



Iterative Sound Source Localization for Unknown Number of Sources

Yanjie Fu¹, Meng Ge^{1,2,*}, Haoran Yin¹, Xinyuan Qian², Longbiao Wang^{1,*},
Gaoyan Zhang¹, Jianwu Dang^{1,3}

¹Tianjin Key Laboratory of Cognitive Computing and Application,
College of Intelligence and Computing, Tianjin University, Tianjin, China

²Department of Electrical and Computer Engineering, National University of Singapore, Singapore

³Japan Advanced Institute of Science and Technology, Ishikawa, Japan

{fuyanjie, gemeng, haoran.yin, longbiao.wang}@tju.edu.cn

Abstract

Sound source localization aims to seek the direction of arrival (DOA) of all sound sources from the observed multi-channel audio. For the practical problem of unknown number of sources, existing localization algorithms attempt to predict a likelihood-based coding (i.e., spatial spectrum) and employ a pre-determined threshold to detect the source number and corresponding DOA value. However, these threshold-based algorithms are not stable since they are limited by the careful choice of threshold. To address this problem, we propose an iterative sound source localization approach called ISSL, which can iteratively extract each source's DOA without threshold until the termination criterion is met. Unlike threshold-based algorithms, ISSL designs an active source detector network based on binary classifier to accept residual spatial spectrum and decide whether to stop the iteration. By doing so, our ISSL can deal with an arbitrary number of sources, even more than the number of sources seen during the training stage. The experimental results show that our ISSL achieves significant performance improvements in both DOA estimation and source number detection compared with the existing threshold-based algorithms.

Index Terms: Sound source localization, direction-of-arrival estimation, spatial spectrum, unknown multiple sound sources

1. Introduction

Sound source localization (SSL) is the technique of estimating the position of one or several sound sources from the multi-channel signals captured by the microphone array. It can be simplified to the estimation of direction-of-arrival (DOA) to the microphone array of each sound source. This localization cue is now an imperative part of many real-world speech applications such as speech separation [1, 2, 3], speaker extraction [4, 5, 6], speech recognition [7, 8] and speaker diarization [9, 10, 11], which makes SSL a valuable research topic.

In the past few decades, most researchers have focused on DOA estimation with a known number of sound sources. They usually borrow the nonlinear ability of neural network to derive source's direction from a low-level or high-level signal representation, including waveform [12, 13, 14], spectrogram [15, 16], acoustic intensity vector [17, 18, 19] and GCC-PHAT [20, 21, 22]. Despite their promising results, localization of sound sources in real-world scenarios with unknown number of sources is still a challenge.

Recently, a lot of works utilize a predefined threshold to deal with the DOA estimation problem when the number of active sources is not known in advance, including [1, 15, 21, 23].

* denotes the corresponding author.

For example, He et al. [21] design a 4-layer fully connected network and a 4-layer convolution neural network to map the GCC-PHAT input feature to desired spatial spectrum coding. In addition, He et al. [24] further the idea and derive the spatial spectrum coding from both the real and imaginary parts of the complex-valued spectrogram. These threshold-based solutions aim to predict this likelihood-based spatial spectrum output at each possible azimuth direction, and then employ an empirically-defined threshold to detect the source number and estimate the corresponding DOA value. However, these methods are extremely dependent on the threshold selection and are not stable in real-world applications, since real-world scenes often contain various noises and different sound source types.

To address the threshold problem for unknown number of sources, we propose an iterative sound source localization approach called ISSL (Fig. 1), which avoids the threshold assumption by using an iterative strategy and doesn't require the knowledge of the number of sources in the observed signal. Inspired by the success that spatial spectrum can encode a signal with an arbitrary number of sources into a likelihood-based vector, we first build a spatial spectrum prediction module using ResNet-based [25] network (Fig. 2(a)) to obtain the spatial spectrum output coding for the following DOA estimation. Unlike the existing threshold-based algorithms, we then design an iterative DOA estimation module using VGG-based [26] network (Fig. 2(b)) to accept spatial spectrum to iteratively extract each source's DOA until the termination criterion is met. This iterative DOA estimation module replaces the original threshold operation with a termination strategy designed by a binary classifier. By doing so, our ISSL system can handle an arbitrary number of sources without a pre-determined threshold, even more than the number of sources seen during training. Compared with the threshold-based algorithm on source number detection, our ISSL system naturally avoids the instability of the threshold, and the experimental results also verify that our system achieves 19.41% and 16.95% absolute improvement in detection accuracy for scenarios containing at most 3 and 4 sources, respectively.

2. ISSL System

In this section, we describe our proposed ISSL system for unknown multiple sound sources localization. Then we introduce how we infer with this system, followed by training methods.

2.1. Overview

Our ISSL system consists of spatial spectrum prediction module and iterative DOA estimation module, as shown in Fig. 1.

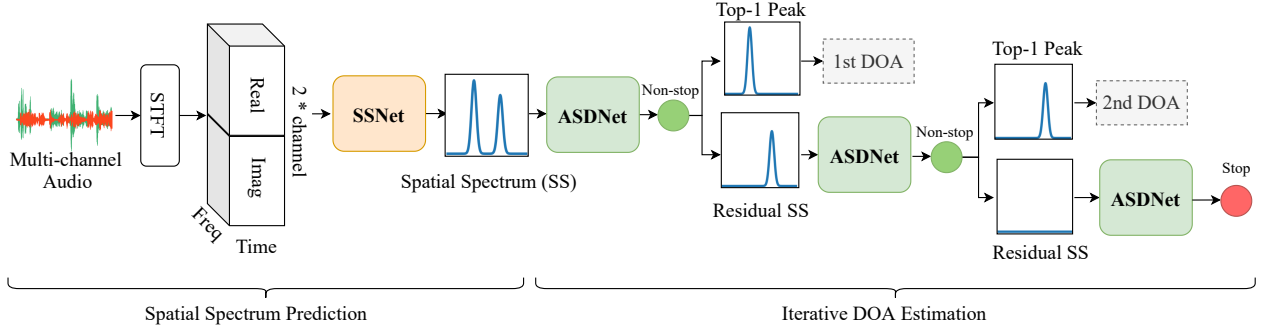


Figure 1: Illustration of our proposed sound source localization framework ISSL when the number of sources is 2. SSNet processes the multi-channel audio feature and produces one spatial spectrum. ASDNet predicts if there still exists active sound source. Iterative DOA Estimation predicts the final DOAs as well as the number of active sound sources.

The former module aims to predict the spatial spectrum from the multi-channel input audio, while the latter one counts the source number and estimates each source’s DOA via the proposed iterative strategy. The whole training and inference process is data-driven and has no threshold operation.

2.2. Spatial Spectrum Prediction

The spatial spectrum prediction module is built by using Spatial Spectrum Network (SSNet) (Fig. 2(a)), which learns the non-linear mapping from the audio segment to the spatial spectrum.

Formally, let X_c be the complex STFT of the c -th channel mixture signal x_c of one segment. We concatenate the real part and imaginary part of STFT along the channel axis to obtain the feature representation v which is defined as:

$$v = \{\mathcal{R}e[X], \mathcal{I}m[X]\} \quad (1)$$

where $X = \{X_c\}_{c=1}^C$ is the input complex-value spectrogram computed from all C -microphones. After that, we use v as the input of SSNet to predict the spatial spectrum:

$$\hat{p} = \text{SSNet}(v) \quad (2)$$

Specifically, the network structure of SSNet is shown in Fig. 2(a). It contains two convolutional layers for downsampling in the frequency dimension, five residual blocks for the high-level feature extraction, one 1×1 convolutional layer for projection onto the DOA space followed by an operation of swapping the DOA and the frequency axes. The last two layers convolve along the time and DOA axes and the final output representation is a 360-dimensional vector.

Inspired by [21], we adopt the likelihood-based spatial spectrum to represent the probability of sound source locating at the 360 individual directions. Specifically, each element of the encoded 360-dimensional vector $p(\theta_i)$ is assigned to a particular azimuth direction $\theta_i \in \{1^\circ, 2^\circ, \dots, 360^\circ\}$. The values of the vector follow a Gaussian distribution that maximizes at the ground truth DOA, which is defined as follow:

$$p(\theta_i) = \begin{cases} \max_{\theta' \in y} \left\{ \exp\left(-\frac{d(\theta_i, \theta')^2}{\sigma^2}\right) \right\} & \text{if } |y| > 0 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where y is the set of the ground truth sound sources, $|y|$ is the number of sources, θ' is the ground truth DOA of one source, σ is a predefined constant that controls the width of the Gaussian function and $d(\cdot, \cdot)$ denotes the angular distance.

We train SSNet by computing the mean squared error (MSE) between the segment-wise output spatial spectrum \hat{p} from SSNet and the ground truth p , which represented as

$$\mathcal{L}_{\text{SSNet}} = \|\hat{p} - p\|^2 \quad (4)$$

2.3. Iterative DOA Estimation

In practical sound source localization scenarios, we will not be informed of the source number beforehand. Thus, we have no idea how many peaks in the spatial spectrum should be taken as predictions during inference. In [24], the peaks in the predicted spatial spectrum above a given threshold are taken as predictions when the number of sources is unknown. This threshold-based method does achieve relatively good results of localization for unknown number of sources. However, the optimal threshold is hard to choose and a pre-determined threshold for all domains is prone to errors due to the domain shift problem.

To perform better localization for unknown number of sources, we propose to estimate DOAs by applying the Active Source Detector Network (ASDNet) iteratively inspired by [27]. Instead of predicting all DOAs at once, the proposed approach predicts only one DOA from a spatial spectrum mixture at a time and the residual spatial spectrum is fed back to the model for the next iteration, as shown in Fig. 1.

Specifically, we use ASDNet in the task of binary classification to estimate if there exists active sound source given a spatial spectrum. If so, we zero the highest peak of the spatial spectrum to separate the sound source from the spatial spectrum mixture and feed the new spatial spectrum into the model to continue the iteration. Otherwise, we stop the iteration and get a DOA estimation as well as an estimation of the source number. The complete algorithm is shown in Algorithm 1.

In our ASDNet architecture presented in Fig. 2(b), we insert a height dimension and a channel dimension into the input 360-dimensional spatial spectrum’s shape, respectively. As a result, the input shape is (1, 1, 360). Our proposed ASDNet adopts VGG-like architecture [26], which consists of 12 convolutional layers and 2 fully connected (FC) layers. The 6 convolutional blocks have 16, 32, 64, 128, 256 and 512 filters, respectively. Each of the convolutional layers is followed by a Batch Normalization layer [28] and Rectified Linear Unit [29]. The first hidden layer has 256 neurons. The second hidden layer is a FC layer with 32 neurons. The last output layer has only one neuron with sigmoid activation function. We label the residual spatial spectrum containing active source as 0 and 1 the other

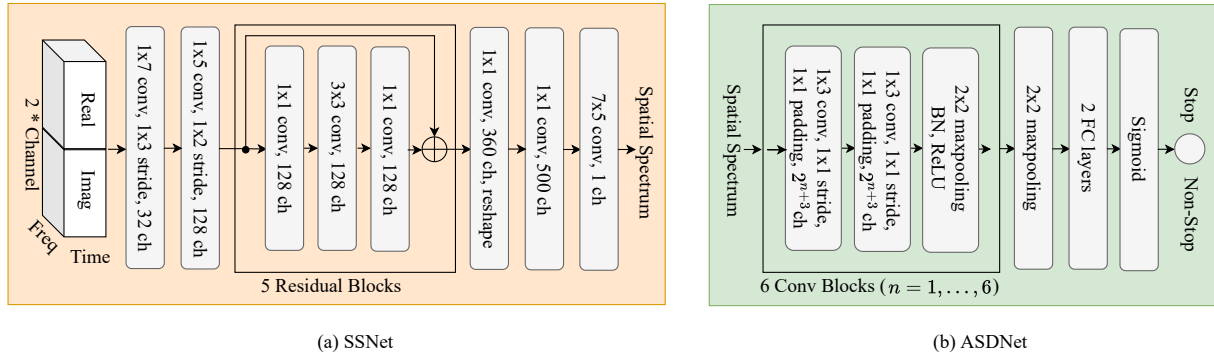


Figure 2: Architectures of Spatial Spectrum Network (SSNet) and Active Source Detector Network (ASDNet).

Algorithm 1: Iterative DOA Estimation for Unknown Number of Sources

Input : One 360-dimensional predicted spatial spectrum \hat{p}

Output: Estimated DOAs *results* & number of sources *nos*

- 1 $results \leftarrow \{\}, nos \leftarrow 0$
 - 2 Find the maximum value of the spectrum and its corresponding DOA: $\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \hat{p}(\theta)$
 - 3 Feed \hat{p} into ASDNet
 - 4 **while** ASDNet detects an active source **do**
 - 5 Update \hat{p} by setting the values at $\hat{\theta}$ and its neighborhood to 0
 - 6 $results \leftarrow results \cup \{\hat{\theta}\}$
 - 7 $nos \leftarrow nos + 1$
 - 8 Find the maximum value of the spectrum and its corresponding DOA: $\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \hat{p}(\theta)$
 - 9 Feed \hat{p} into ASDNet
 - 10 **end**
 - 11 **return** *results*, *nos*
-

way around. During each iteration of inference, we interpret the output value above 0.5 as “stop” and below 0.5 as “non-stop”. We train ASDNet to minimize the binary cross entropy between the label value and the output value as follows:

$$\mathcal{L}_{\text{ASDNet}} = -\frac{1}{N} \sum_{n=1}^N (z_n \log \hat{z}_n + (1 - z_n) \log(1 - \hat{z}_n)) \quad (5)$$

where N is the total amount of samples and $n = 1, \dots, N$ is the sample index, $z_n \in \{0, 1\}$ is the label value of the n^{th} sample, $\hat{z}_n \in \{0, 1\}$ is the model’s output value of the n^{th} sample.

2.4. Network Training

For training ISSL system, we need to train SSNet and ASDNet separately. Once finish training SSNet, we build a spatial spectrum dataset predicted by SSNet for further training ASDNet.

Specifically, we train SSNet with a simulated multi-channel audio dataset in the first step. Then, we train ASDNet with the spatial spectrum dataset containing samples inferred from a dataset which is the same duration as the training set using SSNet. Among the samples, the positive samples which indicate to stop the iteration are the residual spatial spectrum of the 1st

Table 1: The parameters settings of simulated rooms.

Length(L)	Width(W)	Height(H)	RT60
[4m, 6m]	[3m, L]	[2.7m, 3.5m]	[0.2, 0.5]
[6m, 10m]			[0.3, 0.6]
[10m, 15m]			[0.4, 0.7]

iteration’s result of 1 source, 2nd iteration’s result of 2 sources and 3rd iteration’s result of 3 sources since no active sources remain in the residual spatial spectrum. All other combinations such as the 1st iteration of 2 sources mixture are used as negative samples as the residual spatial spectrum still contain sound.

3. Experiment and Discussion

3.1. Dataset

We simulate multi-channel mixture audio of different number of sources from 0 to 4. The clean speech audio data are randomly selected from the VCTK [30] corpus at a sample rate of 48kHz. The simulated database contains 109 speakers and is divided into three sets: training set (89 speakers), validation set (10 speakers), and test set (10 speakers). All multi-channel audio data are generated by convolving clean speech audio with simulated room impulse responses using Pyroomacoustics [31].

Specifically, we first generate spatialized 4-channel audio data of single-source speech in cubic rooms of random sizes. We place the microphone array in the center of the room. The distances between the sound source and the walls are at least 0.5m. The microphone array geometry is set according to that of the microphone array which is used during collecting SSLR database [21]. Then, we respectively mix 3 and 4 randomly selected audios with a minimum angular spacing of 20 degree. We duplicate or truncate the original audio to the guarantee all mixed audios are 5 seconds long. We slice the audios into 170 ms long segments (8192 samples in 48kHz recordings) and label each segment with the active sources’ DOAs. It is worth mentioning that we identify active sources in each segment with the assistance of SpeechBrain [32] VAD results on every single source. Thus, we obtain VCTK-3mix and VCTK-4mix that contain 0,1,2,3 and 0,1,2,3,4 sources segments, respectively. The total duration of training set, validation set and test set are 27.8h, 2.78h and 2.78h, respectively. We simulated 50 rooms for training set and 2 rooms each for validation set and test set. The parameters of the simulated rooms are listed in Table 1.

Table 2: Precision (%), Recall (%) and F1 (%) in a comparative study of DOA estimation for unknown number of sources.

Dataset	Methods	Precision	Recall	F1
VCTK-3mix	Threshold (CNN)	49.40	40.92	44.76
	Threshold (Res)	77.01	81.42	79.15
	Our ISSL	84.41	82.19	83.29
VCTK-4mix	Threshold (CNN)	45.40	34.12	38.96
	Threshold (Res)	77.87	77.04	77.45
	Our ISSL	79.83	78.03	78.92

Table 3: Detection accuracy (%) of number of sound sources.

Methods	VCTK-3mix	VCTK-4mix
Threshold (CNN) [21]	24.74	24.81
Threshold (Res) [24]	62.01	58.50
Our ISSL	81.41	75.45

The dataset in our experiments can be found here¹.

3.2. Experimental Setup

We train all models in our ISSL system for 30 epochs. Adam [33] optimizer is used in training both SSNet and ASDNet. We set the minibatch size to 100 in all experiments. The parameter σ of Gaussian curve in spatial spectrum coding is 8. In our experiments, the segments are 170ms long 4-channel audio (8192 samples with 48kHz sampling rate). We compute the STFT of these segments with a 43ms (2048 samples) long window and 50% long hop. Thus, we have 7 frames of STFT from each segment as the network input. We only use the frequency bins between 100 and 8000 Hz, so the number of frequency bins is reduced to 337. The dimension of the feature representation v is $7 \times 337 \times 8$. Source code in this paper is available¹.

3.3. Comparative study on VCTK-3mix and VCTK-4mix

We compare our ISSL system with the previous state-of-the-art threshold-based baseline systems on VCTK-3mix and VCTK-4mix in terms of Precision, Recall and F1-score [34]. From Table 2, we find that our ISSL system significantly outperforms the ‘‘Threshold (CNN)’’ and ‘‘Threshold (Res)’’ systems on VCTK-3mix with absolute F1-score improvements of 38.53% and 4.14%, respectively. We also achieve 39.96% and 1.47% absolute F1-score improvements on VCTK-4mix, respectively. The results demonstrate that the SSL task with more sources is more challenging, and validate that our ISSL system still maintains superiority on both VCTK-3mix and VCTK-4mix.

We further report the source number detection performance on VCTK-3mix and VCTK-4mix, as shown in Table 3. It is observed that our ISSL system with iterative strategy achieves the best performance for detection of the number of sound sources. Specifically, our system leads up to 31.3% and 29.0% relative improvements than ‘‘Threshold (Res)’’ baseline system on both two databases. This observation further demonstrates that the performance improvements of DOA estimation mainly come from better detection on the number of sound sources.

3.4. A universal study of ISSL system

To prove the motivation that our ISSL system can deal with more sources than that seen during the training stage, we study

¹<https://github.com/FYJNEVERFOLLOWS/ISSL>

Table 4: A universal study on VCTK-4mix. Precision (%), Recall (%) and F1 (%) are the metrics of DOA estimation. Detection accuracy (%) is the metric of source number detection.

Model	Train Data	Infer Data	Precision	Recall	F1	Detection Accuracy
ISSL	4mix	4mix	79.83	78.03	78.92	75.45
	3mix	4mix	81.35	75.84	78.50	74.47

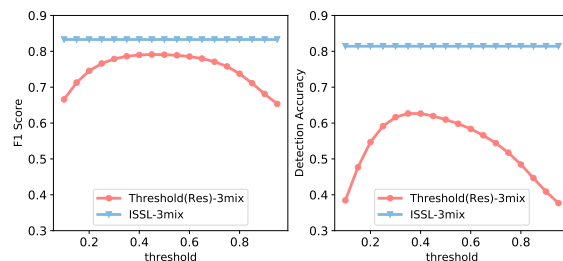


Figure 3: A stability analysis of ISSL and Threshold (Res) with different threshold values on VCTK-3mix. Left part is the results of DOA estimation, and right is source number detection results.

a universal experiment summarized in Table 4. As shown in Table 4, our ISSL system still works well when we directly use the model trained on VCTK-3mix to infer the results on VCTK-4mix. This can be explained that our system uses a data-driven way to produce the terminal criterion automatically, which allows the system to decide whether to continue iteration process and output the DOA result. These observed results verify the universal characteristics of our proposed ISSL system. In addition, we find that the performance of ISSL system is just slightly degraded. This phenomenon reflects our system’s robustness, and also provides a solution to the problem of reducing training costs in SSL task.

3.5. A stability analysis of ISSL system

We further verify the robustness and stability of our ISSL system compared with the threshold-based baseline on different threshold values. The results are shown in Figure 3. We can conclude that the performance of DOA estimation and source number detection vary greatly with the change of the predefined threshold value, which is unsatisfactory for practical speech applications. Benefiting from the advantage of not requiring a threshold, our ISSL is quite stable.

4. Conclusions

We propose a novel iterative sound source localization approach to estimate DOA and detect the number of sound sources. Our novelty lies in designing an iterative strategy to replace the original threshold operation. This strategy allows our approach to handle an arbitrary number of sources, even more than the number of sources in the training step. Experimental results show that our proposed ISSL achieves better DOA estimation and source number detection performance for unknown number of sources compared to the unstable threshold-based algorithms.

5. Acknowledgements

This work was supported partially by the National Key RD Program of China under Grant 2018YFB1305200 and the National Natural Science Foundation of China under Grant 62176182.

6. References

- [1] S. E. Chazan, H. Hammer, G. Hazan, J. Goldberger, and S. Ganot, "Multi-microphone speaker separation based on deep doa estimation," in *2019 27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019, pp. 1–5.
- [2] Z. Wang and D. Wang, "Localization based sequential grouping for continuous speech separation," *CoRR*, vol. abs/2107.06853, 2021. [Online]. Available: <https://arxiv.org/abs/2107.06853>
- [3] X. Li, Y. Xu, M. Yu, S.-X. Zhang, J. Xu, B. Xu, and D. Yu, "Mimo self-attentive rnn beamformer for multi-speaker speech separation," *arXiv preprint arXiv:2104.08450*, 2021.
- [4] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, "L-spex: Localized target speaker extraction," *arXiv preprint arXiv:2202.09995*, 2022.
- [5] Z. Zhang, Y. Xu, M. Yu, S.-X. Zhang, L. Chen, and D. Yu, "Adl-mvdr: All deep learning mvdr beamformer for target speech separation," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6089–6093.
- [6] Y. Xu, Z. Zhang, M. Yu, S.-X. Zhang, and D. Yu, "Generalized spatio-temporal rnn beamformer for target speech separation," *arXiv preprint arXiv:2101.01280*, 2021.
- [7] M. Kim and H.-M. Park, "Efficient online target speech extraction using doa-constrained independent component analysis of stereo data for robust speech recognition," *Signal Processing*, vol. 117, pp. 126–137, 2015.
- [8] A. S. Subramanian, C. Weng, S. Watanabe, M. Yu, and D. Yu, "Deep learning based multi-source localization with source splitting and its effectiveness in multi-talker speech recognition," *Computer Speech & Language*, p. 101360, 2022.
- [9] J. Ajmera, G. Lathoud, and L. McCowan, "Clustering and segmenting speakers and their locations in meetings," in *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 1. IEEE, 2004, pp. I–605.
- [10] S. Otterson, "Improved location features for meeting speaker diarization," in *Eighth Annual Conference of the International Speech Communication Association*, 2007.
- [11] S. Zheng, W. Huang, X. Wang, H. Suo, J. Feng, and Z. Yan, "A real-time speaker diarization system based on spatial spectrum," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7208–7212.
- [12] D. Suvorov, G. Dong, and R. Zhukov, "Deep residual network for sound source localization in the time domain," *arXiv preprint arXiv:1808.06429*, 2018.
- [13] P. Vecchiotti, N. Ma, S. Squartini, and G. J. Brown, "End-to-end binaural sound localisation from the raw waveform," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 451–455.
- [14] G. Bologni, R. Heusdens, and J. Martinez, "Acoustic reflectors localization from stereo recordings using neural networks," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 1–5.
- [15] S. Advanne, A. Politis, and T. Virtanen, "Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network," in *2018 26th European Signal Processing Conference (EUSIPCO)*. IEEE, 2018, pp. 1462–1466.
- [16] W. He, P. Motlicek, and J.-M. Odobez, "Neural network adaptation and data augmentation for multi-speaker direction-of-arrival estimation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1303–1317, 2021.
- [17] N. Liu, H. Chen, K. Songgong, and Y. Li, "Deep learning assisted sound source localization using two orthogonal first-order differential microphone arrays," *The Journal of the Acoustical Society of America*, vol. 149, no. 2, pp. 1069–1084, 2021.
- [18] L. Perotin, R. Serizel, E. Vincent, and A. Guérin, "Crnn-based multiple doa estimation using acoustic intensity features for ambisonics recordings," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 22–33, 2019.
- [19] —, "Crnn-based joint azimuth and elevation localization with the ambisonics intensity vector," *2018 16th International Workshop on Acoustic Signal Enhancement (IWAENC)*, pp. 241–245, 2018.
- [20] X. Xiao, S. Zhao, X. Zhong, D. L. Jones, E. S. Chng, and H. Li, "A learning-based approach to direction of arrival estimation in noisy and reverberant environments," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 2814–2818.
- [21] W. He, P. Motlicek, and J.-M. Odobez, "Deep neural networks for multiple speaker detection and localization," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 74–79.
- [22] Q. Li, X. Zhang, and H. Li, "Online direction of arrival estimation based on deep learning," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 2616–2620.
- [23] R. Takeda and K. Komatani, "Discriminative multiple sound source localization based on deep neural networks using independent location model," in *2016 IEEE Spoken Language Technology Workshop (SLT)*, 2016, pp. 603–609.
- [24] W. He, P. Motlicek, and J.-M. Odobez, "Adaptation of multiple sound source localization neural networks with weak supervision and domain-adversarial training," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 770–774.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [26] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [27] N. Takahashi, S. Parthasarathy, N. Goswami, and Y. Mitsufuji, "Recursive speech separation for unknown number of speakers," *arXiv preprint arXiv:1904.03065*, 2019.
- [28] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International conference on machine learning*. PMLR, 2015, pp. 448–456.
- [29] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2011, pp. 315–323.
- [30] J. Yamagishi, C. Veaux, and K. MacDonald, "Cstr vtck corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92)," <https://doi.org/10.7488/ds/2645>, 2019.
- [31] R. Scheibler, E. Bezzam, and I. Dokmanić, "Pyroomacoustics: A python package for audio room simulation and array processing algorithms," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 351–355.
- [32] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, "SpeechBrain: A general-purpose speech toolkit," 2021, arXiv:2106.04624.
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [34] N. Chinchor, "Muc-4 evaluation metrics," in *Proceedings of the 4th Conference on Message Understanding*, ser. MUC4 '92. USA: Association for Computational Linguistics, 1992, p. 22–29. [Online]. Available: <https://doi.org/10.3115/1072064.1072067>