MASKCYCLEGAN-VC: LEARNING NON-PARALLEL VOICE CONVERSION WITH FILLING IN FRAMES

Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, Nobukatsu Hojo

NTT Communication Science Laboratories, NTT Corporation, Japan

ABSTRACT

Non-parallel voice conversion (VC) is a technique for training voice converters without a parallel corpus. Cycleconsistent adversarial network-based VCs (CycleGAN-VC and CycleGAN-VC2) are widely accepted as benchmark methods. However, owing to their insufficient ability to grasp time-frequency structures, their application is limited to mel-cepstrum conversion and not mel-spectrogram conversion despite recent advances in mel-spectrogram vocoders. To overcome this, CycleGAN-VC3, an improved variant of CycleGAN-VC2 that incorporates an additional module called time-frequency adaptive normalization (TFAN), has been proposed. However, an increase in the number of learned parameters is imposed. As an alternative, we propose MaskCycleGAN-VC, which is another extension of CycleGAN-VC2 and is trained using a novel auxiliary task called filling in frames (FIF). With FIF, we apply a temporal mask to the input mel-spectrogram and encourage the converter to fill in missing frames based on surrounding frames. This task allows the converter to learn time-frequency structures in a self-supervised manner and eliminates the need for an additional module such as TFAN. A subjective evaluation of the naturalness and speaker similarity showed that MaskCycleGAN-VC outperformed both CycleGAN-VC2 and CycleGAN-VC3 with a model size similar to that of CycleGAN-VC2.1

Index Terms— Voice conversion (VC), non-parallel VC, generative adversarial networks (GANs), CycleGAN-VC, mel-spectrogram conversion

1. INTRODUCTION

Voice conversion (VC) is a technique for translating one voice into another without changing the linguistic content, and has been extensively studied owing to its various applications, including speaking assistance [1, 2], speech enhancement [3, 4], and accent conversion [5, 6]. Machine-learning-based approaches have been widely used, ranging from statistical modeling (e.g., Gaussian mixture models [7, 8]) to neural networks (NNs) (e.g., feedforward NNs [9], recurrent NNs [10], convolutional NNs [6], and attention networks [11, 12, 13]).

Many VC methods (including those above) are categorized into parallel VC approaches and train a converter between the source and target speakers using parallel utterances. Parallel VC has the advantage that it can train a converter in a supervised manner; however, it requires a parallel corpus, which is not always easy to collect.

As an alternative, non-parallel VC, a technique for training a converter without a parallel corpus, has attracted attention, and many such methods have thus been proposed. Among them, a promising approach is to utilize linguistic information to compensate for the missing parallel supervision [14, 15, 16, 17]; however, extra data or pretrained models are needed to derive such linguistic information.

To remove such a requirement and solve non-parallel VC without any additional data or pretrained models, deep generative models, such as generative adversarial networks (GANs) [18] and variational autoencoders (VAEs) [19], have been introduced [20, 21, 22, 23, 24]. Among them, the family of CycleGAN-VCs (CycleGAN-VC [22, 25], CycleGAN-VC2 [26], and StarGAN-VCs [27, 28, 29]) are significant achievements and have been widely accepted as benchmark approaches (e.g., [17, 30, 31]). However, owing to their insufficient capacity to capture the time-frequency structure (e.g., the harmonic structure is compromised, as shown in Figure 1 in [32]), their application is limited to mel-cepstrum conversion and not mel-spectrogram conversion despite recent advances in mel-spectrogram vocoders [33, 34, 35, 36, 37].

To overcome this, CycleGAN-VC3 [32], an improved variant of CycleGAN-VC2, was recently proposed, and addresses the problem by incorporating an additional module called time-frequency adaptive normalization (TFAN). Although the performance is superior, an increase in the number of converter parameters is necessary (from 16M to 27M).

As an alternative, we propose *MaskCycleGAN-VC*, which is another extension of CycleGAN-VC2 and is trained using a novel auxiliary task called *filling in frames (FIF)*. With FIF, we apply a temporal mask to the input mel-spectrogram and encourage the converter to fill in the missing frames based on the surrounding frames. FIF is inspired by the success of complementation-based self-supervised learning in other fields, e.g., image inpainting in computer vision [38] and text infilling in natural language processing [39, 40]. Similarly, FIF allows the converter to learn the time-frequency feature structure in a self-supervised manner through a complementation process. This strong property eliminates the need for an additional module such as TFAN, and makes CycleGAN-VC2 applicable to mel-spectrogram conversion with negligibly small network modifications.

We investigated the effectiveness of MaskCycleGAN-VC on the Spoke (i.e., non-parallel VC) task of the Voice Conversion Challenge 2018 (VCC 2018) [41]. A subjective eval-

¹Audio samples are available at http://www.kecl.ntt.co.jp/ people/kaneko.takuhiro/projects/maskcyclegan-vc/ index.html.

uation of the naturalness and speaker similarity showed that MaskCycleGAN-VC outperformed both CycleGAN-VC2 and CycleGAN-VC3 while keeping the model size similar to that of CycleGAN-VC2.

The rest of this paper is organized as follows. In Section 2, we review CycleGAN-VC2, which is the baseline of our model. We then introduce the proposed MaskCycleGAN-VC in Section 3. In Section 4, we describe the experimental results. Finally, we provide some concluding remarks and areas of future study in Section 5.

2. CONVENTIONAL CYCLEGAN-VC2

The purpose of CycleGAN-VC2 is to train a converter $G_{X \to Y}$ that translates source acoustic features $x \in X$ into target acoustic features $y \in Y$ without parallel supervision. Following CycleGAN [42, 43, 44], which was proposed for unpaired image-to-image translation, CycleGAN-VC2 solves this problem using an *adversarial loss* [18], *cycleconsistency loss* [45], and *identity-mapping loss* [46]. In addition, CycleGAN-VC2 uses a *second adversarial loss* [26] to improve the quality of the cyclically reconstructed features. Adversarial loss. An adversarial loss $\mathcal{L}_{adv}^{X \to Y}$ is used to make the converted feature $G_{X \to Y}(x)$ appear to be the target:

$$\mathcal{L}_{adv}^{X \to Y} = \mathbb{E}_{\boldsymbol{y} \sim P_Y}[\log D_Y(\boldsymbol{y})] \\ + \mathbb{E}_{\boldsymbol{x} \sim P_X}[\log(1 - D_Y(G_{X \to Y}(\boldsymbol{x})))], \quad (1)$$

where the discriminator D_Y distinguishes a real y from the generated $G_{X \to Y}(x)$ by maximizing this loss, whereas $G_{X \to Y}$ generates $G_{X \to Y}(x)$, which can deceive D_Y by minimizing this loss. Similarly, the inverse converter $G_{Y \to X}$ is trained with the discriminator D_X using $\mathcal{L}_{adv}^{Y \to X}$.

Cycle-consistency loss. A cycle-consistency loss $\mathcal{L}_{cyc}^{X \to Y \to X}$ is used to determine the pseudo pair within the cycle-consistency constraint without parallel supervision:

$$\mathcal{L}_{cyc}^{X \to Y \to X} = \mathbb{E}_{\boldsymbol{x} \sim P_X}[\|G_{Y \to X}(G_{X \to Y}(\boldsymbol{x})) - \boldsymbol{x}\|_1].$$
(2)

Similarly, $\mathcal{L}_{cyc}^{Y \to X \to Y}$ is used for the inverse-forward mapping (i.e., $G_{X \to Y}(G_{Y \to X}(\boldsymbol{y}))$).

Identity-mapping loss. An identity-mapping loss $\mathcal{L}_{id}^{X \to Y}$ is used to enhance the input preservation:

$$\mathcal{L}_{id}^{X \to Y} = \mathbb{E}_{\boldsymbol{y} \sim P_Y}[\|G_{X \to Y}(\boldsymbol{y}) - \boldsymbol{y}\|_1].$$
(3)

Similarly, $\mathcal{L}_{id}^{Y \to X}$ is used for the inverse converter $G_{Y \to X}$.

Second adversarial loss. A second adversarial loss $\mathcal{L}_{adv2}^{X \to Y \to X}$ is used to mitigate the statistical averaging caused by L1 loss in Eq. 2.

$$\mathcal{L}_{adv2}^{X \to Y \to X} = \mathbb{E}_{\boldsymbol{x} \sim P_X} [\log D'_X(\boldsymbol{x})] \\ + \mathbb{E}_{\boldsymbol{x} \sim P_X} [\log(1 - D'_X(G_{Y \to X}(G_{X \to Y}(\boldsymbol{x}))))], \quad (4)$$

where the discriminator D'_X distinguishes a reconstructed $G_{Y \to X}(G_{X \to Y}(\boldsymbol{x}))$ from a real \boldsymbol{x} . Similarly, $\mathcal{L}_{adv2}^{Y \to X \to Y}$ is used for the inverse-forward mapping with an additional discriminator D'_Y .



Fig. 1. Pipeline of FIF for the forward-inverse mapping. We encourage the converter to fill in the missing frames (surrounded by the red box) based on the surrounding frames through a cyclic conversion process. In practice, a similar procedure is used for the inverse-forward mapping.

Full objective. A full objective \mathcal{L}_{full} is written as follows:

$$\mathcal{L}_{full} = \mathcal{L}_{adv}^{X \to Y} + \mathcal{L}_{adv}^{Y \to X} + \lambda_{cyc} (\mathcal{L}_{cyc}^{X \to Y \to X} + \mathcal{L}_{cyc}^{Y \to X \to Y}) + \lambda_{id} (\mathcal{L}_{id}^{X \to Y} + \mathcal{L}_{id}^{Y \to X}) + \mathcal{L}_{adv2}^{X \to Y \to X} + \mathcal{L}_{adv2}^{Y \to X \to Y},$$
(5)

where λ_{cyc} and λ_{id} are weighing parameters. $G_{X \to Y}$ and $G_{Y \to X}$ are optimized by minimizing this loss, whereas D_X , D_Y , D'_X , and D'_Y are optimized by maximizing this loss.

3. MASKCYCLEGAN-VC

3.1. Training with Filling in Frames (FIF)

As shown in [32], CycleGAN-VC2, which was developed for mel-cepstrum conversion, does not have sufficient ability to capture the time-frequency structure in mel-spectrogram conversion; consequently, the harmonic structure is often compromised. To alleviate this, we devised *MaskCycleGAN-VC*, which is trained using the auxiliary *FIF* task. We present the overall pipeline of FIF in Fig. 1.

Given the source mel-spectrogram x, we first create a temporal mask $m \in M$, which has the same size as x, parts of which have a value of zero (denoted by the black region in Fig. 1), and the remaining parts have a value of 1 (indicated by the white region in Fig. 1). A masked region (i.e., zero region) is randomly determined based on a predetermined rule (the effect of which is examined in Section 4.2).

Subsequently, we apply the mask m to x as follows:

$$\hat{\boldsymbol{x}} = \boldsymbol{x} \cdot \boldsymbol{m},$$
 (6)

where represents an element-wise product. By using this procedure, we artificially create missing frames, as shown in the region surrounded by the red box in Fig. 1.

Next, the MaskCycleGAN-VC converter $G_{X \to Y}^{mask}$ synthesizes y' from \hat{x} and m as follows:

$$\boldsymbol{y}' = G_{X \to Y}^{mask}(\text{concat}(\hat{\boldsymbol{x}}, \boldsymbol{m})), \tag{7}$$

where concat denotes the channel-wise concatenation. By using m as the conditional information, $G_{X \to Y}^{mask}$ can fill in the frames while knowing which frames need to be filled in.

Similar to CycleGAN-VC2, we can ensure that y' is in the target Y by using an adversarial loss (Eq. 1) but cannot compare y' with the ground truth directly owing to the lack of parallel supervision. As an alternative, we aim to fill in the frames through a cyclic conversion process. To do so, we reconstruct x'' using the inverse converter $G_{Y\to X}^{mask}$:

$$\boldsymbol{x}^{\prime\prime} = G_{Y \to X}^{mask}(\text{concat}(\boldsymbol{y}^{\prime}, \boldsymbol{m}^{\prime})), \qquad (8)$$

where m' is represented using an all-ones matrix under the assumption that the missing frames have been filled in ahead of this process. We then apply the cycle-consistency loss for the original and reconstructed mel-spectrograms:

$$\mathcal{L}_{mcyc}^{X \to Y \to X} = \mathbb{E}_{\boldsymbol{x} \sim P_X, \boldsymbol{m} \sim P_M}[\|\boldsymbol{x}'' - \boldsymbol{x}\|_1], \qquad (9)$$

where we simultaneously used a second adversarial loss (Eq. 4) for x''.

To optimize $\mathcal{L}_{mcyc}^{X \to Y \to X}$, $G_{X \to Y}^{mask}$ needs to derive information useful for filling in the missing frames from the surrounding frames. This induction is useful for learning the timefrequency structure in a mel-spectrogram in a self-supervised manner. Note that similar effects have been observed for similar tasks in other fields (e.g., image inpainting [38] and text infilling [39, 40]), as mentioned in Section 1. Finally, it should be noted that (1) unlike CycleGAN-VC3, which uses TFAN, MaskCycleGAN-VC does not need a large increase in the converter parameters (only the input channels are doubled to receive m along with \hat{x}), and (2) FIF is a type of self-supervised learning; therefore, neither extra data nor a pretrained model (e.g., linguistic information) is required.

3.2. Conversion with all-ones mask

As a remaining question, what mask should be used during the conversion process (i.e., test phase)? For this question, we simply use an all-ones mask. Thus, we can convert speech under the assumption that no missing frames exist. This assumption is the same as that used in typical VC.

4. EXPERIMENTS

4.1. Experimental conditions

Dataset. We examined the effectiveness of MaskCycleGAN-VC on the Spoke (i.e., non-parallel VC) task of VCC 2018 [41], which contains recordings of native speakers of American English. We used a subset of speakers that covers all interand intra-gender VC, i.e., VCC2SF3 (SF), VCC2SM3 (SM), VCC2TF1 (TF), and VCC2TM1 (TM), where S, T, F, and M indicate the sources, targets, females, and males, respectively. We used combinations of 2 sources \times 2 targets for the evaluation. For each speaker, we used 81 sentences for training (of approximately 5 min in length, which is relatively short for VC) and 35 sentences for the evaluation. Note that the training set contains no overlapping utterances between the source and target speakers; therefore, we need to train a converter in a fully non-parallel setting. In this dataset, audio clips were down-sampled to 22.05 kHz. Similar to a study on CycleGAN-VC3 [32], we extracted an 80-dimensional log mel-spectrogram with a window length of 1024 and hop length of 256 samples.

Conversion and synthesis process. For a fair comparison with CycleGAN-VC3 [32], we used the same conversion

and synthesis process as CycleGAN-VC3. Namely, we applied MaskCycleGAN-VC to mel-spectrogram conversion and synthesized the waveform using the pretrained MelGAN vocoder [35].² Although for a fair comparison we did not change the parameters of the vocoder, fine-tuning it for each speaker is acceptable.

Network architectures. We used similar network architectures as in CycleGAN-VC2 for mel-spectrogram conversion, which was used as the baseline in the study on CycleGAN-VC3 [32] (see Figure 4 in [26] and Section 4.1 in [32] for the details). The converter consists of a 2-1-2D CNN [26], and the discriminator is PatchGAN [47]. As mentioned in Section 3.1, the only difference between CycleGAN-VC2 and MaskCycleGAN-VC is that the input channels are doubled in the converter to receive m along with \hat{x} .

Training settings. We used the same training settings as in CycleGAN-VC3 [32]. During the preprocessing, we normalized the mel-spectrograms using the training set statistics. We used a least-squares GAN [48] as the GAN objective. We trained the networks for 500k iterations using an Adam optimizer [49], with the learning rates of the converter and discriminator set to 0.0002 and 0.0001, respectively, under momentum terms β_1 and β_2 of 0.5 and 0.999, respectively. The batch size was set to 1, where each training sample consisted of 64 randomly cropped frames (approximately 0.75 s in length). λ_{cyc} and λ_{id} were set to 10 and 5, respectively, and \mathcal{L}_{id} was used for only the first 10k iterations to prevent \mathcal{L}_{id} from disturbing the learning of conversion. Similar to the previous CycleGAN-VCs, we did not use extra data, pretrained models, or a time alignment procedure for training.

4.2. Objective evaluation

We conducted an objective evaluation to examine the differences in performance when using different components. Because a direct comparison between the converted and target mel-spectrograms is difficult owing to the lack of a correct alignment, we used two metrics: (1) *mel-cepstral distortion (MCD)*, which is the most commonly applied measure and calculates the distance within the mel-cepstral domain (particularly, a 35-dimensional mel-cepstrum was extracted from the converted or targeted waveform using the WORLD analyzer [50]), and (2) *Kernel DeepSpeech Distance (KDSD)* [51], which computes the maximum mean discrepancy within the DeepSpeech2 feature space [52] and is shown to be well correlated with human judgement [51]. For both metrics, the smaller the value, the better the performance.

Comparison among different-sized masks. We first examined the effect of the mask size selection. Here, the mask size indicates the size of the zero region (i.e., the black region in Fig. 1). We tested two variations. (1) *FIF X*: The mask size is constantly X% (i.e., $64 \times \frac{X}{100}$ frames). Here, *FIF 0* means that an all-ones mask is used. (2) *FIF 0-X*: The mask size is randomly determined within the range of [0, X%]. We list the results in Table 1(a). We found that (i) FIF with a non-zero-sized mask (Nos. 2–5) outperformed that with a zero-sized mask (No. 1) regardless of the mask size, (ii) the performance

²https://github.com/descriptinc/melgan-neurips

Table 1. Comparison of MCD and KDSD using (a) differentsized masks, (b) different types of masks, and (c) different CycleGAN-VCs. The results are listed as MCD [dB]/KDSD [$\times 10^5$]. Bold numbers indicate the best scores. #param represents the number of converter parameters.

No.	(a) Size	SF-TF	SM-TM	SF-TM	SM-TF	#param
1	FIF 0 FIF 25	7.66/786	7.11/356	6.91/277	8.11/1094	16M
3	FIF 0-25	7.45/489	6.83/103	6.78/206	7.80/605	16M
4 5	FIF 0-50 FIF 0-75	7.37/467 7.40/468	6.77/ 83.8 6.75/89.2	6.73/ 146 6.72/169	7.64/502 7.66/546	16M 16M
No.	(b) Type	SF-TF	SM-TM	SF-TM	SM-TF	#param
6	FIF	7.37/467	6.77/83.8	6.73/146	7.64/502	16M
8	FIS	7.52/727	6.95/437	6.88/418 7.09/774	7.94/974	16M 16M
	111	1.051720	0.777477	1.07114	0.24/2120	10141
No.	(c) Model	SF-TF	SM-TM	SF-TM	SM-TF	#param
10 11 12	Mask V2 V3	7.37 /467 7.66/891 7.54/ 369	6.77/83.8 7.07/509 7.10/227	6.73/146 6.96/494 6.91/311	7.64/502 8.07/1107 7.97/819	16M 16M 27M

is affected by the mask size (Nos. 3–5) and maximizes at approximately X = 50, and (iii) FIF with a random-sized mask (No. 4) outperformed FIF with a constant-sized mask (No. 2) despite the same average size. The possible reason is that, during training, the former includes an all-ones mask, which is used in the test phase, whereas the latter does not.

Comparison among different types of masks. We inspect the effect of the mask type selection. We compared four variations. (1) *FIF*: Subsequent frames are masked, as shown in Fig. 1. (2) *FIF_{NS}*: Non-subsequent frames (i.e., each frame is independently and randomly selected) are masked. (3) *FIS*: Subsequent spectrum bands (e.g., 45th–60th melspectrograms) are masked. (4) *FIP*: Mel-spectrogram was masked in a point-wise manner similar to a dropout [53]. Under all settings, we used a mask size of 0-50, which was the best setting in the previous experiment. We summarize the results in Table 1(b). We found that *FIF* (No. 6) outperformed the others (Nos. 7–9) for all speaker pairs. We consider that, although learning the temporal structure is the most difficult, it is important for CycleGAN-VC2, and *FIF* is the most effective in mitigating this difficulty.

Comparison among CycleGAN-VCs. We examined the differences in performance among (1) MaskCycleGAN-VC (*Mask*, particularly *FIF 0-50*, was used); (2) CycleGAN-VC2 [26] (*V2*), which was the same as *Mask* except FIF was not used; and (3) CycleGAN-VC3 [32] (*V3*), which applied TFAN instead of FIF. The results are listed in Table 1(c). We found that *Mask* (No. 10) outperformed both *V2* (No. 11) and *V3* (No. 12) in most cases, reducing the model size compared to *V3*. Further evidence is provided in the next section.

4.3. Subjective evaluation

We conducted listening tests to investigate the differences in perceptual quality. As the benchmark performance of melspectrogram conversion based on CycleGAN-VCs was previously examined in [32], we investigated the comparative performance between *Mask* and *V2* and that between *Mask* and *V3* using two forced-choice preference tests. In the AB test



Fig. 2. Average preference scores on naturalness with 95% confidence intervals. The numbers in parentheses indicate the p-values computed using a one-tailed binomial test.



Fig. 3. Average preference scores on speaker similarity with 95% confidence intervals. The numbers in parentheses denote the p-values calculated using a one-tailed binomial test.

on naturalness, each listener was presented with two speech samples (A and B) and asked to choose their preferred one (A or B) considering both naturalness and intelligibility. In the XAB test on speaker similarity, each listener was presented with three speech samples, including comparison targets (A and B) and a reference with a different utterance (X), and asked to choose their preferred one (A or B) with speaker characteristics closer to that of X. These tests were conducted online, and 15 and 16 listeners participated in the AB and XAB tests, respectively. Sentences, comparison targets, and the compared order (AB or BA) were randomly chosen from the collection of speech samples. We gathered at least 300 answers for each model pair. Audio samples are available from the link¹ presented in the first page.

We show the results of the AB test on naturalness and the XAB test on speaker similarity in Figs. 2 and 3, respectively. We found that in both tests, *Mask* achieved statistically significantly better scores than V2 and V3 with a p-value of < 0.05 (where a one-tailed binomial test was used).

5. CONCLUSIONS

Motivated by recent advances in mel-spectrogram vocoders, we proposed MaskCycleGAN-VC, which is an improvement of CycleGAN-VC2 for mel-spectrogram conversion. To learn the time-frequency structure in a mel-spectrogram without an additional module such as TFAN, we introduced FIF, which allows the converter to learn such a structure in a self-supervised manner. The experimental results showed that MaskCycleGAN-VC outperformed both CycleGAN-VC2 and CycleGAN-VC3 while maintaining a model size similar to that of CycleGAN-VC2. Examining the generality of FIF is an interesting research topic, and future work includes applications to multi-domain VC [27, 28, 29] and application-side VC [1, 2, 3, 4, 5, 6].

Acknowledgements: This work was supported by JSPS KAKENHI 17H01763 and JST CREST Grant Number JP-MJCR19A3, Japan.

6. REFERENCES

- Alexander B Kain, John-Paul Hosom, Xiaochuan Niu, Jan P. H. van Santen, Melanie Fried-Oken, and Janice Staehely, "Improving the intelligibility of dysarthric speech," Speech Commun., vol. 49, no. 9, pp. 743–759, 2007.
- [2] Keigo Nakamura, Tomoki Toda, Hiroshi Saruwatari, and Kiyohiro Shikano, "Speaking-aid systems using GMM-based voice conversion for electrolaryngeal speech," *Speech Commun.*, vol. 54, no. 1, pp. 134–146, 2012.
- [3] Zeynep Inanoglu and Steve Young, "Data-driven emotion conversion in spoken English," Speech Commun., vol. 51, no. 3, pp. 268–283, 2009.
- [4] Tomoki Toda, Mikihiro Nakagiri, and Kiyohiro Shikano, "Statistical voice conversion techniques for body-conducted unvoiced speech enhancement," *IEEE Trans. Audio Speech Lang. Process.*, vol. 20, no. 9, pp. 2505–2517, 2012.
- [5] Daniel Felps, Heather Bortfeld, and Ricardo Gutierrez-Osuna, "Foreign accent conversion in computer assisted pronunciation training," *Speech Commun.*, vol. 51, no. 10, pp. 920–932, 2009.
- [6] Takuhiro Kaneko, Hirokazu Kameoka, Kaoru Hiramatsu, and Kunio Kashino, "Sequence-to-sequence voice conversion with similarity metric learned using generative adversarial networks," in *Proc. Interspeech*, 2017, pp. 1283–1287.
- [7] Yannis Stylianou, Olivier Cappé, and Eric Moulines, "Continuous probabilistic transform for voice conversion," *IEEE Trans. Speech Audio Process.*, vol. 6, no. 2, pp. 131–142, 1998.
- [8] Tomoki Toda, Alan W Black, and Keiichi Tokuda, "Voice conversion based on maximum-likelihood estimation of spectral parameter trajectory," *IEEE Trans. Audio Speech Lang. Process.*, vol. 15, no. 8, pp. 2222–2235, 2007.
- [9] Srinivas Desai, Alan W Black, B Yegnanarayana, and Kishore Prahallad, "Spectral mapping using artificial neural networks for voice conversion," *IEEE Trans. Audio Speech Lang. Process.*, vol. 18, no. 5, pp. 954–964, 2010.
- [10] Lifa Sun, Shiyin Kang, Kun Li, and Helen Meng, "Voice conversion using deep bidirectional long short-term memory based recurrent neural networks," in *Proc. ICASSP*, 2015, pp. 4869–4873.
- [11] Jing-Xuan Zhang, Zhen-Hua Ling, Li-Juan Liu, Yuan Jiang, and Li-Rong Dai, "Sequence-to-sequence acoustic modeling for voice conversion," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 3, pp. 631–644, 2019.
- [12] Kou Tanaka, Hirokazu Kameoka, Takuhiro Kaneko, and Nobukatsu Hojo, "AttS2S-VC: Sequence-to-sequence voice conversion with attention and context preservation mechanisms," in *Proc. ICASSP*, 2019, pp. 6805–6809.
- [13] Hirokazu Kameoka, Kou Tanaka, Takuhiro Kaneko, and Nobukatsu Hojo, "ConvS2S-VC: Fully convolutional sequence-to-sequence voice conversion," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 1849–1863, 2020.
- [14] Lifa Sun, Kun Li, Hao Wang, Shiyin Kang, and Helen Meng, "Phonetic posteriorgrams for many-to-one voice conversion without parallel data training," in *Proc. ICME*, 2016, pp. 1–6.
- [15] Feng-Long Xie, Frank K Soong, and Haifeng Li, "A KL divergence and DNNbased approach to voice conversion without parallel training sentences," in *Proc. Interspeech*, 2016, pp. 287–291.
- [16] Yuki Saito, Yusuke Ijima, Kyosuke Nishida, and Shinnosuke Takamichi, "Nonparallel voice conversion using variational autoencoders conditioned by phonetic posteriorgrams and d-vectors," in *Proc. ICASSP*, 2018, pp. 5274–5278.
- [17] Jingxuan Zhang, Zhenhua Ling, and Li-Rong Dai, "Non-parallel sequence-tosequence voice conversion with disentangled linguistic and speaker representations," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 540–552, 2020.
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial nets," in *Proc. NIPS*, 2014, pp. 2672–2680.
- [19] Diederik P Kingma and Max Welling, "Auto-encoding variational Bayes," in Proc. ICLR, 2014.
- [20] Chin-Cheng Hsu, Hsin-Te Hwang, Yi-Chiao Wu, Yu Tsao, and Hsin-Min Wang, "Voice conversion from non-parallel corpora using variational auto-encoder," in *Proc. APSIPA ASC*, 2016, pp. 1–6.
- [21] Chin-Cheng Hsu, Hsin-Te Hwang, Yi-Chiao Wu, Yu Tsao, and Hsin-Min Wang, "Voice conversion from unaligned corpora using variational autoencoding Wasserstein generative adversarial networks," in *Proc. Interspeech*, 2017, pp. 3364–3368.
- [22] Takuhiro Kaneko and Hirokazu Kameoka, "Parallel-data-free voice conversion using cycle-consistent adversarial networks," arXiv preprint arXiv:1711.11293, Nov. 2017.
- [23] Hirokazu Kameoka, Takuhiro Kaneko, Kou Tanaka, and Nobukatsu Hojo, "ACVAE-VC: Non-parallel voice conversion with auxiliary classifier variational autoencoder," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 27, no. 9, pp. 1432–1443, 2019.
- [24] Patrick Lumban Tobing, Yi-Chiao Wu, Tomoki Hayashi, Kazuhiro Kobayashi, and Tomoki Toda, "Non-parallel voice conversion with cyclic variational autoencoder," in *Proc. Interspeech*, 2019, pp. 674–678.
- [25] Takuhiro Kaneko and Hirokazu Kameoka, "CycleGAN-VC: Non-parallel voice conversion using cycle-consistent adversarial networks," in *Proc. EUSIPCO*, 2018, pp. 2114–2118.
- [26] Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, and Nobukatsu Hojo, "CycleGAN-VC2: Improved CycleGAN-based non-parallel voice conversion," in *Proc. ICASSP*, 2019, pp. 6820–6824.

- [27] Hirokazu Kameoka, Takuhiro Kaneko, Kou Tanaka, and Nobukatsu Hojo, "StarGAN-VC: Non-parallel many-to-many voice conversion using star generative adversarial networks," in *Proc. SLT*, 2018, pp. 266–273.
- [28] Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, and Nobukatsu Hojo, "StarGAN-VC2: Rethinking conditional methods for StarGAN-based voice conversion," in *Proc. Interspeech*, 2019, pp. 679–683.
- [29] Hirokazu Kameoka, Takuhiro Kaneko, Kou Tanaka, and Nobukatsu Hojo, "Nonparallel voice conversion with augmented classifier star generative adversarial networks," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 2982–2995, 2020.
- [30] Shindong Lee, BongGu Ko, Keonnyeong Lee, In-Chul Yoo, and Dongsuk Yook, "Many-to-many voice conversion using conditional cycle-consistent adversarial networks," in *Proc. ICASSP*, 2020, pp. 6279–6283.
- [31] Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark Hasegawa-Johnson, "Auto-VC: Zero-shot voice style transfer with only autoencoder loss," in *Proc. ICML*, 2019, pp. 5210–5219.
- [32] Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, and Nobukatsu Hojo, "CycleGAN-VC3: Examining and improving CycleGAN-VCs for melspectrogram conversion," in *Proc. Interspeech*, 2020, pp. 2017–2021.
- [33] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al., "Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions," in *Proc. ICASSP*, 2018, pp. 4779–4783.
- [34] Ryan Prenger, Rafael Valle, and Bryan Catanzaro, "WaveGlow: A flow-based generative network for speech synthesis," in *Proc. ICASSP*, 2019, pp. 3617–3621.
- [35] Kundan Kumar, Rithesh Kumar, Thibault de Boissiere, Lucas Gestin, Wei Zhen Teoh, Jose Sotelo, Alexandre de Brébisson, Yoshua Bengio, and Aaron Courville, "MelGAN: Generative adversarial networks for conditional waveform synthesis," in *Proc. NeurIPS*, 2019, pp. 14881–14892.
- [36] Ryuichi Yamamoto, Eunwoo Song, and Jae-Min Kim, "Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *Proc. ICASSP*, 2020, pp. 6199–6203.
- [37] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan, "WaveGrad: Estimating gradients for waveform generation," arXiv preprint arXiv:2009.00713, Sept. 2020.
- [38] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros, "Context Encoders: Feature learning by inpainting," in *Proc. CVPR*, 2016, pp. 2536–2544.
- [39] William Fedus, Ian Goodfellow, and Andrew M Dai, "MaskGAN: Better text generation via filling in the __," in *Proc. ICLR*, 2018.
- [40] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL*, 2019, pp. 4171–4186.
- [41] Jaime Lorenzo-Trueba, Junichi Yamagishi, Tomoki Toda, Daisuke Saito, Fernando Villavicencio, Tomi Kinnunen, and Zhenhua Ling, "The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods," in *Proc. Odyssey*, 2018, pp. 195–202.
- [42] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros, "Unpaired imageto-image translation using cycle-consistent adversarial networks," in *Proc. ICCV*, 2017, pp. 2223–2232.
- [43] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong, "DualGAN: Unsupervised dual learning for image-to-image translation," in *Proc. ICCV*, 2017, pp. 2849–2857.
- [44] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim, "Learning to discover cross-domain relations with generative adversarial networks," in *Proc. ICML*, 2017, pp. 1857–1865.
- [45] Tinghui Zhou, Philipp Krähenbühl, Mathieu Aubry, Qixing Huang, and Alexei A Efros, "Learning dense correspondence via 3D-guided cycle consistency," in *Proc. CVPR*, 2016, pp. 117–126.
- [46] Yaniv Taigman, Adam Polyak, and Lior Wolf, "Unsupervised cross-domain image generation," in Proc. ICLR, 2017.
- [47] Chuan Li and Michael Wand, "Precomputed real-time texture synthesis with Markovian generative adversarial networks," in *Proc. ECCV*, 2016, pp. 702–716.
- [48] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley, "Least squares generative adversarial networks," in *Proc. ICCV*, 2017, pp. 2794–2802.
- [49] Diederik Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," in Proc. ICLR, 2015.
- [50] Masanori Morise, Fumiya Yokomori, and Kenji Ozawa, "WORLD: A vocoderbased high-quality speech synthesis system for real-time applications," *IEICE Trans. Inf. Syst.*, vol. 99, no. 7, pp. 1877–1884, 2016.
- [51] Mikołaj Bińkowski, Jeff Donahue, Sander Dieleman, Aidan Clark, Erich Elsen, Norman Casagrande, Luis C Cobo, and Karen Simonyan, "High fidelity speech synthesis with adversarial networks," in *Proc. ICLR*, 2020.
- [52] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al., "Deep Speech 2: End-to-end speech recognition in English and Mandarin," in *Proc. ICML*, 2016, pp. 173–182.
- [53] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," J. Mach. Learn. Res., vol. 15, no. 1, pp. 1929–1958, 2014.