



An Initialization Scheme for Meeting Separation with Spatial Mixture Models

Christoph Boeddeker, Tobias Cord-Landwehr, Thilo von Neumann, Reinhold Haeb-Umbach

Paderborn University, Germany

{boeddeker, cord, vonneumann, haeb}@nt.upb.de

Abstract

Spatial mixture model (SMM) supported acoustic beamforming has been extensively used for the separation of simultaneously active speakers. However, it has hardly been considered for the separation of meeting data, that are characterized by long recordings and only partially overlapping speech. In this contribution, we show that the fact that often only a single speaker is active can be utilized for a clever initialization of an SMM that employs time-varying class priors. In experiments on LibriCSS we show that the proposed initialization scheme achieves a significantly lower Word Error Rate (WER) on a downstream speech recognition task than a random initialization of the class probabilities by drawing from a Dirichlet distribution. With the only requirement that the number of speakers has to be known, we obtain a WER of 5.9%, which is comparable to the best reported WER on this data set. Furthermore, the estimated speaker activity from the mixture model serves as a diarization based on spatial information.

Index Terms: spatial mixture model, meeting separation, beamforming, LibriCSS

1. Introduction

Acoustic beamforming guided by Spatial Mixture Models (SMMs) has been widely used for blind speech separation given a multi-channel input [1–6]. Building on the sparsity and w-disjoint orthogonality of speech in the Short-Time Fourier Transform (STFT) domain, the vector of channels of a microphone array, which contain the speech of multiple simultaneously active speakers, can be modeled by an SMM. Its parameters are estimated with the Expectation Maximization (EM) algorithm, of which both batch and frame-online variants exist. The resulting class posterior probabilities are used to estimate spatial covariance matrices, for each speaker in the mixture. From those, the coefficients of statistically optimal acoustic beamformers, such as the weighted Minimum Power Distortionless Response (wMPDR) beamformer [7], can be computed, that extract the signals of the individual speakers in the mixture.

Despite being very effective, SMMs have hardly been used on meeting data, with [4, 5] being notable exceptions. With meeting data we here denote comparatively long (1 minute) recordings of conversations among a fixed number of participants, which contain segments of no speech activity, of a single speaker talking or multiple speakers being active concurrently. Meeting separation is concerned with mapping such speech data to output streams in a way such that there is no overlapping speech on those streams. The prevalent approach to this so-called Continuous Speech Separation (CSS) are neural networks [8–12]. However, an SMM has the crucial advantage that its parameters are estimated in an unsupervised fashion, while the supervised neural network training needs both the separate utterances of the individual speakers as targets and their artificial superposition as input to the network. For example, in the

CHiME-5 and CHiME-6 challenges [13, 14], SMMs were the prevailing approach to source separation, because ground truth single-speaker utterances were not available for training.

However, parameter estimation in an SMMs is notoriously sensitive to initialization. This is particularly crucial because multiple SMMs have to be learnt, namely one for each frequency bin. Depending on the initialization, the EM algorithm may settle on an unfavorable local optimum of the likelihood function where the performance of the resulting beamformer is unsatisfactory.

Initialization of SMMs can be achieved by either giving initial values to its parameters or to its latent variables, i.e. the posterior class affiliation probabilities. If prior information about the spatial or temporal properties of the recording is available, it can be used to initialize the parameters of the mixture model, as was done, e.g. in [4] for spatial and [6] for temporal initialization. When initializing the class posteriors instead, the most uninformative initialization of the class posterior probabilities is to draw their initial values at random from a Dirichlet distribution. But often, some prior knowledge is available to improve over this. For example, it is often valid to assume that there are noise-only periods at the beginning and end of a recording [15]. The initial values of the posterior of the noise class can then be set accordingly.

But in case of meeting data, where often only a single speaker is active, one can do better: We suggest splitting the meeting data into short segments. Assuming that each segment only contains a single speaker, a single-component probability density function (PDF) per frequency is fit to the data of each segment. The resulting models are clustered to $K + 1$ clusters, one cluster for each of the K speakers and one for the noise. From each of those clusters a frequency-independent class posterior is computed, which serves as initialization of the SMM.

We conducted experiments on the LibriCSS database [8], which consists of partially overlapping utterances of the LibriSpeech corpus [16] that have been played back in a meeting room and recorded with a 7-element microphone array. Using the above sketched initialization scheme we were able to achieve a Word Error Rate (WER) of 5.9% with a downstream recognizer, which is superior to the 13.3% obtained by random initialization and which comes close to the 5.1% WER obtained by an oracle initialization.

2. Signal model

Let $\mathbf{y}_{t,f} \in \mathbb{C}^M$ denote the vector of the M microphone signals in the STFT domain, where the indices t and f are the time-frame and frequency-bin index, respectively. The observation \mathbf{y} contains the sum of K speaker signals $s_{t,f,k}$, multiplied by their acoustic transfer function $\mathbf{h}_{f,k}$, and noise $\mathbf{n}_{t,f}$. Those signals are assumed to have the same length as \mathbf{y} , i.e. the speaker signal contains zeros if a speaker is inactive. While, in theory, all speakers could be active simultaneously and for the complete duration of the meeting, databases of real conversations [17, 13]

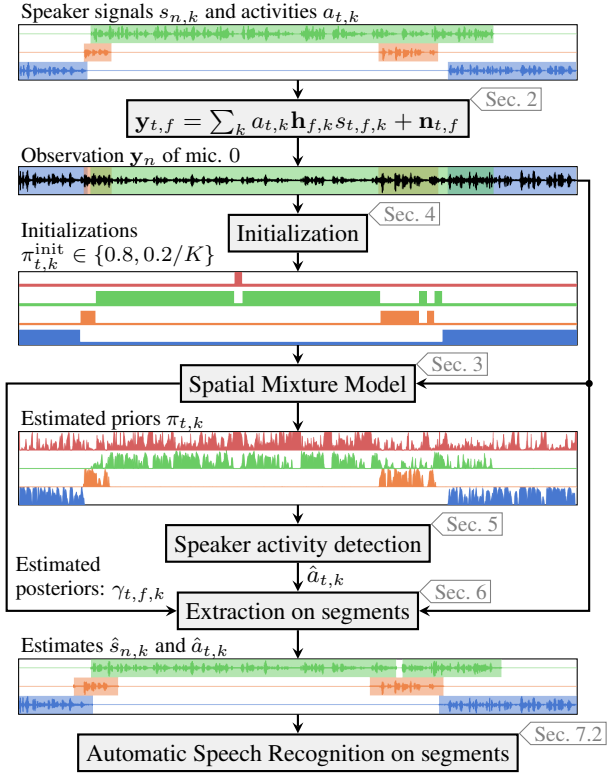


Figure 1: System overview, with images from one LibriCSS example. The x-axis represents time (sample n or frame t) and the y-axis the values and speaker index. For visualization the permutation of the initialization was removed.

show only partial activity of each speaker. This partial activity can be modelled by an additional speaker-dependent activity $a_{t,k} \in \{0, 1\}$ without loss of generality, resulting in the model

$$\mathbf{y}_{t,f} = \sum_{k=1}^K a_{t,k} \mathbf{h}_{f,k} s_{t,f,k} + \mathbf{n}_{t,f} \quad (1)$$

for the observation signal.

3. Spatial Mixture Model

To model the multichannel observation, we use an SMM, in particular the complex Angular Central Gaussian Mixture Model (cACGMM) [2]:

$$p(\mathbf{z}_{t,f}) = \sum_{k=1}^{K+1} \pi_{t,k} \mathcal{A}(\mathbf{z}_{t,f}; \mathbf{B}_{f,k}) \quad \text{with } \mathbf{z}_{t,f} = \frac{\mathbf{y}_{t,f}}{\|\mathbf{y}_{t,f}\|} \quad (2)$$

where $\pi_{t,k}$ is a time-varying a priori probability as proposed in [3]. Using time dependent priors, instead of frequency dependent ones, better fits a meeting scenario because they can reflect the activity of the speakers over time. The number of classes $K+1$ is set to the number of speakers K plus one for the noise. $\mathcal{A}(\cdot)$ is the complex Angular Central Gaussian (cACG) distribution

$$\mathcal{A}(\mathbf{z}_{t,f}; \mathbf{B}_{f,k}) = \frac{(M-1)!}{2\pi^M \det(\mathbf{B}_{f,k})} \left(\mathbf{z}_{t,f}^H \mathbf{B}_{f,k}^{-1} \mathbf{z}_{t,f} \right)^{-M} \quad (3)$$

and $\mathbf{B}_{f,k}$ is a hermitian parameter matrix of this distribution. $(\cdot)^H$ and $(\cdot)^{-1}$ denote the conjugate transpose operation and matrix inversion, respectively.

The parameters of the SMM are trained with the EM algorithm, which consists of alternating E- and M-steps. For the cACGMM the update equations are [2, 3]:

E-step:

$$\gamma_{t,f,k} = \frac{\pi_{t,k} \det(\mathbf{B}_{f,k})^{-1} (\mathbf{z}_{t,f}^H \mathbf{B}_{f,k}^{-1} \mathbf{z}_{t,f})^{-M}}{\sum_{\kappa=1}^{K+1} \pi_{t,\kappa} \det(\mathbf{B}_{f,\kappa})^{-1} (\mathbf{z}_{t,f}^H \mathbf{B}_{f,\kappa}^{-1} \mathbf{z}_{t,f})^{-M}} \quad (4)$$

M-step:

$$\pi_{t,k} = \frac{1}{F} \sum_{f=1}^F \gamma_{t,f,k} \quad (5)$$

$$\mathbf{B}_{f,k} = \frac{M}{\sum_{t=1}^T \gamma_{t,f,k}} \sum_{t=1}^T \gamma_{t,f,k} \frac{\mathbf{z}_{t,f}^H \mathbf{z}_{t,f}}{\mathbf{z}_{t,f}^H \mathbf{B}_{f,k}^{-1} \mathbf{z}_{t,f}}, \quad (6)$$

where $\gamma_{t,f,k}$ is the a posteriori probability that the class k is active at a time-frequency bin t, f . While the time-varying a priori probability [3] was introduced to prevent the permutation problem between frequencies, it was observed [15] that a permutation solver after each EM step is nevertheless beneficial. Here, we use a similarity-based permutation solver. It is based on the work by Tran Vu [1] and is similar to [18].

From now on, we will use the terms prior and posterior as abbreviation for the a priori probability $\pi_{t,k}$ and a posteriori probability $\gamma_{t,f,k}$, respectively.

4. Initialization

The EM algorithm converges to a local optimum of the objective function and is thus susceptible to its initialization. Either the model parameters $(\pi_{t,k}, \mathbf{B}_{f,k})$ or the posteriors $(\gamma_{t,f,k})$ can be initialized. Note, however, that there is no closed-form update equation for the parameter matrix $\mathbf{B}_{f,k}$ of the cACG, see Eq. (6). Thus, an initialization with the posteriors additionally requires an initial value for $\mathbf{B}_{f,k}$. Here, the identity matrix is chosen for the initialization of $\mathbf{B}_{f,k}$. In the following, several strategies for the initialization of the mixture model are investigated.

4.1. Oracle initialization

For the oracle initialization, we follow [19] by initializing the posteriors with the true speaker activities $a_{t,k}$, see Eq. (1), as

$$\gamma_{t,f,k}^{\text{init}} = \frac{a_{t,k}}{\sum_{k=1}^{K+1} a_{t,k}}, \quad (7)$$

The noise class is set to be always active, i.e. $a_{t,K+1} = 1$.

Although this initialization uses oracle information, the SMM could forget the initialization while iterating, which makes this scheme different from the guided source separation proposed in [19].

4.2. Draws from Dirichlet distribution

The most straightforward initialization of the posteriors is to draw their values from the Dirichlet distribution as in [20]:

$$\gamma_{t,f,1}^{\text{init}}, \dots, \gamma_{t,f,K+1}^{\text{init}} \sim \text{Dirichlet}(\boldsymbol{\alpha}) \quad (8)$$

where the $(K+1)$ -dimensional parameter vector $\boldsymbol{\alpha}$ of the Dirichlet distribution is set to one for each class. Instead of sampling the initialization values independently for each frequency, another option is to tie them across the frequency axis, i.e. to use the same set of drawn class posteriors for all frequencies. This frequency-tied initialization can be interpreted

as initializing the prior $\pi_{t,k}$ instead of the posterior. Since $\mathbf{B}_{f,k}$ is initialized as the identity matrix, Eq. (4) simplifies and the first E-step becomes:

$$\pi_{t,1}^{\text{init}}, \dots, \pi_{t,K+1}^{\text{init}} \sim \text{Dirichlet}(\boldsymbol{\alpha}) \quad (9)$$

$$\gamma_{t,f,k}^{\text{init}} = \pi_{t,k}^{\text{init}} \quad \forall f. \quad (10)$$

4.3. Proposed cACG clustering-based initialization

Using the same value for all elements of the parameter vector $\boldsymbol{\alpha}$ of the Dirichlet distribution leads to an initialization that assumes that all speakers are active for roughly the same amount of time.

However, one can do better. In a typical meeting, most speakers are at least once the sole active speaker. For each of these single-speaker regions, a single cACG is sufficient to model the observations. We use this property to find the following clustering-based initialization. First, the meeting is split into short segments. Then, the parameters of a cACG are estimated for each of these segments. If two segments, ℓ and l , have the same active speaker, the estimated parameter matrices $\mathbf{B}_{f,\ell}$, $\mathbf{B}_{f,l}$ of these segments will exhibit a low distance to each other. To measure their distance, the correlation matrix distance [21] is used:

$$d(\mathbf{B}_{f,\ell}, \mathbf{B}_{f,l}) = 1 - \frac{\text{tr}(\mathbf{B}_{f,\ell}, \mathbf{B}_{f,l})}{\|\mathbf{B}_{f,\ell}\|_{\text{Fro}} \|\mathbf{B}_{f,l}\|_{\text{Fro}}} \quad (11)$$

where $\|\cdot\|_{\text{Fro}}$ and $\text{tr}(\cdot)$ are the Frobenius norm and trace operator, respectively. Using this distance metric, the pair-wise distance matrix \mathbf{D} with the entries $D_{\ell,l}$

$$D_{\ell,l} = \frac{1}{F} \sum_{f=1}^F d(\mathbf{B}_{f,\ell}, \mathbf{B}_{f,l}) \quad (12)$$

is constructed. Based on this distance matrix, all similar segments are grouped together to obtain an initialization for the $K + 1$ classes. This grouping is achieved by complete linkage clustering [22], a variant of agglomerative hierarchical clustering (AHC). The clustering starts with each segment forming its own cluster. Next, two clusters are merged if the maximum distance between entries of the one and the other cluster is the smallest compared to all other potential cluster merges. This process is repeated until $K + 1$ clusters remain. We used the complete linkage clustering because it keeps the cluster diameter small, so that it is unlikely that it groups two speakers in one cluster. We then initialize

$$\pi_{t,k}^{\text{init}} = \begin{cases} 0.8 & \text{if segment } \ell \text{ of } t \text{ belongs to } k \\ 0.2/K & \text{otherwise.} \end{cases} \quad (13)$$

Fig. 1 illustrates an example for this initialization, where 30 consecutive frames are used to estimate $\mathbf{B}_{f,\ell}$.

5. Speaker activity detection

After the convergence of the EM algorithm, speaker activity is detected through the learned class prior probabilities $\pi_{t,k}$. This process is divided into three steps: noise class identification, class fusion, and speech segmentation. Note that this delivers the diarization information of which speaker is active and when.

All three steps are done with thresholded dilation and/or erosion operations. Both are well known techniques from image processing that apply as follows to the 1-dimensional case of speech signals: For a dilation, a sliding window is moved over

the signal with a shift of one, and the maximum value inside the window is the new value for its center sample. In erosion, we do the same, just with the minimum operation.

5.1. Noise identification

We observed that a class prior $\pi_{t,k}$ is close to zero when the corresponding speaker is inactive, and the noise is active all the time (See estimates in Fig. 1). We can estimate the active time of a class by smoothing the class prior $\pi_{t,k}$ with a dilation followed by an erosion and then discretizing with a threshold. Thus, the noise class k^{noise} is identified as the class that is the most active:

$$\bar{\pi}_{t,k} = \text{Erode}(\text{Dilate}(\pi_{t,k}, 101), 101) \quad (14)$$

$$k^{\text{noise}} = \underset{k}{\text{argmax}} \left\{ \sum_t (\bar{\pi}_{t,k} > 0.2) \right\} \quad (15)$$

where 101 is the frame length of the sliding window.

5.2. Class fusing

Sometimes, an initialization leads the SMM to use two classes for one speaker while another speaker is ignored. To mitigate this error, the mixture classes that represent the same speaker are fused. This fusion is again based on the prior $\pi_{t,k}$ which is smoothed with dilation and erosion. If the Intersection-over-Union ratio (IoU)

$$\text{IoU}(k, \kappa) = \frac{\sum_t (\bar{\pi}_{t,k} > 0.2) \wedge (\bar{\pi}_{t,\kappa} > 0.2)}{\sum_t (\bar{\pi}_{t,k} > 0.2) \vee (\bar{\pi}_{t,\kappa} > 0.2)} \quad (16)$$

between two classes k and κ is larger than 0.8, they are fused, i.e. their priors and posteriors are summed up.

5.3. Speech segmentation

Finally, we estimate the speech activities $\hat{a}_{t,k}$ from $\pi_{t,k}$, i.e. smooth and discretize $\pi_{t,k}$. Here, only the dilation is applied to overestimate the activity instead of accidentally removing speech regions at the beginning and end of each utterance. Initial and final silence is no issue for an Automatic Speech Recognition (ASR) system and this further ensures that the segments are long enough for the ASR system to operate on. To keep the false alarms rate small, we had to select a larger value for the threshold as was chosen for noise detection and speaker fusion:

$$\hat{a}_{t,k} = \text{Dilate}(\pi_{t,k}, 79) \geq 0.5. \quad (17)$$

This estimate $\hat{a}_{t,k}$ can be interpreted as a diarization based on spatial information.

6. Extraction

The speech extraction is performed on segments, defined by continuous activity in $\hat{a}_{t,k}$, where k is the target speaker. To extract the target speech from one segment, a mask-based convolutional beamformer [7], factorized [23, 24] into a Weighted Prediction Error (WPE) dereverberation [25] and a wMPDR beamformer [7] component, is used. As the target mask for a segment of speaker k , the posterior $\gamma_{t,f,k}$ of this segment is used. For the distortion mask, the remaining speakers, floored by a small value $\max(\sum_{\kappa \neq k} \gamma_{t,f,\kappa}, 10^{-4})$ are used. Because of space limitations, we skip the equations here. They can be found in [23] and [24].

Table 1: *cpWER on LibriCSS segments for different initializations, fusions, and w/ and w/o WPE dereverberation as preprocessing.*

Initialization		Class fusion		WPE	Overlap ratio in %						Average
		intermediate	final		OS	OL	10	20	30	40	
(1)	Dirichlet	-	-	-	11.7	11.1	13.5	15.3	22.9	21.1	16.6
(2)	Dirichlet (frequency-tied)	-	-	-	13.6	12.2	14.4	17.5	23.2	22.5	17.9
(3)	Dirichlet	-	✓	✓	11.3	11.1	10.1	14.3	15.7	15.9	13.4
(4)	Proposed (frequency-tied)	-	-	-	5.5	6.4	6.8	9.6	11.4	15.7	9.7
(5)	Proposed (frequency-tied)	-	-	✓	5.2	4.7	5.6	8.2	10.6	15.4	8.8
(6)	Proposed (frequency-tied)	-	✓	✓	4.8	4.7	4.7	7.0	8.5	10.9	7.1
(7)	Proposed (frequency-tied)	✓	✓	✓	4.2	4.3	4.2	5.7	7.3	8.5	5.9
(8)	Oracle (frequency-tied)	-	-	-	3.9	4.0	4.2	5.3	6.5	6.1	5.1

Table 2: *Literature comparison. [8, 9, 11] reported the asclite WER (no diarization) and we calculated the average WER from the individual ones. [12] reported cpWER with source counting.*

System	Overlap ratio in %						Average
	OS	OL	10	20	30	40	
Chen et al. [8]	11.9	9.7	13.6	15.0	19.9	21.9	16.0
Raj et al. [12]	11.2	8.5	10.6	14.8	15.5	17.5	13.4
Chen et al. [9]	5.2	4.0	5.8	6.8	9.0	10.0	7.1
Wang et al. [11]	5.4	5.0	4.8	5.0	6.6	7.6	5.9
Proposed	4.2	4.3	4.2	5.7	7.3	8.5	5.9

7. Experiments

7.1. Dataset

For the experiments, we use the LibriCSS database [8] with a sample rate of 16 kHz. This database is meant to simulate a realistic meeting scenario: Eight loudspeakers were placed in a conference room and played back utterances from the LibriSpeech corpus [16] with a variable amount of overlap between the speakers. A seven-channel microphone array recorded the playback. To investigate different conversation styles, the database has different setups: short silence (OS), long silence (OL) and different overlaps between 10% and 40%. For each setup, 10 meetings with a duration of 10 minutes were created. Afterwards, the data was split at silence to obtain smaller segments with an average length of 50 seconds containing 1 to 8 speakers.

7.2. Automatic Speech Recognition

We perform Automatic Speech Recognition (ASR) with a pre-trained LibriSpeech end-to-end speech recognizer [26] from the ESPnet framework [27]. It uses a transformer architecture with an RNN-based language model and achieves a WER of 2.6% on the clean test set of LibriSpeech. We did not adapt the model to separation artifacts.

To assess the performance, we use the concatenated minimum-permutation Word Error Rate (cpWER) [14], which is computed as follows. All transcriptions from the same speaker are concatenated for both the reference and the estimation. Since the ordering of the references and estimates are usually arbitrary, the lowest WER among the permutations between references and estimates is reported.

This is different from other works on LibriCSS [8, 9, 11]. Often, the separation systems produce two output streams without diarization. In this case, it is not possible to obtain speaker-wise transcriptions. So, alternatively, the transcriptions from the same output stream are concatenated and the asclite tool [28], which uses temporal information to align multiple reference utterance transcriptions with the estimated stream transcriptions,

is used to compute a (speaker agnostic) WER.

7.3. Results

Table 1 shows our results on LibriCSS, where 100 EM iterations are used and the STFT size and shift are 1024 and 256, respectively. In row (1) the same initialization as in [20] is used. Changing the initialization to frequency-tied (2), the average performance gets worse.

Replacing the initialization with our proposed initialization (4), we see a clear improvement and using WPE [29] as preprocessing (5) further improves the cpWER. With added class fusion (6) after the EM, the performance is improved, although it discards some speakers. Finally, adding an intermediate speaker fusion (7), where we start with $K + 3$ classes (two more than in (6)) and force a fusion after 10 and 20 EM iterations, i.e. fuse the classes with the highest IoU (see Eq. (16)), yields an cpWER of 5.9%. As an upper bound on the performance, we initialized the SMM with the annotated utterance boundaries (8), see Sec. 4.1, and obtained a cpWER of 5.1%.

Since this is a publically available dataset, we list a few other publications in Table 2. We could not find other separation systems without Neural Networks (NNs), hence we can only compare with NNs. The best system that we found that estimates per-speaker transcriptions is [12]. In contrast to us, they also did source counting. Most other publications [9, 11] build on top of the CSS idea from [8], and hence only separate and recognize speech, but are unable to re-identify speakers. We see a clear improvement over [12], the LibriCSS baseline [8] and followup system [9]. Our system is on par with the currently best reported system in [11], albeit with an unsupervised separation approach.

8. Conclusion

We proposed an initialization scheme for a spatial mixture model to cope with meeting separation, where a class density of the SMM is learned on small segments and then clustered. We showed, on LibriCSS, that this initialization is superior to a conventional initialization. Using dereverberation and an also proposed class fusion, the spatial mixture model is comparable to the best reported WER on this data set, produced by an NN-based meeting separation.

9. Acknowledgements

Computational resources were provided by the Paderborn Center for Parallel Computing. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project 448568305.

10. References

- [1] 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.
- [2] 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)*. IEEE, 2016, pp. 1153–1157.
- [3] —, "Permutation-free convolutive blind source separation via full-band clustering based on frequency-independent source presence priors," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 3238–3242.
- [4] S. Araki, M. Okada, T. Higuchi, A. Ogawa, and T. Nakatani, "Spatial correlation model based observation vector clustering and mvdr beamforming for meeting recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 385–389.
- [5] S. Araki, N. Ono, K. Kinoshita, and M. Delcroix, "Meeting recognition with asynchronous distributed microphone array using block-wise refinement of mask-based mvdr beamformer," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5694–5698.
- [6] C. Boeddeker, J. Heitkaemper, J. Schmalenstroerer, L. Drude, J. Heymann, and R. Haeb-Umbach, "Front-end processing for the chime-5 dinner party scenario," in *CHiME5 Workshop, Hyderabad, India*, vol. 1, 2018.
- [7] T. Nakatani and K. Kinoshita, "Maximum likelihood convolutional beamformer for simultaneous denoising and dereverberation," in *2019 27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019, pp. 1–5.
- [8] Z. Chen, T. Yoshioka, L. Lu, T. Zhou, Z. Meng, Y. Luo, J. Wu, X. Xiao, and J. Li, "Continuous speech separation: Dataset and analysis," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7284–7288.
- [9] S. Chen, Y. Wu, Z. Chen, J. Wu, J. Li, T. Yoshioka, C. Wang, S. Liu, and M. Zhou, "Continuous speech separation with conformer," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5749–5753.
- [10] S. Chen, Y. Wu, Z. Chen, T. Yoshioka, S. Liu, J. Li, and X. Yu, "Don't shoot butterfly with rifles: Multi-channel continuous speech separation with early exit transformer," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6139–6143.
- [11] Z.-Q. Wang, P. Wang, and D. Wang, "Multi-microphone complex spectral mapping for utterance-wise and continuous speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 29, pp. 2001–2014, 2021.
- [12] D. Raj, P. Denisov, Z. Chen, H. Erdogan, Z. Huang, M. He, S. Watanabe, J. Du, T. Yoshioka, Y. Luo *et al.*, "Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 897–904.
- [13] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth 'CHiME' speech separation and recognition challenge: dataset, task and baselines," *arXiv preprint arXiv:1803.10609*, 2018.
- [14] S. Watanabe, M. Mandel, J. Barker, E. Vincent, A. Arora, X. Chang, S. Khudanpur, V. Manohar, D. Povey, D. Raj *et al.*, "CHiME-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings," *arXiv preprint arXiv:2004.09249*, 2020.
- [15] L. Drude, R. Haeb-Umbach, and T. Gerkmann, "Integration of neural networks and probabilistic spatial models for acoustic blind source separation," 2020, 18.12.2020, Dissertation: Universitaet Paderborn (2020).
- [16] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [17] J. Carletta, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, M. Kronenthal *et al.*, "The AMI meeting corpus: A pre-announcement," in *International workshop on machine learning for multimodal interaction*. Springer, 2005, pp. 28–39.
- [18] 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, 2010.
- [19] C. Boeddeker, J. Heitkaemper, J. Schmalenstroerer, L. Drude, J. Heymann, and R. Haeb-Umbach, "Front-End Processing for the CHiME-5 Dinner Party Scenario," in *CHiME5 Workshop*, 2018.
- [20] L. Drude, J. Heitkaemper, C. Boeddeker, and R. Haeb-Umbach, "SMS-WJS: Database, performance measures, and baseline recipe for multi-channel source separation and recognition," *arXiv preprint arXiv:1910.13934*, 2019.
- [21] M. Herdin, N. Czink, H. Ozcelik, and E. Bonek, "Correlation matrix distance, a meaningful measure for evaluation of non-stationary mimo channels," in *2005 IEEE 61st Vehicular Technology Conference*, vol. 1. IEEE, 2005, pp. 136–140.
- [22] T. A. Sorensen, "A method of establishing groups of equal amplitude in plant sociology based on similarity of species content and its application to analyses of the vegetation on danish commons," *Biol. Skar.*, vol. 5, pp. 1–34, 1948.
- [23] C. Boeddeker, T. Nakatani, K. Kinoshita, and R. Haeb-Umbach, "Jointly optimal dereverberation and beamforming," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 216–220.
- [24] T. Nakatani, C. Boeddeker, K. Kinoshita, R. Ikeshita, M. Delcroix, and R. Haeb-Umbach, "Jointly optimal denoising, dereverberation, and source separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2267–2282, 2020.
- [25] T. Nakatani, T. Yoshioka, K. Kinoshita, M. Miyoshi, and B.-H. Juang, "Speech dereverberation based on variance-normalized delayed linear prediction," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 7, pp. 1717–1731, 2010.
- [26] S. Watanabe, "ESPnet2 pretrained model, Shinji Watanabe/librispe ech_asr_train_asr_transformer_e18_raw_bpe_sp_valid_acc.best, fs=16k, lang=en," Jul. 2020. [Online]. Available: <https://doi.org/10.5281/zenodo.3966501>
- [27] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-End Speech Processing Toolkit," in *Interspeech*. ISCA, Sep. 2018, pp. 2207–2211.
- [28] <https://github.com/usnistgov/SCTK>.
- [29] L. Drude, J. Heymann, C. Boeddeker, and R. Haeb-Umbach, "NARA-WPE: A python package for weighted prediction error dereverberation in numpy and tensorflow for online and offline processing," in *Speech Communication; 13th ITG-Symposium*. VDE, 2018, pp. 1–5.