

A RECURRENT ENCODER-DECODER APPROACH WITH SKIP-FILTERING CONNECTIONS FOR MONAURAL SINGING VOICE SEPARATION

*Stylianos Ioannis Mimilakis**, *Konstantinos Drossos[†]*, *Tuomas Virtanen[†]*, *Gerald Schuller*[‡]*

*Fraunhofer IDMT, Ilmenau, Germany
 {mis, shl}@idmt.fraunhofer.de

[†]Tampere University of Technology, Tampere, Finland
 firstname.lastname@tut.fi

[‡]Technical University of Ilmenau, Ilmenau, Germany

ABSTRACT

The objective of deep learning methods based on encoder-decoder architectures for music source separation is to approximate either ideal time-frequency masks or spectral representations of the target music source(s). The spectral representations are then used to derive time-frequency masks. In this work we introduce a method to directly learn time-frequency masks from an observed mixture magnitude spectrum. We employ recurrent neural networks and train them using prior knowledge only for the magnitude spectrum of the target source. To assess the performance of the proposed method, we focus on the task of singing voice separation. The results from an objective evaluation show that our proposed method provides comparable results to deep learning based methods which operate over complicated signal representations. Compared to previous methods that approximate time-frequency masks, our method has increased performance of signal to distortion ratio by an average of 3.8 dB.

Index Terms— Music source separation, deep learning, denoising autoencoders

1. INTRODUCTION

The under-determined separation of audio and music signals from mixtures is an active research area in the field of audio signal processing. The main objective is to estimate individual sources contained in an observed single-channel (i.e. monaural) mixture. An important task which has attracted a lot of attention is the estimation of singing voice and background music [1]. The most widely-used strategy to achieve the estimation of individual sources employs time-varying filters, usually in the short-time Fourier transform (STFT) domain. These filters, henceforth denoted as time-frequency masks, are derived from rational models which incorporate prior information about the spectral representation of each source in the mixture [2, 3].

When sources are already known (i.e. informed source separation) [4], the process of source separation is almost trivial by employing an optimal time-frequency masking strategy [2]. On the other hand, when the sources are not known, a prior estimation of each individual source takes place. To that aim, numerous approaches have been proposed. More specifically, in [5] it is proposed to exploit phase information for separating sources that have harmonic and impulsive characteristics, while in [6] the repetitive structure of music sources is studied for separating singing voice and background music. Alternative methods approximate the mixture magnitude via non-negative low rank modeling [7, 8], and/or matrix completion via robust principal component analysis (RPCA) [9].

Studies in music source separation have shown that supervised methods based on deep learning can yield state-of-the-art results in singing voice separation [1]. Such methods can be roughly discriminated into two main categories. In the first category the goal is to train deep neural networks, to approximate an ideal time-frequency mask from observed magnitude spectral representations [10, 11, 12]. The second category exploits denoising auto-encoders (i.e. neural networks trained to map from a noisy observation to a clean one), with the main goal of estimating the magnitude spectral representation of individual music sources from input mixtures [13, 14, 15]. Then, these estimates are combined to derive [13, 15, 16] or optimize [14, 17] a time-frequency mask that filters the input mixture. However, the quality of the time-frequency mask is heavily depended on the computation of the sources [14, 16] and the time-frequency mask approximation is not part of the optimization process (only the estimation of the sources is).

An exception to the latter is the work in [14]. In that work, in order to predict the time-frequency mask the sources are first predicted and then combined in a generalized Wiener filtering process. Although, and according to [14], the approximation of the time-frequency mask is subject to optimization, this optimization is based on the ability of the previous neural network layers to estimate the sources.

In this work we propose a method for predicting a time-frequency mask from the observed mixture magnitude spectrogram, and optimizing it according to its performance on estimating the target source. During training, the only prior knowledge is the mixture and the target source magnitude spectrograms. After training, only the mixture magnitude spectrogram is required. This approach differs from the existing ones because: a) we do not base the prediction of the time-frequency mask on the prior estimation of the source(s), and b) we do not require a prior knowledge of the ideal time-frequency mask for training. We let a recurrent neural network (RNN) to predict the time-frequency mask. Then, based on that time-frequency mask, we estimate the magnitude spectrogram of the target source using skip-filtering connections, a highway network, and a generalized Wiener filtering process. The RNN and the highway network are jointly trained.

The rest of the paper is organized as follows: Section 2 presents the proposed method, followed by Section 3 which provides information about the training and evaluation of the proposed method. Section 4 presents the obtained results from the experimental procedure, followed by discussion. Conclusions are in Section 5.

2. PROPOSED METHOD

Our method accepts as an input the time-domain samples of a monaural mixture vector \mathbf{x} and produces the time-domain samples of the j -th target source vector $\hat{\mathbf{y}}^j$, as illustrated in Fig. 1. We calculate the matrix \mathbf{Y} , which contains the complex-valued time-frequency representation of \mathbf{x} . In order to use short sequences of frames with context information (i.e. previous and next frames), we create the tensor \mathbf{Y}_{in} consisting of overlapping segments of the magnitude of \mathbf{Y} . We use each matrix in \mathbf{Y}_{in} as an input to a single-layered, bi-directional gated recurrent unit (BiGRU), the encoder, in order to let our method learn temporal inter-dependencies of the input data. The output of the encoder is used as an input to a single layered GRU, the decoder, in order to estimate a time-frequency mask. We combine the estimated time-frequency mask and the input to the encoder using skip-filtering connections, in order to estimate the magnitude spectrogram of the target source, $\mathbf{Y}_{\text{fit}}^j$. The encoder and the decoder are optimized by minimizing the generalized Kullback-Leibler divergence, \mathcal{L}_{KL} , between the true and the estimated magnitude spectrogram of the target source.

The $\mathbf{Y}_{\text{fit}}^j$ is used as an input to a highway network in order to reduce the interferences from the rest sources. The encoder, the decoder, and the highway network are optimized using \mathcal{L}_{KL} between the true magnitude spectrogram of the target source and the output of the highway network, plus an L_2 regularization term. The output of the highway network with the complex-valued mixture representation, \mathbf{Y} , are given as an input to a generalized Wiener filtering process. The latter

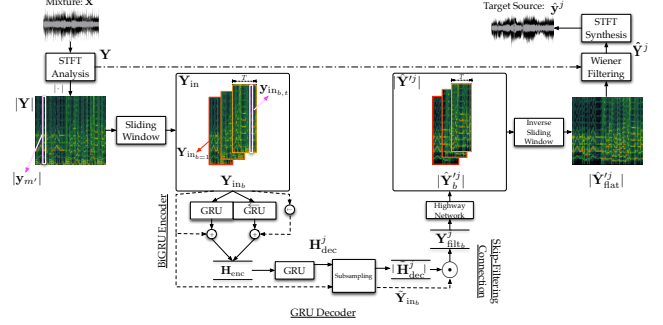


Fig. 1. Illustration of the proposed encoder-decoder approach with skip-filter connections.

produces the complex-valued time-frequency representation of the target source and further reduces the interferences from the rest sources. The final output, $\hat{\mathbf{y}}^j$, is calculated through an overlap and add synthesis procedure [18].

2.1. Input processing

Vector \mathbf{x} is sliced in M overlapping time frames of length N , using a shift/hop-size of H samples. The overlapping time frames are element-wise multiplied by the Hamming windowing function. This yields a matrix of overlapping segments of the original mixture, from which the complex-valued time-frequency representation $\mathbf{Y} \in \mathbb{C}^{M \times N}$ is acquired by applying the short-time Fourier (STFT) analysis for all M time segments. From the STFT representation of the mixture \mathbf{Y} , we compute the magnitude spectrum $|\mathbf{Y}|$, with $|\cdot|$ denoting the element-wise magnitude of the matrix. Using a sliding window over the time-frames of $|\mathbf{Y}|$ we form the tensor $\mathbf{Y}_{\text{in}} \in \mathbb{R}_{\geq 0}^{B \times T \times N}$, where $B = \lceil M/T \rceil$ and T is an integer indicating the amount of STFT time-frames used for encoding. \mathbf{Y}_{in} consists of a B number of matrices that encompass $T \times N$ segments of the mixture magnitude spectrogram $|\mathbf{Y}|$. Each matrix $\mathbf{Y}_{\text{in}_b} \in \mathbb{R}_{\geq 0}^{T \times N}$ is a sequence of T overlapping time-frames, and is acquired from $|\mathbf{Y}|$ as

$$\mathbf{Y}_{\text{in}_b} = [|\mathbf{y}_{m'+1}|, \dots, |\mathbf{y}_{m'+T}|], \text{ where} \quad (1)$$

$$m' = (b-1) \times T - (b-1) \times 2 \times L, \quad (2)$$

and L and m' are integers indicating the amount of time-frames that will be used as context information, and the time-frame location in $|\mathbf{Y}|$ ($|\mathbf{y}_{m'}|$), respectively. In the case that $m' > M$, $|\mathbf{y}_{m'}|$ is a vector of zeros and of length N . The usage of overlapping time-frames is interpreted as having $2 \times L$ context frames (L before and L after) for $T - (2 \times L)$ length sequence of time-frames and is necessary for the sub-sampling operation, imposed by the decoding part of our proposed method. The term $2 \times L$ is derived experimentally.

2.2. Mask and source estimation

Each \mathbf{Y}_{in_b} is used as an input to the BiGRU encoder, which processes it according to the equations describing a GRU and using the proposed non-linearities [19]. The BiGRU encoder consists of two GRUs, the forward GRU which accepts \mathbf{Y}_{in_b} as an input, and the backward one which accepts $\overleftarrow{\mathbf{Y}}_{in_b} = [\mathbf{y}_{in_b,T}, \dots, \mathbf{y}_{in_b,t}, \dots, \mathbf{y}_{in_b,1}]$, where $\mathbf{y}_{in_b,t} \in \mathbb{R}_{\geq 0}^N$ is a vector in \mathbf{Y}_{in_b} at time-step t , and \leftarrow denotes the direction of recursion over the time sequences. The hidden outputs from the bi-directional encoder at time-step t , \mathbf{h}_t and $\overleftarrow{\mathbf{h}}_t$, are concatenated and updated using residual connections [20], as

$$\mathbf{h}_{enc_t} = [(\mathbf{h}_t + \mathbf{y}_{in_b,t})^T, (\overleftarrow{\mathbf{h}}_t + \overleftarrow{\mathbf{y}}_{in_b,t})^T]^T \quad (3)$$

and the output of the encoder is

$$\mathbf{H}_{enc} = [\mathbf{h}_{enc_1}, \dots, \mathbf{h}_{enc_t}, \dots, \mathbf{h}_{enc_T}]. \quad (4)$$

\mathbf{H}_{enc} is used as an input to the decoder, producing an output denoted as \mathbf{H}_{dec}^j . The subscripts “enc” and “dec” stand for encoder and decoder, respectively. The superscript j signifies that the values of the matrix are dependent on the j -th source. The concatenation imposed by Eq. (3) increases the dimensionality of the encoded information by a factor of two (\mathbf{H}_{enc} is of shape $T \times (N \times 2)$). For this reason the decoder utilizes dimensionality reduction, such the \mathbf{H}_{dec}^j matches the dimensionality of \mathbf{Y}_{in_b} . Furthermore, a sub-sampling operation over the time-steps of \mathbf{H}_{dec}^j is applied, defined as

$$\tilde{\mathbf{H}}_{dec}^j = [\mathbf{h}_{dec_{1+L}}^j, \mathbf{h}_{dec_{2+L}}^j, \dots, \mathbf{h}_{dec_{T-L}}^j], \quad (5)$$

where L in Eq. (5) indicates the amount of time steps to be discarded from the start and end of \mathbf{H}_{dec}^j . The intuition behind Eq. (5) is to allow the flow of information from previous and preceding context frames L , without dramatically increasing the learning complexity of longer time interdependencies [21]. After the sub-sampling operation, we apply the skip-filtering connections between $\tilde{\mathbf{H}}_{dec}^j$ and a sub-sampled version of the input data $\tilde{\mathbf{Y}}_{in_b}$ as

$$\mathbf{Y}_{filt_b}^j = \tilde{\mathbf{Y}}_{in_b} \odot |\tilde{\mathbf{H}}_{dec}^j|, \text{ where} \quad (6)$$

$$\tilde{\mathbf{Y}}_{in_b} = [\mathbf{y}_{in_b,L}, \dots, \mathbf{y}_{in_b,T-L}], \quad (7)$$

\odot is the Hadamard product, $\tilde{\mathbf{H}}_{dec}^j$ is interpreted as a time-frequency mask specific to the source j , and $\tilde{\mathbf{Y}}_{in_b}$ is the sub-sampled version of \mathbf{Y}_{in_b} . According to Eq. (6), $\mathbf{Y}_{filt_b}^j$ can be considered as a filtered version of the input sequence $\tilde{\mathbf{Y}}_{in_b}$ approximating the magnitude spectrogram of the j -th target source, and thus the subscript “filt”. The skip-filtering connections directly filter the input mixture magnitude sequence by the output of the decoder. The motivation behind this is to enforce $\tilde{\mathbf{H}}_{dec}^j$ to be a time-frequency mask derived through optimization, based on the prior knowledge of the magnitude spectrogram of the target source and not on the prior knowledge of an ideal time-frequency mask. The skip-filtering

connections are inspired by the denoising source separation framework presented in [22]. We minimize the \mathcal{L}_{KL} between the true and estimated magnitude spectrogram of the j -th target source to train the encoder and the decoder.

2.3. Enhancement of estimated magnitude spectrogram

In order to enhance the outcome of the filtering process of Eq. (6) [15], $\mathbf{Y}_{filt_b}^j$ is used as an input to a single layer of highway neural networks [23], with shared weights in time, as

$$|\hat{\mathbf{Y}}_b^{tj}| = \sigma(\mathbf{Y}_{filt_b}^j \mathbf{W}^h + \mathbf{b}^h) \odot g(\mathbf{Y}_{filt_b}^j \mathbf{W}^{tr} + \mathbf{b}^{tr}) + \mathbf{Y}_{filt_b}^j \odot (1 - \sigma(\mathbf{Y}_{filt_b}^j \mathbf{W}^{tr} + \mathbf{b}^{tr})), \quad (8)$$

where $|\hat{\mathbf{Y}}_b^{tj}|$ is the enhanced estimated magnitude spectrogram of the j -th target source [15], σ is the sigmoid function, g is defined as $g(x) = \max(0, x)$, and \mathbf{W}^{tr} and \mathbf{W}^h are the weight matrices of the transformation and gating operations with their corresponding bias vectors \mathbf{b}^h and \mathbf{b}^{tr} , respectively. All the layers are trained by minimizing the \mathcal{L}_{KL} between the true and the enhanced estimated magnitude spectrogram of the j -th target source, plus an additional L_2 penalty, for regularizing $|\hat{\mathbf{Y}}_b^{tj}|$, scaled by $\lambda = 1e^{-4}$. Given that L_2 penalty is higher for non pseudo-periodic information [9], we use it in order to enforce the highway network to enhance the estimated magnitude spectrogram of the singing voice.

2.4. Final output and implementation details

We iterate through all the above presented equations for all b in \mathbf{Y}_{in_b} and the result is the tensor $|\hat{\mathbf{Y}}^{tj}| \in \mathbb{R}_{\geq 0}^{B \times T' \times N}$, where $T' = T - 2 \times L$. Then, $|\hat{\mathbf{Y}}^{tj}|$ is reshaped to $|\hat{\mathbf{Y}}_{flat}^{tj}| \in \mathbb{R}_{\geq 0}^{M \times N}$, by reversing the process in Eq. (1) and (2). The subscript “flat” is used to refer to the aforementioned reshaping procedure. The final output of the proposed method is the complex-valued spectral representation of the j -th source $\hat{\mathbf{Y}}^j \in \mathbb{C}^{M \times N}$ computed as

$$\hat{\mathbf{Y}}^j = \mathbf{M}^j \odot \mathbf{Y}, \text{ where} \quad (9)$$

$$\mathbf{M}^j = \frac{|\hat{\mathbf{Y}}_{flat}^{tj}|^\alpha}{|\mathbf{Y}|^\alpha}, \quad (10)$$

$\alpha \in (0, 2]$ is an exponent applied element-wise, signifying the statistical assumption(s) about the sources and their additive property in the magnitude time-frequency domain [2]. Furthermore, the division is also performed element-wise. The reasoning behind Eqs. (9)–(10) and the α factor, is to infer the information that the mixture encapsulates about the phase and the magnitude of the target source, practically improving the interference reduction from other concurrently active sources. The reconstruction of the time-domain function of the target source $\hat{\mathbf{y}}^j$ is achieved by the DFT synthesis operation, followed by the overlap and add method presented in [18].

The highway network and the GRUs in the encoder consist of 1025 neurons each and the GRU in the decoder has 2050 neurons. This results in approximately 24 million parameters in total for our method. All weight matrices are randomly initialized using the method presented in [24] and jointly trained using the adam algorithm [25], with a learning rate of $1e^{-3}$, over b batches of 16, and an L_2 based gradient norm clipping equal to 0.35. The training is terminated if after two consecutive iterations over all the available training batches, no minimization took place. The implementation of the proposed method is based on the keras [26] and Theano [27] frameworks. The above parameters are chosen experimentally by informal listening tests and subjectively evaluating the obtained quality of separation, with data drawn from the development subset of DSD100 (presented in the next section).

3. EVALUATION

3.1. Dataset and preprocessing

In order to assess the performance of the proposed method we focus on the task of singing voice and background music separation. The Demixing Secret Dataset (DSD100)¹ is used for training and evaluating our approach. DSD100 consists of 100 professionally produced multi-tracks of various music genres, sampled at 44.1kHz. The dataset is by default divided evenly into development and evaluation subsets, each consisting of 50 multi-tracks, and each multi-track contains four stereo target sources forming the produced mixture.

For each multi-track contained in the development subset, a monaural version of each of the four sources is generated by averaging the two available channels. Afterwards, two signals are generated. One containing the corresponding monaural mixture of all the monaural sources (singing voice, bass, drums, etc) and a second containing the mixture of all the monaural sources but the singing voice. The mixing gain values for each source are not modified (no data augmentation is applied). An analysis operation using the STFT is applied to the two preceding mixture signals and the target singing voice, using $H = 256$ and $N' = 2048$. Since we are concerned with real-valued signals, their time-frequency representation using the DFT is Hermitian and thus the redundant information is discarded, resulting into a dimensionality of $N = 1025$. The length of the sequences is set to $T = 18$, modeling approximately 120 ms, and $L = 3$. To avoid inconsistencies due to the sub-sampling operation of Eq. (5), the considered shifts for acquiring \mathbf{Y}_{in} are overlapping by six time-frames ($2 \times L$). During training, the true source $|\mathbf{Y}^j|$ is the outcome of a generalized Wiener filtering using the a priori knowledge from the dataset for each source and a value of $\alpha = 1$.

3.2. Metrics and evaluation procedure

To evaluate our method we use the standard metrics employed in the music source separation evaluation campaign [1]. These are the signal to distortion ratio (SDR) and the signal to interference ratio (SIR). These metrics are computed by using the output of the proposed method on the evaluation sub-set and the true sources acquired from the multi-tracks of the same sub-set. Results from three different strategies using the proposed method are reported. For the first one, we train our method using the magnitude spectrogram of the singing voice $|\mathbf{Y}^{j=1}|$. After the training, we obtain an estimation of the singing voice using the input mixture spectrogram and an α value equal to 1.7, as according to Eq. (9) and (10). For the rest two strategies, we use two deep neural networks trained separately. Both of them are based on our method. One to predict the singing voice $|\mathbf{Y}^{j=1}|$ and another to predict the background music without the singing voice $|\mathbf{Y}^{j=2}|$. The predicted magnitude spectral estimates (i.e. $|\hat{\mathbf{Y}}^{j=1}|$ and $|\hat{\mathbf{Y}}^{j=2}|$) are used in the Eq. (10) and the Eq. (9) to compute the final estimation of the $|\mathbf{Y}^{j=1}|$. In Eq. (10), the numerator is defined as $|\hat{\mathbf{Y}}^{j=1}|^\alpha$ and we replace the $|\mathbf{Y}|^\alpha$ of the denominator with $|\hat{\mathbf{Y}}|^\alpha = |\hat{\mathbf{Y}}^{j=2}|^\alpha + |\hat{\mathbf{Y}}^{j=1}|^\alpha$. For the second strategy we use $\alpha = 1.7$ and for the third $\alpha = 2$. We will refer to these three strategies as GRU-S (the first), GRU-D (the second), and GRU-DWF (the third). The values for α are chosen according to the generalized Wiener filtering (i.e. $\alpha = 2$) and its extension to heavy-tailed distributions [2].

We compare with state-of-the-art methods dealing with monaural singing voice separation. More specifically, we compare the results from our proposed method against: i) two approaches based on deep feed-forward neural networks of approximately 17 million parameters each, trained using data augmentation to predict both binary and soft time-frequency masks derived from [3], denoted as GRA2 and GRA3 [10] ii) a convolutional based encoder-decoder approach with approximately 27 million parameters, trained using data augmentation to predict a soft time-frequency mask, denoted as CHA [12], and iii) two additional deep denoising auto-encoders operating on the common fate signal representation [28] denoted as STO1, STO2. As for the oracle estimation, the results from ideal binary masking (IBM) are also presented. We report results for the overall performance on the employed dataset and for three clusters of music genres, namely Jazz/Pop/Rock, Electronic/Rap, and Heavy Metal.

4. RESULTS & DISCUSSION

Results from the SDR and SIR metrics for the above methods and the oracle estimation are illustrated in Figures 2 and 3. Examining the differences in SDR and SIR metrics between the methods are trained to approximate the ideal time-frequency masks (GRA3, CHA) and GRU-D, it can be observed that the median SDR has been improved by approximately 5.4 dB when compared to GRA3 and by 2.1 dB when

¹<http://www.sisec17.audiolabs-erlangen.de>

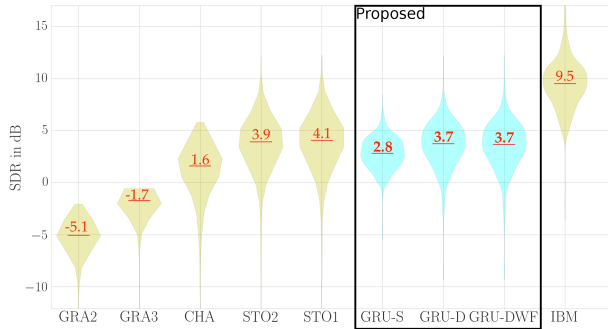


Fig. 2. Analysis of variance of SDR for previous approaches and the proposed ones. Red lines denote the median values.

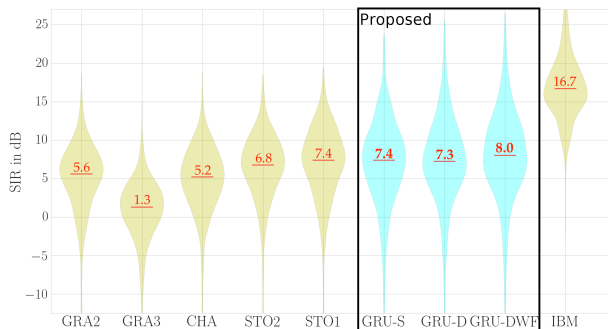


Fig. 3. Analysis of variance of SIR for previous approaches and the proposed ones. Red lines denote the median values.

compared to CHA. An average improvement of 3.5 dB for the SIR can be observed regarding the methods that predict of the time-frequency mask (GRA2, GRA3, CHA). On the other hand, when it comes to comparison with common denoising auto-encoders (STO1, STO2) and our method, a marginal average loss of 0.3 dB in the median SDR and a gain of 0.25 dB and 0.9 dB in the median SIR can be observed according to the results over the GRU-D and GRU-DWF cases, respectively. We believe (but we don't have any evidence) that these marginal differences can be attributed to the more sophisticated signal representations, such as the common fate model [28], used in STO1, STO2. In contrast, our method operates on top of a magnitude spectral representation computed from a STFT. By inspecting the differences between the proposed strategies GRU-S and GRU-D, it can be seen that SDR can be increased by approximately 0.9 dB by incorporating separate deep neural networks for approximating additional sources contained in mixtures. On the other hand, dramatic differences in SIR between the three strategies were not observed, unless the value for α was increased, like in the case of GRU-DWF. This shows, that combining multiple deep neural networks in generalized Wiener filtering leads to improved SDR. Additionally, the additivity property, acknowledged in generalized Wiener filtering [2], of the estimates of deep neural networks might not hold true without an explicit cost objective. Nonetheless, it can be used for improv-

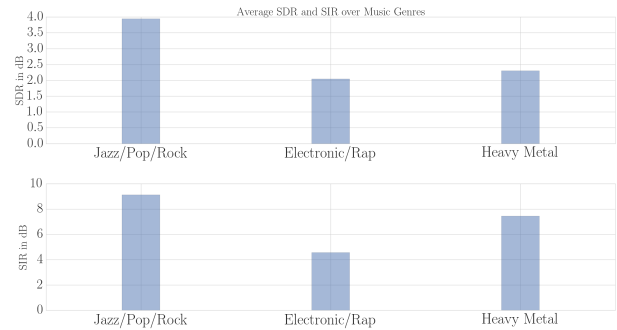


Fig. 4. Average objective performance of the three presented strategies over three clusters of music genres.

ing the interference reduction. In Figure 4 are the average SDR and SIR measures of the three employed strategies over the three employed music genres. As it can be seen, the performance over the cluster Jazz/Pop/Rock is higher than the performance of for the other clusters. An explanation to this is that the DSD100 dataset contains training multi-tracks mainly from pop and rock music genres. This means that the poorer performance of our method can be attributed to the fact that recurrent models might need additional data for modeling more complicated structures of singing voice.

5. CONCLUSIONS & FUTURE WORK

In this work we presented a deep learning method for music source separation. The method used an encoder-decoder configuration based on GRUs and skip-filtering connections between input spectral representations and their hidden latent variables, forcing the GRUs to approximate a time-frequency masking operation. Its application to monaural singing voice separation was studied and assessed objectively. The obtained results signify that the skip-filtering connections can be used for approximating time-frequency masks, providing comparable results to state of the art deep learning approaches. Future work will focus on psycho-acoustically motivated loss minimization and exploring sparsity priors for improving the approximated time-frequency mask. Subjective assessment of the plausible extensions of our methodology are also emerging [29]. Source code and listening examples can be found under: https://github.com/Js-Mim/mlsp2017_svsep_skipfilt.

6. ACKNOWLEDGEMENTS

We would like to thank the authors of [1] for making the results of the evaluation available. The major part of the research leading to these results has received funding from the European Union's H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement no 642685 MacSeNet. Konstantinos Drossos was partially funded from the European Union's H2020 Framework Programme through ERC Grant Agreement 637422 EVERYSOUND. Minor part of the computations leading to these results were performed on a TITAN-X GPU donated by NVIDIA.

7. REFERENCES

- [1] A. Liutkus, F.-R. Stöter, Z. Rafii, D. Kitamura, B. Rivet, N. Ito, N. Ono, and J. Fontecave, “The 2016 signal separation evaluation campaign,” in *Latent Variable Analysis and Signal Separation: 13th International Conference, LVA/ICA 2017*, 2017, pp. 323–332.
- [2] A. Liutkus and R. Badeau, “Generalized wiener filtering with fractional power spectrograms,” in *40th International Conference on Acoustics, Speech and Signal Processing (ICASSP 2015)*, April 2015, pp. 266–270.
- [3] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, “Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks,” in *40th International Conference on Acoustics, Speech and Signal Processing (ICASSP 2015)*, April 2015, pp. 708–712.
- [4] G. Puy, A. Ozerov, N. Q. K. Duong, and P. Perez, “Informed source separation via compressive graph sampling,” in *42nd International Conference on Acoustics, Speech and Signal Processing (ICASSP 2017)*, Mar 2017, pp. 1–5.
- [5] E. Cano, M. Plumbley, and C. Dittmar, “Phase-based harmonic percussive separation,” in *Proceedings of the Annual Conference of the International Speech Communication Association (Interspeech)*, Sept. 2014, pp. 1628–1632.
- [6] Z. Rafii and B. Pardo, “Repeating pattern extraction technique (repet): A simple method for music/voice separation,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 1, pp. 73–84, Jan 2013.
- [7] A. Liutkus, D. Fitzgerald, and R. Badeau, “Cauchy nonnegative matrix factorization,” in *Applications of Signal Processing to Audio and Acoustics (WASPAA 2015)*, Oct 2015, pp. 1–5.
- [8] D. Fitzgerald, A. Liutkus, and R. Badeau, “PROJET - Spatial Audio Separation Using Projections,” in *41st International Conference on Acoustics, Speech and Signal Processing (ICASSP 2016)*, 2016, pp. 36–40.
- [9] I.-Y. Jeong and K. Lee, “Singing voice separation using RPCA with weighted l_1 -norm,” in *Latent Variable Analysis and Signal Separation: 13th International Conference, LVA/ICA 2017*, 2017, pp. 553–562.
- [10] E.-M. Grais, G. Roma, A.J.R. Simpson, and M.-D. Plumbley, “Single-channel audio source separation using deep neural network ensembles,” in *Audio Engineering Society Convention 140*, May 2016.
- [11] E.-M. Grais, G. Roma, A.J.R. Simpson, and M.-D. Plumbley, “Combining mask estimates for single channel audio source separation using deep neural networks,” in *Proceedings of the 17th Annual Conference of the International Speech Communication Association (Interspeech)*, Sept. 8-12 2016, pp. 3339–3343.
- [12] P. Chandna, M. Miron, J. Janer, and E. Gómez, “Monoaural audio source separation using deep convolutional neural networks,” in *Latent Variable Analysis and Signal Separation: 13th International Conference, LVA/ICA 2017*, 2017, pp. 258–266.
- [13] S. Uhlich, F. Giron, and Y. Mitsufuji, “Deep neural network based instrument extraction from music,” in *40th International Conference on Acoustics, Speech and Signal Processing (ICASSP 2015)*, 2015, pp. 2135–2139.
- [14] P.-S. Huang, M. Kim, M. Hasegawa-Johnson, and P. Smaragdis, “Joint optimization of masks and deep recurrent neural networks for monaural source separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 12, pp. 2136–2147, Dec 2015.
- [15] S.-I. Mimilakis, E. Cano, J. Abeßer, and G. Schuller, “New sonorities for jazz recordings: Separation and mixing using deep neural networks,” in *Audio Engineering Society 2nd Workshop on Intelligent Music Production*, 2016.
- [16] S. Uhlich, M. Porcu, F. Giron, M. Enekl, T. Kemp, N. Takahashi, and Y. Mitsufuji, “Improving music source separation based on deep neural networks through data augmentation and network blending,” in *42nd International Conference on Acoustics, Speech and Signal Processing (ICASSP 2017)*, 2017, pp. 261–265.
- [17] A.-A. Nugraha, A. Liutkus, and E. Vincent, “Multichannel music separation with deep neural networks,” in *24th European Signal Processing Conference (EUSIPCO)*, Aug 2016, pp. 1748–1752.
- [18] D. Griffin and J. Lim, “Signal estimation from modified short-time Fourier transform,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 2, pp. 236–243, Apr 1984.
- [19] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [20] Y. Wu et al, “Google’s neural machine translation system: Bridging the gap between human and machine translation,” *CoRR*, vol. abs/1609.08144, 2016.
- [21] A. Graves, *Hierarchical Subsampling Networks*, pp. 109–131, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- [22] J. Särelä and H. Valpola, “Denoising source separation,” *J. Mach. Learn. Res.*, vol. 6, pp. 233–272, Dec. 2005.
- [23] R.-K. Srivastava, K. Greff, and J. Schmidhuber, “Highway networks,” *CoRR*, vol. abs/1505.00387, 2015.
- [24] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS’10)*, 2010, pp. 249–256.
- [25] D.-P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014.
- [26] F. Chollet et al., “Keras (version 1.2.2),” <https://github.com/fchollet/keras>, 2015.
- [27] Theano Development Team, “Theano: A python framework for fast computation of mathematical expressions (version 0.9),” *CoRR*, vol. abs/1605.02688, 2016.
- [28] F.-R. Stöter, A. Liutkus, R. Badeau, B. Edler, and P. Magron, “Common fate model for unison source separation,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP 2016)*, 2016, pp. 126–130.
- [29] E. Cano, D. Fitzgerald, and K. Brandenburg, “Evaluation of quality of sound source separation algorithms: Human perception vs quantitative metrics,” in *24th European Signal Processing Conference (EUSIPCO)*, Aug 2016, pp. 1758–1762.