CONTENT BASED SINGING VOICE EXTRACTION FROM A MUSICAL MIXTURE

Pritish Chandna¹, Merlijn Blaauw¹, Jordi Bonada¹, Emilia Gómez^{1,2}

¹Music Technology Group, Universitat Pompeu Fabra, Barcelona, Spain

²Joint Research Centre, European Commission, Seville, Spain

ABSTRACT

We present a deep learning based methodology for extracting the singing voice signal from a musical mixture based on the underlying linguistic content. Our model follows an encoder-decoder architecture and takes as input the magnitude component of the spectrogram of a musical mixture with vocals. The encoder part of the model is trained via knowledge distillation using a teacher network to learn a content embedding, which is decoded to generate the corresponding vocoder features. Using this methodology, we are able to extract the unprocessed raw vocal signal from the mixture even for a processed mixture dataset with singers not seen during training. While the nature of our system makes it incongruous with traditional objective evaluation metrics, we use subjective evaluation via listening tests to compare the methodology to state-of-the-art deep learning based source separation algorithms. We also provide sound examples and source code for reproducibility.

Index Terms— Source separation, singing voice, content disentangling, knowledge distillation, AutoVC.

1. INTRODUCTION

Source separation, or the process of isolating individual signals from a mixture of such signals has been long studied, traditionally with statistical approaches like non-negative matrix factorization (NMF) [1], principal component analysis [2] and independent component analysis (ICA) [3]. In the last few years, methodologies based on deep learning have been proposed, leading to a significant advance in the field. Source separation has diverse applications across fields; medical imaging, image processing and financial modelling among others. In music, the most useful application is that of separating the lead vocals from a musical mixture. This problem is well researched and numerous deep learning based models have recently been proposed to tackle it [4, 5, 6, 7, 8, 9, 10, 11]. Most of these models use the neural network to predict soft time frequency masks, given

an input magnitude spectrogram of the mixture signal. This mask is then applied to the magnitude spectrogram to isolate the desired signal, which is re-synthesised by using the phase of the mixture spectrogram. The output of such algorithms is the processed vocal signal, sometimes referred to as the *stem* [12].

The *stem* is the raw audio signal with linear and non-linear effects such as reverb, compression, delay and equalization among others. We recently proposed a model [13], which re-synthesises the underlying raw vocal signal present in a musical signal by using a deep neural network architecture as a function approximator that predicts vocoder features pertaining to the vocal signal. The model outperformed state-of-the-art source separation algorithms in terms of isolation of the signal from the backing track.

In this paper, we present a methodology for synthesising raw, unprocessed vocals from a musical mixture based on human cognition; i.e., we first extract a representation of the linguistic content present in the mixture and then generate the vocal signal based on this content. The linguistic content pertains primarily to cognitively relevant features such as phonemes, however we do not explicitly predict the phonemes present in the signal. Although we did explore this possibility, we were hindered by the unavailability of a sufficiently large annotated dataset for training and testing our model.

Instead, we use a representation of the linguistic features as presented in AutoVC [14]. This representation, described in the paper as the content embedding, does not require explicit phoneme annotations and can be extracted from the clean speech signal. We use the AutoVC model as a teacher network for knowledge distillation to train a network to learn the content embedding from an input mixture spectrogram. A decoder is then trained to generate vocoder features given this content embedding, which are then used for vocal synthesis. While we acknowledged in our previous work that the use of a mel-spectrogram based wavenet vocoder [15] would lead to an improvement in synthesis quality, we believe that the ease and speed of use of vocoder features sufficiently offsets the slight degradation in audio quality for research purposes. As such, we consider the raw vocals re-synthesised by the WORLD vocoder as the upper limit for the performance of our algorithm.

The TITANX used for this research was donated by the NVIDIA Corporation. This work is partially supported by the Towards Richer Online Music Public-domain Archives (TROMPA H2020 770376) project. The dataset used for training was provided by Yamaha Corporation.

We tried our algorithm for both singing with musical mixtures and unison choir singing, which includes several singers singing at similar pitches and found the methodology to work effectively for both scenarios. In this paper, we only present the musical mixtures case, but present online examples for both use cases. We evaluate our network in terms of intelligibility of the output, isolation from the backing track and audio quality against a state-of-the-art source separation baseline via a subjective listening test. Since the output of our system is a re-synthesis of the original raw vocal track, we found the use of the standard *SDR* metric [16] to be ineffective for objective evaluation.

The source code for our model is available online¹, as are sound examples², showcasing the robustness of the model with use on real-world examples and on unison choir singing.

2. RELATED WORK

U-net [7] is one of the best performing algorithms for the task singing voice separation. The U-Net model takes the magnitude spectrogram of the mixture as input and applies a series of convolutional operations, resulting in a lower dimensional embedding. This lower dimensional embedding is then up-sampled, following the steps of the encoder, to generate an output with the same dimensions as the input. There are skip connections between the corresponding layers of the encoder and the decoder, leading to the nomenclature of the U-Net. The output is then treated as a mask, which is applied to the input spectrogram, resulting in the desired source, in this case: the singing voice. This model assumes that the mixture is a linear sum of the individual sources and the difference between the input spectrogram and the masked input is a representation of the magnitude spectrum of the backing track. Both sources are re-synthesised using the phase of the original spectrogram.

The Unet architecture has also been applied to directly estimating the waveform [11], thus avoiding the phase estimation problem. This model also estimates the *stem* of the sources and follows the energy conservation assumption that the mixture is a linear sum of the individual *stems*. Unlike these models, our proposed model generates a version of the raw unprocessed vocal track and not the *stem*.

The work closest to ours is [17], which uses an end-to-end sequence-to-sequence model to extract the loudest signal from a mix of speech signal. While it is useful in a source extraction scenario, the primary application of this is voice conversion and it requires a dataset of parallel paired input-output speech utterances.

For our proposed methodology, we require a representation of linguistic information. Such representations are often used for speaker conversion systems or low resource speech synthesis systems. Algorithms such as [17, 18, 19, 20, 21]

use different types of content representations, learned from a clean input voice signal. For our study, we use the speaker-dependent version of AutoVC [14], which is one of the most effective speaker conversion algorithms. This algorithm follows an encoder-decoder architecture and imposes restrictions on the size of the bottleneck of the encoder, thereby constraining it to learn only content based information. As shown in Equation 1, the content encoder E_{avc} , takes as input the mel-spectrogram of the input speech signal, X , along with the speaker identity as a one-hot vector, S to produce the content embedding, C_{avc} . The decoder of the network, D_{avc} then takes the learned content embedding, C_{avc} along with the speaker identity S to produce the corresponding output mel-spectrogram, $\hat{X}$, as shown in Equation 2.

$$C_{avc} = E_{avc}(X, S) \quad (1)$$

$$\hat{X} = D_{avc}(C_{avc}, S) \quad (2)$$

The entire network is trained to minimize the reconstruction loss, $\mathcal{L}_{recon}$, shown in Equation 3.

$$\mathcal{L}_{recon} = \mathbb{E}[\|\hat{X} - X\|_2^2] \quad (3)$$

An additional content loss, $\mathcal{L}_{content}$, shown in Equation 4 is added to ensure that the content embedding for the output matches that of the input.

$$\mathcal{L}_{content} = \mathbb{E}[\|E_{avc}(\hat{X}, S) - C_{avc}\|_1] \quad (4)$$

This leads to the final loss for the model, $\mathcal{L}_{avc}$, shown in Equation 5

$$\mathcal{L}_{avc} = \mathcal{L}_{recon} + \lambda \mathcal{L}_{content} \quad (5)$$

where λ is a weight applied to the content loss, and is set to 1.

3. METHODOLOGY

The basic pipeline of our system is shown in Figure 1. We first train the AutoVC network on clean vocoder features extracted from singing voice signals to encode the content embedding as described in [14]. While the original AutoVC model uses the mel-spectrogram as a representation of the vocal signal, we use WORLD vocoder [22] features. The reason for this is two-fold; firstly re-synthesis from vocoder features is easier than from mel-spectrogram features and the use of WORLD vocoder [22] features allows us to disentangle linguistic content from expressive content, present in the fundamental frequency. The encoder and decoder of the network both take the singer identity as a one-hot vector, following the methodology suggested in [14]. We tested the voice change capability of this network and found it to be effective for changing the timbre but not the expression, thereby fulfilling our criteria. The network is trained with same loss functions as used in AutoVC, shown in Equation 5, with the exception that X

¹https://github.com/pc2752/sep_content/

²https://pc2752.github.io/sep_content/

refers to the vocoder features pertaining to the voice signal, instead of the mel-spectrogram.

Once we have the AutoVC network trained, we use it as a teacher network for training a singer independent encoder, as shown in Figure 1 b). This encoder, E_{spec} has an architecture similar to E_{avc} and takes the magnitude spectrogram, M of the mixture signal as input and outputs a content embedding, C_{spec} , that matches the corresponding content embedding, C_{avc} . As the singer identity is not provided to the encoder, it is singer independent. The corresponding loss function is shown in Equation 6.

$$\begin{aligned} C_{spec} &= E_{spec}(M) \\ \mathcal{L}_{encoder} &= \mathbb{E}[\|C_{spec} - C_{avc}\|_1] \end{aligned} \quad (6)$$

The decoder part of the model, D_{sdn} , takes as input the content embedding, C_{spec} , along with singer identity as a one-hot vector, S , and outputs the clean vocoder features, $\hat{X}_{sdn}$ of the corresponding vocal track. Equation 7 shows the loss function for this network. The architecture for the decoder is the same as D_{avc} . This system is henceforth referred to as the singer dependent network (SDN).

$$\begin{aligned} \hat{X}_{sdn} &= D_{sdn}(C_{spec}, S) \\ \mathcal{L}_{D_{sdn}} &= \mathbb{E}[\|\hat{X}_{sdn} - X\|_2^2] \end{aligned} \quad (7)$$

Finally, we train a singer independent decoder, D_{sin} , which takes the content embedding, C_{spec} , of the previously trained singer independent encoder along with the mixture spectrogram as input and outputs the corresponding vocoder features, $\hat{X}_{sin}$. Figure 1 c) shows the data-flow used in this model, while the loss function is represented by Equation 8. This system is henceforth referred to as the singer independent network (SIN).

$$\begin{aligned} \hat{X}_{sin} &= D_{sin}(C_{spec}, M) \\ \mathcal{L}_{D_{sin}} &= \mathbb{E}[\|\hat{X}_{sin} - X\|_2^2] \end{aligned} \quad (8)$$

While the training part of the networks is inter-related, the feed-forward prediction of an individual network is independent of the other networks

For fundamental frequency estimation, we use a network similar to the one proposed in [23], trained on the same data as the other networks. We use this prediction along with the vocoder features to synthesise the audio signal. We tried both the discrete representation of the fundamental frequency as described in [23] and a continuous representation, normalised to the range 0 to 1 as used in [13] and found that while the discrete representation leads to slightly higher accuracy in the output, the continuous representation produces a pitch contour perceptually more suitable for synthesis of the signal.

Fig. 1. The three networks used in our study. Part a) shows the AutoVC [14] network. The content embedding learned by this network is used to train the encoder part of the second network, SDN, shown in part b). The encoder for this network is singer independent, but the decoder is singer dependent. The embedding from this encoder, along with the mixture spectrogram is passed to a third network, SIN, shown in part c). Both the encoder and decoder for this network are independent of the singer.

4. EXPERIMENTS

4.1. Dataset

We use a proprietary dataset for training the model and the MedleyDB dataset [12] for evaluation of the model. The training set consists of 205 songs by 45 distinct male and female singers, with a total of around 12 hours of data. The songs are mostly pop songs in the English and Japanese languages. We use 90 % of the proprietary dataset for training and 10 % for validation, which we use for early stopping during training of the model. We have access to the raw vocal track in this training set as well as annotation of the singers, which makes it ideal for our proposed model.

For testing, we use the the MedleyDB dataset, which contains 122 songs. The raw audio tracks and the mixing stems for which are present. We use the raw audio vocal tracks of 6 of the songs for computing the vocoder features, which are used to re-synthesise the singing track and are used as a reference during evaluation. To the best of our knowledge, there is no overlap amongst the singers in the training set and the singers present in MedleyDB. Therefore the use of this dataset for evaluation makes sense as we are using both songs and singers not seen by the model during training.

The audio files were all downsampled to a sampling rate of 32 kHz. The short time fourier transform (STFT) was calculated with a Hanning window of size 1024. Both the STFT and the vocoder features were calculated with a hopsize of 5 ms. We use dimensionality reduction for the vocoder features [24, 25, 13], leading to 64 features.

4.2. Training

As suggested in the AutoVC [14] model, we use the Adam [26] optimizer for training the networks. A batch size of 30 is used with an input of length 640 ms, randomly sampled from the

tracks. We augment the input data by applying variable gains to the vocal and backing track spectrograms while creating the mixture during training.

4.3. Evaluation Methodology

There are three aspects of the signal predicted by our model to be evaluated; the intelligibility of the produced signal, isolation from the backing track and the quality of the signal. The prediction of our model is different from that made by most state-of-the-art source separation models in that our model predicts the raw unprocessed vocals whereas other source separation models predict the processed *stem*. Thus, a direct comparison with other models via objective measures like those presented in *bss_eval* [16] is not feasible. Instead we use a subjective listening test for evaluation [27]. We evaluate both the singer dependent network, SDN and the singer independent network, SIN, against a baseline [7], referred to as UNET and against our previously proposed model [13], SS. We adjusted the U-Net architecture to account for the change in the sampling frequency and the window size used in our model. All four models were trained and evaluated on the same datasets as described above. For subjective evaluation, we used an online AB listening test, wherein participants were asked to pick between two examples on the criterion of intelligibility, source isolation and audio quality. Each criteria had 15 questions with 5 for each pair. We also use Mel Cepstral Distortion (MCD) compared to the raw vocal signal, as shown in table 1 as an objective measure for overall audio quality of the vocal synthesis.

Model	MCD MEAN (dB)	MCD STD (dB)
SS	7.39	1.25
SDN	7.55	0.63
SIN	6.45	0.75
UNET	6.58	1.88

Table 1. The Mel Cepstral Distortion (MCD) metric in dB. It can be observed that the singer independent network, SIN, is comparable in quality to UNET.

4.4. Subjective Evaluation

There were 25 participants in our listening test, from various countries. All participants claimed proficiency in the English language, that was used in our survey and 18 had previous musical training. The results of the listening test are shown in Figure 2. Audio quality and intelligibility is observed to be lower for SDN than SIN, this is expected, since the SDN model produces an output given only the content embedding and the singer identity whereas the decoder of the SIN model has access to more information via the magnitude spectrogram of the input.

We can see that both the SDN and SIN are ranked higher on intelligibility than SS, thus showing that the content en-

coder is able to effectively learn the underlying linguistic features from the given input spectrogram in a singer independent manner, even for singers not seen during training and for a mixed and processed vocal track. We also observe that all three vocoder based models outperform the mask based model, UNET on isolating the vocal signal from the mixture. This follows from our previous observation [13] and is due to the approach that we follow of re-synthesising the vocal signal from the mixture rather than using a mask over the input signal. The models still lag behind on audio quality. This can partly be attributed to the degradation by the use of vocoder features and to the errors in estimation of the fundamental frequency, which is done via an external model. However, it can be seen that the proposed SIN model outperforms the baseline of the SS model, which can be explained to the difference in architecture and the additional information from the content embedding that the network receives.

Fig. 2. Results of the listening test. The vocoder based models outperform UNET [7] in terms of source isolation, but lag behind in terms of audio quality. The proposed models, SDN and SIN, are ranked higher than our previous model, SS [13] in terms of intelligibility, showing an improvement over the baseline.

5. CONCLUSIONS

We present a methodology to extract the raw unprocessed vocal signal from a musical mixture based on a linguistic feature extractor. We show that using knowledge distillation from a teacher network, we are able to train an encoder to extract linguistic features from a mixture spectrogram and use these features to synthesise the underlying vocal signal. The model is able to synthesise intelligible vocals, even for an unseen singer, without interference from the backing track. Using a decoder that takes the mixture spectrogram as an input instead of the singer identity leads to a significant improvement in audio quality and intelligibility over the baseline vocoder based model that we presented earlier. An additional application of our system is for isolating voices in unison singing, a task that has not been tackled so far by traditional source separation methodologies.

6. REFERENCES

- [1] JL Durrieu, A Ozerov, and C Févotte, “Main instrument separation from stereophonic audio signals using a source/filter model,” *17th European Signal Processing Conference*, 2009.
- [2] Po-Sen Huang, Scott Deeann Chen, Paris Smaragdis, and Mark Hasegawa-Johnson, “Singing-voice separation from monaural recordings using robust principal component analysis,” in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2012, pp. 57–60.
- [3] Pierre Comon, “Independent component analysis, a new concept?,” *Signal processing*, vol. 36, no. 3, pp. 287–314, 1994.
- [4] Po-Sen Huang, Minje Kim, Mark Hasegawa-Johnson, and Paris Smaragdis, “Singing-voice separation from monaural recordings using deep recurrent neural networks,” in *ISMIR*, 2014, pp. 477–482.
- [5] Pritish Chandna, Marius Miron, Jordi Janer, and Emilia Gómez, “Monaural audio source separation using deep convolutional neural networks,” in *International conference on latent variable analysis and signal separation*. Springer, 2017, pp. 258–266.
- [6] Aditya Arie Nugraha, Antoine Liutkus, and Emmanuel Vincent, “Multichannel music separation with deep neural networks,” in *2016 24th European Signal Processing Conference (EUSIPCO)*. IEEE, 2016, pp. 1748–1752.
- [7] Andreas Jansson, Eric Humphrey, Nicola Montecchio, Rachel Bittner, Aparna Kumar, and Tillman Weyde, “Singing voice separation with deep u-net convolutional networks,” 2017.
- [8] Shrikant Venkataramani, Efthymios Tzinis, and Paris Smaragdis, “A style transfer approach to source separation,” 2019.
- [9] Abhimanyu Sahai, Romann Weber, and Brian McWilliams, “Spectrogram feature losses for music source separation,” 2019.
- [10] Yi Luo and Nima Mesgarani, “Tasnet: time-domain audio separation network for real-time, single-channel speech separation,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 696–700.
- [11] Daniel Stoller, Sebastian Ewert, and Simon Dixon, “Wave-u-net: A multi-scale neural network for end-to-end audio source separation,” 2018.
- [12] Rachel M Bittner, Justin Salamon, Mike Tierney, Matthias Mauch, Chris Cannam, and Juan Pablo Bello, “Medleydb: A multitrack dataset for annotation-intensive mir research,” in *ISMIR*, 2014, vol. 14, pp. 155–160.
- [13] Pritish Chandna, Merlijn Blaauw, Jordi Bonada, and Emilia Gomez, “A vocoder based method for singing voice extraction,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 990–994.
- [14] Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark Hasegawa-Johnson, “Autovc: Zero-shot voice style transfer with only autoencoder loss,” in *International Conference on Machine Learning*, 2019, pp. 5210–5219.
- [15] 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.
- [16] Cédric Févotte, Rémi Gribonval, and Emmanuel Vincent, “Bss_eval toolbox user guide–revision 2.0,” 2005.
- [17] Fadi Biadsy, Ron J Weiss, Pedro J Moreno, Dimitri Kanvesky, and Ye Jia, “Parrotron: An end-to-end speech-to-speech conversion model and its applications to hearing-impaired speech and speech separation,” 2019.
- [18] Ju-chieh Chou, Cheng-chieh Yeh, Hung-yi Lee, and Lin-shan Lee, “Multi-target voice conversion without parallel data by adversarially learning disentangled audio representations,” 2018.
- [19] Ehsan Hosseini-Asl, Yingbo Zhou, Caiming Xiong, and Richard Socher, “A multi-discriminator cyclegan for unsupervised non-parallel speech domain adaptation,” 2018.
- [20] Chin-Cheng Hsu, Hsin-Te Hwang, Yi-Chiao Wu, Yu Tsao, and Hsin-Min Wang, “Voice conversion from non-parallel corpora using variational auto-encoder,” in *2016 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)*. IEEE, 2016, pp. 1–6.
- [21] Aaron van den Oord, Oriol Vinyals, et al., “Neural discrete representation learning,” in *Advances in Neural Information Processing Systems*, 2017, pp. 6306–6315.
- [22] Masanori Morise, Fumiya Yokomori, and Kenji Ozawa, “World: a vocoder-based high-quality speech synthesis system for real-time applications,” *IEICE TRANSACTIONS on Information and Systems*, vol. 99, no. 7, pp. 1877–1884, 2016.
- [23] Andreas Jansson, Rachel M Bittner, Sebastian Ewert, and Tillman Weyde, “Joint singing voice separation and f0 estimation with deep u-net architectures,” .
- [24] Merlijn Blaauw and Jordi Bonada, “A neural parametric singing synthesizer modeling timbre and expression from natural songs,” *Applied Sciences*, vol. 7, no. 12, pp. 1313, 2017.
- [25] Keiichi Tokuda, Takao Kobayashi, Takashi Masuko, and Satoshi Imai, “Mel-generalized cepstral analysis-a unified approach to speech spectral estimation,” in *Third International Conference on Spoken Language Processing*, 1994.
- [26] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” 2014.
- [27] Sebastian Kraft and Udo Zölzer, “Beaglejs: Html5 and javascript based framework for the subjective evaluation of audio quality,” in *Linux Audio Conference, Karlsruhe, DE*, 2014.