

# CycleGAN-VC3: Examining and Improving CycleGAN-VCs for Mel-spectrogram Conversion

Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, Nobukatsu Hojo

NTT Communication Science Laboratories, NTT Corporation, Japan

takuhiro.kaneko.tb@hco.ntt.co.jp

## Abstract

Non-parallel voice conversion (VC) is a technique for learning mappings between source and target speeches without using a parallel corpus. Recently, cycle-consistent adversarial network (CycleGAN)-VC and CycleGAN-VC2 have shown promising results regarding this problem and have been widely used as benchmark methods. However, owing to the ambiguity of the effectiveness of CycleGAN-VC/VC2 for mel-spectrogram conversion, they are typically used for mel-cepstrum conversion even when comparative methods employ mel-spectrogram as a conversion target. To address this, we examined the applicability of CycleGAN-VC/VC2 to mel-spectrogram conversion. Through initial experiments, we discovered that their direct applications compromised the time-frequency structure that should be preserved during conversion. To remedy this, we propose CycleGAN-VC3, an improvement of CycleGAN-VC2 that incorporates time-frequency adaptive normalization (TFAN). Using TFAN, we can adjust the scale and bias of the converted features while reflecting the time-frequency structure of the source mel-spectrogram. We evaluated CycleGAN-VC3 on inter-gender and intra-gender non-parallel VC. A subjective evaluation of naturalness and similarity showed that for every VC pair, CycleGAN-VC3 outperforms or is competitive with the two types of CycleGAN-VC2, one of which was applied to mel-cepstrum and the other to mel-spectrogram.<sup>1</sup>

**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 converting non/paralinguistic information in speech while retaining the linguistic information. VC has been actively studied owing to its potential diverse applications such as in speaking aids [1, 2], speech enhancement [3, 4], and accent conversion [5]. Machine-learning approaches are widely used and include statistical methods based on Gaussian mixture models [6, 7] and neural networks (NNs), including feedforward NNs [8], recurrent NNs [9], convolutional NNs (CNNs) [5], attention networks [10, 11], and generative adversarial networks (GANs) [5].

For ease of learning, many VC methods (including the above-mentioned methods) learn mappings from source to target speeches using a parallel corpus. However, collection of such data is not necessarily easy or practical. Additionally, even if such data are collected, the time alignment procedure required in most VC methods remains a challenge.

As an alternative, non-parallel VC, which does not require a parallel corpus for training, has gained attention recently. Non-

<sup>1</sup>Audio samples are available at <http://www.kecl.ntt.co.jp/people/kaneko.takuhiro/projects/cyclegan-vc3/index.html>.

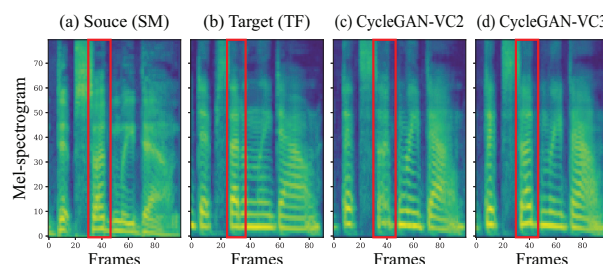


Figure 1: Comparison of source, target, and converted mel-spectrograms. In the red box, CycleGAN-VC2 (c) compromises the harmonic structure that appears in the source (a) and target (b), whereas CycleGAN-VC3 (d) succeeds in retaining it. Note that achieving this is not trivial because parallel data, such as (a) and (b), were not available for training.

parallel VC is attractive in terms of data collection cost; however, its learning is challenging owing to the absence of explicit supervision. To address this, several studies have utilized linguistic information [12, 13, 14, 15]. Although this extra supervision improves the performance, auxiliary data or modules are required to extract linguistic information.

To avoid such a requirement and achieve non-parallel VC using only acoustic data, variational autoencoder-based methods [16, 17, 18] and GAN-based methods [17, 19] have been proposed. Among them, CycleGAN-VC [19, 20] has garnered attention alongside its variants (CycleGAN-VC2 [21] and StarGAN-VCs [22, 23]), and they have been widely used as benchmark methods in several studies (e.g., [15, 24, 25]). However, owing to the ambiguity of their effectiveness for mel-spectrogram conversion, they are typically used for mel-cepstrum conversion even when comparative methods employ mel-spectrogram as a conversion target (e.g., [15, 25]).

These facts motivated us to examine the applicability of CycleGAN-VC [19] and CycleGAN-VC2 [21] to mel-spectrogram conversion. Through initial experiments, we found that when CycleGAN-VC/VC2 is directly applied to a mel-spectrogram, it compromises the time-frequency structure that should be preserved during conversion, as shown in Figure 1. To address this, we propose CycleGAN-VC3 that is an improvement of CycleGAN-VC2 that incorporates *time-frequency adaptive normalization (TFAN)*. TFAN is inspired by spatially adaptive (de)normalization (SPADE) [26] that was originally proposed for semantic image synthesis. We revise SPADE for application to 1D and 2D time-frequency features. Using TFAN, we can adjust the scale and bias of the converted features while reflecting the time-frequency structure of the source mel-spectrogram.

We examined the effectiveness of the CycleGAN-VC3 on inter-gender and intra-gender non-parallel VC using the Voice Conversion Challenge 2018 (VCC 2018) dataset [27]. A subjective evaluation of naturalness and similarity showed that for

every VC pair, CycleGAN-VC3 achieved a better or competitive performance compared with the two types of CycleGAN-VC2, one of which was applied to mel-cepstrum and the other to mel-spectrogram.

The remainder of this paper is organized as follows. In Section 2, we briefly review conventional CycleGAN-VC/VC2. In Section 3, we introduce the proposed CycleGAN-VC3. In Section 4, we describe the experimental results. Section 5 presents a concise summary and description of future work.

## 2. Conventional CycleGAN-VC/VC2

### 2.1. Training objectives

CycleGAN-VC/VC2 aims at learning a mapping  $G_{X \rightarrow Y}$  that converts source acoustic features  $\mathbf{x} \in X$  into target acoustic features  $\mathbf{y} \in Y$  without using a parallel corpus. Inspired by CycleGAN [28], originally proposed for unpaired image-to-image translation, CycleGAN-VC/VC2 learns a mapping using an *adversarial loss* [29], *cycle-consistency loss* [30], and *identity-mapping loss* [31]. Furthermore, CycleGAN-VC2 uses a *second adversarial loss* to improve the details of the reconstructed features.

**Adversarial loss.** To ensure that the converted feature  $G_{X \rightarrow Y}(\mathbf{x})$  is in the target  $Y$ , an adversarial loss  $\mathcal{L}_{adv}^{X \rightarrow Y}$  is used as follows:

$$\begin{aligned} \mathcal{L}_{adv}^{X \rightarrow Y} = & \mathbb{E}_{\mathbf{y} \sim P_Y} [\log D_Y(\mathbf{y})] \\ & + \mathbb{E}_{\mathbf{x} \sim P_X} [\log(1 - D_Y(G_{X \rightarrow Y}(\mathbf{x})))], \end{aligned} \quad (1)$$

where the discriminator  $D_Y$  attempts to distinguish synthesized  $G_{X \rightarrow Y}(\mathbf{x})$  from real  $\mathbf{y}$  by maximizing the loss, whereas  $G_{X \rightarrow Y}$  attempts to synthesize  $G_{X \rightarrow Y}(\mathbf{x})$  that can deceive  $D_Y$  by minimizing the loss. Similarly, the inverse mapping  $G_{Y \rightarrow X}$  and the discriminator  $D_X$  are trained adversarially using  $\mathcal{L}_{adv}^{Y \rightarrow X}$ .

**Cycle-consistency loss.** To preserve the composition in conversion, a cycle-consistency loss  $\mathcal{L}_{cyc}$  is used as follows:

$$\begin{aligned} \mathcal{L}_{cyc} = & \mathbb{E}_{\mathbf{x} \sim P_X} [\|G_{Y \rightarrow X}(G_{X \rightarrow Y}(\mathbf{x})) - \mathbf{x}\|_1] \\ & + \mathbb{E}_{\mathbf{y} \sim P_Y} [\|G_{X \rightarrow Y}(G_{Y \rightarrow X}(\mathbf{y})) - \mathbf{y}\|_1]. \end{aligned} \quad (2)$$

This loss is used with a hyper-parameter  $\lambda_{cyc}$ , which controls its relative importance. The loss aids  $G_{X \rightarrow Y}$  and  $G_{Y \rightarrow X}$  to identify pseudo pairs within the cycle-consistency constraint.

**Identity-mapping loss.** To facilitate input preservation, an identity-mapping loss  $\mathcal{L}_{id}$  is used as follows:

$$\begin{aligned} \mathcal{L}_{id} = & \mathbb{E}_{\mathbf{y} \sim P_Y} [\|G_{X \rightarrow Y}(\mathbf{y}) - \mathbf{y}\|_1] \\ & + \mathbb{E}_{\mathbf{x} \sim P_X} [\|G_{Y \rightarrow X}(\mathbf{x}) - \mathbf{x}\|_1]. \end{aligned} \quad (3)$$

This loss is used with a hyper-parameter  $\lambda_{id}$ , which controls its relative importance.

**Second adversarial loss.** In CycleGAN-VC2, to mitigate the statistical averaging caused by the L1 loss (Equation 2), an additional discriminator  $D'_X$  is introduced, and a second adversarial loss  $\mathcal{L}_{adv2}^{X \rightarrow Y \rightarrow X}$  is imposed on the circularly converted features as follows:

$$\begin{aligned} \mathcal{L}_{adv2}^{X \rightarrow Y \rightarrow X} = & \mathbb{E}_{\mathbf{x} \sim P_X} [\log D'_X(\mathbf{x})] \\ & + \mathbb{E}_{\mathbf{x} \sim P_X} [\log(1 - D'_X(G_{Y \rightarrow X}(G_{X \rightarrow Y}(\mathbf{x})))]. \end{aligned} \quad (4)$$

Similarly, the discriminator  $D'_Y$  is introduced, and  $\mathcal{L}_{adv2}^{Y \rightarrow X \rightarrow Y}$  is imposed on the inverse-forward mapping.

### 2.2. Generator architectures

CycleGAN-VC uses *1D CNN generators* [5] to capture the overall relationship along with the feature direction while preserving the temporal structure. In particular, the network is composed of downsampling, residual [32], and upsampling blocks to capture wide-range temporal relationships effectively, and gated linear units (GLUs) [33] are used as activation to learn sequential and hierarchical structure adaptively.

However, a study on CycleGAN-VC2 [21] showed that 1D CNNs in downsampling and upsampling blocks affect the structure that should be retained in conversion. To alleviate this, CycleGAN-VC2 introduces *2-1-2D CNN* that uses 2D CNNs in upsampling and downsampling blocks and uses 1D CNNs in residual blocks. The former is used for extracting the time-frequency structure while preserving the original structure. The latter is used for performing dynamic changes.

### 2.3. Discriminator architectures

CycleGAN-VC uses *2D CNN discriminators* [34] to discriminate data based on 2D spectral textures. In particular, it uses *FullGAN*, which has a fully connected layer as the last layer, to discriminate data based on the overall input structure. However, in FullGAN, the need to learn many parameters causes learning difficulty. To mitigate this, CycleGAN-VC2 introduces *PatchGAN* [35] that uses convolution at the last layer. This reduces the parameters and stabilizes the GAN training.

## 3. CycleGAN-VC3

### 3.1. TFAN: Time-frequency adaptive normalization

CycleGAN-VC and CycleGAN-VC2 were originally designed for mel-cepstrum conversion, and their effectiveness in mel-spectrogram conversion is not sufficiently examined. We empirically examined their effectiveness and discovered that they compromise the time-frequency structure that should be preserved in conversion, as shown in Figure 1.

Motivated by this finding, we devised TFAN that extends instance normalization (IN) [36] to adjust the scale and bias of the converted features while reflecting the source information (i.e.,  $\mathbf{x}$ ) in a time- and frequency-wise manner. In particular, we designed TFAN for 1D and 2D time-frequency features to be used in 2-1-2D CNN (Section 2.2). Figure 2 illustrates the architectures of TFAN. Given feature  $\mathbf{f}$ , TFAN normalizes it in a *channel-wise manner* similar to IN and then modulates the normalized feature in an *element-wise manner* using scale  $\gamma(\mathbf{x})$  and bias  $\beta(\mathbf{x})$ , which are calculated from  $\mathbf{x}$  using CNNs:

$$\mathbf{f}' = \gamma(\mathbf{x}) \frac{\mathbf{f} - \boldsymbol{\mu}(\mathbf{f})}{\boldsymbol{\sigma}(\mathbf{f})} + \beta(\mathbf{x}), \quad (5)$$

where  $\mathbf{f}'$  is the output feature, and  $\boldsymbol{\mu}(\mathbf{f})$  and  $\boldsymbol{\sigma}(\mathbf{f})$  are the channel-wise average and standard deviation of  $\mathbf{f}$ , respectively. In IN, *x-independent* scale  $\beta$  and bias  $\gamma$  are applied in a *channel-wise manner*, whereas in TFAN, those *calculated from x* (i.e.,  $\beta(\mathbf{x})$  and  $\gamma(\mathbf{x})$ ) are applied in an *element-wise manner*. This difference allows TFAN to adjust the scale and bias of  $\mathbf{f}$  while reflecting  $\mathbf{x}$  in a time- and frequency-wise manner.

Note that TFAN is inspired by SPADE [26], which was originally proposed for semantic image synthesis. The main differences are that (1) SPADE was devised for 2D image features, whereas TFAN was designed for both 1D and 2D time-frequency features, (2) SPADE uses a one-layer CNN, as shown in the dashed-line box in Figure 2, because the drastic changes

## 4. Experiments

### 4.1. Experimental conditions

**Dataset.** We evaluated CycleGAN-VCs on the Spoke (i.e., non-parallel VC) task of VCC 2018 [27], which contains recordings of professional US English speakers. We selected a subset of speakers that considered all inter-gender and intra-gender VC: VCC2SF3 (*SF*), VCC2SM3 (*SM*), VCC2TF1 (*TF*), and VCC2TM1 (*TM*), where *S*, *T*, *F*, and *M* represent source, target, female, and male, respectively. Combinations of 2 sources  $\times$  2 targets were used for evaluation. For each speaker, 81 utterances (approximately 5 min, which is relatively low for VC) and 35 utterances were used for training and evaluation, respectively. In the training set, there were no overlaps between the source and target utterances; therefore, this problem must be solved in a fully non-parallel setting. The recordings were downsampled to 22.05 kHz. Following the study of MelGAN [38], which we used as a vocoder in our experiments, we extracted an 80-dimensional log mel-spectrogram with a window length of 1024 and hop length of 256 samples.

**Conversion process.** One aim in this study was to examine the feasibility of using CycleGAN-VCs for mel-spectrogram conversion. Hence, we used CycleGAN-VCs for mel-spectrogram conversion and synthesized waveforms using the pretrained MelGAN vocoder [38].<sup>2</sup> We did not alter the parameters of the vocoder such that we could focus on the evaluation of mel-spectrogram conversion; however, fine-tuning them for each speaker is a possible means for improvement.

**Network architectures.** As the acoustic feature is changed from mel-cepstrum to mel-spectrogram, the feature dimension increased from 35 to 80. However, the generators of CycleGAN-VCs are fully convolutional; therefore, they can be used without modifying the network architecture. Regarding the discriminators, we used the same network architecture as those for mel-cepstrum conversion, except that in CycleGAN-VC2/VC3, the kernel size in the second-last convolutional layer was doubled in the frequency direction (see Figure 4 in [21] for details of the original network architectures).

**Training settings.** The training settings were similar to those used in CycleGAN-VC/VC2 for mel-cepstrum conversion [19, 21]. For preprocessing, we normalized the mel-spectrograms using the mean and variance of the training data. We used the least square GAN [39] as the GAN objective. We trained the networks for 500k iterations using the Adam optimizer [40] with a batch size of 1. A training sample consisted of randomly cropped 64 frames (approximately 0.75 s). The learning rates were set to 0.0002 for the generators and 0.0001 for the discriminators with momentum terms  $\beta_1$  and  $\beta_2$  of 0.5 and 0.999, respectively.  $\lambda_{cyc}$  and  $\lambda_{id}$  were set to 10 and 5, respectively, and  $\mathcal{L}_{id}$  was used only for the first 10k iterations. Note that similar to the original CycleGAN-VC/VC2, we did not use extra data, modules, or time alignment procedures for training.

### 4.2. Objective evaluation

We conducted an objective evaluation to investigate the effect of TFAN parameter selection and the performance difference among CycleGAN-VCs. Direct measurement of the difference between the target and converted mel-spectrograms is difficult because their alignment is not trivial. As an alternative, we used two evaluation metrics that are commonly used in pre-

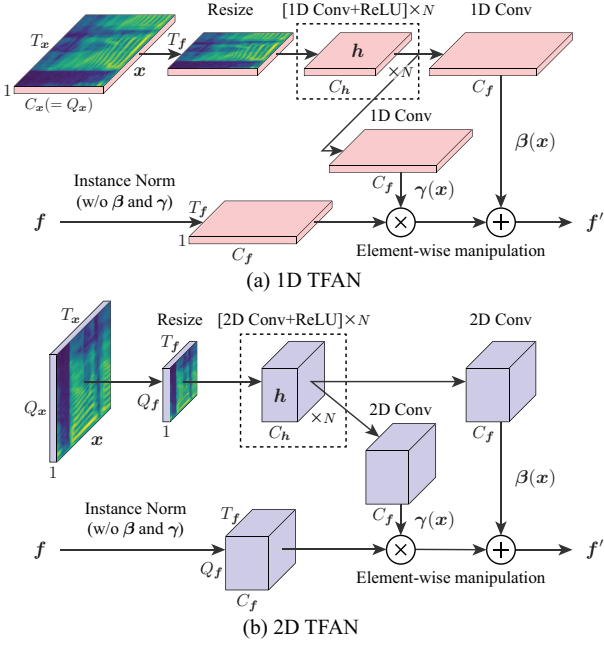


Figure 2: Architectures of 1D and 2D TFAN.  $x$ ,  $f$ ,  $f'$ , and  $h$  indicate the source mel-spectrogram, input feature, output feature, and hidden-layer feature, respectively.  $T$ ,  $Q$ , and  $C$  with subscripts denote the time, frequency, and channel dimensions in the corresponding feature, respectively. We omit them when their values are the same as those in the previous step.  $N$  represents the number of layers in the dashed-line box. After  $f$  is normalized in a channel-wise manner, it is modulated in an element-wise manner using scale  $\gamma(x)$  and bias  $\beta(x)$ , which are calculated using either (a) 1D CNN or (b) 2D CNN.

are not required in semantic image synthesis, whereas TFAN uses a multi-layer CNN to ensure the dynamical change, and (3) SPADE is based on batch normalization [37], whereas TFAN is based on IN. We empirically examine the effects of (1) and (2) and present our findings in Section 4.2.

### 3.2. Implementation

In CycleGAN-VC3, we incorporated TFAN into the CycleGAN-VC2 generator (i.e., 2-1-2D CNN (Section 2.2)). Particularly, IN in the 1D $\rightarrow$ 2D block and that in the upsampling block were replaced with 1D TFAN and 2D TFAN, respectively, as shown in Figure 3. In TFAN, we set the number of channels (i.e.,  $C_h$ ) and kernel size in  $h$  (Figure 2) to 128 and 5, respectively. We examine the performance with varying  $N$  and adjusting the position where the TFAN is inserted and present our findings in Section 4.2. The discriminator was the same as that used in CycleGAN-VC2 (i.e., PatchGAN (Section 2.3)).

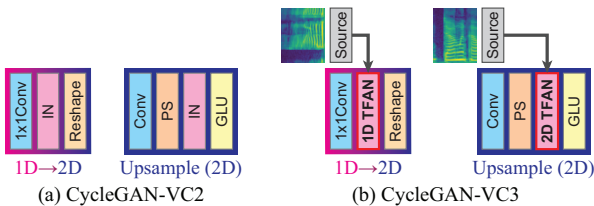


Figure 3: Comparison of 1D $\rightarrow$ 2D and upsampling blocks between (a) CycleGAN-VC2 and (b) CycleGAN-VC3. See Figure 4 in [21] for details of the overall network architectures.

<sup>2</sup><https://github.com/descriptinc/melgan-neurips>

Table 1: Comparison of MCD and MSD using (a) different-depth TFAN, (b) TFAN in different positions, and (c) different models. The results are listed as MCD [dB]/MSD [dB], where bold and italic numbers indicate the best and second-best scores, respectively.

(a) Depth ( $N$ )	SF-TF	SM-TM	SF-TM	SM-TF
1	7.55/ <b>1.54</b>	7.39/1.83	7.61/1.77	9.48/1.81
2	<b>7.53</b> /1.55	7.23/1.79	<b>6.91</b> /1.76	<b>7.97</b> /1.61
3	7.54/1.55	<b>7.10</b> / <b>1.78</b>	<b>6.91</b> / <b>1.74</b>	<b>7.97</b> /1.61
4	7.54/1.55	<b>7.10</b> /1.80	6.92/1.77	8.01/ <b>1.59</b>
(b) Position	SF-TF	SM-TM	SF-TM	SM-TF
1D→2D	7.79/1.58	7.12/1.93	7.05/1.82	8.15/1.65
Upsampling	7.75/1.58	7.51/1.84	6.96/1.80	8.95/1.79
Both	<b>7.54</b> / <b>1.55</b>	<b>7.10</b> / <b>1.78</b>	<b>6.91</b> / <b>1.74</b>	<b>7.97</b> / <b>1.61</b>
(c) Model	SF-TF	SM-TM	SF-TM	SM-TF
V1	7.91/1.60	7.85/2.21	7.56/2.04	8.45/1.74
V2	7.66/1.56	<b>7.07</b> /1.81	6.96/ <b>1.74</b>	8.07/1.63
V3	<b>7.54</b> /1.55	7.10/ <b>1.78</b>	<b>6.91</b> / <b>1.74</b>	<b>7.97</b> / <b>1.61</b>
V2U	7.69/ <b>1.53</b>	7.75/1.95	8.11/2.33	9.43/1.76

vious VC [21, 23, 24]: the *mel-cepstral distortion* (MCD), which measures the global structural difference based on the target and converted mel-cepstra, and *modulation-spectra distance* (MSD), which assesses the local structural difference based on the target and converted modulation spectra of mel-cepstra. For both metrics, the smaller the value, the better the performance. 35-dimensional mel-cepstrum parameters were extracted from the target or converted waveform using WORLD [41].

**Effect of TFAN parameter selection.** We initially examined the effect of the TFAN parameter selection. In particular, we investigated the performance with varying the depth of TFAN ( $N$  in Figure 2) and the position where the TFAN is inserted (1D→2D and/or upsampling blocks (see Figure 3)). Table 1(a) and (b) list the respective results. Our major findings are as follows. (1) *Comparison of different-depth TFAN* (Table 1(a)). We found that (i) 1-depth TFAN shows the worst scores except for SF-TF, where the performance is comparable with that of the others, and (ii) the scores reach their peak at approximately  $N = 3$ . This indicates the importance of performing dynamic changes using a multi-layer CNN. This differs from semantic image synthesis with SPADE [26] applied. (2) *Comparison of TFAN positions* (Table 1(b)). We found that the joint usage of 1D TFAN and 2D TFAN in 1D→2D and upsampling blocks is the most effective. Therefore, we set  $N$  to 3 and used TFAN in both positions in the following experiments.

**Comparison among CycleGAN-VCs.** We analyzed the performance difference among the CycleGAN-VCs. In addition to CycleGAN-VC (V1), -VC2 (V2), and -VC3 (V3), we examined CycleGAN-VC2 with U-net [42] (V2U), which might be useful as an alternative to TFAN to propagate the source information to converted features. Table 1(c) summarizes the results. V2 and V3 generally have the best or second-best scores. Among the two, V3 showed a better performance in most cases in terms of both metrics. Furthermore, we show the qualitative and subjective comparisons in Figure 1 and Section 4.3, respectively.

### 4.3. Subjective evaluation

We conducted listening tests to assess the applicability of CycleGAN-VCs in mel-spectrogram conversion. We compared four models: CycleGAN-VC2 with mel-cepstrum conversion (the current best and benchmark model; denoted by B) and CycleGAN-VC, -VC2, and -VC3 with mel-spectrogram conversion (denoted by V1, V2, and V3, respectively). To

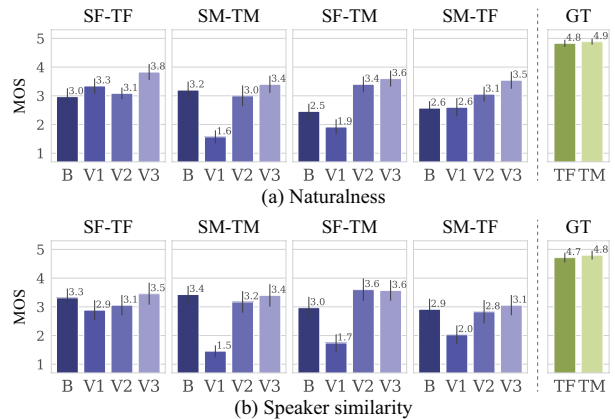


Figure 4: MOS for (a) naturalness and (b) speaker similarity with 95% confidence intervals.

measure the naturalness and speaker similarity, we conducted mean opinion score (MOS) tests,<sup>3</sup> where we included the target ground-truth speech (GT) as anchor samples. In both tests, all evaluation data (35 utterances) were used for evaluation, and each utterance was evaluated once. In the speaker similarity test, we paired the converted speech and target speech (of which utterance contents were different) in random order. Nine and 11 listeners participated in the naturalness and speaker similarity tests, respectively. Audio samples are available online.<sup>1</sup>

Figure 4 shows the results. Our main findings are as follows. (1) *Comparison between CycleGAN-VC2 for mel-cepstrum (B) and that for mel-spectrogram (V2)*. V2 is particularly effective for inter-gender VC (SF-TM and SM-TF) in terms of naturalness, and for the speaker similarity, the results are case-dependent. This indicates that the direct application to the mel-spectrogram (i.e., V2) is not necessarily reasonable. (2) *Comparison among CycleGAN-VCs (V1, V2, and V3)*. In most cases, the performance improves in terms of both metrics as the version increases. This confirms the utility of V2 and V3 compared with the previous versions. (3) *Effectiveness of CycleGAN-VC3 (V3)*. V3 has a better or competitive performance compared with the other models. This indicates the potential use of V3 as a new benchmark method in lieu of B.

## 5. Conclusions

Although CycleGAN-VCs are widely used as benchmark methods, their feasibility for mel-spectrogram conversion is not sufficiently examined. Therefore, following our examination, we proposed CycleGAN-VC3, an improvement of CycleGAN-VC2 that incorporates TFAN. Using this module, we can convert the mel-spectrogram while adaptively reflecting the source mel-spectrogram. The experimental results indicate the potential of using CycleGAN-VC3 as a new benchmark method to replace CycleGAN-VC2. Application of TFAN to advanced tasks, such as multi-domain VC [22, 23] and application-side VC [1, 2, 3, 4, 5], remains an interesting future direction.

## 6. Acknowledgements

This work was supported by JSPS KAKENHI 17H01763 and JST CREST Grant Number JPMJCR19A3, Japan.

<sup>3</sup>In the naturalness test, 1 = bad, 2 = poor, 3 = fair, 4 = good, and 5 = excellent. In the speaker similarity test, 1 = very different, 2 = moderately different, 3 = fair, 4 = moderately similar, and 5 = very similar.

## 7. References

- [1] A. B. Kain, J.-P. Hosom, X. Niu, J. P. H. van Santen, M. Fried-Oken, and J. Staehely, "Improving the intelligibility of dysarthric speech," *Speech Commun.*, vol. 49, no. 9, pp. 743–759, 2007.
- [2] K. Nakamura, T. Toda, H. Saruwatari, and K. Shikano, "Speaking-aid systems using GMM-based voice conversion for electrolaryngeal speech," *Speech Commun.*, vol. 54, no. 1, pp. 134–146, 2012.
- [3] Z. Inanoglu and S. Young, "Data-driven emotion conversion in spoken English," *Speech Commun.*, vol. 51, no. 3, pp. 268–283, 2009.
- [4] T. Toda, M. Nakagiri, and K. 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] T. Kaneko, H. Kameoka, K. Hiramatsu, and K. Kashino, "Sequence-to-sequence voice conversion with similarity metric learned using generative adversarial networks," in *Proc. Interspeech*, 2017, pp. 1283–1287.
- [6] Y. Stylianou, O. Cappé, and E. Moulines, "Continuous probabilistic transform for voice conversion," *IEEE Trans. Speech Audio Process.*, vol. 6, no. 2, pp. 131–142, 1998.
- [7] T. Toda, A. W. Black, and K. 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.
- [8] S. Desai, A. W. Black, B. Yegnanarayana, and K. Prahallad, "Spectral mapping using artificial neural networks for voice conversion," *IEEE Trans. Audio Speech Lang. Process.*, vol. 18, no. 5, pp. 954–964, 2010.
- [9] L. Sun, S. Kang, K. Li, and H. Meng, "Voice conversion using deep bidirectional long short-term memory based recurrent neural networks," in *Proc. ICASSP*, 2015, pp. 4869–4873.
- [10] K. Tanaka, H. Kameoka, T. Kaneko, and N. Hojo, "AttS2S-VC: Sequence-to-sequence voice conversion with attention and context preservation mechanisms," in *Proc. ICASSP*, 2019, pp. 6805–6809.
- [11] H. Kameoka, K. Tanaka, T. Kaneko, and N. Hojo, "ConvS2S-VC: Fully convolutional sequence-to-sequence voice conversion," in *arXiv preprint arXiv:1811.01609*, Nov. 2018.
- [12] L. Sun, K. Li, H. Wang, S. Kang, and H. Meng, "Phonetic posteriorgrams for many-to-one voice conversion without parallel data training," in *Proc. ICME*, 2016, pp. 1–6.
- [13] F.-L. Xie, F. K. Soong, and H. Li, "A KL divergence and DNN-based approach to voice conversion without parallel training sentences," in *Proc. Interspeech*, 2016, pp. 287–291.
- [14] Y. Saito, Y. Ijima, K. Nishida, and S. Takamichi, "Non-parallel voice conversion using variational autoencoders conditioned by phonetic posteriorgrams and d-vectors," in *Proc. ICASSP*, 2018, pp. 5274–5278.
- [15] J. Zhang, Z. Ling, and L.-R. Dai, "Non-parallel sequence-to-sequence voice conversion with disentangled linguistic and speaker representations," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 540–552, 2020.
- [16] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, "Voice conversion from non-parallel corpora using variational auto-encoder," in *Proc. APSIPA ASC*, 2016, pp. 1–6.
- [17] —, "Voice conversion from unaligned corpora using variational autoencoding Wasserstein generative adversarial networks," in *Proc. Interspeech*, 2017, pp. 3364–3368.
- [18] H. Kameoka, T. Kaneko, K. Tanaka, and N. 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.
- [19] T. Kaneko and H. Kameoka, "Parallel-data-free voice conversion using cycle-consistent adversarial networks," in *arXiv preprint arXiv:1711.11293*, Nov. 2017.
- [20] —, "CycleGAN-VC: Non-parallel voice conversion using cycle-consistent adversarial networks," in *Proc. EUSIPCO*, 2018, pp. 2114–2118.
- [21] T. Kaneko, H. Kameoka, K. Tanaka, and N. Hojo, "CycleGAN-VC2: Improved CycleGAN-based non-parallel voice conversion," in *Proc. ICASSP*, 2019, pp. 6820–6824.
- [22] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "StarGAN-VC: Non-parallel many-to-many voice conversion using star generative adversarial networks," in *Proc. SLT*, 2018, pp. 266–273.
- [23] T. Kaneko, H. Kameoka, K. Tanaka, and N. Hojo, "StarGAN-VC2: Rethinking conditional methods for StarGAN-based voice conversion," in *Proc. Interspeech*, 2019, pp. 679–683.
- [24] S. Lee, B. Ko, K. Lee, I.-C. Yoo, and D. Yook, "Many-to-many voice conversion using conditional cycle-consistent adversarial networks," in *Proc. ICASSP*, 2020, pp. 6279–6283.
- [25] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, "Auto-VC: Zero-shot voice style transfer with only autoencoder loss," in *Proc. ICML*, 2019, pp. 5210–5219.
- [26] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *Proc. CVPR*, 2019, pp. 2337–2346.
- [27] J. Lorenzo-Trueba, J. Yamagishi, T. Toda, D. Saito, F. Villavicencio, T. Kinnunen, and Z. Ling, "The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods," in *Proc. Odyssey*, 2018, pp. 195–202.
- [28] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. ICCV*, 2017, pp. 2223–2232.
- [29] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. NIPS*, 2014, pp. 2672–2680.
- [30] T. Zhou, P. Krähenbühl, M. Aubry, Q. Huang, and A. A. Efros, "Learning dense correspondence via 3D-guided cycle consistency," in *Proc. CVPR*, 2016, pp. 117–126.
- [31] Y. Taigman, A. Polyak, and L. Wolf, "Unsupervised cross-domain image generation," in *Proc. ICLR*, 2017, pp. 1–15.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770–778.
- [33] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," in *Proc. ICML*, 2017, pp. 933–941.
- [34] T. Kaneko, H. Kameoka, N. Hojo, Y. Ijima, K. Hiramatsu, and K. Kashino, "Generative adversarial network-based postfilter for statistical parametric speech synthesis," in *Proc. ICASSP*, 2017, pp. 4910–4914.
- [35] C. Li and M. Wand, "Precomputed real-time texture synthesis with Markovian generative adversarial networks," in *Proc. ECCV*, 2016, pp. 702–716.
- [36] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," in *arXiv preprint arXiv:1607.08022*, July 2016.
- [37] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. ICML*, 2015, pp. 448–456.
- [38] K. Kumar, R. Kumar, T. de Boissiere, L. Geste, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, "MelGAN: Generative adversarial networks for conditional waveform synthesis," in *Proc. NeurIPS*, 2019, pp. 14 881–14 892.
- [39] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *Proc. ICCV*, 2017, pp. 2794–2802.
- [40] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015, pp. 1–15.
- [41] M. Morise, F. Yokomori, and K. Ozawa, "WORLD: A vocoder-based high-quality speech synthesis system for real-time applications," *IEICE Trans. Inf. Syst.*, vol. 99, no. 7, pp. 1877–1884, 2016.
- [42] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015, pp. 234–241.