

END-TO-END SOURCE SEPARATION WITH ADAPTIVE FRONT-ENDS

Shrikant Venkataramani, Jonah Casebeer

Paris Smaragdīs

University of Illinois at Urbana Champaign
svnktrm2, jonahmc2@illinois.edu

University of Illinois at Urbana Champaign
Adobe Research

ABSTRACT

Source separation and other audio applications have traditionally relied on the use of short-time Fourier transforms as a front-end frequency domain representation step. The unavailability of a neural network equivalent to forward and inverse transforms hinders the implementation of end-to-end learning systems for these applications. We develop an auto-encoder neural network that can act as an equivalent to short-time front-end transforms. We demonstrate the ability of the network to learn optimal, real-valued basis functions directly from the raw waveform of a signal and further show how it can be used as an adaptive front-end for supervised source separation. In terms of separation performance, these transforms significantly outperform their Fourier counterparts. Finally, we also propose and interpret a novel source to distortion ratio based cost function for end-to-end source separation.

Index Terms— Auto-encoders, adaptive transforms, source separation, deep learning

1. INTRODUCTION

Several neural network (NN) architectures and methods have been proposed for supervised single-channel source separation [1, 2, 3, 4]. These approaches can be grouped together into a common two-step workflow as follows. The first step is to transform the time domain signals into a suitable time-frequency (TF) representation using short-time Fourier transforms (STFTs). These short time spectra are subsequently divided into their magnitude and phase components. The actual separation procedure takes place in the second step of the workflow which often operates on the extracted magnitude components. Common approaches include neural networks which given the noisy magnitudes, either predict a noiseless magnitude spectrum [5, 6, 7], or some type of a masking function [8, 9]. Figure 1 (a) shows the block diagram of such a system using the STFT as a front-end transform. Although they produce very good results, these NN approaches suffer from a couple of drawbacks. First, by restricting the processing to only magnitudes, they do not take full advantage of the information contained in the input signals. Additionally, there is no guarantee that the STFT (or whichever front-end one uses) would be optimal for the task at hand. In this paper, we investigate the use of adaptive front-end transforms for supervised source separation. Using these adaptive front-ends for forward and inverse transformations enables the development of end-to-end learning systems for supervised source separation and potentially for other related NN models that rely on fixed transforms. In section 2, we consider the use of the real and imaginary parts of the DFT as a front-end transform to develop the necessary intuition for using real-valued front-ends. We then develop a neural network equivalent




Fig. 1. Block diagram of generalized NN based source separation system using (a) STFT front-end (top) and (b) the proposed adaptive front-end transform (bottom).

in section 3 and show how it can be used as an adaptive front-end for end-to-end source separation. Our experiments and results are discussed in section 4 and we conclude in section 5.

2. USING REAL-VALUED TRANSFORMS

Our fundamental approach towards learning an adaptive front-end is to replace the front-end transform with a regular convolutional layer. To develop the architecture of the end-to-end separation network, we will first replace the DFT by a real-valued transform. This will allow us to develop an appropriate formulation before we move to an adaptive transform. Given a time domain sequence x , the short-time transform operation of x can be expressed as a generalized equation given by,

$$\mathbf{X}_{nk} = \sum_{t=0}^{N-1} x(nh + t) \cdot w(t) \cdot b(k, t) \quad (1)$$

Here, \mathbf{X}_{nk} represents the coefficient corresponding to the k^{th} component in the n^{th} frame, N represents the size of a window function w and h represents the hop size of the short-time transform.

This work was supported by NSF grant 1453104.

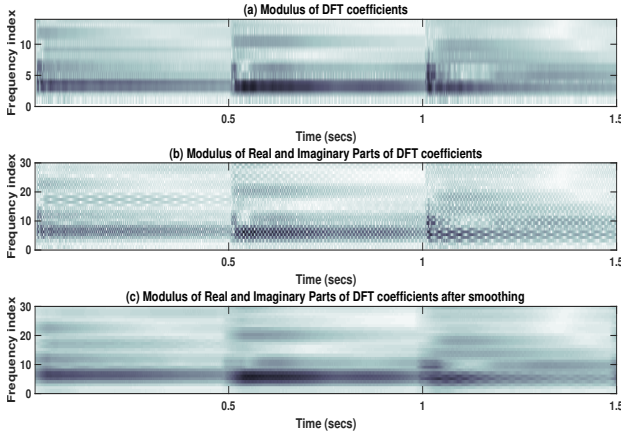


Fig. 2. (a) Modulus of DFT coefficients (first 15 coefficients) for a sequence of piano notes. (b) Modulus of equivalent real and imaginary parts of DFT coefficients. The unsmoothed coefficients oscillate excessively and need to be averaged across time to produce what we would expect as a magnitude spectrum. (c) Modulus of real and imaginary parts of DFT coefficients after smoothing.

The functions $b(k, t)$ form the basis functions of the transformation.

To learn a real-valued representation of the audio signal, we replace the magnitude and phase components of the DFT by its equivalent real and imaginary parts. Thus, the values of $b(k, t)$ are the elements of a matrix obtained by vertically stacking the sine and cosine terms of the DFT bases. As shown in figure 2, the resulting spectral energies do not maintain locality [10] and exhibit a modulation that could be dependent on the frequency, the window size and the hop size parameters. Thus, we need to apply a suitable averaging (smoothing) operation across time. This can be achieved by convolving the absolute values of the resulting coefficients with a suitable averaging function across the time axis given as,

$$\mathbf{M} = |\mathbf{X}| * s \quad (2)$$

Here, $|\cdot|$ represents the element-wise modulus operator, s represents the bank of smoothing filters and $*$ denotes the one-dimensional convolution operation applied only along the time axis. The matrix \mathbf{M} obtained after smoothing can be interpreted as the magnitude spectrogram of the signal in this representation. The variations in the coefficients that are not captured by this magnitude spectrogram \mathbf{M} can be captured in a new matrix given by $\mathbf{P} = \frac{\mathbf{X}}{\mathbf{M}}$ where, the division is also element-wise. This can be interpreted as the corresponding phase component of the sequence x . We can also interpret these two quantities as the results of a demodulation operation in which we extract a smooth amplitude modulator, which modulates by a faster moving carrier that encodes more of the details. Using this approach we can easily match the performance of the STFT front-end, while only performing real-valued computations. Thus, we will refer to \mathbf{M} and \mathbf{P} matrices as “Modulation spectrogram” and “Carrier component” in order to avoid any confusion with the DFT based magnitude spectrogram and phase components. However, the use of a fixed front-end transform continues to pose some challenges. Short-time transforms need to be optimized with respect to the window size, window shape and hop size parameters. In the case of real

transforms, the transformation must be followed by a smoothing operation that depends on the window size, the hop size and coefficient frequency. Thus, we also need to optimize over suitable smoothing function shapes and durations. As described in section 3, we can interpret each step of the forward short-time transform as a neural network. In doing so, we can automatically learn adaptive transforms and smoothing functions directly from the waveform of the audio signal, thereby bypassing the aforementioned issues.

3. AUTO-ENCODER TRANSFORMS FOR SOURCE SEPARATION

Despite recent advances in neural networks, the STFT continues to be the transform of choice for audio applications [1, 2, 3, 4]. Recently, Sainath et.al [11], and Dieleman and Schrauwen [12] have proposed the use of a convolutional layer as an alternative to front-end STFTs. This approach has also been used for source separation in [9]. In this section, we expand upon the premise and develop a real valued, convolutional auto-encoder transform (AET) that can be used as an alternative to front-end short-time transforms. The encoder part of the auto-encoder (AE) acts as the analysis filter bank and produces a spectrogram-equivalent of the input. The decoder performs the synthesis step used to recover the time-domain signal.

3.1. Analysis Encoder

Assuming a unit sample hop size, we can interpret (1) as a filtering operation,

$$\mathbf{X}_{nk} = \sum_{t=0}^{N-1} x(n+t) \cdot \mathbf{F}(k, t) \quad (3)$$

Thus, we may replace the front-end transform by a convolutional neural network (CNN) such that the k^{th} filter of the CNN represents the k^{th} row of \mathbf{F} . The output of the CNN gives a feature space representation of the input signal with a unit hop size. As described in section 2, the transformation stage should be followed by a smoothing operation. This smoothing stage given by (2) can also be interpreted as a CNN applied on $|\mathbf{X}|$. However, since there are no non-negativity constraints applied on the averaging filters, the elements of the smoothed modulation spectrogram \mathbf{M} can potentially assume negative values. To avoid this solution, we include a non-linearity to the convolutional smoothing layer. The non-linearity $g : \mathbb{R} \rightarrow \mathbb{R}^+$ is a mapping from the space of real numbers to the space of positive real numbers. In this paper, we use a softplus non-linearity for this step. As before, the output of this layer \mathbf{M} can be interpreted as the modulation spectrogram of the input signal and $\mathbf{P} = \frac{\mathbf{X}}{\mathbf{M}}$ can be interpreted as the corresponding carrier component. As before, we expect the carrier component captures the variations in the coefficients which cannot be modeled in the modulation spectrogram. In order to use a more economical subsampled representation, we can apply a max-pooling operation [13] that replaces a pool of h frames with a single frame containing the maximum value of each coefficient over the pool. Note that all the convolution and pooling operations are one-dimensional i.e., the operations are applied across the time-axis only. In addition, these operations are independently applied on each filter output of the front-end CNN.

3.2. Synthesis Decoder

Given the modulation spectrograms and the carrier components obtained using the AET, the next step is to synthesize the signal back

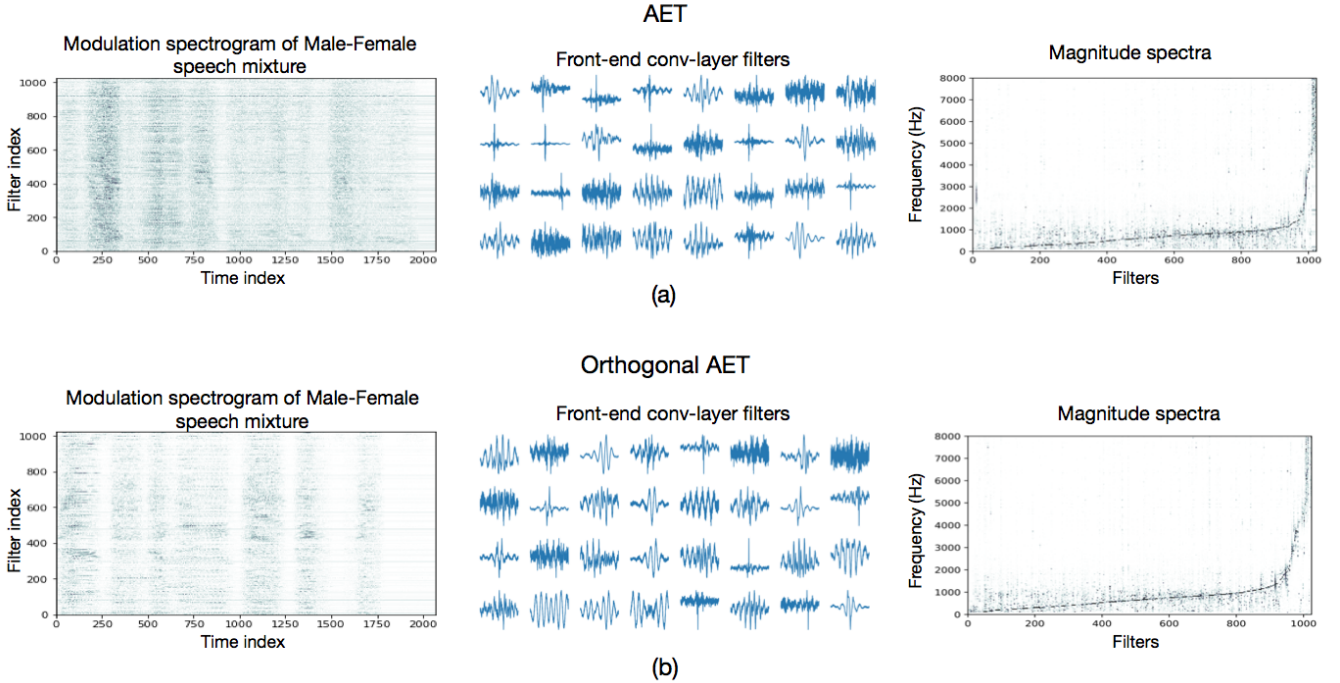


Fig. 3. (a) A sample modulation spectrogram obtained for a speech mixture containing a male and female speaker, front-end convolutional layer filters, and their normalized magnitude spectra using the AET (top) (b) A sample modulation spectrogram obtained for a speech mixture containing a male and female speaker, front-end convolutional layer filters, and their normalized magnitude spectra using the orthogonal-AET (bottom). The orthogonal-AET uses a transposed version of the analysis filters for the synthesis convolutional layer. The filters are ordered according to their dominant frequency component (from low to high). In the middle subplots, we show a subset of the first 32 filters.

into the time domain. This can be achieved by inverting the operations performed by the analysis encoder while computing the forward transform. The first step of the synthesis procedure is to undo the effect of the lossy pooling layer. We use an upsampling operator by inserting as many zeros between the frames as the pooling duration as proposed by Zieler et al., [14]. The unpooled magnitude spectrogram is then multiplied by the phase using an element-wise multiplication to give an approximation $\hat{\mathbf{X}}$ to the matrix \mathbf{X} . We invert the operation of the first transform layer by a convolutional layer to implement the deconvolution operation, using a synthesis convolutional layer. This convolutional layer thus, performs the interpolation between the samples. Essentially, the output of the analysis encoder gives the weights of each basis function in representing the time domain signal. The synthesis layer reconstructs each component by adding multiple shifted versions of the basis functions at appropriate locations in time. This inversion procedure is equivalent to an overlap-and-add strategy applied separately for each basis of the transform, followed by an overall summation step [15]. The weights (filters) of the first convolutional layer give the AET basis functions (see figure 3).

3.3. Examining the learned bases

Having developed the convolutional auto-encoder for AETs, we can now examine and understand the nature of the basis functions obtained. We note that the architecture of our end-to-end network is a natural extension to the DFT separation procedure. We plot the bases

and their corresponding normalized magnitude spectra in figure (3) for AET (top) and orthogonal-AET (bottom) transforms. In the case of orthogonal-AET, the synthesis convolutional layer filters are held to be transposed versions of the front-end layer filters. Thus, the inverse transform is a transpose of the forward transform. In figure 3, the middle figures are the filters obtained from the front-end convolutional layer that operates on the input mixture waveform. The complete network architecture and training-data for obtaining these plots are described in section 4.1. We rank the filters according to their dominant frequency component. Then, we use a 1024-point Fourier transform to compute the magnitude spectra of the filters. The middle figures show the first 32 low-frequency filters obtained after the sorting step. The plots on the right show the corresponding filter magnitude spectra. From the magnitude spectra, it is clear that the filters are frequency selective even though they are noisy and consist of multiple frequencies. The filters are concentrated at the lower frequencies and spaced out at the higher frequencies. The left figures show the output of the front-end layer for a sample mixture input waveform, with respect to the corresponding transform bases. These observations hold for the AET and the orthogonal-AET. In other words, we see that the adaptive front-ends learn a representation that is tailored to the input waveform.

3.4. End-to-end Source Separation

Figure 1(b) shows the application of AET analysis and synthesis layers for end-to-end supervised source separation. The forward and

inverse transforms can be directly replaced by the AET analysis and synthesis networks in a straightforward manner. We train the network by giving the mixture waveform at the input and minimize a suitable distance metric between the network output and the clean waveform of the source. Thus, the network learns to estimate the contribution of the source given the raw waveform of the mixture. Since the basis and smoothing functions are automatically learned, it is reasonable to expect the network to learn optimal, task specific basis and smoothing functions. The advantages of interpreting the forward and inverse transforms as a neural network now begins to come through. We can propose interesting extensions to these adaptive front-ends by exploiting the available diversity of neural networks. For example, we can propose recurrent auto-encoder alternatives as in [16] or multilayer or recurrent extensions to adaptive front-ends. We can also experiment with better pooling strategies for the adaptive front-end. These generalizations to adaptive front-ends are not easily explored with fixed front-end transforms.

4. EXPERIMENTS

We now present some experiments aimed at comparing the effect of different front-ends for a supervised source separation task. We evaluate the separation performance for three types of front-ends: STFT, AET and orthogonal-AET. We compare the results based on the BSS_EVAL metrics. [17].

4.1. Experimental setup

For training the neural networks, we randomly selected 10 male-female speaker pairs from the TIMIT database [18]. In the database, each speaker has a total of 10 recorded sentences. For each pair, we mix the sentences at 0 dB. This gives 10 mixture sentences per pair and a total of 100 mixture sentences overall. For each pair, we train on 8 sentences and test on the remaining two. Thus, the network is trained on 80 mixture sentences and evaluated on the remaining 20 mixture sentences. For training the NN we use a batch size of 16 and a dropout-rate of 0.2. The separation network consisted of a cascade of 3 dense layers with 512 hidden units each, each followed by a softplus non-linearity, which is the architecture used in [5]. As seen in figure 1, the STFT and AET magnitude spectrograms were given as inputs to the separator network. The STFT was computed at a window-size of 1024 samples at a hop of 16 samples, using a Hann window. For the STFT front-end, the separation results were inverted into the time domain using the inverse STFT operation and the mixture phase. To have a fair comparison with adaptive front-ends, the CNN filters were chosen to have a width of 1024, a stride of 16 samples was selected for the pooling operation. The smoothing layer was selected to have a length of 5 frames. We used two different cost functions for this task. First we used the mean-squared error (MSE) between the target waveform and the target estimate. Second, we used a cost function that directly optimizes the signal to distortion ratio (SDR) instead [19]. To do the latter, we note that for a reference signal y and an output x we would maximize:

$$\begin{aligned} \max \text{SDR}(x, y) &= \max \frac{\langle xy \rangle^2}{\langle yy \rangle \langle xx \rangle - \langle xy \rangle^2} \equiv \\ \min \frac{\langle yy \rangle \langle xx \rangle - \langle xy \rangle^2}{\langle xy \rangle^2} &= \min \frac{\langle yy \rangle \langle xx \rangle}{\langle xy \rangle^2} - \frac{\langle xy \rangle^2}{\langle xy \rangle^2} \propto \min \frac{\langle xx \rangle}{\langle xy \rangle^2} \end{aligned} \quad (4)$$

Here, the inner product $\langle yy \rangle$ is neglected to simplify the cost-function since it is a constant with respect to the output of the net-




Fig. 4. Comparison of source separation performance on 20 speech on speech mixtures in terms of BSS_EVAL parameters. We compare the separation performance for multiple front end transforms viz., STFT, AET and orthogonal AET. The dashed line in the centre indicates the median value and the dotted lines above and below indicate the interquartile range. We see that opting for an adaptive front-end results in a significant improvement in source separation performance over STFT front-ends. Comparing the cost-functions we see that SDR (left) is a more appropriate cost-function to MSE (right) for end-to-end source separation.

work x . Thus, maximizing the SDR is equivalent to maximizing the correlation between x and y , while producing a minimum energy output.

4.2. Results and Discussion

The corresponding violin plots that show the distribution of the BSS_EVAL metrics from our experiments are shown in figure 4. We see that the use of AETs improves the separation performance in terms of all metrics compared to an STFT front-end. We additionally see that when using the orthogonal AET we obtain additional performance gains, overall in the range of 2dB for SDR, 5dB in SIR and 3dB in SAR. One possible reason for the increased performance of the orthogonal AET could be the reduction in the number of trainable parameters caused by forcing the synthesis transform to be the transpose of the analysis transform, which in turn reduces the possibility of over-fitting to the training data. The above trends appear consistent for both the cost-functions considered. Additionally we can compare the use of the two cost-functions for our networks. We see that the use of SDR as a cost function (expectedly) results in a significant improvement over using MSE. This is observed for all the front-end options considered in this paper. We also note that the use of MSE increases the variance of the separation results, whereas the SDR is more consistent. We thus conclude that the SDR is a better choice of cost-function for end-to-end source separation.

5. CONCLUSION AND FUTURE WORK

In this paper, we developed and investigated a convolutional auto-encoder based front-end transform that can be used as an alterna-

tive to using STFT front-ends. The adaptive front-end comprises a cascade of three layers viz., a convolutional front-end transform layer, a convolutional smoothing layer and a pooling layer. We have shown that AETs are capable of automatically learning adaptive basis functions and discovering data-specific frequency domain representations directly from the raw waveform of the data. The use of AETs significantly improves the separation performance over fixed front-ends and also enables the development of end-to-end source separation. Finally, we have also demonstrated that the SDR is a superior alternative as a cost-function for end-to-end source separation systems.

6. REFERENCES

- [1] P. Smaragdis and S. Venkataramani, "A neural network alternative to non-negative audio models," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 86–90.
- [2] P. S. Huang, M. Kim, M. Hasegawa-Johnson, and P. Smaragdis, "Deep learning for monaural speech separation," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 1562–1566.
- [3] P. Chandna, M. Miron, J. Janer, and E. Gómez, "Monoaural audio source separation using deep convolutional neural networks," in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2017, pp. 258–266.
- [4] S. Uhlich, M. Porcu, F. Giron, M. Enekl, T. Kemp, N. Takahashi, and Y. Mitsufuji, "Improving music source separation based on deep neural networks through data augmentation and network blending," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017.
- [5] M. Kim and P. Smaragdis, "Adaptive denoising autoencoders: A fine-tuning scheme to learn from test mixtures," in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2015, pp. 100–107.
- [6] E. M. Grais, M. U. Sen, and H. Erdogan, "Deep neural networks for single channel source separation," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 3734–3738.
- [7] E. M. Grais and M. D. Plumbley, "Single channel audio source separation using convolutional denoising autoencoders," *arXiv preprint arXiv:1703.08019*, 2017.
- [8] X.-L. Zhang and D. Wang, "A deep ensemble learning method for monaural speech separation," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 24, no. 5, pp. 967–977, 2016.
- [9] Y. Luo and N. Mesgarani, "Tasnet: Time-domain audio separation network for real-time single-channel speech separation," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. IEEE, 2018.
- [10] O. Abdel-Hamid, A.-r. Mohamed, H. Jiang, L. Deng, G. Penn, and D. Yu, "Convolutional neural networks for speech recognition," *IEEE/ACM Transactions on audio, speech, and language processing*, vol. 22, no. 10, pp. 1533–1545, 2014.
- [11] T. N. Sainath, R. J. Weiss, A. W. Senior, K. W. Wilson, and O. Vinyals, "Learning the speech front-end with raw waveform cldnns," in *INTERSPEECH*. ISCA, 2015, pp. 1–5.
- [12] S. Dieleman and B. Schrauwen, "End-to-end learning for music audio," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 6964–6968.
- [13] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," *arXiv preprint arXiv:1412.6806*, 2014.
- [14] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *European conference on computer vision*. Springer, 2014, pp. 818–833.
- [15] J. O. Smith, *Spectral Audio Signal Processing*. <http://ccrma.stanford.edu/jos/saspl/>.
- [16] S. Venkataramani, C. Subakan, and P. Smaragdis, "Neural network alternatives to convolutive models for source separation," in *IEEE Workshop on Machine Learning for Signal Processing (MLSP)*, 2017.
- [17] C. Févotte, R. Gribonval, and E. Vincent, "Bss_eval toolbox user guide—revision 2.0," 2005.
- [18] J. S. Garofolo, L. F. Lamel, J. G. F. William M Fisher, D. S. Pallett, N. L. Dahlgren, and V. Zue, "Timit acoustic phonetic continuous speech corpus," Philadelphia, 1993.
- [19] S. Venkataramani, J. Casebeer, and P. Smaragdis, "Adaptive front-ends for end-to-end source separation." [Online]. Available: http://media.aau.dk/smc/wp-content/uploads/2017/12/ML4AudioNIPS17_paper_39.pdf