



Multi-Channel Far-Field Speaker Verification with Large-Scale Ad-hoc Microphone Arrays

Chengdong Liang, Yijiang Chen, Jiadi Yao, Xiao-Lei Zhang

CIAIC, School of Marine Science and Technology, Northwestern Polytechnical University, Xi'an, China

liangchengdong@mail.nwpu.edu.cn, xiaolei.zhang@nwpu.edu.cn

Abstract

Speaker verification based on ad-hoc microphone arrays has the potential of reducing the error significantly in adverse acoustic environments. However, existing approaches extract utterance-level speaker embeddings from each channel of an ad-hoc microphone array, which does not consider fully the spatial-temporal information across the devices. In this paper, we propose to aggregate the multichannel signals of the ad-hoc microphone array at the frame-level by exploring the cross-channel information deeply with two attention mechanisms. The first one is a self-attention method. It consists of a cross-frame self-attention layer and a cross-channel self-attention layer successively, both working at the frame level. The second one learns the cross-frame and cross-channel information via two graph attention layers. Experimental results demonstrate that the proposed methods reach the state-of-the-art performance. Moreover, the graph-attention method is better than the self-attention method in most cases.

Index Terms: ad-hoc microphone arrays, multi-channel speaker verification, far-field speaker verification

1. Introduction

With the increasing popularity of smart terminal devices, such as vehicle-mounted devices, smart speakers at home, etc., far-field automatic speaker verification (ASV) is commonly required. However, due to the noise, reverberation and speech signal attenuation in far-field environments, the performance of single-channel ASV drops sharply. To make the smart devices robust in the noise and reverberant environments, many single-channel and multi-channel speech enhancement based on a single device has been extensively studied [1, 2, 3, 4, 5]. For example, beamforming, which explores the spectral and spatial diversity of the target and interference signals via multiple microphones at a single device [6, 7, 8], leads to substantial performance improvement of ASV.

Recently, a novel kind of front-ends, named ad-hoc microphone arrays, have received attention as an alternative solution to the far-field speech processing. Unlike the fixed microphone array at a single device, an ad-hoc microphone array is composed of a group of microphone devices/nodes randomly distributed in space, where the number and arrangement of the nodes may be unknown in practice [9]. It allows users to use their own devices to virtually form a microphone array system flexibly. It has the potential to reduce the occurrence probability of the far-field problem. It also captures more spatial and temporal information than the fixed microphone array at a single device.

Recently, several studies on ad-hoc microphone arrays have been conducted. In [10], the authors proposed deep ad-hoc

beamforming based on speaker extraction for speech separation. It uses a supervised channel selection framework and a deep learning based MVDR algorithm to extract the targeted speech from a mixed speech of multiple speakers. In [11, 12], the authors used the attention mechanism to obtain the relevant information between and within channels for multi-channel speech separation. In [13], the authors released a large-scale ad-hoc data set and a baseline system for speech recognition. In [14, 15], the authors applied neural transformer architectures for multi-channel speech recognition systems, where the multichannel speech signals are aggregated by attention layers. In [16], the authors further conducted channel selection by a scaling sparsemax operator for speech recognition with large ad-hoc arrays.

With regard to the ASV study with ad-hoc microphone arrays, [17, 18] independently proposed to aggregate the multichannel speech signals at the utterance-level. They add an aggregation layer above the pooling layer to fuse the speaker embeddings from different channels. The difference between them is mainly on how to aggregate the speaker embeddings. [17] used cross-channel self-attention mechanism, while [18] adopted attentive pooling layer. An apparent drawback of the utterance-level aggregation methods in [17, 18] is that it misses the spatial-temporal information between the channels.

To address this problem, in this paper, we propose to aggregate the multichannel signals of ad-hoc microphone arrays at the frame-level by adding the multi-channel aggregation module before the pooling layer. We propose two multi-channel aggregation modules. The first one is the *self-attention based aggregation* (SA-aggregation), in which cross-frame and cross-channel self-attentions are stacked to capture the temporal and spatial information respectively. Unlike [12, 11], our SA-aggregation does not adopt feed-forward networks (FFN), since that they empirically degrade performance. A weakness of SA-aggregation is that it builds the connections between acoustic features indirectly, where the attention weight is calculated as the dot product of a feature and a projection vector instead of the calculation between the features. As we know, the direct connections between the features are important information for the multi-channel aggregation.

Inspired by the recent success of graph neural networks [19, 20], the second multi-channel aggregation module is based on graph attention (GAT-aggregation). It consists of a stack of the cross-frame and cross-channel graph attention layers. It obtains a single representation from the correlation between all possible microphone node pairs. It overcomes the weakness of SA-aggregation.

We conducted an extensive experiment on a simulated corpus generated from Librispeech as well as the semi-real Libri-adhoc40 [13] corpus. Experimental results with ad-hoc microphone arrays of as many as 40 nodes demonstrate that the

proposed SA-aggregation and GAT-aggregation outperform the utterance-level methods in both data. With the increase of the number of ad-hoc nodes, the advantage of the proposed frame-level methods over the utterance-level methods is enlarged. Moreover, GAT-aggregation achieves better performance than SA-aggregation. For example, the best equal error rate (EER) of GAT-aggregation is 6.86%, which is 6.81% relatively lower than that of SA-aggregation. The proposed methods trained with the simulated data show better generalization to real scenarios than the utterance-level methods.

2. Proposed method

Figure 1 shows the architecture of the proposed multi-channel frame-level speaker verification model with ad-hoc microphone arrays. It consists of two elements: a frame-level feature processor and a multi-channel aggregation module.

Note that, although some mathematical symbols in Sections 2.1 and 2.2 are used duplicately across the subsections, they have different definitions in different subsections, since that the two subsections present the SA-aggregation and GAT-aggregation independently. The reason why we used the same symbols is to make readers easily understand the two algorithms comparatively which share a similar procedure.

We denote T as the number of frames, C as the number of input channels, and D as the feature dimension of each channel. The input of the multi-channel aggregation module is denoted as $\mathbf{X} = [\mathbf{X}_0, \dots, \mathbf{X}_{C-1}]$, where $\mathbf{X}_c \in \mathbb{R}^{T \times D}$ is the input feature matrix of the c -th channel. Let M denote the number of the attention heads.

2.1. Aggregation module based on self-attention

As shown in Figure 2, the aggregation module based on self-attention consists of a cross-frame self-attention layer and a cross-channel self-attention layer. The entire network can make use of information across channels and frames. The detailed calculation process is as follows.

2.1.1. Cross-frame self-attention layer

For each attention head, the input features are transformed into query (Q), key (K) and value (V) subspaces of dimension E as follows:

$$\mathbf{Q}_c^m = \mathbf{X}_c \mathbf{W}_Q^m, \mathbf{K}_c^m = \mathbf{X}_c \mathbf{W}_K^m, \mathbf{V}_c^m = \mathbf{X}_c \mathbf{W}_V^m \quad (1)$$

where $d_k = E/K$, the matrices \mathbf{Q}_c^m , \mathbf{K}_c^m , \mathbf{V}_c^m denote the query, key, and value embeddings respectively, all of which are in $\mathbb{R}^{T \times d_k}$. For the m -th attention head at channel c , $\mathbf{W}_*^m \in \mathbb{R}^{D \times d_k}$ are the trainable parameters where $*$ $\in \{K, Q, V\}$. The similarity matrix is computed as the product of the query and key matrices. The output of the m -th attention head is then computed by:

$$\mathbf{H}_c^m = \text{softmax} \left(\frac{\mathbf{Q}_c^m \cdot (\mathbf{K}_c^m)^\top}{\sqrt{d_k}} \right) \mathbf{V}_c^m \quad (2)$$

where $\mathbf{H}_c^m \in \mathbb{R}^{T \times d_k}$. Finally, the output of all attention heads are concatenated across the subspaces by:

$$\mathbf{Y}_c = \text{concat} [\mathbf{H}_c^1, \mathbf{H}_c^2, \dots, \mathbf{H}_c^M] \mathbf{W}_O \quad (3)$$

where $\mathbf{Y}_c \in \mathbb{R}^{T \times D}$ and $\mathbf{W}_O \in \mathbb{R}^{E \times D}$ is a weight matrix of the linear projection layer.

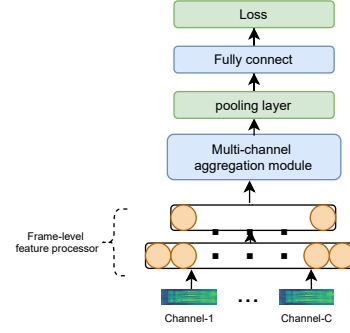


Figure 1: Architecture of the proposed multi-channel ASV system. The blue block is fine-tuned with ad-hoc data.

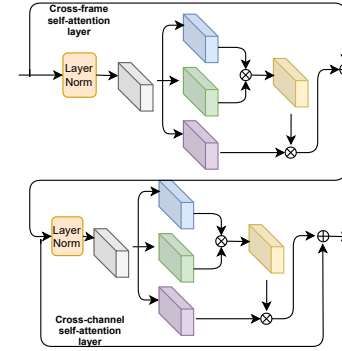


Figure 2: Architecture of the multi-channel aggregation based on self-attention.

2.1.2. Cross-channel self-attention layer

Then, we transform $\mathbf{Y} = [\mathbf{Y}_0, \dots, \mathbf{Y}_{C-1}]$ to $\hat{\mathbf{X}} = [\hat{\mathbf{X}}_0, \dots, \hat{\mathbf{X}}_{T-1}]$ as the input of the cross-channel self-attention, where $\hat{\mathbf{X}}_t \in \mathbb{R}^{C \times D}$. Similar to the cross-frame self-attention layer, the cross-channel self-attention can be denoted as follows:

$$\begin{aligned} \hat{\mathbf{H}}_t^m &= \text{softmax} \left(\frac{\mathbf{Q}_t^m \cdot (\mathbf{K}_t^m)^\top}{\sqrt{d_k}} \right) \mathbf{V}_t^m \\ \hat{\mathbf{Y}}_t &= \text{concat} [\hat{\mathbf{H}}_t^1, \hat{\mathbf{H}}_t^2, \dots, \hat{\mathbf{H}}_t^M] \hat{\mathbf{W}}_O \end{aligned} \quad (4)$$

where

$$\mathbf{Q}_t^m = \hat{\mathbf{X}}_t \hat{\mathbf{W}}_Q^m, \mathbf{K}_t^m = \hat{\mathbf{X}}_t \hat{\mathbf{W}}_K^m, \mathbf{V}_t^m = \hat{\mathbf{X}}_t \hat{\mathbf{W}}_V^m \quad (5)$$

As shown in Figure 2, Layer Normalization [21] is applied to the input of the two self-attention modules separately. A residual connection [22] is applied between the input and output of the self-attention module to alleviate the vanishing gradient problem [17].

2.2. Aggregation module based on graph attention

As shown in Figure 3, we use GAT-aggregation instead of SA-aggregation. GAT-aggregation consists of a cross-frame GAT layer and a cross-channel GAT layer. The detailed calculation process is as follows.

2.2.1. Cross-frame graph attention layer

For the cross-frame GAT layer, we first formulate a graph using the frame-level features. Specially, the input of the cross-frame

GAT layer is $\mathbf{X} = [\mathbf{X}_0, \dots, \mathbf{X}_{C-1}]$, where $\mathbf{X}_c \in \mathbb{R}^{T \times D}$ and $\mathbf{X}_c = [\mathbf{x}_0, \dots, \mathbf{x}_{T-1}]$. Let \mathbb{G}_1 be a complete graph comprising of T nodes. A set of nodes in \mathbb{G}_1 is defined as \mathbf{X}_c . Each node is represented as a row vector \mathbf{x}_t , which is a D -dimensional feature. For m -th attention head, \mathbf{X}_c is first projected into a d_k -dimensional space as follows:

$$\mathbf{g}_l^m = \mathbf{X}_c \mathbf{W}_l^m, \mathbf{g}_r^m = \mathbf{X}_c \mathbf{W}_r^m, \quad (6)$$

where $\mathbf{g}_l^m, \mathbf{g}_r^m \in \mathbb{R}^{T \times d_k}$, and $\mathbf{W}_l^m, \mathbf{W}_r^m \in \mathbb{R}^{D \times d_k}$. Then, we calculate the attention scores by:

$$e_{jk}^m = \beta^\top \text{LeakyReLU}(\text{concat}(\mathbf{g}_{li}^m, \mathbf{g}_{rj}^m)), \quad (7)$$

$$a_{ij}^m = \text{softmax}(e_{ij}^m) = \frac{\exp(e_{ij}^m)}{\sum_{n=1}^N \exp(e_{in}^m)},$$

where a_{ij}^m is the attention score between the i -th node and the j -th node ($\mathbf{g}_{li}^m, \mathbf{g}_{rj}^m \in \mathbb{R}^{d_k}, i \neq j$). $\beta \in \mathbb{R}^{d_k \times 1}$ is a learnable parameter. Then we calculate the output of each head as follows:

$$\mathbf{h}_i^m = \sum_{j=1}^N a_{ij}^m \mathbf{g}_{rj}^m. \quad (8)$$

where $\mathbf{h}_i^m \in \mathbb{R}^{d_k}$. Finally, we denote $\mathbf{H}_c^k = [\mathbf{h}_0^k, \mathbf{h}_1^k, \dots, \mathbf{h}_{T-1}^k]$, and concatenate the outputs of all heads as follows:

$$\mathbf{Y}_c = \text{concat}[\mathbf{H}_c^1, \mathbf{H}_c^2, \dots, \mathbf{H}_c^M]. \quad (9)$$

2.2.2. Cross-channel graph attention layer

We transform $\mathbf{Y} = [\mathbf{Y}_0, \dots, \mathbf{Y}_{C-1}]$ to $\hat{\mathbf{X}} = [\hat{\mathbf{X}}_0, \dots, \hat{\mathbf{X}}_{T-1}]$ as the input of the cross-channel GAT layer. Similar to the cross-frame GAT layer, we conduct the following successively steps:

$$\mathbf{g}_l^m = \hat{\mathbf{X}}_t \mathbf{W}_l^m, \mathbf{g}_r^m = \hat{\mathbf{X}}_t \mathbf{W}_r^m,$$

$$e_{jk}^m = \beta^\top \text{LeakyReLU}(\text{concat}(\mathbf{g}_{li}^m, \mathbf{g}_{rj}^m)),$$

$$a_{ij}^m = \text{softmax}(e_{ij}^m) = \frac{\exp(e_{ij}^m)}{\sum_{n=1}^N \exp(e_{in}^m)}, \quad (10)$$

$$\mathbf{h}_i^m = \sum_{j=1}^N a_{ij}^m \mathbf{g}_{rj}^m.$$

Then, we get the output of the cross-channel GAT layer as follows:

$$\mathbf{Y}_c = \text{concat}[\mathbf{H}_t^1, \mathbf{H}_t^2, \dots, \mathbf{H}_t^h] \quad (11)$$

As shown in Figure 3, similar to SA-aggregation, we also add the Layer Normalization [21] before the two graph attention layers separately. The residual connection [22] is applied between the input and output of the graph attention layers.

3. Experiments

3.1. Dataset

Our experiments use three data sets, which are the Librispeech corpus [25], Librispeech simulated with ad-hoc microphone arrays (Libri-adhoc-simu), and Libri-adhoc40 [13]. Each node of the ad-hoc microphone arrays of Libri-adhoc-simu and Libri-adhoc40 has only one microphone. Therefore, a channel refers to a node in the remaining of the paper.

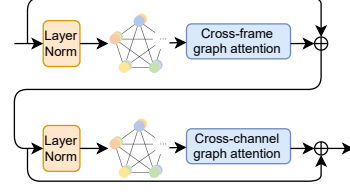


Figure 3: Architecture of the multi-channel aggregation based on graph attention.

The Libri-adhoc-simu corpus is a simulation database of Librispeech. We used 'train-clean-100', 'dev-clean' and 'test-clean' as the training, validation and test sets of the simulated ad-hoc data respectively. The training set contains 251 speakers. The validation and test sets contain 40 speakers respectively. We added room impulse response and noise to the clean speech data [17]. The size of the simulation room is [10, 10, 4] meters and the range of the reverberation time is [0.5 – 1.2] seconds. We randomly placed a speaker and forty microphones in the room. The noise for the training and validation sets was randomly selected from a large-scale noise segments library [26]. The noise of the test set comes from the CHiME-3 dataset [27] and NOISEX-92 [28]. The RIR-Generator¹ and ANF-Generator² were used for the data simulation. The Libri-adhoc40 corpus was collected via playing the speech data of Librispeech in a large room, in which forty microphones and a speaker were placed [13].

Considering that a large amount of data from a massive ad-hoc array, denoted as ad-hoc data for short, leads to a large memory requirement for the model training, all comparison multi-channel ASV models based on ad-hoc microphone arrays first trained a single-channel ASV with clean speech data, then used the single channel ASV to initialize the multi-channel ASV model, and finally used the ad-hoc data to fine-tune the multi-channel aggregation block. In our experiments, the single channel ASV systems were trained with 960 hours of the clean data of Librispeech, with another 10 hours of the clean data was used for development. To study the generalization ability of the proposed systems, we fine-tuned the systems using Libri-adhoc-simu and Libri-adhoc40 respectively. For each epoch of fine-tuning with a dataset, we randomly selected 20 channels from the training and validation sets respectively. In the test stage, we only considered the Libri-adhoc40 corpus, and randomly selected 8, 16, 32 and 40 channels to construct four test scenarios.

3.2. Experimental setup

For the proposed SA-aggregation and GAT-aggregation, the network structure of their initial single-channel ASV is the same as in [29], which contains three main components: a front-end residual convolution neural network (ResNet) [22], a self-attentive pooling (SAP) [30] layer and a fully-connected layer. It was trained for 200 epochs on the Librispeech corpus. Then, the parameters of the ResNet layer and SAP layer were fixed and sent to the proposed multi-channel ASV. Finally, we fine-tuned the multi-channel aggregation block with Libri-adhoc-simu and Libri-adhoc40 data respectively, where the number of the aggregation block is 1, and the number of the attention

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

²<https://github.com/ehabets/ANF-Generator>

Table 1: *EER (%) comparison on the test set of the Libri-adhoc40 corpus.*

Method	Number of parameters	Train set	Test set of Libri-adhoc40			
			8 channels	16 channels	32 channels	40 channels
Oracle one-best	1.437 M	—	26.2217	21.2862	16.1878	14.5232
Mean pooling	1.437 M	—	20.5158	19.7459	19.6453	19.5376
EV [23]	1.437 M	—	24.1581	20.5435	17.2928	15.7186
Beamforming [24]	1.437 M	—	19.6441	15.8513	12.9207	12.0551
Attentive pooling [18]	1.454 M	Simulation	14.4540	13.1235	12.1465	11.9208
		Real	10.4251	9.7030	9.2164	9.1536
UCSA [17]	1.503 M	Simulation	13.9136	12.5733	11.2091	10.8437
		Real	10.2522	9.4263	8.5960	8.4308
SA-aggregation (proposed)	1.570 M	Simulation	13.6264	12.6282	11.4266	11.2482
		Real	9.7846	8.5505	7.5891	7.3640
GAT-aggregation (proposed)	1.503 M	Simulation	13.2089	11.7986	10.5637	10.2163
		Real	10.0407	8.3622	7.1318	6.8627

heads is 4. We used *voxceleb_trainer*³ to build our models. The preprocessing of the data and training setting of the proposed models are the same as [17]. The following six baselines were used for comparison:

- **Oracle one-best:** We pick the channel that is physically closest to the speaker source as the input of the single-channel ASV model. Note that, for the *oracle one-best* baseline, the distances between the speaker and the microphones are known beforehand.
- **Utterance-level cross-channel self-attention (UCSA) [17]:** It adds an utterance-level cross-channel self-attention layer after the pooling layer. Unlike the original UCSA in [17], we removed FFN which leads to better performance than the original UCSA. The number of SA layers and attention heads are 1 and 4 respectively.
- **Attentive pooling [18]:** It adds an attentive pooling layer above the pooling layer of the single-channel ASV.
- **Mean pooling:** It sets equal weights to all speaker embeddings from different channels.
- **EV [23]:** It selects the signal of a microphone with the highest envelope variance as the input of the single-channel ASV.
- **Beamforming [24]:** It aggregates the multi-channel signals into a single-channel enhanced signal by classic acoustic beamforming. The enhanced speech is taken as the input of the single channel ASV.

3.3. Results

Table 1 lists the performance of the comparison methods on the Libri-adhoc40 test set. From the table, we see that, as the number of channels increases, almost all methods achieve better performance. The EV method, which tends to select a channel that is least affected by reverberation, achieves better performance than the *Oracle one-best* method in the 8-channel and 16-channel test environments. It indicates that reverberation affects ASV performance. From the results of Beamforming, we find that adding a multichannel speech enhancement front-end is a good way of improving the ASV performance when the number of the channels is large enough. From the table, we see that the attentive-pooling and UCSA methods trained with the real data outperform the *Oracle one-best*, EV, and Beamforming, which shows that the attention scheme is good at capturing the global information across channels.

³https://github.com/clovaai/voxceleb_trainer

More importantly, from the table, we also see that all of the proposed methods perform well on the Libri-adhoc40 test set. Specially, compared with the attentive-pooling and UCSA, the SA-aggregation fine-tuned with the real data achieves a relative EER reduction of 19.55% over the attentive-pooling, and 12.65% over UCSA in the 40-channel test scenario. It demonstrates the effectiveness of the frame-level modeling strategy over the utterance-level one. Moreover, GAT-aggregation achieves better performance than SA-aggregation, in the 16-, 32- and 40-channel test scenarios. For example, the GAT-aggregation fine-tuned with the real data achieves a relative EER reduction of 6.02% and 6.81% over SA-aggregation in the 32- and 40-channels test scenarios, respectively. This result shows that the connections between the microphone nodes provide important information for the performance improvement of the multi-channel ASV with a large number of ad-hoc microphone nodes. Although the proposed models fine-tuned with the simulated data are inferior to those fine-tuned with the real data, they still outperform most of the other baseline systems, showing that the proposed methods have a strong generalization ability.

4. Conclusion

In this paper, we propose a novel multi-channel ASV framework with ad-hoc microphone arrays. It conducts channel aggregation at the frame-level by adding a multi-channel aggregation module before the pooling layer. We propose two types of multi-channel aggregation modules: SA-aggregation and GAT-aggregation. The aggregation module can be trained in a way that is independent of the number and permutation of the microphones. Experimental results show that the proposed frame-level modules are able to mine more spatial and temporal information than the utterance-level modules. Moreover, GAT-aggregation performs better than SA-aggregation, when the number of nodes of the ad-hoc microphone arrays is large. Finally, the models fine-tuned with the simulated data have a good generalization ability to the real test scenarios.

5. Acknowledgement

This work was supported in part by the National Science Foundation of China (NSFC) under Grant 62176211, and in part by the Project of the Science, Technology, and Innovation Commission of Shenzhen Municipality, China under Grant JCYJ20210324143006016 and JSGG20210802152546026.

6. References

- [1] H. Taherian, Z.-Q. Wang, J. Chang, and D. Wang, "Robust speaker recognition based on single-channel and multi-channel speech enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1293–1302, 2020.
- [2] Y. Jiang, D. Wang, R. Liu, and Z. Feng, "Binaural classification for reverberant speech segregation using deep neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 12, pp. 2112–2121, 2014.
- [3] Z.-Q. Wang and D. Wang, "All-neural multi-channel speech enhancement," in *Interspeech*, 2018, pp. 3234–3238.
- [4] S. Abd El-Moneim, M. Nassar, M. I. Dessouky, N. A. Ismail, A. S. El-Fishawy, and F. E. Abd El-Samie, "Text-independent speaker recognition using lstm-rnn and speech enhancement," *Multimedia Tools and Applications*, vol. 79, no. 33, pp. 24 013–24 028, 2020.
- [5] L. Moner, O. Plchot, J. Rohdin, L. Burget, and J. ernock, "Speaker verification with application-aware beamforming," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 411–418.
- [6] L. Mošner, P. Matějka, O. Novotný, and J. H. Černocký, "Dereverberation and beamforming in far-field speaker recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5254–5258.
- [7] Q. Jin, R. Li, Q. Yang, K. Laskowski, and T. Schultz, "Speaker identification with distant microphone speech," in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2010, pp. 4518–4521.
- [8] H. Taherian, Z.-Q. Wang, and D. Wang, "Deep learning based multi-channel speaker recognition in noisy and reverberant environments," in *Interspeech*, 2019.
- [9] X.-L. Zhang, "Deep ad-hoc beamforming," *Computer Speech & Language*, vol. 68, p. 101201, 2021.
- [10] Z. Yang, S. Guan, and X.-L. Zhang, "Deep ad-hoc beamforming based on speaker extraction for target-dependent speech separation," *arXiv preprint arXiv:2012.00403*, 2020.
- [11] D. Wang, T. Yoshioka, Z. Chen, X. Wang, T. Zhou, and Z. Meng, "Continuous speech separation with ad hoc microphone arrays," in *2021 29th European Signal Processing Conference (EUSIPCO)*. IEEE, 2021, pp. 1100–1104.
- [12] D. Wang, Z. Chen, and T. Yoshioka, "Neural speech separation using spatially distributed microphones," *Proc. INTERSPEECH*, pp. 339–343, 2020.
- [13] S. Guan, S. Liu, J. Chen, W. Zhu, S. Li, X. Tan, Z. Yang, M. Xu, Y. Chen, C. Liang *et al.*, "Libri-adhoc40: A dataset collected from synchronized ad-hoc microphone arrays," in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2021, pp. 1116–1120.
- [14] F.-J. Chang, M. Radfar, A. Mouchtaris, B. King, and S. Kunzmann, "End-to-end multi-channel transformer for speech recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5884–5888.
- [15] F.-J. Chang, M. Radfar, A. Mouchtaris, and M. Omologo, "Multi-channel transformer transducer for speech recognition," *Proc. INTERSPEECH*, 2021.
- [16] J. Chen and X.-L. Zhang, "Scaling sparsemax based channel selection for speech recognition with ad-hoc microphone arrays," in *Interspeech*, 2021, pp. 291–295.
- [17] C. Liang, J. Chen, S. Guan, and X.-L. Zhang, "Attention-based multi-channel speaker verification with ad-hoc microphone arrays," in *2021 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2021, pp. 1111–1115.
- [18] D. Cai and M. Li, "Embedding aggregation for far-field speaker verification with distributed microphone arrays," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 308–315.
- [19] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, "Graph attention networks," *ICLR*, 2018.
- [20] S. Brody, U. Alon, and E. Yahav, "How attentive are graph attention networks?" *ICLR*, 2022.
- [21] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [22] 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.
- [23] M. Wolf and C. Nadeu, "Channel selection measures for multi-microphone speech recognition," *Speech Communication*, vol. 57, pp. 170–180, 2014.
- [24] X. Anguera, C. Wooters, and J. Hernando, "Acoustic beamforming for speaker diarization of meetings," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 7, pp. 2011–2022, 2007.
- [25] 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.
- [26] X. Tan and X.-L. Zhang, "Speech enhancement aided end-to-end multi-task learning for voice activity detection," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6823–6827.
- [27] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, "The third chimespeech separation and recognition challenge: Dataset, task and baselines," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 504–511.
- [28] A. Varga and H. J. Steeneken, "Assessment for automatic speech recognition: Ii. noisex-92: A database and an experiment to study the effect of additive noise on speech recognition systems," *Speech communication*, vol. 12, no. 3, pp. 247–251, 1993.
- [29] J. S. Chung, J. Huh, S. Mun, M. Lee, H. S. Heo, S. Choe, C. Ham, S. Jung, B.-J. Lee, and I. Han, "In defence of metric learning for speaker recognition," *INTERSPEECH*, pp. 2977–2981, 2020.
- [30] Y. Zhu, T. Ko, D. Snyder, B. Mak, and D. Povey, "Self-attentive speaker embeddings for text-independent speaker verification," in *Interspeech*, vol. 2018, 2018, pp. 3573–3577.