

CoDiff-VC: A Codec-Assisted Diffusion Model for Zero-shot Voice Conversion

1stYuke Li

Northwestern Polytechnical University
Xi'an, China
liyik22@mail.nwpu.edu.cn

2ndXinfa Zhu

Northwestern Polytechnical University
Xi'an, China
xfzhu@mail.nwpu.edu.cn

3rdHanzhao Li

Northwestern Polytechnical University
Xi'an, China
lihanzhao.mail@gmail.com

4thJiXun Yao

Northwestern Polytechnical University
Xi'an, China
yaojx@mail.nwpu.edu.cn

5th WenJie Tian

Northwestern Polytechnical University
Xi'an, China
twj@mail.nwpu.edu.cn

6th YunLin Chen

Shanghai Mobvoi Information Technology
Shanghai, China
yunlinchen@mobvoi.com

7th XiPeng Yang

Shanghai Mobvoi Information Technology
Shanghai, China
xipeng.yang@mobvoi.com

8th Zhifei Li

Shanghai Mobvoi Information Technology
Shanghai, China
zfli@mobvoi.com

9th Lei Xie*

Northwestern Polytechnical University
Xi'an, China
lxie@nwpu.edu.cn

Abstract—Zero-shot voice conversion (VC) aims to convert the original speaker’s timbre to any target speaker while keeping the linguistic content. Current mainstream zero-shot voice conversion approaches depend on pre-trained recognition models to disentangle linguistic content and speaker representation. This results in a timbre residue within the decoupled linguistic content and inadequacies in speaker representation modeling. In this study, we propose CoDiff-VC, an end-to-end framework for zero-shot voice conversion that integrates a speech codec and a diffusion model to produce high-fidelity waveforms. Our approach involves employing a single-codebook codec to separate linguistic content from the source speech. To enhance content disentanglement, we introduce Mix-Style layer normalization (MSLN) to perturb the original timbre. Additionally, we incorporate a multi-scale speaker timbre modeling approach to ensure timbre consistency and improve voice detail similarity. To improve speech quality and speaker similarity, we introduce dual classifier-free guidance, providing both content and timbre guidance during the generation process. Objective and subjective experiments affirm that CoDiff-VC significantly improves speaker similarity, generating natural and higher-quality speech. Audio samples are available on the demo page¹.

Index Terms—zero-shot voice conversion, multi-scale timbre modeling, diffusion model.

I. INTRODUCTION

Zero-shot voice conversion (VC) aims to transfer the timbre of a source speaker to the timbre of an unseen target speaker while maintaining the original linguistic content [1]. This approach requires only one utterance from the target speaker, making it applicable in various scenarios like movie dubbing [2], speech translation [3], and speech anonymization [4]. In real-world applications, zero-shot voice conversion can enhance personalized voice interactions [5], [6] and improve entertainment experiences by converting speaker timbre characteristics. However, since only one target speaker utterance is available, disentangling speech components and simultaneously converting to the target speaker’s timbre becomes more challenging.

In the zero-shot VC task, there are two primary challenges to address: firstly, how to disentangle the linguistic content and speaker timbre from the source speech; secondly, how to model the speaker representation precisely. To solve the first challenge, many previous approaches [7]–[10] utilize pre-trained automatic speech recognition (ASR) or self-supervised learning (SSL) models [11], [12] to extract

bottleneck features as linguistic content decoupled from the source speech. Simultaneously, a speaker verification model is employed to extract the speaker representation. The VC model then combines the original linguistic content with the target speaker representation to reconstruct the converted speech. Despite the previous approach achieving some success in zero-shot VC [13], [14], the extracted bottleneck features still contain speaker-related information, leading to poor speaker similarity in the converted speech. Meanwhile, most of the mentioned approaches depend on an acoustic model for predicting the mel-spectrogram-like latent representations and employ a vocoder to reconstruct the representations into speech waveform [10], [15], [16]. This two-stage framework introduces cascading errors, thereby degrading the quality of the converted speech.

For the second challenge, the previous studies on speaker representation modeling can be broadly divided into two categories: coarse-grained modeling approach [17]–[19] and fine-grained modeling approach [20]–[22]. In coarse-grained modeling, a pre-trained speaker verification model is utilized to extract a global speaker embedding as the coarse-grained speaker representation. While effective in controlling the overall timbre characteristics of the converted speech, these approaches fall short in capturing detailed speaker timbre information and semantic-related timbre changes, leading to limited speaker similarity between the converted speech and target speaker speech. Conversely, other studies employ an attention mechanism to capture fine-grained speaker representation from multiple reference utterances, with Mega-TTS2 [21] being the most prominent example, which allows for the generation of more natural speech for the target speaker. However, the speaker similarity in converted speech using these approaches relies on the duration of the reference, resulting in a notable degradation in similarity performance when the reference speech duration is excessively short.

In this study, with particular consideration of the above two challenges, we propose CoDiff-VC, a codec-assisted end-to-end framework for zero-shot voice conversion, which can generate high-fidelity waveforms without any auxiliary losses and avoid cascading error issues. We employ a pre-trained codec model, featuring only a single codebook, to extract discrete tokens from the source speech as the linguistic content. The single codebook architecture can partially disentangle the speaker’s timbre by introducing a speaker reference encoder while retaining accurate linguistic content information. Mean-

* Corresponding authors.

¹<https://aries457.github.io/CoDiff-VC/>

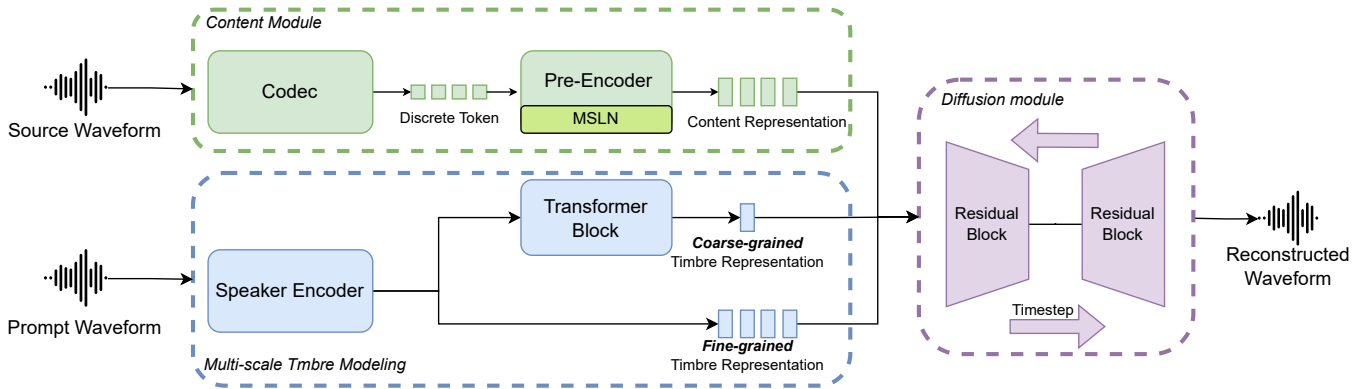


Fig. 1. Overall architecture of the CoDiff-VC. The content module within the green dashed line is used to extract linguistic content from the source speech, while we employ a multi-scale timbre modeling module to capture the details of speaker timbre. The diffusion module within the purple dashed line reconstructs the speech waveform conditioned on the linguistic content and speaker timbre.

while, we incorporate Mix-Style layer normalization (MSLN) [23] to perturb the timbre information within the discrete tokens, facilitating a more thorough disentanglement of timbre characteristics from the tokens. To improve timbre similarity and consistency, we introduce a multi-scale speaker timbre modeling approach to recover voice details when reconstructing the waveform. Finally, we propose a dual classifier-free guidance strategy to train unconditional models of content and timbre for guiding the reverse process to generate high-quality waveforms. Objective and subjective experiments demonstrate that CoDiff-VC outperforms the baseline systems in both speech quality and speaker similarity. Ablation studies further demonstrate the effectiveness of each component in our proposed approach.

II. PROPOSED APPROACH

The overall architecture of our proposed system is shown in Figure 1, which consists of three parts: content module, multi-scale timbre modeling, and diffusion module. For the content module, we utilize a pre-trained codec model with a reference encoder to encode audio into discrete tokens and employ mix-style layer normalization further to reduce the residual timbre in the discrete tokens. To achieve more comprehensive timbre modeling, we introduce multi-scale speaker timbre modeling, encompassing modeling of both coarse-grained and fine-grained timbre representation. Following this, a denoising diffusion probabilistic model serves as the backbone of the conversion model for an end-to-end reconstruction process. To enhance the quality of the reconstructed speech, we propose a dual classifier-free guidance strategy to effectively guide content and timbre modeling, thereby improving speech quality and speaker similarity.

A. Linguistic Content Extraction

To disentangle the linguistic content from the source speech, we utilize a pre-trained codec to extract discrete tokens as linguistic content. The architecture closely resembles the codec in [24]. Since the majority of information is preserved in the first codebook and lower bit rates make it easier to decouple timbres [25], we employ a single-codebook speech codec, comprising only 4096 tokens, to reduce bit rates. And, to disentangle timbre characteristics within the tokens in the codec, we incorporate a reference encoder into the codec. The output of the reference encoder, representing global representation related to acoustic information, is subtracted from the hidden representation of the encoder. This subtraction process imposes implicit constraints on the tokens, leading them to spontaneously disentangle global representations, such as timbre and acoustic features, at lower bit rates.

We use an embedding layer to transform the discrete token into a high-level representation. The token embedding is upsampled to match the target audio’s shape, serving as the content condition for the diffusion process. To eliminate the speaker timbre, we incorporate MSLN into the pre-encoder to perturb the original timbre information within the tokens by introducing mismatched timbre information, preventing the generation of timbre-consistent representations.

B. Multi-scale Speaker Modeling

Many previous studies [26], [27] assume that speaker timbre is a time-invariant global representation, neglecting the potential timbre variations due to intense emotions or specific semantics. In scenarios of zero-shot voice conversion, relying solely on global representation to model speaker timbre may result in a degradation of timbre similarity in the converted speech. To address this limitation and leverage the powerful reconstruction capability of diffusion, we introduce multi-scale speaker timbre modeling within the diffusion model. This approach enables the simultaneous capture of coarse-grained and fine-grained timbre information from reference speech, thereby enhancing the speaker similarity in the converted speech sample. As shown in Figure 1, the speaker timbre modeling module comprises a speaker encoder and a transformer module. Initially, we randomly select an utterance and clip it to a fixed-length segment as the reference speech. The speaker encoder processes this reference to produce frame-level speaker representations, which serve as fine-grained timbre representations. The transformer module then takes these frame-level representations and generates coarse-grained timbre representations through global pooling. To integrate timbre characteristics at different levels, we use distinct approaches within the diffusion model. As depicted in Figure 2, the coarse-grained timbre information is concatenated with the timestep embedding and then fed into the residual block. This setup enables the supervision of diffusion by timbre at each timestep when predicting noise. For fine-grained timbre modeling, we introduce a cross-attention convolution module every 4 U-net blocks, associating linguistic content information with timbre. Specifically, the latent representation undergoes 256-time downsampling to match the same dimension as the mel spectrogram to be employed as the query in the cross-attention module, and the fine-grained timbre serves as the keys and values. Ultimately, the upsampling module restores it to the same shape as the audio.

C. Dual Classifier-free Guidance

To guide the reverse diffusion process, we propose a novel dual classifier-free guidance strategy. We employ classifier-free guid-

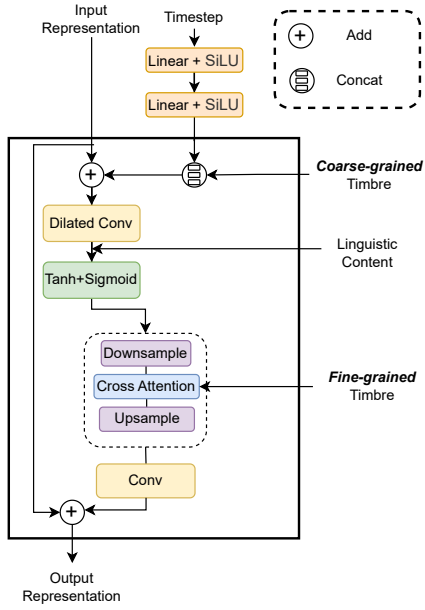


Fig. 2. The structure of U-net block.

ance [28] to content representation R_c and speaker timbre R_s , respectively. By introducing two distinctive implicit classifiers, both content information and timbre information can guide noise generation in unconditional scenarios. During training, we discard the content condition and timbre condition at a rate of 15%. The inference process is shown in Eq.(1), where the $\hat{\epsilon}_\theta$ denotes the denoising network which reverses original sample x_0 from Gaussian noise x_t . Meanwhile, w_c and w_s represent the guidance scales of text conditions and timbre conditions for noise prediction, respectively.

$$\hat{\epsilon}_\theta(x_t, t, R_c, R_s) = (1 + w_c + w_s) \cdot \epsilon(x_t, t, R_c, R_s) - w_c \cdot \epsilon(x_t, t, R_s) - w_s \cdot \epsilon(x_t, t, R_c) \quad (1)$$

During inference, we need to assign a guidance scale to unconditional estimation, which can determine the impact of real conditional estimation and unconditional estimation on the synthesis process. With each timestep of inference, the coarse-grained information is first modeled before modeling fine-grained information. Therefore, we use an annealing strategy for content representation, where the scale w_c gradually decreases over timesteps. Meanwhile, we set the guidance scale w_s following a hardening strategy, which gradually increases over timesteps, indicating that fine-grained speaker information is gradually restored to improve speaker similarity.

D. Training and Inference

We first train the speech codec and then extract quantized discrete tokens for conversion model training. We select audio clips of the same speaker as the reference speech to get a multi-scale timbre representation. Then, with the content representation and timbre representation as condition information, the proposed CoDiff-VC is optimized by minimizing the following unweighted variant of the ELBO [29] as shown in Eq.(3), which has proved to be effective for high-quality generation [24].

$$\bar{\alpha}_t = \prod_{t=1}^T (1 - \beta_t) \quad (2)$$

$$\min_{\theta} L(\theta) = \mathbb{E}_{x_0, \epsilon, t} \|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t, \text{cond})\|_2^2 \quad (3)$$

where the $\epsilon \sim N(0, 1)$, β_t is linear noise schedule and *cond* means condition information. For the reverse process, it can also be defined as a Markov chain from the noisy data x_t to the original data x_0 as shown in Eq.(4), where μ_θ and Σ_θ are the mean and variance terms respectively. Finally, in order to ensure the speed and quality of inference, we adopt fast sampling [30] and $T_{\text{infer}} = 10$.

$$p_\theta(x_{t-1} | x_t) = N(x_{t-1}; \mu_\theta(x_t, t, \text{cond}), \Sigma_\theta(x_t, t)) \quad (4)$$

III. EXPERIMENTS

A. Datasets and Model Setup

We conduct our experiments using the LibriTTS dataset [33], which comprises 580 hours of speech from 2400 speakers. To standardize the data, we crop the audio samples to the same length and add silence at the end of shorter samples. The dataset is divided into a training set with speech data from 2390 speakers and a test set with data from the remaining 10 speakers. All audio samples are resampled to 24000 Hz.

In our implementation, we employ a codec based on the Tortoise approach [24]. We trained the codec model on 4 NVIDIA V100 GPUs with a batch size of 1024 for 300k training steps. For the conversion model, CoDiff-VC adopts a U-net structure with 30 blocks, each having a dimension of 128. It utilizes a linear noise schedule for $\beta_t \in [1 \times 10^{-4}, 0.02]$ with $T = 200$ in the diffusion model, following the approach in DiffWave [34]. CoDiff-VC is trained on 4 NVIDIA V100 GPUs in the training setup with a batch size of 16 for 1M steps, using the Adam optimizer with an initial learning rate of $1e-4$.

B. Baseline Systems

We adopt YourTTS [18], SEF-VC [31], LM-VC [8] and Diff-VC [32] as the baseline systems for comparison with our proposed CoDiff-VC in the zero-shot voice conversion scenario:

- **YourTTS**: A TTS model with a speaker voiceprint model to extract speaker embedding and employ HiFi-GAN vocoder to reconstruct the waveform. The speaker embedding served as a condition of the flow-based decoder and posterior encoder for enhancing the zero-shot multi-speaker generation capability.
- **SEF-VC**: A VC framework with a semantic encoder and a mel encoder. The semantic encoder reconstructs the discrete representations extracted by Hubert into mel spectrograms as the backbone model. The mel encoder processes the mel spectrogram to obtain speaker timbre representation via a powerful position-agnostic cross-attention mechanism.
- **LM-VC**: A two-stage language modeling approach that generates acoustic tokens to recover linguistic content and timbre and reconstructs the fine for acoustic details as converted speech.
- **Diff-VC**: A diffusion-based VC model that employs the average voice encoder to transform mel features into a speaker-average representation in the dataset for the disentanglement problem.

C. Subjective Evaluation

We use Mean Opinion Score (MOS) as the subjective metric to evaluate performance between the baseline and proposed systems regarding speaker similarity, naturalness, and speech quality. The test set comprises 20 utterances involving 10 unseen speakers (5 females and 5 males). Using 10 speakers as references, each of the 20 sentences serves as input, generating 200 audio samples. Ten listeners are invited to participate in the subjective evaluation, scoring the similarity between the converted and reference speech. Additionally, Naturalness MOS and Speech Quality MOS assessments are conducted to measure the intelligibility of the synthesized voice.

TABLE I

RESULTS OF CoDIFF-VC COMPARED WITH BASELINE SYSTEMS AND ABLATION MODELS REGARDING SPEAKER SIMILARITY MOS, NATURALNESS MOS AND SPEECH QUALITY MOS WITH CONFIDENCE INTERVALS OF 95%. SCOS AND WER MEAN SPEAKER COSINE SIMILARITY AND WORD ERROR RATE.

Model	Speaker Similarity (MOS \uparrow)	Naturalness (MOS \uparrow)	Speech Quality (MOS \uparrow)	SCOS \uparrow	WER \downarrow
Ground Truth	-	4.39 \pm 0.09	4.53 \pm 0.12	-	2.176
YourTTS [18]	3.43 \pm 0.10	3.37 \pm 0.09	3.61 \pm 0.11	0.574	4.850
SEF-VC [31]	3.67 \pm 0.12	3.77 \pm 0.10	3.76 \pm 0.10	0.782	5.436
LM-VC [8]	3.70 \pm 0.09	3.83 \pm 0.11	4.03 \pm 0.11	0.806	5.150
Diff-VC [32]	3.72 \pm 0.09	3.65 \pm 0.08	4.10 \pm 0.08	0.815	4.326
CoDiff-VC	3.94\pm0.10	4.07\pm0.08	4.26\pm0.09	0.839	3.121
-MSLN	3.79 \pm 0.10	3.95 \pm 0.11	4.22 \pm 0.10	0.830	3.170
-Fine-grained	3.73 \pm 0.10	3.88 \pm 0.10	4.20 \pm 0.08	0.824	3.240
-Dual CFG	3.85 \pm 0.08	4.01 \pm 0.11	4.10 \pm 0.09	0.831	3.344

As shown in Table I, the subjective evaluation results demonstrate the superior performance of our proposed approach in zero-shot voice conversion. YourTTS solely relies on global timbre representation to model speaker timbre, while SEF-VC utilizes only local timbre representation and the self-supervised model extracting semantic information exhibits residual timbre information, leading to poor timbre similarity. In addition, Diff-VC exhibits lower naturalness MOS than CoDiff-VC, suggesting that our content module extracts better content representations. Notably, these three models are constrained by a two-stage method, first reconstructing the mel and then restoring it through the vocoder, resulting in degradation of generation quality. Limited to coarse acoustic tokens, LM-VC faces challenges in restoring unseen speaker timbres. In contrast, CoDiff-VC modeling both coarse and fine-grained timbre, coupled with the application of dual CFG on the diffusion model, enhances the quality of converted speech while effectively capturing timbre information.

D. Objective Evaluation

To evaluate speaker similarity, we calculate the cosine similarity between the converted speech and the reference speech using a pre-trained ECAPA-TDNN [35]. For intelligibility evaluation, we calculate the Word Error Rate (WER) for the converted speech. The WER assessment is computed using an open-source Wenet toolkit [36] trained on the LibriSpeech dataset.

As shown in Table I, the objective evaluation results align with the subjective evaluation, showing that CoDiff-VC achieved the highest cosine similarity. SEF-VC, LM-VC, and Diff-VC recorded higher WER scores than CoDiff-VC, suggesting that the discrete representation of content through timbre decoupling in Co-DiffVC facilitates more accurate modeling of content information. Meanwhile, the lower WER results indicate that our proposed CoDiff-VC’s converted speech outperforms all baseline intelligibility systems.

To demonstrate the effectiveness of timbre modeling, we visualize the coarse-grained timbre representations through t-SNE [37]. We select 100 audio samples from five random speakers in the test set and extract coarse-grained timbre representations for clustering using the model’s coarse-grained timbre modeling structure. The clustering results, shown in Figure 3, indicate that the timbre representations of each speaker are strongly related to the corresponding speakers.

E. Ablation Study

To evaluate the effectiveness of each component in CoDiff-VC, we conduct an ablation study to analyze the impact of three key components: MSLN, multi-scale timbre modeling, and dual classifier-free guidance. We compare CoDiff-VC with three variants: first, without the MSLN module (-MSLN); second, using only coarse-grained speaker timbre instead of multi-scale ones (-fine-grained);




Fig. 3. T-SNE visualization of the coarse-grained timbre in multi-speakers settings.

and third, without dual classifier-free guidance (-dual cfg), utilizing only classifier-free guidance for content representation.

As illustrated in Table I, the speaker similarity and speech naturalness are significantly degraded when removing the MSLN module. These results suggest that the MSLN module is crucial for retaining content information while eliminating residual timbre information, playing an important role in disentangling linguistic content. In terms of timbre modeling, relying solely on single-scale speaker representation results in a notable decline in speaker similarity, indicating the critical role of the multi-scale timbre modeling module in capturing the details of speaker timbre. Additionally, naturalness experiences degradation with only coarse-grained timbre modeling, highlighting the importance of timbral details in synthesizing natural speech. Furthermore, the model without dual classifier-free guidance exhibits lower speaker similarity and speech quality, emphasizing the contribution of additional regular timbre information guidance to improved timbre modeling and speech quality.

IV. CONCLUSIONS

In this paper, we propose CoDiff-VC, a codec-assisted end-to-end framework designed for zero-shot voice conversion. We employ a pre-trained codec and a pre-encoder to decouple discrete content representations from the source speech. Meanwhile, we introduce a multi-scale speaker timbre modeling module within the diffusion model for modeling fine-grained timbre details. Our proposed CoDiff-VC excels in generating converted audio for unseen target speakers with more similar voice characteristics, higher naturalness, and improved speech quality. Additionally, we present a dual classifier-free guidance method, creating implicit classifiers for content and speaker representations to enhance the inference results of the diffusion model. However, it’s worth noting that the inference speed of diffusion remains relatively slow. We aim to address and improve the reverse process of diffusion in our future work.

REFERENCES

- [1] Berrak Sisman, Junichi Yamagishi, Simon King, and Haizhou Li, “An overview of voice conversion and its challenges: From statistical modeling to deep learning,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 29, pp. 132–157, 2020.
- [2] Jixun Yao, Yi Lei, Qing Wang, Pengcheng Guo, Ziqian Ning, Lei Xie, Hai Li, Junhui Liu, and Danming Xie, “Preserving background sound in noise-robust voice conversion via multi-task learning,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [3] David Sündermann, Hermann Ney, and Harald Höge, “Vtln-based cross-language voice conversion,” *Proc. ASRU*, pp. 676–681, 2003.
- [4] In-Chul Yoo, Keonnyeong Lee, Seong-Gyun Leem, Hyunwoo Oh, Bonggu Ko, and Dongsuk Yook, “Speaker anonymization for personal information protection using voice conversion techniques,” *IEEE Access*, vol. 8, pp. 198637–198645, 2020.
- [5] Yi-Chin Huang, Chung-Hsien Wu, and Yu-Ting Chao, “Personalized spectral and prosody conversion using frame-based codeword distribution and adaptive crf,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 21, pp. 51–62, 2013.
- [6] Berrak Sisman, Haizhou Li, and Kay Chen Tan, “Transformation of prosody in voice conversion,” *Proc. APSIPA ASC*, pp. 1537–1546, 2017.
- [7] Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark A. Hasegawa-Johnson, “Autovc: Zero-shot voice style transfer with only autoencoder loss,” *Proc. ICML*, pp. 5210–5219, 2019.
- [8] Zhichao Wang, Yuan-Jui Chen, Linfu Xie, Qiao Tian, and Yuping Wang, “LM-VC: Zero-shot voice conversion via speech generation based on language models,” *IEEE Sig. Proc. Lett.*, pp. 1157–1161, 2023.
- [9] Guanlong Zhao, Sinem Sonsaat, John M. Levis, Evgeny Chukharev-Hudilainen, and Ricardo Gutierrez-Osuna, “Accent conversion using phonetic posteriorgrams,” *Proc. ICASSP*, pp. 5314–5318, 2018.
- [10] Songxiang Liu, Yuewen Cao, Dan Su, and Helen Meng, “Diffsvc: A diffusion probabilistic model for singing voice conversion,” in *Proc. ASRU*. IEEE, 2021, pp. 741–748.
- [11] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, “Wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Proc. NIPS*, vol. 33, pp. 12449–12460, 2020.
- [12] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdel rahman Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
- [13] Ruitong Xiao, Xiaofen Xing, Jichen Yang, and Xiangmin Xu, “CA-VC: A novel zero-shot voice conversion method with channel attention,” *Proc. APSIPA ASC*, pp. 800–807, 2021.
- [14] Haozhe Zhang, Zexin Cai, Xiaoyi Qin, and Ming Li, “SIG-VC: A speaker information guided zero-shot voice conversion system for both human beings and machines,” *Proc. ICASSP*, pp. 6567–65571, 2021.
- [15] Trung Dang, Dung T. Tran, Peter Chin, and Kazuhito Koishida, “Training robust zero-shot voice conversion models with self-supervised features,” *Proc. ICASSP*, pp. 6557–6561, 2022.
- [16] Kai Shen, Zeqian Ju, Xu Tan, Eric Liu, Yichong Leng, Lei He, Tao Qin, Sheng Zhao, and Jiang Bian, “Naturalspeech 2: Latent diffusion models are natural and zero-shot speech and singing synthesizers,” in *Proc. ICLR*, 2024.
- [17] Rama Sanand Doddipatla, Norbert Braunschweiler, and Ranniery Maia, “Speaker adaptation in dnn-based speech synthesis using d-vectors,” in *Proc. INTERSPEECH*, 2017, pp. 3404–3408.
- [18] Edresson Casanova, Julian Weber, Christopher Dane Shulby, Arnaldo Cândido Júnior, Eren Gölge, and Moacir Antonelli Ponti, “YourTTS: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone,” in *Proc. ICML*, 2021.
- [19] Tomi H. Kinnunen, Lauri Juvela, Paavo Alku, and Junichi Yamagishi, “Non-parallel voice conversion using i-vector plda: towards unifying speaker verification and transformation,” *Proc. ICASSP*, pp. 5535–5539, 2017.
- [20] Yixuan Zhou, Changhe Song, Xiang Li, Luwen Zhang, Zhiyong Wu, Yanyao Bian, Dan Su, and Helen Meng, “Content-dependent fine-grained speaker embedding for zero-shot speaker adaptation in text-to-speech synthesis,” in *Proc. INTERSPEECH*, 2022, pp. 2573–2577.
- [21] Ziyue Jiang, Jinglin Liu, Yi Ren, Jinzheng He, Chen Zhang, Zhe Ye, Pengfei Wei, Chunfeng Wang, Xiang Yin, Zejun Ma, and Zhou Zhao, “Mega-TTS 2: Zero-shot text-to-speech with arbitrary length speech prompts,” *ArXiv*, vol. abs/2307.07218, 2023.
- [22] Ji-Hyun Lee, Sang-Hoon Lee, Ji-Hoon Kim, and Seong-Whan Lee, “PVAE-TTS: Adaptive text-to-speech via progressive style adaptation,” *Proc. ICASSP*, pp. 6312–6316, 2022.
- [23] Rongjie Huang, Yi Ren, Jinglin Liu, Chenye Cui, and Zhou Zhao, “Generspeech: Towards style transfer for generalizable out-of-domain text-to-speech,” *Proc. NIPS*, vol. 35, pp. 10970–10983, 2022.
- [24] James Betker, “Better speech synthesis through scaling,” *ArXiv*, vol. abs/2305.07243, 2023.
- [25] Yong Ren, Tao Wang, Jiangyan Yi, Le Xu, Jianhua Tao, Chuyuan Zhang, and Jun Zhou, “Fewer-token neural speech codec with time-invariant codes,” *ArXiv*, vol. abs/2310.00014, 2023.
- [26] Minki Kang, Dong Min, and Sung Ju Hwang, “Grad-stylespeech: Any-speaker adaptive text-to-speech synthesis with diffusion models,” *Proc. ICASSP*, pp. 1–5, 2023.
- [27] Dongchan Min, Dong Bok Lee, Eunho Yang, and Sung Ju Hwang, “Meta-stylespeech: Multi-speaker adaptive text-to-speech generation,” in *Proc. ICML*, 2021, pp. 7748–7759.
- [28] Jonathan Ho, “Classifier-free diffusion guidance,” *ArXiv*, vol. abs/2207.12598, 2022.
- [29] Jonathan Ho, Ajay Jain, and Pieter Abbeel, “Denosing diffusion probabilistic models,” *Proc. NIPS*, vol. 33, pp. 6840–6851, 2020.
- [30] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J. Weiss, Mohammad Norouzi, and William Chan, “WaveGrad: Estimating gradients for waveform generation,” in *Proc. ICLR*, 2021.
- [31] Junjie Li, Yiwei Guo, Xie Chen, and Kai Yu, “SEF-VC: Speaker embedding free zero-shot voice conversion with cross attention,” *ArXiv*, vol. abs/2312.08676, 2023.
- [32] Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, Mikhail Sergeevich Kudinov, and Jiansheng Wei, “Diffusion-based voice conversion with fast maximum likelihood sampling scheme,” in *Proc. ICLR*, 2022.
- [33] Heiga Zen, Viet-Trung Dang, Robert A. J. Clark, Yu Zhang, Ron J. Weiss, Ye Jia, Z. Chen, and Yonghui Wu, “Libritts: A corpus derived from librispeech for text-to-speech,” in *Proc. INTERSPEECH*, 2019, pp. 1526–1530.
- [34] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro, “DiffWave: A versatile diffusion model for audio synthesis,” *Proc. ICLR*, pp. 562–571, 2020.
- [35] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck, “Ecapattnn: Emphasized channel attention, propagation and aggregation in tdn based speaker verification,” in *Proc. INTERSPEECH*, 2020, pp. 3830–3834.
- [36] Zhuoyuan Yao, Di Wu, Xiong Wang, Binbin Zhang, Fan Yu, Chao Yang, Zhendong Peng, Xiaoyu Chen, Lei Xie, and Xin Lei, “WeNet: Production oriented streaming and non-streaming end-to-end speech recognition toolkit,” in *Proc. INTERSPEECH*, 2021, pp. 4054–4058.
- [37] Laurens van der Maaten and Geoffrey E. Hinton, “Visualizing data using t-sne,” 2008, vol. 9, pp. 2579–2605.