

End-to-end Two-dimensional Sound Source Localization With Ad-hoc Microphone Arrays

Yijun Gong, Shupeil Liu and Xiao-Lei Zhang

School of Marine Science and Technology, Northwestern Polytechnical University, China

E-mail: {gongyj, shupeil.liu}@mail.nwpu.edu.cn, xiaolei.zhang@nwpu.edu.cn

Abstract—Conventional sound source localization methods are mostly based on a single microphone array that consists of multiple microphones. They are usually formulated as the estimation of the direction of arrival problem. In this paper, we propose a deep-learning-based end-to-end sound source localization method with ad-hoc microphone arrays, where an ad-hoc microphone array is a set of randomly distributed microphone arrays that collaborate with each other. It can produce two-dimensional locations of speakers with only a single microphone per node. Specifically, we divide a targeted indoor space into multiple local areas. We encode each local area by a one-hot code, therefore, the node and speaker locations can be represented by the one-hot codes. Accordingly, the sound source localization problem is formulated as such a classification task of recognizing the one-hot code of the speaker given the one-hot codes of the microphone nodes and their speech recordings. An end-to-end spatial-temporal deep model is designed for the classification problem. It utilizes a spatial-temporal attention architecture with a fusion layer inserted in the middle of the architecture, which is able to handle arbitrarily different numbers of microphone nodes during the model training and test. Experimental results show that the proposed method yields good performance in highly reverberant and noisy environments.

I. INTRODUCTION

Sound source localization aims to find the relative position of a sound source to a microphone array by using the multi-channel acoustic signals collected by the microphone array [1]. In most cases, sound source localization with a single microphone array could be described as the estimation of the direction of arrival (DOA) of sound sources. The location information is crucial in many applications such as teleconferencing [2], [3], robot audition [4], [5], speech separation [6], [7], speech enhancement [8] and automatic speech recognition [9]. Conventional sound source localization methods mainly include time difference of arrival (TDOA) [10], steered response power with phase transform (SRP-PHAT) [11], subspace based multiple signal classification (MUSIC) [12], etc. Nevertheless, the performance of these conventional methods drops significantly in reverberant and noisy environments. Nowadays, with the development of neural networks, deep-learning-based sound source localization methods e.g. [7], [13], [14], which show promising performance in the adverse environments, become popular.

However, the aforementioned sound source localization methods are all based on a single microphone array, which

mainly obtain merely the directions of sound sources. Moreover, their performance may be unsatisfactory when the sound sources are far away from the microphone array. Ad-hoc microphone arrays may alleviate this problem. An ad-hoc microphone array is a set of randomly distributed microphone nodes with each node containing a single microphone or a microphone array. It reduces the probability of far-field speech processing by grouping the microphone nodes that are close to the sound sources into local microphone arrays. Particularly, deep-learning-based ad-hoc array speech processing shows promising performance in plenty of tasks like speech enhancement [15], speech separation [16], speech recognition [17], and speaker verification [18].

As for the sound source localization problem, an ad-hoc microphone array may be able to estimate the two- or three-dimensional coordinates of sound sources beyond the DOA estimation. Early studies of this direction apply deep learning to distributed microphone arrays [19]-[20]. Specifically, by fixing a given number of microphone nodes on walls, [19] proposed to predict the 2D coordinates of speaker positions directly. [21] partitions a room space into grids, and formulated 2D speaker localization as a problem of identifying which spatial grids the speakers belong to. [20] applies deep learning based DOA to two distributed microphone nodes respectively, and utilizes triangulation to get the 2D coordinate of a speaker. Although the above methods could cope with the two-dimensional localization task, the methods are limited to small number of microphone nodes such as two nodes, and require matched positions of the microphone nodes in both training and test. These limitations make the methods hard to generalize to mismatched scenarios, and therefore cannot reflect the merits of ad-hoc microphone arrays fully.

In order to promote the advantage of ad-hoc microphone arrays with strong generalization performance, in [22], a deep-learning-based stage-wise sound source localization with ad-hoc microphone arrays which could contain unfixed number of nodes has been proposed. It first exploits deep neural networks to get the DOA information for each ad-hoc node. Then, it uses a triangulation method to estimate a candidate sound source location from every two nodes. Finally, it conducts clustering on all of the candidate positions for the final accurate location estimation. Although the stage-wise sound source localization method has demonstrated its effectiveness in many challenging environments, it consists

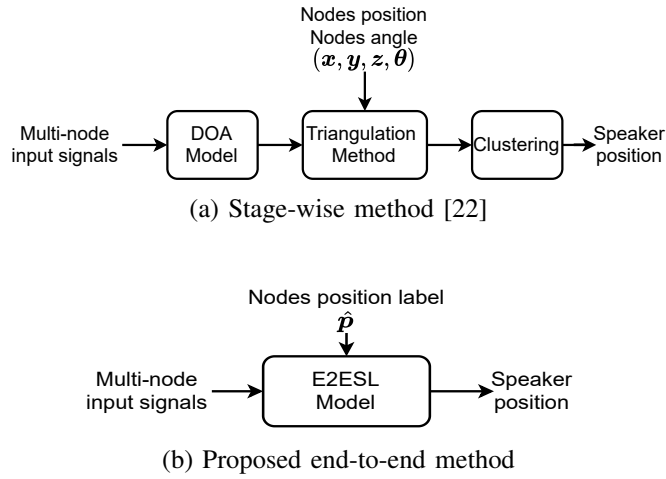


Fig. 1. Comparison of the architectures of the sound source localization methods with ad-hoc microphone arrays.

of many independent modules, which makes it inconvenient to apply in practice. Moreover, some modules are even in unsupervised manner, which do not optimize the localization performance directly. Besides, each microphone node of the ad-hoc array needs to be a conventional microphone array of multiple microphones. The directions and accurate positions of the ad-hoc nodes need to be manually labeled precisely, which is complex and time-consuming.

To deal with the aforementioned problems, in this paper, we propose an end-to-end sound source localization method (E2ESL) based on ad-hoc microphone arrays to produce the two-dimensional locations of sound sources directly. Similar to [20] and [21], we also treat the sound source localization problem as a classification problem, which is suitable to the end-to-end formulation of sound source localization. To relax the constraint of the previous studies [20], [21] that the node positions in the training and test scenarios have to be the same, we use one-hot codes to encode spatial grids, and use the one-hot codes of the spatial grids where the ad-hoc microphone nodes locate as an input feature of a deep model. The deep model adopts a spatial-temporal attention architecture. It predicts the one-hot codes of the positions of the sound sources. Fig. 1 summarizes the difference between the stage-wise method in [22] and the proposed end-to-end method. We have proved the effectiveness of our method on simulated data with noise and reverberation. Our main contributions are as follows:

- We propose a novel sound source localization method, named E2ESL, which encodes the positions of the ad-hoc microphone nodes as an input feature. This novelty makes E2ESL able to deal with the generalization problem that the locations of the ad-hoc nodes at the training and test scenarios are different.
- E2ESL is able to work fine even if each node contains only a single microphone, which is quite different from

conventional sound source localization methods where each node has to contain a conventional microphone array, e.g. a linear array. This property may simplify the deployment of E2ESL, since that it does not have to know the direction of the array at each node.

- The deep model of E2ESL applies the attention mechanism to fuse the features produced from the ad-hoc nodes, which is able to handle the case that the numbers of the ad-hoc nodes in training and test are different.

The rest of this paper is organized as follows. Section II introduces the proposed framework and the architecture of the end-to-end deep model. Experimental results are shown in Section III. Finally, Section IV draws a conclusion.

II. PROPOSED METHOD

A. Framework

We formulate the E2ESL method with ad-hoc microphone arrays as a classification problem. Given a room of size $\{X, Y, Z\}$. Suppose there are N randomly distributed nodes in the room, where each node contains a single microphone. Suppose the cartesian coordinate of the n th node in the room is $\mathbf{p}_n = \{x_n, y_n, z_n\}$, and $\mathbf{P} = \{\mathbf{p}_n | \forall n = 1, \dots, N\}$ represents the set of the coordinates of the nodes. Suppose the coordinate of a speaker is $\mathbf{p}_{\text{spkr}} = \{x_{\text{spkr}}, y_{\text{spkr}}, z_{\text{spkr}}\}$. Here we only consider the 2-dimensional (2D) coordinate estimation of the speaker, though the proposed framework is not limited to the 2D coordinate estimation. In the following, the height coordinates Z , z_n , and z_{spkr} are omitted. Suppose each node contains a single microphone, though the proposed method is not limited to this assumption. Suppose the speech signal recordings from the N nodes are $\mathbf{S} = \{\mathbf{s}_n(t) | \forall t = 1, \dots, T, \forall n = 1, \dots, N\}$, where $\mathbf{s}_n(t)$ denotes the speech signal recorded by the n th node at the t th frame.

We establish a cartesian coordinate system for the room, where we select a corner of the room as the origin. We first split the room uniformly into M local areas, as shown in Fig. 2. Then, we assign each local area an index. Given the size of the local area denoted as $\{r_x, r_y\}$, M could be obtained by:

$$M = M_x M_y = \frac{X}{r_x} \frac{Y}{r_y}. \quad (1)$$

The coordinates that the m th rectangle local area covers are in the region surrounded by the four corners of the rectangle:

$$b_{start}^{mx} = ((m-1) // M_y) \times r_x \quad (2)$$

$$b_{end}^{mx} = (((m-1) // M_y) + 1) \times r_x \quad (3)$$

$$b_{start}^{my} = ((m-1) \% M_y) \times r_y \quad (4)$$

$$b_{end}^{my} = (((m-1) \% M_y) + 1) \times r_y \quad (5)$$

where the symbols “//” and “%” are the floor division and modulo operators respectively. Given the above formulation, the position of an ad-hoc node \mathbf{p}_n is simplified to the index of the local area where the node is located, which is denoted as $\hat{p}_n \in \{1, \dots, M\}$, so as to the sound sources.

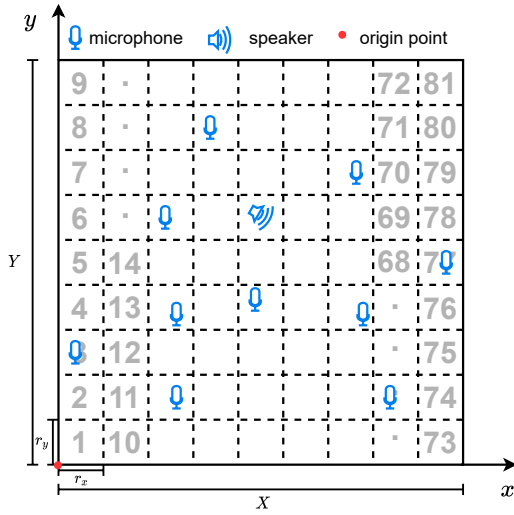


Fig. 2. Formulation of the E2ESL method with ad-hoc microphone arrays as a classification problem. The symbols r_x and r_y in the figure are the resolutions along the x and y axes which determine the size of each local area.

E2ESL takes speech recordings of the ad-hoc array S and the indexes of the positions of the ad-hoc nodes $\hat{P} = \{\hat{p}_1, \dots, \hat{p}_N\}$ as the input to estimate the index of the local area where the speaker locates $\hat{p}_{\text{spk}} \in \{1, \dots, M\}$:

$$\hat{p}_{\text{spk}} = f(\hat{P}, S) \quad (6)$$

where $f(\cdot)$ is the deep model described in the next subsections. Finally, the center of the predicted local area is regarded as the coordinates of the estimated speaker position.

B. Implementation

The general pipeline of the proposed E2ESL is given in Fig. 3. It first takes the one-hot code and short-term Fourier transform (STFT) acoustic feature of the ad-hoc node as the inputs of two encoders respectively. Then, it concatenates the outputs of the two encoders, and takes the concatenated feature as the input of a spatial-temporal attention network. The attention network conducts cross-channel attention, channel fusion, and cross-frame attention successively. Finally, it averages the output of the spatial-temporal attention network along the time axis for the estimation of the one-hot code of the local area of the speaker. The details of the proposed E2ESL are described as follow:

1) *One-hot encoding for the local areas*: Because the position indexes \hat{P} of the ad-hoc nodes are not suitable to be a feature, here we encode the indexes by the one-hot codes, though other coding methods can be applied as well. For a local area of index m , the one-hot code is an M -dimensional sparse vector with the m th element set to 1 and the other elements set to 0. By doing this, we can represent the positions of the ad-hoc nodes and speakers as the one-hot vectors. We denote the one-hot codes of the ad-hoc nodes

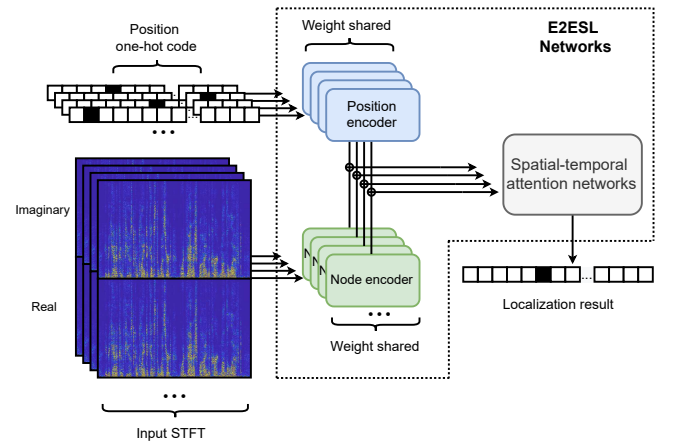


Fig. 3. Diagram of the end-to-end sound source localization algorithm.

as $\mathbf{U} = \{\mathbf{u}_n | \forall n = 1, \dots, N\}$, and the one-hot code of the speaker as \mathbf{l} .

2) *Acoustic feature extraction*: We extract time-frequency spectrograms from the multi-channel signals of the ad-hoc microphone array by STFT, and then concatenate the real and imaginary parts as the acoustic feature:

$$\begin{aligned} \hat{\mathbf{S}} &= \text{Concat}(\text{Imag}(\text{STFT}(\mathbf{S})), \text{Real}(\text{STFT}(\mathbf{S}))) \\ &= \{\hat{s}_n(f, t) | \forall f = 1, \dots, 2F, \forall n = 1, \dots, N, \forall t = 1, \dots, T\} \end{aligned} \quad (7)$$

where $s_n(f, t)$ denotes the concatenated feature of the n th node at time t and frequency f .

3) *Encoders*: We feed the acoustic feature $\hat{\mathbf{S}}$ and the one-hot vectors \mathbf{U} of the ad-hoc nodes to a *node encoder* and a *position encoder* respectively:

$$\tilde{\mathbf{s}}_n = \text{NodeEncoder}(\hat{\mathbf{s}}_n), \quad \forall n = 1, \dots, N \quad (8)$$

$$\tilde{\mathbf{u}}_n = \text{PositionEncoder}(\mathbf{u}_n), \quad \forall n = 1, \dots, N \quad (9)$$

where $\tilde{\mathbf{s}}_n$ and $\tilde{\mathbf{u}}_n$ are the *acoustic embedding* and *position embedding* of the n th ad-hoc node. The two encoders, which consist of multiple linear layers, are shared by all ad-hoc nodes.

Then, for each ad-hoc node, the position embedding $\tilde{\mathbf{u}}_n$ is first duplicated by T times, and then concatenated with $\tilde{\mathbf{s}}_n$:

$$\mathbf{e}_n = \text{Concat}(\text{Duplicate}(\tilde{\mathbf{u}}_n), \tilde{\mathbf{s}}_n), \quad \forall n = 1, \dots, N. \quad (10)$$

We denote $\mathbf{E} = \{\mathbf{e}_n \in \mathbb{R}^{D \times T} | \forall n = 1, \dots, N\}$ where D is the dimension of \mathbf{e}_n .

4) *Spatial-temporal attention network*: As shown in Fig. 4(a), the spatial-temporal attention network consists of three main parts: two stacked spatial layers, a fusion layer and a temporal layer. Each layer consists a multi-head attention sub-layer and a feed forward sub-layer. Besides, the layer normalization and residual connection are also applied.

Spatial layer: We first feed \mathbf{E} into the spatial layer frame by frame for the cross-channel attention, which learns a new

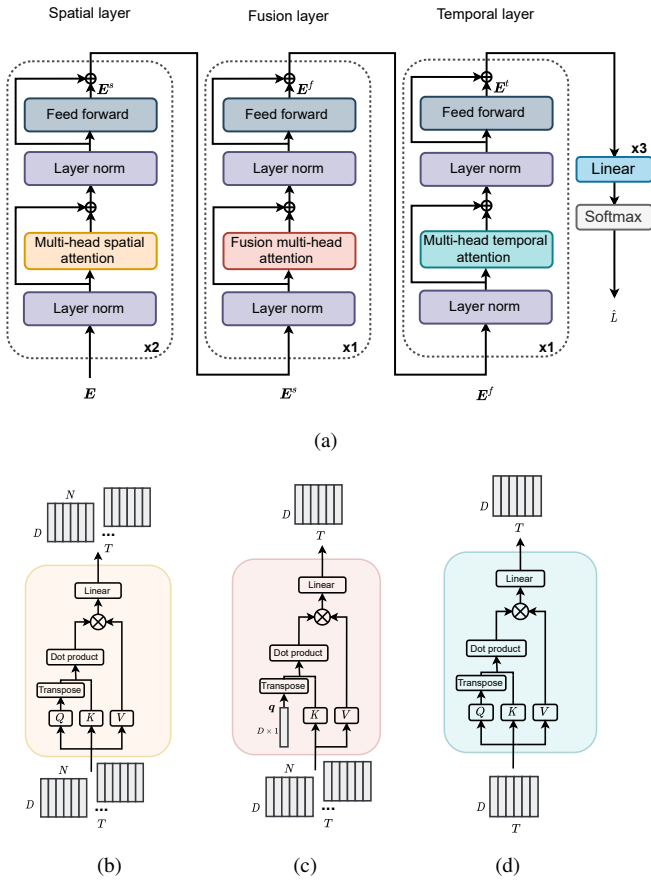


Fig. 4. Spatial-temporal attention network. (a) General architecture. (b) Spatial layer. (c) Fusion layer. (d) Temporal layer.

multi-channel feature that contains the global spatial connection between the ad-hoc nodes, denoted as $\mathbf{E}^s \in \mathbb{R}^{T \times D \times N}$:

$$\begin{aligned} \mathbf{E}^s &= \text{FFN}(\mathbf{E} + \text{MHA}(\mathbf{E})) \\ &= \text{Linear}(\text{Linear}(\mathbf{E} + \text{MHA}(\mathbf{E}))) \end{aligned} \quad (11)$$

where FFN and MHA denote the feed-forward module and multi-head attention networks. FFN consists of two linear layers as shown in (11), and MHA is defined as:

$$\text{MHA}(\mathbf{E}) = \text{Concat}(\mathbf{O}_1^s, \dots, \mathbf{O}_h^s, \dots, \mathbf{O}_H^s) \mathbf{W}^o \quad (12)$$

where H denotes the number of heads, $\mathbf{W}^o \in \mathbb{R}^{D \times D}$ is the output projection matrix, and \mathbf{O}_h^s with $h = 1, \dots, H$ is the output of each attention head:

$$\mathbf{O}_h^s = \text{Softmax} \left(\frac{(\mathbf{W}_h^q \mathbf{E})^T (\mathbf{W}_h^k \mathbf{E})}{\sqrt{D/H}} \right) (\mathbf{W}_h^v \mathbf{E}), \quad \forall h = 1, \dots, H \quad (13)$$

where $\mathbf{W}_h^q \in \mathbb{R}^{D/H \times N}$, $\mathbf{W}_h^k \in \mathbb{R}^{D/H \times N}$, and $\mathbf{W}_h^v \in \mathbb{R}^{D/H \times N}$ are learnable projection matrices.

Fusion layer: After conducting the multi-head attention across channels by the spatial layer, the multi-channel features produced from the spatial layer are fused into a single-channel feature through the fusion layer. As shown in Fig. 4(c), the

attention module in the fusion layer is the same as the standard multi-head attention in the spatial layer except that $\mathbf{W}_h^q \mathbf{E}^s$ is replaced by a learnable vector $\mathbf{q} \in \mathbb{R}^{D \times 1}$. The output of the fusion layer \mathbf{E}^f

$$\mathbf{E}^f = \text{FFN}(\mathbf{E}^s + \text{FMHA}(\mathbf{E}^s)) \quad (14)$$

where FMHA is the fusion multi-head attention:

$$\text{FMHA}(\mathbf{E}^s) = \text{Concat}(\mathbf{O}_1^f, \dots, \mathbf{O}_h^f, \dots, \mathbf{O}_H^f) \mathbf{W}^o \quad (15)$$

with \mathbf{O}_h^f calculated by:

$$\mathbf{O}_h^f = \text{Softmax} \left(\frac{\mathbf{q}^T (\mathbf{W}_h^k \mathbf{E}^s)}{\sqrt{D/H}} \right) (\mathbf{W}_h^v \mathbf{E}^s), \quad \forall h = 1, \dots, H \quad (16)$$

Temporal layer: We feed \mathbf{E}^f into the temporal layer for the cross-frame attention. The architecture of the multi-head attention module in the temporal layer is the same as that in the spatial layer. The output of the temporal layer $\mathbf{E}^t \in \mathbb{R}^{D \times T}$ is calculated by:

$$\mathbf{E}^t = \text{FFN}(\mathbf{E}^f + \text{MHA}(\mathbf{E}^f)). \quad (17)$$

Finally, we obtain the classification result $\hat{\mathbf{L}} \in \mathbb{R}^{T \times M}$ through a stack of linear layers and a softmax output layer:

$$\hat{\mathbf{L}} = \text{Softmax}(\text{Linear}(\mathbf{E}^t)). \quad (18)$$

5) Postprocessing and coordinate estimation of speakers:

We average $\hat{\mathbf{L}}$ over the time dimension for an estimation of the one-hot code of the speaker $\hat{\mathbf{l}} \in \mathbb{R}^M$. We conduct a hard decision to $\hat{\mathbf{l}} \in \mathbb{R}^M$, which selects the index of the class with the highest probability as the estimated local area \hat{l} . The estimated coordinate of the speaker $\tilde{\mathbf{p}}_{\text{spkr}} = \{\tilde{x}_{\text{spkr}}, \tilde{y}_{\text{spkr}}\}$ is obtained as the center of the local area:

$$\tilde{x}_{\text{spkr}} = \frac{1}{2} [(((\hat{l}-1)/M_y)+1) \times r_x + ((\hat{l}-1)/M_y) \times r_x] \quad (19)$$

$$\tilde{y}_{\text{spkr}} = \frac{1}{2} [(((\hat{l}-1)\%M_y)+1) \times r_y + ((\hat{l}-1)\%M_y) \times r_y]. \quad (20)$$

III. EXPERIMENTS

A. Experimental setup

We conducted an experiment on a simulated dataset. The source speech is from the TIMIT corpus [23], which consists of 462 speakers in the training set and 168 speakers in the test set. We first removed silent segments of the speech recordings by voice activity detection. Then, we cut each recording into segments of two-second long, and discard the segments that are shorter than two seconds. For each simulated speech recording for training, the source speech is a randomly chosen segment from a speaker in the training set of TIMIT, so as to the simulated test data. The image source model [24]¹ is exploited to generate reverberant environments. Diffused white noise is added by the noise generators.² The

¹ <https://github.com/ehabets/RIR-Generator>

² <https://www.audiolabs-erlangen.de/fau/professor/habets/software/noise-generators>

TABLE I
ROOM SIZE IN TRAINING AND TEST SET

Dataset	Room	Room size
Training set	Room1	5 × 7 × 3 m
	Room2	10 × 10 × 3 m
	Room3	7 × 9 × 3 m
	Room4	8 × 8 × 3 m
	Room5	8 × 6 × 3 m
	Room6	5 × 5 × 3 m
	Room7	7 × 4 × 3 m
	Room8	6 × 7 × 3 m
	Room9	5 × 4 × 3 m
	Room10	9 × 8 × 3 m
Test set	Room1	5 × 6 × 3 m
	Room2	7 × 8 × 3 m
	Room3	5 × 8 × 3 m
	Room4	6 × 6 × 3 m
	Room5	6 × 9 × 3 m

reverberation time T_{60} for each sample is chosen randomly from a range of [0.3, 1.0] second. The SNR at the position of the speaker is randomly selected from a range of [20, 50] dB.

For the simulated training data, we generated 10 rooms with different sizes. Each room was divided into 225 local areas of equal size. We randomly selected 36 positions in each local area as the speaker position. For each speaker position, we randomly picked a segment from the training set of TIMIT. Eventually, we generated 81,000 source samples for training with each sample being 2 seconds long. The duration of the training set is 45 hours long. For each sample, we generated an ad-hoc microphone array of 30 microphone nodes, where the 30 nodes were randomly put into 30 local areas. The microphone position was randomly chosen inside the local area.

The TIMIT test set was partitioned into two parts, which were used to construct the simulated validation data and test data respectively. For the simulated test data, we simulated 5 rooms, each of which also contains 225 local areas. We randomly selected 5 speaker positions in each local area, each of which corresponds to a randomly selected test sample. The other settings were the same as the training set. The room sizes of the simulated training and test data are listed in Table I, which as we can see are different.

The sampling rate of the simulated data is 16KHz. We conducted 512 points STFT for each frame with a frame overlap of 50%. The details of the network structure are summarized in Table II. For each layer, ReLU was applied as the activation function. In the training process, we used minimum cross entropy as the loss function, and used Adam as the optimizer.

We used mean absolute error (MAE) to evaluate the performance of the proposed algorithm:

$$MAE = \frac{1}{I} \sum_{i=1}^I \sqrt{(y_{\text{spkr},i} - \tilde{y}_{\text{spkr},i})^2 + (x_{\text{spkr},i} - \tilde{x}_{\text{spkr},i})^2} \quad (21)$$

where I is the number of the test samples. The lower the MAE is, the better the performance is.

TABLE II
MODEL STRUCTURE OF THE PROPOSED E2ESL.

Module	Network	Input size	Output size
Position encoder	Linear	M	512
	Linear	512	256
Node encoder	Linear	512	1024
	Linear	1024	512
	Linear	512	256
Spatial layer	Self-attention*	512	512
	Feed forward	512	512
Fusion layer	Self-attention*	512	512
	Feed forward	512	512
Temporal layer	Self-attention*	512	512
	Feed forward	512	512
Output layer	Linear	512	1024
	Linear	1024	1024
	Linear	1024	M

* We used 4 heads attention in our experiments.

TABLE III
MAE (IN METERS) OF THE PROPOSED METHOD ON THE SIMULATED TEST DATA WITH DIFFERENT NUMBERS OF AD-HOC NODES.

Nodes	Room1	Room2	Room3	Room4	Room5	Average
30	0.2271	0.2838	0.2785	0.2292	0.3127	0.2662
29	0.2208	0.2873	0.2804	0.2282	0.3168	0.2667
28	0.2221	0.2857	0.2831	0.2350	0.3227	0.2697
27	0.2221	0.2917	0.2907	0.2360	0.3278	0.2736
26	0.2336	0.3012	0.2986	0.2413	0.3357	0.2820
25	0.2454	0.3073	0.3024	0.2467	0.3391	0.2881
24	0.2487	0.3120	0.3167	0.2519	0.3432	0.2945
23	0.2567	0.3187	0.3174	0.2572	0.3577	0.3015
22	0.2575	0.3123	0.3258	0.2685	0.3738	0.3075
21	0.2703	0.3337	0.3320	0.2736	0.3837	0.3186
20	0.2735	0.3327	0.3437	0.2887	0.4075	0.3292

B. Main results

We trained E2ESL by the 45 hours training data with 30 nodes, and tested the model on the 5 test rooms with the number of the ad-hoc nodes varying from 20 to 30. When the number of the ad-hoc nodes in a test scenario is less than 30, the nodes are randomly selected from the 30 nodes.

The results are shown in Table III. From the table, we see that, with the increasing of the number of nodes, MAE decreases in all 5 test rooms. For example, the average MAE in the scenario with 30 nodes is 6 cm lower than that with 20 nodes.

Note that the reason why we do not compare with other sound source location methods, e.g. [22], is that our test scenario where each ad-hoc node contains only a single microphone makes the methods unapplicable.

C. Effect of the size of the local area

It is obvious that the size of the local area affect the performance of E2ESL. To exam its effect, we split a room into 100 or 64 local areas with an equal size respectively. We compared the E2ESL models trained with 225, 100 and 64 local areas in room1 of the simulated test data. Table IV shows the experimental result where ACC denotes the classification accuracy. From the table, we find interestingly that, when the resolution of the room is improved from 64 to 225 local areas,

TABLE IV
RESULTS ON TEST ROOM1 WITH DIFFERENT LOCAL AREA NUMBERS

225 local areas		100 local areas		64 local areas	
MAE (m)	ACC (%)	MAE (m)	ACC (%)	MAE (m)	ACC (%)
0.2271	52.44	0.2828	62.25	0.3285	70.01

although the classification accuracy decreases from 70.01% to 52.44%, the localization performance by MAE is improved from 0.3285 meters to 0.2271 meters.

IV. CONCLUSION

In this paper, we have proposed an end-to-end sound source localization method with ad-hoc microphone arrays where each ad-hoc node contains a single microphone. With a room splitting strategy and one-hot encoding, we leverage location information into the classification based end-to-end sound source localization system, which only needs rough indexes of the local areas as the position features of the ad-hoc nodes. A spatial-temporal attention network is utilized to capture the global information of speakers and acoustic environments in the dimensions of both space and time. We demonstrated the effectiveness of the proposed algorithm in the simulated data with additive noise and strong reverberation, given the situation where each ad-hoc node contains only a single microphone.

V. ACKNOWLEDGEMENT

This work was supported in part by National Science Foundation of China under Grant No. 62176211, in part by Project of the Science, Technology, and Innovation Commission of Shenzhen Municipality under grant No. JSGG20210802152546026 and JCYJ20210324143006016.

REFERENCES

- [1] Pierre-Amaury Grumiaux, Srđan Kitić, Laurent Girin, and Alexandre Guérin, "A survey of sound source localization with deep learning methods," *The Journal of the Acoustical Society of America*, vol. 152, no. 1, pp. 107–151, 2022.
- [2] Hong Wang and Peter Chu, "Voice source localization for automatic camera pointing system in videoconferencing," in *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*. IEEE, 1997, vol. 1, pp. 187–190.
- [3] Marius H Hennecke and Gernot A Fink, "Towards acoustic self-localization of ad hoc smartphone arrays," in *2011 Joint Workshop on Hands-free Speech Communication and Microphone Arrays*. IEEE, 2011, pp. 127–132.
- [4] J-M Valin, François Michaud, Brahim Hadjou, and Jean Rouat, "Localization of simultaneous moving sound sources for mobile robot using a frequency-domain steered beamformer approach," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*. IEEE, 2004, vol. 1, pp. 1033–1038.
- [5] Xiaofei Li, Laurent Girin, Fabien Badeig, and Radu Horaud, "Reverberant sound localization with a robot head based on direct-path relative transfer function," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 2819–2826.
- [6] Shlomo E Chazan, Hodaya Hammer, Gershon Hazan, Jacob Goldberger, and Sharon Gannot, "Multi-microphone speaker separation based on deep doa estimation," in *2019 27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019, pp. 1–5.

- [7] Yu Chen, Bowen Liu, Zijian Zhang, and Hun-Seok Kim, "An end-to-end deep learning framework for multiple audio source separation and localization," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 736–740.
- [8] Angeliki Xenaki, Jesper Bünsow Boldt, and Mads Græsbøll Christensen, "Sound source localization and speech enhancement with sparse bayesian learning beamforming," *The Journal of the Acoustical Society of America*, vol. 143, no. 6, pp. 3912–3921, 2018.
- [9] Ho-Yong Lee, Ji-Won Cho, Minook Kim, and Hyung-Min Park, "Dnn-based feature enhancement using doa-constrained ica for robust speech recognition," *IEEE Signal Processing Letters*, vol. 23, no. 8, pp. 1091–1095, 2016.
- [10] Charles Knapp and Glifford Carter, "The generalized correlation method for estimation of time delay," *IEEE transactions on acoustics, speech, and signal processing*, vol. 24, no. 4, pp. 320–327, 1976.
- [11] Joseph Hector DiBiase, *A high-accuracy, low-latency technique for talker localization in reverberant environments using microphone arrays*. Brown University, 2000.
- [12] Ralph Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE transactions on antennas and propagation*, vol. 34, no. 3, pp. 276–280, 1986.
- [13] Soumitro Chakrabarty and Emanuël A. P. Habets, "Multi-speaker doa estimation using deep convolutional networks trained with noise signals," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 8–21, 2019.
- [14] Thi Ngoc Tho Nguyen, Woon-Seng Gan, Rishabh Ranjan, and Douglas L Jones, "Robust source counting and doa estimation using spatial pseudo-spectrum and convolutional neural network," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2626–2637, 2020.
- [15] Xiao-Lei Zhang, "Deep ad-hoc beamforming," *Computer Speech & Language*, vol. 68, pp. 101201, 2021.
- [16] Ziye Yang, Shanzheng Guan, and Xiao-Lei Zhang, "Deep ad-hoc beamforming based on speaker extraction for target-dependent speech separation," *Speech Communication*, vol. 140, pp. 87–97, 2022.
- [17] Junqi Chen and Xiao-Lei Zhang, "Scaling Sparsemax Based Channel Selection for Speech Recognition with ad-hoc Microphone Arrays," in *Proc. Interspeech 2021*, 2021, pp. 291–295.
- [18] Chengdong Liang, Junqi Chen, Shanzheng Guan, and Xiao-Lei Zhang, "Attention-based multi-channel speaker verification with ad-hoc microphone arrays," in *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference, APSIPA ASC 2021*. 2021, pp. 1111–1115, IEEE.
- [19] Fabio Vesperini, Paolo Vecchiotti, Emanuele Principi, Stefano Squartini, and Francesco Piazza, "Localizing speakers in multiple rooms by using deep neural networks," *Computer Speech & Language*, vol. 49, pp. 83–106, 2018.
- [20] Stijn Kindt, Alexander Bohlender, and Nilesh Madhu, "2d acoustic source localisation using decentralised deep neural networks on distributed microphone arrays," in *Speech Communication; 14th ITG Conference*. VDE, 2021, pp. 1–5.
- [21] Guillaume Le Moing, Phongtharin Vinayavekhin, Tadanobu Inoue, Jayakorn Vongkulbhisal, Asim Munawar, Ryuki Tachibana, and Don Joven Agravante, "Learning multiple sound source 2d localization," in *2019 IEEE 21st International Workshop on Multimedia Signal Processing (MMSP)*. IEEE, 2019, pp. 1–6.
- [22] Shupe Li, Yijun Gong, Xiao-Lei Zhang, and Xuelong Li, "Deep learning based two-dimensional speaker localization with large ad-hoc microphone arrays," submitted for publication, to be available online.
- [23] John S Garofolo, Lori F Lamel, William M Fisher, Jonathan G Fiscus, and David S Pallett, "Darpa timit acoustic-phonetic continuous speech corpus cd-rom. nist speech disc 1-1.1," *NASA STI/Recon technical report n*, vol. 93, pp. 27403, 1993.
- [24] Jont B Allen and David A Berkley, "Image method for efficiently simulating small-room acoustics," *The Journal of the Acoustical Society of America*, vol. 65, no. 4, pp. 943–950, 1979.