
Unsupervised Speech Decomposition via Triple Information Bottleneck

Kaizhi Qian^{*12} Yang Zhang^{*1} Shiyu Chang¹ David Cox¹ Mark Hasegawa-Johnson²

Abstract

Speech information can be roughly decomposed into four components: language content, timbre, pitch, and rhythm. Obtaining disentangled representations of these components is useful in many speech analysis and generation applications. Recently, state-of-the-art voice conversion systems have led to speech representations that can disentangle speaker-dependent and independent information. However, these systems can only disentangle timbre, while information about pitch, rhythm and content is still mixed together. Further disentangling the remaining speech components is an under-determined problem in the absence of explicit annotations for each component, which are difficult and expensive to obtain. In this paper, we propose SPEECHFLOW, which can blindly decompose speech into its four components by introducing three carefully designed information bottlenecks. SPEECHFLOW is among the first algorithms that can separately perform style transfer on timbre, pitch and rhythm without text labels. Our code is publicly available at <https://github.com/auspicious3000/SpeechSplit>.

1. Introduction

Human speech conveys a rich stream of information, which can be roughly decomposed into four important components: *content*, *timbre*, *pitch* and *rhythm*. The language content of speech comprises the primary information in speech, which can also be transcribed to text. Timbre carries information about the voice characteristics of a speaker, which is closely connected with the speaker’s identity. Pitch and rhythm are the two major components of *prosody*, which expresses the emotion of the speaker. Pitch variation conveys the aspects

^{*}Equal contribution ¹MIT-IBM Watson AI Lab, USA ²University of Illinois at Urbana-Champaign, USA. Correspondence to: Yang Zhang <yang.zhang2@ibm.com>, Kaizhi Qian <kaizhiqian@gmail.com>.

of the tone of the speaker, and rhythm characterizes how fast the speaker utters each word or syllable.

For decades, speech researchers have sought to obtain disentangled representations of these speech components, which are useful in many speech applications. In speech analysis tasks, the disentanglement of speech components helps to remove interference introduced by irrelevant components. In speech generation tasks, disentanglement is the foundation of many applications, such as voice conversion (Chou & Lee, 2019), prosody modification (Shechtman & Sorin, 2019), emotional speech synthesis (Pell et al., 2011), and low bit-rate speech encoding (Schroeder & Atal, 1985), to name a few.

Recently, state-of-the-art voice conversion systems have been able to obtain a speaker-invariant representation of speech, which disentangles the speaker-dependent information (Qian et al., 2019; Chou et al., 2018; Chou & Lee, 2019). However, these algorithms are only able to disentangle timbre. The remaining aspects, *i.e.* content, pitch, and rhythm are still lumped together. As a result, the converted speech produced by these algorithms differs from the source speech only in terms of timbre. The pitch contour and rhythm remain largely the same.

From an information-theoretic perspective, the success in timbre disentanglement can be ascribed to the availability of a speaker identity label, which preserves almost all the information of timbre, such that voice conversion systems can ‘subtract’ such information from speech. For example, AUTOVC (Qian et al., 2019), a voice conversion system, constructs an autoencoder for speech and feeds the speaker identity label to the decoder. As shown in figure 2(a), by constructing an information bottleneck between the encoder and decoder, AUTOVC can force the encoder to remove the timbre information, because the equivalent information is supplied to the decoder directly. It is worth noting that although the speaker identity is also correlated with the pitch and timbre information, the information overlap is relatively small, so the speaker identity cannot serve as labels for pitch and rhythm. If we had analogous information-preserving labels for timbre, rhythm or pitch, the disentanglement of these aspects would be straightforward, simply by utilizing these labels the same way voice conversion algorithms use the speaker identity label.

However, obtaining annotations for these other speech components is challenging. First, although language content annotation is effectively provided by text transcriptions, obtaining a large number of text transcriptions is expensive, especially for low-resourced languages. Therefore, here, we will focus on unsupervised methods that do not rely on text transcriptions. Second, the rhythm annotation, which is essentially the length of each syllable, can only be obtained with the help of text transcriptions, which are again unavailable under our unsupervised setting. Finally, for pitch annotation, although the pitch information can be extracted as pitch contour using pitch extraction algorithms, the pitch contour itself is entangled with rhythm information, because it contains the information of how long each speech segment is. Without the information preserving labels, disentangling content, rhythm and pitch becomes an under-determined problem. Hence, here we ask: is it possible to decompose these remaining speech components without relying on text transcriptions and other information-preserving labels?

In this paper, we propose SPEECHFLOW, a speech generative model that can blindly decompose speech into content, timbre, pitch, and rhythm, and generate speech from these disentangled representations. Thus, SPEECHFLOW is among the first algorithms that can enable flexible conversion of different aspects to different styles without relying on any text transcription. To achieve unsupervised decomposition, SPEECHFLOW introduces an encoder-decoder structure with three encoder channels, each with a different, carefully-crafted information bottleneck design. The information bottleneck is imposed by two mechanisms: first, a constraint on the physical dimension of the representation, which has been shown effective in AUTOVC, and second, the introduction of noise by randomly resampling along the time dimension, which has been shown effective in (Polyak & Wolf, 2019). We find that subtle differences in the information bottleneck design can force different channels to pass different information, such that one passes language content, one passes rhythm, and one passes pitch information, thereby achieving the blind disentanglement of all speech components.

Besides direct value in speech applications, SPEECHFLOW also provides insight into a powerful design principle that can be broadly applied to any disentangled representation learning problem: in the presence of an information bottleneck, a neural network will prioritize passing through the information that cannot be provided elsewhere. This observation inspires a generic approach to disentanglement.

2. Related Work

The Source-Filter Model Early research on speech generation proposed the source-filter model (Quatieri, 2006), and many subsequent research efforts try to decompose

speech into the source that includes pitch and the filter that includes content, using signal processing approaches, such as linear predictive coding (Atal & Schroeder, 1979), cepstral analysis (Mermelstein, 1976), temporally stable power spectral analysis (Kawahara et al., 2008) and probabilistic approaches (Zhang et al., 2014). However, these approaches do not consider the prosody aspects of speech.

Voice Conversion Inspired by the style transfer and disentanglement techniques in computer vision (Lample et al., 2017; Kaneko & Kameoka, 2017; Choi et al., 2018), many approaches based on variational autoencoders (VAEs) and generative adversarial networks (GANs) have been proposed in the field of voice conversion to disentangle the timbre information from the speech. VAE-VC (Hsu et al., 2016) directly applies VAE for voice conversion, where the encoder produces a speaker-independent content embedding. After that, VAE-GAN (Hsu et al., 2017) replaces the decoder of VAE and a GAN when generating the converted speech to improve the quality of the conversion results. CDVAE-VC (Huang et al., 2018) uses two VAEs working on different speech features, one on STRAIGHT spectra (Kawahara et al., 2008), and one on mel-cepstral coefficients (MCCs), and encourages that the latent representation can reconstruct both features well. ACVAE-VC (Kameoka et al., 2018a) introduces an auxiliary classifier for the conversion outputs, and encourages the converted speech to be correctly classified as the source speaker. Chou et al. (2018) introduced a classifier for the latent code, and discourages the latent code to be correctly classified as the target speaker. Inspired by image style transfer frameworks, Gao et al. (2018) and Kameoka et al. (2018b) adapted CycleGan (Kaneko & Kameoka, 2017) and StarGan (Choi et al., 2018) respectively for voice conversion. Later, CDVAE-VC was extended by directly applying GAN (Huang et al., 2020) to improve the degree of disentanglement. Chou & Lee (2019) used instance normalization to further disentangle speaker from content, and thus can convert to speakers that are not seen during training. StarGan-VC2 (Kaneko et al., 2019) refined the adversarial framework by conditioning the generator and discriminator on the source speaker label, in addition to the target speaker label. Recently, Qian et al. (2019) proposed AUTOVC, a simple autoencoder based method that disentangles the timbre and content using information-constraining bottlenecks. Later, Qian et al. (2020) fixed the pitch jump problem of AUTOVC by F0 conditioning. Besides, the time-domain deep generative model is gaining more research attention for voice conversion (Niwa et al., 2018; Nachmani & Wolf, 2019; Serrà et al., 2019). However, these methods only focus on converting timbre, which is only one of the speech components.

Prosody Disentanglement There are recently many text-to-speech (TTS) systems that seek to disentangle the prosody information to generate expressive speech. Skerry-Ryan




Figure 1. Spectrograms (left) and pitch contours (right) of two utterances of the same sentence ‘Please call Stella’. The left rectangle marks highlight the formant structures of the phone ‘ea’. The arrows mark the frequencies of the second, third and fourth formants. The right rectangle marks highlight the pitch tones of the word ‘Stella’.

et al. (2018) introduced a Tacotron based speech synthesizer that can disentangle prosody from speech content by having an encoder that can extract the prosody information from the original speech. Mellotron (Valle et al., 2020) is a speech synthesizer conditional on both explicit prosody labels and latent prosody code to capture and disentangle different aspects of the prosody information. CHiVE (Kenter et al., 2019) introduces a hierarchical encoder-decoder structure that is conditioned on a set of prosodic features and linguistic features. However, these TTS systems all require text transcriptions, which, as discussed, makes the task easier but limits their applications to high-resource language. Besides TTS systems, Parrotron (Biadys et al., 2019) disentangles prosody by encouraging the latent codes to be the same as the corresponding phone representation of the input speech. However, Parrotron still requires text transcriptions to label the phone representation, as well as to generate the synthetic parallel dataset. Polyak & Wolf (2019) proposed, to the best of our knowledge, the only prosody disentanglement algorithm that does not rely on text transcriptions, which attempts to remove the rhythm information by randomly resampling the input speech. However, the effect of their prosody conversion is not very pronounced. In this paper, we would like to achieve effective prosody conversion without using text transcriptions, which is more flexible for low-resource languages.

3. Background: Information in Speech

Since this paper focuses on the decomposition of speech information into rhythm, pitch, timbre, and content, we provide here a brief primer on each of these components. Figure 1 shows the spectrograms (left) and pitch contours (right) of utterances of the sentence ‘Please call Stella’. Throughout this paper, the term ‘spectrogram’ refers to the magnitude spectrogram.

Rhythm Rhythm characterizes how fast the speaker utters each syllable. As shown in figure 1, each spectrogram is divided into segments, which correspond to each word, as marked on the horizontal axis. So the lengths of these segments reflect the rhythm information. In the top spectrogram, each segment is long, indicating a slow speaker; in the bottom spectrogram, each segment is short, indicating a fast speaker.

Pitch Pitch is an important component of intonation. One popular representation of the pitch information is the *pitch targets* (Xu & Wang, 2001), which is defined as the intended pitch, e.g. rise or fall, high or low etc., of each syllable. The pitch information, or pitch target information, is contained in the pitch contour, because the pitch contour is generally considered as the result of a constant attempt to hit the pitch targets of each syllable, subject to the physical constraints (Xu & Wang, 2001). However, the pitch contour also entangles other information. First, the pitch contour contains the rhythm information, because each nonzero segment of the pitch contour represents a voiced segment, which typically corresponds to a word or a syllable. So the length of each voiced segment indicates how fast the speaker speaks. Second, the pitch range reflects certain speaker identity information – female speakers tend to have a high pitch range, as shown in the upper panel of figure 1, and male speakers tend to have a low pitch range, as shown in the lower panel of figure 1. Here, we assume that the impact of the speaker identity on the pitch contour is linear. In other words, if we normalize the pitch contour to a common mean and standard deviation, the speaker identity information will be removed. To sum up, the pitch contour entangles the information of speaker identity, rhythm and pitch; the normalized pitch contour only contains the information of the latter two.

Timbre Timbre is perceived as the voice characteristics of a speaker. It is reflected by the *formant frequencies*, which are the resonant frequency components in the vocal tract. In a spectrogram, the formants are shown as the salient frequency components of the spectral envelope. In figure 1, the rectangles and arrows on the spectrogram highlight three formants. As can be seen, the top spectrogram has a higher formant frequency range, indicating a bright voice; the bottom spectrogram has a lower formant frequency range, indicating a deep voice.

Content In English and many other languages, the basic unit of content is the phone. Each phone comes with a particular formant pattern. For example, the three formants highlighted in figure 1 are the second, third and fourth lowest formants of the phone ‘ea’ as in ‘please’. Although their formant frequencies have different ranges, which indicates their timbre difference, they have the same pattern – they tend to cluster together and are far away from the lowest formant (which is at around 100 Hz).

4. SPEECHFLOW

This section introduces SPEECHFLOW. For notation, upper-cased letters, X and \mathbf{X} , denote random scalars and vectors respectively; lower-cased letters, x and \mathbf{x} , denote deterministic scalars and vectors respectively; $H(\mathbf{X})$ denotes the Shannon entropy of \mathbf{X} ; $H(Y|\mathbf{X})$ denotes the entropy of Y conditional on \mathbf{X} ; $I(Y; \mathbf{X})$ denotes the mutual information.

4.1. Problem Formulation

Denote $\mathbf{S} = \{\mathbf{S}_t\}$ as a speech spectrogram, where t is the time index. Denote the speaker’s identity as U . We assume that \mathbf{S} and U are generated through the following random generative processes

$$\mathbf{S} = g_s(\mathbf{C}, \mathbf{R}, \mathbf{F}, \mathbf{V}), \quad U = g_u(\mathbf{V}), \quad (1)$$

where \mathbf{C} denotes content; \mathbf{R} denotes rhythm; \mathbf{F} denotes pitch target; \mathbf{V} denotes timbre. $g_s(\cdot)$ and $g_u(\cdot)$ are assumed to be a one-to-one mapping. Note that here we assume \mathbf{C} also accounts for the residual information that is not included in rhythm, pitch or timbre.

Our goal is to construct an autoencoder-based generative model for speech, such that the hidden code contains disentangled representations of the speech components. We formally denote the representations as \mathbf{Z}_c , \mathbf{Z}_r and \mathbf{Z}_f , and these representations should satisfy

$$\mathbf{Z}_c = h_c(\mathbf{C}), \quad \mathbf{Z}_r = h_r(\mathbf{R}), \quad \mathbf{Z}_f = h_f(\mathbf{F}), \quad (2)$$

where $h_c(\cdot)$, $h_r(\cdot)$ and $h_f(\cdot)$ are all one-to-one mappings.

4.2. AUTOVC and Its Limitations

Since SPEECHFLOW inherits the information bottleneck mechanism proposed in AUTOVC, it is necessary to first review its framework and limitations. Figure 2(a) shows the framework of AUTOVC, which consists of an encoder and a decoder. The encoder has an information bottleneck at the end (shown as the grey tip), which is implemented as hard constraint on code dimensions. The input to the encoder is speech spectrogram \mathbf{S} , and the output of the encoder is called the speech code, denoted as \mathbf{Z} . The decoder takes \mathbf{Z} and the speaker identity label U as its inputs, and produces a speech spectrogram $\hat{\mathbf{S}}$ as output. We formally denote the encoder as $E(\cdot)$, and the decoder as $D(\cdot, \cdot)$. The AUTOVC pipeline can be expressed as

$$\mathbf{Z} = E(\mathbf{S}), \quad \hat{\mathbf{S}} = D(\mathbf{Z}, U). \quad (3)$$

During training, the output of the decoder tries to reconstruct the input spectrogram:

$$\min_{\theta} \mathbb{E}[\|\hat{\mathbf{S}} - \mathbf{S}\|_2^2], \quad (4)$$

where θ denotes all the trainable parameters.

It can be shown that if the information bottleneck is tuned to the right size, this simple scheme can achieve disentanglement of the timbre information as

$$\mathbf{Z} = h(\mathbf{C}, \mathbf{R}, \mathbf{F}). \quad (5)$$

Figure 2(a) provides an intuitive explanation of why this is possible. As can be seen, speech is represented as a concatenation of different blocks, indicating the content, rhythm, pitch and timbre information. Note that speaker identity is represented with the same block style as timbre because it is assumed to preserve equivalent information to timbre according to equation (1). Since the speaker identity is separately fed to the decoder, the decoder can still have access to all the information to perform self-reconstruction even if the encoder does not preserve the timbre information in its output. Therefore, when the information bottleneck is binding, the encoder will remove the timbre information. However, \mathbf{Z} still lumps content, rhythm, and pitch together. As a result, AUTOVC can only convert timbre.

4.3. The SPEECHFLOW Framework

Figure 2(b) illustrates the SPEECHFLOW framework. SPEECHFLOW is also an autoencoder with an information bottleneck. However, in order to further decompose the remaining speech components, SPEECHFLOW introduces three encoders with heterogeneous information bottleneck, which are a *content encoder*, a *rhythm encoder*, and a *pitch encoder*. Below are the details of the encoders and the decoder of SPEECHFLOW.

The Encoders As shown in figure 2(b), all three encoders are almost the same, but with two subtle differences. First, the input to the content encoder and rhythm encoder is speech \mathbf{S} , whereas the input to the pitch encoder is the normalized pitch contour, which we denote as \mathbf{P} . As discussed in section 3, the normalized pitch contour \mathbf{P} refers to the pitch contour that is normalized to have the same mean and variance across all the speakers, so the normalized pitch contour only contains the pitch information, \mathbf{F} , and rhythm information, \mathbf{R} , but no speaker ID information, U .

Second, the content encoder and pitch encoder perform a random resampling operation along the time dimension of the input. Random resampling involves two steps of operations. The first step is to divide the input into segments of random lengths. The second step is to randomly stretch or squeeze each segment along the time dimension. Therefore, random resampling can be regarded as an information bottleneck on rhythm. More details of random resampling can be found in appendix B.2. All the encoders have the physical information bottleneck at the output. The final outputs of the encoders are called content code, rhythm code and pitch code, which are denoted as \mathbf{Z}_c , \mathbf{Z}_r and \mathbf{Z}_f respectively. Formally, denote the content encoder as $E_c(\cdot)$, rhythm encoder

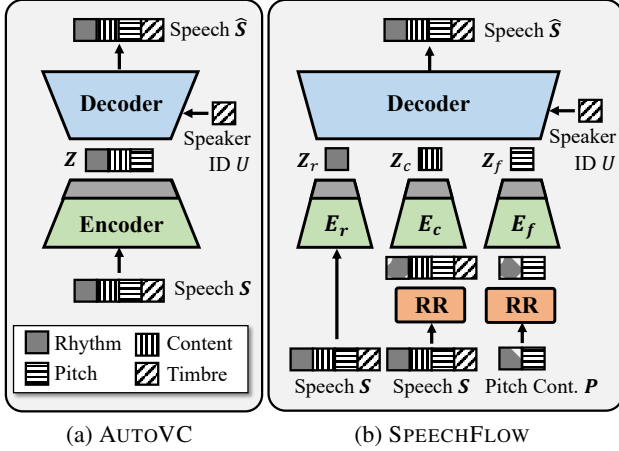


Figure 2. Frameworks and AUTOVC and SPEECHFLOW and illustration of why they can perform disentanglement. Signals are represented as blocks to denote their information components. E_r denotes the rhythm encoder; E_c denotes the content encoder; E_f denotes the pitch encoder. ‘RR’ denotes random resampling. ‘Pitch Cont.’ is short for the normalized pitch contour. The grey block at the tip of the encoders denotes the information bottleneck. Some rhythm blocks have some holes in them, which represents that a portion of the rhythm information is lost. The bottlenecks force the encoders to pass only the information that other encoders cannot supply, hence achieving the disentanglement.

as $E_r(\cdot)$ and pitch encoder as $E_f(\cdot)$, and denote the random resampling operation as $A(\cdot)$. Then we have

$$\mathbf{Z}_c = E_c(A(\mathbf{S})), \quad \mathbf{Z}_r = E_r(\mathbf{S}), \quad \mathbf{Z}_f = E_f(A(\mathbf{P})). \quad (6)$$

The Decoder The decoder takes all the speech code and the speaker identity label (or embedding) as its inputs, and produce a speech spectrogram as output, *i.e.*,

$$\hat{\mathbf{S}} = D(\mathbf{Z}_c, \mathbf{Z}_r, \mathbf{Z}_f, U). \quad (7)$$

During training, the output of the decoder tries to reconstruct the input spectrogram, which is the same as in equation (4).

Counter-intuitive as it may sound, we claim that when all the information bottlenecks are appropriately set and the network representation power is sufficient, a minimizer of equation (4) will satisfy the disentanglement condition as in equation (2). In what follows, we will explain why such decomposition is possible.

4.4. Why Does It Force Speech Decomposition?

Figure 2 provides an intuitive illustration of how SPEECHFLOW achieves speech decomposition, where a few important assumptions are made.

Assumption 1: The random resampling operation will contaminate the rhythm information \mathbf{R} , *i.e.* $\forall r_1 \neq r_2$

$$\Pr[A(g_s(\mathbf{C}, r_1, \mathbf{F}, \mathbf{V})) = A(g_s(\mathbf{C}, r_2, \mathbf{F}, \mathbf{V}))] > 0. \quad (8)$$

Assumption 2: The random resampling operation will not contaminate the other speech components, *i.e.*

$$I(\mathbf{C}; A(\mathbf{S})) = H(\mathbf{C}), \quad I(\mathbf{F}; A(\mathbf{S})) = H(\mathbf{F}). \quad (9)$$

Assumption 3: The pitch contour \mathbf{P} contains all the pitch information and a portion of rhythm information.

$$\mathbf{P} = g_p(\mathbf{F}, \mathbf{R}), \quad I(\mathbf{F}; \mathbf{P}) = H(\mathbf{F}). \quad (10)$$

As shown in figure 2(b), speech contains four blocks of information. When it passes through the random resampling operation, a random portion of the rhythm block is wiped (shown as the holes in the rhythm block at the output of the RR module), but the other blocks remain intact. On the other hand, the normalized pitch contour mainly contains two blocks, the pitch block, and the rhythm block. The rhythm block is missing a corner because the normalized pitch contour does not contain all the rhythm information, and it misses even more when it passes through the random resampling module.

Similar to the AUTOVC claim, the timbre information is directly fed to the decoder, so all the encoders do not need to encode the timbre information. Therefore, this section focuses on explaining why SPEECHFLOW can force the encoders to separately encode the content, pitch, and timbre.

First, the rhythm encoder $E_r(\cdot)$ is the only encoder that has access to the complete rhythm information \mathbf{R} . The other two encoders only preserve a random portion of \mathbf{R} , and there is no way for $E_r(\cdot)$ to guess which part is lost and thus only supply the lost part. Therefore, $E_r(\cdot)$ must pass all the rhythm information. Meanwhile, the other aspects are available in the other two encoders. So if $E_r(\cdot)$ is forced to lose some information by its information bottleneck, it will prioritize removing the content, pitch, and timbre.

Second, given that $E_r(\cdot)$ only encodes \mathbf{R} , then the content encoder $E_c(\cdot)$ becomes the only encoder that can encode all the content information \mathbf{C} , because the pitch encoder does not have access to \mathbf{C} . Therefore, $E_c(\cdot)$ must pass all the content information. Meanwhile, the other aspects can be supplied elsewhere, so the rhythm encoder will remove the other aspects if the information bottleneck is binding.

Finally, with $E_r(\cdot)$ encoding only \mathbf{R} and $E_c(\cdot)$ encoding only \mathbf{C} , the pitch encoder $E_f(\cdot)$ must encode the pitch information. All the other aspects are supplied in other channels, so $E_f(\cdot)$ will prioritize removing these aspects if the information bottleneck is binding.

Simply put, if each encoder is only allowed to pass one block, then the arrangement in figure 2 is the only way to ensure full recovery of the speech information.

Now we are ready to give our formal result.




Figure 3. The architecture of SPEECHFLOW. ‘GNorm’ denotes group normalization; ‘RR’ denotes random resampling; ‘Down’ and ‘Up’ denote downsampling and upsampling operations respectively. ‘Linear’ denotes linear projection layer. $\times n$ denotes the module above is repeated n times.

Table 1. Hyperparameter settings of the encoders.

| | Rhythm | Content | Pitch |
|-------------------|--------|---------|-------|
| Conv Layers | 1 | 3 | 3 |
| Conv Dim | 128 | 512 | 256 |
| Norm Groups | 8 | 32 | 16 |
| BLSTM Layers | 1 | 2 | 1 |
| BLSTM Dim | 1 | 8 | 32 |
| Downsample Factor | 8 | 8 | 8 |

Theorem 1. Assume C , R , F , U are independent, and that the information bottleneck is precisely set such that

$$H(\mathbf{Z}_c) = H(C), \quad H(\mathbf{Z}_r) = H(R), \quad H(\mathbf{Z}_f) = H(F). \quad (11)$$

Assume equations (1), (8), (9) and (10) hold. Then the global optimum of equation (4) would produce the disentangled representation as in equation (2).

The proof is presented in the appendix. Although theorem 1 is contingent on a set of relatively stringent conditions, which may not hold in practice, we will empirically verify the disentanglement capabilities in section 5.

4.5. Network Architecture

Figure 3 shows the architecture of SPEECHFLOW. The left module corresponds to the encoders and the right to the decoder. All three encoders share a similar architecture, which consists of a stack of 5×1 convolutional layers followed by group normalization (Wu & He, 2018). For the content encoder, the output of each convolutional layer is passed to a random resampling module to further contaminate rhythm. The final output of the convolutional layers is fed to a stack of bidirectional-LSTM layers to reduce the feature dimension, which is then passed through a downsampling operation to reduce the temporal dimension, producing the hidden representations. Table 1 shows the hyperparameter settings of each encoder. More details of the downsampling operation are provided in appendix B.2.

The decoder first upsamples the hidden representation to restore the original sampling rate. The speaker identity

label U , which is a one-hot vector, is also repeated along the time dimension to match the temporal dimension of the other upsampled representations. All the representations are then concatenated along the channel dimension and fed to a stack of three bidirectional-LSTM layers with an output linear layer to produce the final output. The spectrogram is converted back to the speech waveform using the same wavenet-vocoder as in AUTOVC. Additional architecture and implementation details are provided in appendix B.

5. Experiments

In this section, we will empirically verify the disentanglement capability of SPEECHFLOW. We will be visualizing our speech results using spectrogram and pitch contour. However, to fully appreciate the performance of SPEECHFLOW, we strongly encourage readers to refer to our online demo¹. Additional experiment results can be found in appendix C. The frequency axis units of all the spectrograms are in kHz, and those of the pitch contour plots are in Hz.

5.1. Configurations

The experiments are performed on the VCTK dataset (Veaux et al., 2016). The training set contains 20 speakers where each speaker has about 15 minutes of speech. The test set contains the same 20 speakers but with different utterances, which is the conventional voice conversion setting. SPEECHFLOW is trained using the ADAM optimizer (Kingma & Ba, 2014) with a batch size of 16 for 800k steps. Since there are no other algorithms that can perform blind decomposition so far, we will be comparing our result with AUTOVC, a conventional voice conversion baseline.

The model selection is performed on the training dataset. Specifically, the physical bottleneck dimensions are tuned based on the criterion: when the input to one of the encoders or the speaker embedding is set to zero, the output reconstruction should not contain the corresponding information. As will be shown in section 5.4, setting the inputs and speaker embedding to zero can measure the degree of disentanglement. From the models that satisfy this criterion, we pick the one with the lowest training error. Appendix B.4 provides additional guidance on how to tune the bottlenecks.

5.2. Rhythm, Pitch and Timbre Conversions

If SPEECHFLOW can decompose the speech into different components, then it should be able to separately perform style transfer on each aspect, which is achieved by replacing the input to the respective encoder with that of the target utterance. For example, if we want to convert pitch, we feed the target pitch contour to the pitch encoder. To convert

¹<https://auspicious3000.github.io/SpeechSplit-Demo/>

timbre, we feed the target speaker id to the decoder.

We construct parallel speech pairs from the test set, where both the source and target speakers read the same utterances. Please note that we use the parallel pairs *only for testing*. During training, SPEECHFLOW is trained without parallel speech data. For each parallel pair, we set one utterance as the source and one as the target, and perform seven different types of conversions, including three single-aspect conversions (rhythm-only, pitch-only and timbre-only), three double-aspect conversions (rhythm+pitch, rhythm+timbre, and pitch+timbre), and one all-aspect conversion.

Conversion Visualization Figure 4 shows the single-aspect conversion results on a speech pair uttering ‘Please call Stella’. The source speaker is a slow female speaker, and the target speaker is a fast male speaker. As shown in figure 4, SPEECHFLOW can separately convert each aspect. First, in terms of rhythm, note that the rhythm-only conversion is perfectly aligned with the target utterance in time, whereas the timbre-only and pitch-only conversions are perfectly aligned with the source utterance in time. Second, in terms of pitch, notice that the timbre-only and rhythm-only conversions have a falling tone on the word ‘Stella’, which is the same as the source utterance, as highlighted by the dashed rectangle. The pitch-only conversion has a rising tone on ‘Stella’, which is the same as the target utterance, as highlighted by the solid rectangles. Third, in terms of timbre, as highlighted by the rectangles on the spectrograms, the formants of pitch-only and rhythm-only conversions are as high as those of the source speech, and the formants of timbre-only conversions are as high as those in the target.

Subjective Evaluation We also perform a subjective evaluation on *Amazon Mechanical Turk* on whether the conversion of each aspect is successful. For example, to evaluate whether the different conversions convert pitch, we select 20 speech pairs that are perceptually distinct in pitch, and generate all the seven types of conversions, plus the AUTOVC conversion and the source utterance as baselines. Each test is assigned to five subjects. In the test, the subject is presented with two reference utterances, which are the source and target utterances in a random order, and then with one of the nine conversion results. The subject is asked to select which reference utterance has a more similar pitch tone to the converted utterance. We compute the *pitch conversion rate* as the percentage of answers that choose the target utterance. We would expect the utterances with pitch converted to have a high pitch conversion rate; otherwise, the pitch conversion rate should be low. The *rhythm conversion rate* and *timbre conversion rate* are computed in a similar way. More details of the test token generation process can found in appendix B.5.

Figure 5 shows the conversion rates of different types of conversions. As expected, the conversion rate is high when

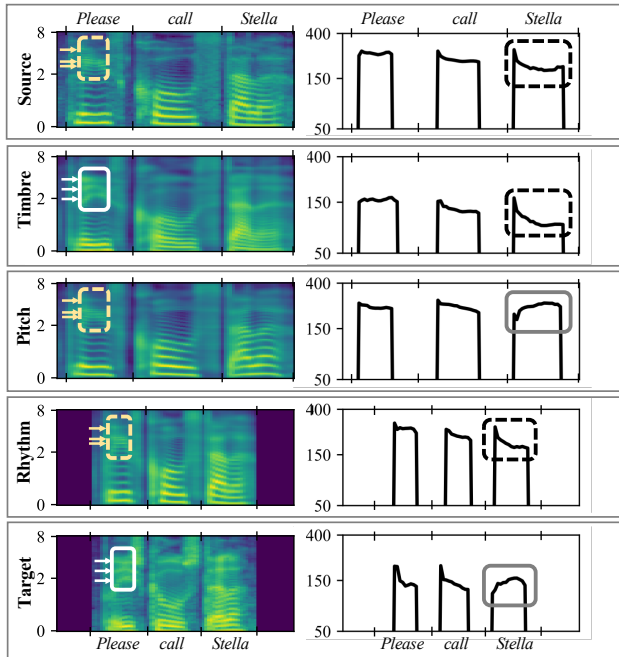


Figure 4. Spectrogram (left) and pitch contours (right) of single-aspect conversion results of the utterance ‘Please call Stella’. The left rectangle marks highlight the formant structures of the phone ‘ea’. The arrows mark the frequencies of the second, third and fourth formants. The right rectangle marks highlight the pitch tones of the word ‘Stella’.

Table 2. MOS of different conversion types/algorithms.

| Rhythm Only | Pitch Only | Timbre Only |
|--------------|---------------|--------------|
| 3.21 | 3.79 | 3.40 |
| Rhythm+Pitch | Rhythm+Timbre | Pitch+Timbre |
| 3.04 | 2.73 | 3.35 |
| All Three | AUTOVC | Source |
| 2.79 | 3.24 | 4.65 |

the corresponding aspect is converted, and low otherwise. For example, the pitch-only conversion has a high pitch conversion rate but low rhythm and timbre conversion rates; whereas the rhythm+timbre conversion has a high rhythm and timbre conversion rates but a low pitch conversion rate. It is worth noting that AUTOVC has a high timbre conversion rate, but low in the other, indicating that it only converts timbre. In short, both the visualization results and our subjective evaluation verifies that each conversion can successfully convert the intended aspects, without altering the other aspects, whereas AUTOVC only converts timbre.

We also evaluate the MOS (mean opinion score), ranging from one to five, on the quality of the conversion, as shown in table 5.2. There are a few interesting observations. First, the MOS of pitch conversion is higher than that of timbre and rhythm conversions, which implies that timbre and rhythm conversions are the more challenging tasks. Second,




Figure 5. Subjective conversion rates of different conversion types. Each bar group corresponds to a conversion type/algorithm. The three bars within each group represent the rhythm, pitch and timbre conversion rates respectively.

as the number of converted aspects increases, the MOS gets lower, because the conversion task gets more challenging.

Objective Evaluation Due to the lack of explicit labels of the speech components, it is difficult to fully evaluate the disentanglement results using objective metrics. However, we can still objectively evaluate the pitch-only conversion performance by comparing the pitch contour of the converted speech and the target pitch contour. Following Valle et al. (2020), we use three metrics for the comparison: Gross Pitch Error (GPE) (Nakatani et al., 2008), Voice Decision Error (VDE) (Nakatani et al., 2008), and F0 Frame Error (FFE) (Chu & Alwan, 2009). SPEECHFLOW achieves a GPE of 1.04%, a VDE of 8.14%, and an FFE of 8.86%. As a reference, these results are comparable with the results reported in Valle et al. (2020), with a slightly higher GPE and lower VDE and FFE. Note that these two sets of results cannot be directly compared, because the datasets are different, but they show the effectiveness of the SPEECHFLOW in disentangling pitch.

5.3. Mismatched Conversion Target

Since utterances with mismatched contents have different numbers of syllables and lengths, we would like to find out how SPEECHFLOW converts rhythm when the source and target utterances read different content. Figure 7 shows the rhythm-only conversion between a long utterance, ‘And we will go meet her Wednesday’ (top left panel), and a short




Figure 6. Subjective conversion rates of single-aspect conversions of SPEECHFLOW when the information bottleneck of the rhythm encoder (top panel) or the content encoder (bottom panel) is too wide. Each bar group corresponds to a conversion type/algorithm. The three bars within each group represent the rhythm, pitch and timbre conversion rates respectively.

utterance, ‘Please call Stella’ (top right panel).

The short to long conversion is shown in the bottom left panel. It can be observed that the conversion tries to match the syllable structure of the long utterance by stretching its limited words. In particular, ‘please’ is stretched to cover ‘and we will’, ‘call’ to cover ‘go meet’, and ‘Stella’ to cover ‘her Wednesday’. On the contrary, the long to short conversion, as shown in the bottom right panel, tries to squeeze everything to the limited syllable slots in the short utterance. Intriguingly still, the word mapping between the long utterance and the short utterance is *exactly the same* as in the short to long conversion. In both cases, the word boundaries between the converted speech and the target speech are surprisingly aligned.

These observations suggest that SPEECHFLOW has an intricate ‘fill in the blank’ mechanism when combining the rhythm information with content and pitch. The rhythm code provides a number of blanks, and the decoder fills the blanks with the content information and pitch information provided by the respective encoders. Furthermore, there seems to be an anchoring mechanism that associates the content and pitch with the right blank, which functions stably even if the blanks and the content are mismatched.

5.4. Removing Speech Components

To further understand the disentanglement mechanism of SPEECHFLOW, we generate spectrograms with one of the four components removed. To remove rhythm, content or pitch, we respectively set the input to the rhythm encoder, content encoder or pitch encoder to zero. To remove timbre, we set the speaker embedding to zero. Figure 8 shows




Figure 7. Rhythm-only conversion when the source and target speech have mismatched content.

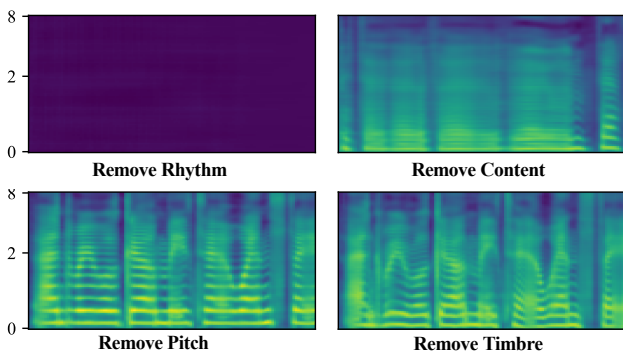


Figure 8. Reconstructed speech when one speech component is removed. The ground truth speech is in figure 7 top left panel.

the output spectrograms with one component removed. As can be observed, when the rhythm is removed, the output becomes zero, and when the content is removed, the output becomes a set of slots with no informative spectral shape. These findings are consistent with our ‘fill in the blank’ hypothesis in section 5.4. When rhythm code is removed, there is no slot to fill, and hence the output spectrogram is blank. When content is removed, there is nothing to fill in the blanks, resulting in a spectrogram with uninformative blanks. When the pitch is removed, the pitch of the output becomes completely flat, as can be seen from the flat harmonics. Finally, when timbre is removed, the formant positions of the output spectrogram shift, which indicates that the timbre has changed, possibly to an average speaker. These results further verify that SPEECHFLOW can separately model different speech components.

5.5. Varying the Information Bottleneck

In this section, we would like to verify our theoretical explanation in section 4.4 by varying the information bottleneck and see if SPEECHFLOW will still act as our theory predicts.

According to figure 2, if the physical information bottleneck

of the rhythm encoder is too wide, then the rhythm encoder will pass all the information through, and the content encoder, pitch encoder and speaker identity will be useless. As a result, the rhythm-only conversion will convert all the aspects. On the other hand, the pitch-only and timbre-only conversions will alter nothing. Similarly, if the physical information bottleneck of the content encoder is too wide, but random sampling is still present, then the content encoder will pass almost all the information through, except for the rhythm information, because the random resampling operations still contaminate the rhythm information and SPEECHFLOW would still rely on the rhythm encoder to recover the rhythm information. As a result, the rhythm-only conversion would still convert rhythm, but the pitch-only and timbre-only conversions would barely alter anything.

Figure 6 shows the subjective conversion rates of single-aspect conversions when the physical bottleneck of rhythm encoder or the content encoder is too wide. These results agree with our theoretical predictions. When the rhythm encoder physical bottleneck is too wide, the rhythm-only conversion converts all the aspects, while other conversions convert nothing. When the content encoder physical bottleneck is too wide, the rhythm-only conversion still converts rhythm. Notably, the timbre-only conversion still converts timbre to some degree, possibly due to the random resampling operation of the content encoder. These results verify our theoretical explanation of SPEECHFLOW.

6. Conclusion

We have demonstrated that SPEECHFLOW has powerful disentanglement capabilities by having multiple intricately designed information bottlenecks. There are three take-aways. First, we have shown that the physical dimension of the hidden representations can effectively limit the information flow. Second, we have verified that when information bottleneck is binding, neural autoencoder will only pass the information that other channels cannot provide. Third, even if we only have a partial disentanglement algorithm, *e.g.* the random resampling, we can still design a complete disentanglement algorithm by having multiple channels with different information bottleneck. These intriguing observations inspire a generic approach to disentanglement. As future directions, we will seek to refine the bottleneck design of SPEECHFLOW using more information-theory-guided approaches, as well as explore the application of SPEECHFLOW to low-resource speech processing systems.

Acknowledgment

We would like to give special thanks to Gaoyuan Zhang from MIT-IBM Watson AI Lab, who has helped us a lot with building our demo webpage.

References

- Atal, B. and Schroeder, M. Predictive coding of speech signals and subjective error criteria. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 27(3):247–254, 1979.
- Biadsy, F., Weiss, R. J., Moreno, P. J., Kanvesky, D., and Jia, Y. Parrottron: An end-to-end speech-to-speech conversion model and its applications to hearing-impaired speech and speech separation. *Proc. Interspeech 2019*, pp. 4115–4119, 2019.
- Choi, Y., Choi, M., Kim, M., Ha, J.-W., Kim, S., and Choo, J. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8789–8797, 2018.
- Chou, J.-c. and Lee, H.-Y. One-shot voice conversion by separating speaker and content representations with instance normalization. *Proc. Interspeech 2019*, pp. 664–668, 2019.
- Chou, J.-c., Yeh, C.-c., Lee, H.-y., and Lee, L.-s. Multi-target voice conversion without parallel data by adversarially learning disentangled audio representations. *Proc. Interspeech 2018*, pp. 501–505, 2018.
- Chu, W. and Alwan, A. Reducing F0 frame error of F0 tracking algorithms under noisy conditions with an unvoiced/voiced classification frontend. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3969–3972. IEEE, 2009.
- Gao, Y., Singh, R., and Raj, B. Voice impersonation using generative adversarial networks. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2506–2510. IEEE, 2018.
- Hsu, C.-C., Hwang, H.-T., Wu, Y.-C., Tsao, Y., and Wang, H.-M. Voice conversion from non-parallel corpora using variational auto-encoder. In *Signal and Information Processing Association Annual Summit and Conference (APSIPA), 2016 Asia-Pacific*, pp. 1–6. IEEE, 2016.
- Hsu, C.-C., Hwang, H.-T., Wu, Y.-C., Tsao, Y., and Wang, H.-M. Voice conversion from unaligned corpora using variational autoencoding Wasserstein generative adversarial networks. *Proc. Interspeech 2017*, pp. 3364–3368, 2017.
- Huang, W.-C., Hwang, H.-T., Peng, Y.-H., Tsao, Y., and Wang, H.-M. Voice conversion based on cross-domain features using variational auto encoders. In *2018 11th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, pp. 51–55. IEEE, 2018.
- Huang, W.-C., Luo, H., Hwang, H.-T., Lo, C.-C., Peng, Y.-H., Tsao, Y., and Wang, H.-M. Unsupervised representation disentanglement using cross domain features and adversarial learning in variational autoencoder based voice conversion. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2020.
- Kameoka, H., Kaneko, T., Tanaka, K., and Hojo, N. ACVAE-VC: Non-parallel many-to-many voice conversion with auxiliary classifier variational autoencoder. *arXiv preprint arXiv:1808.05092*, 2018a.
- Kameoka, H., Kaneko, T., Tanaka, K., and Hojo, N. StarGAN-VC: Non-parallel many-to-many voice conversion using star generative adversarial networks. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pp. 266–273. IEEE, 2018b.
- Kaneko, T. and Kameoka, H. Parallel-data-free voice conversion using cycle-consistent adversarial networks. *arXiv preprint arXiv:1711.11293*, 2017.
- Kaneko, T., Kameoka, H., Tanaka, K., and Hojo, N. StarGAN-VC2: Rethinking conditional methods for StarGAN-based voice conversion. *Proc. Interspeech 2019*, pp. 679–683, 2019.
- Kawahara, H., Morise, M., Takahashi, T., Nisimura, R., Irino, T., and Banno, H. TANDEM-STRAIGHT: A temporally stable power spectral representation for periodic signals and applications to interference-free spectrum, F0, and aperiodicity estimation. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3933–3936. IEEE, 2008.
- Kenter, T., Wan, V., Chan, C.-A., Clark, R., and Vit, J. CHiVE: Varying prosody in speech synthesis with a linguistically driven dynamic hierarchical conditional variational network. In *International Conference on Machine Learning*, pp. 3331–3340, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Lample, G., Zeghidour, N., Usunier, N., Bordes, A., Denoyer, L., and Ranzato, M. Fader networks: Manipulating images by sliding attributes. In *Advances in neural information processing systems*, pp. 5967–5976, 2017.
- Mermelstein, P. Distance measures for speech recognition, psychological and instrumental. *Pattern recognition and artificial intelligence*, 116:374–388, 1976.
- Nachmani, E. and Wolf, L. Unsupervised singing voice conversion. *Proc. Interspeech 2019*, pp. 2583–2587, 2019.
- Nakatani, T., Amano, S., Irino, T., Ishizuka, K., and Kondo, T. A method for fundamental frequency estimation

- and voicing decision: Application to infant utterances recorded in real acoustical environments. *Speech Communication*, 50(3):203–214, 2008.
- Niwa, J., Yoshimura, T., Hashimoto, K., Oura, K., Nankaku, Y., and Tokuda, K. Statistical voice conversion based on WaveNet. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5289–5293. IEEE, 2018.
- Pell, M. D., Jaywant, A., Monetta, L., and Kotz, S. A. Emotional speech processing: Disentangling the effects of prosody and semantic cues. *Cognition & Emotion*, 25(5): 834–853, 2011.
- Polyak, A. and Wolf, L. Attention-based WaveNet autoencoder for universal voice conversion. In *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6800–6804. IEEE, 2019.
- Qian, K., Zhang, Y., Chang, S., Yang, X., and Hasegawa-Johnson, M. AutoVC: Zero-shot voice style transfer with only autoencoder loss. In *International Conference on Machine Learning*, pp. 5210–5219, 2019.
- Qian, K., Jin, Z., Hasegawa-Johnson, M., and Mysore, G. J. F0-consistent many-to-many non-parallel voice conversion via conditional autoencoder. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6284–6288. IEEE, 2020.
- Quatieri, T. F. *Discrete-time speech signal processing: principles and practice*. Pearson Education India, 2006.
- Schroeder, M. and Atal, B. Code-excited linear prediction (CELP): High-quality speech at very low bit rates. In *1985 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 10, pp. 937–940. IEEE, 1985.
- Serrà, J., Pascual, S., and Perales, C. S. Blow: a single-scale hyperconditioned flow for non-parallel raw-audio voice conversion. In *Advances in Neural Information Processing Systems*, pp. 6790–6800, 2019.
- Shechtman, S. and Sorin, A. Sequence to sequence neural speech synthesis with prosody modification capabilities. In *Proc. 10th ISCA Speech Synthesis Workshop*, pp. 275–280, 2019.
- Skerry-Ryan, R., Battenberg, E., Xiao, Y., Wang, Y., Stanton, D., Shor, J., Weiss, R., Clark, R., and Saurous, R. A. Towards end-to-end prosody transfer for expressive speech synthesis with Tacotron. In *International Conference on Machine Learning*, pp. 4693–4702, 2018.
- Valle, R., Li, J., Prenger, R., and Catanzaro, B. Mellotron: Multispeaker expressive voice synthesis by conditioning on rhythm, pitch and global style tokens. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6189–6193. IEEE, 2020.
- Veaux, C., Yamagishi, J., MacDonald, K., et al. Superseded-CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit. *University of Edinburgh. The Centre for Speech Technology Research (CSTR)*, 2016.
- Wu, Y. and He, K. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 3–19, 2018.
- Xu, Y. and Wang, Q. E. Pitch targets and their realization: Evidence from mandarin chinese. *Speech communication*, 33(4):319–337, 2001.
- Yamamoto, R., Felipe, J., and Blaauw, M. r9y9/pysptk: 0.1. 14, 2019.
- Zhang, Y., Ou, Z., and Hasegawa-Johnson, M. Improvement of probabilistic acoustic tube model for speech decomposition. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7929–7933. IEEE, 2014.