# SINGING VOICE CONVERSION WITH DISENTANGLED REPRESENTATIONS OF SINGER AND VOCAL TECHNIQUE USING VARIATIONAL AUTOENCODERS 

Yin-Jyun Luo ${ }^{1,3}$ Chin-Cheng Hsu ${ }^{2} \quad$ Kat Agres ${ }^{3,4} \quad$ Dorien Herremans ${ }^{1,3}$

${ }^{1}$ Singapore University of Technology and Design<br>${ }^{2}$ University of Southern California, Los Angeles, United States<br>${ }^{3}$ Institute of High Performance Computing, A*STAR, Singapore<br>${ }^{4}$ Yong Siew Toh Conservatory of Music, National University of Singapore


#### Abstract

We propose a flexible framework that deals with both singer conversion and singers vocal technique conversion. The proposed model is trained on non-parallel corpora, accommodates many-to-many conversion, and leverages recent advances of variational autoencoders. It employs separate encoders to learn disentangled latent representations of singer identity and vocal technique separately, with a joint decoder for reconstruction. Conversion is carried out by simple vector arithmetic in the learned latent spaces. Both a quantitative analysis as well as a visualization of the converted spectrograms show that our model is able to disentangle singer identity and vocal technique and successfully perform conversion of these attributes. To the best of our knowledge, this is the first work to jointly tackle conversion of singer identity and vocal technique based on a deep learning approach.


Index Terms- singing voice conversion, vocal technique, variational autoencoders, disentangled representations

## 1. INTRODUCTION

Singing voice conversion (SVC) comprehensively refers to tasks that convert an attribute of singing. Converting from one singer's voice to that of another without affecting linguistic content has been the focus in SVC [1, 2, 3, 4, 5]. Converting between different vocal techniques, however, is a worthwhile line of research that has lacked attention. Such an approach would allow one to convert a singing voice into a different timbre or pitch that was originally infeasible due to physical constraints or lack of singing skills, thereby facilitating applications in entertainment and pedagogy.

Vocal techniques, such as 'breathy' and 'vibrato', enrich the sound and are an integral part of singing. Singers perform different techniques at different points in time so as to create emotional ebbs and flows. Modeling vocal techniques

[^0]

Fig. 1: The proposed framework, fully detailed in Section 3.2. The blue, red and green blocks correspond to the singer encoder, vocal technique encoder, and the joint decoder, respectively.
through data-driven models is challenging due to lack of labeled and balanced data, together with intrinsic ambiguity, just to name a few.

We propose a framework that deals with the conversion of both singer identity and vocal technique. We augment the model based on the Deep Bi-directional Long Short-Term Memory (DBLSTM) from [6] with latent variables, such that it learns disentangled representations for both singer and vocal technique through Gaussian mixture variational autoencoders (GMVAEs) [7, 8]. Unlike typical SVC models that condition generation of singing voice on an utterance-level singer label [9, 5, 4], our model is conditioned on timedependent singer/technique variables on a shorter temporal scale, accommodating cases in which the vocal technique varies across time. The proposed model can be trained on non-parallel (i.e., unpaired) corpora, and allows for many-tomany conversion of singer identities and vocal techniques.

We describe our modified implementation of the GMVAE model $[10,11]$ along with our singer/vocal technique conversion strategy in Section 2. Next, we elaborate on our experimental setup using VocalSet [12], a dataset featuring signing techniques, in Section 3. Finally, we report and discuss the experimental results in Section 4.

## 2. METHOD

### 2.1. Variational Autoencoders

Our proposed singing voice generation process $(\mathbf{z} \rightarrow \mathbf{X})$ based on VAE includes the generation of a chunk of spectrogram $\mathbf{X} \in \mathcal{R}^{T \times F}$ which is generated from a latent variable $\mathbf{z} \in \mathcal{R}^{D}$. This simple dependency structure allows us to apply variational inference, which optimizes the evidence lower bound (ELBO) of $p(\mathbf{X})$ :

$$
\begin{equation*}
\mathcal{L}(p, q ; \mathbf{X})=\mathbb{E}_{q(\mathbf{z} \mid \mathbf{X})}[\log p(\mathbf{X} \mid \mathbf{z})]-D_{K L}(q(\mathbf{z} \mid \mathbf{X}) \| p(\mathbf{z})), \tag{1}
\end{equation*}
$$

where we assume $p(\mathbf{z})=\mathcal{N}(0, I), p(\mathbf{X} \mid \mathbf{z})=\mathcal{N}\left(\mu_{x}, \mathcal{I}\right)$ and $q(\mathbf{z} \mid \mathbf{X})=\mathcal{N}\left(\mu_{z}, \operatorname{diag}\left(\boldsymbol{\sigma}_{z}\right)\right)$, in which $\mu_{z}$ and $\boldsymbol{\sigma}_{z}$ are inferred from $\mathbf{X}$ by an encoder, and $\mu_{x}$ is predicted by a decoder. Reconstructing $\mathbf{X}$ from $\mathbf{z}$ which is inferred from $\mathbf{X}$ itself using variational methods thus concludes a VAE.

### 2.2. Increasing Expressivity using a GMM Prior

The above made assumption $p(\mathbf{z})=\mathcal{N}(0, I)$ reflects the preference for a simple distribution of data, which in turn sacrifices model expressivity. Replacing $p(\mathbf{z})$ with a Gaussian mixture model (GMM), known as GMVAEs, has been proven effective in increasing the expressivity and controllability [7, $8,10,11]$. It has an additional layer of dependency: $\mathbf{y} \rightarrow \mathbf{z}$ that enables us to utilize categorical attributes $\mathbf{y}$ that may be available in data. The ELBO of a GMVAE then becomes:
$\mathcal{L}(p, q ; \mathbf{X}, \mathbf{y})=\mathbb{E}_{q(\mathbf{z} \mid \mathbf{X})}[\log p(\mathbf{X} \mid \mathbf{z})]-D_{K L}(q(\mathbf{z} \mid \mathbf{X})| | p(\mathbf{z} \mid \mathbf{y}))$,
where the prior $p(\mathbf{z})$ is now multi-modal (GMM), more likely to model data with higher diversity. In addition, introducing y endows the model with direct controllability over attributes and flexibility for generation, as will be elaborated next.

### 2.3. Controlling Singer Identities and Vocal Techniques

Using the proposed GMVAE, the generation process for singing voice is as follows: given a singer $\mathbf{y}_{s}$ and a vocal technique $\mathbf{y}_{t}$ (collectively referred to as attributes), the model first infers latent representations ( $\mathbf{z}_{s}$ and $\mathbf{z}_{t}$ ) of each of the attributes, and then combines these two to generate a spectrogram of the singing voice. Mathematically, the joint probability, given the attributes can be factorized as follows:

$$
\begin{equation*}
p\left(\mathbf{X}, \mathbf{z}_{s}, \mathbf{z}_{t} \mid \mathbf{y}_{s}, \mathbf{y}_{t}\right)=p\left(\mathbf{X} \mid \mathbf{z}_{s}, \mathbf{z}_{t}\right) p\left(\mathbf{z}_{s} \mid \mathbf{y}_{s}\right) p\left(\mathbf{z}_{t} \mid \mathbf{y}_{t}\right) \tag{3}
\end{equation*}
$$

where both the conditional distributions $p\left(\mathbf{z}_{s} \mid \mathbf{y}_{s}\right)$ and $p\left(\mathbf{z}_{t} \mid \mathbf{y}_{t}\right)$ are assumed to be Gaussian with learnable means and diagonal covariances. This GMVAE model now takes into consideration singer and vocal technique and thus can directly control them during the conversion phase.

### 2.4. Learning an Attribute-discriminative Space

We incorporate two classifiers, one for vocal techniques and the other for singers, to encourage the learned spaces $\mathbf{z}_{t}$ and $\mathbf{z}_{s}$ to be discriminative. Each classifier learns to predict $\mathbf{y}_{*}$ from the sequence of $\mathbf{z}_{*}=\left\{\mathbf{z}_{* 1}, \mathbf{z}_{* 2}, \ldots, \mathbf{z}_{* N}\right\}$, where $N$ is the number of chunks of a recording and $*$ denotes singers or vocal techniques. The classifier receives a sequence level representation that can be summarized by a simplified attention mechanism [13, 14]:

$$
\begin{equation*}
\alpha_{* n}=\frac{e^{f\left(\mathbf{z}_{* n}\right)}}{\sum_{m=1}^{N} e^{f\left(\mathbf{z}_{* m}\right)}}, \mathbf{c}_{*}=\sum_{n=1}^{N} \alpha_{* n} \mathbf{z}_{* n} \tag{4}
\end{equation*}
$$

where $f(\cdot)$ is a learnable function and $\mathbf{c}_{*}$ denotes the summarization (and thus the representation) of the input sequence X. $\alpha_{* n}$ is uniformly distributed without the attention module. The auxiliary objective of maximizing $p\left(\mathbf{y}_{*} \mid \mathbf{X}\right)$ will thus be added to the overall objective.

### 2.5. Weighting KLDs

A singer might be unable to maintain the same level of expressiveness of a technique throughout a recording. Similarly, voice timbre and pitch also vary across a recording even though a singer is asked to perform the same vocal technique. Based on these observations, we may benefit from weighting the KLD terms of (2) with $\alpha_{* n}$ obtained from Section 2.4. Consummating our training objective, we have

$$
\begin{align*}
\mathcal{L}\left(p, q ; \mathbf{X}, \mathbf{y}_{s},\right. & \left.\mathbf{y}_{t}\right) \\
\quad & =\mathbb{E}_{q\left(\mathbf{z}_{s} \mid \mathbf{X}\right) q\left(\mathbf{z}_{t} \mid \mathbf{X}\right)}\left[\log p\left(\mathbf{X} \mid \mathbf{z}_{s}, \mathbf{z}_{t}\right)\right] \\
& -\sum_{n=1}^{N} \alpha_{t n} D_{K L}\left(q\left(\mathbf{z}_{t n} \mid \mathbf{X}\right) \| p\left(\mathbf{z}_{t} \mid \mathbf{y}_{t}\right)\right)  \tag{5}\\
& -\sum_{n=1}^{N} \alpha_{t s} D_{K L}\left(q\left(\mathbf{z}_{s n} \mid \mathbf{X}\right) \| p\left(\mathbf{z}_{s} \mid \mathbf{y}_{s}\right)\right) \\
+ & \beta \log p\left(\mathbf{y}_{s} \mid \mathbf{X}\right)+\gamma \log p\left(\mathbf{y}_{t} \mid \mathbf{X}\right)
\end{align*}
$$

where $\beta$ and $\gamma$ are weights for the discriminative objectives.

### 2.6. Conversion Strategies

We accommodate the time-varying singing attributes in expressive singing voices by learning latent variables at a shorter temporal scale. Consequently, we adopt the model to infer the attributes, rather than assigning attributes directly during the conversion phase. ${ }^{1}$ Generally, conversion is done by

[^1]adding a conversion vector $d \mu_{* n}$ to $\mathbf{z}_{* n}$ at the chunk level. We define $d \mu_{* n}=\mu_{*}^{\text {target }}-\mu_{* n}^{\text {source }}$, where $\mu_{*}$ represents the mean of a Gaussian mixture component. $\mu_{* n}^{\text {source }}$ can be determined by either Gaussian likelihood $p\left(\mathbf{y}_{* n} \mid \mathbf{z}_{* n}\right)$ or the auxiliary sequence-level classifier $p\left(\mathbf{y}_{*} \mid \mathbf{X}\right)$, referred to as $C$ chunk and $C$-sequence, respectively. Note that the latter computes a common $\mu_{* n}^{\text {source }}$ (and hence $d \mu_{* n}$ ) over all $N$ chunks. We report the result from both methods below.

## 3. EXPERIMENTAL SETTINGS

### 3.1. Dataset

We use the VocalSet [12] to evaluate our framework. A subset of audio files we selected has 20 singers, 6 vocal techniques that were most distinguishable, (straight, breathy, vibrato, belt, lip trill and vocal fry), and 5 vowels. Each recording was sung as either a scale or arpeggios. The subset was divided into a training set of 1,065 recordings, and a testing set of 118 ( 17 out of the $20 \times 6 \times 2 \times 5=1200$ combinations are missing). The length of recordings ranges from 3.5 to 23 seconds. This subset approximates a balanced number of instances over classes for singers, techniques, and vowels.

We re-sampled the recordings to $22,050 \mathrm{~Hz}$, normalized the waveform w.r.t. the largest magnitude, computed the log magnitude Mel-spectrogram (MEL) with $F=96$ filter banks, and then rescaled it to $[-1,1]$. We further segmented the MEL into chunks of $T=43$ frames ( $\mathbf{X} \in \mathcal{R}^{43 \times 96}$ ). A frame shift of 256 was used for computing the MELs, so that 43 frames amount to 0.5 seconds.

### 3.2. Architecture

Our SVC framework encompasses seven components: a feature extraction network (FEN), two encoders, a decoder, a post-processor, and two sequence-level classifiers. The overall architecture is shown Fig. 1

The FEN is composed of a two-layer one-dimensional convolutional neural network (CNN), each with 512 filters ( $3 \times 1$ ), followed by two fully-connected (FC) layers with 512 and 256 units respectively. Batch normalization followed by ReLu is used for every layer. The FEN produces a 256 dimension bottleneck feature for a given input MEL chunk, and is shared and consumed by both encoders that follow.

Both of the encoders are parameterized as two-layer Recurrent Neural Networks (RNNs): a BLSTM with 256 hidden units, followed by two FC layers shared across time which predict $\mu_{*}$ and $\log \sigma_{*}$. $\mathbf{z}_{*}$ is then sampled using the reparameterization trick from [15]. The joint decoder has a similar architecture as the encoders but in reverse order. At each time step of the output sequence, a CNN with an architecture that is symmetric to the FEN is employed to reconstruct the MEL chunk. Finally, the refinement network, a three-layer onedimensional CNN with 512 filters ( $3 \times 1$ ), is used to refine the reconstructed MELs.

### 3.3. Hyperparameters

The mean vectors of $p\left(\mathbf{z}_{*} \mid \mathbf{y}_{*}\right)$ were all randomly initialized, whereas the variance vectors were kept fixed with value $e^{-2}$. The number of mixtures for singers is set to 20, and set to 6 for vocal techniques. Both are equal to the number of classes they correspond to. We set the batch size to 128 , and initialized the model parameters with Xavier initialization [16]. The Adam optimizer [17] was used with a learning rate of $10^{-4}$.

### 3.4. Evaluation Metrics

We evaluate our model by how well a classifier correctly recognizes the attributes in the converted MELs. The main idea is that the converted MELs should be accurately classified as the target class, and attributes that are not intended to be converted should be predicted the same. The classification results thereby reveal the effects on the output MELs caused by conversion. Each recording in the test set is first converted to all possible target attributes and then evaluated by three classifiers that recognize singer, singing technique, or vowel. These classifiers have the same architecture as the combination of FEN, RNNs, and the attention module, and are trained independently from unconverted MELs.

## 4. RESULTS

### 4.1. Recognizing Attributes from Converted MELs

We compare three variants of the proposed models: M1 denotes the model trained with neither the attention module nor the the discriminative objectives $(\beta=\gamma=0) . M 2$ is similar to M1 but with $\beta=\gamma=1$. Finally, M3 is the model equipped with the attention module. The performance of M0 servers as the upper bound for classification results using unconverted MELs as input. The results are listed in Table 1. Higher numbers represent better performance for all cases.

We summarize our findings as follows: First, the attributes in the MELs by M2 can be recognized with higher accuracy than those from M1 in all cases; this supports our belief that including discriminative objectives helps the model to disentangle certain attributes from the input. Second, results from M2 and M3 are similar, but M2 is better at singer conversion, whereas M3 is slightly better at converting vocal techniques. Third, there is no noticeable difference in performance between the two conversion strategies C -chunk and C -sequence. Fourth, converting vocal techniques is much more challenging than converting singer identities as the accuracy drops from $90.68 \%$ to below $67.94 \%$ after conversion, even though the number of techniques is fewer.

### 4.2. Visualization

We visualise some examples of converted MELs in Fig. 2. In the upper panel (a), it is clear that the overall pitch level

| Strategy | Model | Effect of Singer Conversion |  |  |  |  |  | Effect of Technique Conversion |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | *Singer |  | Technique |  | Vowel |  | Singer |  | *Technique |  | Vowel |  |
|  |  | Before | After | Before | After | Before | After | Before | After | Before | After | Before | After |
|  | M0 | 89.83 | NA | 90.68 | NA | 77.97 | NA | 89.83 | NA | 90.68 | NA | 77.97 | NA |
| C-chunk | M1 | 80.51 | 63.35 | 83.05 | 75.51 | 77.97 | 69.24 | 80.51 | 75.99 | 83.05 | 54.38 | 77.97 | 72.74 |
|  | M2 | 87.29 | 76.95 | 83.90 | 76.99 | 72.03 | 66.78 | 87.29 | 81.92 | 83.90 | 65.82 | 72.03 | 71.33 |
|  | M3 | 83.05 | 75.68 | 88.98 | 79.32 | 73.73 | 72.88 | 83.05 | 84.18 | 88.98 | 67.66 | 73.73 | 71.47 |
| C-sequence | M2 | 87.29 | 76.65 | 83.90 | 76.64 | 72.03 | 66.69 | 87.29 | 82.06 | 83.90 | 65.64 | 72.03 | 71.33 |
|  | M3 | 83.05 | 75.47 | 88.98 | 79.62 | 73.73 | 72.83 | 83.05 | 84.04 | 88.98 | 67.94 | 73.73 | 72.03 |

Table 1: The classification accuracy (\%) derived by the three attribute classifiers, given different models. * denotes the converted attributes.


Fig. 2: Examples of singer conversion (a) and vocal technique conversion (b), converted by the model M3. The first column refers to source, and the rest correspond to different targets. Targets that are the same as the sources are faded.
changes when doing cross-gender conversion. For technique conversion (b), we can see that converting the lip trill to, e.g., straight, makes the spectrogram less flattened. On the other hand, we can decorate a straight tone by converting it to a bright and energized vocalization as can be seen in straightbelt conversion, or to one with periodic frequency modulation as shown in straight-vibrato conversion. Noticeable effects are also observed when the targets are lip trill and vocal fry.

Despite the change of spectral distribution, the overall pitch contours are retained in all source-target pairs. This hints towards the model's ability to perform many-to-many singer identity and vocal technique conversion.

Conversion at chunk-level enables us to morph from source to target by gradually increasing the conversion vector from 0 to $d \mu_{* n}$. As such, we can, e.g., convert a straight tone to gradually express another technique over time. This has not yet been seen in other existing SVC frameworks, and we leave further investigation for future research.

## 5. RELATED WORK

Recent advances in deep learning have brought great success to VC [18, 19, 20]; SVC, on the other hand, have not benefited from it until recently. SINGAN [9], an SVC framework based on deep generative models has been proposed to map acous-
tic features of a source singer to that of a target one; however, the model is restricted to temporally-aligned singing recordings. In contrast, [5] proposed to combine automatic speech recognition for SVC that is trainable from non-parallel corpora. The model, however, only allows for converting a handful of source singers to a single target. Recently, an encoderdecoder model is proposed which incorporates a domain confusion network [21] to learn singer-agnostic features [4].

We distinguish ourselves by jointly modeling singer and technique with a principled probabilistic generative model, and conditioning the generation of singing voice on timedependent latent variables of the aforementioned attributes. To the best of our knowledge, this is the first study on jointly modeling/converting singer identity and vocal technique with a single deep learning model.

## 6. CONCLUSION AND FUTURE WORK

We have proposed a flexible framework based on GMVAEs to tackle non-parallel, many-to-many SVC for singer identity and vocal technique. Audio samples are available from https://reurl.cc/oD5vjQ. Analyzing the temporal dynamics of the latent variables, as well as accommodating the dependency between singer identity and vocal technique variables will be the focus of our future work.

## 7. REFERENCES

[1] F. Villavicencio and J. Bonada, "Applying voice conversion to concatenative singing-voice synthesis," in $I N$ TERSPEECH, 2010.
[2] H. Doi, T. Toda, T. Nakano, M. Goto, and S. Nakamura, "Singing voice conversion method based on many-tomany eigenvoice conversion and training data generation using a singing-to-singing synthesis system," in APSIPA, 2012, pp. 1-6.
[3] K. Kobayashi, T. Toda, G. Neubig, S. Sakti, and S. Nakamura, "Statistical singing voice conversion with direct waveform modification based on the spectrum differential," in INTERSPEECH, 2014.
[4] E. Nachmani and L. Wolf, "Unsupervised singing voice conversion," arXiv preprint, 2019.
[5] X. Chen, W. Chu, Guo J, and Xu N, "Singing voice conversion with non-parallel data," in IEEE Conf. on Multimedia Information Processing and Retrieval. IEEE, 2019, pp. 292-296.
[6] L. Sun, S. Kang, K. Li, and H. Meng, "Voice conversion using deep bidirectional long short-term memory based recurrent neural networks," in ICASSP, 2015, pp. 48694873.
[7] N. Dilokthanakul, P. A. M. Mediano, M. Garnelo, M. C.-H. Lee, H. Salimbeni, K. Arulkumaran, and M. Shanahan, "Deep unsupervised clustering with gaussian mixture variational autoencoders," arXiv preprint arXiv:1611.02648, 2016.
[8] Z. Jiang, Y. Zheng, H. Tan, B. Tang, and H. Zhou, "Variational deep embedding: An unsupervised and generative approach to clustering," in Int. Joint Conf. on Artificial Intelligence, 2017.
[9] B. Sisman, K. Vijayan, M. Dong, and H. Li, "SINGAN: Singing voice conversion with generative adversarial networks," in APSIPAA, 2019.
[10] W.-N. Hsu, R. J. Weiss Y. Zhang, H. Zen, Y. Wu, Y. Wang, Y. Cao, Y. Jia, Z. Chen, J. Shen, P. Nguyen, and R. Pang, "Hierarchical generative modeling for controllable speech synthesis," in $I C L R, 2019$.
[11] Y.-J. Luo, K. Agres, and D. Herremans, "Learning disentangled representations of timbre and pitch for musical instrument sounds using gaussian mixture variational autoencoders," in ISMIR, 2019.
[12] J. Wilkins, P. Seetharaman, A. Wahl, and B. Pardo, "Vocalset: A singing voice dataset," in ISMIR, 2018, pp. 468-474.
[13] S. K. Sønderby, C. K. Sønderby, H. Nielsen, and O. Winther, "Convolutional lstm networks for subcellular localization of proteins," in Int. Conf. on Algorithms for Computational Biology. Springer, 2015, pp. 68-80.
[14] C. Raffel and D. P. Ellis, "Feed-forward networks with attention can solve some long-term memory problems," in ICLR, workshop track, 2016.
[15] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in ICLR, 2014.
[16] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in $A I S$ TATS, 2010, pp. 249-256.
[17] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in $I C L R, 2015$.
[18] L. Sun, S. Kang, K. Li, and H. Meng, "Voice conversion using deep bidirectional long short-term memory based recurrent neural networks," in ICASSP. IEEE, 2015, pp. 4869-4873.
[19] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.M. Wang, "Voice conversion from non-parallel corpora using variational auto-encoder," in APSIPA, 2016, pp. 1-6.
[20] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "Stargan-vc: non-parallel many-to-many voice conversion using star generative adversarial networks," in 2018 IEEE Spoken Language Technology Workshop (SLT), 2018, pp. 266-273.
[21] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Mario, and V. Lempitsky, "Domain-adversarial training of neural networks," The Journal of Machine Learning Research, vol. 17, no. 1, pp. 2096-2030, 2016.


[^0]:    This work is supported by a SINGA provided by the A*STAR, under reference number SING-2018-01-1270.

[^1]:    ${ }^{1}$ Empirically, we found training a model conditioned on an utterance-level technique label did not work well, and we resorted to the proposed method.

