

# MUSIC SOURCE SEPARATION USING STACKED HOURGLASS NETWORKS

Sungheon Park      Taehoon Kim      Kyogu Lee      Nojun Kwak  
Graduate School of Convergence Science and Technology, Seoul National University, Korea  
{sungheonpark, kcjs55, kglee, nojunk}@snu.ac.kr

## ABSTRACT

In this paper, we propose a simple yet effective method for multiple music source separation using convolutional neural networks. Stacked hourglass network, which was originally designed for human pose estimation in natural images, is applied to a music source separation task. The network learns features from a spectrogram image across multiple scales and generates masks for each music source. The estimated mask is refined as it passes over stacked hourglass modules. The proposed framework is able to separate multiple music sources using a single network. Experimental results on MIR-1K and DSD100 datasets validate that the proposed method achieves competitive results comparable to the state-of-the-art methods in multiple music source separation and singing voice separation tasks.

## 1. INTRODUCTION

Music source separation is one of the fundamental research areas for music information retrieval. Separating singing voice or sounds of individual instruments from a mixture has grabbed a lot of attention in recent years. The separated sources can be further used for applications such as automatic music transcription, instrument identification, lyrics recognition, and so on.

Recent improvements on deep neural networks (DNNs) have been blurring the boundaries between many application domains, including computer vision and audio signal processing. Due to its end-to-end learning characteristic, deep neural networks that are used in computer vision research can be directly applied to audio signal processing area with minor modifications. Since the magnitude spectrogram of an audio signal can be treated as a 2D single-channel image, convolutional neural networks (CNNs) have been successfully used in various music applications, including the source separation task [1, 8]. While very deep CNNs are typically used in computer vision literature with very large datasets [4, 25], CNNs used

for audio source separation so far have relatively shallow architectures.

In this paper, we propose a novel music source separation framework using CNNs. We used stacked hourglass network [18] which was originally proposed to solve human pose estimation in natural images. The CNNs take spectrogram images of a music signal as inputs, and generate masks for each music source to separate. An hourglass module captures both holistic features from low resolution feature maps and fine details from high resolution feature maps. The module outputs 3D volumetric data which has the same width and height as those of the input spectrogram. The number of output channels equals the number of music sources to separate. The module is stacked for multiple times by taking the results of the previous module. As passing multiple modules, the results are refined and intermediate supervision helps faster learning in the initial state. We used a single network to separate multiple music sources, which reduces both time and space complexity for training as well as testing.

We evaluated our framework on a couple of source separation tasks: 1) separating singing voice and accompaniments, and 2) separating bass, drum, vocal, and other sounds from music. The results show that our method outperforms existing methods on MIR-1K dataset [5] and achieves competitive results comparable to state-of-the-art methods on DSD100 dataset [30] despite its simplicity.

The rest of the paper is organized as follows. In Section 2, we briefly review the literature of audio source separation focusing on DNN based methods. The proposed source separation framework and the architecture of the network are explained in Section 3. Experimental results are provided in Section 4, and the paper is concluded in Section 5.

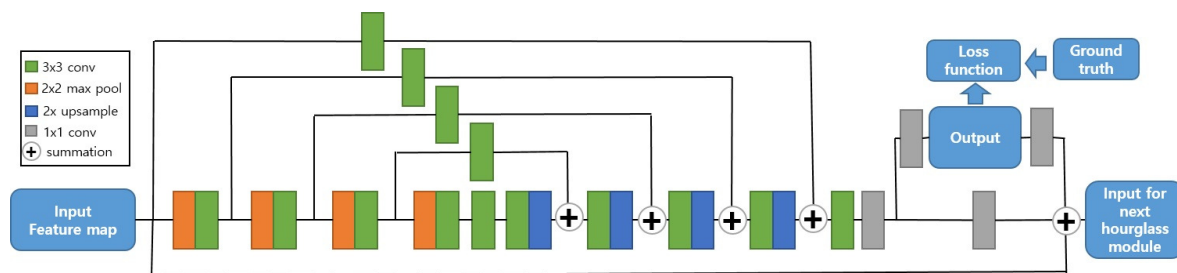
## 2. RELATED WORK

Non-negative matrix factorization (NMF) [12] is one of the most widely-used algorithms for audio source separation. It has been successfully applied to monaural source separation [32] and singing voice separation [29, 38]. However, despite its generality and flexibility, NMF is inferior to recently proposed DNN-based methods in terms of performance and time complexity.

Simple deep feed-forward networks consisting of multiple fully-connected layers showed reasonable performance for supervised audio source separation tasks [27]. Wang *et*



© Sungheon Park, Taehoon Kim, Kyogu Lee, Nojun Kwak. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Sungheon Park, Taehoon Kim, Kyogu Lee, Nojun Kwak. "Music Source Separation Using Stacked Hourglass Networks", 19th International Society for Music Information Retrieval Conference, Paris, France, 2018.



**Figure 1.** Structure of the hourglass module used in this paper. We follow the structure proposed in [17] except that the number of feature maps are set to 256 for all convolutional layers.

*al.* [34] used DNNs to learn an ideal binary mask which boils the source separation problem down to a binary classification problem. Simpson *et al.* [24] proposed a convolutional DNN to predict a probabilistic binary mask for singing voice separation. Recently, a fully complex-valued DNN [13] is proposed to integrate phase information into the magnitude spectrograms. Deep NMF [11] combined DNN and NMF by designing non-negative deep network and its back-propagation algorithm.

Since an audio signal is time series data, it is natural to use a sequence model like recurrent neural networks (RNNs) for music source separation tasks to learn temporal information. Huang *et al.* [6] proposed an RNN framework that jointly optimizes masks of foreground and background sources, which showed promising results for various source separation tasks. Other approaches include a recurrent encoder-decoder that exploits gated recurrent unit [15] or discriminative RNN [33].

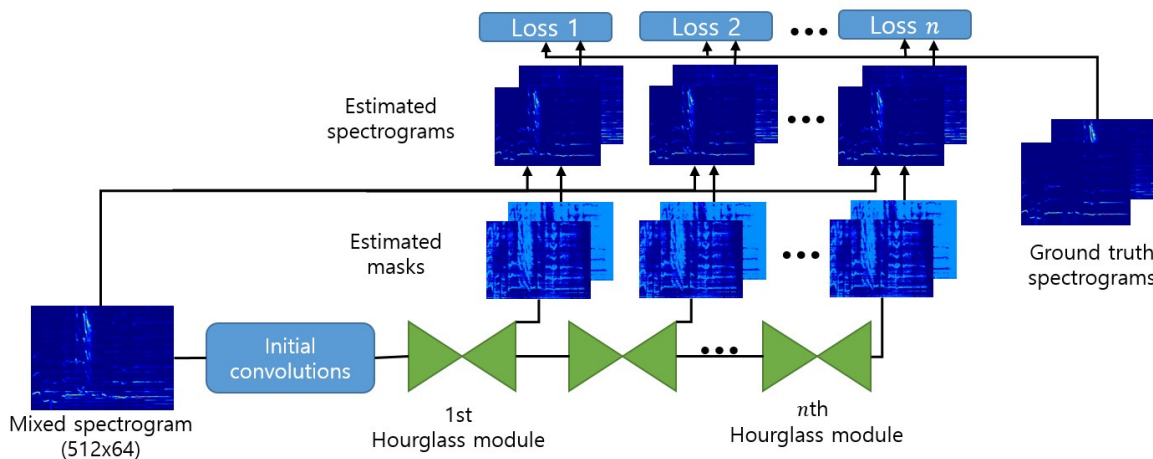
CNNs are also an effective tool for audio signal analysis when the magnitude spectrogram is used as an input. Fully convolutional networks (FCNs) [14] are initially proposed for semantic segmentation in the computer vision area, which is also effective for solving human pose estimation [18, 35] or super-resolution [2]. FCNs usually contain downsampling and upsampling layers to learn meaningful features at multiple scales. Strided convolution or pooling is used for downsampling, while transposed convolution or nearest neighbor interpolation is mainly used for upsampling. It is proven that FCNs are also effective in signal processing. Chandna *et al.* [1] proposed encoder-decoder style FCN for monoaural audio source separation. Recently, singing voice separation using an U-Net architecture [8] showed impressive performance. U-Net [22] is a FCN which consists of a series of convolutional layers and upsampling layers. There is a skip connection which connects the convolutional layers of the same resolution. They trained vocal and accompaniment parts separately on different networks. Miron *et al.* [16] proposed the method that separates multiple sources using a single CNN. They used score-filtered spectrograms as inputs and generated masks for each source via an encoder-decoder CNN. Multi-resolution FCN [3] was proposed for monoaural audio source separation. Recently proposed CNN architecture [26] based on DenseNet [7] achieved state-of-the-art performance on DSD100 dataset.

### 3. METHOD

#### 3.1 Network Architecture

The stacked hourglass network [18] was originally proposed to solve human pose estimation in RGB images. It is an FCN consisting of multiple hourglass modules. The hourglass module is similar to U-Net [22], of which feature maps at lower (coarse) resolution are obtained by repeatedly applying convolution and pooling operations. Then, the feature maps at the lowest resolution are upsampled via nearest neighbor interpolation with a preceding convolutional layer. Feature maps at the same resolution in the downsampling and the upsampling steps are connected with an additional convolutional layer. The hourglass module captures features at different scales by repeating pooling and upsampling with convolutional layers at each resolution. In addition, multiple hourglass modules are stacked to make the network deeper. As more hourglass modules are stacked, the network learns more powerful and informative features which refine the estimation results. Loss functions are applied at the end of each module. This intermediate supervision improves training speed and performance of the network.

The structure of a single hourglass module used in this paper is illustrated in Fig 1. Considering the efficiency and the size of the network, we adopt the hourglass module used in [17] which is a smaller network than the originally proposed one in [18]. A notable difference is that the residual blocks [4] used in [18] are replaced with a single convolutional layer. This light-weight structure showed competitive performance to the original network in human pose estimation with much smaller number of parameters. In the module, there are four downsampling and upsampling steps. All convolutional layers in downsampling and upsampling steps have filter size of  $3 \times 3$ . The  $2 \times 2$  max pooling is used to halve the size of the feature maps, and the nearest neighbor interpolation is used to double the size of the feature maps in the upsampling steps. We fixed the size of the maximum feature maps in convolutional layers to 256 which is different from [17]. After the last upsampling layer, a single  $3 \times 3$  convolution and two  $1 \times 1$  convolution is performed to generate network outputs. Then, an  $1 \times 1$  convolution is applied to the outputs to match the number of channels to that of the input feature maps. Another  $1 \times 1$  convolution is also applied to the feature maps



**Figure 2.** Overall music source separation framework proposed in this paper. Multiple hourglass modules are stacked, and each module outputs masks for each music source. The masks are multiplied with the input spectrogram to generate predicted spectrograms. Differences between the estimated spectrograms and the ground truth ones are used as loss functions of the network.

which used for output generation. Finally, the two feature maps that passed the respective  $1 \times 1$  convolution and the input of the hourglass module is added together, and the resulting feature map is used as an input to the next hourglass module.

In the network used in this paper, input image firstly passes through initial convolutional layers that consist of a  $7 \times 7$  convolutional layer and four  $3 \times 3$  convolutional layers where the number of output feature maps for each layer is 64, 128, 128, 128, and 256 respectively. To make the output mask and the input spectrogram have the same size, we did not use the pooling operations in the initial convolutional layers before the hourglass module. The feature maps generated from the initial layers are fed to the first hourglass module. The proposed overall music source separation framework is depicted in Fig. 2.

### 3.2 Music Source Separation

As shown in Fig. 2, to apply the stacked hourglass network to music source separation, we aim to train the network to output soft masks for each music source given the magnitude spectrogram of the mixed source. Hence, the output dimension of the network is  $H \times W \times C$  where  $H$  and  $W$  are the height and width of the input spectrogram respectively, and  $C$  is the number of music sources to separate. The magnitude spectrogram of separated music source is obtained by multiplying the mask and the input spectrogram. Our framework is scalable in that it requires almost no additional operation as the number of sources increases.

The input for the network is the magnitude of spectrogram obtained from Short-Time Fourier Transform (STFT) with a window size of 1024 and a hop size of 256. The input source is downsampled to 8kHz to increase the duration of spectrograms in a batch and to speed up training. For each sample, magnitude spectrograms of mixed and

separated sources are generated, which are divided by the maximum value of the mixed spectrogram for data normalization. The spectrograms have 512 frequency bins and the width of the spectrogram depends on the duration of the music sources. For all the music sources, the width of the spectrogram is at least 64. Thus, we fix the size of an input spectrogram to  $512 \times 64$ . Hence, the size of the feature maps at the lowest resolution is  $32 \times 4$ . Starting time index is randomly chosen when the input batches are created.

Following [22], we designed the loss function as an  $L_{1,1}$  norm of the difference between the ground truth spectrogram and the estimated spectrogram. More concretely, given an input spectrogram  $\mathbf{X}$ ,  $i$ th ground truth music source  $\mathbf{Y}_i$ , and the generated mask for the  $i$ th source in the  $j$ th hourglass module  $\hat{\mathbf{M}}_{ij}$ , the loss for the  $i$ th source is defined as

$$\mathcal{J}(i, j) = \|\mathbf{Y}_i - \mathbf{X} \odot \hat{\mathbf{M}}_{ij}\|_{1,1}, \tag{1}$$

where  $\odot$  denotes element-wise multiplication of the matrix.  $L_{1,1}$  norm is calculated as the sum of absolute values of matrix elements. The loss function of the network becomes

$$\mathcal{J} = \sum_{i=1}^C \sum_{j=1}^D \mathcal{J}(i, j), \tag{2}$$

where  $D$  is the number of hourglass modules stacked in the network. We directly used the output of the last  $1 \times 1$  convolutional layer as the mask, which is different from [22] where they used the sigmoid activation to generate masks. While it is natural to use the sigmoid function to restrict the value of the mask to  $[0,1]$ , we empirically found that not applying the sigmoid function boosts the training speed and improves the performance. Since sigmoid activations vanish the gradient of the inputs that have large absolute values, they may diminish the effect of intermediate supervision.

We have stacked hourglass modules up to four and provide analysis of the effect of stacking multiple modules in Section 4. The network is trained using Adam optimizer [10] with a starting learning rate of  $10^{-4}$  and a batch size of 4. We trained the network for 15,000 and 150,000 iterations for MIR-1K dataset and DSD100 dataset respectively, and the learning rate is decreased to  $2 \times 10^{-5}$  when 80% of the training is finished. No data augmentation is applied during training. The training took 3 hours for MIR-1K dataset and 31 hours for DSD100 dataset using a single GPU when the biggest model is used. For the singing voice separation task,  $C$  is set to 2 which corresponds to *vocal* and *accompaniments*. For the music source separation task in DSD100 dataset,  $C = 4$  is used where each output mask corresponds to *drum*, *bass*, *vocal*, and *others*. While it can be advantageous in terms of performance to train a network for a single source individually, it is computationally expensive to train a deep CNN for each source. Therefore, we trained a single network for each task.

In the test phase, the magnitude spectrogram of the input source is cropped to network input size and fed to the network sequentially. The output of the last hourglass module is used for testing. We set the negative values of output masks to 0 in order to avoid negative magnitude values. The masks are multiplied by the normalized magnitude spectrogram of the test source and unnormalized to generate spectrograms of separated sources. We did not change the phase spectrogram of the input source, and it is combined with the estimated magnitude spectrogram to retrieve signals for separated sources via inverse STFT.

#### 4. EXPERIMENTS

We evaluated performance of the proposed method on MIR-1K and DSD100 datasets. For quantitative evaluation, we measured signal-to-distortion ratio (SDR), source-to-interference ratio (SIR), and source-to-artifacts ratio (SAR) based on BSS-EVAL metrics [31]. Normalized SDR (NSDR) [20] is also measured for the singing voice separation task which measures improvement between the mixture and the separated source. The values are obtained using mir-eval toolbox [21]. Global NSDR (GNSDR), global SIR (GSIR), and global SAR (GSAR) are calculated as a weighted mean of NSDR, SIR, and SAR respectively whose weights are length of the source. The separated sources generated from the network are upsampled to the original sampling rate of the dataset and compared with ground truth sources for all experiments.

##### 4.1 MIR-1K dataset

MIR-1K dataset is designed for singing voice separation research. It contains a thousand song clips extracted from 110 Chinese karaoke songs at a sampling rate of 16kHz. Following the previous works [6, 37], we used one male and one female (*abjones* and *amy*) as a training set which contains 175 clips in total. The remaining 825 clips are used for evaluation. For the baseline CNN, we trained the FCN that has U-Net [22]-like structure and evaluated its

Singing voice			
Method	GNSDR	GSIR	GSAR
MLRR [37]	3.85	5.63	10.70
DRNN [6]	7.45	13.08	9.68
ModGD [23]	7.50	13.73	9.45
U-Net [8]	7.43	11.79	10.42
SH-1stack	10.29	15.51	12.46
SH-2stack	10.45	15.89	12.49
SH-4stack	<b>10.51</b>	<b>16.01</b>	<b>12.53</b>

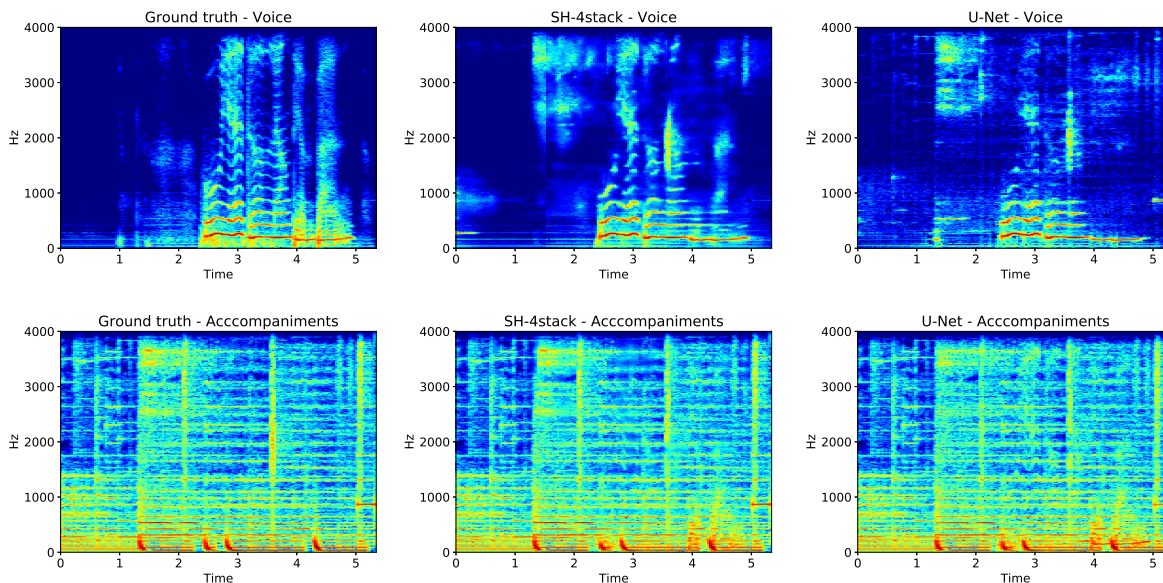
Accompaniments			
Method	GNSDR	GSIR	GSAR
MLRR [37]	4.19	7.80	8.22
U-Net [8]	7.45	11.43	10.41
SH-1stack	9.65	13.90	12.27
SH-2stack	9.64	13.69	<b>12.39</b>
SH-4stack	<b>9.88</b>	<b>14.24</b>	12.36

**Table 1.** Quantitative evaluation of singing voice separation on MIR-1K dataset.

performance. We followed the structure of [8], in which singing voice and accompaniments are trained on different networks. For the stacked hourglass networks, both singing voice and accompaniments are obtained from a single network.

The evaluation results on test sets are shown in Table 1. We trained the networks with varying number of stacked hourglass modules 1, 2, and 4. It is proven that our stacked hourglass network (SH) significantly outperforms existing methods in all evaluation criteria. Our method gains 3.01 dB in GNSDR, 2.28 dB in GSIR, and 1.83 dB in GSAR compared to the best results of the existing methods. It is also proven that the structure of the stacked hourglass module is more efficient and beneficial than U-Net [8] for music source separation. U-Net has 9.82 million parameters while single stack hourglass network has 8.99 million parameters considering only convolutional layers. Even with the absence of batch normalization, smaller number of parameters, and multi-source separation in a single network, the stacked hourglass network showed superior performance to U-Net. While the network with a single hourglass module shows outstanding source separation performance, even better results are provided when multiple hourglass modules are stacked. This indicates that SH network does not overfit even when the network gets deeper despite small amount of the training data. Our method provides good performance on separating both singing voice and accompaniments with a single forward step.

Qualitative results of our method and comparison with U-Net are shown in Fig. 3. The estimated log spectrograms of singing voice and accompaniments from SH-4stack and U-Net and the ground truth log spectrograms are provided. It can be seen that our method captures fine details and harmonics compared to the U-Net. The voice spectrogram from U-Net has more artifacts in the time slot of 0~1 and 4~5 compared to the result of SH-4stack. On the other



**Figure 3.** Qualitative comparison of our method (SH-4stack) and U-Net for singing voice and accompaniments separation on *annar\_3\_05* in MIR-1K dataset. Ground truth and estimated spectrograms are displayed in a log-scale. Our method is superior in capturing fine details compared to U-Net.

hand, harmonics from voice signals can be clearly seen in the spectrogram of SH-4stack. For accompaniments spectrogram, it is observed that U-Net contains voice signals around the time slot of 3.

### 4.2 DSD100 dataset

DSD100 dataset consists of 100 songs that are divided into 50 training sets and 50 test sets. For each song, four different music sources, bass, drums, vocals, and other as well as their mixtures are provided. The sources are stereophonic sound with a sampling rate of 44.1kHz. We converted all sources to monophonic and performed single channel source separation using stacked hourglass networks. We used a 4-stacked hourglass network (SH-4stack) for the experiments.

The performance of music source separation using stacked hourglass network is provided in Table 2. We measured SDR of the separated sources for all test songs and report median values for comparison with existing methods. The methods that use single channel inputs are compared to our method. While the stacked hourglass network gives second-best performance following the state-of-the-art methods [26] for drums and vocals, it shows poor performance for separating bass and other. This is mainly due to the similarity between bass and guitar sound in other sources, which confuses the network especially when trained together in a single network. Since the losses for all sources are summed up with equal weights, the network tends to be trained to improve the separation performance of vocal and drum, which is easier than separating bass and other sources.

Next, we trained the stacked hourglass network for a

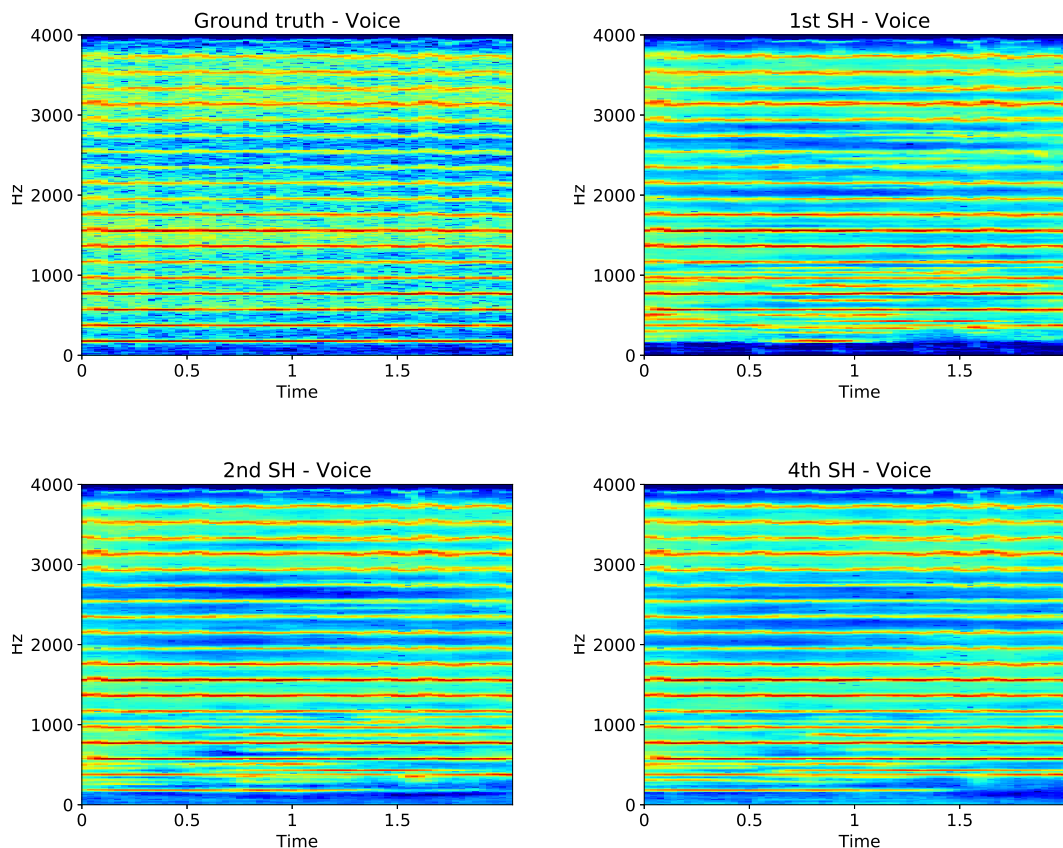
Method	Bass	Drums	Other	Vocals
dNMF [36]	0.91	1.87	2.43	2.56
DeepNMF [11]	1.88	2.11	2.64	2.75
BLEND [28]	2.76	3.93	3.37	5.13
MM-DenseNet [26]	<b>3.91</b>	<b>5.37</b>	<b>3.81</b>	<b>6.00</b>
SH-4stack	1.77	4.11	2.36	5.16

**Table 2.** Median SDR values for music source separation on DSD100 dataset.

Method	Vocals	Accompaniments
DeepNMF [11]	2.75	8.90
wRPCA [9]	3.92	9.45
NUG [19]	4.55	10.29
BLEND [28]	5.23	11.70
MM-DenseNet [26]	<b>6.00</b>	12.10
SH-4stack	5.45	<b>12.14</b>

**Table 3.** Median SDR values for singing voice separation on DSD100 dataset.

singing voice separation task. The three sources except vocals are mixed together to form accompaniments source. The median SDR values for each source are reported in Table 3. Our method achieved best result for accompaniments separation and second-best for vocal separation. Separation performance of vocals is improved compared to the music source separation setting. It can be inferred that the stacked hourglass network provides better results as number of sources are smaller and the separating sources are more distinguishable from each other.



**Figure 4.** Examples showing the effectiveness of stacking multiple hourglass modules. Ground truth and estimated spectrograms of the part of the song *Schoolboy Fascination* in DSD100 dataset are shown. SDR values of the source generated from the spectrograms obtained from first, second, fourth hourglass module are 10.90, 12.50, 13.30 respectively. Especially, it is observed that the estimated spectrogram captures fine details of spectrogram at low frequency range (0~500 Hz) as more hourglass modules are stacked.

Lastly, we investigate how the stacked hourglass network improves the output masks as they pass through the hourglass modules within the network. The example illustrated in Fig. 4 shows the estimated voice spectrogram of first, second, and fourth hourglass module with the ground truth spectrogram from one of the test sets of DSD 100 dataset. It is observed that the estimated spectrogram becomes more similar to the ground truth as it is generated from a deeper side of the network. In the result of the fourth hourglass module, spectrograms at low frequency are clearly recovered compared to the result of the first hourglass module. The artifacts in the range of 2000~3000 Hz are also removed. Although it is hard to recognize the difference in the spectrogram image, the difference of SDR between the source estimated from the first hourglass module and the last hourglass module is about 2.4dB which is a significant performance gain.

## 5. CONCLUSION

In this paper, we proposed music source separation algorithm using stacked hourglass networks. The network suc-

cessfully captures features at both coarse and fine resolution, and it produces masks that are applied to the input spectrograms. Multiple hourglass modules refines the estimation results and outputs the better results. Experimental results has proven the effectiveness of the proposed framework for music source separation. We implemented the framework in its simplest form, and there is a lot of room for performance improvements including data augmentation, regularization of CNNs, and ensemble learning of multiple models. Designing a loss function that considers correlation of different sources may further improves the performance.

## 6. ACKNOWLEDGEMENT

This work was supported by Next-Generation Information Computing Development Program through the National Research Foundation of Korea (2017M3C4A7077582).

## 7. REFERENCES

- [1] Pritish Chandna, Marius Miron, Jordi Janer, and Emilia Gómez. Monoaural audio source separation using deep

- convolutional neural networks. In *International Conference on Latent Variable Analysis and Signal Separation*, pages 258–266. Springer, 2017.
- [2] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2016.
- [3] Emad M Grais, Hagen Wierstorf, Dominic Ward, and Mark D Plumbley. Multi-resolution fully convolutional neural networks for monaural audio source separation. *arXiv preprint arXiv:1710.11473*, 2017.
- [4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [5] C. L. Hsu and J. S. R. Jang. On the improvement of singing voice separation for monaural recordings using the mir-1k dataset. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(2):310–319, Feb 2010.
- [6] Po-Sen Huang, Minje Kim, Mark Hasegawa-Johnson, and Paris Smaragdis. Joint optimization of masks and deep recurrent neural networks for monaural source separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(12):2136–2147, 2015.
- [7] Forrest Iandola, Matt Moskewicz, Sergey Karayev, Ross Girshick, Trevor Darrell, and Kurt Keutzer. Densenet: Implementing efficient convnet descriptor pyramids. *arXiv preprint arXiv:1404.1869*, 2014.
- [8] Andreas Jansson, Eric Humphrey, Nicola Montecchio, Rachel Bittner, Aparna Kumar, and Tillman Weyde. Singing voice separation with deep u-net convolutional networks. *18th International Society for Music Information Retrieval Conference, Suzhou, China*, 2017.
- [9] Il-Young Jeong and Kyogu Lee. Singing voice separation using rpca with weighted  $L_{1/2}$ -norm. In *International Conference on Latent Variable Analysis and Signal Separation*, pages 553–562. Springer, 2017.
- [10] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [11] Jonathan Le Roux, John R Hershey, and Felix Weninger. Deep nmf for speech separation. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 66–70. IEEE, 2015.
- [12] Daniel D Lee and H Sebastian Seung. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, pages 556–562, 2001.
- [13] Yuan-Shan Lee, Chien-Yao Wang, Shu-Fan Wang, Jia-Ching Wang, and Chung-Hsien Wu. Fully complex deep neural network for phase-incorporating monaural source separation. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 281–285. IEEE, 2017.
- [14] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [15] Stylianos Ioannis Mimilakis, Konstantinos Drossos, Tuomas Virtanen, and Gerald Schuller. A recurrent encoder-decoder approach with skip-filtering connections for monaural singing voice separation. *CoRR*, abs/1709.00611, 2017.
- [16] Marius Miron, Jordi Janer, and Emilia Gómez. Monaural score-informed source separation for classical music using convolutional neural networks. In *18th International Society for Music Information Retrieval Conference, Suzhou, China*, 2017.
- [17] Alejandro Newell, Zhiao Huang, and Jia Deng. Associative embedding: End-to-end learning for joint detection and grouping. In *Advances in Neural Information Processing Systems*, pages 2274–2284, 2017.
- [18] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision*, pages 483–499. Springer, 2016.
- [19] Aditya Arie Nugraha, Antoine Liutkus, and Emmanuel Vincent. Multichannel music separation with deep neural networks. In *Signal Processing Conference (EUSIPCO), 2016 24th European*, pages 1748–1752. IEEE, 2016.
- [20] Alexey Ozerov, Pierrick Philippe, Frédéric Bimbot, and Rmi Gribonval. Adaptation of bayesian models for single-channel source separation and its application to voice/music separation in popular songs. *IEEE Transactions on Audio, Speech, and Language Processing*, 15(5):1564–1578, 2007.
- [21] Colin Raffel, Brian McFee, Eric J Humphrey, Justin Salamon, Oriol Nieto, Dawen Liang, Daniel PW Ellis, and C Colin Raffel. mir\_eval: A transparent implementation of common mir metrics. In *In Proceedings of the 15th International Society for Music Information Retrieval Conference, ISMIR*. Citeseer, 2014.
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

- [23] Jilt Sebastian and Hema A Murthy. Group delay based music source separation using deep recurrent neural networks. In *Signal Processing and Communications (SPCOM), 2016 International Conference on*, pages 1–5. IEEE, 2016.
- [24] Andrew JR Simpson, Gerard Roma, and Mark D Plumbley. Deep karaoke: Extracting vocals from musical mixtures using a convolutional deep neural network. In *International Conference on Latent Variable Analysis and Signal Separation*, pages 429–436. Springer, 2015.
- [25] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, volume 4, page 12, 2017.
- [26] Naoya Takahashi and Yuki Mitsufuji. Multi-scale multi-band densenets for audio source separation. In *Applications of Signal Processing to Audio and Acoustics (WASPAA), 2017 IEEE Workshop on*, pages 21–25. IEEE, 2017.
- [27] Stefan Uhlich, Franck Giron, and Yuki Mitsufuji. Deep neural network based instrument extraction from music. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 2135–2139. IEEE, 2015.
- [28] Stefan Uhlich, Marcello Porcu, Franck Giron, Michael Enekl, Thomas Kemp, Naoya Takahashi, and Yuki Mitsufuji. Improving music source separation based on deep neural networks through data augmentation and network blending. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 261–265. IEEE, 2017.
- [29] Shankar Vembu and Stephan Baumann. Separation of vocals from polyphonic audio recordings. In *ISMIR*, pages 337–344. Citeseer, 2005.
- [30] Emmanuel Vincent, Shoko Araki, Fabian Theis, Guido Nolte, Pau Bofill, Hiroshi Sawada, Alexey Ozerov, Vikram Gowreesunker, Dominik Lutter, and Ngoc QK Duong. The signal separation evaluation campaign (2007–2010): Achievements and remaining challenges. *Signal Processing*, 92(8):1928–1936, 2012.
- [31] Emmanuel Vincent, Rémi Gribonval, and Cédric Févotte. Performance measurement in blind audio source separation. *IEEE transactions on audio, speech, and language processing*, 14(4):1462–1469, 2006.
- [32] Tuomas Virtanen. Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria. *IEEE transactions on audio, speech, and language processing*, 15(3):1066–1074, 2007.
- [33] Guan-Xiang Wang, Chung-Chien Hsu, and Jen-Tzung Chien. Discriminative deep recurrent neural networks for monaural speech separation. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 2544–2548. IEEE, 2016.
- [34] Yuxuan Wang, Arun Narayanan, and DeLiang Wang. On training targets for supervised speech separation. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 22(12):1849–1858, 2014.
- [35] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4724–4732, 2016.
- [36] Felix Weninger, Jonathan Le Roux, John R Hershey, and Shinji Watanabe. Discriminative nmf and its application to single-channel source separation. In *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [37] Yi-Hsuan Yang. Low-rank representation of both singing voice and music accompaniment via learned dictionaries. In *ISMIR*, pages 427–432, 2013.
- [38] Xiu Zhang, Wei Li, and Bilei Zhu. Latent time-frequency component analysis: A novel pitch-based approach for singing voice separation. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 131–135. IEEE, 2015.