

POLYPHONIC SOUND EVENT DETECTION AND LOCALIZATION USING A TWO-STAGE STRATEGY

Yin Cao,^{1*} Qiuqiang Kong,^{1*} Turab Iqbal,¹ Fengyan An,² Wenwu Wang,¹ Mark D. Plumbley¹

¹Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, UK
{yin.cao, q.kong, t.iqbal, w.wang, m.plumbley}@surrey.ac.uk

²School of Mechanical & Automotive Engineering, Qingdao University of Technology, China
anfy@qut.edu.cn

ABSTRACT

Sound event detection (SED) and localization refer to recognizing sound events and estimating their spatial and temporal locations. Using neural networks has become the prevailing method for SED. In the area of sound localization, which is usually performed by estimating the direction of arrival (DOA), learning-based methods have recently been developed. In this paper, it is experimentally shown that the trained SED model is able to contribute to the direction of arrival estimation (DOAE). However, joint training of SED and DOAE degrades the performance of both. Based on these results, a two-stage polyphonic sound event detection and localization method is proposed. The method learns SED first, after which the learned feature layers are transferred for DOAE. It then uses the SED ground truth as a mask to train DOAE. The proposed method is evaluated on the DCASE 2019 Task 3 dataset, which contains different overlapping sound events in different environments. Experimental results show that the proposed method is able to improve the performance of both SED and DOAE, and also performs significantly better than the baseline method.

Index Terms— Sound event detection, source localization, direction of arrival, convolutional recurrent neural networks

1. INTRODUCTION

Sound event detection is a rapidly developing research area that aims to analyze and recognize a variety of sounds in urban and natural environments. Compared to sound tagging, event detection also involves estimating the time of occurrence of sounds. Automatic recognition of sound events would have a major impact in a number of applications [1]. For instance, sound indexing and sharing, bioacoustic scene analysis for animal ecology, smart home automatic audio event recognition (baby cry detection, window break alarm), and sound analysis in smart cities (security surveillance).

Recently, approaches based on neural networks have been shown to be especially effective for SED [2]. Unlike audio tagging problems [3–5], which only aim to detect whether the sound events are present in a sound clip, SED also involves predicting temporal information of events. Early neural network architectures utilized fully-connected layers to detect temporally-overlapping sound events [6]. More recently, due to their success in image recognition, convolutional neural networks (CNNs) have become the prevailing architecture in this area [7–10]. Such methods use suitable time-frequency representations of audio, which are analogous to the image inputs

in computer vision. Another popular type of neural network is the recurrent neural network (RNN), which has the ability to learn long temporal patterns present in the data, making it suitable for SED [11]. Hybrids containing both CNN and RNN layers, known as convolutional recurrent neural networks (CRNNs), have also been proposed, which have led to state-of-the-art performance in SED [4, 12].

Sound source localization, which focuses on identifying the locations of sound sources, on the other hand, has been an active research topic for decades [13]. It plays an important role in applications such as robotic listening, speech enhancement, source separation, and acoustic visualization. Unlike the dominance of neural-network-based techniques in SED, DOAE is mainly studied using two methods: parametric-based methods and learning-based methods.

Parametric-based DOAE methods can be divided into three categories [13]: time difference of arrival (TDOA) estimation, maximized steered response power (SRP) of a beamformer, and high-resolution spectral estimation. Generalized cross-correlation (GCC) methods are the most widely-used approaches for TDOA estimation [14, 15]. Since the TDOA information is conveyed in the phase rather than the amplitude of the cross-spectrum, the GCC Phase Transform (GCC-PHAT) was proposed, which eliminates the effect of the amplitude while leaving only the phase [14]. The primary limitation of parametric-based GCC methods is the inability to accommodate multi-source scenarios.

Learning-based DOAE methods have the advantages of good generalization under different levels of reverberation and noise. They are designed to enable the system to learn the connections between input features and the DOA. There has already been a series of research addressing DOAE using deep neural networks [16–24]. Results show that they are promising and comparable to parametric methods. However, these neural-network-based methods are mainly based on static sources. In addition to spectrum-based features, GCC-based features, which can effectively supply time difference information, have also been used as the input features [16, 17, 21–23]. In order to further improve learning-based methods, more practical real-world sources need to be considered.

In real-world applications, a sound event is always transmitted in one or several directions. Given this fact, it is reasonable to combine sound event detection and localization with not only estimating their respective associated spatial location, but also identifying the type and temporal information of sound. Therefore, it is worthwhile to study them together and investigate the effects and potential connections between them. Recently, DCASE 2019 introduced Task 3, which is Sound Event Localization and Detection (SELD) for overlapping sound sources [25]. A recently-developed system known as

* Equal contribution.

SELDnet was used as the baseline system. SELDnet uses magnitude and phase spectrograms as input features and trains the SED and DOAE objectives jointly [26]. However, phase spectrograms are hard for neural networks to learn from, and further relationships between SED and DOAE have not been revealed.

In this paper, joint training of SED and DOAE is implemented first with log mel spectrograms and GCC-PHAT as the input features. According to the experimental results, SED is able to contribute to the performance of DOAE, while joint training of SED and DOAE degrades the performance of both. To solve this problem, a new two-stage method for polyphonic sound event detection and localization is proposed. This method deals with sound event detection and localization in two stages: the SED stage and the DOAE stage, corresponding to the SED branch and the DOAE branch in the model, respectively. During training, the SED branch is trained first only for SED, after which the learned feature layers are transferred to the DOAE branch. The DOAE branch fine-tunes the transferred feature layers and uses the SED ground truth as a mask to learn only DOAE. During inference, the SED branch estimates the SED predictions first, which are used as the mask for the DOAE branch to infer predictions. The experimental results show that by using the proposed method, DOAE can benefit from the SED predictions; both SED and DOAE can be improved at the same time. The proposed method performs significantly better than the baseline method.

The rest of the paper is organized as follows. In Section 2, the proposed learning method is described in detail. Section 3 introduces the dataset used, other methods for comparison, and experimental results. Finally, conclusions are summarized in Section 4.

2. LEARNING METHOD

Joint training of SED and DOAE was first proposed in [26]. Their system is also used as the baseline system for DCASE 2019 Task 3. In this baseline, temporal consecutive magnitude and phase spectrograms are extracted as the input features from the time-domain audio waveform, which are then fed into a CRNN. Its loss is a weighted combination of the SED loss and the DOAE loss. Therefore, it can be imagined that this baseline system has an intrinsic trade-off between SED and DOAE according to the loss weight selected. In this paper, a two-stage polyphonic sound event detection and localization network is proposed to exploit their mutual strength.

2.1. Features

Selecting which features to use is an important factor for audio-related neural network applications. In this paper, the input signal format is of two types: First-Order of Ambisonics (FOA) or tetrahedral microphone array [25]. Log mel spectrograms and GCC-PHAT, which contains phase difference information between all microphone pairs, are chosen as the input features. Ambisonics and GCC-PHAT are explained in this section.

2.1.1. Ambisonics

Ambisonics was developed as a spatial sound encoding approach several decades ago [27]. It is based on the spherical harmonic (SH) decomposition of the sound field. Ambisonics encoding for plane-wave sound fields can be expressed as

$$\mathbf{b}(t) = \sum_{n=0}^N \mathbf{y}_n s_n(t), \quad (1)$$

where $s_n(t)$ is the n -th plane-wave source signal, N is the total number of sources, and \mathbf{y}_n is the vector of the spherical harmonic function values for direction (θ_n, ϕ_n) , and can be expressed as

$$\mathbf{y}_n = [Y_0^0(\theta_n, \phi_n), Y_1^{-1}(\theta_n, \phi_n), Y_1^0(\theta_n, \phi_n), Y_1^1(\theta_n, \phi_n), \dots, Y_L^{-L}(\theta_n, \phi_n), \dots, Y_L^0(\theta_n, \phi_n), \dots, Y_L^L(\theta_n, \phi_n)]^T, \quad (2)$$

where L indicates the order of Ambisonics. It can be seen that Ambisonics contains the information of the source DOA. In addition, a higher directional resolution relates to a higher order of Ambisonics. Order- L of Ambisonics needs at least $(L+1)^2$ microphones to encode. In real applications, the sound field is recorded using a spherical microphone array and converted into Ambisonics.

2.1.2. Generalized Cross-Correlation

GCC is widely used in TDOA estimation by means of maximizing the cross-correlation function to obtain the lag time between two microphones. The cross-correlation function is usually calculated through the inverse-FFT of the cross power spectrum. GCC-PHAT is the phase-transformed version of GCC, which whitens the cross power spectrum to eliminate the influence of the amplitude, leaving only the phase. GCC-PHAT can be expressed as

$$GCC_{ij}(t, \tau) = \mathcal{F}_{f \rightarrow \tau}^{-1} \frac{X_i(f, t) X_j^*(f, t)}{|X_i(f, t)| |X_j(f, t)|}, \quad (3)$$

where $\mathcal{F}_{f \rightarrow \tau}^{-1}$ is the inverse-FFT from f to τ , $X_i(f, t)$ is the Short-Time Fourier Transform (STFT) of the i -th microphone signal, and $*$ denotes the conjugate. TDOA, which is the lag time $\Delta\tau$ between two microphones, can then be estimated by maximizing GCC with respect to τ . Nevertheless, this estimation is usually not stable, especially in high reverberation and low SNR environments, and does not directly work for multiple sources. However, $GCC_{ij}(t, \tau)$ contains all of the time delay information and is generally short-time stationary. $GCC_{ij}(t, \tau)$ can also be considered as a GCC spectrogram with τ corresponding to the number of mel-band filters. That is, GCC-PHAT can be stacked with a log mel spectrogram as the input features. In order to determine the size of GCC-PHAT, the largest distance between two microphones d_{\max} needs to be used. The maximum delayed samples corresponding to $\Delta\tau_{\max}$ can be estimated by $d_{\max}/c \cdot f_s$, where c is the sound speed and f_s is the sample rate. In this paper, log mel and GCC-PHAT are stacked as the input features, considering the possibility of the advance and the delay of GCC. The number of mel-bands, therefore, should be no smaller than the doubled number of delayed samples plus one [14].

2.2. Network architecture

The network is shown in Fig. 1, and has two branches, the SED branch and the DOAE branch. During training, the extracted features, which have shape $C \times T \times F$, are first sent to the SED branch. C indicates the number of feature maps, T is the size of time bins, and F is the number of mel-band filters or delayed samples of GCC-PHAT. The CNN layers, which are also named as feature layers in this paper, are constructed with 4 groups of 2D CNN layers (Convs) with 2×2 average-pooling after each of them. Each Convs' group consists of two 2D Convs, with a receptive field of 3×3 , a stride of 1×1 , and a padding size of 1×1 [10]. The Convs' kernels are able to filter across all of the channels of the input features or

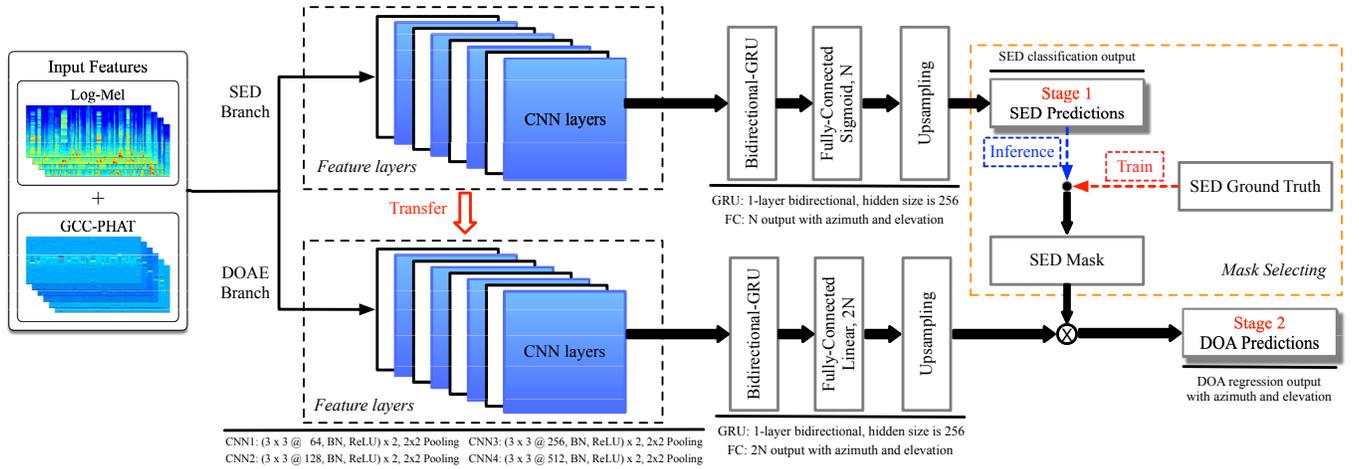


Figure 1: The diagram of the proposed two-stage sound event detection and localization network. SED ground truth is used as the mask to train DOAE branch. SED predictions are used as the mask to infer DOAE.

the feature maps from the last layer, hence are able to learn inter-channel information. CNN layers are capable of learning local temporal and frequency information to better abstract the event-level information. Each single CNN layer is followed by a Batch Normalization layer [28] and a ReLU activation. After the CNN layers, the data has shape $C_{out} \times T/16 \times F/16$, where C_{out} is the number of output feature maps of the last CNN layer. It is then sent to a global average-pooling layer to reduce the dimension of F . After this, the data is reshaped to have shape $T/16 \times C_{out}$ and is fed to a bidirectional GRU. The output size is maintained and is sent to a fully-connected layer with output size $T/16 \times N$, where N is the number of event classes. The sigmoid activation function is used afterwards with an upsampling in the temporal dimension to ensure the output size is consistent with T . The SED predictions can then be obtained through an activation threshold. Binary cross-entropy is used for this multi-label classification task.

The DOAE branch is then trained. The CNN layers are transferred from the SED branch and are fine-tuned. The output of the fully-connected layer for the DOAE branch is a vector of $N \times 2$ linear values, which are azimuth and elevation angles for N events. They are then masked by the SED ground truth during training to determine if the corresponding angles are currently active. Finally, the mean absolute error is chosen as the DOAE regression loss.

During inference, the SED branch first computes the SED predictions, which are then used as the SED mask to obtain the DOAE. The reason for building this network architecture is to enhance the representation ability of a single network so that each branch is only responsible for one task, while the DOAE branch can still incorporate the benefits contributed from SED.

3. EXPERIMENTAL STUDY

The proposed two-stage polyphonic sound event detection and localization method is compared with other methods described in Section 3.2. They are evaluated on the DCASE 2019 Task 3 dataset [25]. This task is for sound event detection and localization. The dataset provides two formats of data: 1) First-Order of Ambisonics; 2) tetrahedral microphone array. The development set consists of 400 one minute long recordings, divided into four cross-validation splits.

There are 11 kinds of isolated sound events in total. The audio recordings are mixtures of isolated sound events and natural ambient noise. The sound events are convolved with impulse responses collected from five indoor locations, resulting in 324 unique combinations of azimuth-elevation angles. One challenging problem in this dataset is that the sound events in the audio recordings have a polyphony of up to two, which means sound events from different locations may overlap. The source code for this paper is released on GitHub¹.

3.1. Evaluation metrics

Polyphonic sound event detection and localization are evaluated with individual metrics for SED and DOAE. For SED, segment-based error rate (ER) and F-score [29] are calculated in one-second lengths. A lower ER or a higher F-score indicates better performance. In addition, mean average precision (mAP), which is the area under the precision and recall curve, is used to evaluate the frame-level tagging performance. The mAP is used here because it does not depend on the threshold selection, hence is able to better objectively evaluate the performance. A higher mAP indicates better performance. For DOAE, DOA error and frame recall are used [24]. A lower DOA error or a higher frame recall indicates better performance.

3.2. Methods for comparison

In order to show the effectiveness of the proposed method, several other methods are compared, including

- **Baseline**, which is the baseline method used in DCASE 2019 Task 3, uses magnitude and phase spectrograms as the input features. The features are then fed to a CRNN network. The loss of SED and DOAE are combined and jointly trained.
- **SELDnet**, which has the same architecture with the baseline but using log mel and GCC-PHAT spectrograms as the input features.
- **DOA**, which uses log mel and GCC-PHAT spectrograms as the input features to only estimate DOA. It transfers the CNN layers from the SED network and utilizes SED ground truth as the mask.

¹ <https://github.com/yinkalarior/Two-Stage-Polyphonic-Sound-Event-Detection-and-Localization>

Table 1: Performance for the development dataset.

Methods	Net	MIC-ARRAY					FOA				
		ER	F	mAP	DOA	FR	ER	F	mAP	DOA	FR
Baseline	CRNN	0.350	0.800	—	30.8°	0.840	0.340	0.799	—	28.5°	0.854
SELDnet	CNN	0.277	0.844	0.718	11.0°	0.827	0.281	0.843	0.718	10.9°	0.828
	CRNN	0.213	0.879	0.770	11.3°	0.847	0.221	0.876	0.768	12.6°	0.844
DOA	CNN	—	—	—	13.3°	—	—	—	—	13.1°	—
	CRNN	—	—	—	11.9°	—	—	—	—	11.9°	—
DOA-NT	CNN	—	—	—	14.7°	—	—	—	—	14.5°	—
	CRNN	—	—	—	14.0°	—	—	—	—	14.3°	—
Two-Stage	CNN	0.251	0.862	0.749	10.9°	0.832	0.248	0.864	0.756	10.8°	0.832
	CRNN	0.167	0.909	0.819	9.85°	0.863	0.181	0.898	0.800	9.84°	0.857

- **DOA-NT**, is the same as DOA method except for not transferring CNN layers. Both DOA and DOA-NT only estimate DOAs.

All of the above-mentioned methods are evaluated on both CNNs and CRNNs. The CNN has the same architecture as the CRNN but without the recurrent layer. Furthermore, microphone array signals do not need extra encoding processes. It is more convenient to use in practice, whereas the encoding of FOA may contain extra spatial information. Therefore, it is worthwhile to evaluate these methods with both the microphone array and FOA data.

3.3. Hyper-parameters

To extract the input features, the sample rate of STFT is set to 32kHz. A 1024-point Hanning window with a hop size of 320 points is utilized. In the DCASE 2019 Task 3 dataset, the largest microphone distance is 4.82cm [25]. According to Section 2.1.2, the number of mel-band filters and the delays of GCC-PHAT is set to be 64. For 4 channels of signals, up to 10 input channels of signals are sent to the network. The audio clips are segmented to have a fixed length of 2 seconds with a 1-second overlap for training. The learning rate is set to 0.001 for the first 30 epochs and is then decayed by 10% every epoch. The final results are calculated after 50 epochs.

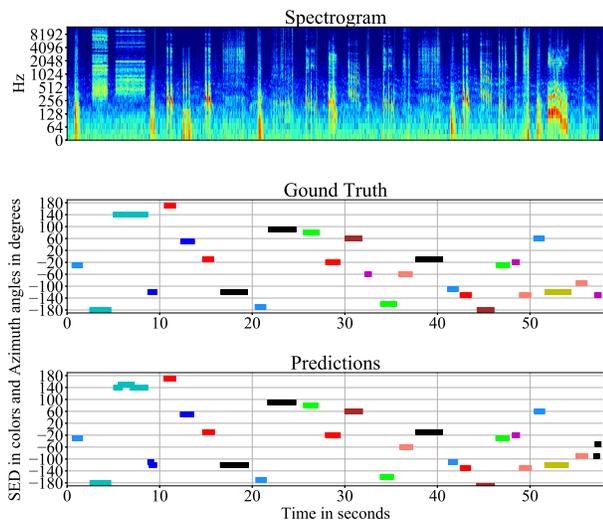


Figure 2: SED and DOAE Azimuth results for proposed two-stage method. Different colors indicate different classes of events.

3.4. Results

The experimental results are shown in Table 1. SELDnet with log mel and GCC-PHAT spectrograms as the input features was implemented first to compare with the baseline method. It can be seen from both microphone array data and FOA data that with log mel and GCC-PHAT spectrograms as the input features, SELDnet outperforms the baseline system using magnitude and phase spectrograms. Log mel spectrograms are more effective input features than magnitude spectrograms, not only due to their better performance, but they are also more compact. GCC-PHAT spectrograms, which mainly contain the time difference information, show their advantages over phase spectrograms. The results of DOA and DOA-NT show that with trained CNN layers transferred, DOA error is consistently lower than not transferring, which indicates that SED information contributes to the DOAE performance; it can also be observed that the convergence speed is much faster with CNN layers transferred. Comparing SELDnet with DOA-NT, it also shows that the joint training is better than the training of DOAE without CNN layers transferred, which also proves SED contributes to DOAE. The proposed two-stage method is presented in the end. The metrics scores are the best among all the methods. Compared with SELDnet, it indicates that the joint training of SED and DOAE degrades the performance of both. This two-stage method minimizes each loss individually, hence the network representation ability is enhanced for each sub-task, while the contribution from SED to DOAE is still preserved by transferring CNN layers to the DOAE branch.

Comparing microphone array data and FOA data, the results do not show FOA is better, which means FOA does not necessarily contain more spatial information than microphone array signals with four channels. On the other hand, in most cases, CRNNs perform better than CNNs, which indicates that long temporal information may be useful for both SED and DOAE. A visualization of SED and DOAE using the proposed method for one clip is shown in Fig. 2. It can be seen that most of the SED and DOAE predictions are accurate in both temporal and spatial dimensions.

4. CONCLUSION

Treating sound event detection and localization as a combined task is reasonable. In this paper, it shows that SED information can be used to improve the performance of DOAE. However, joint training of SED and DOAE degrades the performance of both. A two-stage polyphonic sound event detection and localization method is proposed to solve this problem. The proposed method uses log mel and GCC-PHAT spectrograms as the input features and has two branches of SED and DOAE. The SED branch is trained first, after which the trained feature layers are transferred to the DOAE branch. The DOAE branch then uses the SED ground truth as a mask to train DOAE. Experimental results show that the proposed method is able to enhance the network representation ability for each branch, while still keeping the contributions from SED to DOAE. The proposed method is shown to significantly outperform the baseline method.

5. ACKNOWLEDGMENT

This research was supported by EPSRC grant EP/N014111/1 “Making Sense of Sounds” and was partially supported by the H2020 Project entitled AudioCommons funded by the European Commission with Grand Agreement number 688382.

6. REFERENCES

- [1] T. Virtanen, M. D. Plumbley, and D. Ellis, *Computational Analysis of Sound Scenes and Events*. Springer, 2018, ch. 1, pp. 3–12.
- [2] A. Mesaros, T. Heittola, E. Benetos, P. Foster, M. Lagrange, T. Virtanen, and M. D. Plumbley, “Detection and classification of acoustic scenes and events: Outcome of the DCASE 2016 challenge,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 2, pp. 379–393, Feb. 2018.
- [3] K. Choi, G. Fazekas, and M. Sandler, “Automatic tagging using deep convolutional neural networks,” in *17th International Society for Music Information Retrieval Conference*, 2016.
- [4] Y. Xu, Q. Kong, Q. Huang, W. Wang, and M. D. Plumbley, “Convolutional gated recurrent neural network incorporating spatial features for audio tagging,” in *International Joint Conference on Neural Networks (IJCNN)*, 2017, pp. 3461–3466.
- [5] E. Fonseca, M. Plakal, F. Font, D. P. W. Ellis, X. Favory, J. Pons, and X. Serra, “General-purpose tagging of Freesound audio with AudioSet labels: task description, dataset, and baseline,” in *In Proceedings of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop (DCASE2018)*, 2018.
- [6] E. Çakır, T. Heittola, H. Huttunen, and T. Virtanen, “Polyphonic sound event detection using multi label deep neural networks,” in *International Joint Conference on Neural Networks (IJCNN)*, 2015.
- [7] K. J. Piczak, “Environmental sound classification with convolutional neural networks,” in *IEEE 25th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2015.
- [8] H. Zhang, I. McLoughlin, and Y. Song, “Robust sound event recognition using convolutional neural networks,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [9] J. Salamon and J. P. Bello, “Deep convolutional neural networks and data augmentation for environmental sound classification,” *IEEE Signal Processing Letters*, vol. 24, no. 3, pp. 279–283, Mar. 2017.
- [10] Q. Kong, Y. Cao, T. Iqbal, Y. Xu, W. Wang, and M. D. Plumbley, “Cross-task learning for audio tagging, sound event detection and spatial localization: Dcase 2019 baseline systems,” *arXiv preprint arXiv:1904.03476*, 2019.
- [11] G. Parascandolo, H. Huttunen, and T. Virtanen, “Recurrent neural networks for polyphonic sound event detection in real life recordings,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016.
- [12] E. Çakır, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen, “Convolutional recurrent neural networks for polyphonic sound event detection,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1291–1303, 2017.
- [13] M. Brandstein and D. Ward, *Microphone Arrays: Signal Processing Techniques and Applications*. Springer Science & Business Media, 2013, ch. 8, pp. 157–180.
- [14] C. Knapp and G. 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.
- [15] J. Scheuing and B. Yang, “Correlation-based TDOA-estimation for multiple sources in reverberant environments,” in *Speech and Audio Processing in Adverse Environments*. Springer Berlin Heidelberg, 2008, pp. 381–416.
- [16] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [17] F. Vesperini, P. Vecchiotti, E. Principi, S. Squartini, and F. Piazza, “A neural network based algorithm for speaker localization in a multi-room environment,” in *IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2016.
- [18] N. Ma, T. May, and G. J. Brown, “Exploiting deep neural networks and head movements for robust binaural localization of multiple sources in reverberant environments,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 12, pp. 2444–2453, 2017.
- [19] S. Chakrabarty and E. A. P. Habets, “Broadband DOA estimation using convolutional neural networks trained with noise signals,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2017.
- [20] S. Chakrabarty and E. A. P. Habets, “Multi-speaker localization using convolutional neural network trained with noise,” *arXiv preprint arXiv:1712.04276*, 2017.
- [21] W. He, P. Motlicek, and J. Odobez, “Deep neural networks for multiple speaker detection and localization,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [22] E. L. Ferguson, S. B. Williams, and C. T. Jin, “Sound source localization in a multipath environment using convolutional neural networks,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [23] Y. Sun, J. Chen, C. Yuen, and S. Rahardja, “Indoor sound source localization with probabilistic neural network,” *IEEE Transactions on Industrial Electronics*, vol. 65, no. 8, pp. 6403–6413, 2018.
- [24] S. Adavanne, A. Politis, and T. Virtanen, “Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network,” in *26th European Signal Processing Conference (EUSIPCO)*, 2018.
- [25] <http://dcase.community/challenge2019/task-sound-event-localization-and-detection>.
- [26] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, “Sound event localization and detection of overlapping sources using convolutional recurrent neural networks,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 34–48, 2019.
- [27] M. A. Gerzon, “Ambisonics in multichannel broadcasting and video,” *Journal of the Audio Engineering Society*, vol. 33, no. 11, pp. 859–871, 1985.
- [28] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International Conference on Machine Learning*, 2015.
- [29] A. Mesaros, T. Heittola, and T. Virtanen, “Metrics for polyphonic sound event detection,” *Applied Sciences*, vol. 6, no. 6, pp. 162–178, 2016.