

INTEGRATION OF VARIATIONAL AUTOENCODER AND SPATIAL CLUSTERING FOR ADAPTIVE MULTI-CHANNEL NEURAL SPEECH SEPARATION

Katerina Zmolikova¹, Marc Delcroix², Lukáš Burget¹, Tomohiro Nakatani², Jan "Honza" Černocký¹

¹Brno University of Technology, Faculty of IT, IT4I Centre of Excellence
²NTT Communication Science Laboratories, NTT Corporation, Kyoto, Japan

ABSTRACT

In this paper, we propose a method combining variational autoencoder model of speech with a spatial clustering approach for multi-channel speech separation. The advantage of integrating spatial clustering with a spectral model was shown in several works. As the spectral model, previous works used either factorial generative models of the mixed speech or discriminative neural networks. In our work, we combine the strengths of both approaches, by building a factorial model based on a generative neural network, a variational autoencoder. By doing so, we can exploit the modeling power of neural networks, but at the same time, keep a structured model. Such a model can be advantageous when adapting to new noise conditions as only the noise part of the model needs to be modified. We show experimentally, that our model significantly outperforms previous factorial model based on Gaussian mixture model (DOLPHIN), performs comparably to integration of permutation invariant training with spatial clustering, and enables us to easily adapt to new noise conditions.

Index Terms— Multi-channel speech separation, variational autoencoder, spatial clustering, DOLPHIN

1. INTRODUCTION

Multi-channel speech separation is a problem of great importance for applications of speech technologies, such as personal devices or transcription of meetings. A popular class of approaches to tackle this problem is based on spatial clustering [1–4]. In these approaches, a mixture model of spatial features (e.g. phase differences [4] or normalized observation vectors [1]) is built. By inferring the affiliations of observed spatial features to the mixture components, we obtain time-frequency masks, which can then be used for separating the signal. The spatial clustering approaches have proven their strength in many tasks, e.g. recent CHiME challenges [5, 6].

The spatial clustering approaches make use of the spatial properties of the speech signal, however, they do not use any information about the spectral structure of speech. To overcome this drawback, methods integrating the spatial model with a spectral model have been proposed [7–10]. We can observe two main directions of spectral models for such integration: first, factorial generative models of the mixed speech [7, 8] and second, discriminative neural networks [9, 10].

In *factorial generative models*, models of clean speech and noise signals are trained and combined into a factorial model of the mixture. In the past, such models have been applied also to single-

channel separation [11] and have been integrated with spatial models for multi-channel case [7, 8]. Notably, in DOminance based Locational and Power-spectral cHaracteristics INtegration (DOLPHIN) [7], Gaussian mixture models are used to model speech and noise signals, and the index of the dominant source is shared with a spatial model, to provide joint inference in both models. Factorial models have demonstrated great potential for separation at early challenges [12, 13]. A clear advantage of the factorial model is the explicit structure, which enables us to modify different parts of the model. For instance, if noise conditions in the recordings change, we can only modify the noise model, without changing the other parts, which enables to adapt these models easily to the recording conditions. On the other hand, models such as Gaussian mixture models are not flexible enough to capture the structure of speech well, as has become apparent in many speech tasks.

More recently, approaches based on *discriminative neural networks* have been applied. The success of approaches such as deep clustering [14] or permutation invariant training (PIT) [15, 16] started a new era in speech separation, where neural networks are considered state-of-the-art. In these approaches, a neural network is trained to map from the mixed signal to the individual sources (represented either directly, as time-frequency masks, or embeddings). The neural network-based approaches have been also combined with spatial clustering in [9, 10]. There, the outputs of the network are used as priors on the mixture weights in a spatial model. While the neural networks offer great modeling power, they suffer from problems in mismatched conditions. In contrast with the factorial models, we have no insight into how the neural network works with the speech and noise components of the signal. To adapt to different noise conditions, we thus need to retrain the full model on in-domain data.

To combine the strengths of both stated directions, we here propose to use a neural network as a generative model of the speech signal in a factorial model. For this, we use the variational autoencoder (VAE) [17, 18], which has been previously successfully used to model speech [19, 20]. We train VAE on clean single-speaker speech and create a simple Gaussian model of noise. Using the lifted-max-model of interaction [11], we combine the model of individual sources into a model of the mixture. As in DOLPHIN, we integrate the factorial model with a spatial model through the dominance index. At inference, to estimate the time-frequency masks of each source, we iteratively infer the time-frequency masks and the spectral properties of the signals, represented as latent variables.

Experimentally, we show that our proposed model 1) is stronger than DOLPHIN with GMM as a spectral model, 2) provides comparable results to PIT combined with spatial clustering on matched conditions, and 3) can adapt to different noise conditions by modifying the noise model. The last property can be useful for example in the scenario of long recordings, where we could progressively adapt the noise model with more incoming data.

The work was partly supported by Czech National Science Foundation (GACR) project NEUREM3 No. 19-26934X, and by Czech Ministry of Education, Youth and Sports from project no. LTAI19087 "Multi-linguality in speech technologies".

2. PROBLEM FORMULATION

In this paper, we process an observed multi-channel mixture of speech from multiple speakers and noise with the goal of separating the speech signals. In the time-domain, the mixture signal at the m -th microphone can be expressed as,

$$y_m[\tau] = \sum_i x_{m,i}[\tau] + n_m[\tau], \quad (1)$$

where τ is the time index, and $y_m[\tau]$ $x_{m,i}[\tau]$ $n_m[\tau]$ are the time domain signals of the mixture, the speech sources and the noise respectively. In the following we consider the case of two speakers, i.e. $i \in \{1, 2\}$ and work with all signals in Short-time Fourier transform (STFT) domain. Let us introduce following notations which will be used throughout the paper:

- $\mathcal{Y}_{m,t,f}$: STFT coefficient of the observed mixed signal at m -th microphone for time frame t and frequency bin f .
- $\mathcal{Y}_{t,f} = [\mathcal{Y}_{m,t,f}]_{m=1..M}^\top$: vector of STFT coefficients of the observed mixed signal at all microphones, where M is the number of microphones.
- $\Psi = [\psi_t]_{t=1..T}$, with $\psi_t = [\psi_{t,f}]_{f=1..F}^\top$: tensor of spatial features at all channels and all time-frequency bins (t, f) . The specific form of spatial features we use, is introduced in Eq. (3).
- $y_{t,f}$, $x_{t,f}^{(i)}$ and $n_{t,f}$: the log-magnitude of STFT coefficient at the first microphone and the (t, f) time-frequency bin of the observed mixture signal, the single-speaker speech signal of speaker $i \in \{1, 2\}$ and the noise signal, respectively.
- $\mathbf{Y} = [\mathbf{y}_t]_{t=1..T}$, with $\mathbf{y}_t = [y_{t,f}]_{f=1..F}^\top$: matrix of log-magnitude of STFT coefficients of the observed mixed signal at the first microphone and at all time-frequency bins (t, f) .
- $\mathbf{X}^{(i)} = [\mathbf{x}_t^{(i)}]_{t=1..T}$, with $\mathbf{x}_t^{(i)} = [x_{t,f}^{(i)}]_{f=1..F}^\top$: matrix of log-magnitude of STFT coefficients of single-speaker speech signal of speaker $i \in \{1, 2\}$ at the first microphone and at all time-frequency bins (t, f) . We omit the i index in cases where we refer to single-speaker speech in general, not in the mixture.
- $\mathbf{N} = [\mathbf{n}_t]_{t=1..T}$, with $\mathbf{n}_t = [n_{t,f}]_{f=1..F}^\top$: matrix of log-magnitude of STFT coefficients of noise signal at the first microphone and at all time-frequency bins (t, f) .

The aim of all conventional and proposed approaches in this paper is to estimate time-frequency (T-F) masks for all sources. These masks are then used to extract all speech sources using mask-based beamforming [21].

3. CONVENTIONAL APPROACHES

In this section, we introduce the previously proposed approaches that we build upon, i.e. 1) spatial clustering with CACGMM, 2) DOLPHIN which combines spatial clustering with the factorial spectral model based on GMMs, and 3) integration of NN-based approach, PIT, with spatial clustering.

3.1. Spatial clustering with Complex angular central Gaussian mixture model

Spatial clustering approaches in general work on observed spatial features $\psi_{t,f}$ and aim to obtain soft assignments of the features at all time-frequency points to individual sources. For that, a mixture model is built, where one component corresponds to one source

$$p(\psi_{t,f}) = \sum_c \pi_{c,t} p(\psi_{t,f} | d_{t,f} = c), \quad (2)$$

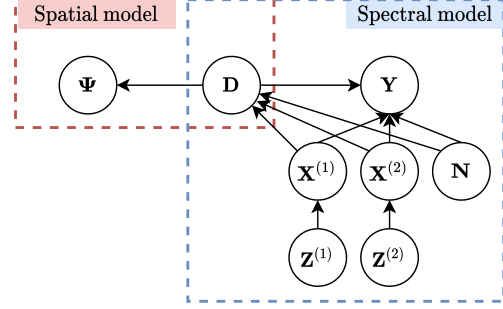


Fig. 1: Graphical model of the integration of spatial and spectral model. The model applies to both DOLPHIN and the proposed model. The difference lies in the speech models $p(\mathbf{X}^{(1)}, \mathbf{Z}^{(1)})$ and $p(\mathbf{X}^{(2)}, \mathbf{Z}^{(2)})$.

where c is the mixture component and $\pi_{c,t}$ is the mixture weight. The assignment of the features to the components is represented by a latent categorical variable $\mathbf{D} = [d_{t,f}]_{t=1..T, f=1..F}$. Through the E-M algorithm, we can infer the model and estimate the soft assignments as posterior distribution $p(\mathbf{D} | \Psi)$, which indicates the probability of the source being active for each time-frequency bin and can thus be interpreted as a time-frequency mask.

There are different choices of spatial features and their models. Here, we use normalized observation vectors as spatial features and complex angular central Gaussian mixture model (CACGMM) [1, 22]. This combination has proven to be strong in recent work [5]. The normalized observation vectors are defined as

$$\psi_{t,f} = \frac{\mathcal{Y}_{t,f}}{|\mathcal{Y}_{t,f}|}. \quad (3)$$

Complex angular central Gaussian is defined as

$$p(\psi_{t,f} | d_{t,f} = c) = \frac{(M-1)!}{2\pi^M \det \mathbf{B}_{c,f}} \frac{1}{(\psi_{t,f}^H \mathbf{B}_{c,f}^{-1} \psi_{t,f})^M}, \quad (4)$$

where M is the number of channels (dimensionality of the spatial features) and \mathbf{B} is the parameter of the distribution. For E-M algorithm in CACGMM, we refer the reader to previous work [1] and the implementation in `pb_bss`¹ [10].

3.2. DOLPHIN

Spatial clustering exploits only spatial information in the signal and does not use the spectral structure of speech, such as the periodicity caused by the fundamental frequency of the voice, or temporal continuity of the speech signal. To improve upon this, integration of the spatial model with a spectral model was proposed [7], called DOLPHIN. There, the speech signal is modeled by a GMM

$$p(\mathbf{x}_t) = \sum_{\mathbf{z}_t} p(\mathbf{z}_t) p(\mathbf{x}_t | \mathbf{z}_t) = \sum_c \pi_c^{(g)} \mathcal{N}(\mathbf{x}_t; \boldsymbol{\mu}_c^{(g)}, \text{diag}(\boldsymbol{\sigma}_c^{(g)})), \quad (5)$$

where c is the mixture component and $\boldsymbol{\mu}_c^{(g)}$, $\text{diag}(\boldsymbol{\sigma}_c^{(g)})$ are the mean and diagonal covariance matrix of the component c . The latent variable $\mathbf{Z} = [\mathbf{z}_t]_{t=1..T}$ models the assignment of time-frames to GMM mixture components.

¹https://github.com/fgnt/pb_bss

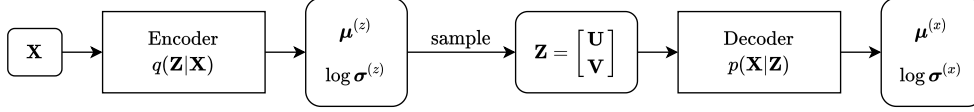


Fig. 2: Scheme of the variational autoencoder used as a spectral model. Square blocks denote neural networks and round blocks their inputs and outputs.

The speech and noise models are pre-trained on a corpus of speech and noise and combined into the factorial model of the mixed signal, using the lifted max-model, which assumes that the log-magnitude STFT coefficients of the mixture is equal to that of the dominant source at each time-frequency bin [11]

$$d_{t,f} = \operatorname{argmax}[n_{t,f}, x_{t,f}^{(1)}, x_{t,f}^{(2)}] \quad (6)$$

$$y_{t,f} = \begin{cases} n_{t,f} & \text{if } d_{t,f} = 0 \\ x_{t,f}^{(d_{t,f})} & \text{otherwise} \end{cases}, \quad (7)$$

where the latent variable $\mathbf{D} = [d_{t,f}]_{t=1..T, f=1..F}$ is the index of the dominant source at time-frequency bin (t, f) . The lifted-max model has been previously used as an approximation of log-sum interaction [11]. In contrast with the max-model, the lifted max-model explicitly uses the index of the dominant source as a latent variable \mathbf{D} . This enables the integration with a spatial model by sharing \mathbf{D} , which is used to assign T-F point to mixture components in the spatial model. The integration of the factorial model with the spatial model is depicted in Figure 1.

The inference in DOLPHIN iteratively estimates the latent variables \mathbf{D} (the time-frequency masks of the sources), latent variables \mathbf{Z} (spectral patterns of the speech sources) and parameters of the spatial model (\mathbf{B} in our case). The procedure is analogous to inference in our proposed model, described in detail in section 4.3, with the difference of the estimation of $q(\mathbf{Z})$, which has an analytical solution for DOLPHIN with GMM. For details, we refer to the original paper [7].

3.3. Combination of PIT and spatial clustering

Neural networks became widely used for speech separation with the emergence of techniques such as deep clustering [14] and PIT [15]. In these, neural networks are discriminatively trained models that map the mixed speech to the individual sources. These models are single-channel models, but they can be integrated with multi-channel processing in several ways. The naive way consists of simply using the time-frequency mask provided by the separation approach to design a beamformer [23]. A tighter integration was proposed in [9], where T-F masks produced by PIT were used as priors π for the source affiliations \mathbf{D} in the spatial clustering model (Section 3.1). This way, the spatial model is pulled towards the solution provided by PIT, which uses spectral information.

Inference in the integrated model is then analogous to inference in the spatial model by itself, with the difference of the mixture weights being fixed to the T-F masks estimated by PIT.

4. PROPOSED METHOD

In this work, we follow the integration of the spectral and spatial model used in DOLPHIN, however, we replace the GMM speech model by a variational autoencoder (VAE), which is a neural network-based generative model. We believe that VAE has better modeling power than GMM due to its non-linear transformation and possibility to model correlations over time. In this section, we

explain the construction and training of the VAE as a spectral model, revise the structure of the full integrated model, and introduce how the inference in the integrated model works.

4.1. Variational autoencoder as spectral model

Among the different models rising in the field in deep generative modeling, variational auto-encoder [17, 18] has become probably the most popular in speech processing [19, 20]. Variational auto-encoder is a deep latent variable model, where the probability density of the observation depends on the output of a highly non-linear transformation of the latent variable, where the transformation is modeled by a neural network (known as a decoder). The parameters of the model are trained by maximizing the variational lower bound of the log-likelihood of the observed data. To make the training fast, another neural network (known as encoder) is used for inference of the approximate posterior of the latent variables given observed data [18].

In our work, we use VAE to model single-speaker speech in the log-magnitude STFT domain. We use a (slightly modified) model from [20], which was inspired by [19]. In this model, part of the latent variables is trained to model speaker variability. As the speaker variability is quite an important factor for speech separation, we consider explicitly disentangling this information in the latent space useful. We denote $\mathbf{U} = [\mathbf{u}_t]_{t=1..T}$ the speaker-independent part of the latent variable, $\mathbf{V} = [\mathbf{v}_t]_{t=1..T}$ the speaker-dependent part, and $\mathbf{Z} = [\mathbf{z}_t]_{t=1..T}$ their concatenation². The model of the latent variable is as follows:

$$p(\mathbf{u}_t) = \mathcal{N}(\mathbf{u}_t; \mathbf{0}, \mathbf{I}) \quad (8)$$

$$p(\mathbf{v}_t) = \mathcal{N}(\mathbf{v}_t; \boldsymbol{\mu}_s^{(spk)}, \mathbf{I}) \quad (9)$$

$$\mathbf{z}_t = [\mathbf{u}_t, \mathbf{v}_t], \quad (10)$$

$$p(\mathbf{z}_t) = \mathcal{N}(\mathbf{z}_t; [\mathbf{0}, \boldsymbol{\mu}_s^{(spk)}], \mathbf{I}), \quad (11)$$

where $\boldsymbol{\mu}_s^{(spk)}$ is a speaker-dependent mean of speaker s . The speech signal is then modeled as a normal distribution with parameters dependent on neural network transformation of the latent variable

$$p(\mathbf{x}_t | \mathbf{Z}) = \mathcal{N}(\mathbf{x}_t; \boldsymbol{\mu}_t^{(x)}(\mathbf{Z}), \operatorname{diag}(\boldsymbol{\sigma}_t^{(x)}(\mathbf{Z}))), \quad (12)$$

where $\boldsymbol{\mu}_t^{(x)}(\cdot)$ and $\boldsymbol{\sigma}_t^{(x)}(\cdot)$ are functions modeled by the VAE decoder. The variational posterior $q(\mathbf{Z})$ inferred by the encoder is assumed to be Gaussian with parameters $\boldsymbol{\mu}_t^{(z)}$, $\boldsymbol{\sigma}_t^{(z)}$ for every time frame t

$$q(\mathbf{Z}) = \prod_t q(\mathbf{z}_t) = \prod_t \mathcal{N}(\mathbf{z}_t; \boldsymbol{\mu}_t^{(z)}, \operatorname{diag}(\boldsymbol{\sigma}_t^{(z)})). \quad (13)$$

The scheme of the described VAE is depicted in Figure 2. The parameters of the encoder, decoder, and speaker means are trained with an objective function, which combines variational lower-bound with speaker classification objective. The speaker classification objective encourages \mathbf{V} to model speaker information. The full objective is:

²To emphasize the common structure of DOLPHIN with GMM and DOLPHIN with VAE, we slightly abuse notation and use \mathbf{Z} as the latent variable in both. In Section 4, \mathbf{Z} refers to VAE latent variable as defined in Eq. (10).

$$\mathcal{J}_{\text{VAE}} = \underbrace{\mathbb{E}_{\mathbf{Z} \sim q(\mathbf{Z})} [\ln p(\mathbf{X}|\mathbf{Z}) - \text{KL}(q(\mathbf{Z})||p(\mathbf{Z}))]}_{\text{variational lower-bound}} + \underbrace{\sum_s \ln p(s|\boldsymbol{\mu}_s^{(spk)})}_{\text{speaker classification}}. \quad (14)$$

Note that the speaker classification objective differs from [19, 20], where the aim was to learn to disentangle the speaker information in an unsupervised way. Here, we assume having the speaker identities during training, and we can thus directly use a supervised objective.

4.2. Full integrated model

To compose the full integrated model, we follow the model of DOLPHIN. The factorial model of mixed signal is composed of the VAE speech model for each speaker, and a noise model. In our work, we used a simple Gaussian model for noise

$$p(\mathbf{n}_t) = \mathcal{N}(\mathbf{n}_t; \boldsymbol{\mu}^{(n)}, \text{diag}(\boldsymbol{\sigma}^{(n)})). \quad (15)$$

As there are two VAE speech models in the integrated model, for two speakers, we distinguish between the variables and distribution parameters by upper index ⁽ⁱ⁾, e.g. $\mathbf{Z}^{(1)} = [\mathbf{z}_t^{(1)}]_{t=1..T}$ for the latent variable of the first speaker in the mixture, and $\boldsymbol{\mu}_t^{(x,1)}$ for the mean of $p(\mathbf{x}_t^{(1)}|\mathbf{Z}^{(1)})$.

The interaction model is the lifted-max model introduced in Section 3.2. As for conventional DOLPHIN, the integration with the spatial model is done through the variable \mathbf{D} used as mixture component affiliation in the spatial model, and index of dominant source in the spectral model. The graphical model of the full integrated model is depicted in Figure 1.

4.3. Inference

The speech and the noise models are pre-trained in advance on non-overlapped speech and noise data. During the inference time, given an observed mixture, we want to obtain an estimate of $p(\mathbf{D}|\mathbf{Y}, \boldsymbol{\Psi})$, i.e. time-frequency masks of the individual sources. For that, we employ Variational Bayes inference, assuming the approximate posterior to factorize as follows

$$q(\mathbf{D}, \mathbf{Z}^{(1)}, \mathbf{Z}^{(2)}) = q(\mathbf{D})q(\mathbf{Z}^{(1)})q(\mathbf{Z}^{(2)}) \\ = \prod_{t,f} q(d_{t,f}) q(\mathbf{Z}^{(1)})q(\mathbf{Z}^{(2)}). \quad (16)$$

Further, the approximate posterior over the latent variables of the spectral models, $q(\mathbf{Z}^{(i)})$, are assumed to be Gaussian as in training of the VAE in Eq. (13)

$$q(\mathbf{Z}^{(1)}) = \prod_t q(\mathbf{z}_t^{(1)}) = \prod_t \mathcal{N}(\mathbf{z}_t^{(1)}; \boldsymbol{\mu}_t^{(z,1)}, \text{diag}(\boldsymbol{\sigma}_t^{(z,1)})) \quad (17)$$

$$q(\mathbf{Z}^{(2)}) = \prod_t q(\mathbf{z}_t^{(2)}) = \prod_t \mathcal{N}(\mathbf{z}_t^{(2)}; \boldsymbol{\mu}_t^{(z,2)}, \text{diag}(\boldsymbol{\sigma}_t^{(z,2)})). \quad (18)$$

The parameters of $q(\mathbf{Z}^{(1)})$, $q(\mathbf{Z}^{(2)})$ are thus $\theta^{(z,1)} = (\boldsymbol{\mu}^{(z,1)}, \boldsymbol{\sigma}^{(z,1)})$ and $\theta^{(z,2)} = (\boldsymbol{\mu}^{(z,2)}, \boldsymbol{\sigma}^{(z,2)})$.

During the inference, we need to infer 1) the approximate posterior $q(\mathbf{D})$, 2) the parameters $\theta^{(z,1)}$, $\theta^{(z,2)}$ of the approximate posterior $q(\mathbf{Z}^{(1)})q(\mathbf{Z}^{(2)})$ and 3) the parameters of the spatial model \mathbf{B} . These three inference steps are performed iteratively. In the beginning of the inference, the parameters may be initialized randomly. For faster convergence, we initialize $q(\mathbf{Z}^{(1)})q(\mathbf{Z}^{(2)})$ by forwarding the observed mixed speech through the VAE encoder.

4.3.1. Inference of $q(\mathbf{D})$

Approximate posterior distribution $q(\mathbf{D})$ represents the T-F masks of speech and noise sources. Estimation of the masks for source k should consider how well the observations at different T-F points fit under the spatial and spectral model of the source k . For the spatial model, it is the k -th component of the mixture. For the spectral model, it is either the noise model or VAE speech model with the current estimate of its latent variables $q(\mathbf{Z}^{(i)})$.

Formally, the update follows Variational Bayes, we thus maximize the variational lower-bound. This has an analytical solution (following Eq. 10.9 in [24])

$$\ln q(d_{t,f} = k) = \overbrace{\ln p(\boldsymbol{\psi}_{t,f} | d_{t,f} = k)}^{\text{spatial model}} + \underbrace{\mathbb{E}_{\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)} \sim q(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})} [\ln p(y_{t,f}, d_{t,f} = k | \mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})]}_{\text{spectral model}} + C \quad (19)$$

where C is constant w.r.t $d_{t,f}$. We obtain the constant C by ensuring that $\sum_k q(d_{t,f} = k) = 1$. The expectation is approximated by sampling N_z samples from $q(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})$. The value inside the expectation expresses how likely the source k is dominant and has the value of the observed mixed speech $y_{t,f}$. For example of $k = 1$, this can be expressed as

$$\ln p(y_{t,f}, d_{t,f} = 1 | \mathbf{Z}^{(1)}, \mathbf{Z}^{(2)}) = \\ \ln p(x_{t,f}^{(1)} = y_{t,f} | \mathbf{Z}^{(1)}) + \ln p(x_{t,f}^{(2)} < y_{t,f} | \mathbf{Z}^{(2)}) \\ + \ln p(n_{t,f} < y_{t,f}) = \\ = \ln \mathcal{N}(y_{t,f}; \boldsymbol{\mu}_{t,f}^{(x,1)}, \boldsymbol{\sigma}_{t,f}^{(x,1)}) + \ln \Phi(y_{t,f}; \boldsymbol{\mu}_{t,f}^{(x,2)}, \boldsymbol{\sigma}_{t,f}^{(x,2)}) \\ + \ln \Phi(y_{t,f}; \boldsymbol{\mu}_{t,f}^{(n)}, \boldsymbol{\sigma}_{t,f}^{(n)}), \quad (20)$$

where $\Phi(\cdot)$ is the cumulative normal distribution. This solution can be also found in previous literature on lifted-max model [7, 11].

4.3.2. Inference of $q(\mathbf{Z}^{(1)})q(\mathbf{Z}^{(2)})$

Given the current estimate of the masks $q(\mathbf{D})$, we estimate the spectral characteristics of the speech of the two speakers, represented as latent variables of the VAE speech models $\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)}$. These should be optimized in such a way that they correspond to the current estimate of $q(\mathbf{D})$ (source estimated as likely dominant should be likely dominant) and also represent proper speech signal (they are close to the prior distribution of $\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)}$).

Formally, we optimize the variational lower bound. In this case, the optimal value for the parameters of $q(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})$ cannot be found analytically, we thus employ a few steps of gradient ascent optimization of the lower-bound (following Eq. 10.6 in [24])

$$\mathbb{E}_{\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)} \sim q(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})} \mathbb{E}_{\mathbf{D} \sim q(\mathbf{D})} [\ln p(\mathbf{Y}, \mathbf{D} | \mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})] - \\ \lambda \text{KL}(q(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)}) || p(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})), \quad (21)$$

where λ is a hyper-parameter weighting two parts of the objective. The theoretically correct value of λ is 1, it may however help to improve the inference to choose a different value. The expectation over \mathbf{Z} is approximated by Monte-Carlo samples. The expectation over \mathbf{D} can be further factorized into individual T-F points. The KL divergence can be computed analytically as both $q(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})$ and $p(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})$ are Gaussian distributed. The speaker means $\boldsymbol{\mu}_s^{(spk)}$, which are part of $p(\mathbf{Z}^{(i)})$ are set to their optimal value w.r.t. the lower-bound, which is an average of $\boldsymbol{\mu}_t^{(z,i)}$ over time. To compute gradient w.r.t. the parameters of $q(\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)})$ that we sample from, we employ the reparametrization trick [17].

4.3.3. Inference of spatial model parameters

Finally, the parameters of the spatial model can be updated using an M-step of the E-M algorithm following the inference procedure in the spatial model itself and using the current estimate of $q(\mathbf{D})$. For details on M-step on CACGMM, we refer to [1].

4.3.4. Inference summary

For clarity, we summarize the inference process in the following algorithm. Details on the initialization steps are given in experimental settings in Section 5.2.

Algorithm 1 Inference in the proposed integrated model.

1. Initialize $q(\mathbf{D})$.
 2. Initialize parameters $\theta^{z,1}, \theta^{z,2}$ of $q(\mathbf{Z}^{(1)}), q(\mathbf{Z}^{(2)})$.
 3. Iterate the following for N_{iter} iterations.
 - (a) Update $q(\mathbf{D})$ using Eq. (19), (20).
 - (b) Update parameters $\theta^{z,1}, \theta^{z,2}$ of $q(\mathbf{Z}^{(1)}), q(\mathbf{Z}^{(2)})$ by iterating the following for N_{up} iterations.
 - i. Sample $\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)}$ from $q(\mathbf{Z}^{(1)}), q(\mathbf{Z}^{(2)})$.
 - ii. Forward $\mathbf{Z}^{(1)}, \mathbf{Z}^{(2)}$ through VAE decoder.
 - iii. Compute gradients of Eq. (21) w.r.t. $\theta^{z,1}, \theta^{z,2}$ by back-propagating through the VAE decoder.
 - iv. Update $\theta^{z,1}, \theta^{z,2}$ with gradient ascent.
 - (c) Update the parameters \mathbf{B} of the spatial model [1].
-

5. EXPERIMENTS

5.1. Datasets

As a dataset for our experiments, we used the spatialized version of WSJ0-2mix³ sampled at 8kHz, which was introduced in [25]. The original WSJ0-2mix dataset [14] was created by mixing utterances of WSJ0 [26] with SNR of one source to the other of -5-5 dB. It consists of training, validation, and evaluation sets of 20000, 5000, and 3000 utterances, respectively. The evaluation set contains a disjoint set of speakers from training and validation. The spatialized version was created by using artificially generated room impulse responses with 8 microphones, randomized setting of microphone array size (15-25 cm), T60s (0.2 - 0.6 s), and speaker distance to the array (1.3m with 0.4m standard deviation). In our work, we additionally added a white diffuse noise to the data, to analyze the possibility of the methods to adapt to different noise conditions. We created 5 different conditions with SNRs of the mixed speech to the noise of 15-20, 10-15, 5-10, 0-5, and -5-0 dB.

5.2. Settings

Signal representation All experiments use Short-time Fourier transform (STFT) with a window of 512 samples (64 ms), a shift of 128 samples (16ms) and a Hanning window. For VAE and GMM, we use the log-magnitude of STFT coefficients, for PIT training we use the magnitude of the STFT coefficients.

GMM settings We use 256-component GMM. The GMM was trained by progressively splitting the mixture by doubling the number of components. We used 1-4-4-10-10-10-10-10 E-M

iterations for each stage with 1-2-4-8-16-32-64-128-256 components.

PIT architecture and training For PIT, we used an available implementation⁴. The architecture consists of 3 BLSTM layers with hidden size 896 and ReLU activation functions. The training uses Adam optimizer with learning rate 1e-3, momentum 0.9, weight decay 1e-5, dropout rate 0.5. We train the PIT network on mixtures with 15-20 dB noise and use three output layers (for two speakers and noise). The outputs are T-F masks and the objective function is the phase-sensitive mean-squared error of the masked mixture [27].

VAE architecture and training The architecture of the encoder and decoder roughly follows [20]. The encoder has 5 1-d convolutional layers and 4 feed-forward layers with hidden dimensionality of 512. The kernel size in all convolutional layers is 3, the stride is 1-1-2-1-1. ReLU activations are used at all layers. Every layer which does not change temporal resolution has an additional residual connection. The final feed-forward layer outputs the parameters of $q(\mathbf{Z})$: $\mu^{(u)}, \mu^{(v)}$ and $\log \sigma^{(u)}, \log \sigma^{(v)}$, which all have a dimension of 20. The dimension of the input to the encoder is 257. The decoder has the same architecture as the encoder, but all convolutions are replaced with transposed convolutions. The dimension of the input to the decoder is 40 (dimension of \mathbf{u} + dimension of \mathbf{v}). The output of the decoder are the parameters of $p(\mathbf{X}|\mathbf{Z})$: $\mu^{(x)}, \log \sigma^{(x)}$, which both have a dimension of 257. The VAE was trained for 100 epochs with Adam optimizer, learning rate 1e-4, and gradient norm clipping 10.

Inference We initialize both $q(\mathbf{Z}^{(1)})$ and $q(\mathbf{Z}^{(2)})$ by forwarding the mixed speech through encoder of the VAE, and $q(\mathbf{D})$ randomly from uniform distribution with subsequent normalization. The inference runs for $N_{\text{iter}} = 100$ iterations (this is common for all tested methods). For updating parameters of $q(\mathbf{Z}^{(1)})q(\mathbf{Z}^{(2)})$, we used $N_{\text{up}} = 5$ updates in each iteration with learning rate 1e-3. We set the weight of the KL divergence part of the objective in Eq. (21) to 10, which was optimal for validation set. The expectations over $q(\mathbf{Z}^{(1)})q(\mathbf{Z}^{(2)})$ are approximated using one sample.

Beamforming To obtain the separated sources, we use a mask-based MVDR beamformer [21] as defined in [28, 29].

5.3. Results

The results of the experiments are summarized in Table 1. We report signal-to-distortion-ratio⁵ [30] (SDR) improvements over the original mixtures. The SDR values of the mixtures for each condition are shown in the header of the table. All experiments use mask-based beamforming to obtain the separated signals. The tested methods are *spatial* clustering (Section 3.1), *spatial+GMM* (Section 3.2), *spatial+PIT* (Section 3.3) and proposed *spatial+VAE* (Section 4). In addition, we also show results of using the masks provided by PIT directly to estimate beamformer without using spatial clustering, denoted as *PIT-bf*, which has been widely used [23, 31]. Note that PIT network is trained on noise condition with SNR 15-20 dB, in order to analyze the effect of mismatched noise conditions. Lastly, *Ideal binary mask* shows an oracle result.

First, let us focus on the comparison between between *spatial*, *spatial+GMM* and *spatial+VAE* (rows 1,2,4). The combination with GMM is consistently better than spatial clustering by itself which confirms the advantage of using the spectral model, as proposed in DOLPHIN [7]. Replacing GMM with VAE spectral model (as pro-

³<http://www.merl.com/demos/deep-clustering>

⁴<https://github.com/funcwj/uPIT-for-speech-separation>

⁵https://github.com/craffel/mir_eval

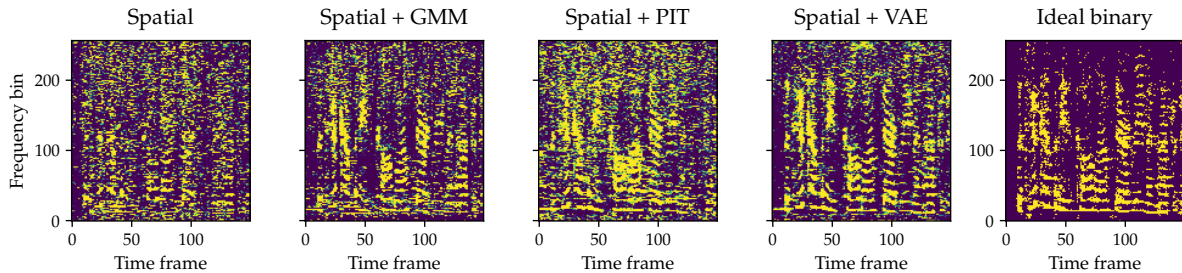


Fig. 3: Time-frequency masks obtained by the tested methods, and oracle ideal binary mask. The color scale in all masks ranges from 0 to 1.

posed here) significantly improves the performance, showing that VAE is a stronger model of the speech signal.

Second, we compare *spatial+PIT* with *spatial+VAE* on matched condition (rows 3,4, column 1). The results show comparable performance of both methods. This shows the validity of our approach and the potential of using deep generative modeling for speech separation, in contrast with discriminatively trained networks, which are today primarily used. Interestingly, no part of our model was trained on mixed speech, but nevertheless, it can achieve the performance of PIT which was trained directly for separation. Note that the commonly used beamforming based on PIT masks (row 5) under-performs the integrated *spatial+PIT* model by a significant margin.

Finally, we explore the possibility of adapting the *spatial+VAE* model to a different noise condition. For every condition, we estimated the noise model parameters ($\mu^{(n)}$, $\sigma^{(n)}$) on the noise signals of corresponding SNRs. The rest of the model remains fixed. We can see the effect by comparing with *spatial+PIT* (rows 3,4), which is not adapted. As the mismatch grows, the margin between *PIT* and *VAE* increases. This shows that the model is successfully adapting by re-training the noise model. In practice, we could for example estimate the noise model on non-speech parts of recordings and possibly vary the complexity of the noise model based on how much data we have available.

Figure 3 shows the time-frequency masks obtained with the spatial model and combination of spatial and spectral models with GMM, PIT, and VAE. We also show the ideal binary mask. Looking at the figure, we observe that the combination of spatial and spectral models lead to masks that better represent the harmonic structure of speech. Comparing the masks obtained with GMM and with VAE, we see the better modeling capability of the VAE, especially at lower frequencies. Moreover, compared to *spatial+PIT* masks, the VAE mask is less noisy and much closer to the ideal binary mask. This improved mask estimation performance translates in better speech enhancement performance after beamforming.

6. RELATED WORKS

The concept of using a variational auto-encoder as the speech model for speech enhancement or separation has been explored in several previous works. For single-channel speech enhancement, authors in [32–34] use VAE speech model and low-rank model of noise together with Gibbs sampling to estimate clean speech spectra from noisy spectrogram. This was further extended to multi-channel scenario [35–37] by integrating the models into multi-channel nonnegative matrix factorization (MNMF) framework.

Similar methods have been explored for multi-channel speech separation in [38, 39]. The main difference to our proposal lies in the multi-channel modeling — authors in [38, 39] extend the MNMF

Table 1: Signal-to-distortion ratio gains (dB) of the tested methods on the spatialized WSJ0-2mix corpus with different levels of added noise.

noise SNR [dB]	15-20	10-15	5-10	0-5	-5-0
mixture SDR [dB]	-0.5	-0.9	-1.8	-3.8	-6.9
Δ SDR [dB]					
1 Spatial	11.6	10.0	8.1	6.3	5.5
2 + GMM (DOLPHIN)	12.1	10.6	8.8	7.3	6.8
3 + PIT	12.6	11.2	9.6	7.7	7.0
4 + VAE (Proposed)	12.7	11.4	10.0	8.9	8.1
5 PIT-bf	10.8	9.7	8.2	6.2	3.3
6 Ideal binary mask	13.5	12.2	10.9	10.1	10.1

framework, while we extend DOLPHIN, which is based on spatial clustering. The works are thus complementary to each other and pursue similar ideas in different settings, which both have their pros and cons. The MNMF+VAE approach uses linear spectra which can be advantageous because the mixture can be accurately modeled by the sum of component signals. Further, due to their use of Multi-channel Wiener filter (MWF), the optimization and estimation use the same objective. On the other hand, the use of CACGMM and spatial clustering has been shown very effective for far-field ASR and is computationally much less demanding than MWF. The MNMF+VAE based methods have been so far tested on a relatively small dataset with a limited number of speakers, direct comparison is thus difficult.

Several ideas from the previous works could be also incorporated into our proposal in future extensions, such as the low-rank model of noise, re-using the VAE encoder for faster inference [34, 39] or more flexible speech modeling with flows [37].

7. CONCLUSION

In this paper, we combined a VAE spectral model of the speech signal with CACGMM spatial clustering for multi-channel speech separation. This extends previously proposed DOLPHIN method which uses GMM as the spectral model. We showed that VAE, as a more flexible model, outperforms GMM on all tested conditions. Furthermore, we showed comparable performance to PIT integrated with spatial clustering and the possibility of adaptation to different noise conditions enabled by an explicit noise model. Future work includes investigations on datasets with more complex noisy conditions and adaptation of the noise models within the inference step. The method could be also extended on the model side, e.g. with more structured latent space or a different choice of deep generative models. The code for the method is released on Github [40].

8. REFERENCES

- [1] N. Ito, S. Araki, and T. Nakatani, "Complex angular central Gaussian mixture model for directional statistics in mask-based microphone array signal processing," in *2016 24th European Signal Processing Conference (EUSIPCO)*, 2016, pp. 1153–1157.
- [2] H. Sawada, S. Araki, and S. Makino, "Underdetermined convolutive blind source separation via frequency bin-wise clustering and permutation alignment," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 3, pp. 516–527, 2011.
- [3] D. H. T. Vu and R. Haeb-Umbach, "Blind speech separation employing directional statistics in an expectation maximization framework," in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2010, pp. 241–244.
- [4] M. I. Mandel, R. J. Weiss, and D. P. W. Ellis, "Model-based expectation maximization source separation and localization," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 2, pp. 382–394, Feb. 2010.
- [5] C. Boeddeker, J. Heitkaemper, J. Schmalenstroer, L. Drude, J. Heymann, and R. Haeb-Umbach, "Front-end processing for the CHiME-5 dinner party scenario," in *CHiME5 Workshop, Hyderabad, India*, 2018.
- [6] S. Watanabe, M. Mandel, J. Barker, and E. Vincent, "Chime-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings," *arXiv preprint arXiv:2004.09249*, 2020.
- [7] T. Nakatani, S. Araki, T. Yoshioka, M. Delcroix, and M. Fujimoto, "Dominance based integration of spatial and spectral features for speech enhancement," *IEEE transactions on audio, speech, and language processing*, vol. 21, no. 12, pp. 2516–2531, 2013.
- [8] M. J. Reyes-Gomez, B. Raj, and D. R. W. Ellis, "Multi-channel source separation by factorial HMMs," in *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03)*, 2003, vol. 1, pp. I–I.
- [9] T. Nakatani, N. Ito, T. Higuchi, S. Araki, and K. Kinoshita, "Integrating DNN-based and spatial clustering-based mask estimation for robust mvdr beamforming," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 286–290.
- [10] L. Drude and R. Haeb-Umbach, "Integration of neural networks and probabilistic spatial models for acoustic blind source separation," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 4, pp. 815–826, 2019.
- [11] S. J. Rennie and P. A. Hershey, J. and Olsen, "Single-channel multitalker speech recognition," *IEEE Signal Processing Magazine*, vol. 27, no. 6, pp. 66–80, 2010.
- [12] M. Delcroix, K. Kinoshita, T. Nakatani, S. Araki, A. Ogawa, T. Hori, S. Watanabe, M. Fujimoto, T. Yoshioka, T. Oba, et al., "Speech recognition in the presence of highly non-stationary noise based on spatial, spectral and temporal speech/noise modeling combined with dynamic variance adaptation," in *Machine Listening in Multisource Environments*, 2011.
- [13] J. Hershey, S. Rennie, P. Olsen, and T. Kristjansson, "Superhuman multi-talker speech recognition: A graphical modeling approach," *Computer Speech & Language*, vol. 24, no. 1, pp. 45–66, 2010.
- [14] J. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [15] D. Yu, M. Kolbæk, Z. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 241–245.
- [16] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [17] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [18] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," *arXiv preprint arXiv:1401.4082*, 2014.
- [19] W. Hsu, Y. Zhang, and J. Glass, "Unsupervised learning of disentangled and interpretable representations from sequential data," in *Advances in neural information processing systems*, 2017, pp. 1878–1889.
- [20] T. Gburrek, T. Glarner, J. Ebberts, R. Haeb-Umbach, and P. Wagner, "Unsupervised learning of a disentangled speech representation for voice conversion," in *Proceedings of the 10 Speech Synthesis Workshop (SSW10)*, 2019.
- [21] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 196–200.
- [22] D. E. Tyler, "Statistical analysis for the angular central Gaussian distribution on the sphere," *Biometrika*, vol. 74, no. 3, pp. 579–589, 1987.
- [23] T. Higuchi, K. Kinoshita, M. Delcroix, K. Zmolíková, and T. Nakatani, "Deep clustering-based beamforming for separation with unknown number of sources," in *Interspeech*, 2017, pp. 1183–1187.
- [24] C. M. Bishop, *Pattern recognition and machine learning*, Springer, 2006.
- [25] Z. Wang, J. Le Roux, and J. Hershey, "Multi-channel deep clustering: Discriminative spectral and spatial embeddings for speaker-independent speech separation," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 1–5.
- [26] D. B. Paul and J. Baker, "The design for the Wall Street Journal-based CSR corpus," in *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*, 1992.
- [27] H. Erdogan, J. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 708–712.
- [28] M. Souden, J. Benesty, and S. Affes, "On optimal frequency-domain multichannel linear filtering for noise reduction," *IEEE Transactions on audio, speech, and language processing*, vol. 18, no. 2, pp. 260–276, 2010.

- [29] H. Erdogan, J. Hershey, S. Watanabe, M. Mandel, and J. Le Roux, "Improved MVDR beamforming using single-channel mask prediction networks.," in *Interspeech*, 2016, pp. 1981–1985.
- [30] C. Raffel, B. McFee, E. Humphrey, J. Salamon, O. Nieto, D. Liang, D. Ellis, and C. 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.
- [31] T. Yoshioka, H. Erdogan, Z. Chen, X. Xiao, and F. Alleva, "Recognizing overlapped speech in meetings: A multichannel separation approach using neural networks," *arXiv preprint arXiv:1810.03655*, 2018.
- [32] 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 *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 716–720.
- [33] S. Leglaive, L. Girin, and R. Horaud, "A variance modeling framework based on variational autoencoders for speech enhancement," in *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2018, pp. 1–6.
- [34] M. Pariente, A. Deleforge, and E. Vincent, "A statistically principled and computationally efficient approach to speech enhancement using variational autoencoders," *arXiv preprint arXiv:1905.01209*, 2019.
- [35] K. Sekiguchi, Y. Bando, K. Yoshii, and T. Kawahara, "Bayesian multichannel speech enhancement with a deep speech prior," in *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (AP-SIPA ASC)*. IEEE, 2018, pp. 1233–1239.
- [36] S. Leglaive, L. Girin, and R. Horaud, "Semi-supervised multichannel speech enhancement with variational autoencoders and non-negative matrix factorization," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 101–105.
- [37] A. A. Nugraha, K. Sekiguchi, and K. Yoshii, "A flow-based deep latent variable model for speech spectrogram modeling and enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1104–1117, 2020.
- [38] S. Seki, H. Kameoka, L. Li, T. Toda, and K. Takeda, "Generalized multichannel variational autoencoder for underdetermined source separation," in *2019 27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019, pp. 1–5.
- [39] L. Li, H. Kameoka, and S. Makino, "Fast MVAE: Joint separation and classification of mixed sources based on multichannel variational autoencoder with auxiliary classifier," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 546–550.
- [40] Katerina Zmolikova, "VAE-DOLPHIN," 2020, https://github.com/BUTSpeechFIT/vae_dolphin.