

Depthwise Separable Convolutions Versus Recurrent Neural Networks for Monaural Singing Voice Separation

Pyry Pyykkönen*, Stylianos I. Mimilakis†, Konstantinos Drossos‡, and Tuomas Virtanen‡

*3D Media Research Group, Tampere University, Tampere, Finland

Email: {firstname.lastname}@tuni.fi

†Semantic Music Technologies Group, Fraunhofer-IDMT, Ilmenau, Germany

Email: mis@idmt.fraunhofer.de

‡Audio Research Group, Tampere University, Tampere, Finland

Email: {firstname.lastname}@tuni.fi

Abstract—Recent approaches for music source separation are almost exclusively based on deep neural networks, mostly employing recurrent neural networks (RNNs). Although RNNs are in many cases superior than other types of deep neural networks for sequence processing, they are known to have specific difficulties in training and parallelization, especially for the typically long sequences encountered in music source separation. In this paper we present a use-case of replacing RNNs with depthwise separable (DWS) convolutions, which are a lightweight and faster variant of the typical convolutions. We focus on singing voice separation, employing an RNN architecture, and we replace the RNNs with DWS convolutions (DWS-CNNs). We conduct an ablation study and examine the effect of the number of channels and layers of DWS-CNNs on the source separation performance, by utilizing the standard metrics of signal-to-artifacts, signal-to-interference, and signal-to-distortion ratio. Our results show that by replacing RNNs with DWS-CNNs yields an improvement of 1.20, 0.06, 0.37 dB, respectively, while using only 20.57% of the amount of parameters of the RNN architecture.

Index Terms—Depthwise separable convolutions, recurrent neural networks, mad, madtwinnet, monaural singing voice separation

I. INTRODUCTION

The task of audio source separation is to extract the underlying audio sources from an observed audio mixture. A particular problem that has attracted great attention in audio and music source separation, is the estimation of the singing voice and accompaniment sources [1]. To address this problem, a common and successfully employed work flow, consists of computing non-negative signal representations, and employing deep neural networks (DNNs) to estimate the target sources.

Although different methods have been recently proposed for computing and learning signal adaptive/dependent representations for source separation [2]–[4], the short-time Fourier transform (STFT) remains a popular choice among state-of-the-art (SOTA) approaches in music source separation [5]–

[9]. Specifically, by using the STFT, the complex valued representation of the mixture signal is computed. Then, the corresponding magnitude information of the mixture signal is processed by an appropriated method, e.g. DNNs, yielding the magnitude information of the target source. Using the phase information of the mixture, the time-domain signals of the estimated sources are recovered by means of the inverse STFT (ISTFT).

Focusing on the DNNs that estimate the target source in the STFT domain, a certain approach that state-of-the-art methods employ is that of filtering/masking. This approach, enforces DNNs to output filters that are optimized for separating audio and music sources, and has led to good results for both separation quality [5], [7], [8] and computational costs [9]. In more details, DNNs are conditioned on the mixture signal magnitude spectrogram and are optimized, in a supervised fashion, to yield a time-varying filter, i.e., a time-frequency mask. The time-frequency mask is applied to the input mixture spectrogram, resulting into a filtered version of the input mixture. The parameters of the DNNs are optimized to minimize the difference between the filtered and the targeted source spectrograms, available in the training dataset. The main benefit of employing such approach versus other approaches is that the DNNs are more efficient in learning the spectrogram structure of the target music source [10].

Typical DNN masking-based approaches for music source separation rely on recurrent neural networks (RNNs) to encode information from the mixture magnitude spectrogram [5], [6], [8], that is then decoded to obtain the source-dependent mask. However, many previous works have highlighted that the optimization of the DNNs could be difficult, due to the involved RNNs, resulting into a very slow, or even sub-optimal learning process. A few reasons to that are improper gradient norms of the RNN parameters during training [11], and the large number of parameters RNNs require to efficiently process long sequences [12]. Although techniques, such as skip-connections [12], bi-directional sequence sampling [12], and regularization schemes [13] have been proposed to al-

K. Drossos and T. Virtanen would like to acknowledge CSC Finland for computational resources. Stylianos I. Mimilakis is supported in part by the German Research Foundation (AB 675/2-1, MU 2686/11-1).

leviate the above severe issues, CNNs have an increased popularity [7], [14]–[16]. In contrast to RNNs, CNNs have fewer parameters and can be easily parallelised, resulting into a faster learning process. Furthermore, recent works have shown that depth-wise separable CNNs can even perform better than typical CNNs in a wide range of applications spanning from image recognition [17] to sound event detection [16] and speech [18] and music source separation [15].

Because of the above, in this work we conduct an ablation study and examine the objective performance differences in singing voice separation, by replacing the RNNs with depth-wise separable CNNs. To that aim, we particularly focus on the Masker and Denoiser (MaD) architecture presented in the following works [5], [6], [19]. We do so because MaD architecture incorporates the RNN techniques that have been previously presented in [6] and in [5], serving a fair, yet competitive baseline for the scope of this work.

The rest of the paper is organized as follows. In Section II we present our proposed method, consisting of the replacement of RNNs with depth-wise separable convolutions at the MaD architecture. In Section III we presented the followed evaluation procedure, and the obtained results are presented in Section IV. Section V concludes the paper.

II. PROPOSED METHOD

Our method accepts as an input the magnitude spectrogram $\mathbf{V} \in \mathbb{R}_{>0}^{T+L \times F}$ of the musical mixture, consisting of $T+L$ time frames with F frequency bands, and outputs the magnitude spectrogram $\hat{\mathbf{V}}_j \in \mathbb{R}_{>0}^{T \times F}$ of the j -th targeted source, by applying a two-step process. First, our method filters \mathbf{V} , producing an initial estimate of the magnitude spectrogram of the j -th source, $\hat{\mathbf{V}}'_j \in \mathbb{R}_{>0}^{T \times N}$, where the extra L vectors of \mathbf{V} are used as temporal context for the initial estimate $\hat{\mathbf{V}}'_j$. Then, our method enhances $\hat{\mathbf{V}}'_j$, producing the final estimate of the magnitude spectrogram of the j -th source, $\hat{\mathbf{V}}_j$.

Our proposed method is based on the MaD system [5], [6], [19], which takes as an input \mathbf{V} and employs two denoising auto-encoders (DAEs), one for estimating $\hat{\mathbf{V}}'_j$, and one for calculating $\hat{\mathbf{V}}_j$. The first DAE in MaD is based on RNNs, which are known to be hard to use for parallelized training, and more hard to optimize than CNNs [16], [20], [21].

A. MaD system

MaD consists of two modules; the masker and the denoiser. The masker accepts as an input \mathbf{V} and outputs $\hat{\mathbf{V}}'_j$, and it consists of a trimming operation, Tr , a bi-directional RNN encoder, RNN_{enc} , a unidirectional RNN decoder, RNN_{dec} , and a feed-forward layer, FNN_m .

The trimming operation, Tr , takes as an input \mathbf{V} and reduces the amount of frequency bands from F to N , resulting in $\mathbf{V}_{\text{tr}} \in \mathbb{R}_{>0}^{T+L \times N}$. This is done in order to reduce the input dimensionality of RNN_{enc} , consequently reducing the amount of parameters of RNN_{enc} . Though, the complete \mathbf{V} will be used later on, after the RNN_{enc} . The bi-directional RNN_{enc}

consists of a forward RNN, $\overrightarrow{\text{RNN}}_{\text{enc}}$, and a backward RNN, $\overleftarrow{\text{RNN}}_{\text{enc}}$, takes as an input \mathbf{V}_{tr} and processes it according to

$$\overrightarrow{\mathbf{h}}_{\text{enc}}^{t'} = \overrightarrow{\text{RNN}}_{\text{enc}}(\mathbf{v}_{\text{tr}}^{t'}, \overrightarrow{\mathbf{h}}_{\text{enc}}^{t'-1}) \quad \text{and} \quad (1)$$

$$\overleftarrow{\mathbf{h}}_{\text{enc}}^{t'} = \overleftarrow{\text{RNN}}_{\text{enc}}(\overleftarrow{\mathbf{v}}_{\text{tr}}^{t'}, \overleftarrow{\mathbf{h}}_{\text{enc}}^{t'-1}), \quad (2)$$

where $\overrightarrow{\mathbf{h}}_{\text{enc}}^{t'}, \overleftarrow{\mathbf{h}}_{\text{enc}}^{t'} \in [-1, 1]^{2N}$ are the latent outputs of $\overrightarrow{\text{RNN}}_{\text{enc}}$ and $\overleftarrow{\text{RNN}}_{\text{enc}}$, respectively, at the t' -th time frame, $t' = 1, \dots, T+L$, $\overrightarrow{\mathbf{h}}_{\text{enc}}^0 = \overleftarrow{\mathbf{h}}_{\text{enc}}^0 = \{0\}^N$, $\overleftarrow{\mathbf{v}}_{\text{tr}}^{t'}$ is the time-flipped (i.e. backwards) version of \mathbf{V}_{tr} , and $\mathbf{H}'_{\text{enc}} = [\mathbf{h}'_{\text{enc}}^1, \dots, \mathbf{h}'_{\text{enc}}^{T+L}]$. Bi-directional RNN_{enc} is used to encode the input magnitude spectrogram, using extra information from the L temporal content vectors. The output of the encoder \mathbf{H}'_{enc} , is summed with the input \mathbf{V} , using residual connections as

$$\mathbf{H}_{\text{enc}} = \mathbf{H}'_{\text{enc}} + [\mathbf{V}_{\text{tr}}^{\top}, \overleftarrow{\mathbf{V}}_{\text{tr}}^{\top}]^{\top}, \quad (3)$$

where $\overleftarrow{\mathbf{V}}_{\text{tr}}$ is the magnitude spectrogram \mathbf{V}_{tr} flipped in time (i.e. backwards) and $\mathbf{H}_{\text{enc}} \in \mathbb{R}^{T+L \times 2N}$. Finally, the extra L time-frames are dropped from \mathbf{H}_{enc} , so the subsequent decoder will be able to focus on the time frames that correspond to the targeted output, as

$$\mathbf{H}_{\text{enc-tr}} = [\mathbf{h}_{\text{enc}}^{\lfloor L/2 \rfloor}, \dots, \mathbf{h}_{\text{enc}}^{T+\lfloor L/2 \rfloor}], \quad (4)$$

where $\mathbf{h}_{\text{enc}}^i \in [-1, 1]^N$ is the i -th vector of \mathbf{H}_{enc} and $\lfloor \cdot \rfloor$ is the floor function. $\mathbf{H}_{\text{enc-tr}}$ is used as an input to RNN_{dec} of masker, obtaining \mathbf{H}_{dec} as

$$\mathbf{h}_{\text{dec}}^t = \text{RNN}_{\text{dec}}(\mathbf{h}_{\text{enc-tr}}^t, \mathbf{h}_{\text{dec}}^{t-1}), \quad (5)$$

where $\mathbf{h}_{\text{dec}}^t$ is the latent output of the RNN_{dec} at the t -th time-frame, $t = 1, \dots, T$, $\mathbf{h}_{\text{dec}}^0 = \{0\}^N$, and $\mathbf{H}_{\text{dec}} = [\mathbf{h}_{\text{dec}}^1, \dots, \mathbf{h}_{\text{dec}}^T]$. \mathbf{H}_{dec} is given as an input to a feed-forward linear layer with shared weights through time, followed by a rectified linear unit (ReLU) as

$$\mathbf{h}_m^t = \text{ReLU}(\text{FNN}_m(\mathbf{h}_{\text{dec}}^t)), \quad (6)$$

where $\mathbf{h}_m^t \in \mathbb{R}_{>0}^F$ and $\mathbf{H}_m = [\mathbf{h}_m^1, \dots, \mathbf{h}_m^T]$. Finally, the output of the masker, $\hat{\mathbf{V}}'_j$, is calculated as

$$\hat{\mathbf{V}}'_j = \mathbf{V}' \odot \mathbf{H}_m, \quad (7)$$

where “ \odot ” is the Hadamard product and $\mathbf{V}' = [\mathbf{v}^{\lfloor L/2 \rfloor}, \dots, \mathbf{v}^{T+\lfloor L/2 \rfloor}]$ is a time-trimmed version of the input magnitude spectrogram \mathbf{V} (i.e. before the trimming process Tr).

The denoiser, accepts as an input the $\hat{\mathbf{V}}'_j$ and outputs $\hat{\mathbf{V}}_j$, and it consists of two feed-forward layers with shared weights through time and functioning as an auto-encoder, FNN_{d1} and FNN_{d2} , where each one is followed by a ReLU. Specifically, the first layer, FNN_{d1} , process the input to the decoder as

$$\mathbf{h}_{\text{d1}}^t = \text{ReLU}(\text{FNN}_{\text{d1}}(\hat{\mathbf{v}}_j^t)), \quad (8)$$

where $\mathbf{h}_{\text{d1}}^t \in \mathbb{R}_{>0}^{\lfloor F/2 \rfloor}$ and $\mathbf{H}_{\text{d1}} = [\mathbf{h}_{\text{d1}}^1, \dots, \mathbf{h}_{\text{d1}}^T]$. Then, the second layer, FNN_{d2} , process \mathbf{H}_{d1} as

$$\mathbf{h}_{\text{d2}}^t = \text{ReLU}(\text{FNN}_{\text{d2}}(\mathbf{h}_{\text{d1}}^t)), \quad (9)$$

where $\mathbf{h}_{d2}^t \in \mathbb{R}_{\geq 0}^F$ and $\mathbf{H}_{d2} = [\mathbf{h}_{d2}^1, \dots, \mathbf{h}_{d2}^T]$. The output of the denoiser, $\hat{\mathbf{V}}_j$ is calculated as

$$\hat{\mathbf{V}}_j = \hat{\mathbf{V}}_j' \odot \mathbf{H}_{d2}. \quad (10)$$

Finally, the masker and the denoiser are jointly optimized by minimizing

$$\begin{aligned} \mathcal{L} = & D_{\text{KL}}(\mathbf{V}_j \parallel \hat{\mathbf{V}}_j') + D_{\text{KL}}(\mathbf{V}_j \parallel \hat{\mathbf{V}}_j) \\ & + \lambda_1 \|\text{diag}\{\mathbf{W}_{\text{FNN}_m}\}\|_1 + \lambda_2 \|\mathbf{W}_{\text{FNN}_{d2}}\|_2^2, \end{aligned} \quad (11)$$

where \mathbf{V}_j is the targeted magnitude spectrogram of the j -th source, D_{KL} is the generalized Kullback-Leibler divergence, $\lambda_1 = 1 \times 10^{-2}$ and $\lambda_2 = 1 \times 10^{-4}$ are regularization terms, $\|\cdot\|_1$ is the ℓ_1 vector norm, and $\|\cdot\|_2^2$ is the L_2 matrix norm. $\text{diag}\{\mathbf{W}_{\text{FNN}_m}\}$ is the main diagonal of the weight matrix of the FNN_m (i.e. the elements w_{ij} of $\mathbf{W}_{\text{FNN}_m}$ with $i = j$). More information about the specific regularizations terms and optimization process, can be found at the original MaD and the MaDTwinNet papers [5], [19].

B. Replacing RNNs

In our proposed method, we replace the bi-directional RNN_{enc} and the unidirectional RNN_{dec} with two sets of convolutional blocks, CNN_{enc} and CNN_{dec} , respectively. Following recent and SOTA published work [16], we opt to employ depth-wise separable (DWS) convolutions and not typical convolutions for our CNN blocks. The DWS convolution is a factorized version of the typical convolution, that first applies a spatial-wise convolution, and then a channel-wise convolution. The spatial-wise convolution learns spatial relationships in the input features to the convolution. The channel-wise convolution, learns cross-channel relationships between the channels of the spatial-wise convolution.

Specifically, each DWS convolution block of our method consists of a CNN (the spatial-wise convolution CNN_d), followed by a leaky ReLU (LReLU), a batch-normalization process, another CNN (the channel-wise convolution CNN_s), and a ReLU, as

$$\mathbf{H} = \text{ReLU}(\text{CNN}_s(\text{BN}(\text{LReLU}(\text{CNN}_d(\mathbf{X}))))), \text{ where} \quad (12)$$

$$\begin{aligned} \mathbf{D}^{c_i, x_h - K_h, x_w - K_w} &= \text{CNN}_d(\mathbf{X}^{c_i}; \mathbf{K}_d^{c_i}) \\ &= (\mathbf{K}_d^{c_i} * \mathbf{X}^{c_i})(x_h - K_h, x_w - K_w) \\ &= \sum_{k_h=1}^{K_h} \sum_{k_w=1}^{K_w} \mathbf{X}^{c_i, x_h - k_h, x_w - k_w} \mathbf{K}_d^{c_i, k_h, k_w}, \end{aligned} \quad (13)$$

$$\begin{aligned} \mathbf{H}^{c_o, \phi_h, \phi_w} &= \text{CNN}_s(\mathbf{D}^{:, \phi_h, \phi_w}; \mathbf{K}_s^{c_o}) \\ &= \sum_{c_i=1}^{C_i} \mathbf{D}^{c_i, \phi_h, \phi_w} \mathbf{K}_s^{c_o, c_i}, \end{aligned} \quad (14)$$

$$\text{LReLU}(x) = \begin{cases} x, & \text{if } x \geq 0, \\ \beta x & \text{otherwise} \end{cases}, \quad (15)$$

BN is the batch normalization process, * indicates convolution, $\mathbf{D} \in \mathbb{R}^{C_i \times \Phi_h \times \Phi_w}$ and $\mathbf{K}_d \in \mathbb{R}^{C_i \times K_{dh} \times K_{dw}}$ are the output

and kernel tensors of CNN_d , respectively, $\mathbf{H} \in \mathbb{R}^{C_o \times \Phi_h' \times \Phi_w'}$ and $\mathbf{K}_s \in \mathbb{R}^{C_o \times C_i}$ are the output tensor and kernel matrix of CNN_s , respectively, and $\beta < 1$ is a hyper-parameter. Eq. (13) is used to learn the spatial relationships of the data $\mathbf{X} \in \mathbb{R}^{C_i \times X_h \times X_w}$, and Eq. (14) is used to learn the cross-channel relationships. We employ LReLU according to previous studies using depth-wise separable convolutions [16].

Our CNN_{enc} consists of one DWS convolution block that is followed by a batch normalization process, a max-pooling operation, and a dropout with p_{enc} probability, and then of L_{enc} DWS convolution blocks, with each block followed by a batch normalization process and a dropout with probability p_{enc} (but no max-pooling operation). The output of each of the L_{enc} DWS convolution blocks has the same dimensionality as the input. That is, at each of the L_{enc} DWS convolution blocks, we utilize proper zero padding (i.e. depending on the kernel size) in order not to alter the dimensions of the input. Each block of CNN_{enc} gets as an input the output of the previous one, the first gets as an input \mathbf{V} , and the last outputs the tensor $\mathbf{H}^{L_{\text{enc}}} \in \mathbb{R}_{\geq 0}^{C_{\text{enc}} \times H_{\text{enc}} \times W_{\text{enc}}}$.

CNN_{dec} consists of a transposed convolution, followed by two DWS convolution blocks, batch-normalization and max-pooling processes, a dropout with probability p_{dec} , a CNN, and the FNN_{enc} . The transposed convolution of CNN_{dec} gets as an input the $\mathbf{H}^{L_{\text{enc}}}$, and the FNN_{enc} outputs \mathbf{H}_m . Finally, the output of the masker of our method is calculated according to Eq. (7). The final audio signal of the output is calculated according to the original MaD paper [5].

III. EXPERIMENTAL PROCEDURE

A. Dataset and pre-processing

We use the development sub-set of Demixing Secret Dataset¹ (DSD100) for optimizing the parameters of the proposed method, in a supervised fashion. From each multi-track we compute a monaural version of each of the four sources, by averaging the two available channels. Then, we compute the STFT of each monaural signal using a Hamming window of 2049 samples (46ms) over a step size of 384 samples (8ms). Each windowed segment is zero-padded to 4096 samples. After the STFT, we remove the redundant information of the STFT retaining the first $N = 2049$ frequency bands, and then compute the absolute values. Then the magnitude spectrogram of the mixture and singing voice are segmented into $B = \lceil M/T \rceil$ sequences, with T being the length of the sequence, and $\lceil \cdot \rceil$ is the ceiling function. Each sequence b is employed as our \mathbf{V} and \mathbf{V}_j , for the mixture and target source respectively, and overlaps with the preceding one by an empirical factor of $L \times 2$. The overlap factor is used for aggregating context information in the previously described stages of encoding.

B. Hyperparameters and training of proposed method

We evaluate our method by conducting an ablation study, employing different amounts of CNN_{enc} blocks, L_{enc} , and different number of channels, C_o , for our convolutional kernels.

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

TABLE I
SDR, SIR, AND SAR VALUES, AND AMOUNT OF PARAMETERS (N_{PARAMS}) FOR THE DIFFERENT AMOUNTS OF CNN_{enc} BLOCKS (L_{enc}) AND CHANNELS OF THE CORRESPONDING KERNEL (C_0). VALUES OF SDR, SIR, AND SAR ARE PRESENTED IN dB. WITH BOLD FONTS ARE THE VALUES FOR THE COMBINATION OF L_{enc} AND C_0 , THAT YIELDS THE BIGGER SDR.

Value of L_{enc}	Value of C_0	SDR			SIR			SAR			N_{params}		
		64	128	256	64	128	256	64	128	256	64	128	256
5		4.47	4.84	4.91	8.11	8.07	8.59	6.71	6.74	6.98	4 783 426	4 922 754	5 447 170
7		4.44	4.65	4.94	7.83	8.23	8.23	6.67	6.91	7.15	4 795 586	4 963 458	5 594 114
9		4.46	4.57	4.88	8.14	8.06	8.44	6.45	6.59	6.95	4 807 746	5 004 162	5 741 058
11		4.46	4.64	4.83	8.03	8.77	8.98	6.39	6.84	8.98	4 819 906	5 044 866	5 888 002
13		4.59	4.80	4.72	8.17	9.02	8.41	6.40	6.56	6.97	4 832 066	5 085 570	6 034 946
15		4.39	4.58	4.76	8.77	8.41	8.60	6.20	6.61	7.01	4 844 226	5 126 274	6 181 890

Specifically, we employ six different number of L_{enc} , namely 5, 7, 9, 11, 13, and 15, and three different C_0 , namely 64, 128, and 256. We indicate the amount of L_{enc} and C_0 , using a subscript, e.g. $\text{CNN}_{\text{enc-5,64}}$ for the $L_{\text{enc}} = 5$ and $C_0 = 64$ combination. All DWS convolution blocks of CNN_{enc} have a square kernel of $K_{dh} = K_{dw} = 5$. At the CNN_{dec} we utilize the same C_0 amount of channels with the CNN_{enc} , $K_{dh} = K_{dw} = 5$, and a unit stride, and $\beta = 1e-2$. The values for K_{dh} and K_{dw} are chosen according to previous work that employed DWS convolutions [16] and the value for β as the default value for the LReLU in the PyTorch framework.

We optimize the parameters of our method following the approach in the original papers of MaD [6], [19], using 100 epochs on the training dataset, with a batch size of 4. We utilized the Adam optimizer for updating the weights of our method, with a learning rate of $1e-4$ and for betas we used the values proposed in the original corresponding paper [22]. Additionally, we employ a clipping of the gradient L_2 norm equal to 0.5, similar to the training process of the original MaD system. The above are implemented using the PyTorch framework, and our code is freely available online².

C. Objective Evaluation

We compare our method with an established masking based approach to singing voice separation, denoted as the Masker and Denoiser (MaD) architecture and friends, namely the MaD TwinNet [5] and the MaD architecture with the recurrent inference algorithm [6]. The length of the sequences for the MaD and friends is set to $T = 60$ timeframes, according to the corresponding papers [5]. We focus on those two particular approaches because to the best of our knowledge those approaches are the only ones that *do not* estimate all the other music sources in an attempt to re-fine the estimated singing voice signal [6]–[8], [16]. This allows us to clearly examine the potentials of using depth-wise separable convolutional networks for masking based approaches to singing voice separation. For assessment, the evaluation subset of DSD100 (50 mixtures and corresponding sources) is used for measuring the objective performance of our method, in terms of signal-to-distortion (SDR), signal-to-interference (SIR), and signal-to-artifacts (SAR) ratios. The computation of SDR, SIR, and SAR for all the compared methods is performed over overlapping

TABLE II
COMPARISON OF OUR PROPOSED METHOD WITH MaD, ON DSD DATASET. VALUES OF SDR, SIR, AND SAR ARE PRESENTED IN DB. $N_{\text{PARAMS-M}}$ IS THE AMOUNT OF PARAMETERS FOR MASKER. RESULTS OF MaD ARE TAKEN FROM THE LITERATURE.

Approach	SDR	SIR	SAR	$N_{\text{params-M}}$
MaD [6]	3.62	7.06	5.88	
MaD-Rec.Inference [6]	4.20	7.94	5.91	22 996 113
MaDTwinNet [5]	4.57	8.17	5.95	
$\text{CNN}_{\text{enc-7,256}}$	4.94	8.23	7.15	1 394 689

signal segments, following the proposed rules of the official Signal Separation and Evaluation Campaign (SiSEC) [23].

IV. RESULTS AND DISCUSSION

In Table I are amount of parameters and the obtained values for the SDR, SIR, and SAR versus the different L_{enc} and C_0 . From that table, it can be seen that the increase of C_0 has a bigger impact to the obtained SDR, SIR, and SAR, compared to the increase of L_{enc} . That is, the increase at the amount of channels benefits more the obtained SDR, SIR, and SAR, than the increase of the depth of the CNNs. Though, this benefit from C_0 could be attributed to the more pronounced effect that C_0 has on N_{params} . From Table I, it can be seen that the increase of C_0 has more impact on the total amount of parameters N_{params} , than increasing L_{enc} . Regarding the best performing combination, we focus on the SDR and we consider as best performing the combination of $L_{\text{enc}} = 7$ and $C_0 = 256$.

To evaluate the benefit of our proposed method compared to the usage of RNNs, we compare our results with the vanilla, with recurrent inference, and with twin networks variants of the MaD system. In Table II are the SDR, SIR, and SAR values of the best performing combination according to SDR and Table I (i.e. $\text{CNN}_{\text{enc-7,256}}$), compared to the values for the same metrics obtained from MaD system. Additionally, since one of the main benefits of DWS convolutions is that they have quite few parameters, we also list in Table I the amount of parameters of the Masker. We do not list the parameters for the Denoiser, since the Denoiser is the same in all the listed systems in Table I. For reference, the amount of parameters of the Denoiser is 4 199 425.

As can be seen from Table II, our proposed method surpassed all variants of the MaD system, while, at the same time, it has only 6% of the parameters at the Masker (i.e.

²<https://github.com/pppyyknkn/mad-twinnet>

94% reduction) compared to MaD. Specifically, we achieve an increase of 0.37 dB, 0.06 dB, and 1.20 dB for SRD, SIR, and SAR, respectively, when using our method and compared to the MaD system trained with the TwinNet regularization (i.e. MaD TwinNet), which is the best performing variant of MaD. As can be seen, the improvement is mainly attributed on the reduction of artifacts in the separated signal (i.e. increase in the SAR). This indicates that the replacement of the RNNs with DWS convolutions can result in signals that have less distortion from the separation method [24].

Finally, comparing Tables I and II, we can see that with our method, with $L_{enc} = 5$ and $C_o = 64$, we can use the 2.54% of the parameters of the Masker, and still have an increase of 0.76 dB at the SAR, while having a marginal reduction of 0.10 dB and 0.06 dB at SDR and SIR, respectively. Basically, this means that with our method we can significantly reduce the parameters of the Masker by 97.5%, while still getting some improvement at the reduction of distortion from the separation method (i.e. increase at the SAR). In terms of the amount of total parameters, with the best performing combination of our method, $CNN_{enc-7,256}$, we get a reduction of 79.43% (i.e. we use only 20.57% of the total MaD parameters), and with the $CNN_{enc-5,64}$ we get a reduction of 82.41% (i.e. we use only 17.59% of the total MaD parameters).

V. CONCLUSIONS

In this work we examined the effect in objective separation performance of replacing RNNs with depth-wise separable (DWS) CNNs. To assess our proposed approach, we focused on the singing voice separation task and we employed a SOTA performing architecture for monaural singing voice separation that is based on RNNs, we implemented our proposed replacements, and we evaluated the performance of the method with the replacements using an established and freely available dataset for music source separation. We evaluated the performance of the singing voice separation using the widely employed source separation metrics of signal-to-distortion (SDR), signal-to-artifacts (SAR), and signal-to-interference (SIR) ratios. The results show a clear benefit of using our approach, both in the performance and the total amount of parameters needed. Specifically, with our approach we managed to reduce the amount of total parameters by 79.43%, and achieve an increase of 0.37 dB, 0.06 dB, and 1.20 dB at SDR, SIR, and SAR, compared to the original method with RNNs. For future work, we intend to examine the usage of dilated convolutions, in order to exploit the strong temporal context of music (e.g. melody). Additionally, further investigation could be carried, regarding the benefit or having a bigger kernel at the channel-wise convolution, in the depth-wise separable convolution.

REFERENCES

[1] Z. Rafii, A. Liutkus, F. R. Stöter, S. I. Mimilakis, D. FitzGerald, and B. Pardo, "An Overview of Lead and Accompaniment Separation in Music," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 8, pp. 1307–1335, Aug. 2018.

[2] S. Venkataramani, J. Casebeer, and P. Smaragdis, "End-To-End Source Separation With Adaptive Front-Ends," in *2018 52nd Asilomar Conference on Signals, Systems, and Computers*, Oct. 2018, pp. 684–688.

[3] E. Tzinis and S. Venkataramani and Z. Wang and C. Subakan and P. Smaragdis, "Two-Step Sound Source Separation: Training on Learned Latent Targets," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020.

[4] S. I. Mimilakis, K. Drossos, and G. Schuller, "Unsupervised Interpretable Representation Learning for Singing Voice Separation," in *2020 28th European Signal Processing Conference (EUSIPCO)*, Aug. 2020.

[5] K. Drossos, S. I. Mimilakis, D. Serdyuk, G. Schuller, T. Virtanen, and Y. Bengio, "Mad twinnet: Masker-denoiser architecture with twin networks for monaural sound source separation," in *2018 International Joint Conference on Neural Networks (IJCNN)*, 2018, pp. 1–8.

[6] S. I. Mimilakis, K. Drossos, J. F. Santos, G. Schuller, T. Virtanen, and Y. Bengio, "Monaural Singing Voice Separation with Skip-Filtering Connections and Recurrent Inference of Time-Frequency Mask," in *Proceedings of the 43rd International Conference on Acoustics, Speech and Signal Processing (ICASSP 2018)*, Apr. 2018.

[7] L. Pr  tet, R. Hennequin, J. Royo-Letelier, and A. Vaglio, "Singing voice separation: A study on training data," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 506–510.

[8] F.-R. St  ter, S. Uhlich, A. Liutkus, and Y. Mitsufuji, "Open-Unmix - A Reference Implementation for Music Source Separation," *Journal of Open Source Software*, 2019. [Online]. Available: <https://doi.org/10.21105/joss.01667>

[9] M. Huber, G. Schindler, C. Sch  rkhuber, W. Roth, F. Pernkopf, and H. Fr  ning, "Towards real-time single-channel singing-voice separation with pruned multi-scaled densenets," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020, pp. 806–810.

[10] S. I. Mimilakis, K. Drossos, E. Cano, and G. Schuller, "Examining the Mapping Functions of Denoising Autoencoders in Singing Voice Separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 266–278, 2020.

[11] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *Proceedings of the 30th International Conference on Machine Learning - Volume 28*, ser. ICML13. JMLR.org, 2013, p. III1310III1318.

[12] A. Graves, "Generating sequences with recurrent neural networks," *CoRR*, vol. abs/1308.0850, 2013.

[13] D. Serdyuk, N.-R. Ke, A. Sordani, A. Trischler, C. Pal, and Y. Bengio, "Twin Networks: Matching the future for sequence generation," *CoRR*, vol. abs/1708.06742, 2017. [Online]. Available: <http://arxiv.org/abs/1708.06742>

[14] N. Takahashi and Y. Mitsufuji, "Multi-scale multi-band densenets for audio source separation," in *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, Oct. 2017.

[15] A. D  fossez, N. Usunier, L. Bottou, and F. Bach, "Music Source Separation in the Waveform Domain," HAL, Tech. Rep. 02379796v1, 2019.

[16] K. Drossos, S. I. Mimilakis, S. Gharib, Y. Li, and T. Virtanen, "Sound Event Detection with Depthwise Separable and Dilated Convolutions," in *Proceedings of the 2020 IEEE International Joint Conference on Neural Networks (IJCNN)*, Jul. 2020.

[17] J. Guo, Y. Li, W. Lin, Y. Chen, and J. Li, "Network decoupling: From regular to depthwise separable convolutions," in *British Machine Vision Conference (BMVC)*, Sep. 2018.

[18] Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 8, pp. 1256–1266, Aug. 2019.

[19] S. I. Mimilakis, K. Drossos, T. Virtanen, and G. Schuller, "A recurrent encoder-decoder approach with skip-filtering connections for monaural singing voice separation," in *2017 IEEE 27th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2017, pp. 1–6.

[20] Y. He and J. Zhao, "Temporal convolutional networks for anomaly detection in time series," *Journal of Physics: Conference Series*, vol. 1213, p. 042050, Jun. 2019.

[21] C. Lea, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks: A unified approach to action segmentation," in *Computer Vision – ECCV 2016 Workshops*, G. Hua and H. J  gou, Eds. Springer International Publishing, 2016, pp. 47–54.

- [22] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proceedings of the International Conference on Learning Representations (ICLR-15)*, 2015.
- [23] 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*, Feb. 2017, pp. 323–332.
- [24] E. Vincent, R. Gribonval, and C. Fevotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 4, pp. 1462–1469, 2006.

This figure "fig1.png" is available in "png" format from:

<http://arxiv.org/ps/2007.02683v1>