

Unsupervised Representation Disentanglement Using Cross Domain Features and Adversarial Learning in Variational Autoencoder Based Voice Conversion

Wen-Chin Huang , *Student Member, IEEE*, Hao Luo , Hsin-Te Hwang, Chen-Chou Lo, Yu-Huai Peng, Yu Tsao , *Member, IEEE*, and Hsin-Min Wang , *Senior Member, IEEE*

Abstract—An effective approach for voice conversion (VC) is to disentangle linguistic content from other components in the speech signal. The effectiveness of variational autoencoder (VAE) based VC (VAE-VC), for instance, strongly relies on this principle. In our prior work, we proposed a cross-domain VAE-VC (CDVAE-VC) framework, which utilized acoustic features of different properties, to improve the performance of VAE-VC. We believed that the success came from more disentangled latent representations. In this article, we extend the CDVAE-VC framework by incorporating the concept of adversarial learning, in order to further increase the degree of disentanglement, thereby improving the quality and similarity of converted speech. More specifically, we first investigate the effectiveness of incorporating the generative adversarial networks (GANs) with CDVAE-VC. Then, we consider the concept of domain adversarial training and add an explicit constraint to the latent representation, realized by a speaker classifier, to explicitly eliminate the speaker information that resides in the latent code. Experimental results confirm that the degree of disentanglement of the learned latent representation can be enhanced by both GANs and the speaker classifier. Meanwhile, subjective evaluation results in terms of quality and similarity scores demonstrate the effectiveness of our proposed methods.

Index Terms—Voice conversion, unsupervised learning, disentangled representation, variational autoencoder, adversarial learning, cross domain features.

I. INTRODUCTION

VOICE conversion (VC) aims to convert the speech from a source to that of a target without changing the linguistic content [1]. Speaker voice conversion [2] is a typical type of VC and refers to the process of converting speech from a

source speaker to a target speaker. In addition, a wide variety of applications could be solved by applying VC, such as accent conversion [3], personalized speech synthesis [4], [5], and speaking-aid device support [6]–[8]. Since the spectral property plays an important role in characterizing speaker individuality, spectral conversion has been intensively studied in VC. In this work, we focus on spectral mapping in speaker voice conversion.

Numerous VC approaches have been proposed. The Gaussian mixture model (GMM)-based method [9], [10] has been a popular statistical approach that estimates the joint density of the source-target feature vectors, which requires a training procedure and has a well-known disadvantage that the converted outputs generally suffer from an over-smoothing issue. Frequency warping methods, such as vocal tract length normalization [11], weighted frequency warping [12] and dynamic frequency warping [13], are able to keep spectral details while providing inferior speaker identity conversion quality to that of statistical approaches. Exemplar-based methods [14]–[18] require much less training data and are capable of modeling the high-dimensional spectra. In recent years, deep neural networks (DNNs) have established supremacy in a wide range of research fields, including VC [19]–[22]. DNNs have been utilized for not only spectral mapping but also neural vocoding [23]–[25]. It has been shown that employing neural vocoders as the waveform generation module can greatly improve the performance of VC systems [26]–[31]. It has also been shown that VC systems, whether implemented in high-dimensional or low-dimensional features, benefit from spectral detail compensation [15], [18], [32].

Nonetheless, most of the approaches described above rely on the availability of parallel training data, which is often not accessible in real world scenarios. Thus, the development of non-parallel VC methods has been gaining attention [33]. One approach is to construct a pseudo parallel dataset from a non-parallel corpus [34]. Another family of approaches utilizes a pre-trained automatic speech recognition model to compute the phonetic posteriorgram (PPG) as the speaker-independent linguistic feature, followed by a PPG-to-acoustic mapping to generate converted features [35], [36]. A recently popular approach is to use DNNs to model the probability distribution of the target features; state-of-the-art models such as variational autoencoders (VAEs) [37] and generative adversarial networks (GANs) [38] have been successfully applied to non-parallel VC [36], [39]–[46].

Manuscript received June 14, 2019; revised November 5, 2019 and December 28, 2019; accepted January 18, 2020. Date of publication April 6, 2020; date of current version July 22, 2020. This work was supported in part by the MOST-Taiwan under Grant MOST 107-2221-E-001-008-MY3 and Grant MOST 108-2634-F-001-004. (Corresponding author: Wen-Chin Huang.)

Wen-Chin Huang is with the Institute of Information Science, Academia Sinica, Taipei 11529, Taiwan, and also with the Graduate School of Informatics, Nagoya University, Nagoya 464-8601, Japan (e-mail: wen.chinhuang@g.sp.m.is.nagoya-u.ac.jp).

Hao Luo, Hsin-Te Hwang, Chen-Chou Lo, Yu-Huai Peng, and Hsin-Min Wang are with the Institute of Information Science, Academia Sinica, Taipei 11529, Taiwan (e-mail: b04901050@ntu.edu.tw; hwanght@iis.sinica.edu.tw; lochenchou@gmail.com; roland19930601@gmail.com; whm@iis.sinica.edu.tw).

Yu Tsao is with the Research Center of Information Technology Institute of Information Science, Academia Sinica, Taipei 11529, Taiwan (e-mail: yu.tsao@citi.sinica.edu.tw).

Digital Object Identifier 10.1109/TETCI.2020.2977678




Fig. 1. Illustration of how entangled latent representation affects the conversion performance in a general VAE-VC framework. The residual source speaker information in the latent code will be mixed with the given target speaker code, resulting in a mixed speaker identity in the converted feature. Thus, the performance might be harmed.

In this work, we focus on VAE-based VC (VAE-VC) [39]. Specifically, the spectral conversion function is composed of an encoder-decoder pair. The encoder encodes the input spectral feature into a latent code; the decoder mixes the latent code and a specified target speaker code to generate the converted feature. The encoder-decoder network and the speaker codes are trained by back-propagation of the reconstruction error, along with a Kullback-Leibler (KL)-divergence loss that regularizes the distribution of the latent code.

The degree of disentanglement of the latent representation is crucial to the success of many speech processing frameworks [47]–[51], including VAE-VC. Since we focus on the task of speaker voice conversion, the degree of disentanglement is defined as the amount of (source) speaker information residing in the latent code, i.e., the independence of the latent code and the speaker code [52]. An illustration is given in Figure 1. If the latent code is entangled by multiple components (e.g., in the VC task, the source speaker information remains in the latent code), during conversion, the decoder will draw the speaker information from both the given target speaker code and the residual source speaker information in the latent code, which harms the conversion performance. From the success of VAE-VC, we can infer that, at least to some extent, the decoder is trained to use more information in the given speaker code, rather than the speaker characteristics remained in the latent code, otherwise conversion made by changing the speaker code will not work. Although the success may be a natural result of model optimization, we doubt whether the performance is robust enough. For instance, in [53], it was demonstrated that the performance of autoencoder-based VC models was sensitive to the latent space dimension. This raises the need to design better schemes for making the latent code more independent of the speaker.

In our prior work [54], we proposed a cross-domain VAE-based VC framework (referred to as CDVAE-VC in the following discussion). The motivations of CDVAE-VC are: (1) although the effectiveness of VAE-VC using vocoder spectra (e.g., the STRAIGHT spectra, SPs [55]) has been confirmed, the use of other types of spectral features, such as mel-cepstral coefficients (MCCs) [56] that are related to human perception and have been widely used in VC, have not been properly investigated; (2) since modeling the low- and high-dimensional features alone has their respective shortcomings, based on multi-target/task learning [57], [58], it is believed that a model capable of simultaneously modeling two types of spectral features can yield better performance even if they are from the same feature

domain. To this end, CDVAE-VC [54] extended the VAE-VC framework to jointly consider two kinds of spectral features, namely SPs and MCCs. By introducing two additional cross-domain reconstruction losses and a latent similarity constraint into the training objective, the latent representations encoded from the input SPs and MCCs are biased to each other and capable of self- or cross-reconstructing the input features. We speculated that the success of CDVAE-VC came from the fact that a more disentangled latent representation was learned. Furthermore, we observed a positive correlation between the conversion performance and the extent to which the latent code was disentangled.

In this work, we extend the CDVAE-VC framework by incorporating the concept of adversarial training to improve the degree of disentanglement as well as the conversion performance. First, we directly combine CDVAE-VC with GANs. GANs have shown the ability to enhance the output of the decoder in encoder-decoder network based VC frameworks [45]. Therefore, it is expected that such a combination can improve the quality of converted speech. Second, inspired from the idea of domain adversarial training (DAT) [59], we add a speaker classification training objective to the latent variables, in order to explicitly project away speaker-related information. A similar idea has been applied to several speech processing tasks, such as speech recognition [60]–[62], speech enhancement [59], VC [45], [63] and singing VC [64]. Here, we utilize DAT by considering cross-domain features to further facilitate a more disentangled latent representation.

Designing a clear evaluation metric for degree of disentanglement has long been an open problem in the field of machine learning. In image modeling, visual inspection has been a standard and intuitive approach [65], [66]. However, the visual inspection is not perfectly feasible for speech processing tasks since it is hard to quantify the difference in voices as a specific latent variable changes. In previous works [45], [53], [67], a classifier-based metric has been proposed. Since the metric is also based on a trained classifier, it has limitations in comparing the disentanglement between different latent codes obtained by different models due to different training conditions and dynamics. Following [68], we utilize the parallel data that exist in most benchmark VC datasets and derive a novel metric for measuring disentanglement. The key assumption is that an ideal encoder should encode a pair of parallel sentences uttered by two different speakers to similar latent codes. We measure the cosine similarity between such latent codes to evaluate how well the encoder disentangles the latent codes.

The remainder of this paper is organized as follows. In Section II, we first review the VAE-VC and its extended version, CDVAE-VC. Section III introduces how to combine GANs with CDVAE-VC. Then, we describe how to add an adversarial speaker classifier objective to the latent code in Section IV. In Section V, we first examine our proposed mechanisms one by one, using conventional objective and subjective evaluation metrics adopted in VC. Disentanglement measurements of our proposed methods and how they are related to the VC performance are presented afterwards. Finally, we conclude the paper with discussions in Section VI.

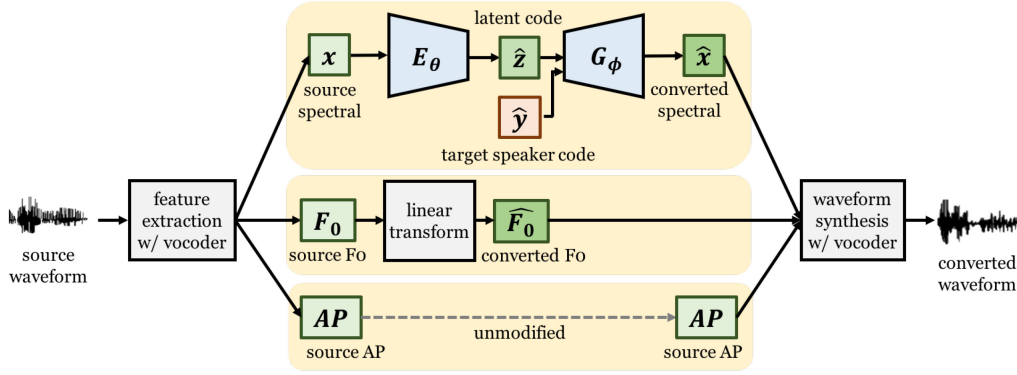


Fig. 2. Illustration of the conversion phase of the VAE-VC [39] framework. Following traditional VC systems, a vocoder first parameterizes the waveform into acoustic features, which are then converted in different streams, and finally the converted features are used to synthesize the converted waveform by a vocoder.

II. BACKGROUND

In conventional VC frameworks, the acoustic features of the source speaker are converted to those of the target speaker in different feature streams. Many researches focus on the conversion of spectral features [10] and thus formulate VC as follows. Given N source speaker's spectral frames $\mathbf{X}_s = \{\mathbf{x}_{s,1}, \dots, \mathbf{x}_{s,N}\}$, the goal is to find a conversion function f such that

$$\hat{\mathbf{x}}_{t,n} = f(\mathbf{x}_{s,n}). \quad (1)$$

Note that the second subindices in both sides of the equation are both n , which means that the converted spectral feature sequence has the same length with that of the source. In the rest of the article, we drop the frame or the speaker indices for simplicity.

In the following subsections, we describe two VAE based VC frameworks. Throughout the paper, we use “bar” to indicate the reconstructed features, and “hat” to indicate the converted features.

A. VAE-VC

Figure 2 depicts the conversion process of a typical VAE-VC system [39]. The core of VAE-VC is an encoder-decoder network. During training, given an observed (source or target) spectral frame \mathbf{x} , a speaker-independent encoder E_θ with parameter set θ encodes \mathbf{x} into a latent code: $\bar{\mathbf{z}} = E_\theta(\mathbf{x})$. The speaker code \mathbf{y} of the input frame is then concatenated with the latent code, and passed to a conditional decoder G_ϕ with parameter set ϕ to reconstruct the input. This reconstruction process can be expressed as:

$$\bar{\mathbf{x}} = G_\phi(\bar{\mathbf{z}}, \mathbf{y}) = G_\phi(E_\theta(\mathbf{x}), \mathbf{y}). \quad (2)$$

The model parameters can be obtained by maximizing the variational lower bound:

$$\mathcal{L}_{vae}(\theta, \phi; \mathbf{x}, \mathbf{y}) = \mathcal{L}_{recon}(\mathbf{x}, \mathbf{y}) + \mathcal{L}_{lat}(\mathbf{x}), \quad (3)$$

$$\mathcal{L}_{recon}(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{\mathbf{z} \sim q_\theta(\bar{\mathbf{z}}|\mathbf{x})} [\log p_\phi(\bar{\mathbf{x}}|\mathbf{z}, \mathbf{y})], \quad (4)$$

$$\mathcal{L}_{lat}(\mathbf{x}) = -D_{KL}(q_\theta(\bar{\mathbf{z}}|\mathbf{x}) \| p(\mathbf{z})), \quad (5)$$

where $q_\theta(\bar{\mathbf{z}}|\mathbf{x})$ is the approximate posterior, $p_\phi(\bar{\mathbf{x}}|\mathbf{z}, \mathbf{y})$ is the data likelihood, and $p(\mathbf{z})$ is the prior distribution of the latent

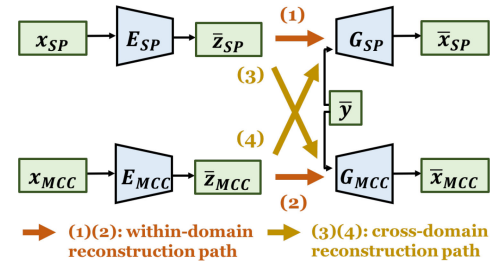


Fig. 3. Illustration of the training phase of the CDVAE-VC [54] framework. In this framework, each feature has its own set of encoder and decoder. During training, by minimizing the loss derived from the within- and cross-domain reconstruction paths, the latent codes \mathbf{z}_{SP} and \mathbf{z}_{MCC} learn to reconstruct not only corresponding input features but also the cross-domain features.

space. \mathcal{L}_{recon} is simply a reconstruction term as in any vanilla autoencoder, whereas \mathcal{L}_{lat} regularizes the encoder to align the approximate posterior with the prior distribution.

In the conversion phase, one could use (2) to formulate the conversion function f :

$$\hat{\mathbf{x}} = f(\mathbf{x}, \hat{\mathbf{y}}) = G_\phi(\hat{\mathbf{z}}, \hat{\mathbf{y}}) = G_\phi(E_\theta(\mathbf{x}), \hat{\mathbf{y}}), \quad (6)$$

where $\hat{\mathbf{y}}$ is the target speaker code.

The VAE framework makes several assumptions. First, $p_\phi(\bar{\mathbf{x}}|\mathbf{z}, \mathbf{y})$ is assumed to follow a normal distribution whose covariance is an identity matrix. Second, $p(\mathbf{z})$ is set to be a standard normal distribution. Third, the expectation over \mathbf{z} is approximated by sampling via a linear-transformation based re-parameterization trick [37]. With these simplifications, we can avoid intractability and optimize the autoencoder parameter sets $\theta \cup \phi$ and the speaker codes via back-propagation.

B. CDVAE-VC

In [54], we proposed the CDVAE-VC framework to utilize spectral features of different properties extracted from the same observed speech frame. As depicted in Figure 3, the CDVAE framework is formed by a collection of encoder-decoder pairs, one for each kind of spectral feature. Considering the SPs and MCCs as two kinds of spectral features (denoted as \mathbf{x}_{SP} and

\mathbf{x}_{MCC}), the following losses are defined:

$$\bar{\mathbf{z}}_{SP} = E_{SP}(\mathbf{x}_{SP}), \bar{\mathbf{z}}_{MCC} = E_{MCC}(\mathbf{x}_{MCC}), \quad (7)$$

$$\bar{\mathbf{x}}_{S-S} = G_{SP}(\bar{\mathbf{z}}_{SP}, \mathbf{y}), \bar{\mathbf{x}}_{M-M} = G_{MCC}(\bar{\mathbf{z}}_{MCC}, \mathbf{y}), \quad (8)$$

$$\bar{\mathbf{x}}_{S-M} = G_{MCC}(\bar{\mathbf{z}}_{SP}, \mathbf{y}), \bar{\mathbf{x}}_{M-S} = G_{SP}(\bar{\mathbf{z}}_{MCC}, \mathbf{y}), \quad (9)$$

$$\mathcal{L}_{in} = \mathcal{L}_{recon}(\bar{\mathbf{x}}_{S-S}, \mathbf{y}) + \mathcal{L}_{recon}(\bar{\mathbf{x}}_{M-M}, \mathbf{y}), \quad (10)$$

$$\mathcal{L}_{KLD} = \mathcal{L}_{lat}(\mathbf{x}_{SP}) + \mathcal{L}_{lat}(\mathbf{x}_{MCC}), \quad (11)$$

$$\mathcal{L}_{cross} = \mathcal{L}_{recon}(\bar{\mathbf{x}}_{S-M}, \mathbf{y}) + \mathcal{L}_{recon}(\bar{\mathbf{x}}_{M-S}, \mathbf{y}), \quad (12)$$

where E_{SP} and G_{SP} are the encoder and decoder for SPs, and E_{MCC} and G_{MCC} are the encoder and decoder for MCCs; $\bar{\mathbf{x}}_{S-S}$ and $\bar{\mathbf{x}}_{M-M}$, respectively, denote the generated SPs and MCCs from the within-domain reconstruction paths; $\bar{\mathbf{x}}_{M-S}$ and $\bar{\mathbf{x}}_{S-M}$, respectively, denote the generated SPs and MCCs from the cross-domain reconstruction paths. Note that $\mathcal{L}_{recon}(\cdot, \mathbf{y})$ calculates the reconstruction loss between the first argument and the corresponding input feature.

In short, we introduce two extra reconstruction streams. By minimizing the cross-domain reconstruction loss, we enforce \mathbf{z}_{SP} to contain enough information to reconstruct \mathbf{x}_{MCC} , and vice versa. As a result, the behavior of the encoders for both feature domains are constrained to be the same, i.e., they are expected to extract similar latent information from different types of input spectral features. To explicitly reinforce this constraint, a latent similarity L1 loss defined as

$$\mathcal{L}_{sim} = -\|\bar{\mathbf{z}}_{SP} - \bar{\mathbf{z}}_{MCC}\|_1, \quad (13)$$

can be included in the final objective expressed as:

$$\mathcal{L}_{cdvae} = \mathcal{L}_{in} + \mathcal{L}_{KLD} + \mathcal{L}_{cross} + \mathcal{L}_{sim}. \quad (14)$$

The model parameters can be learned by maximizing (14). In the conversion phase, there are four conversion paths (i.e., two within-domain and two cross-domain paths). As reported in [54], the CDVAE MCC-MCC path gave the best performance in terms of subjective evaluation, which matched the assumption that MCCs are more related to human perception.

III. INCORPORATING CDVAE-VC WITH GANS

Minimizing the reconstruction loss in VAE-VC and CDVAE-VC tends to result in blurry spectra, similar to the over-smoothing effects in other VC frameworks. It is expected that introducing a GAN objective [38] can guide the output spectra to be more realistic. In this section, we present the main concepts and system architectures of the combination of GANs and the VAE-VC and CDVAE-VC frameworks.

A. The GAN Objective in the General VAE-VC

We follow [69] and incorporate a GAN objective into the decoder in the original VAE-VC. Assume that the real data distribution of any spectral frame admits density p^* , and the autoencoding process defined in (2) induces a conditional probability $p_{\bar{\mathbf{x}}}$. From the data distribution prospective of view, the goal is to enhance the decoder network G in (2) such that $p_{\bar{\mathbf{x}}}$

best approximates the real data distribution p^* :

$$p_{\bar{\mathbf{x}}}(G_{\phi}(E_{\theta}(\mathbf{x}), \mathbf{y})) \approx p^*(\mathbf{x}). \quad (15)$$

A typical GAN [38] realizes the above-mentioned probability approximation by introducing a discriminator D_{ψ} with parameter set ψ that judges whether an input follows a true and natural probability distribution or an artificial one. Together with a generator G that tries to produce realistic output features, these two components play a min-max game and seek an equilibrium with the Jensen-Shannon divergence \mathcal{D}_{JS} as the objective, which is defined as follows:

$$\begin{aligned} \mathcal{L}_{gan}(\theta, \psi; x) &= 2\mathcal{D}_{JS}(p^*||p_{\bar{\mathbf{x}}}) + 2\log 2 \\ &= \mathbb{E}_{\mathbf{x} \sim p^*}[\log D_{\psi}(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim q_{\theta}}[\log D_{\psi}(G_{\phi}(\mathbf{z}))]. \end{aligned} \quad (16)$$

To facilitate stable training, in this work we adopt a Wasserstein GAN (WGAN) [70], [71]. In the WGAN, the following Wasserstein distance is derived:

$$W(p^*, p_{\bar{\mathbf{x}}}) = \sup_{\|D\|_L \leq 1} \mathbb{E}_{\mathbf{x} \sim p^*}[D(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \sim p_{\bar{\mathbf{x}}}}[D(\mathbf{x})], \quad (17)$$

where the supremum is over all 1-Lipschitz functions $D: \mathcal{X} \rightarrow \mathbb{R}$. Based on the above distance, the following WGAN loss can be defined:

$$\mathcal{L}_{wgan}(\mathbf{x}) = \mathbb{E}_{\mathbf{x} \sim p^*}[D_{\psi}(\mathbf{x})] - \mathbb{E}_{\mathbf{z} \sim q_{\theta}(\bar{\mathbf{z}}|\mathbf{x})}[D_{\psi}(G_{\phi}(\mathbf{z}, \mathbf{y}))], \quad (18)$$

where D_{ψ} is now a 1-Lipschitz discriminator. Finally, we can combine the objectives of VAE and WGAN by assigning the decoder of VAE as the generator of WGAN. As a result, combining the WGAN loss (18) and the VAE loss (3) results in a VAEGAN objective:

$$\mathcal{L}_{vaeagan}(\theta, \phi, \psi; \mathbf{x}, \mathbf{y}) = \mathcal{L}_{vae}(\mathbf{x}, \mathbf{y}) + \alpha \mathcal{L}_{wgan}(\mathbf{x}), \quad (19)$$

where α is the weight of the WGAN loss. This objective is shared across the encoder, decoder, and discriminator. As in standard GAN training, the discriminator is first updated by maximizing this objective, and the encoder and decoder are updated by minimizing the objective. Therefore, the components are optimized in an alternating order. GANs produce more realistic (in our case, sharper) outputs because they optimize a loss function between two distributions in a more direct fashion.

The VAW-GAN-VC method in [41] has a similar motivation to better model spectral features to improve feature generation. However, there is a fundamental difference between the training procedures of VAW-GAN-VC and the training procedures here. In VAW-GAN-VC, the objective of WGANs is to minimize the Wasserstein distance of the two distributions of the converted features and the real target features. Although this is a strong objective, it also brings some limitations. The original VAE-VC and CDVAE-VC consider only auto-encoding in the training phase, and perform conversion by changing the speaker code in the conversion phase. In other words, multiple conversion pairs are integrated into one model, sometimes referred to as ‘‘multi-target’’ training in VC. VAW-GAN-VC, in contrast, needs

to consider not only auto-encoding but also conversion in the training phase, since the discriminator needs to discriminate the real target features and the converted features in order to align the distribution of the latter to that of the former. As a result, VAW-GAN-VC is trained to convert from one source to one target, which limits the flexibility of the model. In this work, we intend to maintain the multi-target flexibility in CDVAE-VC and thus design the WGAN objective to match the distributions of the real features and the reconstructed features. Considering this fundamental difference and to avoid confusion, we focus on multi-target VC and do not take VAW-GAN-VC into discussion and comparison in this paper.

B. CDVAE-VC With GANs (CDVAE-GAN)

Now we can combine the GAN objective with CDVAE-VC, which we will refer to as CDVAE-GAN, where the derivation of the objective is as simple as replacing the VAE loss in (19) with the CDVAE objective defined in (14). However, in practice, combining CDVAE-VC with GANs is not as trivial as replacing the encoder and decoder in VAE-GAN with CDVAE. For each kind of feature, a separate discriminator should be trained, i.e., D_{SP} and D_{MCC} should be considered. It seems natural to train two discriminators jointly with the whole network. However, as mentioned above, the MCC-MCC path in CDVAE-VC performs best in four paths in the conversion phase. Introducing a discriminator for SPs might not necessarily benefit the quality of the output MCCs. To determine the best architecture, we examine the effect of three settings, including combining CDVAE with only D_{SP} , only D_{MCC} , and both D_{SP} and D_{MCC} . Detailed experimental results will be shown in Sections V-C and V-D.

IV. ADVERSARIAL SPEAKER CLASSIFIER (CLS)

As discussed above, the viability of the family of VAE-VC frameworks relies on the decomposition of input, which is assumed to be composed of phonetic representation and speaker information. Ideally, the latent code extracted using the encoder should contain solely phonetic information and free from any speaker information. However, this decomposition is not explicitly guaranteed. To this end, we investigate the effect of an adversarial speaker classifier to explicitly force the latent code to be speaker independent.

A. The Classifier Loss

An adversarial speaker classifier C_Ψ with parameter set Ψ tries to classify which speaker the latent code comes from. We will refer to this classifier as CLS. Specifically, given a latent code z , the CLS predicts a posterior probability $P(\mathbf{y} = y|z)$, which is the probability that z is extracted from an input frame produced by speaker y . Therefore, we can define the CLS loss as the negative cross-entropy between the predicted posterior and the one-hot ground truth vector:

$$\mathcal{L}_{cls}(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{z \sim q_\theta(\mathbf{z}|\mathbf{x})} [-\log P(\mathbf{y}|z)]. \quad (20)$$

B. CDVAE-GAN With CLS (CDVAE-CLS-GAN)

We now augment the CDVAE-GAN framework with the adversarial speaker classifier, which we will refer to as CDVAE-CLS-GAN. Adding the CLS loss (20) to the CDVAE-GAN loss, we obtain the final objective:

$$\mathcal{L}_{all}(\theta, \phi, \psi, \Psi; \mathbf{x}, \mathbf{y}) = \mathcal{L}_{cdvae}(\mathbf{x}, \mathbf{y}) + \alpha \mathcal{L}_{wgan}(\mathbf{x}) + \lambda \mathcal{L}_{cls}(\mathbf{x}, \mathbf{y}), \quad (21)$$

where λ is the weight of the classifier loss. This objective is shared, again, across the encoder, decoder, discriminator, and classifier.

The training process is divided into three phases, as depicted in Figure 4. Phase one involves the training of the VAE. In phase two, to pre-train the classifier, we first use the trained VAE obtained in phase one to extract latent codes from the same training set. The classifier is then trained with these latent codes to minimize (20). In the third phase, we train the whole network using an alternating update schedule, similar to the one described in Section III-A. Specifically, the encoder and the decoder are first frozen and the discriminator and classifier are trained to maximize \mathcal{L}_{wgan} and minimize \mathcal{L}_{cls} defined in (18) and (20), respectively, and thus they can discriminate self-reconstructed features and classify latent codes correctly. Then, we freeze these modules and train the encoder and decoder to not only minimize \mathcal{L}_{cdvae} in (14), but also optimize \mathcal{L}_{wgan} and \mathcal{L}_{cls} so that they can fool the frozen components.

The described training scheme also plays a min-max game between {encoders, decoders} and {discriminator, classifier}. An ideally trained model should contain encoders that learns to project away as much speaker information as possible and decoders that can generate realistic and natural output spectra given an inferred latent code with a specific speaker code. Algorithm 1 summarizes the training procedure of CDVAE-CLS-GAN.

V. EXPERIMENTAL EVALUATIONS

A. Experimental Settings

We conducted all experiments on the Voice Conversion Challenge (VCC) 2018 dataset, which contained recordings of 12 professional US English speakers with a sampling rate of 22050 Hz. The training and testing sets, respectively, consisted of 81 utterances and 35 utterances per speaker. We further divided the training utterances into 70/11 training/validation sets. The WORLD vocoder was used to extract acoustic features, including 513-dimensional SPs, 513-dimensional aperiodicity signals (APs), and fundamental frequency (F_0). 35-dimensional MCCs were then extracted from the SPs, which were then normalized to unit-sum, and the normalizing factor was used as the energy of SPs. The 0-th coefficient of MCCs was taken out as the energy of MCCs. We further applied Min-Max normalization to SPs and MCCs. In the conversion phase, the converted SPs in VAE systems and the converted MCCs in CDVAE systems (excluding CDVAE-GAN with D_{SP}) were obtained. The energy and AP were kept unmodified, and F_0 was converted using a linear mean-variance transformation in the $\log-F_0$ domain.




Fig. 4. Illustration of the training procedure of our proposed CDVAE-CLS-GAN model. *Phase 1*: A CDVAE is trained. *Phase 2*: The latent codes are used to train the CLS. *Phase 3-A and 3-B*: The encoders, decoders and the CLS, discriminators are trained in an alternating order.

Algorithm 1: Training Procedure of CDVAE-CLS-GAN.

```

function AUTOENCODE  $x, y$ 
   $z \leftarrow$  sample using  $E(x)$ 
   $\bar{x} \leftarrow G(z, y)$ 
  return  $\bar{x}, z$ 
 $\phi, \theta, \psi, \Psi \leftarrow$  initialization
// Phase 1: train the VAE
while not converged do
   $X, Y \leftarrow$  mini-batch of samples from the training set
   $\bar{X}, Z \leftarrow$  AUTOENCODE( $X, Y$ )
   $\phi, \theta \xleftarrow{\text{update}} -\nabla \mathcal{L}_{cdvae}(\bar{X}, Y)$ 
// Phase 2: train the CLS
while not converged do
   $X, Y \leftarrow$  mini-batch of samples from the training set
   $\mathcal{L}_{cls} \leftarrow \mathcal{L}_{cls}(X, Y)$ 
   $\Psi \xleftarrow{\text{update}} -\nabla \mathcal{L}_{cls}(X, Y)$ 
// Phase 3: train the whole network
while not converged do
   $X, Y \leftarrow$  mini-batch of samples from the training set
   $\bar{X}, Z \leftarrow$  AUTOENCODE( $X, Y$ )
   $\mathcal{L}_{cdvae} \leftarrow \mathcal{L}_{cdvae}(\bar{X}, Y)$ 
   $\mathcal{L}_{wgan} \leftarrow \mathcal{L}_{wgan}(X)$ 
   $\mathcal{L}_{cls} \leftarrow \mathcal{L}_{cls}(X, Y)$ 
// Update the discriminator and classifier
while not converged do
   $\psi \xleftarrow{\text{update}} -\nabla_{\psi}(-\mathcal{L}_{wgan})$ 
   $\Psi \xleftarrow{\text{update}} -\nabla_{\Psi} \mathcal{L}_{cls}$ 
// Update the encoder and generator
while not converged do
   $\theta \xleftarrow{\text{update}} -\nabla_{\theta}(\mathcal{L}_{cdvae} + \alpha \mathcal{L}_{wgan} - \lambda \mathcal{L}_{cls})$ 
   $\phi \xleftarrow{\text{update}} -\nabla_{\phi}(\mathcal{L}_{cdvae} + \alpha \mathcal{L}_{wgan})$ 

```

The detailed network architectures are shown in Table I. We adopted the fully convolutional network (FCN) [72] based CDVAE-VC as our baseline system [68], which consumes continuous spectral frames extracted from the whole utterance and outputs a sequence of converted frames of the same length. This model has been confirmed to outperform the frame-wise CDVAE-VC counterpart. We also adopted a gradient penalty regularization [71] in the WGAN objective to stabilize the training. Layer normalization [73], the gated linear units activation

TABLE I

MODEL ARCHITECTURES. CONV- $H \times W \times N$ INDICATES A CONVOLUTIONAL LAYER WITH KERNEL SIZE $H \times W$ AND N OUTPUT CHANNELS. LReLU INDICATES THE LEAKY RELU ACTIVATION FUNCTION. FC INDICATES FULLY-CONNECTED LINEAR LAYER. LN INDICATES THE LAYER NORMALIZATION LAYER

	ConvLReLU	Conv-3x1-n, LN, LReLU
	ConvGLU	(Conv-3x1-n, LN, sigmoid) \odot (Conv-3x1-n, LN, tanh)
E_{SP}	ConvLReLU $\times 5$ (n=1024, 512, 256, 64, 32), FC-16 (μ), FC-16 (σ)	
E_{MCC}	ConvLReLU $\times 5$ (n=512, 256, 128, 64, 32), FC-16 (μ), FC-16 (σ)	
G_{SP}	(Concat with y , ConvGLU) $\times 4$ (n=128, 256, 512, 1024), Concat with y , Conv-3x1-513	
G_{MCC}	(Concat with y , ConvGLU) $\times 4$ (n=64, 128, 256, 512), Concat with y , Conv-3x1-513	
D_{SP}	ConvLReLU $\times 5$ (n=1024, 512, 256, 64, 32), FC-1	
D_{MCC}	ConvLReLU $\times 5$ (n=512, 256, 128, 64, 32), FC-1	
C	(Concat with y , ConvGLU) $\times 4$ (n=128, 256, 512, 1024), Concat with y , Conv-3x1-513	

function, and skip connections were also used to more effectively propagate the conditional information.

Following [68], the latent space and speaker representation were set to 16-dimensional. We used a mini-batch of 16 and the Adam optimizer with a fixed learning rate of 0.0001. The hyper-parameters α and λ were set to be 50 and 1000, respectively, according to a held-out validation set. For CDVAE-GAN, we first pre-trained the CDVAE for 100000 steps. Then, we adversarially trained the discriminator(s) with the whole network for 10000 steps. We followed a common WGAN training scheme [70], [71] such that the discriminator(s) were updated for 5 iterations followed by 1 iteration of encoder and decoder update. For CDVAE-CLS-GAN, after training the CDVAE for 100000 steps, we pre-trained the classifier with the latent code extracted from the encoders for 30000 steps. Then, we trained the whole network for 10000 steps. After experimenting with different training schemes, here we updated the discriminator and the classifier for 1 iteration followed by 5 iterations of encoder and decoder update.

The following models are compared in order to examine the effectiveness of our proposed methods.

- **VAE**: The FCN version of the VAE-VC model introduced in [39]. This model is only used to evaluate the impact of cross domain features on the degree of disentanglement.

- **CDVAE**: The FCN model in [68], which is the baseline model in our experiments.
- **CDVAE-GAN_{SP}**: The CDVAE with D_{SP} .
- **CDVAE-GAN_{MCC}**: The CDVAE with D_{MCC} .
- **CDVAE-GAN_{BOTH}**: The CDVAE with D_{SP} and D_{MCC} .
- **CDVAE-CLS**: The CDVAE with CLS.
- **CDVAE-CLS-GAN_{SP}**: The CDVAE with D_{SP} and CLS.
- **CDVAE-CLS-GAN_{MCC}**: The CDVAE with D_{MCC} and CLS.

For simplicity, in the rest of the paper, we use brackets to surround the type of feature used during conversion, and that path will be used in CDVAE-based methods. For instance, **CDVAE-GAN_{MCC}** [MCC] uses the MCC and the MCC-MCC path. In addition, if MCC is used in **CDVAE** and **CDVAE-CLS**, we additionally compare systems incorporating the global variance (GV) post-filter [74] to enhance the output, as in the original CDVAE [54].

B. Evaluation Methodology

1) Objective Evaluation Metrics:

- **Mel-Cepstrum distortion (MCD)**: MCD measures the spectral distortion in the MCC domain, and is a commonly adopted objective metric in the field of VC. It is calculated as:

$$MCD[dB] = \frac{10}{\log 10} \sqrt{2 \sum_{d=1}^K \left(mcc_d^{(c)} - mcc_d^{(t)} \right)^2}, \quad (22)$$

where K is the dimension of the MCCs and $mcc_d^{(c)}$ and $mcc_d^{(t)}$ represent the d -th dimensional coefficient of the converted MCCs and the target MCCs, respectively. In practice, MCD is calculate in a utterance-wise manner. A dynamic time warping (DTW) based alignment is performed to find the corresponding frame pairs between the non-silent converted and target MCC sequences beforehand.

- **Global variance (GV)**: GV serves as a metric for the over-smoothness of the output features. GV is usually calculated dimension-wise over all non-silent frames in the evaluation set. The d -dimensional GV value is calculated as follows:

$$GV[d] = \frac{1}{N} \sum_{n=1}^N \left(mcc_{n,d}^{(c)} - \overline{mcc}_d^{(c)} \right)^2, \quad (23)$$

where $\overline{mcc}_d^{(c)}$ is the mean of all converted d -th dimensional MCC coefficients.

- **Modulation Spectrum (MS)**: MS [75] is defined as the log-scaled power spectrum of a given feature sequence. The temporal fluctuation of the sequence is first decomposed into individual modulation frequency components, and their power values are represented as the MS. In this work we measure the MS of MCCs. Different from previous works that measured the MS of specific dimension of the MCC sequence, here we report the average of all dimensions. We also measure a MS distortion (MSD), where the

TABLE II
MEAN MEL-CEPSTRAL DISTORTIONS [DB] OF ALL NON-SILENT FRAMES IN THE EVALUATION SET FOR THE COMPARED MODELS

Model	1				
	F-F	M-M	M-F	F-M	Avg.
CDVAE [MCC]	6.56	5.76	6.96	6.27	6.39
CDVAE-GAN _{SP} [SP]	7.09	6.38	7.38	6.93	6.94
CDVAE-GAN _{MCC} [MCC]	7.52	6.65	7.87	7.27	7.33
CDVAE-GAN _{BOTH} [MCC]	7.44	6.85	7.90	7.30	7.37
CDVAE-CLS [MCC]	6.65	6.29	7.02	6.49	6.61
CDVAE-CLS-GAN _{SP} [SP]	7.23	6.91	7.64	7.08	7.21
CDVAE-CLS-GAN _{MCC} [MCC]	7.71	6.62	7.76	7.05	7.29
CDVAE-CLS-GAN _{BOTH} [MCC]	7.57	7.13	8.12	7.30	7.53

MSD for the d -dimension is calculated by:

$$MSD[d] = \sqrt{\frac{1}{N} \sum_{n=1}^N \left(mcc_{n,d}^{(t)} - mcc_{n,d}^{(c)} \right)^2}. \quad (24)$$

2) *Subjective Evaluation Methods*: We recruited 14 participants for the following two subjective evaluations.¹

- **The mean opinion score (MOS) test on naturalness**: Subjects were asked to evaluate the naturalness of the converted and natural speech samples on a scale from 1 (completely unnatural) to 5 (completely natural).
- **The VCC [33] style test on similarity**: This paradigm was adopted by the VCC organizing committee. Listeners were given a pair of speech utterances consisting of a natural speech sample from a target speaker and a converted speech sample. Then, they were asked to determine whether the pair of utterances can be produced by the same speaker, with a 4-level confidence of their decision, i.e., sure or not sure.

C. Applying GANs to Different Features

We first compare **CDVAE-GAN_{SP}**, **CDVAE-GAN_{MCC}**, **CDVAE-GAN_{BOTH}** and **CDVAE-CLS-GAN_{SP}**, **CDVAE-CLS-GAN_{MCC}**, **CDVAE-CLS-GAN_{BOTH}**, respectively. As in Table II, **CDVAE-GAN_{BOTH}** and **CDVAE-CLS-GAN_{BOTH}** gave the highest MCD, while in Figures 5(b), 5(c), 6(b) and 6(c), we can see that in terms of GV and MS, **CDVAE-GAN_{MCC}** and **CDVAE-CLS-GAN_{MCC}** yielded curves closer to the target curves, where the curves of the other models deviated more from the target curves. Meanwhile, consistent with a common observation in the VC literature that MCD, which measures the sample mean, often yields opposite results to GV and MS, both presenting the sample variance [10], [76]. This result suggests that modeling both feature domains simultaneously does not always yield better results. As for perceptual performance, our internal listening tests revealed that **CDVAE-GAN_{MCC}** gave the best results among the three models. Note that although **CDVAE-GAN_{SP}** and **CDVAE-CLS-GAN_{SP}** gave the lowest

¹A demo web page with samples used for subjective evaluation is available at <https://unilight.github.io/CDVAE-GAN-CLS-Demo/>




Fig. 5. Global variance curves of all non-silent frames averaged over all conversion pairs for the compared models.

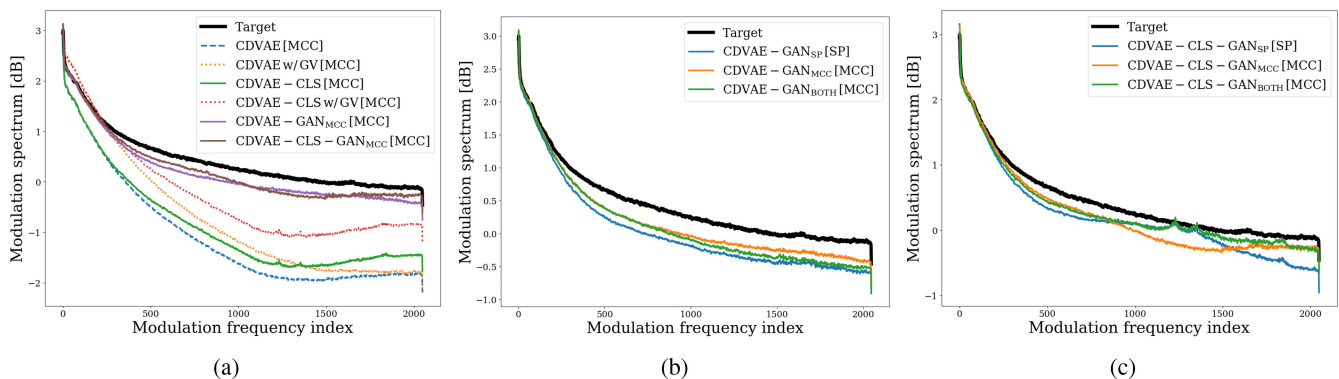


Fig. 6. Average modulation spectrum curves over all dimensions of all non-silent frames over all conversion pairs for the compared models.

TABLE III
MEAN OPINION SCORES ON NATURALNESS FOR THE COMPARED MODELS AND THE NATURAL TARGET VOICE WITH 95% CONFIDENCE INTERVALS

Model	F-F	M-M	M-F	F-M	Avg.
CDVAE [MCC]	2.50 ± 0.27	2.42 ± 0.21	2.31 ± 0.25	2.28 ± 0.31	2.40 ± 0.28
CDVAE-CLS [MCC]	2.72 ± 0.37	2.61 ± 0.48	2.44 ± 0.30	2.17 ± 0.29	2.55 ± 0.39
CDVAE w/ GV [MCC]	3.36 ± 0.44	3.36 ± 0.26	2.94 ± 0.30	2.89 ± 0.50	3.12 ± 0.30
CDVAE-CLS w/ GV [MCC]	3.50 ± 0.53	3.53 ± 0.49	3.06 ± 0.31	3.17 ± 0.38	3.30 ± 0.38
CDVAE-GAN _{MCC} [MCC]	3.19 ± 0.40	3.06 ± 0.38	2.61 ± 0.36	2.58 ± 0.27	2.95 ± 0.30
CDVAE-CLS-GAN _{MCC} [MCC]	2.94 ± 0.37	3.58 ± 0.48	3.06 ± 0.26	3.11 ± 0.57	3.15 ± 0.36
Target	-	-	-	-	4.75 ± 0.25

MCD compared with the other two models, they do not necessarily outperform their MCC counterparts in listening tests. We speculate that fitting the SP domain tends to give more over-smoothed output features, resulting in low MCDs but not beneficial for improving perceptual performance. The result is reasonable since the MCC-MCC path is used when performing conversion.

D. Effectiveness of GANs

Next, we examine the effectiveness of combining GANs with CDVAE and CDVAE-CLS. Based on the discussion in the previous subsection, we focus on CDVAE-GAN_{MCC} and

CDVAE-CLS-GAN_{MCC} here. As in Figures 5(a), 6(a) and 7, we can see that CDVAE-GAN_{MCC} and CDVAE-CLS-GAN_{MCC} fit the GV and MS statistics to the target much better than CDVAE and CDVAE-CLS, respectively. Also, models with GAN yield very small MSD comparing to the rest of the models. This confirms that involving a GAN objective in training indeed improves the modeling of the statistics, especially the variance of real speech data. Table III and Figure 8 show the subjective evaluation results. The *t*-test showed that CDVAE-GAN_{MCC} significantly outperformed CDVAE with a *p*-value of 4.21×10^{-5} . Meanwhile, CDVAE-CLS-GAN_{MCC} significantly outperformed CDVAE-CLS with a *p*-value of 4.02×10^{-4} . These results confirm the effectiveness of GAN. On the other hand,

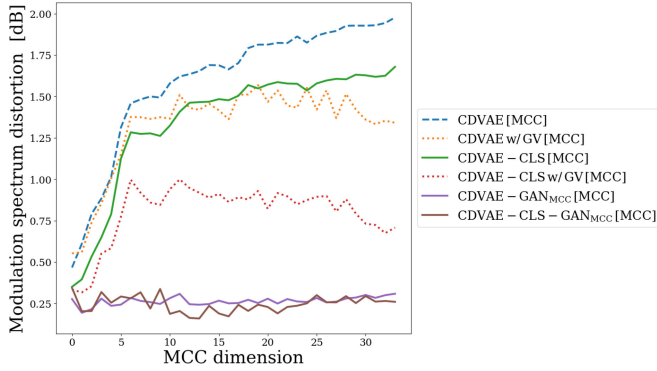


Fig. 7. Modulation spectrum distortion curves of all non-silent frames over all conversion pairs for the compared models.




Fig. 8. Similarity results over all speaker pairs for the compared models.

CDVAE-GAN_{MCC} performed comparably with **CDVAE** with GV post-processing (p -value = 4.03×10^{-2}), while **CDVAE-CLS-GAN_{MCC}** performed comparably with **CDVAE-CLS** with GV post-processing (p -value = 7.64×10^{-2}). These results are consistent with our findings in the objective evaluations, suggesting that GANs enhance the variance of output features, thus have the potential to replace the GV post-filtering process commonly involved in traditional MCC-based VC systems [10]. This is advantageous since the model can then be freed from the post-filtering process in the online conversion phase, which may benefit real-time applications.

E. Effectiveness of CLS

Next, we evaluate the effectiveness of the adversarial speaker classifier. Looking at the **CDVAE**, **CDVAE-GAN** models and their counterparts with CLS, a trend of increase in MCD values can be observed in Table II. On the other hand, Figures 5(a), 6(a) and 7 show that applying CLS to **CDVAE** and **CDVAE-CLS-GAN_{MCC}** yields similar GV values, but with MS values closer to those of the target, as well as a smaller MSD. These results imply that CLS can improve objective statistics.

Table III and Figure 8 show the subjective evaluation results. The effectiveness of CLS can be confirmed by the following observations: The speech naturalness was improved in all conversion pairs, by adding CLS to **CDVAE**, **CDVAE w/ GV**, and **CDVAE-GAN_{MCC}**. This is consistent with our aforementioned findings from the objective evaluations. Furthermore, the conversion similarity is greatly improved when incorporating CLS in **CDVAE** and **CDVAE w/ GV**, and is slightly improved when added to **CDVAE-GAN_{MCC}**. This confirms our initial

TABLE IV
THE RESULTS OF DEM: THE COSINE SIMILARITY OF THE LATENT CODES EXTRACTED FROM NON-SILENT FRAMES OF PARALLEL UTTERANCES OF SOURCE-TARGET PAIRS

Model	F-F	M-M	M-F	F-M	Avg.
VAE [SP]	.568	.633	.534	.552	.571
CDVAE [SP]	.597	.658	.557	.577	.597
CDVAE-GAN _{SP} [SP]	.605	.677	.565	.594	.610
CDVAE-CLS [SP]	.573	.582	.508	.535	.550
CDVAE-CLS-GAN _{SP} [SP]	.629	.638	.573	.602	.610
CDVAE [MCC]	.530	.588	.476	.502	.524
CDVAE-GAN _{MCC} [MCC]	.559	.609	.502	.534	.551
CDVAE-CLS [MCC]	.575	.581	.507	.533	.549
CDVAE-CLS-GAN _{MCC} [MCC]	.583	.621	.561	.563	.584

motivation of CLS, which is to increase speaker similarity by eliminating source speaker identity in the latent code.

F. Disentanglement Measure

In this section, we investigate the degree of disentanglement of the VC models involved in this study. We use a novel metric that was recently proposed in [68] as the disentanglement measurement, termed DEM. The main design concept of DEM is that a pair of sentences of the same content uttered by the source and target speakers should have similar latent codes since the phonetic contents are the same. Therefore, we can use the cosine similarity to measure the distance of the latent codes obtained from the paired utterances. Specifically, the procedure to calculate DEM is as follows:

- 1) extracting the latent codes of a pair of parallel utterances spoken by the source and target speakers;
- 2) aligning the frame sequences of the pair of utterances using DTW;
- 3) calculating the frame-wise cosine similarity, and then taking the average of the entire sequence.

As with other popular evaluation metrics, e.g., MCD and MSD, computing DEM requires parallel data. Since parallel data are usually available in standardized VC datasets, DEM is a simple but effective measure of the degree of disentanglement of the latent codes.

Table IV shows the evaluation results of DEM. First, we observe that **CDVAE [SP]** yields higher DEM scores than **VAE [SP]**. This confirms that introducing cross domain features indeed increases the degree of disentanglement. Next, comparing the corresponding methods in the upper and lower half of the table, which used SP and MCC as input features respectively, the DEM scores of the upper is consistently higher than those of the latter. This result is somehow reasonable because here SPs (513-dimensional) are of higher dimensions than MCCs (35-dimensional) and carry much detailed information. As a result, in terms of cosine similarity measure, higher DEM could be observed in the upper half methods than the lower half.

One interesting finding here is that when incorporating GANs in **CDVAE** and **CDVAE-CLS** models, the DEM scores are consistently and significantly improved. This result indicates

that during training of **CDVAE-GAN**_{MCC}, although not in our original expectations, the discriminator not only benefits the decoders, but also indirectly guides the latent codes to be better disentangled.

As for CLS, we first observe that including CLS in **CDVAE** improves the DEM score when using MCC yet degrades when using SP. Although this somewhat makes the effectiveness of CLS inconvincing, we note that **CDVAE-CLS** [SP] and **CDVAE-CLS** [MCC] have nearly identical DEM scores. This interesting finding shows that the CLS forces the encoders to encode different features into similar contents. On the other hand, including CLS in **CDVAE-GAN** models boosts the DEM scores of cross gender pairs, which confirms that CLS can help the encoders eliminate speaker independent information, such as gender.

Finally, we compare the results of similarity tests of **CDVAE** [MCC], **CDVAE-GAN**_{MCC}, and **CDVAE-CLS-GAN**_{MCC} in Figure 8 and the DEM results in Table IV. **CDVAE-CLS-GAN**_{MCC} achieves the highest similarity scores in Figure 8 and gives the highest DEM scores in Table IV. The result verifies the positive correlation between the conversion performance and the degree of disentanglement of the latent codes.

VI. CONCLUSIONS

In this paper, we have extended the cross-domain VAE based VC framework by integrating GANs and CLS into the training phase. The GAN objective was used to better approximate the distribution of real speech signals. The CLS, on the other hand, was applied to the latent code as an explicit constraint to eliminate speaker-dependent factors. Objective and subjective evaluations confirmed the effectiveness of the GAN and CLS objectives. We have also investigated the correlation between the degree of disentanglement and the conversion performance. A novel evaluation metric, DEM, that measures the degree of disentanglement in VC was derived. Experimental results confirmed a positive correlation between the degree of disentanglement and the conversion performance.

In the future, we will exploit more acoustic features in the **CDVAE** system, including rawer features, such as the magnitude spectrum, and hand-crafted features, such as line-spectral pairs. An effective algorithm that can optimally determine the latent space dimension is also worthy of study. Finally, it is worthwhile to generalize this disentanglement framework to extract speaker-invariant latent representation from unknown source speakers in order to achieve many-to-one VC.

We have made the source code publicly accessible so that readers can reproduce our results.²

REFERENCES

- [1] S. H. Mohammadi and A. Kain, "An overview of voice conversion systems," *Speech Commun.*, vol. 88, pp. 65–82, 2017.
- [2] D. Childers, B. Yegnanarayana, and K. Wu, "Voice conversion: Factors responsible for quality," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, vol. 10, 1985, pp. 748–751.
- [3] S. Aryal, D. Felps, and R. Gutierrez-Osuna, "Foreign accent conversion through voice morphing," in *Proc. Interspeech*, 2013, pp. 3077–3081.
- [4] A. Kain and M. W. Macon, "Spectral voice conversion for text-to-speech synthesis," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, vol. 1, 1998, pp. 285–288.
- [5] J. Latorre, V. Wan, and K. Yanagisawa, "Voice expression conversion with factorised hmm-tts models," in *Proc. Interspeech*, 2014, pp. 1514–1518.
- [6] 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.
- [7] H. Doi, T. Toda, K. Nakamura, H. Saruwatari, and K. Shikano, "Alaryngeal speech enhancement based on one-to-many eigenvoice conversion," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 22, no. 1, pp. 172–183, Jan. 2014.
- [8] K. Tanaka, T. Toda, G. Neubig, S. Sakti, and S. Nakamura, "A hybrid approach to electrolaryngeal speech enhancement based on spectral subtraction and statistical voice conversion," in *Proc. Interspeech*, 2013, pp. 3067–3071.
- [9] Y. Stylianou, O. Cappe, and E. Moulines, "Continuous probabilistic transform for voice conversion," *IEEE Trans. Speech Audio Process.*, vol. 6, no. 2, pp. 131–142, Mar. 1998.
- [10] 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, Nov. 2007.
- [11] D. Sundermann and H. Ney, "Vtl-n-based voice conversion," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2003, pp. 556–559.
- [12] D. Erro, A. Moreno, and A. Bonafonte, "Voice conversion based on weighted frequency warping," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 18, no. 5, pp. 922–931, Jul. 2010.
- [13] E. Godoy, O. Rossec, and T. Chonavel, "Voice conversion using dynamic frequency warping with amplitude scaling, for parallel or nonparallel corpora," *IEEE/ACM Trans. Audio, Speech and Language Process.*, vol. 20, no. 4, pp. 1313–1323, May 2012.
- [14] R. Takashima, T. Takiguchi, and Y. Ariki, "Exemplar-based voice conversion in noisy environment," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2012, pp. 313–317.
- [15] Z. Wu, T. Virtanen, E. S. Chng, and H. Li, "Exemplar-based sparse representation with residual compensation for voice conversion," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 22, no. 10, pp. 1506–1521, Oct. 2014.
- [16] Y.-C. Wu, H.-T. Hwang, C.-C. Hsu, Y. Tsao, and H.-M. Wang, "Locally linear embedding for exemplar-based spectral conversion," in *Proc. Interspeech*, 2016, pp. 1652–1656.
- [17] B. Sisman, H. Li, and K. C. Tan, "Sparse representation of phonetic features for voice conversion with and without parallel data," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2017, pp. 677–684.
- [18] X. Tian, S. W. Lee, Z. Wu, E. S. Chng, and H. Li, "An exemplar-based approach to frequency warping for voice conversion," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 25, no. 10, pp. 1863–1876, Oct. 2017.
- [19] M. Narendranath, H. A. Murthy, S. Rajendran, and B. Yegnanarayana, "Transformation of formants for voice conversion using artificial neural networks," *Speech Commun.*, vol. 16, no. 2, pp. 207–216, 1995.
- [20] 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, Jul. 2010.
- [21] L. H. Chen, Z. H. Ling, L. J. Liu, and L. R. Dai, "Voice conversion using deep neural networks with layer-wise generative training," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 22, no. 12, pp. 1859–1872, Dec. 2014.
- [22] L. Sun, S. Kang, K. Li, and H. Meng, "Voice conversion using deep bidirectional long short-term memory based recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2015, pp. 4869–4873.
- [23] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," 2016, *arXiv:1609.03499*.
- [24] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, "Speaker-dependent wavenet vocoder," in *Proc. Interspeech*, 2017, pp. 1118–1122.
- [25] T. Hayashi, A. Tamamori, K. Kobayashi, K. Takeda, and T. Toda, "An investigation of multi-speaker training for wavenet vocoder," in *Proc. IEEE Autom. Speech Recognit. Understanding Workshop*, 2017, pp. 712–718.
- [26] K. Kobayashi, T. Hayashi, A. Tamamori, and T. Toda, "Statistical voice conversion with wavenet-based waveform generation," in *Proc. Interspeech*, 2017, pp. 1138–1142.

²<https://github.com/unilight/cdvae-vc>

- [27] B. Sisman, M. Zhang, and H. Li, "A voice conversion framework with tandem feature sparse representation and speaker-adapted "wavenet" vocoder," in *Proc. Interspeech*, 2018, pp. 1978–1982.
- [28] L.-J. Liu, Z.-H. Ling, Y. Jiang, M. Zhou, and L.-R. Dai, "Wavenet vocoder with limited training data for voice conversion," in *Proc. Interspeech*, 2018, pp. 1983–1987.
- [29] P. L. Tobing, T. Hayashi, Y.-C. Wu, K. Kobayashi, and T. Toda, "An evaluation of deep spectral mappings and wavenet vocoder for voice conversion," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2018, pp. 297–303.
- [30] W.-C. Huang *et al.*, "Refined wavenet vocoder for variational autoencoder based voice conversion," in *Proc. 27th Eur. Signal Process. Conf.*, 2019.
- [31] B. Sisman, M. Zhang, S. Sakti, H. Li, and S. Nakamura, "Adaptive wavenet vocoder for residual compensation in gan-based voice conversion," in *Proc. IEEE Spoken Lang. Technol. Workshop*, 2018, pp. 282–289.
- [32] Y.-H. Peng, H.-T. Hwang, Y. Wu, Y. Tsao, and H.-M. Wang, "Exemplar-based spectral detail compensation for voice conversion," in *Proc. Interspeech*, 2018, pp. 486–490.
- [33] J. Lorenzo-Trueba *et al.*, "The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods," in *Proc. Odyssey Speaker Lang. Recognit. Workshop*, 2018, pp. 195–202.
- [34] D. Erro, A. Moreno, and A. Bonafonte, "Inca algorithm for training voice conversion systems from nonparallel corpora," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 18, no. 5, pp. 944–953, Jul. 2010.
- [35] 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. IEEE Int. Conf. Multimedia Expo*, 2016, pp. 1–6.
- [36] 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. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 5274–5278.
- [37] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2013, *arXiv:1312.6114*.
- [38] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [39] 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. APISPA ASC*, 2016, pp. 1–6.
- [40] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "Acvae-vc: Non-parallel voice conversion with auxiliary classifier variational auto-encoder," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 9, pp. 1432–1443, Sep. 2019.
- [41] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, "Voice conversion from unaligned corpora using variational autoencoding wasserstein generative adversarial networks," in *Proc. Interspeech*, 2017, pp. 3364–3368.
- [42] T. Kaneko and H. Kameoka, "Parallel-data-free voice conversion using cycle-consistent adversarial networks," 2017, *arXiv:1711.11293*.
- [43] T. Kaneko, H. Kameoka, K. Tanaka, and N. Hojo, "Cyclegan-vc2: Improved cyclegan-based non-parallel voice conversion," in *Proc. IEEE Int. Conf. Acoust., Speech and Signal Process.*, 2019, pp. 6820–6824.
- [44] 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. IEEE Spoken Lang. Technol. Workshop*, 2018, pp. 266–273.
- [45] J. C. Chou, C. C. Yeh, H. Y. Lee, and L. S. Lee, "Multi-target voice conversion without parallel data by adversarially learning disentangled audio representations," in *Proc. Interspeech*, 2018, pp. 501–505.
- [46] J. Lorenzo-Trueba, F. Fang, X. Wang, I. Echizen, J. Yamagishi, and T. Kinnunen, "Can we steal your vocal identity from the internet?: Initial investigation of cloning obamas voice using gan, wavenet and low-quality found data," in *Proc. Odyssey Speaker Lang. Recognit. Workshop*, 2018, pp. 240–247.
- [47] W.-N. Hsu, Y. Zhang, and J. Glass, "Unsupervised learning of disentangled and interpretable representations from sequential data," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 1878–1889.
- [48] W.-N. Hsu *et al.*, "Hierarchical generative modeling for controllable speech synthesis," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [49] W.-N. Hsu *et al.*, "Disentangling correlated speaker and noise for speech synthesis via data augmentation and adversarial factorization," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2019, pp. 5901–5905.
- [50] Y.-A. Chung and J. Glass, "Speech2vec: A sequence-to-sequence framework for learning word embeddings from speech," in *Proc. Interspeech*, 2018, pp. 811–815.
- [51] J. Chorowski, R. J. Weiss, S. Bengio, and A. V. D. Oord, "Unsupervised speech representation learning using wavenet autoencoders," 2019, *arXiv:1901.08810*.
- [52] A. Achille and S. Soatto, "Emergence of invariance and disentanglement in deep representations," *J. Mach. Learn. Res.*, vol. 19, no. 1, pp. 1947–1980, 2018.
- [53] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, "AutoVC: Zero-shot voice style transfer with only autoencoder loss," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 5210–5219.
- [54] W.-C. Huang, H.-T. Hwang, Y.-H. Peng, Y. Tsao, and H.-M. Wang, "Voice conversion based on cross-domain features using variational auto encoders," in *Proc. 11th Int. Symp. Chin. Spoken Lang. Process.*, 2018, pp. 51–55.
- [55] H. Kawahara, I. Masuda-Katsuse, and A. De Cheveigne, "Restructuring speech representations using a pitch-adaptive time–frequency smoothing and an instantaneous-frequency-based f0 extraction: Possible role of a repetitive structure in sounds," *Speech Commun.*, vol. 27, no. 3–4, pp. 187–207, 1999.
- [56] T. Fukada, K. Tokuda, T. Kobayashi, and S. Imai, "An adaptive algorithm for mel-cepstral analysis of speech," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 1992, pp. 137–140.
- [57] Y. Xu, J. Du, Z. Huang, L.-R. Dai, and C.-H. Lee, "Multi-objective learning and mask-based post-processing for deep neural network based speech enhancement," in *Proc. Interspeech*, 2015, pp. 1508–1512.
- [58] R. Caruana, "Multitask learning," *Mach. Learn.*, vol. 28, no. 1, pp. 41–75, 1997.
- [59] Y. Ganin *et al.*, "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [60] Z. Meng *et al.*, "Speaker-invariant training via adversarial learning," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 5969–5973.
- [61] S. Sun, B. Zhang, L. Xie, and Y. Zhang, "An unsupervised deep domain adaptation approach for robust speech recognition," *Neurocomputing*, vol. 257, pp. 79–87, 2017.
- [62] S. Sun, C.-F. Yeh, M.-Y. Hwang, M. Ostendorf, and L. Xie, "Domain adversarial training for accented speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 4854–4858.
- [63] O. Ocal, O. H. Elibol, G. Keskin, C. Stephenson, A. Thomas, and K. Ramchandran, "Adversarially trained autoencoders for parallel-data-free voice conversion," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2019, pp. 2777–2781.
- [64] E. Nachmani and L. Wolf, "Unsupervised singing voice conversion," in *Proc. Interspeech*, 2019, pp. 2583–2587.
- [65] T. Q. Chen, X. Li, R. B. Grosse, and D. K. Duvenaud, "Isolating sources of disentanglement in variational autoencoders," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 2610–2620.
- [66] C. Eastwood and C. K. I. Williams, "A framework for the quantitative evaluation of disentangled representations," in *Int. Conf. Learn. Representations*, 2018.
- [67] J. chieh Chou and H.-Y. Lee, "One-shot voice conversion by separating speaker and content representations with instance normalization," in *Proc. Interspeech*, 2019, pp. 664–668.
- [68] W.-C. Huang *et al.*, "Investigation of f0 conditioning and fully convolutional networks in variational autoencoder based voice conversion," in *Proc. Interspeech*, 2019, pp. 709–713.
- [69] A. B. L. Larsen, S. K. Snderby, H. Larochelle, and O. Winther, "Autoencoding beyond pixels using a learned similarity metric," in *Proc. 33rd Int. Conf. Mach. Learn.*, vol. 48, 2016, pp. 1558–1566.
- [70] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein gan," 2017, *arXiv:1701.07875*.
- [71] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5767–5777.
- [72] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, 2015, pp. 3431–3440.
- [73] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, *arXiv:1607.06450*.
- [74] H. Siln, E. Hel, J. Nurminen, and M. Gabbouj, "Ways to implement global variance in statistical speech synthesis," in *Proc. Interspeech*, 2012, pp. 1436–1439.
- [75] S. Takamichi, T. Toda, A. W. Black, G. Neubig, S. Sakti, and S. Nakamura, "Postfilters to modify the modulation spectrum for statistical parametric speech synthesis," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 24, no. 4, pp. 755–767, Apr. 2016.

- [76] T. Kaneko, H. Kameoka, N. Hojo, Y. Ijima, K. Hiratsmu, and K. Kashino, "Generative adversarial network-based postfilter for statistical parametric speech synthesis," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2017, pp. 4910–4914.



Wen-Chin Huang (Student Member, IEEE) received the B.S. degree in computer science and information engineering from National Taiwan University, Taipei, Taiwan, in 2018. He is currently working toward the M.S. degree with Nagoya University, Nagoya, Japan. He was a Research Assistant with the Institute of Information Science, Academia Sinica, Taipei, Taiwan, from 2017 to 2019. His research interests include voice conversion, speech synthesis and deep learning. He was the recipient of the Best Student Paper Award in ISCSLP 2018.



Hao Luo received the B.S. degree in electrical engineering from National Taiwan University, Taipei, Taiwan. He is currently working toward the M.S. degree in electrical engineering from National Taiwan University, Taipei, Taiwan. He is an Intern in Institute of Information Science, Academia Sinica, Taipei, Taiwan. His research interest is voice conversion.



Hsin-Te Hwang received the Ph.D degree from the Institute of Communications Engineering, National Chiao Tung University, Hsinchu, Taiwan, in 2017. From 2010 to 2017, he was a Research Assistant with the Institute of Information Science, Academia Sinica, Taipei, Taiwan. He was a Postdoctoral Fellow with the Institute of Information Science, Academia Sinica, Taipei, Taiwan, from 2017 to 2018. He joined the Intelligo Technology Inc. Taiwan Branch in 2019, where he is currently a Senior Engineer. His research interests include speech signal processing, particularly,

voice conversion, speech enhancement, and speech synthesis.



Chen-Chou Lo received the M.S. degree in electrical engineering from National Central University, Taoyuan, Taiwan. He is currently working toward the Ph.D. degree with the Department of Electrical Engineering, Katholieke Universiteit Leuven, Leuven, Belgium. From 2018 to 2019, he was a Research Assistant with the Institute of Information Science, Academia Sinica, Taipei, Taiwan. He joined the research group EAVISE with Katholieke Universiteit Leuven in October 2019. His research interests are in audiovisual signal processing, speech enhancement,

and converted speech assessments.



Yu-Huai Peng received the M.S. degree from the Department of Electrical Engineering, National Tsing Hua University, Hsinchu, Taiwan. He is a Research Assistant with the Institute of Information Science, Academia Sinica, Taipei, Taiwan. His research interests include speech signal processing, particularly, voice conversion, and speaker diarization.



Yu Tsao (Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from National Taiwan University, Taipei, Taiwan, in 1999 and 2001, respectively, and the Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2008. From 2009 to 2011, he was a Researcher with the National Institute of Information and Communications Technology, Tokyo, Japan, where he engaged in research and product development in automatic speech recognition for multilingual speech-to-speech translation. He is currently an Associate Research Fellow with the Research Center for Information Technology Innovation, Academia Sinica, Taipei, Taiwan. His research interests include speech and speaker recognition, acoustic and language modeling, audio coding, and biosignal processing. He is currently an Associate Editor for the *IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING* and *IEICE transactions on Information and Systems*. He was the recipient of the Academia Sinica Career Development Award in 2017, National Innovation Award in 2018 and 2019, and Outstanding Elite Award, Chung Hwa Rotary Educational Foundation 2019–2020.

He is currently an Associate Editor for the *IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING* and *IEICE transactions on Information and Systems*. He was the recipient of the Academia Sinica Career Development Award in 2017, National Innovation Award in 2018 and 2019, and Outstanding Elite Award, Chung Hwa Rotary Educational Foundation 2019–2020.



Hsin-Min Wang (Senior Member, IEEE) received the B.S. and Ph.D. degrees in electrical engineering from National Taiwan University, Taipei, Taiwan, in 1989 and 1995, respectively. In October 1995, he joined the Institute of Information Science, Academia Sinica, Taipei, Taiwan, where he is currently a Research Fellow. He also holds a joint appointment as a Professor with the Department of Computer Science and Information Engineering, National Cheng Kung University. He currently serves as an Editorial Board Member of *IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH AND LANGUAGE PROCESSING* and *APSIPA Transactions on Signal and Information Processing*. His major research interests include spoken language processing, natural language processing, multimedia information retrieval, machine learning and pattern recognition. He is a member of the International Speech Communication Association and ACM. He was the General Co-Chair of ISCSLP 2016 and ISCSLP 2018 and a Technical Co-Chair of ISCSLP 2010, O-COCOSDA 2011, APSIPAASC 2013, ISMIR 2014, and ASRU 2019. He was an APSIPA Distinguished Lecturer for 2014–2015. He was the recipient of Chinese Institute of Engineers Technical Paper Award in 1995 and the ACM Multimedia Grand Challenge First Prize in 2012.

He is currently an Associate Editor for the *IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH AND LANGUAGE PROCESSING* and *APSIPA Transactions on Signal and Information Processing*. His major research interests include spoken language processing, natural language processing, multimedia information retrieval, machine learning and pattern recognition. He is a member of the International Speech Communication Association and ACM. He was the General Co-Chair of ISCSLP 2016 and ISCSLP 2018 and a Technical Co-Chair of ISCSLP 2010, O-COCOSDA 2011, APSIPAASC 2013, ISMIR 2014, and ASRU 2019. He was an APSIPA Distinguished Lecturer for 2014–2015. He was the recipient of Chinese Institute of Engineers Technical Paper Award in 1995 and the ACM Multimedia Grand Challenge First Prize in 2012.