SEMI-BLIND SOURCE SEPARATION WITH MULTICHANNEL VARIATIONAL AUTOENCODER

Hirokazu Kameoka[1], Li Li[2], Shota Inoue[2], Shoji Makino[2]

[1] NTT Communication Science Laboratories, NTT Corporation, Japan
[2] University of Tsukuba, Japan

ABSTRACT

This paper proposes a multichannel source separation technique called the multichannel variational autoencoder (MVAE) method, which uses a conditional VAE (CVAE) to model and estimate the power spectrograms of the sources in a mixture. By training the CVAE using the spectrograms of training examples with source-class labels, we can use the trained decoder distribution as a universal generative model capable of generating spectrograms conditioned on a specified class label. By treating the latent space variables and the class label as the unknown parameters of this generative model, we can develop a convergence-guaranteed semi-blind source separation algorithm that consists of iteratively estimating the power spectrograms of the underlying sources as well as the separation matrices. In experimental evaluations, our MVAE produced better separation performance than a baseline method.

Index Terms— Blind source separation, multichannel non-negative matrix factorization, variational autoencoders (VAEs)

1. INTRODUCTION

Blind source separation (BSS) is a technique for separating out individual source signals from microphone array inputs when the transfer characteristics between the sources and microphones are unknown. The frequency-domain BSS approach provides the flexibility of allowing us to utilize various models for the time-frequency representations of source signals and/or array responses. For example, independent vector analysis (IVA) [1, 2] allows us to efficiently solve frequency-wise source separation and permutation alignment in a joint manner by assuming that the magnitudes of the frequency components originating from the same source tend to vary coherently over time.

With a different approach, multichannel extensions of nonnegative matrix factorization (NMF) have attracted a lot of attention in recent years [3–7]. NMF was originally applied to music transcription and monaural source separation tasks [8, 9]. The idea is to approximate the power (or magnitude) spectrogram of a mixture signal, interpreted as a non-negative matrix, as the product of two non-negative matrices. This amounts to assuming that the power spectrum of a mixture signal observed at each time frame can be approximated by the linear sum of a limited number of basis spectra scaled by time-varying amplitudes. Multichannel NMF (MNMF) is an extension of this approach to a multichannel case to allow the use of spatial information as an additional clue to separation. It can also be viewed as an extension of frequency-domain BSS that allows the use of spectral templates as a clue for jointly solving frequency-wise source separation and permutation alignment.

The original MNMF [3] was formulated under a general problem setting where sources can outnumber microphones and a determined version of MNMF was subsequently proposed in [4]. While the determined version is applicable only to determined cases, it allows the implementation of a significantly faster algorithm than the general version. The determined MNMF framework was later called "independent low-rank matrix analysis (ILRMA)" [7]. In [6], the theoretical relation of MNMF to IVA was discussed, which has naturally allowed for the incorporation of the fast update rule of the separation matrix developed for IVA, called "iterative projection (IP)" [10], into the parameter optimization process in ILRMA. It has been shown that this has contributed not only to further accelerating the entire optimization process but also to improving the separation performance. While ILRMA is notable in that the optimization algorithm is guaranteed to converge, it can fail to work for sources with spectrograms that do not comply with the NMF model.

As an alternative to the NMF model, some attempts have recently been made to use deep neural networks (DNNs) for modeling the spectrograms of sources for multichannel source separation [11, 12]. The idea is to replace the process for estimating the power spectra of source signals in a source separation algorithm with the forward computations of pre-trained DNNs. This can be viewed as a process of refining the estimates of the power spectra of the source signals at each iteration of the algorithm. While this approach is particularly appealing in that it can take advantage of the strong representation power of DNNs for estimating the power spectra of source signals, one weakness is that the convergence of an algorithm devised in this way will not be guaranteed.

To address the drawbacks of the methods mentioned above, this paper proposes a multichannel source separation

method using variational autoencoders (VAEs) [13, 14] for source spectrogram modeling. We call our approach the "multichannel VAE (MVAE)" method.

2. PROBLEM FORMULATION

We consider a situation where I source signals are captured by I microphones. Let $x_i(f, n)$ and $s_j(f, n)$ be the short-time Fourier transform (STFT) coefficients of the signal observed at the i-th microphone and the j-th source signal, where f and n are the frequency and time indices, respectively. We denote the vectors containing $x_1(f, n), \ldots, x_I(f, n)$ and $s_1(f, n), \ldots, s_I(f, n)$ by

$$\mathbf{x}(f, n) = [x_1(f, n), \ldots, x_I(f, n)]^\mathsf{T} \in \mathbb{C}^I, \quad (1)$$

$$\mathbf{s}(f, n) = [s_1(f, n), \ldots, s_I(f, n)]^\mathsf{T} \in \mathbb{C}^I, \quad (2)$$

where $(\cdot)^\mathsf{T}$ denotes transpose. Now, we use a separation system of the form

$$\mathbf{s}(f, n) = \mathbf{W}^\mathsf{H}(f)\mathbf{x}(f, n), \quad (3)$$

$$\mathbf{W}(f) = [\mathbf{w}_1(f), \ldots, \mathbf{w}_I(f)], \quad (4)$$

to describe the relationship between $\mathbf{x}(f, n)$ and $\mathbf{s}(f, n)$ where $\mathbf{W}^\mathsf{H}(f)$ is usually called the separation matrix. $(\cdot)^\mathsf{H}$ denotes Hermitian transpose. The aim of BSS methods is to estimate $\mathbf{W}^\mathsf{H}(f)$ solely from the observation $\mathbf{x}(f, n)$.

Let us now assume that $s_j(f, n)$ independently follows a zero-mean complex Gaussian distribution with variance $v_j(f, n) = \mathbb{E}[|s_j(f, n)|^2]$

$$s_j(f, n) \sim \mathcal{N}_{\mathbb{C}}(s_j(f, n)|0, v_j(f, n)). \quad (5)$$

We call (5) the local Gaussian model (LGM). When $s_j(f, n)$ and $s_{j'}(f, n)$ $(j \neq j')$ are independent, $\mathbf{s}(f, n)$ follows

$$\mathbf{s}(f, n) \sim \mathcal{N}_{\mathbb{C}}(\mathbf{s}(f, n)|\mathbf{0}, \mathbf{V}(f, n)), \quad (6)$$

where $\mathbf{V}(f, n)$ is a diagonal matrix with diagonal entries $v_1(f, n), \ldots, v_I(f, n)$. From (3) and (5), we can show that $\mathbf{x}(f, n)$ follows

$$\mathbf{x}(f, n) \sim \mathcal{N}_{\mathbb{C}}(\mathbf{x}(f, n)|\mathbf{0}, (\mathbf{W}^\mathsf{H}(f))^{-1}\mathbf{V}(f, n)\mathbf{W}(f)^{-1}). \quad (7)$$

Hence, the log-likelihood of the separation matrices $\mathcal{W} = \{\mathbf{W}(f)\}_f$ given the observed mixture signals $\mathcal{X} = \{\mathbf{x}(f, n)\}_{f,n}$ is given by

$$\log p(\mathcal{X}|\mathcal{W}, \mathcal{V}) \overset{c}{=} 2N \sum_f \log |\det \mathbf{W}^\mathsf{H}(f)|$$

$$- \sum_{f,n} \sum_j \left(\log v_j(f, n) + \frac{|\mathbf{w}_j^\mathsf{H}(f)\mathbf{x}(f, n)|^2}{v_j(f, n)} \right), \quad (8)$$

where $\overset{c}{=}$ denotes equality up to constant terms. If we individually treat $v_j(f, n)$ as a free parameter, all the variables

in (8) will be indexed by frequency f. The optimization problem will thus be split into frequency-wise source separation problems. Under this problem setting, the permutation of the separated components in each frequency cannot be uniquely determined and so permutation alignment must be performed after $\mathcal{W}$ has been obtained. However, it is preferable to solve permutation alignment and source separation jointly since the clues used for permutation alignment can also be helpful for source separation. If there is a certain assumption, constraint or structure that we can incorporate into $v_j(f, n)$, it can help eliminate the permutation ambiguity during the estimation of $\mathcal{W}$. One such example is the NMF model, which expresses $v_j(f, n)$ as the linear sum of spectral templates $b_{j,1}(f), \ldots, b_{j,K_j}(f) \geq 0$ scaled by time-varying magnitudes $h_{j,1}(n), \ldots, h_{j,K_j}(n) \geq 0$:

$$v_j(f, n) = \sum_{k=1}^{K_j} b_{j,k}(f) h_{j,k}(n). \quad (9)$$

ILRMA is a BSS framework that incorporates this model into the log-likelihood (8) [4,6,7]. Here, in a particular case where $K_j = 1$ and $b_{j,k}(f) = 1$ for all j in (9), which means each source has only one flat-shaped spectral template, assuming $s_j(0, n), \ldots, s_j(F, n)$ independently follow (5) is equivalent to assuming the norm $r_j(n) = \sqrt{\sum_f |s_j(f, n)|^2}$ follows a Gaussian distribution with time-varying variance $h_j(n)$. This is analogous to the assumption employed by IVA where the norm $r_j(n)$ is assumed to follow a supergaussian distribution. [6] showed that ILRMA can significantly outperform IVA in terms of source separation ability. This fact implies that within the LGM-based BSS framework, the stronger the representation power of a power spectrogram model becomes, the better the source separation performance we can expect to obtain.

3. RELATED WORK

3.1. ILRMA

The optimization algorithm of ILRMA consists of iteratively updating $\mathcal{W}$, $\mathcal{B} = \{b_{j,k}(f)\}_{j,k,f}$ and $\mathcal{H} = \{h_{j,k}(n)\}_{j,k,n}$ so that (8) is guaranteed to be non-decreasing at each iteration [4,6,7]. To update $\mathcal{W}$, we can use the natural gradient method or IP. The IP-based update rule for $\mathcal{W}$ [10] is given as

$$\mathbf{w}_j(f) \leftarrow (\mathbf{W}^\mathsf{H}(f)\mathbf{\Sigma}_j(f))^{-1}\mathbf{e}_j, \quad (10)$$

$$\mathbf{w}_j(f) \leftarrow \frac{\mathbf{w}_j(f)}{\mathbf{w}_j^\mathsf{H}(f)\mathbf{\Sigma}_j(f)\mathbf{w}_j(f)}, \quad (11)$$

where $\mathbf{\Sigma}_j(f) = \frac{1}{N}\sum_n \mathbf{x}(f, n)\mathbf{x}^\mathsf{H}(f, n)/v_j(f, n)$ and $\mathbf{e}_j$ denotes the j-th column of the $I \times I$ identity matrix. To update $\mathcal{B}$ and $\mathcal{H}$, we can employ the expectation-maximization (EM) algorithm or the majorization-minimization (MM) algorithm. The MM-based update rules for $\mathcal{B}$ and $\mathcal{H}$ can be
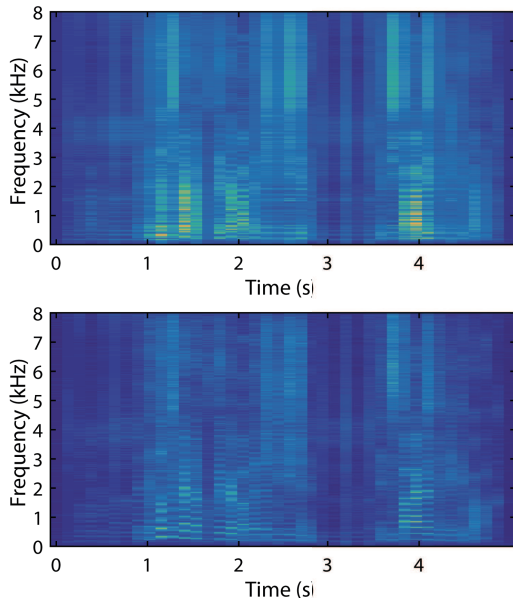
Fig. 1. Example of the NMF model (top) fitted to a speech spectrogram (bottom).

derived [15–17] as

$$b_{j,k}(f) \leftarrow b_{j,k}(f)\sqrt{\frac{\sum_n |y_j(f,n)|^2 h_{j,k}(n)/v_j^2(f,n)}{\sum_n h_{j,k}(n)/v_j(f,n)}}, \quad (12)$$

$$h_{j,k}(n) \leftarrow h_{j,k}(n)\sqrt{\frac{\sum_f |y_j(f,n)|^2 b_{j,k}(f)/v_j^2(f,n)}{\sum_f b_{j,k}(f)/v_j(f,n)}}, \quad (13)$$

where $y_j(f,n) = \mathbf{w}_j^{\mathsf{H}}(f)\mathbf{x}(f,n)$.

ILRMA is notable in that the optimization algorithm is guaranteed to converge to a stationary point of (8) and is shown experimentally to converge quickly. However, one limitation is that since $v_j(f,n)$ is restricted to (9), it can fail to work for sources with spectrograms that do not actually follow (9). Fig. 1 shows an example of the NMF model optimally fitted to a speech spectrogram. As can be seen from this example, there is still plenty of room for improvement in the model design.

3.2. DNN approach

As an alternative to the NMF model, some attempts have recently been made to combine deep neural networks (DNNs) with the LGM-based multichannel source separation framework [11, 12]. [11, 12] propose algorithms where $v_j(f,n)$ is updated at each iteration to the output of pretrained DNNs

$$\tilde{\mathbf{v}}_j(n) \leftarrow \mathrm{DNN}(\tilde{\mathbf{y}}_j(n); \theta_j) \quad (n = 1, \ldots, N). \quad (14)$$

Here, $\mathrm{DNN}(\cdot; \theta_j)$ indicates the output of the pretrained DNN, θ_j is the set of NN parameters, $\tilde{\mathbf{y}}_j(n) = \{|y_j(f, n \pm n')|\}_{f,n'}$ denotes the magnitude spectra of the estimate of the j-th

separated signal around the n-th time frame and $\tilde{\mathbf{v}}_j(n) = \{\sqrt{v_j(f,n)}\}_f$. With this approach, multiple DNNs are trained, and the j-th DNN is trained so that it produces only spectra related to source j in noisy input spectra. (14) can thus be seen as a process of refining the magnitude spectra of the separated signals according to the training examples of the known sources.

While this approach is noteworthy in that it can exploit the benefits of the representation power of DNNs for source power spectrum modeling, one drawback is that the devised iterative algorithm is not guaranteed to converge to a stationary point of the log-likelihood since updating $v_j(f,n)$ in this way does not guarantee an increase in the log-likelihood.

3.3. Source separation using deep generative models

It is worth noting that there have been some attempts to apply deep generative models including VAEs [13, 14] and generative adversarial networks (GANs) [18] to monaural speech enhancement and source separation [19, 20]. However, to the best of our knowledge, their applications to multichannel source separation has yet to be proposed.

4. PROPOSED METHOD

To address the limitations and drawbacks of the conventional methods, this paper proposes a multichannel source separation method using VAEs for source spectrogram modeling. We briefly review the idea behind the VAEs in 4.1 and present the proposed source separation algorithm in 4.2, which we call the multichannel VAE (MVAE).

4.1. Variational autoencoder (VAE)

VAEs [13, 14] are stochastic neural network models consisting of encoder and decoder networks. The encoder network generates a set of parameters for the conditional distribution $q_\phi(\mathbf{z}|\mathbf{s})$ of a latent space variable $\mathbf{z}$ given input data $\mathbf{s}$ whereas the decoder network generates a set of parameters for the conditional distribution $p_\theta(\mathbf{s}|\mathbf{z})$ of the data $\mathbf{s}$ given the latent space variable $\mathbf{z}$. Given a training dataset $\mathcal{S} = \{\mathbf{s}_m\}_{m=1}^M$, VAEs learn the parameters of the entire network so that the encoder distribution $q_\phi(\mathbf{z}|\mathbf{s})$ becomes consistent with the posterior $p_\theta(\mathbf{z}|\mathbf{s}) \propto p_\theta(\mathbf{s}|\mathbf{z})p(\mathbf{z})$. By using Jensen's inequality, the log marginal distribution of the data $\mathbf{s}$ can be lower-bounded by

$$\log p_\theta(\mathbf{s}) = \log \int q_\phi(\mathbf{z}|\mathbf{s})\frac{p_\theta(\mathbf{s}|\mathbf{z})p(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{s})}d\mathbf{z}$$

$$\geq \int q_\phi(\mathbf{z}|\mathbf{s}) \log \frac{p_\theta(\mathbf{s}|\mathbf{z})p(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{s})}d\mathbf{z} \quad (15)$$

$$= \mathbb{E}_{\mathbf{z}\sim q_\phi(\mathbf{z}|\mathbf{s})}[\log p_\theta(\mathbf{s}|\mathbf{z})] - \mathrm{KL}[q_\phi(\mathbf{z}|\mathbf{s})\|p(\mathbf{z})],$$

where the difference between the left- and right-hand sides of this inequality is equal to the Kullback-Leibler divergence
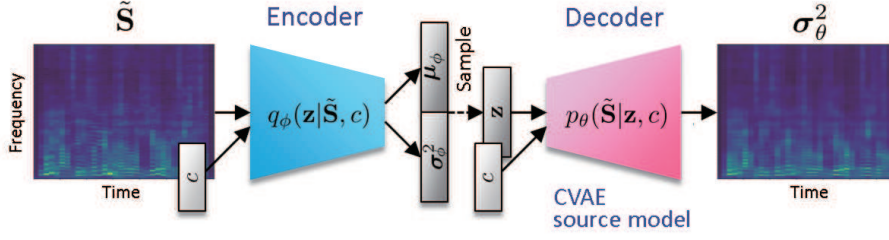
Fig. 2. Illustration of the present CVAE.

KL$[q_\phi(\mathbf{z}|\mathbf{s})\|p_\theta(\mathbf{z}|\mathbf{s})]$, which is minimized when

$$q_\phi(\mathbf{z}|\mathbf{s}) = p_\theta(\mathbf{z}|\mathbf{s}). \qquad (16)$$

This means we can make $q_\phi(\mathbf{z}|\mathbf{s})$ and $p_\theta(\mathbf{z}|\mathbf{s}) \propto p_\theta(\mathbf{s}|\mathbf{z})p(\mathbf{z})$ consistent by maximizing the lower bound of (15). One typical way of modeling $q_\phi(\mathbf{z}|\mathbf{s})$, $p_\theta(\mathbf{s}|\mathbf{z})$ and $p(\mathbf{z})$ is to assume Gaussian distributions

$$q_\phi(\mathbf{z}|\mathbf{s}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_\phi(\mathbf{s}), \mathrm{diag}(\boldsymbol{\sigma}_\phi^2(\mathbf{s}))), \qquad (17)$$

$$p_\theta(\mathbf{s}|\mathbf{z}) = \mathcal{N}(\mathbf{s}|\boldsymbol{\mu}_\theta(\mathbf{z}), \mathrm{diag}(\boldsymbol{\sigma}_\theta^2(\mathbf{z}))), \qquad (18)$$

$$p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \mathbf{I}), \qquad (19)$$

where $\boldsymbol{\mu}_\phi(\mathbf{s})$ and $\boldsymbol{\sigma}_\phi^2(\mathbf{s})$ are the outputs of an encoder network with parameter ϕ, and $\boldsymbol{\mu}_\theta(\mathbf{z})$ and $\boldsymbol{\sigma}_\theta^2(\mathbf{z})$ are the outputs of a decoder network with parameter θ. The first term of the lower bound can be interpreted as an autoencoder reconstruction error since it can be written as

$$\mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{s})}[\log p(\mathbf{s}|\mathbf{z})]$$
$$=\mathbb{E}_{\boldsymbol{\epsilon}\sim\mathcal{N}(\boldsymbol{\epsilon}|\mathbf{0},\mathbf{I})}\bigg[-\frac{1}{2}\sum_i \log 2\pi[\boldsymbol{\sigma}_\theta^2(\boldsymbol{\mu}_\phi(\mathbf{s}) + \boldsymbol{\sigma}_\phi(\mathbf{s})\odot\boldsymbol{\epsilon})]_n$$
$$-\sum_n \frac{(s_n - [\boldsymbol{\mu}_\theta(\boldsymbol{\mu}_\phi(\mathbf{s}) + \boldsymbol{\sigma}_\phi(\mathbf{s})\odot\boldsymbol{\epsilon})]_n)^2}{2[\boldsymbol{\sigma}_\theta^2(\boldsymbol{\mu}_\phi(\mathbf{s}) + \boldsymbol{\sigma}_\phi(\mathbf{s})\odot\boldsymbol{\epsilon})]_n}\bigg], \qquad (20)$$

which reduces to a negative weighted squared error between $\mathbf{s}$ and $\boldsymbol{\mu}_\theta(\boldsymbol{\mu}_\phi(\mathbf{s}))$ if we exclude all the stochastic terms related to $\boldsymbol{\epsilon}$. Here, we have used a reparameterization $\mathbf{z} = \boldsymbol{\mu}_\phi(\mathbf{s}) + \boldsymbol{\sigma}_\phi(\mathbf{s})\odot\boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon}\sim\mathcal{N}(\boldsymbol{\epsilon}|\mathbf{0},\mathbf{I})$ where $\odot$ indicates the element-wise product and $[\cdot]_n$ denotes the n-th element of a vector. On the other hand, the second term is given as the negative KL divergence between $q_\phi(\mathbf{z}|\mathbf{s})$ and $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\mathbf{0},\mathbf{I})$. This term can be interpreted as a regularization term that forces each element of the encoder output to be independent and normally distributed.

Conditional VAEs (CVAEs) [14] are an extended version of VAEs where the only difference is that the encoder and decoder networks can take an auxiliary variable c as an additional input. With CVAEs, (17) and (18) are replaced with

$$q_\phi(\mathbf{z}|\mathbf{s},c) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_\phi(\mathbf{s},c), \mathrm{diag}(\boldsymbol{\sigma}_\phi^2(\mathbf{s},c))), \qquad (21)$$

$$p_\theta(\mathbf{s}|\mathbf{z},c) = \mathcal{N}(\mathbf{s}|\boldsymbol{\mu}_\theta(\mathbf{z},c), \mathrm{diag}(\boldsymbol{\sigma}_\theta^2(\mathbf{z},c))), \qquad (22)$$

and the variational lower bound to be maximized becomes

$$\mathcal{J}(\phi,\theta) = \mathbb{E}_{(\mathbf{s},c)\sim p_\mathrm{D}(\mathbf{s},c)}\big[\mathbb{E}_{\mathbf{z}\sim q(\mathbf{z}|\mathbf{s},c)}[\log p(\mathbf{s}|\mathbf{z},c)]$$

$$-\mathrm{KL}[q(\mathbf{z}|\mathbf{s},c)\|p(\mathbf{z})]\big], \qquad (23)$$

where $\mathbb{E}_{(\mathbf{s},c)\sim p_\mathrm{D}(\mathbf{s},c)}[\cdot]$ denotes the sample mean over the training examples $\{\mathbf{s}_m, c_m\}_{m=1}^M$.

One notable feature as regards CVAEs is that they are able to learn a "disentangled" latent representation underlying the data of interest. For example, when a CVAE is trained using the MNIST dataset of handwritten digits and c as the digit class label, $\mathbf{z}$ and c are disentangled so that $\mathbf{z}$ represents the factors of variation corresponding to handwriting styles. We can thus generate images of a desired digit with random handwriting styles from the trained decoder by specifying c and randomly sampling $\mathbf{z}$. Analogously, we would be able to obtain a generative model that can represent the spectrograms of a variety of sound sources if we could train a CVAE using class-labeled training examples.

4.2. Multichannel VAE

Let $\tilde{\mathbf{S}} = \{s(f,n)\}_{f,n}$ be the complex spectrogram of a particular sound source and c be the class label of that source. Here, we assume that a class label comprises one or more categories, each consisting of multiple classes. We thus represent c as a concatenation of one-hot vectors, each of which is filled with 1 at the index of a class in a certain category and with 0 everywhere else. For example, if we consider speaker identities as the only class category, c will be represented as a single one-hot vector, where each element is associated with a different speaker.

We now model the generative model of $\tilde{\mathbf{S}}$ using a CVAE with an auxiliary input c. So that the decoder distribution has the same form as the LGM (5), we define it as a zero-mean complex Gaussian distribution

$$p_\theta(\tilde{\mathbf{S}}|\mathbf{z},c,g) = \prod_{f,n}\mathcal{N}_{\mathbb{C}}(s(f,n)|0, v(f,n)), \qquad (24)$$

$$v(f,n) = g\cdot\sigma_\theta^2(f,n;\mathbf{z},c), \qquad (25)$$

where $\sigma_\theta^2(f,n;\mathbf{z},c)$ denotes the (f,n)-th element of the decoder output $\boldsymbol{\sigma}_\theta^2(\mathbf{z},c)$ and g represents the global scale of the generated spectrogram. As regards the encoder distribution $q_\phi(\mathbf{z}|\tilde{\mathbf{S}},c)$, we adopt a regular Gaussian distribution

$$q_\phi(\mathbf{z}|\tilde{\mathbf{S}},c) = \prod_k \mathcal{N}(z(k)|\mu_\phi(k;\tilde{\mathbf{S}},c), \sigma_\phi^2(k;\tilde{\mathbf{S}},c)), \qquad (26)$$

where $z(k)$, $\mu_\phi(k; \tilde{\mathbf{S}}, c)$ and $\sigma_\phi^2(k; \tilde{\mathbf{S}}, c)$ represent the k-th elements of the latent space variable $\mathbf{z}$ and the encoder outputs $\boldsymbol{\mu}_\phi(\tilde{\mathbf{S}}, c)$ and $\boldsymbol{\sigma}_\phi^2(\tilde{\mathbf{S}}, c)$, respectively. Given a set of labeled training examples $\{\tilde{\mathbf{S}}_m, c_m\}_{m=1}^M$, we train the decoder and encoder NN parameters θ and ϕ, respectively, prior to source separation, using the training objective

$$\mathcal{J}(\phi, \theta) = \mathbb{E}_{(\tilde{\mathbf{S}}, c) \sim p_D(\tilde{\mathbf{S}}, c)} \big[\mathbb{E}_{\mathbf{z} \sim q(\mathbf{z} | \tilde{\mathbf{S}}, c)} [\log p(\tilde{\mathbf{S}} | \mathbf{z}, c)] \\ - \mathrm{KL}[q(\mathbf{z} | \tilde{\mathbf{S}}, c) \| p(\mathbf{z})] \big], \quad (27)$$

where $\mathbb{E}_{(\tilde{\mathbf{S}}, c) \sim p_D(\tilde{\mathbf{S}}, c)}[\cdot]$ denotes the sample mean over the training examples $\{\tilde{\mathbf{S}}_m, c_m\}_{m=1}^M$. Fig. 2 shows the illustration of the present CVAE.

The trained decoder distribution $p_\theta(\tilde{\mathbf{S}} | \mathbf{z}, c, g)$ can be used as a universal generative model that is able to generate spectrograms of all the sources involved in the training examples where the latent space variable $\mathbf{z}$, the auxiliary input c and the global scale g can be interpreted as the model parameters. According to the properties of CVAEs, we consider that the CVAE training promotes disentanglement between $\mathbf{z}$ and c where $\mathbf{z}$ characterizes the factors of intra-class variation whereas c characterizes the factors of categorical variation that represent source identities. We call $p_\theta(\tilde{\mathbf{S}} | \mathbf{z}, c, g)$ the CVAE source model.

Since the CVAE source model is given in the same form as the LGM given by (5), we can develop a log-likelihood that has the same expression as (8) if we use $p_\theta(\tilde{\mathbf{S}}_j | \mathbf{z}_j, c_j, g_j)$ to express the generative model of the complex spectrogram of source j. Hence, we can search for a stationary point of the log-likelihood by iteratively updating the separation matrices $\mathcal{W}$, the global scale parameter $\mathcal{G} = \{g_j\}_j$ and the VAE source model parameters $\Psi = \{\mathbf{z}_j, c_j\}_j$ so that the log-likelihood is guaranteed to be non-decreasing at each iteration. We can use (10) and (11) to update $\mathcal{W}$, backpropagation to update Ψ and

$$g_j \leftarrow \frac{1}{FN} \sum_{f,n} \frac{|y_j(f, n)|^2}{\sigma_\theta^2(f, n; \mathbf{z}_j, c_j)}, \quad (28)$$

to update $\mathcal{G}$ where $y_j(f, n) = \mathbf{w}_j^{\mathsf{H}}(f) \mathbf{x}(f, n)$. Note that (28) maximizes (8) with respect to g_j when $\mathcal{W}$ and Ψ are fixed.

The proposed algorithm is thus summarized as follows:

1. Train θ and ϕ using (27).

2. Initialize $\mathcal{W}$, $\mathcal{G}$ and $\Psi = \{\mathbf{z}_j, c_j\}_j$.

3. Iterate the following steps for each j:

   (a) Update $\mathbf{w}_j(0), \dots, \mathbf{w}_j(F)$ using (10) and (11).

   (b) Update $\psi_j = \{\mathbf{z}_j, c_j\}$ using backpropagation.

   (c) Update g_j using (28).

The proposed MVAE is noteworthy in that it offers the advantages of the conventional methods concurrently. Namely, (1) it takes full advantage of the strong representation power of DNNs for source power spectrogram modeling, (2) the
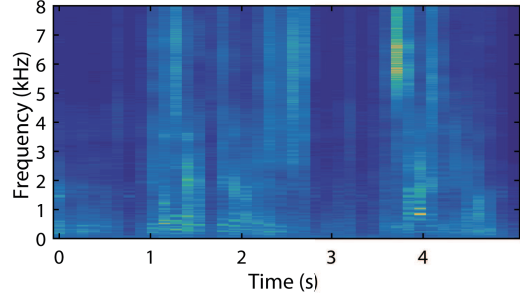

Fig. 3. Example of the CVAE source model fitted to the speech spectrogram shown in Fig. 1.

convergence of the source separation algorithm is guaranteed, and (3) the criteria for CVAE training and source separation are consistent, thanks to the consistency between the expressions of the CVAE source model and the LGM. Fig. 3 shows an example of the CVAE source model fitted to the speech spectrogram shown in Fig. 1. We can confirm from this example that the CVAE source model is able to approximate the speech spectrogram somewhat better than the NMF model.

4.3. Network architectures

We propose designing the encoder and decoder networks using fully convolutional architectures to allow the encoder to take a spectrogram as an input and allow the decoder to output a spectrogram of the same length instead of a single-frame spectrum. This allows the networks to capture time dependencies in spectral sequences. While RNN-based architectures are a natural choice for modeling time series data, we use convolutional neural network (CNN)-based architectures to design the encoder and decoder as detailed below.

We use 1D CNNs to design the encoder and the decoder networks by treating $\tilde{\mathbf{S}}$ as an image of size $1 \times N$ with F channels. Specifically, we use a gated CNN [21], which was originally introduced to model word sequences for language modeling and was shown to outperform long short-term memory (LSTM) language models trained in a similar setting. We previously employed gated CNN architectures for voice conversion [22–24] and monaural audio source separation [25], and have already confirmed their effectiveness. In the encoder, the output of the l-th hidden layer, $\mathbf{h}_l$, is described as a linear projection modulated by an output gate

$$\mathbf{h}_{l-1}' = [\mathbf{h}_{l-1}; \mathbf{c}_{l-1}], \quad (29)$$
$$\mathbf{h}_l = (\mathbf{W}_l * \mathbf{h}_{l-1}' + \mathbf{b}_l) \odot \sigma(\mathbf{V}_l * \mathbf{h}_{l-1}' + \mathbf{d}_l), \quad (30)$$

where $\mathbf{W}_l \in \mathbb{R}^{D_l \times D_{l-1} \times 1 \times N_l}$, $\mathbf{b}_l \in \mathbb{R}^{D_l}$, $\mathbf{V}_l \in \mathbb{R}^{D_l \times D_{l-1} \times 1 \times N_l}$ and $\mathbf{d}_l \in \mathbb{R}^{D_l}$ are the encoder network parameters ϕ, and σ denotes the elementwise sigmoid function. Similar to LSTMs, the output gate multiplies each element of $\mathbf{W}_l * \mathbf{h}_{l-1} + \mathbf{b}_l$ and controls what information should be propagated through the hierarchy of layers. This gating mechanism is called a gated linear unit (GLU). Here, $[\mathbf{h}_l; \mathbf{c}_l]$
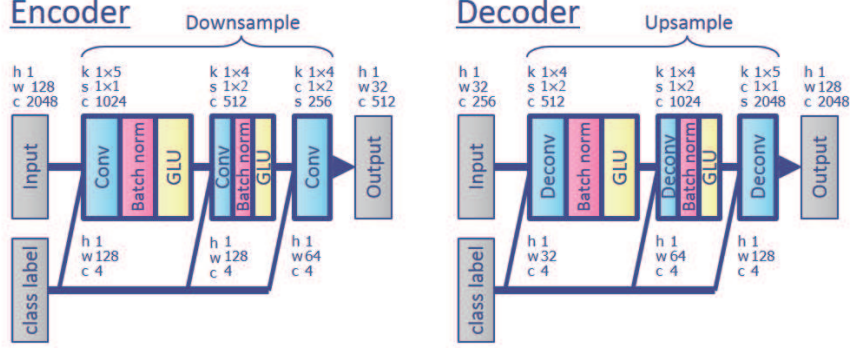
Fig. 4. Network architectures of the encoder and decoder. Here, the inputs and outputs of the encoder and decoder are interpreted as images, where "h", "w" and "c" denote the height, width and channel number, respectively. "Conv", "Batch norm", "GLU", "Deconv" denote convolution, batch normalization, gated linear unit, and transposed convolution layers, respectively. "k", "s" and "c" denote the kernel size, stride size and output channel number of a convolution layer, respectively. Note that all the networks are fully convolutional with no fully connected layers, thus allowing inputs to have arbitrary lengths.

means the concatenation of $\mathbf{h}_l$ and $\mathbf{c}_l$ along the channel dimension, and $\mathbf{c}_l$ is a 2D array consisting of a N_l tiling of copies of c in the time dimensions. The input into the 1st layer of the encoder is $\mathbf{h}_0 = \tilde{\mathbf{S}}$. The outputs of the final layer are given as regular linear projections

$$\boldsymbol{\mu}_\phi = \mathbf{W}_L * \mathbf{h}'_{L-1} + \mathbf{b}_L, \tag{31}$$

$$\log \boldsymbol{\sigma}_\phi^2 = \mathbf{V}_L * \mathbf{h}'_{L-1} + \mathbf{d}_L. \tag{32}$$

The decoder network is devised in the same way as below with the only difference being that $\boldsymbol{\mu}_\theta = \mathbf{0}$:

$$\mathbf{h}_0 = \mathbf{z},$$
$$\mathbf{h}'_{l-1} = [\mathbf{h}_{l-1}; \mathbf{c}_{l-1}],$$
$$\mathbf{h}'_l = (\mathbf{W}'_l * \mathbf{h}'_{l-1} + \mathbf{b}'_l) \odot \sigma(\mathbf{V}'_l * \mathbf{h}'_{l-1} + \mathbf{d}'_l),$$
$$\boldsymbol{\mu}_\theta = \mathbf{0},$$
$$\log \boldsymbol{\sigma}_\theta^2 = \mathbf{V}'_L * \mathbf{h}'_{L-1} + \mathbf{d}'_L,$$

where $\mathbf{W}'_l \in \mathbb{R}^{D_l \times D_{l-1} \times 1 \times N_l}$, $\mathbf{b}'_l \in \mathbb{R}^{D_l}$, $\mathbf{V}'_l \in \mathbb{R}^{D_l \times D_{l-1} \times 1 \times N_l}$ and $\mathbf{d}'_l \in \mathbb{R}^{D_l}$ are the decoder network parameters θ. It should be noted that the entire architecture is fully convolutional with no fully-connected layers. The trained decoder can therefore be used a generative model of spectrograms with arbitrary lengths. This is particularly convenient when designing source separation systems since they can allow signals of any length.

5. EXPERIMENTS

To confirm the effect of the incorporation of the CVAE source model, we conducted experiments involving a semi-blind source separation task using speech mixtures. We excerpted speech utterances from the Voice Conversion Challenge (VCC) 2018 dataset [26], which consists of recordings
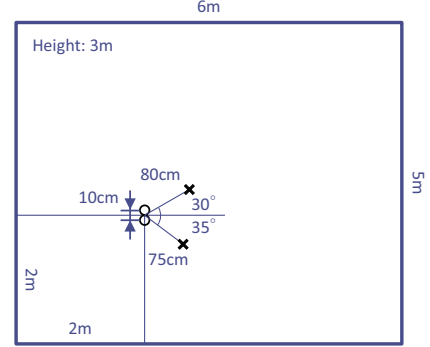

Fig. 5. Simulated room configuration.

of six female and six male US English speakers. Specifically, we used the utterances of two female speakers, 'SF1' and 'SF2', and two male speakers, 'SM1' and 'SM2' for CVAE training and source separation. We considered speaker identities as the only source class category. Thus, c was a four-dimensional one-hot vector. The audio files for each speaker were manually segmented into 116 short sentences (each about 7 minutes long) where 81 and 35 sentences (about 5 and 2 minutes long, respectively) were provided as training and evaluation sets, respectively. We used simulated two-channel recordings of two sources as the test data where the impulse responses were synthesized by using the image method. Fig. 5 shows the two-dimensional configuration of the room. $\circ$ and $\times$ represent the positions of microphones and sources, respectively. The reverberation time (RT_{60}) of the simulated signals could be controlled according to the setting of the reflection coefficient of the walls. To simulate anechoic and echoic environments, we created test signals with the reflection coefficients set at 0.20 and 0.80, respectively. The corresponding RT_{60}s were 78 [ms] and 351 [ms], respectively. We generated 10 speech mixtures
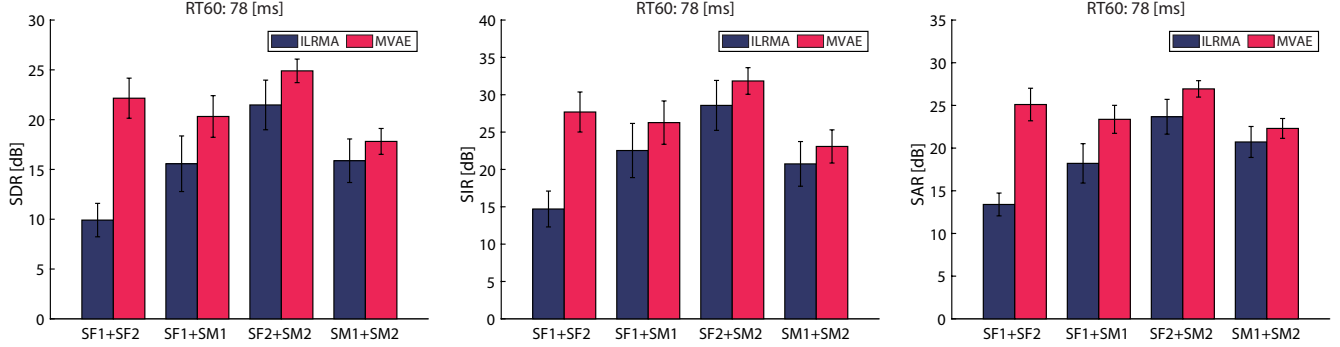
Fig. 6. Average SDRs, SIRs and SARs obtained with the baseline and proposed methods for RT_{60} of 78 [ms].
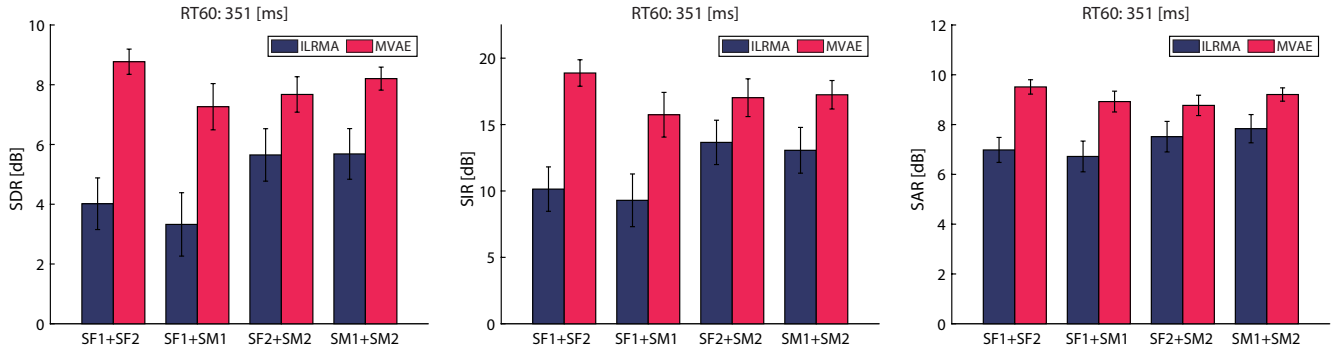

Fig. 7. Average SDRs, SIRs and SARs obtained with the baseline and proposed methods for RT_{60} of 351 [ms].

for each speaker pair, SF1+SF2, SF1+SM1, SM1+SM2, and SF2+SM2. Hence, there were 40 test signals in total, each of which was about 4 to 7 [s] long. All the speech signals were re-sampled at 16000 [Hz]. The STFT frame length was set at 256 [ms] and a Hamming window was used with an overlap length of 128 [ms].

We chose ILRMA [4, 6, 7] as a baseline method for comparison. The source separation algorithms were run for 40 iterations for the proposed method and 100 iterations for the baseline method. For the proposed method, $\mathcal{W}$ was initialized using the baseline method run for 30 iterations and Adam optimization [27] was used for CVAE training and the estimation of Ψ in the source separation algorithm. The network configuration we used for the proposed method is shown in detail in Fig. 4. Note that we must take account of the sum-to-one constraints when updating c_j. This can be easily implemented by inserting an appropriately designed softmax layer that outputs c_j

$$c_j = \mathrm{softmax}(u_j), \qquad (33)$$

and treating u_j as the parameter to be estimated instead.

To evaluate the source separation performance, we took the averages of the signal-to-distortion ratio (SDR), signal-to-interference ratio (SIR) and signal-to-artifact ratio (SAR) [28] of the separated signals obtained with the baseline and proposed methods using 10 test signals for each speaker pair. Figs. 6 and 7 show the average SDRs, SIRs and SARs ob-

tained with the baseline and proposed methods under different RT_{60} conditions. As the results show, the proposed method significantly outperformed the baseline method, revealing the advantage of the CVAE source model. Audio samples are provided at: http://www.kecl.ntt.co.jp/people/kameoka.hirokazu/Demos/mvae-ass/.

As can be seen from a comparison between the results in Figs. 6 and 7, there were noticeable performance degradations with both the baseline and proposed methods when the reverberation became relatively long. We hope that these degradations can be overcome by introducing the idea of jointly solving dereverberation and source separation, as in [4, 29, 30].

6. CONCLUSIONS

This paper proposed a multichannel source separation technique called the multichannel variational autoencoder (MVAE) method. The method used VAEs to model and estimate the power spectrograms of the sources in mixture signals. The key features of the MVAE are that (1) it takes full advantage of the strong representation power of deep neural networks for source power spectrogram modeling, (2) the convergence of the source separation algorithm is guaranteed, and (3) the criteria for the VAE training and source separation are consistent, which contributed to obtaining better separations than with conventional methods.

7. REFERENCES

[1] T. Kim, T. Eltoft, and T.-W. Lee, "Independent vector analysis: An extension of ICA to multivariate components," in *Proc. ICA*, 2006, pp. 165–172.

[2] A. Hiroe, "Solution of permutation problem in frequency domain ICA using multivariate probability density functions," in *Proc. ICA*, 2006, pp. 601–608.

[3] A. Ozerov and C. Févotte, "Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation," *IEEE Trans. ASLP*, vol. 18, no. 3, pp. 550–563, 2010.

[4] H. Kameoka, T. Yoshioka, M. Hamamura, J. Le Roux, and K. Kashino, "Statistical model of speech signals based on composite autoregressive system with application to blind source separation," in *Proc. LVA/ICA*, 2010, pp. 245–253.

[5] H. Sawada, H. Kameoka, S. Araki, and N. Ueda, "Multichannel extensions of non-negative matrix factorization with complex-valued data," *IEEE Trans. ASLP*, vol. 21, no. 5, pp. 971–982, 2013.

[6] D. Kitamura, N. Ono, H. Sawada, H. Kameoka, and H. Saruwatari, "Determined blind source separation unifying independent vector analysis and nonnegative matrix factorization," *IEEE/ACM Trans. ASLP*, vol. 24, no. 9, pp. 1626–1641, 2016.

[7] D. Kitamura, N. Ono, H. Sawada, H. Kameoka, and H. Saruwatari, "Determined blind source separation with independent low-rank matrix analysis," in *Audio Source Separation*, Shoji Makino, Ed., pp. 125–155. Springer, Mar. 2018.

[8] P. Smaragdis, "Non-negative matrix factorization for polyphonic music transcription," in *Proc. WASPAA*, 2003, pp. 177–180.

[9] C. Févotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the Itakura-Saito divergence," *Neural Comput.*, vol. 21, no. 3, pp. 793–830, 2009.

[10] N. Ono, "Stable and fast update rules for independent vector analysis based on auxiliary function technique," in *Proc. WASPAA*, 2011, pp. 189–192.

[11] A. A. Nugraha, A. Liutkus, and E. Vincent, "Multichannel audio source separation with deep neural networks," *IEEE/ACM Trans. ASLP*, vol. 24, no. 9, pp. 1652–1664, 2016.

[12] D. Kitamura, H. Sumino, N. Takamune, S. Takamichi, H. Saruwatari, and N. Ono, "Experimental evaluation of multichannel audio source separation based on IDLMA," in *IEICE Tech. Rep.*, 2018, vol. 117, pp. 13–20, in Japanese.

[13] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. ICLR*, 2014.

[14] D. P. Kingma and D. J. Rezendey, S. Mohamedy, and M. Welling, "Semi-supervised learning with deep generative models," in *Adv. Neural Information Processing Systems (NIPS)*, 2014, pp. 3581–3589.

[15] H. Kameoka, M. Goto, and S. Sagayama, "Selective amplifier of periodic and non-periodic components in concurrent audio signals with spectral control envelopes," in *IPSJ SIG Tech. Rep.*, 2006, vol. 2006-MUS-66-13, in Japanese.

[16] M. Nakano, H. Kameoka, J. Le Roux, N. Ono, and S. Sagayama, "Convergence-guaranteed multiplicative algorithms for non-negative matrix factorization with beta-divergence," in *Proc. MLSP*, 2010.

[17] C. Févotte and J. Idier, "Algorithms for nonnegative matrix factorization with the β-divergence," *Neural Comput.*, vol. 23, no. 9, pp. 2421–2456, 2011.

[18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial nets," in *Adv. Neural Information Processing Systems (NIPS)*, 2014, pp. 2672–2680.

[19] Y. Bando, M. Mimura, K. Itoyama, K. Yoshii, and T. Kawahara, "Statistical speech enhancement based on probabilistic integration of variational autoencoder and non-negative matrix factorization," in *Proc. ICASSP*, 2018, pp. 716–720.

[20] Y. Subakan and P. Smaragdis, "Generative adversarial source separation," in *Proc. ICASSP*, 2018, pp. 26–30.

[21] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," in *Proc. ICML*, 2017, pp. 933–941.

[22] T. Kaneko, H. Kameoka, K. Hiramatsu, and K. Kashino, "Sequence-to-sequence voice conversion with similarity metric learned using generative adversarial networks," in *Proc. Interspeech*, 2017, pp. 1283–1287.

[23] T. Kaneko and H. Kameoka, "Parallel-data-free voice conversion using cycle-consistent adversarial networks," *eprint arXiv:1711.11293*, Nov. 2017.

[24] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "StarGAN-VC: Non-parallel many-to-many voice conversion with star generative adversarial networks," *arXiv:1806.02169 [cs.SD]*, June 2018.

[25] L. Li and H. Kameoka, "Deep clustering with gated convolutional networks," in *Proc. ICASSP*, 2018, pp. 16–20.

[26] J. Lorenzo-Trueba, J. Yamagishi, T. Toda, D. Saito, F. Villavicencio, T. Kinnunen, and Z. Ling, "The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods," *eprint arXiv:1804.04262*, Apr. 2018.

[27] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.

[28] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE Trans. ASLP*, vol. 14, no. 4, pp. 1462–1469, 2006.

[29] T. Yoshioka, T. Nakatani, M. Miyoshi, and H. G. Okuno, "Blind separation and dereverberation of speech mixtures by joint optimization," *IEEE Trans. ASLP*, vol. 19, no. 1, pp. 69–84, 2011.

[30] H. Kagami, H. Kameoka, and M. Yukawa, "Joint separation and dereverberation of reverberant mixtures with determined multichannel non-negative matrix factorization," in *Proc. ICASSP*, 2018, pp. 31–35.