



Distributed speech separation in spatially unconstrained microphone arrays

Nicolas Furnon, Romain Serizel, Irina Illina, Slim Essid

► To cite this version:

Nicolas Furnon, Romain Serizel, Irina Illina, Slim Essid. Distributed speech separation in spatially unconstrained microphone arrays. ICASSP 2021 - 46th International Conference on Acoustics, Speech, and Signal Processing, Jun 2021, Toronto / Virtual, Canada. 10.1109/ICASSP39728.2021.9414758 . hal-02985794v3

HAL Id: hal-02985794

<https://hal.science/hal-02985794v3>

Submitted on 15 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DISTRIBUTED SPEECH SEPARATION IN SPATIALLY UNCONSTRAINED MICROPHONE ARRAYS

Nicolas Furnon, Romain Serizel, Irina Illina

Université de Lorraine, CNRS, Inria, Loria
F-54000 Nancy, France
{firstname.lastname}@loria.fr

Slim Essid

LTCI, Télécom Paris,
Institut Polytechnique de Paris,
Palaiseau, France
slim.essid@telecom-paristech.fr

ABSTRACT

Speech separation with several speakers is a challenging task because of the non-stationarity of the speech and the strong signal similarity between interferent sources. Current state-of-the-art solutions can separate well the different sources using sophisticated deep neural networks which are very tedious to train. When several microphones are available, spatial information can be exploited to design much simpler algorithms to discriminate speakers. We propose a distributed algorithm that can process spatial information in a spatially unconstrained microphone array. The algorithm relies on a convolutional recurrent neural network that can exploit the signal diversity from the distributed nodes. In a typical case of a meeting room, this algorithm can capture an estimate of each source in a first step and propagate it over the microphone array in order to increase the separation performance in a second step. We show that this approach performs even better when the number of sources and nodes increases. We also study the influence of a mismatch in the number of sources between the training and testing conditions.

Index Terms— Speech separation, microphone arrays, distributed processing.

1. INTRODUCTION

Speech separation aims at extracting the speech signals of each speaker in a noisy mixture. It has many applications, for example in automatic speech recognition [1], hearing aids [2] or music processing [3]. In recent years, deep neural network (DNN)-based solutions have replaced model-based approaches because of the great progress they enabled [4–8]. However, most of these DNN-based solutions are developed in “clean” contexts, where the speech signals are not corrupted by noise or by reverberation, which makes them quite unrealistic for real-life applications. Besides, the recent trend shows that state-of-the-art results are mostly achieved with complex DNNs [6–9], with very high model sizes [10]. They also operate on single-channel data, and hence neglect the spatial information which could be accessible in everyday life scenarios, where the majority of recording devices are embedded with multiple microphones. Some solutions have been designed to address the case

of multichannel scenarios [11, 12] or reverberant conditions [13], but they result in even more complex DNNs than those used in the single-channel context. In everyday life scenarios, where several people speak in presence of their personal devices, these solutions are not (yet) applicable because the computational requirements are too demanding for the memory, computing and energy capacities of the devices.

One way to reduce the computational cost of the DNN-based methodologies while exploiting spatial information is to use ad-hoc microphone arrays and to distribute the processing over all the devices of the array. In a previous article, we introduced a solution that proved to efficiently process multichannel data in a distributed microphone array in the context of speech enhancement [14]. This approach was based on a two-step version of the distributed adaptive node-specific signal estimation (DANSE) algorithm by Bertrand and Moonen, where so-called compressed signals are sent among the devices [15]. It achieved, with a simple convolutional recurrent neural network (CRNN), competitive results compared to a filter computed from oracle time-frequency (TF) masks. In an extended study, we have shown that exchanging the estimates of both the target and the noise across the devices could potentially increase the final performance [16].

In this paper, we show that this solution can be adapted to speech separation in a typical use case of a meeting, where all the sources must be estimated in a reverberant but noise-free environment. We redesign the solution such that each device sends an estimate of a different source in a first step. This way, in a second step, each device has access to proper insight into the whole acoustic scene and can separate the signals more accurately. We analyse the performance over the number of sources present and devices available.

This paper is organised as follows. We formulate the problem in Section 2 and detail our solution in Section 3. The experimental setup on which we test the solution is described in Section 4. The results are given in Section 5. We conclude the paper in Section 6.

2. PROBLEM FORMULATION

We consider a scenario where N speakers $\{s_n\}_{n=1..N}$ are recorded by K devices, thereafter called nodes. Each node k contains M_k microphones. We assume that no noise is present, so the m -th micro-

This work was made with the support of the French National Research Agency, in the framework of the project DiSCogs “Distant speech communication with heterogeneous unconstrained microphone arrays” (ANR-17-CE23-0026-01). Experiments presented in this paper were partially carried out using the Grid5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000>).

phone of the k -th node records the following signal at a time stamp t :

$$\begin{aligned} y_{m_k}(t) &= \sum_{n=1}^N c_{n,m_k}(t) * s_n(t) \\ &= \sum_{n=1}^N \hat{s}_{n,m_k}(t) \end{aligned}$$

where $*$ stands for the convolution operator and c_{n,m_k} is the room impulse response (RIR) from the source n to the microphone m_k such that $\hat{s}_{n,m_k} = c_{n,m_k} * s_n$ is the reverberated image of the source n captured by the m_k -th microphone. Speech separation aims at recovering all the speech signals $\{\hat{s}_{n,\mu^{(n)}}\}_{n=1..N}$ where $\mu^{(n)}$ is the reference microphone for source s_n .

In the short-time Fourier transform (STFT) domain, under the narrowband approximation, we can write:

$$y_{m_k}(f, \tau) = \sum_{n=1}^N \hat{s}_{n,m_k}(f, \tau)$$

where f is the frequency index and τ the frame index, and where y and \hat{s} are the STFTs of y and \hat{s} respectively. In the sequel, for the sake of conciseness, we will omit the time and frequency indexes. The under bar \cdot indicates a signal in the STFT domain.

In the context of ad-hoc microphone arrays, we can gather all the microphones of one node into the column vector:

$$\mathbf{y}_k = [\mathbf{y}_{1k}, \dots, \mathbf{y}_{M_k k}]^T,$$

and the vectors of the signals of all nodes into another vector:

$$\mathbf{y} = [\mathbf{y}_1^T, \dots, \mathbf{y}_K^T]^T.$$

3. DISTRIBUTED SPEECH SEPARATION ALGORITHM

To recover the separated signals, we apply a two-step algorithm derived from the DANSE algorithm introduced by Bertrand and Moonen [15], and illustrated in Figure 1 for two nodes. In the first step, at node k , the local signals \mathbf{y}_k are pre-filtered by a local multichannel Wiener filter (MWF) \mathbf{w}_{kk} minimising the mean squared error (MSE) between a single desired speech source $\hat{s}_{k,\mu^{(k)}}$ and the filtered signal:

$$\mathbf{w}_{kk} = \arg \min_{\mathbf{w}} \mathbb{E}\{|\hat{s}_{k,\mu^{(k)}} - \mathbf{w}^H \mathbf{y}_k|^2\}, \quad (1)$$

where \cdot^H is the Hermitian transpose operator. The solution to Equation (1) yields a so-called compressed signal $\mathbf{z}_k = \mathbf{w}_{kk}^H \mathbf{y}_k$ which is a first estimate of the source and sent to all the other nodes. Similarly, the node k receives $K - 1$ compressed signals \mathbf{z}_{-k} :

$$\mathbf{z}_{-k} = [\mathbf{z}_1, \dots, \mathbf{z}_{k-1}, \mathbf{z}_{k+1}, \dots, \mathbf{z}_K]^T,$$

which are stacked with its local signals into:

$$\tilde{\mathbf{y}}_k = [\mathbf{y}_k^T, \mathbf{z}_{-k}^T]^T.$$

In the second step, a second MWF \mathbf{w}_k is applied on $\tilde{\mathbf{y}}_k$ to finally estimate the source $\hat{s}_{k,\mu^{(k)}}$:

$$\hat{s}_{\mu_k} = \mathbf{w}_k^H \tilde{\mathbf{y}}_k. \quad (2)$$

The advantages of this algorithm are twofold. First, it removes the dependency on a central node that would gather all the signals of

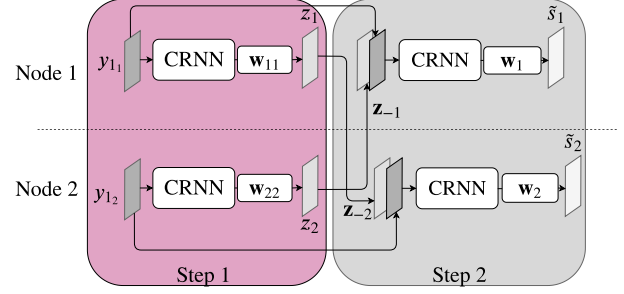


Fig. 1. Illustration of our solution in a two-node context. "CRNN" refers to convolutional recurrent neural network.

all nodes. Second, each node has access to the spatial information of the whole microphone array, but in a condensed and pre-filtered form as each node sends only the compressed signal.

If the number of nodes is at least equal to the number of speakers, at the scale of one node k , the speech separation problem can be viewed as a speech enhancement problem in the presence of interfering speakers, where the target source is $\hat{s}_{k,\mu^{(k)}}$. Under the assumption that the sources are uncorrelated, the solution to Equation (1) is given by:

$$\mathbf{w}_{kk} = \mathbf{R}_{y_k}^{-1} \mathbf{R}_{s_k} \mathbf{e}_1. \quad (3)$$

\mathbf{R}_{y_k} is the spatial covariance matrix of the mixture \mathbf{y}_k ; \mathbf{R}_{s_k} is the spatial covariance matrix of the target signal and we have $\mathbf{e}_1 = [1, 0, \dots, 0]$. The filter in Equation (2) can be obtained in a similar way by replacing the covariance matrix of \mathbf{y}_k by the covariance matrix of $\tilde{\mathbf{y}}_k$ and by computing the target covariance matrix out of the target components of $\tilde{\mathbf{y}}_k$.

These different covariance matrices are computed from signals estimated with a TF mask applied on the mixture. We use one common mask for all the signals of one node, as we noticed that taking a specific mask for the compressed signals does not influence the final performance [16]. In the first step, the mask is estimated by a single-node CRNN, which predicts the mask from a single (local) mixture signal of the node it operates on. In the second step, the TF mask is estimated by a multi-node CRNN, which predicts the masks from the local signal together with the compressed signals sent by all the other nodes [16].

As a consequence of this methodology, in a setup where $K \geq N$, i.e. where there are at least as many nodes as sources, each node can automatically estimate a different source and send a different compressed signal to the other nodes (see Figure 2). That way, at the second step, each node has an estimate of all the sources, which helps the multi-node neural network to better predict the TF mask.

4. SETUP

4.1. Dataset

We simulate a typical situation of a meeting, where N persons are talking around a round table, each of the speakers having laid their smartphone, tablet or computer in front of them on the table. We assume that all the nodes have $M_k = 4$ microphones. The shoebox-like room has a random length, width and height uniformly drawn in [3, 9] meters, [3, 7] meters and [2.5, 3] meters, respectively. The table has a radius randomly drawn between 0.3 m and 2.5 m. The height of the table is randomly selected between 0.8 m and 0.9 m.

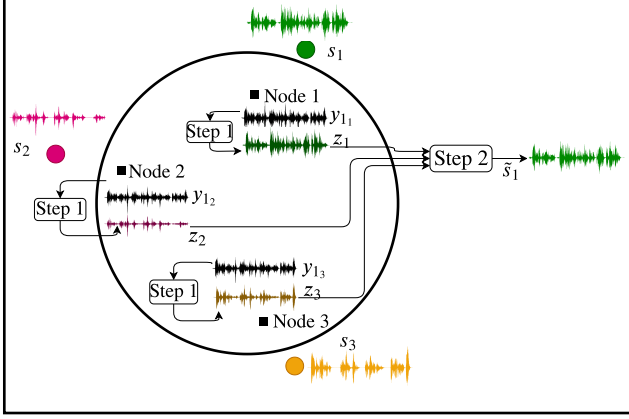


Fig. 2. Illustration of our solution in a three-node context, focusing on the first node for the second step.

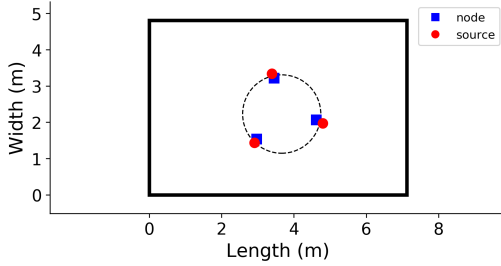


Fig. 3. Illustration of one configuration with three sources (hence three nodes).

The sources are evenly placed around the table, so that the angle between two sources is equal for all the pairs of sources. Their distance to the table edge is randomly selected between 0 cm and 50 cm, and their height between 1.15 m and 1.80 m, as if people were sitting or standing close to the table. The reverberation time is randomly selected between 150 ms and 400 ms. The level of all the sources is set to the same power. The RIRs of the room with N equal to 2, 3 and 4 are computed with Pyroomacoustics [17]. The acoustical effect of the table is not simulated. An example with $N = 3$ is shown in Figure 3.

All the speech files are from LibriSpeech [18]. The repartition of the `train-clean-360`, `dev-clean` and `test-clean` subfolders is kept for our split between training, validation and test datasets. Within a mixture, the different speech signals fully overlap in time. We created around 30 hours of training data, 3 hours of validation data and 3 hours of test data.¹

4.2. Experimental settings

All the signals are sampled at a frequency of 16 kHz. The STFT is computed with a Hanning window of 32 ms with an overlap of 16 ms. The same CRNN model as the one used in our previous experiments is used [16]. It is made of three convolutional layers, followed by a recurrent layer and a fully-connected layer. The convolutional layers have 32, 64 and 64 filters respectively, with kernel

size 3×3 and stride 1×1 . The recurrent layer is a 256-unit gated recurrent unit, and the activation function of the fully-connected layer is a sigmoid. The network was trained with the RMSprop optimizer [19]. The input of the model are STFT windows of 21 frames and the ground truth targetted are the corresponding frames of the ideal ratio mask (IRM).

5. RESULTS

We compare four methods in terms of scale-invariant signal to distortion ratio (SI-SDR) [20]. The first method uses IRMs to compute the signal statistics. In the scenarios we designed, unless a source is removed, node k is always in front of the k -th source. At the scale of this node, the speech separation problem is a speech enhancement problem where all the sources $j \neq k$ sum to the noise component. Hence we can compute the IRM at node k as:

$$\text{IRM}_k = \frac{|\hat{s}_{k,\mu^{(k)}}|}{|\hat{s}_{k,\mu^{(k)}}| + |\underline{n}_{k,\mu^{(k)}}|} \quad (4)$$

where $\mu^{(k)}$ is the reference microphone for source s_k and

$$\underline{n}_{k,\mu^{(k)}} = \sum_{j \neq k} \hat{s}_{j,\mu^{(k)}}.$$

This oracle method is denoted “IRM” in the legends of the following figures. The second method, denoted “MN”, is our multi-node solution where the CRNN sees at the second step the compressed signals to predict the masks. The third method, denoted “SN”, is the single-node solution where the same CRNN sees only the local signal to predict the masks at both filtering steps. The fourth method, denoted “MWF”, is a MWF applied on each node without exchanging the signals. We analyse the behaviour of these methods when the number of sources and nodes varies.

5.1. Performance with an equal number of sources

In this section, the four methods are compared in the scenarios where the number of sources and nodes is equal. The performance in terms of SI-SDR are reported in Figure 4. First, although the single-node solution and the MWF do not differ much, the single-node distributed processing shows a significant improvement over the MWF when the number of nodes increases. This shows that exploiting the spatial information conveyed by the whole microphone array helps improving the separation performance. Besides, the multi-node solution significantly outperforms both the single-node solution and the MWF. This shows that the compressed signals are useful not only for the beamforming but also for the mask prediction. Finally, the $\Delta\text{SI-SDR}$ increases when the number of nodes and sources increases, even if the task gets more challenging, and achieves less than 0.5 dB worse than the oracle performance. This is because the output performance remains constant while the input SI-SDR decreases. This shows the robustness of our solution to spatial diversity.

5.2. Performance in over-determined cases

In this section, we analyse the performance of the previous methods in the case where the number of nodes is greater than the number of sources. This could happen in a real situation, for example if a person gets out of the room while leaving their phone on the table. In such a scenario, each node estimates the source in front of which it is placed, which means that the node left without speaker has no target speaker. At this node, the target and noise components of

¹A Python implementation to generate the dataset is available at https://github.com/nfurnon/disco/tree/master/dataset_generation/gen_meetit

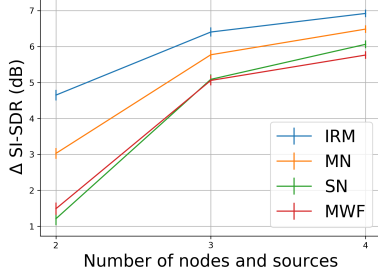


Fig. 4. Speech separation performance of the different methods when the number of sources and nodes is equal. The bars correspond to the 95% confidence interval.

Equation (4) are not defined, and it is not obvious to determine the compressed signal to send with the IRM method. Because of this, we omit the results obtained with the IRM in this section. The results of the three other methods are reported in Figure 5 where we recall the performance of the equally-determined case ($N = K = 2$ and $N = K = 3$) for an easier comparison.

In over-determined cases as well, the multi-node solution outperforms almost always the other two methods. An interesting exception can be noticed when there are two more nodes than sources ($N = 2, K = 4$). We observed that the nodes placed in front of no source estimate completely silent compressed signals, because the masks predicted by the single-node CRNN are close to 0 in almost all TF bins. This means that the filter at the second step is applied on silent signals, which degrades the final performance. Since the multi-node solution still outperforms the single-node solution, the silent compressed signals do not degrade the mask prediction. However, they degrade the distributed beamforming output. This effect is not dominant when there is only one more node than the number of sources, which means that our solution is robust to a source mismatch between training and testing, but only to a limited extent. One solution to cope with this performance drop, could be to automatically ignore the recording devices which start to send silent signals.

Finally, the performance consistently increases for all methods when the number of nodes increases from $K = 2$ to $K = 3$ with two sources. This shows that the neural networks trained with more sources (so on harder conditions) perform better than those trained with a lesser number of sources.




Fig. 5. Speech separation performance in over-determined cases with (a) $N = 2$ sources and (b) $N = 3$ sources.

5.3. Performance in under-determined cases

We now focus on the performance of the proposed methods in the case where the number of nodes is smaller than the number of sources. This could happen in a real situation, for example if a phone shuts down during the meeting. The performance for the scenario with $N = 3$ sources is presented in Figure 6 where we recall the performance of the equally-determined case ($N = K = 3$) for an easier comparison.

In this under-determined case, even if there is only one more source than nodes, the multi-node solution performs worse than the two other ones. The mismatch between training and testing leads to worse performance. Since the single-node solution, as well as the MWF, performs quite well, it means that the drop of performance of the multi-node solution is due to the multi-node CRNN that is trained on mixtures with only one interferent speaker while tested on mixtures with two interferent speakers. A similar behaviour was observed in the under-determined cases with $N = 4$ sources.

This indicates that dealing with under-determined cases probably requires to train specific networks with the proper number of interferent sources. Training and testing on a variable number of sources remains an open challenge. This could be addressed within our spatially distributed framework by adapting a dedicated strategy, e.g. in estimating the sources iteratively [21] or by adapting the loss function [22, 23].

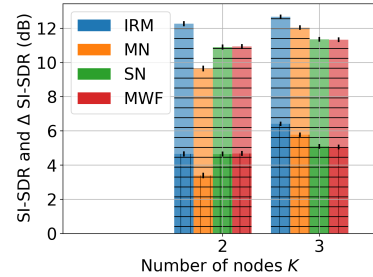


Fig. 6. Speech separation performance in an under-determined case with $N = 3$ sources. The lighter bars with horizontal hatches correspond to the SI-SDR. The darker bars with vertical hatches correspond to the Δ SI-SDR.

6. CONCLUSION

We introduced a distributed algorithm that can process spatial information captured by a spatially unconstrained microphone array. The spatial information is propagated over the microphone array to deliver to all the nodes a global insight into the whole acoustic scene. We evaluated the algorithm in typical meeting configurations and showed that thanks to the spatial information, a CRNN can predict accurate TF masks which lead to almost oracle performance. In scenarios where the number of nodes match the number of sources, we showed that the performance increases when the number of sources (and nodes) increases. We also analysed the limits of this approach when the number of nodes does not match the number of sources. Solving this problem could require the adaptation of techniques developed in the case of fixed microphone arrays.

7. REFERENCES

- [1] Jon Barker, Shinji Watanabe, Emmanuel Vincent, and Jan Trmal, “The fifth ‘CHiME’ speech separation and recognition challenge: dataset, task and baselines,” *Interpseech 2018 – 19th Annual Conference of the International Speech Communication Association*, 2018.
- [2] Kostas Kokkinakis and Philipos C. Loizou, “Using blind source separation techniques to improve speech recognition in bilateral cochlear implant patients,” *The Journal of the Acoustical Society of America*, vol. 123, no. 4, pp. 2379–2390, 2008.
- [3] Cemil Demir, Murat Saraclar, and Ali Taylan Cemgil, “Single-channel speech-music separation for robust ASR with mixture models,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 4, pp. 725–736, 2012.
- [4] Hakan Erdogan, John R. Hershey, Shinji Watanabe, and Jonathan 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)*, 2015, pp. 708–712.
- [5] John R. Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 31–35.
- [6] Yi Luo and Nima 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.
- [7] Liwen Zhang, Ziqiang Shi, Jiqing Han, Anyan Shi, and Ding Ma, “FurcaNeXt: End-to-end monaural speech separation with dynamic gated dilated temporal convolutional networks,” in *International Conference on Multimedia Modeling*. Springer, 2020, pp. 653–665.
- [8] Neil Zeghidour and David Grangier, “Wavesplit: End-to-end speech separation by speaker clustering,” *arXiv preprint arXiv:2002.08933*, 2020.
- [9] Jingjing Chen, Qirong Mao, and Dong Liu, “Dual-path transformer network: Direct context-aware modeling for end-to-end monaural speech separation,” *arXiv preprint arXiv:2007.13975*, 2020.
- [10] Yi Luo, Zhuo Chen, and Takuya Yoshioka, “Dual-path RNN: Efficient long sequence modeling for time-domain single-channel speech separation,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 46–50.
- [11] Rongzhi Gu, Jian Wu, Shi-Xiong Zhang, Lianwu Chen, Yong Xu, Meng Yu, Dan Su, Yuexian Zou, and Dong Yu, “End-to-end multi-channel speech separation,” *arXiv preprint arXiv:1905.06286*, 2019.
- [12] Dongmei Wang, Zhuo Chen, and Takuya Yoshioka, “Neural speech separation using spatially distributed microphones,” *arXiv preprint arXiv:2004.13670*, 2020.
- [13] Masood Delfarah and DeLiang Wang, “Deep learning for talker-dependent reverberant speaker separation: An empirical study,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1839–1848, 2019.
- [14] Nicolas Furnon, Romain Serizel, Irina Illina, and Slim Essid, “DNN-based distributed multichannel mask estimation for speech enhancement in microphone arrays,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 4672–4676.
- [15] Alexander Bertrand and Marc Moonen, “Distributed adaptive node-specific signal estimation in fully connected sensor networks — Part I: Sequential node updating,” Oct 2010.
- [16] Nicolas Furnon, Romain Serizel, Irina Illina, and Slim Essid, “DNN-based mask estimation for distributed speech enhancement in spatially unconstrained microphone arrays,” *submitted to IEEE/ACM Transactions on Audio Speech and Language Processing*, 2020.
- [17] Robin Scheibler, Eric Bezzam, and Ivan Dokmanic, “Py-roomacoustics: A python package for audio room simulation and array processing algorithms,” *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 351–355, 2018.
- [18] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5206–5210, 2015.
- [19] Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky, “Lecture 6.5–rmsprop: Divide the gradient by a running average of its recent magnitude,” *COURSERA: Neural networks for machine learning*, 2012.
- [20] Jonathan Le Roux, Scott Wisdom, Hakan Erdogan, and John R. Hershey, “SDR–half-baked or well done?,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 626–630.
- [21] Keisuke Kinoshita, Lukas Drude, Marc Delcroix, and Tomohiro Nakatani, “Listening to each speaker one by one with recurrent selective hearing networks,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5064–5068.
- [22] Yi Luo and Nima Mesgarani, “Separating varying numbers of sources with auxiliary autoencoding loss,” *arXiv preprint arXiv:2003.12326*, 2020.
- [23] Nicolas Turpault, Scott Wisdom, Hakan Erdogan, John Hershey, Romain Serizel, Eduardo Fonseca, Prem Seetharaman, and Justin Salamon, “Improving sound event detection in domestic environments using sound separation,” *arXiv preprint arXiv:2007.03932*, 2020.