

DISENTANGLED SPEECH REPRESENTATION LEARNING FOR ONE-SHOT CROSS-LINGUAL VOICE CONVERSION USING β -VAE

Hui Lu^{1,2}, Disong Wang¹, Xixin Wu^{1,2}, Zhiyong Wu^{1,3}, Xunying Liu¹, Helen Meng^{1,2,3}

¹The Chinese University of Hong Kong, Hong Kong SAR, China

²Centre for Perceptual and Interactive Intelligence, CUHK, Hong Kong SAR, China

³Tsinghua-CUHK Joint Research Center for Media Sciences, Technologies and Systems, Shenzhen International Graduate School, Tsinghua University, Shenzhen, China

ABSTRACT

We propose an unsupervised learning method to disentangle speech into content representation and speaker identity representation. We apply this method to the challenging one-shot cross-lingual voice conversion task to demonstrate the effectiveness of the disentanglement. Inspired by β -VAE, we introduce a learning objective that balances between the information captured by the content and speaker representations. In addition, the inductive biases from the architectural design and the training dataset further encourage the desired disentanglement. Both objective and subjective evaluations show the effectiveness of the proposed method in speech disentanglement and in one-shot cross-lingual voice conversion.

Index Terms— Speech disentanglement, Voice conversion, cross-lingual, one-shot, unsupervised learning

1. INTRODUCTION

Voice conversion (VC) is a technique that converts a source speaker’s speech to make it sound as though it was uttered by the designated target speaker, while the spoken content remains unchanged [1]. Cross-lingual VC refers to the scenario where the source and the target speakers do not speak the same language [2]. Cross-lingual VC is more challenging than intra-lingual VC because generally only a monolingual corpus is available from the target speaker, which presents the problem of training-inference mismatch. Furthermore, one-shot cross-lingual VC is even more challenging, because there is only a single utterance from the target speaker, who speaks a language that is different from the source speaker. We focus on one-shot cross-lingual VC in this paper.

Speech generation embeds different information elements into the signal: content, speaker identity, emotion, language, etc., where the first two elements tend to be more prominent. Generally speaking, language can be considered as part of content. We may regard that a universal phoneme set can cover pronunciation patterns of all languages, and encode different languages in the spoken content. Following this rationale, we can solve the one-shot cross-lingual VC by disen-

tangling speech into a general content representation and a speaker representation. Then the content representation from the source speech and the speaker representation from the target speech can be combined to generate the speech with the source content and the target speaker identity.

Many methods have been proposed to disentangle the speech into content and speaker representations for one-shot VC. Commonly adopted architectures are the auto-encoder (AE) and variational auto-encoder (VAE) [3], with encoder(s) to extract a frame-level feature sequence and a single vector separately from the speech, aiming for the former to encode spoken content, and the latter to encode speaker identity. This design is based on the fact that the speaker identity is a sequence-level element while the content is a frame-level element. However, as we will discuss in Sec 3.1, this alone is not enough to ensure the disentanglement of the content and speaker representations. Many methods have been proposed to further regularize the representation learning to facilitate disentanglement.

AutoVC [4] proposes to apply down-sampling and dimension restriction to the content representation to remove the speaker identity. VQVC and VQVC+ [5, 6] apply vector quantization (VQ) to the content representation to eliminate the speaker information. However, through explicit down-sampling or VQ operations, these methods may hurt the fine-grained temporal information and cause conversion quality degradation. The mutual information (MI) between the content and the speaker representations has also been imposed as a learning objective [7, 8] to encourage them to be statistically independent, but the estimation of MI is difficult and may make the training process more complex. Other methods apply instance normalization [9] and activation regularization [10] to the content representation, which can induce less information loss and are easier to implement.

Recently, β -VAE [11, 12] has been proposed as a variant of VAEs for better latent variable disentanglement. With an extra weight parameter β being imposed on the Kullback-Leibler (KL) divergence term, β -VAE restricts the channel capacity of the latent variable to encode the information of

the input data. The original β -VAE is more suitable for static data such as images. In this paper, we propose a variant of β -VAE that is specifically for disentanglement of sequential data into time-variant and time-invariant representations. Different from β -VAE which only has one encoder, we adopt two encoders that are respectively designed for content and speaker representation learning. In addition, two weight parameters β_c and β_s are imposed on the KL divergence terms for content and speaker representations, respectively. We show theoretically that the KL divergence term is an upper bound of the MI between the latent variable and the input data. Thus, β_c and β_s function as two gates to restrict the amounts of information of the data that can be captured by the content and speaker representations, respectively. With proper β_c and β_s imposed, the information captured by the two representations are complementary to each other, while the content information and speaker identity information are precisely allocated to them, respectively.

Compared to existing one-shot VC methods, the proposed method imposes much simpler restrictions on the latent variable (only two weight parameters on the KL divergence terms) which is easier to implement and train. We do not apply intentionally bottleneck operations such as dimension reduction or down-sampling to the latent variable, thus less information loss is induced. Furthermore, the proposed method can be explained by information theory, which is theoretically more solid.

2. RELATED WORK

Traditional cross-lingual VC methods rely extensively on the phonetic posteriorgrams (PPGs) [2, 13, 14, 15], which is a speaker-independent content representation explicitly extracted via a speaker-independent automatic speech recognition (SI-ASR) model. VC is achieved by first extracting PPGs from the source speech, then converting the PPGs into the target speech with a synthesis model, which can be trained only on the target speaker’s speech corpus. Though PPGs can be considered as speaker-independent, they are language-dependent, which means that PPGs of one language may not be a good content representation of another language. Besides, since the PPGs extractor is fixed during the training of the synthesis model, it cannot be optimized for all interested languages. In contrast, the proposed method learns to extract the content representation for all interested languages jointly with other parts of the model, thus can in theory learn a more generalized content representation.

For unsupervised disentangled speech representation learning, FHVAE [16] propose hierarchical prior distributions for VAE to encourage elements of different time scales to be factorized, while the sequence-level representation is further regularized by a discriminative objective. Based on FHVAE, a method is proposed for one-shot cross-lingual VC [17]. The method proposed in this paper is different in terms

of the restriction imposed, as we emphasize the effects of two weights on the KL divergence terms and no other learning objective is introduced. Another recently proposed method for one-shot intra-lingual VC [18] also stresses the effects of imposing weights on two KL divergence terms to encourage disentanglement. However, the proposed method is different from [18] in the following ways: First, the prior distribution for the content representation of [18] is a trainable autoregressive one, while ours is an fixed isotropic Gaussian. Second, as we will show in Section 3.2, the current paper presents a different derivation (especially for the MI). Besides, it is claimed in [18] that the right choice of the ratio between the two weights can yield the desired disentanglement, while we argue that the absolute values of both weights can also be important for good disentanglement. While some methods [7, 8] estimate the MI between different representations as a loss term, the derivation of the MI in this paper only serves as an explanation of the proposed method, we do not need to estimate it numerically.

3. METHOD

Let \mathbf{x} be the acoustic feature extracted from speech, and the objective is to find latent variables \mathbf{z}_c and \mathbf{z}_s , such that \mathbf{z}_c encodes only spoken content information, while \mathbf{z}_s exclusively contains speaker identity information. With the assumption that \mathbf{z}_c and \mathbf{z}_s are statistically independent, we can derive the objective function as shown in Eqn. (1), following the formalization of VAEs. In Eqn. (1), $q_\phi(\mathbf{z}_c|\mathbf{x})$ and $q_\phi(\mathbf{z}_s|\mathbf{x})$, are posterior distributions of \mathbf{z}_c and \mathbf{z}_s given \mathbf{x} , respectively, parameterized by ϕ . $p_\theta(\mathbf{x}|\mathbf{z}_c, \mathbf{z}_s)$ denotes the conditional distribution of \mathbf{x} given \mathbf{z}_s and \mathbf{z}_c and is parameterized by θ . $p(\mathbf{z}_c)$ and $p(\mathbf{z}_s)$ are prior distributions of \mathbf{z}_c and \mathbf{z}_s , respectively, and are defined as isotropic Gaussians in this paper, $p_d(\mathbf{x})$ is the data distribution of \mathbf{x} . For simplicity, we refer to the three terms in Eqn. (1) as the reconstruction loss, content KL, and speaker KL, respectively.

$$\begin{aligned} \mathcal{L}_{\text{vanilla}} = & -\mathbb{E}_{p_d(\mathbf{x})q_\phi(\mathbf{z}_c|\mathbf{x})q_\phi(\mathbf{z}_s|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z}_c, \mathbf{z}_s)] \\ & + \mathbb{E}_{p_d(\mathbf{x})} [\text{KL} [q_\phi(\mathbf{z}_c|\mathbf{x}) \parallel p(\mathbf{z}_c)]] \\ & + \mathbb{E}_{p_d(\mathbf{x})} [\text{KL} [q_\phi(\mathbf{z}_s|\mathbf{x}) \parallel p(\mathbf{z}_s)]] \end{aligned} \quad (1)$$

The general unsupervised learning model defined above without any other inductive biases cannot ensure the disentanglement, as is commonly believed [19]. To achieve the disentanglement of content and speaker representations with the above formalization, our solutions are as follows: We first introduce the architectural design that facilitate the separation of time-variant and time-invariant elements in Section 3.1; In Section 3.2 we present the learning objective that can encourage \mathbf{z}_c and \mathbf{z}_s to be complementary, while exactly the content information and speaker identity information can be assigned separately to them. The effect of the training dataset is also emphasized in Section 3.3.

3.1. Architectural Design

A basic observation about speech is that the content varies constantly across an utterance while the speaker identity remains the same for the whole utterance. That is, the content is a time-variant feature while speaker identity is a time-invariant one. This informs us about the structures of the content and speaker representations. Suppose that the acoustic feature \mathbf{x} is of the shape $\mathbb{R}^{T \times D_x}$, where T and D_x denote the number of frames and the dimension of \mathbf{x} , respectively. Then \mathbf{z}_c and \mathbf{z}_s should have the corresponding shapes of $\mathbb{R}^{T \times D_c}$ and \mathbb{R}^{D_s} to separately capture the time-variant content information and time-invariant speaker identity information, where D_c and D_s are the dimensions of \mathbf{z}_c and \mathbf{z}_s , respectively. A content encoder transforming the input frame-wisely and a speaker encoder aggregating the whole input sequence into a single vector should thus be adopted. Many methods introduced in Section 1 have adopted similar architectures.

However, this architecture alone cannot guarantee the disentanglement of content and speaker representations. The reasons are two-fold: First, the architecture cannot ensure that the two representations learned are complementary to each other. Intuitively, the two representations will strive to capture as much information about the input data as possible to reduce the reconstruction loss. Thus, it is possible that the two representations have overlap in the captured information. Second, even if we can restrict the two representations to be complementary (e.g., via bottleneck), then we can reasonably assume that \mathbf{z}_c and \mathbf{z}_s respectively captures complementary time-variant and time-invariant information of speech, but it is not guaranteed that the learned time-variant and time-invariant elements are exactly content and speaker identity. The reason is that there are many combinations of time-variant and time-invariant information elements other than the targeted one. For example, background noise can (mostly) be considered as a time-invariant element, against the time-variant element of clean speech, but the architecture alone cannot distinguish this pair of elements from the content-speaker identity pair.

3.2. Learning Objective

Aside from the architectural design to facilitate the separation of time-variant and time-invariant feature into \mathbf{z}_c and \mathbf{z}_s , we also need to ensure that: 1) the information captured by \mathbf{z}_c and \mathbf{z}_s are complementary and 2) the complementary elements captured are content and speaker identity instead of other element pairs – and the key is to restrict the amounts of information of \mathbf{x} that can be captured by \mathbf{z}_c and \mathbf{z}_s . To solve issue 1), we can restrict the overall amount of information captured by \mathbf{z}_c and \mathbf{z}_s together to a low level. With no redundant information encoded, the two representations will be compact and contain complementary information. To further tackle issue 2), it is important to restrict the relative amounts

of information captured in the two representations, such that \mathbf{z}_c and \mathbf{z}_s captures exactly the content part and the speaker identity part of information in \mathbf{x} , respectively.

The information captured by a latent variable can be quantified as its MI with the data. Following the typical practice by approximating the joint distribution $p(\mathbf{x}, \mathbf{z})$ with its variational version $q_\phi(\mathbf{z}|\mathbf{x})p_d(\mathbf{x})$ [20, 21], we can derive the variational MI between \mathbf{x} and \mathbf{z} as shown in Eqn. (2), where \mathbf{z} is a general variable name covering \mathbf{z}_c and \mathbf{z}_s .

$$\begin{aligned} \mathbb{I}_v(\mathbf{x}, \mathbf{z}) &= \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})p_d(\mathbf{x})} \left[\log \frac{q_\phi(\mathbf{z}|\mathbf{x})p_d(\mathbf{x})}{p_d(\mathbf{x})q_\phi(\mathbf{z})} \right] \\ &= \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})p_d(\mathbf{x})} \left[\log \frac{q_\phi(\mathbf{z}|\mathbf{x})}{p(\mathbf{z})} + \log \frac{p(\mathbf{z})}{q_\phi(\mathbf{z})} \right] \\ &= \mathbb{E}_{p_d(\mathbf{x})} [\text{KL} [q_\phi(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})]] \\ &\quad - \text{KL} [q_\phi(\mathbf{z}) \parallel p(\mathbf{z})] \end{aligned} \quad (2)$$

Here $q_\phi(\mathbf{z})$ is the marginal distribution of \mathbf{z} and is defined as $q_\phi(\mathbf{z}) = \int_{\mathbf{x}} q_\phi(\mathbf{z}|\mathbf{x})p_d(\mathbf{x})d\mathbf{x}$. The results in Eqn. (2) is different from the MI derived in [18], which claim that $\mathbb{I}_v(\mathbf{x}, \mathbf{z}) = \mathbb{E}_{p_d(\mathbf{x})} [\text{KL} [q_\phi(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})]]$. This derivation is resulted from approximating the marginal distribution of \mathbf{z} with its prior $p(\mathbf{z})$, which we consider is generally not applicable.

Eqn. (2) denotes that $\mathbb{E}_{p_d(\mathbf{x})} [\text{KL} [q_\phi(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})]]$ is an upper-bound of $\mathbb{I}_v(\mathbf{x}, \mathbf{z})$. In this sense, an approximate method to restrict the amount of information of x captured by the latent variable \mathbf{z} is to restrict the KL divergence term $\mathbb{E}_{p_d(\mathbf{x})} [\text{KL} [q_\phi(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z})]]$. Though the term $\text{KL} [q_\phi(\mathbf{z}) \parallel p(\mathbf{z})]$ is also penalized as a side-effect, which encourages the marginal distribution of \mathbf{z} to be more like an isotropic Gaussian and does not other harm. Motivated by this, we introduce two weight parameters β_c and β_s to the content KL and speaker KL terms, respectively, with the goal to restrict the information that can be captured by \mathbf{z}_c and \mathbf{z}_s . The restricted objective function is shown in Eqn. (3), while Eqn. (4) explicitly denotes the relationship between two KL terms and their corresponding MI terms.

$$\begin{aligned} \mathcal{L}_\beta &= -\mathbb{E}_{p_d(\mathbf{x})q_\phi(\mathbf{z}_c|\mathbf{x})q_\phi(\mathbf{z}_s|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z}_c, \mathbf{z}_s)] \\ &\quad + \beta_c \cdot \mathbb{E}_{p_d(\mathbf{x})} [\text{KL} [q_\phi(\mathbf{z}_c|\mathbf{x}) \parallel p(\mathbf{z}_c)]] \\ &\quad + \beta_s \cdot \mathbb{E}_{p_d(\mathbf{x})} [\text{KL} [q_\phi(\mathbf{z}_s|\mathbf{x}) \parallel p(\mathbf{z}_s)]] \end{aligned} \quad (3)$$

$$\begin{aligned} &= -\mathbb{E}_{p_d(\mathbf{x})q_\phi(\mathbf{z}_c|\mathbf{x})q_\phi(\mathbf{z}_s|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z}_c, \mathbf{z}_s)] \\ &\quad + \beta_c \cdot [\mathbb{I}_v(\mathbf{x}, \mathbf{z}_c) + \text{KL} [q_\phi(\mathbf{z}_c) \parallel p(\mathbf{z}_c)]] \\ &\quad + \beta_s \cdot [\mathbb{I}_v(\mathbf{x}, \mathbf{z}_s) + \text{KL} [q_\phi(\mathbf{z}_s) \parallel p(\mathbf{z}_s)]] \end{aligned} \quad (4)$$

While the reconstruction loss in Eqn. (3) encourages both \mathbf{z}_c and \mathbf{z}_s to capture as much information about \mathbf{x} as possible, the weight parameters β_c and β_s act as two gates to control the amounts of information of \mathbf{x} going through \mathbf{z}_c and \mathbf{z}_s , as indicated by Eqn. (4). The learning objective (3) can help solve the problems of the architectural design in the following ways: By choosing the proper absolute values of both β_c

and β_s , the information captured by \mathbf{z}_c and \mathbf{z}_s together will be compact and thus the information contained in each latent variable will be complementary. On the other hand, by tuning relative values of β_c and β_s , the amounts of information in \mathbf{x} allocated to \mathbf{z}_c and \mathbf{z}_s will change accordingly. In this case, we can find the optimal parameters pair that yields exactly the separation of content and speaker identity.

Note that if we let $\beta_c = \beta_s$ and combine \mathbf{z}_c and \mathbf{z}_s together, Eqn. (3) will become exactly the learning objective of the standard β -VAE [11]. In this sense, Eqn. (3) can be considered as a generalized form of β -VAE.

3.3. Dataset bias

The dataset used to train the model is also a very important inductive bias to facilitate the learning of disentangled representations. One consideration is that the two information elements to be disentangled should vary independently in the dataset. Besides, other elements of variation that may interfere the interested elements should not appear in the dataset. For example, when disentangling the content and speaker identity elements, too much variation in emotions may disturb the learning of speaker identity element since they are both sequence-level elements. In our case, we combine two monolingual corpora as the training dataset, in which the speaker and general content are the main variations. More details are given in Section 5.1.

4. IMPLEMENTATION

The proposed model consists of a speaker encoder, a content encoder, and a decoder, to respectively model $q_\phi(\mathbf{z}_s|\mathbf{x})$, $q_\phi(\mathbf{z}_c|\mathbf{x})$, and $p_\theta(\mathbf{x}|\mathbf{z}_c, \mathbf{z}_s)$ defined in Section 3. The structure of each component as well as other implementation details are described as follows.

Speaker encoder: The speaker encoder consists of one fully-connected (FC) layer with 256 hidden units, whose output is activated by ReLU [22] and fed into 4 down-sampling convolutional blocks. Each down-sampling convolutional block includes two 1D convolution (Conv) layers with 256 filters and ReLU activations, while the output of the second Conv layer is down-sampled by a factor of 2 using average pooling along the time axis. A residual connection is adopted in each convolutional block to connect the input and the output. The kernel sizes of Conv layers in the 4 convolutional blocks are 3, 3, 5 and 5, respectively. The output of convolutional blocks is further aggregated through global average-pooling and fed into a fully connected layer to get the 128-dimensional mean and variance vector of the speaker posterior.

Content encoder: The content encoder consists of two 1D Conv layers each with 256 hidden units and a kernel size of 3. Dropout [23] with the drop rate of 0.2 is embedded after each Conv layer. Two self-attention blocks [24] with

256 hidden units, 4 attention heads, and 1024 feed-forward network hidden dimensions are further stacked. The output of self-attention blocks is projected to the 128-dimensional mean and the variance for each frame.

Decoder: The decoder takes in the concatenation of the speaker representation and the content representation to reconstruct the acoustic feature. Conv layers and self-attention blocks with the same configurations as those in the content encoder are first adopted. An FC layer is then used to predict the acoustic feature. We further apply a PostNet [25] module which consists of 5 layers of 1D Conv with the kernel size of 5, to predict the residual of the acoustic feature. Dropout with probability 0.2 is used to regularize the PostNet. Both acoustic features predicted before and after the PostNet are taken out for the loss computation.

Loss function: While Eqn. (3) defines the theoretical form of the learning objective, which consists the negative log-likelihood of the conditional distribution of the acoustic feature as the reconstruction loss. In practice, we adopt the sum of the mean-squared error (MSE) loss and mean-absolute error (MAE) as the reconstruction loss. Since the conditional distribution defined by our reconstruction loss cannot be explicitly expressed, the β_c and β_s values are not normalized [11] with respect to the standard form of VAE.

Training and inference: During training, the inputs to the speaker encoder and content encoder are from the same utterance, while the input to the speaker encoder is firstly segmented and shuffled along the time axis before being fed into the following network. This operation can help avoid the content information being leaked into the speaker representation. The speaker latent and the content latent are sampled though re-parameterization during training, while only the mean vector and mean sequence of the two latent are used during inference. For model optimization, we use Adam optimizer [26] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$ and $\epsilon = 10^{-7}$, the learning rate is fixed at 1.25×10^{-4} . The training batch size is set as 32. During inference, the speech from the target speaker is fed into the speaker encoder to obtain the speaker representation, which is combined with content extracted from the source speech in the decoder to generated the converted speech.

Hyper-parameters tuning: The two most important hyper-parameters for the proposed model are β_c and β_s , which directly determines the disentanglement performance. To be honest, the tuning of these two parameters relies largely on trial and error, i.e., one can find the best setting through grid search. But we do notice one simple trick that can help find the good parameters faster. One can start with a relatively large β_c (e.g., 0.1, if the same loss function as ours is used) and a small β_s (e.g., 10^{-3}), this can generally yield the content and speaker disentanglement, but the generation quality may not be that good since the large β_c causes too much information loss of the content. Then we can decrease β_c gradually to find the value that yields also good generation quality. At the same time, β_s should first be decreased

proportionally with β_c and then separately tuned.

5. EXPERIMENTS

5.1. Dataset

We combine two openly available corpora, which are VCTK [27] and AISHELL-3 [28] together for training and evaluation of the proposed model. VCTK contains 110 speakers’ English speech data while AISHELL-3 consists of speech uttered by 218 native Mandarin speakers. 88 speakers’ data from VCTK and 116 speakers’ data from AISHELL-3 are combined as a bilingual training set. Another 20 speakers’ data from VCTK are evenly split for validation and testing, while for AISHELL-3 15 and 16 other speakers’ data are used so. Note that there are no common speakers among different splits of the dataset. We down-sample all speech into 16kHz and extract 80-dimensional logarithm Mel-Spectrograms with 200ms window length and 50ms window shift as the feature.

5.2. Evaluation on disentanglement

Following prior works [16, 18, 29], we conduct speaker verification (SV) on the learned content and speaker representations and report the equal error rate (EER) as a metric of the disentanglement. Intuitively, a high EER produced by \mathbf{z}_c and a low EER yielded by the \mathbf{z}_s can denote a good disentanglement of content and speaker representations. To compute the EER, we randomly select 4 utterances from each speaker in the test set as the enrolled samples, which are used to compute the speaker embedding (by averaging their \mathbf{z}_c ’s or \mathbf{z}_s ’s). The remaining utterances of the speaker are used as the positive trials while all other speakers’ utterances are negative trials. Cosine similarity is computed as the score. The EERs are evaluated separately for English and Mandarin test sets.

As discussed in Section 3.2, the weight parameters β_c and β_s restrict the respective amounts of information contained in \mathbf{z}_c and \mathbf{z}_s , and thus are very important for disentanglement. We show EERs with respect to \mathbf{z}_c and \mathbf{z}_s for both English (EN) and Mandarin (CN) test set in Table 1, where β_c varies among $\{10^{-3}, 10^{-2}\}$ and β_s is taken from $\{10^{-5}, 10^{-4}, 10^{-3}\}$.

We can first observe the effect of decreasing the absolute values of both β_c and β_s by comparing the results yielded by cases $(\beta_c, \beta_s) = (10^{-2}, 10^{-4})$ and $(\beta_s, \beta_c) = (10^{-3}, 10^{-5})$. While the ratios between β_c and β_s for both cases are 100, they produce quite different representations in terms of disentanglement. As we can observe from results of English test set, the EERs computed using \mathbf{z}_c and \mathbf{z}_s are 0.369 and 0.115 respectively for the former case, while those values become 0.198 and 0.069 for the latter case. While the former case shows more desirable disentanglement, the second case yields a \mathbf{z}_c with too much speaker information. This is because with restrictions that are too loose on both \mathbf{z}_c and \mathbf{z}_s when setting $(\beta_s, \beta_c) = (10^{-3}, 10^{-5})$, it cannot be ensured that the

information captured by \mathbf{z}_c and \mathbf{z}_s are complementary. Thus, different from [18] which claims that a proper ratio $\frac{\beta_c}{\beta_s}$ can induce the desired disentanglement, we argue that it is also important to first set proper absolute values for β_c and β_s , which restricts \mathbf{z}_c and \mathbf{z}_s to be more complementary.

Furthermore, setting proper relative values of β_c and β_s is also important for disentanglement. As can be observed from Table 1, increasing the value of β_c causes significant increases of EERs for both EN and CN content representations and all β_s values, denoting the reduction of the speaker information captured by \mathbf{z}_c . On the other hand, the increase of β_s puts more penalization on the speaker representation and thus allows more speaker information to leak into \mathbf{z}_c , which is indicated by the overall declining trend of EERs for \mathbf{z}_c and the opposite trend for \mathbf{z}_s horizontally. Meanwhile, larger β_c (10^{-2}) produces a stabler speaker-independent \mathbf{z}_c as the changes in EERs caused by the increasing β_s are limited.

Table 1. SV results on content and speaker representations

Rep.		$\beta_s = 10^{-5}$	$\beta_s = 10^{-4}$	$\beta_s = 10^{-3}$
\mathbf{z}_c (EN) \uparrow	$\beta_c = 10^{-3}$	0.198	0.179	0.151
	$\beta_c = 10^{-2}$	0.366	0.369	0.319
\mathbf{z}_c (CN) \uparrow	$\beta_c = 10^{-3}$	0.253	0.223	0.155
	$\beta_c = 10^{-2}$	0.379	0.375	0.328
\mathbf{z}_s (EN) \downarrow	$\beta_c = 10^{-3}$	0.069	0.131	0.372
	$\beta_c = 10^{-2}$	0.075	0.115	0.357
\mathbf{z}_s (CN) \downarrow	$\beta_c = 10^{-3}$	0.061	0.105	0.273
	$\beta_c = 10^{-2}$	0.043	0.113	0.285

5.3. Evaluation on one-shot VC

We conduct subjective evaluations on both intra-lingual and cross-lingual VC. We select 4 Mandarin speakers and 4 English speakers, both consisting of 2 male speakers and 2 female speakers. One utterance of each speaker is randomly chosen as the source and also the reference speech. All utterances are converted to all other speakers, thus in total we obtain 56 converted samples, while 24 of them are intra-lingual (EN2EN and CN2CN) and 32 are cross-lingual cases (EN2CN and CN2EN). Twelve subjects are asked to score these converted samples based on their naturalness and speaker similarity. 5-scale mean opinion score (MOS) is applied for both speech naturalness and speaker similarity.

We adopt two competitive unsupervised one-shot VC models AdIN-VC [9] and VQMIVC [8] as our baselines, while Hifi-GAN [30] is used as the vocoder. We denote the proposed model as β -VAEVC. We train the two baseline models as well as the vocoder on the same acoustic features and training set as ours. The copy-synthesized speech of the source speech is included in the naturalness evaluation, while those of randomly selected speech from the target speakers are scored in the speaker similarity evaluation. Some samples are available on <https://beta-vaevc.github.io>.

We first show the EERs obtained by content and speaker representations of three compared models in Table 2. For β -

VAEVC the β_c and β_s are set as 3×10^{-3} and 10^{-7} respectively. However, as shown in Table 2, this setting does not yield the best disentanglement results in terms of EERs. As \mathbf{z}_c of β -VAEVC seems to contain more speaker information than two baselines. Though we can achieve more speaker-independent \mathbf{z}_c for β -VAEVC via further increasing β_c , we find it can also decrease the generation quality. Thus, we choose the parameter setting that works better on the validation set in terms of speech generation quality. The naturalness

Table 2. SV results of three compared models

Model	\mathbf{z}_c (EN) \uparrow	\mathbf{z}_s (EN) \downarrow	\mathbf{z}_c (CN) \uparrow	\mathbf{z}_s (CN) \downarrow
AdIN-VC	0.373	0.065	0.371	0.074
VQMIVC	0.398	0.106	0.376	0.091
β -VAEVC	0.279	0.054	0.322	0.053

and similarity MOS results are shown in Table 3 and Table 4, respectively. We can observe that the proposed method achieves overall both better naturalness and speaker similarity than the two baselines. While AdIN-VC works comparably well for English intra-lingual conversion, its performance is much worse when it comes to Mandarin intra-lingual conversion and two directions of cross-lingual conversion. While VQMIVC archives overall good conversion naturalness, the speaker similarity performance is not as satisfying as that for speech naturalness, especially for cross-lingual cases.

Table 3. Speech naturalness MOS results ($\pm 95\%$ CI)

Model	EN2EN	EN2CN	CN2CN	CN2EN	Overall
Hifi-GAN	4.32 \pm 0.08	4.32 \pm 0.08	4.30 \pm 0.08	4.30 \pm 0.08	4.31 \pm 0.06
AdIN-VC	3.41 \pm 0.15	2.85 \pm 0.14	2.94 \pm 0.17	2.74 \pm 0.15	2.96 \pm 0.08
VQMIVC	3.56 \pm 0.15	3.15 \pm 0.12	3.27 \pm 0.13	3.18 \pm 0.14	3.27 \pm 0.07
β -VAEVC	3.71 \pm 0.14	3.53\pm0.13	3.71\pm0.14	3.35 \pm 0.15	3.56\pm0.07

Table 4. Speaker similarity MOS results ($\pm 95\%$ CI)

Model	EN2EN	EN2CN	CN2CN	CN2EN	Overall
Hifi-GAN	4.30 \pm 0.13	4.46 \pm 0.10	4.46 \pm 0.10	4.30 \pm 0.13	4.36 \pm 0.08
AdIN-VC	3.36 \pm 0.23	2.83 \pm 0.18	2.75 \pm 0.24	2.91 \pm 0.21	2.95 \pm 0.11
VQMIVC	3.04 \pm 0.25	2.60 \pm 0.19	3.26 \pm 0.22	2.55 \pm 0.20	2.82 \pm 0.11
β -VAEVC	3.54 \pm 0.19	3.00 \pm 0.16	3.32 \pm 0.24	3.31 \pm 0.18	3.27\pm0.10

In addition, we utilize the transcription error obtained by open-source pre-trained ASR models [31] as indicators of the conversion intelligibility. We conduct both intra-lingual and cross-lingual VC on the whole test set, that is, all utterances are converted to all the other speakers. There are in total 42,340 converted utterances for English intra-lingual VC, 77,920 for Mandarin intra-lingual VC, 67,744 for English-to-Mandarin conversion and 48,700 for Mandarin-to-English conversion. Pre-trained English and Mandarin ASR models are used to transcribe the corresponding utterances, then the word error rate (WER) and character error rate (CER) are computed respectively for English and Mandarin. The recognition results on re-synthesized utterances of all samples

in the test set by Hifi-GAN are also included as reference. The results are shown in Table 5. The proposed model surpasses baseline models by large margins for two intra-lingual VC cases and the Mandarin-to-English conversion case, and shows comparable performance with other two methods for the English-to-Mandarin conversion. Besides, we notice that VQMIVC achieves better WER / CER than AdIN-VC for cross-lingual conversion, while the latter model is better for intra-lingual cases.

Though all compared models can realize cross-lingual VC, we can see that there are gaps between the performance of intra-lingual VC and cross-lingual VC for all metrics, as shown in Table 3, 4 and 5. This is reasonable since the inputs to the content encoder and speaker encoder are from the different domain for cross-lingual VC cases, which is not the case during training. Besides, the difference in the recording conditions of two corpora VCTK and AISHELL-3 can make the train-inference mismatch issue even worse. We will tackle this problem in our future work.

Table 5. Speech recognition error results

Model	EN2EN	EN2CN	CN2CN	CN2EN
Hifi-GAN	5.11%		2.70%	
AdIN-VC	30.61%	45.88%	23.01%	51.31%
VQMIVC	35.09%	43.60%	33.37%	44.70%
β -VAEVC	23.41%	46.33%	10.58%	32.24%

6. CONCLUSION

We propose a method to disentangle speech into content and speaker representations, which can be applied to the challenging task of one-shot cross-lingual VC. Our method is based on a VAE with two encoders to extract speaker and content representations respectively. With the speaker encoder compressing the whole speech into a single vector and the content encoder extracting the frame-level representation out of speech, the time-variant and time-invariant elements of speech can be more easily separated into two representations. Furthermore, inspired by β -VAE, we propose a learning objective that incorporates two weight parameters to restrict the amount of information that can be captured by the two representations. With proper weight parameters imposed, the disentanglement can be ensured to be with respect to content and speaker information. We apply the proposed method to one-shot cross-lingual VC, through which we show the effectiveness of the proposed method in achieving content and speaker disentanglement.

7. ACKNOWLEDGMENTS

This research is supported by the Centre for Perceptual and Interactive Intelligence, a CUHK InnoCentre.

8. REFERENCES

- [1] Berrak Sisman, Junichi Yamagishi, Simon King, and Haizhou Li, “An overview of voice conversion and its challenges: From statistical modeling to deep learning,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 132–157, 2020.
- [2] Yi Zhou, Xiaohai Tian, Haihua Xu, Rohan Kumar Das, and Haizhou Li, “Cross-lingual voice conversion with bilingual phonetic posteriorgram and average modeling,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6790–6794.
- [3] Diederik P. Kingma and Max Welling, “Auto-encoding variational bayes,” in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun, Eds., 2014.
- [4] Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark Hasegawa-Johnson, “Autovc: Zero-shot voice style transfer with only autoencoder loss,” in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, Kamalika Chaudhuri and Ruslan Salakhutdinov, Eds. 2019, vol. 97 of *Proceedings of Machine Learning Research*, pp. 5210–5219, PMLR.
- [5] Da-Yi Wu and Hung-Yi Lee, “One-shot voice conversion by vector quantization,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*. 2020, pp. 7734–7738, IEEE.
- [6] Da-Yi Wu, Yen-Hao Chen, and Hung-yi Lee, “VQVC+: one-shot voice conversion by vector quantization and u-net architecture,” in *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, Helen Meng, Bo Xu, and Thomas Fang Zheng, Eds. 2020, pp. 4691–4695, ISCA.
- [7] Siyang Yuan, Pengyu Cheng, Ruiyi Zhang, Weituo Hao, Zhe Gan, and Lawrence Carin, “Improving zero-shot voice style transfer via disentangled representation learning,” in *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. 2021, OpenReview.net.
- [8] Disong Wang, Liquan Deng, Yu Ting Yeung, Xiao Chen, Xunying Liu, and Helen Meng, “VQMIVC: vector quantization and mutual information-based unsupervised speech representation disentanglement for one-shot voice conversion,” in *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August - 3 September 2021*, Hynek Hermansky, Honza Cernocký, Lukás Burget, Lori Lamel, Odette Scharenborg, and Petr Motlíček, Eds. 2021, pp. 1344–1348, ISCA.
- [9] Ju-Chieh Chou and Hung-yi Lee, “One-shot voice conversion by separating speaker and content representations with instance normalization,” in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, Gernot Kubin and Zdravko Kacic, Eds. 2019, pp. 664–668, ISCA.
- [10] Yen-Hao Chen, Da-Yi Wu, Tsung-Han Wu, and Hung-yi Lee, “Again-vc: A one-shot voice conversion using activation guidance and adaptive instance normalization,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5954–5958.
- [11] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner, “beta-vae: Learning basic visual concepts with a constrained variational framework,” in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. 2017, OpenReview.net.
- [12] Christopher P. Burgess, Irina Higgins, Arka Pal, Loïc Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner, “Understanding disentangling in β -vae,” *CoRR*, vol. abs/1804.03599, 2018.
- [13] Lifa Sun, Kun Li, Hao Wang, Shiyin Kang, and Helen Meng, “Phonetic posteriorgrams for be-to-one voice conversion without parallel data training,” in *2016 IEEE International Conference on Multimedia and Expo (ICME)*, 2016, pp. 1–6.
- [14] Yi Zhou, Xiaohai Tian, Emre Yilmaz, Rohan Kumar Das, and Haizhou Li, “A modularized neural network with language-specific output layers for cross-lingual voice conversion,” *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 160–167, 2019.
- [15] Shengkui Zhao, Hao Wang, Trung Hieu Nguyen, and Bin Ma, “Towards natural and controllable cross-lingual voice conversion based on neural tts model and phonetic posteriorgram,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5969–5973.
- [16] Wei-Ning Hsu, Yu Zhang, and James R. Glass, “Unsupervised learning of disentangled and interpretable rep-

- representations from sequential data,” in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, Eds., 2017, pp. 1878–1889.
- [17] Seyed Hamidreza Mohammadi and Taehwan Kim, “Investigation of using disentangled and interpretable representations for one-shot cross-lingual voice conversion,” in *INTERSPEECH*, 2018.
- [18] Jiachen Lian, Chunlei Zhang, and Dong Yu, “Robust disentangled variational speech representation learning for zero-shot voice conversion,” in *IEEE ICASSP*. IEEE, 2022.
- [19] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem, “Challenging common assumptions in the unsupervised learning of disentangled representations,” in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, Kamalika Chaudhuri and Ruslan Salakhutdinov, Eds. 2019, vol. 97 of *Proceedings of Machine Learning Research*, pp. 4114–4124, PMLR.
- [20] Ricky TQ Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud, “Isolating sources of disentanglement in variational autoencoders,” *Advances in neural information processing systems*, vol. 31, 2018.
- [21] Hyunjik Kim and Andriy Mnih, “Disentangling by factorising,” in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, Jennifer G. Dy and Andreas Krause, Eds. 2018, vol. 80 of *Proceedings of Machine Learning Research*, pp. 2654–2663, PMLR.
- [22] Abien Fred Agarap, “Deep learning using rectified linear units (relu),” *arXiv preprint arXiv:1803.08375*, 2018.
- [23] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [24] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [25] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al., “Natural tts synthesis by conditioning wavenet on mel spectrogram predictions,” in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [26] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2015.
- [27] Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald, “Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92),” 2019.
- [28] Yao Shi, Hui Bu, Xin Xu, Shaoji Zhang, and Ming Li, “AISHELL-3: A Multi-Speaker Mandarin TTS Corpus,” in *Proc. Interspeech 2021*, 2021, pp. 2756–2760.
- [29] Yingzhen Li and Stephan Mandt, “Disentangled sequential autoencoder,” in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, Jennifer G. Dy and Andreas Krause, Eds. 2018, vol. 80 of *Proceedings of Machine Learning Research*, pp. 5656–5665, PMLR.
- [30] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae, “Hifigan: Generative adversarial networks for efficient and high fidelity speech synthesis,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 17022–17033, 2020.
- [31] Oleksii Kuchaiev, Jason Li, Huyen Nguyen, Oleksii Hrinchuk, Ryan Leary, Boris Ginsburg, Samuel Kriman, Stanislav Beliaev, Vitaly Lavrukhin, Jack Cook, et al., “Nemo: a toolkit for building ai applications using neural modules,” *arXiv preprint arXiv:1909.09577*, 2019.