \times	0.048	0.808	0.570	8.767	0.889	3.686
DDIM	1 \times	0.203	0.801	0.617	8.652	0.888	3.789
	5 \times	0.067	0.801	0.613	8.745	0.889	3.773
	10 \times	0.060	0.801	0.612	8.754	0.889	3.763
	20 \times	0.043	0.799	0.618	8.769	0.889	3.782

Table 4: Performance comparisons in different sampling methods at different acceleration in singing voice conversion tasks.

it better preserve low-frequency details like pronunciation and F0. In contrast, HQ-SVC, as a zero-shot singing voice conversion model, focuses on controllability and generates high-sample-rate audio by decoupling features and directly reconstructing them into Mel spectrograms. While this approach offers more flexibility and control, it introduces minor and acceptable errors in pronunciation and pitch.

To evaluate the subjective quality of audio generated by HQ-SVC, we conducted an MOS evaluation, as presented in Table 3. The results show that HQ-SVC outperforms AudioSR in both NMOS and SMOS, confirming its ability to produce high-quality and natural voice super-resolution in low-resource scenarios.

The spectral comparison in Fig. 3 shows that HQ-SVC up-scales low-sample-rate speech to 44.1 kHz with clearer mid-to-high frequency harmonics (yellow boxes) and less high-frequency noise (red boxes) than AudioSR, generating more natural voices. This is mainly due to the HQ-SVC, which does not simply rely on increasing the signal-to-noise ratio to reconstruct the audio. Instead, it leverages multi-feature fusion, making the audio sound more like natural voice. This aligns with our previous objective and subjective results.

Sampler and Acceleration Rate Comparison in HQ-SVC

Table 4 presents a comparison of zero-shot singing voice conversion performance under various acceleration rates (1 \times , 2 \times , 10 \times , 20 \times) using three diffusion-based samplers: Dpm-solver++ (Lu et al. 2022), UniPC (Zhao et al. 2023), and DDIM (Song, Meng, and Ermon 2020).

We compare diffusion samplers and acceleration rates using RTF, STOI, SECS, F0 RMSE, FPC, and NISQA. In this study, which focuses on comparing the performance of dif-

Method	STOI(\uparrow)	SECS(\uparrow)	F0 RMSE(\downarrow)	FPC(\uparrow)	NISQA(\uparrow)
HQ-SVC	0.799	0.627	8.681	0.891	3.841
w/o Diffusion	0.820	0.420	8.547	0.888	3.175
w/o DDSP	0.847	0.608	8.224	0.901	3.960
HQ-SVC-SE	0.738	0.668	9.803	0.873	3.880
w/o \mathcal{L}_{spk}	0.735	0.660	9.747	0.873	3.780
w/o \mathcal{L}_{f_0}	0.739	0.663	9.814	0.871	3.790

Table 5: Ablation studies in singing voice conversion.

ferent samplers and acceleration rates, we define RTF as the ratio of the time taken to convert the embeddings into audio to the duration of the generated audio. The results show that Dpm-solver++ at 10 \times gives the best speed-quality trade-off.

Ablation Studies We did more ablation studies to check how well each part works and to better understand the results. As shown in Table 5, taking out the diffusion model made speech a bit clearer but much worse in terms of matching the original voice (timbre) and natural sound. This shows the diffusion model is important for keeping the speaker’s voice and making converted speech sound natural. On the other hand, removing the DDSP module improved some scores but also made timbre matching worse. This might look like it’s better for singing, but because the diffusion model uses the original audio’s Mel spectrogram, the output stays close to the original. So even with higher scores, it doesn’t actually convert voices well.

We also compared HQ-SVC with a version using separate encoders (HQ-SVC-SE). In HQ-SVC-SE, we replaced FA-Codec with two separate encoders: CAM++ (for speakers) and ContentVec (for content). ContentVec is based on HuBERT and is designed to remove timbre info. The comparison showed HQ-SVC’s single encoder is better at combining different types of information and controlling rhythm. But HQ-SVC still isn’t as good as HQ-SVC-SE at matching timbre, which means we need to get better at separating out the speaker’s identity. Future work could use methods like changing pitch or speaker characteristics to improve this. Also, removing either the speaker loss or pitch loss hurt how well the model handled pitch (F0) and timbre matching, showing these are important for good voice conversion.

In short, HQ-SVC balances performance across different measures. It might not be the best in every single area, but its complete design, reliability, and ability to work in different situations make it the best choice as the core of our system.

Conclusions

This paper introduces HQ-SVC, a zero-shot singing voice conversion method for low-resource scenarios. It uses a decoupled codec to disentangle content and speaker features and integrates additional acoustic features with the proposed Enhanced Voice Adaption module. Both objective and subjective experiments show that HQ-SVC achieves high naturalness and similarity in zero-shot singing voice conversion, and even performs well in zero-shot voice super-resolution, despite limited GPU resources and a small training dataset. Future work will focus on singing voice style conversion.

Acknowledgments

This work is supported by the National Key R&D Program of China under Grant No.2024YFB2808802.

References

- Bai, B.; Wang, F.; Gao, Y.; and Li, Y. 2024. SPA-SVC: Self-supervised Pitch Augmentation for Singing Voice Conversion. In *Interspeech 2024*, 4353–4357.
- Chen, S.; Gu, Y.; Zhang, J.; Li, N.; Chen, R.; Chen, L.; and Dai, L. 2024. LDM-SVC: Latent Diffusion Model Based Zero-Shot Any-to-Any Singing Voice Conversion with Singer Guidance. In *Proc. Interspeech 2024*, 2770–2774.
- Chen, S.; Wang, C.; Chen, Z.; Wu, Y.; Liu, S.; Chen, Z.; Li, J.; Kanda, N.; Yoshioka, T.; Xiao, X.; et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6): 1505–1518.
- Doi, H.; Toda, T.; Nakano, T.; Goto, M.; and Nakamura, S. 2012. Singing voice conversion method based on many-to-many eigenvoice conversion and training data generation using a singing-to-singing synthesis system. In *Proceedings of The 2012 Asia Pacific Signal and Information Processing Association Annual Summit and Conference*, 1–6. IEEE.
- Du, C.; Guo, Y.; Shen, F.; Liu, Z.; Liang, Z.; Chen, X.; Wang, S.; Zhang, H.; and Yu, K. 2024. UniCATS: A unified context-aware text-to-speech framework with contextual vq-diffusion and vocoding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17924–17932.
- Elfwing, S.; Uchibe, E.; and Doya, K. 2018. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks*, 107: 3–11.
- Engel, J.; Gu, C.; Roberts, A.; et al. 2020. DDSF: Differentiable Digital Signal Processing. In *International Conference on Learning Representations*.
- Ferreira, A. I.; Gris, L. R.; Da Rosa, A.; Oliveira, F.; Casanova, E.; Sousa, R.; Junior, A.; Soares, A.; and Galvão Filho, A. 2025. FreeSVC: Towards Zero-shot Multilingual Singing Voice Conversion. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- gil Lee, S.; Ping, W.; Ginsburg, B.; Catanzaro, B.; and Yoon, S. 2023. BigVGAN: A Universal Neural Vocoder with Large-Scale Training. In *The Eleventh International Conference on Learning Representations*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Gray, A.; and Markel, J. 1976. Distance measures for speech processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 24(5): 380–391.
- Gulati, A.; Qin, J.; Chiu, C.-C.; Parmar, N.; Zhang, Y.; Yu, J.; Han, W.; Wang, S.; Zhang, Z.; Wu, Y.; and Pang, R. 2020. Conformer: Convolution-augmented Transformer for Speech Recognition. In *Interspeech 2020*, 5036–5040.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hsu, W.-N.; Bolte, B.; Tsai, Y.-H. H.; Lakhotia, K.; Salakhutdinov, R.; and Mohamed, A. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29: 3451–3460.
- Huang, R.; Chen, F.; Ren, Y.; Liu, J.; Cui, C.; and Zhao, Z. 2021. Multi-singer: Fast multi-singer singing voice vocoder with a large-scale corpus. In *Proceedings of the 29th ACM International Conference on Multimedia*, 3945–3954.
- Huang, Y.; Lai, X.; Ye, M.; Zhu, A.; Wang, Z.; Xu, J.; Zhang, S.; Zhou, Z.; and Niu, W. 2025. LHQ-SVC: Lightweight and High Quality Singing Voice Conversion Modeling. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Ju, Z.; Wang, Y.; Shen, K.; Tan, X.; Xin, D.; Yang, D.; Liu, E.; Leng, Y.; Song, K.; Tang, S.; et al. 2024. NaturalSpeech 3: Zero-Shot Speech Synthesis with Factorized Codec and Diffusion Models. In *Forty-first International Conference on Machine Learning*.
- Kawamura, M.; Yamamoto, R.; Shirahata, Y.; Hasumi, T.; and Tachibana, K. 2024. LibriTTS-P: A Corpus with Speaking Style and Speaker Identity Prompts for Text-to-Speech and Style Captioning. In *Interspeech 2024*, 1850–1854.
- Kim, J.; Kong, J.; and Son, J. 2021. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In *International Conference on Machine Learning*, 5530–5540. PMLR.
- Kobayashi, K.; Toda, T.; Neubig, G.; Sakti, S.; and Nakamura, S. 2014. Statistical singing voice conversion with direct waveform modification based on the spectrum differential. In *Interspeech*, 2514–2518.
- Kobayashi, K.; Toda, T.; Neubig, G.; Sakti, S.; and Nakamura, S. 2015. Statistical singing voice conversion based on direct waveform modification with global variance. In *Interspeech*, 2754–2758. Citeseer.
- Kong, J.; Kim, J.; and Bae, J. 2020. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in neural information processing systems*, 33: 17022–17033.
- Kumar, R.; Seetharaman, P.; Luebs, A.; Kumar, I.; and Kumar, K. 2024. High-fidelity audio compression with improved rvqgan. *Advances in Neural Information Processing Systems*, 36.
- Liu, H.; Chen, K.; Tian, Q.; Wang, W.; and Plumbley, M. D. 2024. AudioSR: Versatile audio super-resolution at scale. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1076–1080. IEEE.
- Liu, S.; Cao, Y.; Hu, N.; Su, D.; and Meng, H. 2021a. Fastsvc: Fast cross-domain singing voice conversion with feature-wise linear modulation. In *2021 IEEE International Conference on Multimedia and Expo (icme)*, 1–6. IEEE.

- Liu, S.; Cao, Y.; Su, D.; and Meng, H. 2021b. Diffsvc: A diffusion probabilistic model for singing voice conversion. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 741–748. IEEE.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*.
- Lu, Y.; Ye, Z.; Xue, W.; Tan, X.; Liu, Q.; and Guo, Y. 2024. CoMoSVC: Consistency model-based singing voice conversion. In *2024 IEEE 14th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, 184–188. IEEE.
- Mittag, G.; Naderi, B.; Chehadi, A.; and Möller, S. 2021. NISQA: A Deep CNN-Self-Attention Model for Multidimensional Speech Quality Prediction with Crowdsourced Datasets. *Interspeech 2021*.
- Oord, A. v. d.; Dieleman, S.; Zen, H.; Simonyan, K.; Vinyals, O.; Graves, A.; Kalchbrenner, N.; Senior, A.; and Kavukcuoglu, K. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- OpenVPI. 2022. DiffSinger (OpenVPI maintained version). <https://github.com/openvpi/DiffSinger>. Accessed: 2023-08-23.
- Perez, E.; Strub, F.; De Vries, H.; Dumoulin, V.; and Courville, A. 2018. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Qian, K.; Zhang, Y.; Gao, H.; Ni, J.; Lai, C.-I.; Cox, D.; Hasegawa-Johnson, M.; and Chang, S. 2022. Contentvec: An improved self-supervised speech representation by disentangling speakers. In *International Conference on Machine Learning*, 18003–18017. PMLR.
- Ren, Y.; Hu, C.; Tan, X.; Qin, T.; Zhao, S.; Zhao, Z.; and Liu, T.-Y. 2020. FastSpeech 2: Fast and High-Quality End-to-End Text to Speech.
- Sharma, B.; Gao, X.; Vijayan, K.; Tian, X.; and Li, H. 2021. NHSS: A speech and singing parallel database. *Speech Communication*, 133: 9–22.
- Sisman, B.; and Li, H. 2020. Generative Adversarial Networks for Singing Voice Conversion with and without Parallel Data. In *Odyssey*, 238–244.
- Sisman, B.; Vijayan, K.; Dong, M.; and Li, H. 2019. Singan: Singing voice conversion with generative adversarial networks. In *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (AP-SIPA ASC)*, 112–118. IEEE.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*.
- SVC-Develop-Team. 2023. SoftVC VITS Singing Voice Conversion. <https://github.com/svc-develop-team/so-vits-svc>. Accessed: 2024-10-12.
- Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Wang, H.; Zheng, S.; Chen, Y.; Cheng, L.; and Chen, Q. 2023. CAM++: A Fast and Efficient Network for Speaker Verification Using Context-Aware Masking. In *INTER-SPEECH 2023*, 5301–5305.
- Wang, X.; Takaki, S.; and Yamagishi, J. 2019. Neural source-filter-based waveform model for statistical parametric speech synthesis. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5916–5920. IEEE.
- Wang, Z.; Ma, L.; Feng, Y.; Pan, X.; Jin, Y.; and Zhang, K. 2024. SaMoye: Zero-shot Singing Voice Conversion Model Based on Feature Disentanglement and Enhancement. <https://github.com/CarlWangChina/SaMoye-SVC>. Accessed: 2024-10-12.
- Wei, H.; Cao, X.; Dan, T.; and Chen, Y. 2023. RMVPE: A Robust Model for Vocal Pitch Estimation in Polyphonic Music. *arXiv preprint arXiv:2306.15412*.
- Yxlllc. 2023. Real-time end-to-end singing voice conversion system based on DDSP. <https://github.com/yxlllc/DDSP-SVC>. Accessed: 2023-08-23.
- Zhang, L.; Li, R.; Wang, S.; et al. 2022. M4singer: A multi-style, multi-singer and musical score provided mandarin singing corpus. *Advances in Neural Information Processing Systems*, 35: 6914–6926.
- Zhang, X.; Xue, L.; Gu, Y.; Wang, Y.; Li, J.; He, H.; Wang, C.; Song, T.; Chen, X.; Fang, Z.; Chen, H.; Zhang, J.; Tang, T. Y.; Zou, L.; Wang, M.; Han, J.; Chen, K.; Li, H.; and Wu, Z. 2024. Amphion: An Open-Source Audio, Music and Speech Generation Toolkit. In *IEEE Spoken Language Technology Workshop, SLT 2024*.
- Zhao, W.; Bai, L.; Rao, Y.; Zhou, J.; and Lu, J. 2023. Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. *Advances in Neural Information Processing Systems*, 36: 49842–49869.
- Zhou, Y.; Chen, M.; Lei, Y.; Zhu, J.; and Zhao, W. 2023. VITS-based Singing Voice Conversion System with DSP-GAN post-processing for SVCC2023. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 1–8. IEEE.