

Unsupervised Interpretable Representation Learning for Singing Voice Separation

Stylios I. Mimitakis
Semantic Music Techn. Group
Fraunhofer-IDMT
Ilmenau, Germany
mis@idmt.fraunhofer.de

Konstantinos Drossos
Audio Research Group
Tampere University
Tampere, Finland
konstantinos.drossos@tuni.fi

Gerald Schuller
Applied Media Systems Group
Technical University of Ilmenau
Ilmenau, Germany
gerald.schuller@tu-ilmenau.de

Abstract—In this work, we present a method for learning interpretable music signal representations directly from waveform signals. Our method can be trained using unsupervised objectives and relies on the denoising auto-encoder model that uses a simple sinusoidal model as decoding functions to reconstruct the singing voice. To demonstrate the benefits of our method, we employ the obtained representations to the task of *informed* singing voice separation via binary masking, and measure the obtained separation quality by means of scale-invariant signal to distortion ratio. Our findings suggest that our method is capable of learning meaningful representations for singing voice separation, while preserving conveniences of the short-time Fourier transform like non-negativity, smoothness, and reconstruction subject to time-frequency masking, that are desired in audio and music source separation.

Index Terms—representation learning, unsupervised learning, denoising auto-encoders, singing voice separation

I. INTRODUCTION

A particular task in music signal processing that has attracted a lot of research interest is the estimation of the singing voice source from within an observed mixture signal [1]. To that aim, deep supervised learning is shown to yield remarkable results. Approaches that rely on deep supervised learning can be discriminated in two categories, the ones that operate in the short-time Fourier transform (STFT) domain [2], [3], and we denote as spectral-based approaches, and the ones that operate directly on the waveform signals [4], [5], that we denote as waveform-based approaches. Spectral and waveform based approaches have in common that they implicitly compute source-dependent masks that are applied to the mixture signal, prior to the reconstruction of the target signals [2]–[5]¹.

Although the implicit masking is shown to be a simple and robust method to learn source dependent patterns for source separation [6], one could expect that waveform based approaches would significantly outperform the spectral ones. That is because waveform based approaches are optimized using time-domain signals that also contain the phase information, that unarguably carries important signal information [7], [8] and has been neglected by many spectral based approaches [2], [3], [9], [10]. However, experimental evidence shows that spectral based approaches have comparable or marginally better separation performance to the waveform

ones [2], [4], [5]. Since the state-of-the-art (SOTA) methods for both waveform and spectral approaches rely on deep neural networks, and in both spectral and waveform approaches a considerable engineering effort has been directed to the employed neural architecture, it is evident that the difference in the performance between the two different approaches can be attributed to the utilized signal representation. For the waveform-based ones this is the output of an encoder, but for the spectral-based it is the non-negative signal representation offered by the magnitude of the STFT. Thus, we believe that learning generalized signal representations for music signals is an intriguing direction for music source separation research.

In this work, we focus on representation learning [11] for singing voice separation in an attempt to bridge the gap between spectral and waveform based approaches. To this aim, we propose a simple method for unsupervised representation learning from waveform signals, alleviating the need of having paired training data (i.e., matched multi-track audio data). However, the method still requires isolated source’s audio signals. More specifically, our method is based on the denoising auto-encoder (DAE) model [12], but for the decoding functions our model for representation learning inherits a simple and real-valued sinusoidal model. The sinusoidal model consists of amplitude-modulated cosine functions, and whose parameters are jointly optimized with the rest of the DAE. The motivation behind using a sinusoidal model as a decoding function is to guide (via back-propagation) the encoding layers to learn and convey information regarding the energy of specific cosine functions that compose the audio signal, leading to interpretable representations.

Our method is inspired by the concept of differentiable digital signal processing [13] where the parameters of common digital signal processing functions are optimized by means of back-propagation, and in our case we back-propagate through the parameters of a simple signal model. Furthermore, our method is similar to the Sinc-Network presented in [14], that uses sinc functions in the encoding layers of convolutional kernels for interpretable deep learning, and its extension to complex-valued representations for speaker separation [15]. However, our method differs from [15] as the representation of the proposed method is real-valued, alleviating the cumbersome signal processing operations on complex numbers. It also differs from approaches that initialize the front-end parts of the networks with cosine functions [16] that are

Stylios I. Mimitakis is supported in part by the German Research Foundation (AB 675/2-1, MU 2686/11-1).

¹Regarding the masking strategy, we are referring to the adaptation of Conv-TasNet for music signals also presented in [4].

then updated by means of back-propagation, by inheriting the cosine functions as a part of the model to be optimized. Finally, our method provides an unsupervised alternative to the source informed method for representation learning presented in [17]. The rest of the document is organized as follows: Section II presents the proposed method, Section III describes the followed experimental procedure, Section IV discusses the obtained results, and Section V concludes this work.

II. PROPOSED METHOD

Our proposed method employs two functions, the encoder E , and the decoder D . The input to our method is a music signal $\mathbf{x} \in \mathbb{R}^N$ of N time-domain samples, and the output is the learned representation of \mathbf{x} , denoted as $\mathbf{A} \in \mathbb{R}^{C \times T}$. C is the number of templates and T is the temporal length of each template (similarly as the time-frames in STFT-related representations). The encoder E learns the representation \mathbf{A} with the help of the decoder D . D is responsible for reconstructing the signal, given the representation computed by E . The reconstructed signal can then be used to optimize E and D using a reconstruction objective. To enforce interpretability for the representation \mathbf{A} , we use a differentiable sinusoidal synthesis model for the decoder D . An illustration of the proposed method is given in Figure 1.

A. The Encoder

During inference, the encoder E gets as an input any music signal \mathbf{x} and outputs its representation \mathbf{A} . In order for E to yield the representation \mathbf{A} , an initial stage of training is performed. During training, two synthetic signals are used. Each synthetic signal employs the singing voice signal ($\mathbf{x}_v \in \mathbb{R}^N$). The first synthetic signal is termed as $\tilde{\mathbf{x}}_m \in \mathbb{R}^N$ and is the result of a corruption process for \mathbf{x}_v with an additive generic multi-modal distribution-based noise (e.g. a randomly selected signal that contains accompaniment music, like a mixture of drums, guitars, synthesizers, and bass). The second signal is termed as $\tilde{\mathbf{x}}_v$ and is the result of a corruption process for \mathbf{x}_v using additive Gaussian noise.

Both signals $\tilde{\mathbf{x}}_m$, $\tilde{\mathbf{x}}_v$ are used independently as an input to E , resulting into two representations \mathbf{A}_m , $\mathbf{A}_v \in \mathbb{R}^{C \times T}$, respectively. To compute each representation, E consists of two one-dimensional strided convolutions, with appropriate zero-padding. The first operation involves a convolution of each signal with a set of C number of kernels of temporal length L and a stride S . The stride S is a hyper-parameter and affects the expected number of time-frames T by $T = \lceil N/S \rceil$, where $\lceil \cdot \rceil$ is the ceiling function. The resulting latent signal is given to the second convolution, which is a *dilated* one-dimensional convolution [18] with C number of kernels, a smaller temporal length $L' \ll L$, and a stride equal to 1. The output of the second convolution is updated by means of residual connections using the output from the first convolution, followed by the rectified linear unit (ReLU) activation function [19]. The ReLU function promotes a non-negative and sparse representation by preserving positive values and

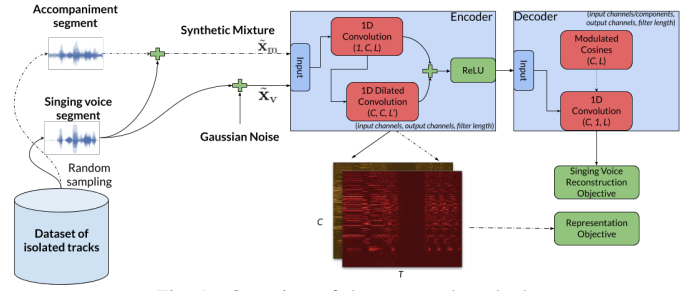


Fig. 1. Overview of the proposed method.

setting the rest to zero [20], and is shown to be particularly useful in general modelling of audio signals [21].

Another targeted (and useful) property of the representation is that of smoothness [16], [21], especially useful when real-valued cosine functions are involved in auto-encoding or separation models [16]. That is because audio signal modelling based on cosine functions requires the phase information for reconstruction. Phase information is usually encoded as the sign (positive or negative value) of the real-valued representation that varies along the time-frames of the representation. Since the negative values are nullified by the application of the ReLU function, neighbouring time-frames, that convey similar information for music signals are expected to be non-smooth. To compensate for that, the second convolution operation of E is using dilated convolutions that aggregate temporal information from neighboring time-frames [18], [22].

In order to enforce the learning of smooth representations, we employ a representation objective that the encoder has to minimize. Specifically, we use the representation of $\tilde{\mathbf{x}}_m$, \mathbf{A}_m , to compute the total variation denoising [23] (\mathcal{L}_{TV}) as

$$\mathcal{L}_{TV}(\mathbf{A}_m) = \frac{1}{CT} \left(\sum_{c=1}^{C-1} \|\mathbf{A}_m[c;\dots] - \mathbf{A}_m[c-1;\dots]\| + \sum_{t=1}^{T-1} \|\mathbf{A}_m[\dots;t] - \mathbf{A}_m[\dots;t-1]\| \right), \quad (1)$$

where $\mathbf{A}_m[c;\dots]$ and $\mathbf{A}_m[\dots;t]$ are the c -th row and t -th column vectors (respectively) of the matrix \mathbf{A}_m , and $\|\cdot\|$ is the ℓ_1 vector norm. Eq. (1) penalizes E by the norm of the first order difference across both time-frames T and templates C . The former promotes slow time varying representations as the magnitude of the STFT representation, and the latter promotes a grouping of the template activity. We use \mathbf{A}_m only to compute \mathcal{L}_{TV} , to enforce the encoder E to yield smooth representations on the most realistic corruption scenario. This scenario is the additive generic multi-modal distribution-based noise $\tilde{\mathbf{x}}_m$ that contains also the information regarding the singing voice signal \mathbf{x}_v . Thus, the smoothness for the representation of the singing voice is implicitly enforced.

B. The Decoder

The decoder D accepts the representation \mathbf{A}_v , and yields $\hat{\mathbf{x}}_v$ which is the approximation of the clean singing voice \mathbf{x}_v . Specifically, D models \mathbf{x}_v as a sum of C signal components that overlap in \mathbb{R}^N . The components are computed by a strided

convolution² between the representation template $\mathbf{A}_{v_{[c:\dots]}}$ and the kernel $\mathbf{w}_c \in \mathbb{R}^L$ of temporal length L as

$$\mathbf{x}_v \approx \hat{\mathbf{x}}_v := \sum_{c=0}^{C-1} \mathbf{A}_{v_{[c:\dots]}} * \mathbf{w}_c. \quad (2)$$

Similar to Sinc-Net [14] and its complex-valued extension for speech enhancement [15], we do not allow each \mathbf{w}_c to be updated directly using back-propagation. Instead, we re-parameterize each \mathbf{w}_c using sinusoidal functions and back-propagate through their corresponding parameters. More specifically, we compute each \mathbf{w}_c using

$$\mathbf{w}_c = \cos(2\pi f_c^2 \odot \mathbf{t} + \phi_c) \odot \mathbf{b}_c, \quad (3)$$

where \cos and \odot are the element-wise cosine function and product, respectively, and $\mathbf{t} \in \mathbb{Z}^L$ is a vector denoting the integer time indices $[0, \dots, L-1]$ of the kernels. These parameters of the cosine function are considered constants and are shared between the kernels. The sampling-rate-normalized carrier frequency f_c , the phase ϕ_c (in radians), and the modulating signal \mathbf{b}_c are learnable and different for each kernel. The non-linear squaring operation applied to f_c is motivated by the increased frequency resolution in lower frequencies that music signals commonly have [1], and is an experimental finding that is studied in Section IV. Using Eq. (3) for all C , our method constructs $\mathbf{W} \in \mathbb{R}^{C \times L}$ by stacking the corresponding outcome. After the stacking, a sorting operation is applied to \mathbf{W} , which sorts the kernels \mathbf{w}_c in ascending order based on the normalized and squared carrier frequency f_c , promoting an intuitive representation. Then the decoding operation for \mathbf{A}_v takes place using Eq. (2).

There are three reasons for using modulated cosine functions for decoding \mathbf{A}_v : a) cosine functions promote interpretability [14], i.e., the representation \mathbf{A} is expected to convey amplitude related information for driving a well established synthesis model based on sinusoidal functions [24], b) the auto-encoding operation shares many similarities with the STFT yet without having to deal directly with the phase information, for which supervised based separation works remarkably well [2], [3], and c) amplitude modulations allow an extra degree of freedom in reconstructing signals that cannot be described by pure sinusoidal functions [24]. The latter statement is supported by the convolution theorem which states that the element-wise product of two vectors can be expressed in the Fourier domain as their corresponding convolution. Since in our re-parameterization scheme (i.e., Eq. (3)) one of the signals is a cosine function, then \mathbf{b}_c is expected to convey information regarding fricatives and/or formants of the singing voice signal \mathbf{x}_v . Regarding on whether the proposed decoder is efficient in reconstructing the singing voice compared to either cosine functions or commonly employed convolutional layers, the reader is kindly referred to Section IV.

The optimization objective for D is the negative signal-to-

noise ratio (neg-SNR) [25], defined as:

$$\mathcal{L}_{\text{neg-SNR}}(\mathbf{x}_v, \hat{\mathbf{x}}_v) = -10 \log_{10} \left(\frac{\|\mathbf{x}_v\|_2^2}{\|\mathbf{x}_v - \hat{\mathbf{x}}_v\|_2^2} \right), \text{ where} \quad (4)$$

$\|\cdot\|_2^2$ is the squared ℓ_2 vector norm, and the negative sign is used to cast the logarithmic SNR as a minimization problem. Using Eq.(4) and Eq.(1) the overall minimization objective is

$$\mathcal{L} = \mathcal{L}_{\text{neg-SNR}} + \lambda \mathcal{L}_{\text{TV}} \quad (5)$$

where λ is a scalar for weighting the impact of Eq.(1) in the learning signal. The decoder D computes $\hat{\mathbf{x}}_v$ only from the signing voice representation \mathbf{A}_v . That is because we aim at learning general representations in an unsupervised and not discriminative fashion. To achieve that by means of the DAE model [12], we assume that the distribution of the corruption process is constant for all segments in the data-set [26]. This cannot be assumed for music signal mixtures, as even the distribution of the accompaniment instruments can vary dramatically from one segment to another. Consequently, by making such an assumption it could lead to degenerate representations for singing voice.

III. EXPERIMENTAL PROCEDURE

For training and evaluating the proposed method we use the MUSDB18 data-set [27] that consists of 150 two-channel multi-tracks, sampled at 44100Hz and split into training (100 multi-tracks) and testing (50 multi-tracks) subsets. During training we sample a set of four multi-tracks from which we use the vocals and the accompaniment sources. Each sampled multi-track is down-mixed to a single channel and is partitioned into overlapping segments of $N = 44100$ samples with an overlap of 22050 samples. We then randomly shuffle the segments for each source and corrupt the singing voice signal as described in Section II. The standard deviation of the additive Gaussian noise corruption is set to $1e-4$ and is independent from the signal's amplitude. A batch of 8 segments is used for optimizing the parameters of the proposed method, minimizing Eq. (5) using adam algorithm [28] with a learning rate of $1e-4$. For choosing the convolution hyper-parameters we conducted an informal experiment employing 20 tracks from the training subset, followed by informal listening tests. This resulted into the following hyper-parameters: $C = 800$, $S = 256$, $L = 2048$, $L' = 5$, $D = 10$, and $\lambda = 0.5$. During the informal experiments, we observed that the method converges fast so we set the total number of iterations throughout the whole data to 10. The choice for $N = 44100$ samples was based on the available computational resources.

For evaluation we use the rest 50 tracks, that are down-mixed and partitioned into non-overlapping segments. The shuffling and random mixing are not considered in the evaluation stage, but silent segments are discarded. We test the usefulness of the representation by performing informed and masking-based singing voice separation, following the recently proposed framework for assessing latent representations for audio source separation [17]. To that aim, we employ the trained decoder (according to the previously described procedure), and reconstruct the time-domain signals of the un-

²Appropriate zero-padding is assumed to be applied in order to deal with the differences between T and L .

corrupted singing voice representation and the binary masked mixture representation, respectively. The binary mask is computed using the encoded singing voice, accompaniment, and their corresponding mixture signals, that are available in the test sub-set. The reconstructed time-domain signals are used for computing the scale-invariant signal-to-distortion ratio (SI-SDR) [29] defined as

$$\text{SI-SDR}(\mathbf{x}_v, \hat{\mathbf{x}}_v) = 10 \log_{10} \left(\frac{\|\alpha \mathbf{x}_v\|_2^2}{\|\alpha \mathbf{x}_v - \hat{\mathbf{x}}_v\|_2^2} \right), \text{ for } \alpha = \frac{\hat{\mathbf{x}}_v^T \mathbf{x}_v}{\|\mathbf{x}_v\|_2^2}, \quad (6)$$

and is used computed for each segment. In the following section, we report the median value of SI-SDR across segments and three experimental runs.

Using the above described procedure, we conduct two experiments. In the first experiment, we examine whether the modulated cosine functions (`mod-cos`) are a good synthesis model by measuring the reconstruction performance, after being optimized for the denoising task. We optimize various models that use the proposed training scheme presented in Section II *without* the random stopping corruption processes, and by employing the early stopping mechanism to terminate the training procedure if the model has stopped decreasing the loss expressed in Eq. (4) during the updates of the previous epoch. We consider various decoding strategies such as non-modulated cosine functions (`cos`), and common one-dimensional convolutional networks (`conv`) with and without the \tanh non-linearity at the output. We also examine Sinc-Net [14] (`sinc`) as the first encoding stage as proposed in [14]. In this experiment C is adapted so that each model uses approximately the same number of parameter.

For the second experiment we re-train the best combination of the above, using various values for the number of components $C \in [400, 800, 1600]$ and perform the reconstruction of the binary masked mixture signal. To examine the regularization effect of the total-variation (Eq. (1)) computed using the random mixing corruption process, we report each model’s performance by using Eq. (1) for both \mathbf{A}_v and \mathbf{A}_m , respectively. For comparison, we employ the STFT and perform the above described operations of analysis, masking, and synthesis. The STFT uses a hop-size of 384 samples, a window size of 2048 samples, and the hamming windowing function. The difference between the first and the second set of experiments is that for the second set of experiments the modulated cosine functions are sorted after each gradient update, as explained in Section II-B, whereas in the first they are not. The sorting is performed for the representation to have information analogous other cosine related transforms.

IV. RESULTS & DISCUSSION

The obtained results from the two experiments are presented in Tables I and II. Additional results, illustrations underlining the interpretability of the representations, and audio examples can be found online³. Table I demonstrates the median SI-SDR expressed in dB (the higher the better) yielded by the first experiment, along with additional information regarding

³https://github.com/Js-Mim/rl_singing_voice

TABLE I
RESULTS REFLECTING THE DECODING PERFORMANCE, BY MEANS OF SI-SDR. BOLD-FACED NUMBERS DENOTE THE BEST PERFORMANCE.

E/D Setup	Non-linearity	C	SI-SDR	N_P
<code>conv/cos</code>	N/A	952	20.83	6.483M
	f_c^2		22.34	
<code>conv/conv</code>	N/A	800	31.25	6.476M
	$\tanh(\text{decoder})$		30.50	
<code>conv/mod-cos</code>	N/A	800	28.72	6.478M
	f_c^2		32.62	
<code>sinc/mod-cos</code>	f_c^2	952	26.82	6.487M

TABLE II
SI-SDR FOR INFORMED SEPARATION BY BINARY MASKING (BM). BOLD-FACED NUMBERS DENOTE THE BEST PERFORMANCE.

E/D Setup	C	$\mathcal{L}_{TV}(\ast)$	SI-SDR	BM SI-SDR	N_P
<code>conv/mod-cos</code>	400	\mathbf{A}_v	30.46	3.66	2.439M
		\mathbf{A}_m	30.73	5.93	
	800	\mathbf{A}_v	32.28	4.39	6.478M
		\mathbf{A}_m	32.11	6.28	
	1600	\mathbf{A}_v	31.94	4.68	19.356M
		\mathbf{A}_m	31.54	6.68	
<code>STFT/iSTFT</code>	1025	N/A	N/A	8.80	N/A

the various setups for the encoder E and the decoder D , the number of parameters N_P (in millions M), the used number of components C , and the employed non-linearities. The results in Table I highlight three trends. First, the application of the non-linearity to the normalized frequencies f_c results into better reconstruction performance compared to the linear case. The observed improvement is of ~ 5 dB on average across experimental configurations. Secondly, the modulated cosine functions serve as a good differentiable synthesis model for singing voice signals, outperforming simple cosine functions by approximately 8 dB on average, with respect to the two experimental configurations (with and without frequency scaling of the normalized frequency), and by 1.4 dB the best configuration of convolution based model (`conv`). Since SI-SDR is invariant to scale modifications of the assessed signal, 1.4 dB is a significant improvement of signal quality and does not imply a simple matching of the gain that the model based on modulated cosine functions might have exploited. Thirdly, Sinc-Net [14] does not bring further improvements to the proposed method.

Focusing on the separation performance of the obtained representations, Table II presents the median SI-SDR values of the binary masking separation scenario for three values for the hyper-parameter C and two regularization strategies including two different signal representations, the corrupted by Gaussian noise \mathbf{A}_v , and the synthetic mixtures using the accompaniment signals \mathbf{A}_m . The obtained results are compared to the STFT that has perfect reconstruction properties and masking techniques work very well in practice [1]. The results of Table II underline two main experimental findings. The first finding is that the binary masking can be used to separate sources using the proposed approach for representation learning. This can be seen from the $C = 1600$ model that uses the synthetic mixtures as an input to the unsupervised representation objective and achieves a SI-SDR median value of 6.68 dB. The second finding is that the proposed unsupervised representation objective, i.e., Eq. (1) with the synthetic mixtures, can be used to improve

the reconstruction of the masked mixture signals without additional supervision, as previous studies suggest [17]. This claim is supported by the observed improvement of ~ 2 dB, on average across models of various components C , when the synthetic mixtures are used for the unsupervised representation objective. Nonetheless, there is much room for improvements in order to obtain the quality of the STFT/iSTFT approach that outperforms the best masked approximation of the proposed method by 2.12 dB.

V. CONCLUSIONS

In this work we presented a method for learning music signal representations in an unsupervised way. Our method is based on the denoising autoencoder model [12] and the differential digital signal processing concept [13]. The benefits of our method are interpretability, non-negativity for real-valued music signal representations for driving an established synthesis model, based on cosine functions. We conducted a series of experiments where we investigated the reconstruction capabilities of the proposed method subject to auto-encoding and informed source separation using binary masks. Our results demonstrate a reconstruction above 30 dB of scale-invariant signal-to-distortion ratio, and that separation by masking is possible using the obtained representation. The latter, opens up directions for supervised approaches to masking-based separation. However, compared to the short-time Fourier transform and its inverse counterpart our results suggest that there is much room for improvements in order to achieve the benefits of the STFT.

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] F.-R. Stöter, 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>
- [3] R. Hennequin, A. Khlif, F. Voituret, and M. Moussallam, "Spleeter: A Fast And State-of-the Art Music Source Separation Tool With Pre-trained Models," Late-Breaking/Demo ISMIR 2019, November 2019, deezer Research.
- [4] A. Défossez, N. Usunier, L. Bottou, and F. Bach, "Music Source Separation in the Waveform Domain," HAL, Tech. Rep. 02379796v1, 2019.
- [5] D. Samuel, A. Ganeshan, and J. Naradowsky, "Meta-learning Extractors for Music Source Separation," in *Proceedings of the 45th International Conference on Acoustics, Speech and Signal Processing (ICASSP 2020)*, May 2020.
- [6] 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.
- [7] P. Magron and T. Virtanen, "Online Spectrogram Inversion for Low-Latency Audio Source Separation," *IEEE Signal Processing Letters*, vol. 27, pp. 306–310, 2020.
- [8] P. Magron, K. Drossos, S. I. Mimilakis, and T. Virtanen, "Reducing Interference with Phase Recovery in DNN-based Monaural Singing Voice Separation," in *Proc. Interspeech 2018*, 2018, pp. 332–336.
- [9] 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 *Proceedings of the 2018 IEEE International Joint Conference on Neural Networks (IJCNN)*, July 2018.
- [10] 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)*, 2018.
- [11] Y. Bengio, A. Courville, and P. Vincent, "Representation Learning: A Review and New Perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, Aug 2013.
- [12] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion," *Journal of Machine Learning Research*, vol. 11, pp. 3371–3408, 2010.
- [13] J. Engel, L. H. Hantrakul, C. Gu, and A. Roberts, "DDSP: Differentiable Digital Signal Processing," in *International Conference on Learning Representations*, 2020.
- [14] M. Ravanelli and Y. Bengio, "Interpretable Convolutional Filters with SincNet," in *International Conference on Neural Information Processing Systems: Workshop on Interpretability and Robustness for Audio, Speech and Language*, 2018.
- [15] M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, "Filterbank Design for End-to-end Speech Separation," in *Proceedings of the 45th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2020)*, 2020.
- [16] 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.
- [17] 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 *Proceedings of the 45th International Conference on Acoustics, Speech and Signal Processing (ICASSP 2020)*, May 2020.
- [18] Y. Fisher and K. Vladlen, "Multi-Scale Context Aggregation by Dilated Convolutions," in *Proceedings of the International Conference on Learning Representations (ICLR)*, May 2016.
- [19] V. Nair and G. E. Hinton, "Rectified Linear Units Improve Restricted Boltzmann Machines," in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ser. ICML'10. Madison, WI, USA: Omnipress, 2010, p. 807–814.
- [20] V. Pappas, Y. Romano, and M. Elad, "Convolutional Neural Networks Analyzed via Convolutional Sparse Coding," *Journal of Machine Learning Research*, vol. 18, no. 83, pp. 1–52, 2017. [Online]. Available: <http://jmlr.org/papers/v18/16-505.html>
- [21] P. Smaragdis and S. Venkataramani, "A Neural Network Alternative to Non-Negative Audio Models," in *Proceedings of the 42nd IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2017)*, March 2017, pp. 86–90.
- [22] 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)*, 2020.
- [23] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear Total Variation Based Noise Removal Algorithms," in *Proceedings of the Eleventh Annual International Conference of the Center for Nonlinear Studies on Experimental Mathematics: Computational Issues in Nonlinear Science: Computational Issues in Nonlinear Science*, 1992, p. 259–268.
- [24] X. Serra, "A System for Sound Analysis/Transformation/Synthesis based on a Deterministic plus Stochastic Decomposition," Ph.D. dissertation, Stanford University, 1989.
- [25] I. Kavalerov, S. Wisdom, H. Erdogan, B. Patton, K. Wilson, J. Le Roux, and J. R. Hershey, "Universal Sound Separation," in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, Oct 2019, pp. 175–179.
- [26] P. Vincent, "A Connection between Score Matching and Denoising Autoencoders," *Neural Computation*, vol. 23, no. 7, p. 1661–1674, Jul. 2011.
- [27] Z. Rafii, A. Liutkus, F. Stöter, S. I. Mimilakis, and R. Bittner, "The MUSDB18 Corpus for Music Separation," Dec 2017. [Online]. Available: <https://doi.org/10.5281/zenodo.1117372>
- [28] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proceedings of the International Conference on Learning Representations (ICLR-15)*, 2015.
- [29] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR – Half-baked or Well Done?" in *Proceedings of the 44th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2019)*, May 2019, pp. 626–630.