


# Multi-Speaker DOA Estimation Using Deep Convolutional Networks Trained With Noise Signals

Soumitro Chakrabarty , *Student Member, IEEE*, and Emanuël A. P. Habets , *Senior Member, IEEE*

**Abstract**—Supervised learning-based methods for source localization, being data driven, can be adapted to different acoustic conditions via training and have been shown to be robust to adverse acoustic environments. In this paper, a convolutional neural network (CNN) based supervised learning method for estimating the direction of arrival (DOA) of multiple speakers is proposed. Multi-speaker DOA estimation is formulated as a multi-class multi-label classification problem, where the assignment of each DOA label to the input feature is treated as a separate binary classification problem. The phase component of the short-time Fourier transform (STFT) coefficients of the received microphone signals are directly fed into the CNN, and the features for DOA estimation are learnt during training. Utilizing the assumption of disjoint speaker activity in the STFT domain, a novel method is proposed to train the CNN with synthesized noise signals. Through experimental evaluation with both simulated and measured acoustic impulse responses, the ability of the proposed DOA estimation approach to adapt to unseen acoustic conditions and its robustness to unseen noise type is demonstrated. Through additional empirical investigation, it is also shown that with an array of  $M$  microphone our proposed framework yields the best localization performance with  $M-1$  convolution layers. The ability of the proposed method to accurately localize speakers in a dynamic acoustic scenario with varying number of sources is also shown.

**Index Terms**—Source localization, multiple speakers, convolutional neural networks.

## I. INTRODUCTION

**M**ANY applications such as hands-free communication, teleconferencing, robot audition and distant speech recognition require information on the location of sound sources in the acoustic environment. Information regarding the source location can be utilized for the task of enhancing the signal coming from a specific location while suppressing the undesired signal components. In some applications, the information is used for camera steering whereas in applications like robot audition the source location information is used for navigation purposes. The relative direction of a sound source with respect to a microphone array is generally given in terms of the direction of arrival (DOA) of the sound wave originating from the source

position. As this information is not available in most practical scenarios, the DOA of the sound source needs to be estimated. However, accurate DOA estimation is a challenging task in the presence of noise and reverberation. The task becomes even more difficult when the DOAs of multiple sound sources need to be estimated.

In the literature related to DOA estimation, there exist two kinds of estimation paradigms: broadband and narrowband DOA estimation. In narrowband DOA estimation, the task of DOA estimation is performed separately for each frequency sub-band, whereas in broadband DOA estimation the task is performed for the whole input spectrum. In this work, the focus is on broadband DOA estimation of multiple simultaneously active sources.

Over the years, several approaches have been developed for the task of broadband DOA estimation. Some popular approaches are: i) subspace based approaches such as multiple signal classification (MUSIC) [1], [2], ii) time difference of arrival (TDOA) based approaches that use the family of generalized cross correlation (GCC) methods [3], iii) generalizations of the cross-correlation methods to microphone arrays with more than two channels such as steered response power (SRP), SRP with phase transform (SRP-PHAT) [4], and multichannel cross correlation coefficient (MCCC) [5], iv) adaptive multichannel time delay estimation using blind system identification based methods [6], v) probabilistic model based methods such as maximum likelihood method [7] and vi) methods based on histogram analysis of narrowband DOA estimates [8], [9]. The methods in categories i-iii, are developed under the free-field propagation model, which leads to severe degradation in their performance in reverberant and noisy scenarios. The method presented in [6] is designed for single-source localization specifically in reverberant scenarios, however, its extension to multi-source case is not straightforward [5]. In recent years, different methods have been proposed that aim to localize multiple simultaneously active sources in adverse acoustic environments [10]–[12]. The methods presented in [10] and [11] were developed particularly for binaural and stereo channels, respectively. An extension of these works to arbitrary microphone arrays require pair-wise processing that would lead to a combinatorial growth in computational cost as the number of microphones increase. Similarly, the method in [12] was formulated for a spherical microphone array. A common aspect of these methods is the inclusion of involved data pre- and post-processing techniques that require manual adjustment of specific parameters to achieve good localization performance.

Manuscript received June 28, 2018; revised December 6, 2018; accepted February 19, 2019. Date of publication February 25, 2019; date of current version April 11, 2019. The guest editor coordinating the review of this paper and approving it for publication was Prof. Peter K. Willett. (*Corresponding author: Soumitro Chakrabarty.*)

The authors are with the International Audio Laboratories Erlangen (a joint institution between the University of Erlangen-Nuremberg and Fraunhofer IIS), Erlangen 91058, Germany (e-mail: soumitro.chakrabarty@audiolabs-erlangen.de; emanuel.habets@audiolabs-erlangen.de).

Digital Object Identifier 10.1109/JSTSP.2019.2901664

In comparison, supervised learning approaches, being data driven, have the advantage that they can be adapted to different acoustic conditions via training. Also, if training data from diverse acoustic conditions are available, then these approaches can be made robust against noise and reverberation. Following the recent success of deep learning based supervised learning methods in various signal processing related tasks [13], [14], different methods for DOA estimation have been proposed [15]–[25]. A common aspect of the methods proposed in [16]–[20] is that they all involve an explicit feature extraction step. In [17], [19], GCC vectors, computed from the microphone signals, are provided as input to the learning framework. In [18], [20], similar to the computations involved in the MUSIC method for localization, the eigenvalue decomposition of the spatial correlation matrix is performed to get the eigenvectors corresponding to the noise subspace, and is provided as input to a neural network. In [16], a binaural setup is considered and binaural cues at different frequency sub-bands are computed and given as input. Such feature extraction steps generally lead to a high computational cost. Additionally, when features computed from the microphone signals are given as input the neural network mainly just learns the functional mapping from the features to the final DOA, which can possibly lead to a lack of robustness against adverse acoustic conditions.

One of the main reasons for the success of deep learning has been the encapsulation of the feature extraction step into the learning framework. In [15], a convolutional neural network (CNN) based approach for single-source localization was proposed where the power spectrograms of all the microphone signals for each time frame were concatenated to form the input vector which was then provided as an input to a CNN. The input feature extraction in the proposed work did not preserve the phase information, and only utilized the power of the microphone signals for DOA estimation. In a similar work [21], a special CNN architecture, known as ResNet, was used for single-source localization with power spectrograms of the microphone signals for multiple time frames as the input. By studying the traditional signal processing based methods for DOA estimation, it can be seen, especially for far-field scenarios, that most methods exploit the phase difference information between the microphone signals to perform localization. Based on this observation, in [22], the current authors proposed a CNN based supervised learning method for broadband DOA estimation of a single active speaker per short-time Fourier transform (STFT) time frame. Rather than involving an explicit feature extraction step, the phase component of STFT coefficients of the input signal were directly provided as input to the neural network. Another contribution of the work was to show the possibility of training the system using synthesized noise signals, which made the creation of training data much simpler compared to using real world signals like speech.

Following that, in [23], the previously proposed framework was extended to estimate multiple speaker DOAs. There, a novel method was developed to generate the training data using synthesized noise signals for multi-speaker localization. One of the main challenges of using noise signals for the multi-speaker case is that, for overlapping signals, the phase of the STFT

coefficients get combined non-linearly, and depend on the magnitude of the individual signals. This makes the learning procedure for the CNN difficult. To overcome this problem, the property of W-disjoint orthogonality [26], which holds approximately for speech signals, was utilized. In terms of evaluation, only preliminary results with simulated data for a single acoustic setup was shown in [23]. Recently, two other CNN based methods for multi-source localization have been proposed [24], [25]. In contrast to [23], the method proposed in [24] involves the extraction of microphone pair-wise GCC vectors which, as mentioned earlier, can lead to a combinatorial increase in input feature dimensions as the number of microphones in the array increases. In [25], a convolutional recurrent neural network (CRNN) based localization method is proposed that takes the magnitude and phase spectrograms of the microphone signals as input to first learn to estimate the MUSIC pseudo-spectrum followed by learning the mapping from the spatial pseudo-spectrum to the DOAs. The input feature in this work corresponds to 2 s audio, making it difficult to apply this method in an online framework.

In this paper, we extend the initial work on DOA estimation of multiple speakers presented in [23]. The formulation of the task of multi-speaker DOA estimation as a multi-label multi-class classification problem is presented, where first the probabilities of the active source DOAs are estimated at the frame-level. This makes the proposed system flexible in terms of the trade-off between localization accuracy vs adaptability in dynamic acoustic scenarios as the complete post-processing technique for the obtained frame-level probabilities can be determined based on the application. In this work, the frame-level probabilities are averaged over multiple time frames depending on the chosen block length over which the final DOA estimates are to be obtained. From these averaged probabilities, assuming the number of speakers,  $L$ , within that block is known, the DOAs corresponding to the classes with the  $L$  highest probabilities are chosen as the final DOA estimates. To build robustness to adverse acoustic conditions, multi-condition training in the form of training data from diverse acoustic scenarios is performed. A detailed description of the previously proposed method for generating training data using synthesized noise signals along with pseudo code for a better understanding of the entire training data generation procedure is also presented.

With respect to the proposed CNN architecture, we first posit that due to the small filters chosen to learn the phase correlations between neighboring microphones,  $M - 1$  convolution layers are required to learn from the phase correlation between all the microphone pairs, where  $M$  is the number of microphones in the array. Through experimental evaluation, the requirement of  $M - 1$  layers is shown in terms of both localization performance as well as number of trainable parameters. This mainly demonstrates a relation between the number of microphones in the array and the number of convolution layers required to achieve the best performance, and gives a structured approach to how a network parameter can be modified due to change in the number of microphones. The influence of distance between the sources and the microphone array is also investigated experimentally. Through further experiments with simulated and

measured room impulse responses (RIRs), the robustness of the proposed method to unseen acoustic conditions and noise types is investigated. Additionally, we also show that even when the CNN is trained to estimate the probabilities of maximum two DOA classes per STFT time frame, at a block level the proposed method can be used to localize more than two speakers also. The practical applicability of the proposed method is further demonstrated by experiments with real recorded speech produced by human speakers.

The remainder of this paper is organized as follows. In Section II, the formulation of the problem as a multi-class multi-label classification is described. In Section III, we review the input feature representation used in our framework. The task of obtaining the final DOA estimates in our proposed system is described in Section IV. Section V presents a detailed description of the proposed method for generating training data using synthesized noise signals. Experimental evaluation of the proposed method is presented in Section VI. Section VII concludes the paper.

## II. PROBLEM FORMULATION

We want to utilize a CNN based supervised learning framework for estimating the DOAs of multiple simultaneously active sources by learning the mapping from the recorded microphone signals to the DOA of the active speech sources using a large set of labeled data. The DOA estimation in this work is performed for signal blocks that consist of multiple time frames of the STFT representation of the observed signals. The block length can be chosen depending on the application scenario. For example, for dynamic sound scenes it might be preferable to choose shorter block lengths compared to a scenario when it is known that the sources would be static.

The problem of multi-source DOA estimation is formulated as an  $I$ -class multi-label classification problem. As the first step, the whole DOA range is discretized to form a set of possible DOA values,  $\Theta = \{\theta_1, \dots, \theta_I\}$ . A class vector of length  $I$  is then formed where each class corresponds to a possible DOA value in the set  $\Theta$ . In this work, we assume an independent source DOA model, i.e., the spatial location of the sources are independent of each other. Due to this assumption, multi-label classification can be tackled using the binary relevance method [27], where the assignment of each DOA class label to the input is treated as a separate binary classification problem. As stated earlier, the aim is to obtain the DOA estimates of multiple speakers for a signal block, however, the input to the system is a feature representation for each STFT time frame separately.

As shown in Fig. 1, a supervised learning framework consists of a training and a test phase. In the training phase, the CNN is trained with a training data set that consists of pairs of fixed dimension feature vectors for each STFT time frame and the corresponding true DOA class labels. In the test phase, given the input feature representation corresponding to a single STFT time frame, the first task is to estimate the posterior probability of each DOA class. Following this, depending on the chosen block length, the frame-level probabilities are averaged over all the time frames in the block. Finally, considering  $L$  sources, the

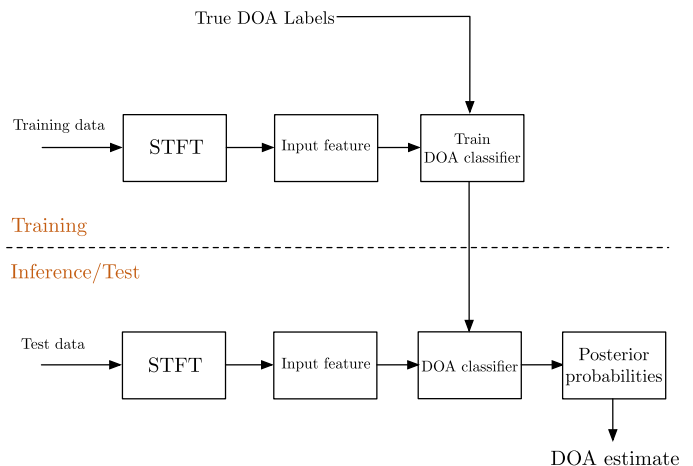


Fig. 1. Block diagram of the proposed system.

DOA estimates are given by selecting the  $L$  DOA classes with the highest probabilities.

In this work, we consider the number of sources  $L$  to be known. As an alternative, the number of active sources can be estimated based on the number of clear peaks in the averaged probabilities for a signal block. Also, the recorded signal from a reference microphone can also be used for speaker count estimation using the method proposed in [28]. Investigating the best strategy for this problem would be part of future work.

## III. INPUT REPRESENTATION

In this work, the aim is to learn the relevant features for the task of DOA estimation via training rather than have an explicit feature extraction step to compute the input to be given to the system. Therefore we use the *phase map* [22], [23] as the input feature representation in this work. For the sake of completeness, we give a brief description of this representation.

As described earlier, the input to the DNN framework is a feature representation corresponding to each STFT time frame. Let us consider that the received microphone signals are transformed to the STFT domain using an  $N_f$  point discrete Fourier transform (DFT). In the STFT domain, the observed signals at each time-frequency (TF) instance are represented by complex numbers. Therefore, the observed signal can be expressed as

$$Y_m(n, k) = A_m(n, k)e^{j\phi_m(n, k)}, \quad (1)$$

where  $A_m(n, k)$  represents the magnitude component and  $\phi_m(n, k)$  denotes the phase component of the STFT coefficient of the received signal at the  $m$ -th microphone for the  $n$ -th time frame and  $k$ -th frequency bin. In this work, we directly provide the phase component of the STFT coefficients of the received signals as input to our system. Note that this phase term consists of the phase of the source along with the effect of the propagation path. The idea is to make the system learn the relevant feature for DOA estimation from the phase component through training.

Since the aim is to compute the probabilities of the DOA classes at each time frame, the input feature for the  $n$ -th time frame is formed by arranging  $\phi_m(n, k)$  for each time-frequency




Fig. 2. Proposed architecture.

bin  $(n, k)$  and each microphone  $m$  into a matrix of size  $M \times K$ , where  $K = N_f/2 + 1$  is the total number of frequency bins, upto the Nyquist frequency, at each time frame and  $M$  is the total number of microphones in the array. We refer to this feature representation as the *phase map*. For example, if we consider a microphone array with  $M = 4$  microphones and  $N_f = 256$ , then the input feature matrix is of size  $4 \times 129$ .

Given the input representations, the next task is to estimate the probabilities of the  $I$  DOA classes for each time frame. For this, we propose a CNN based supervised learning method, which is described in the following section.

#### IV. DOA ESTIMATION WITH CNNs

CNNs are a variant of the standard fully-connected neural network, where the architecture typically consists of one or more convolution layers followed by fully-connected networks leading to the output [29]. In this work, the main motivation behind using CNNs is to learn the discriminative features for DOA estimation from the phase map input by applying small local filters to learn the phase correlations at the different frequency sub-bands.

Given the phase map as the input, the CNN generates the posterior probability for each of the DOA classes. Let us denote the phase map for the  $n$ -th time frame as  $\Phi_n$ . Then the posterior probability generated by the CNN at the output is given by  $p(\theta_i | \Phi_n)$ , where  $\theta_i$  is a binary random variable that determines source activity from the DOA corresponding to the  $i$ -th class. In Fig. 2, the CNN architecture used in this work is shown. In the convolution layers, small filters of size  $2 \times 1$  are applied to learn the phase correlations between neighboring microphones at each frequency sub-band separately. This is in contrast to [22], where square filters of size  $2 \times 2$  were used to learn the features from the neighboring frequency bins also. However, in the case of multiple speakers neighboring frequency bins might contain dominant activity from different speakers, therefore in this work we use  $2 \times 1$  filters. These learned features for each sub-band are then aggregated by the fully connected layers for the classification task. The proposed architecture consists of at most  $M - 1$  convolution layers, where  $M$  is the number of microphones, since after  $M - 1$  layers performing 2D convolutions is no longer possible as the feature maps become vectors. While taking memory and computational constraints

into account, we also performed validation experiments in different acoustic scenarios to determine the filter size as well as the number of feature maps per layer. Increasing the filter size along the frequency dimension was found to lack robustness in acoustic conditions different from the training setup. Increasing the size along the microphone dimension, beginning from the first layer, led to degradation in performance across all conditions. Increasing the number of feature maps significantly increased the computation and memory cost, but did not exhibit any significant performance improvement.

In terms of the design choice related to the number of convolution layers, we posit that by using small filters of size  $2 \times 1$ , with each subsequent convolution layer after the first one, for each sub-band, the phase correlation information from different microphone pairs are aggregated due to the growing receptive field of the filters, and to learn from the correlation between all microphone pairs,  $M - 1$  convolution layers would be required to incorporate this information into the learned features. In Section VI-B4, we experimentally demonstrate that indeed  $M - 1$  convolution layers are required to obtain the best DOA estimation performance for a given microphone array and also show the efficiency of this design choice in terms of number of free parameters.

As stated earlier, we utilize the binary relevance method [27] to tackle the multi-label classification problem, therefore the output layer of the CNN consists of  $I$  sigmoid units, each corresponding to a DOA class. During training, the optimization of the network weights are done in terms of each output neuron separately, using binary cross-entropy as the loss function.

Here, the task of multi-source DOA estimation is performed for a signal block consisting of  $N$  time frames. The block-level posterior probability is obtained by averaging  $N$  frame-level probabilities for each  $\theta_i$ , given by

$$p_n(\theta_i) = \frac{1}{N} \sum_n^{n+N-1} p(\theta_i | \Phi_n). \quad (2)$$

From these computed average probabilities, the  $L$  DOAs corresponding to the  $L$  classes with the highest probabilities are selected as the DOA estimates. In this work, we chose this simple method to demonstrate the effectiveness of the proposed algorithm. Using more advanced post-processing methods, such as automatic peak detection [30], is beyond the scope of this paper.

## V. TRAINING DATA GENERATION

In this section, we describe the training data generation method employed in this work. Please recall that though the task of DOA estimation is performed for a block of multiple time frames, in the proposed system the probabilities of the DOA classes are estimated at each time frame. Therefore, using speech as training signals can be problematic since we would require an extremely accurate voice activity detection method to avoid including silent time frames in the training data, which can adversely affect the training. To avoid this problem, in [22], we proposed to use synthesized noise signals to generate the training data for the single speaker scenario. The problem of assigning accurate ground truth DOA labels at the STFT time frame level becomes even more difficult when the aim is to localize simultaneously active speakers. However, for such a case, using overlapping noise signals for training is also not suitable, since at each TF bin, the phase component of the observed microphone signals' STFT coefficient is a non-linear combination of the phase of the individual directional sources. This combination is dependent on the magnitude of the individual sources at each TF bin. Therefore, if overlapping noise signals are used, the CNN needs to learn from different non-linear combinations of the phases of the individual sources which would make the learning process difficult for the CNN.

To effectively use synthesized noise signals to generate the training data for more than one source per time frame, and taking into account that the aim is to localize speech sources, we utilize the assumption two simultaneously active speech sources do not overlap in the TF domain. This is known as W-disjoint orthogonality which has been shown to hold approximately for speech signals with an appropriate choice of the time and frequency resolutions [26]. In the following, we explain the procedure for generating the training data for a scenario with two active speakers.

As a first step, we generate the training signals for a single-source case by convolving the room impulse responses (RIRs) corresponding to two different directions, for a specific acoustic condition considered for training, with two synthesized white noise signals. Then, for a specific source-array setup, the STFT representation of these two multi-channel signals, corresponding to two different DOAs, are concatenated along the time frame axis. Following this, for each frequency sub-band separately, the time-frequency bins for all microphones are randomized to get a single multi-channel training signal. The phase map input is extracted from this signal by selecting the  $M \times K$  matrix for each time frame. The randomization process is included here such that at each time frame of the resulting training signal, the individual frequency sub-bands correspond to the individual activity of one of the sources. Since the activity at each sub-band is randomized, the phase map input for each frame would now include the activity from two directional sources, however, at different frequency sub-bands. This mimics the W-disjoint orthogonality property mentioned earlier. This also avoids the problem of learning from different non-linear combination of the phases of the individual signals which occurs if completely overlapping signals are used.

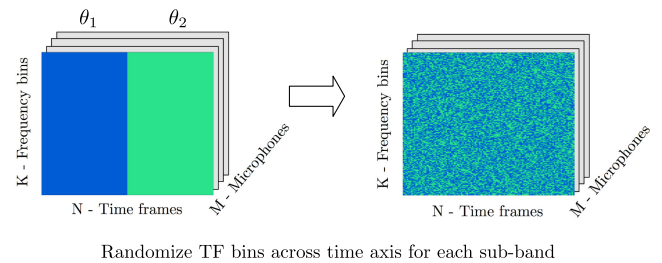


Fig. 3. Illustration of the method used for generating the training data.

There are two important things to note regarding the randomization process. First, it is essential that the randomization of the TF bins is done separately for each frequency sub-band, such that the TF bins remain in the same frequency sub-band. This is essential since phase correlations are frequency dependent. Secondly, it is essential that for each frequency sub-band, the TF bins for all the microphones are randomized together, such that phase relations between the microphones for the individual TF bins are preserved.

An illustration of this procedure is shown in Fig. 3. The figure on the left illustrates the concatenated TF representation of two directional signals, originating from two different directions,  $\theta_1$  and  $\theta_2$ . Following the randomization procedure, it can be seen that at each time frame there are approximately equal number of TF bins with activity corresponding to each of the two DOAs, assuming the length of the concatenated signals was the same. The phase map is extracted from this signal at each time frame. Thus, at each frequency sub-band of the phase map input to the CNN, the phase of the STFT coefficients for all microphones correspond to one of the sources,  $\theta_1$  or  $\theta_2$ . This makes the assumption of disjoint activity of signals implicit within the training data generation. From this phase map, the CNN can learn the relevant features for localizing multiple speakers at each time frame from the individual TF bins that contain the phase relations across the microphones for each source DOA separately. By repeating the above mentioned procedure for all possible angular combinations and acoustic conditions, we obtain the complete training dataset. The entire training data generation procedure for the two speaker case is presented in the form of two-part pseudo code in Pseudo code Part 1 and 2. In the first part, Pseudo code Part 1, the procedure for generating the simulated RIRs for training is described, where different rooms with different array positions in each room along with different source-array distances are selected, and the RIRs for each angular position at the specified distance is simulated for all the microphones. The different acoustic conditions considered for the multi-condition training of the CNN is given in Table. I. In this work, we select  $R = 5$  rooms,  $P = 7$  array positions in each room and  $D = 2$  different source-array distances for each array position. The different rooms as well as positions inside each room are considered to develop robustness in various acoustic conditions. Given the RIRs for all the different acoustic conditions, Pseudo code Part 2 describes the procedure for obtaining the training data using synthesized noise signals for all different acoustic scenarios and angular combinations.

TABLE I  
CONFIGURATION FOR TRAINING DATA GENERATION. ALL ROOMS ARE 2.7 m HIGH

Signal	Synthesized noise signals
Room size	R1: (6 × 6), R2: (5 × 4), R3: (10 × 6), R4: (8 × 3), R5: (8 × 5)
Array positions in room	7 different positions in each room
Source-array distance	1 m and 2 m for each position
RT <sub>60</sub>	R1: 0.3 s, R2: 0.2 s, R3: 0.8 s, R4: 0.4 s, R5: 0.6 s
SNR	Uniformly sampled from 0 to 30 dB

---

**Pseudo code Part 1: Generate RIRs.**


---

**Require:** Array with  $M$  microphones

- 1: Discretize DOA space for  $I$  discrete DOA values
  - 2: Select  $R$  rooms with randomly assigned RT<sub>60</sub>  $\in$  [0.2, 0.8]
  - 3: **for** each  $r \in R$  **do**
  - 4:     Randomly select  $P$  array positions
  - 5:     Select  $D$  source-array distance values
  - 6:     **for** each  $p \in P$  **do**
  - 7:         **for** each  $d \in D$  **do**
  - 8:             Simulate RIRs for  $I$  DOAs and  $M$  mics  
              using [34]
  - 9:             **Save:**  $I \times M$  RIRs
  - 10:         **end for**
  - 11:     **end for**
  - 12: **end for**
  - 13: **Save:**  $N_A$  RIR files  $\triangleright N_A = R \times P \times D$
- 

In total, the training data consisted of  $N_T \approx 12.4$  million time frames. The CNN was trained using the Adam gradient-based optimizer [31], with mini-batches of 512 time frames and a learning rate of 0.001. During training, at the end of the convolution layers and after each fully connected layer, a dropout procedure [32] with a rate of 0.5 was used to avoid over fitting. All the implementations were done in Keras [33].

Please note that in this work, the CNN is trained to estimate the probabilities of DOA classes of only two speakers given the phase map input for each STFT time frame. By following the same procedure as described above the method can be extended for estimating the DOA class probabilities of more than two speakers per time frame. In Section VI-C1, it is shown that despite such a training procedure the proposed method can estimate the DOAs of more than two speakers for a signal block with multiple time frames.

## VI. EXPERIMENTAL EVALUATION

In this section, different experiments with simulated and measured RIRs as well as real recorded data are presented to objectively evaluate the performance of the proposed system. For all the experimental evaluations except the one presented in Section VI-B4, we consider a uniform linear array (ULA) with  $M = 4$  microphones with inter-microphone distance of 8 cm, and the input signals are transformed to the STFT domain using a DFT length of  $N_f = 512$ , with 50% overlap, resulting in  $K = 257$ . The sampling frequency of the signals is  $F_s = 16$  kHz. To form the classes, we discretize the whole

---

**Pseudo code Part 2: Generate Phase Maps.**


---

**Require:**  $N_A$  RIR files with  $I \times M$  RIRs each

- 1: **for** each  $n_a \in N_A$  **do**
  - 2:     **for** each  $i \in I$  **do**
  - 3:         **for**  $j \in I$  and  $j \neq i$  **do**
  - 4:              $s_i, s_j = 2$  s long white Gaussian noise signals
  - 5:              $\mathbf{x}_i = s_i$  convolved with  $M$  RIRs for DOA  $\theta_i$
  - 6:              $\mathbf{x}_j = s_j$  convolved with  $M$  RIRs for DOA  $\theta_j$
  - 7:              $\mathbf{x}_i = \mathbf{x}_i + \mathbf{v}_i$
  - 8:              $\mathbf{x}_j = \mathbf{x}_j + \mathbf{v}_j$   $\triangleright \mathbf{v}_i, \mathbf{v}_j$ : Uncorrelated noise
  - 9:              $\mathbf{X}_i = \text{STFT}(\mathbf{x}_i)$
  - 10:              $\mathbf{X}_j = \text{STFT}(\mathbf{x}_j)$
  - 11:              $\mathbf{Z}_{i,j} =$  Concatenate  $\mathbf{X}_i$  and  $\mathbf{X}_j$  along time  
                  axis and randomize TF bins as  
                  described in Section V
  - 12:              $\mathbf{W}_{i,j} = \text{Phase}(\mathbf{Z}_{i,j})$   $\triangleright \text{Size}(\mathbf{W}_{i,j}) = M \times$   
 $K \times N_{i,j}$
  - 13:             **for** each  $n \in N_{i,j}$  **do**
  - 14:                  $\Phi_n = M \times K$  matrix for each time frame  $n$   
                  from  $\mathbf{W}_{i,j}$
  - 15:                  $\mathbf{t}_n = I \times 1$  vector with  $i$  and  $j$ -th element as  
                  1, rest 0
  - 16:             **end for**
  - 17:         **end for**
  - 18:     **end for**
  - 19: **end for**
  - 20: Concatenate all the computed phase maps  $\Phi_n$  into a  
   tensor  $\Phi$  of size  $M \times K \times N_T$
  - 21: Concatenate all target vectors  $\mathbf{t}_n$  into a matrix  $\mathbf{T}$  of size  
    $I \times N_T$   $\triangleright N_T$ : Total time frames in training dataset
- Training data:**  $\{\Phi | \mathbf{T}\}$
- 

DOA range of a ULA with a  $5^\circ$  resolution to get  $I = 37$  DOA classes, for both training and testing. All the presented objective evaluations are for the two speakers scenario. However, in Section VI-C1, we also demonstrate the ability of the proposed method to deal with scenarios with varying number of speakers.

The speech signals used for evaluation are taken from the LIBRI speech corpus. With random selected speech utterances, five different two speaker mixtures, each of length 2 s, were used. For a specific source-array setup in a room, each two speaker mixture is considered for each possible angular combination. This was done to avoid influence of signal variation on the difference in performance for different acoustic conditions. Also, since the angular space is discretized with a  $5^\circ$  resolution, and

for a proper objective evaluation (described in Section VI-A), for the experiments with simulated RIRs in Section VI-B, it was ensured that the angular distance between the two speakers in the different mixtures is at least  $10^\circ$  during test.

Since the speech utterances can have different lengths of silence at the beginning, the central 0.8 s segment of the mixtures was selected for evaluation. Considering an STFT window length of 32 ms with 50% overlap, this resulted in a signal block of  $N = 50$  time frames over which the frame-level probabilities are averaged for the final DOA estimation, as shown in (2).

### A. Baselines and Objective Measures

The performance of the proposed method is compared to two commonly used signal processing based methods: Steered Response Power with PHase Transform (SRP-PHAT) [35], and broadband MUltiple SIgnal Classification (MUSIC) [2]. For the broadband MUSIC method, to keep the comparison similar with the other methods, the MUSIC pseudo-spectrum is computed at each frequency sub-band for each STFT time frame, with an angular resolution of  $5^\circ$  over the whole DOA space, and then it is averaged over all the frequency sub-bands to get the broadband pseudo-spectrum. This is then averaged over all the time frames considered in a signal block and similar to the proposed method, the  $L$  DOAs with the highest values are selected as the final DOA estimates. Similar post-processing is also performed for the computed SRP-PHAT pseudo-likelihoods at each time frame to get the final DOA estimates for a signal block. Please note that the expectation operator in the computation of the correlation matrix for MUSIC as well as the expectation operator for computing the cross-correlations in SRP-PHAT were approximated via recursive temporal averaging with a time constant of 64 ms for both cases.

For the objective evaluation, two different measures were used: Mean Absolute Error (MAE) and localization accuracy (Acc.). The mean absolute error computed between the true and estimated DOAs for each evaluated acoustic condition is given by

$$\text{MAE}(\circ) = \frac{1}{LC} \sum_{c=1}^C \sum_{l=1}^L |\theta_l^c - \hat{\theta}_l^c|, \quad (3)$$

where  $L$  is the number of simultaneously active speakers and  $C$  is the total number of speech mixture segments considered for evaluation for a specific acoustic condition. The true and estimated DOAs for the  $l$ -th speaker in the  $c$ -th mixture are denoted by  $\theta_l^c$  and  $\hat{\theta}_l^c$ , respectively. The speaker-DOA association is based on the ascending absolute value of the true DOA of the source, i.e., the speaker with the smallest DOA value is the first speaker and so on. Similarly, the lowest estimated DOA value is considered to correspond to the first speaker.

The localization accuracy is given by

$$\text{Acc.}(\%) = \frac{\hat{C}_{\text{acc.}}}{C} \times 100, \quad (4)$$

where  $\hat{C}_{\text{acc.}}$  denotes the number of speech mixtures for which the localization of the speakers is accurate. In our evaluation, the localization of speakers for a speech segment is considered

TABLE II  
CONFIGURATION FOR GENERATING TEST DATA. ALL ROOMS ARE 3 m HIGH

Signal	Speech signals from LIBRI
Room size	Room 1: (5 × 7) m , Room 2: (9 × 4) m
Array positions in room	4 arbitrary positions in each room
Source-array distance	1.3 m for Room 1, 1.7 m for Room 2
RT <sub>60</sub> (s)	Room 1: 0.38 , Room 2: 0.70

accurate if the distance between the estimated and the true DOA for all the speakers is less than or equal to  $5^\circ$ .

### B. Experiments With Simulated RIRs

In this section, first, the performance of the proposed method is evaluated for acoustic conditions different from those considered during training, in the presence of varying levels of spatially uncorrelated white noise in Section VI-B1. Then, we evaluate the performance in the presence of varying levels of diffuse babble noise, a noise type which was unseen during training, along with a constant level of spatially white noise in Section VI-B2. In Section VI-B4, we study the influence of the number of convolution layers on the performance of the proposed method and empirically demonstrate the optimal choice for the number of convolution layers for the proposed method.

1) *Generalization to Unseen Acoustic Conditions:* To evaluate the performance of the methods for unseen acoustic conditions, we consider two rooms with different reverberation times as shown in Table II. In each room, the ULA is placed at four different positions and for each of these array positions, the two speakers from each of the five considered mixtures are placed at different angular positions at the same specified source-array distance. For each array position, the total number of mixtures considered for evaluation is  $C = 630 * 5 = 3150$ , where 630 corresponds to the number of possible angular combinations with the constraint of  $10^\circ$  angular separation between the two speakers for each of the five mixtures.

The performance of the three methods under test is evaluated for three different levels of spatially white noise, with input SNRs 10, 20 and 30 dB, for both the rooms and the results in terms of the two considered objective measures are presented in Table III. The shown results for each input SNR was averaged over the four different array positions considered in each room.

From the results, it can be seen that the proposed method is able to provide accurate localization performance in acoustic environments that were not part of the training data. For input SNR of 30 dB, it manages to localize both sources accurately in 98% of the speech mixtures and shows a very low MAE. As the noise level increases, the performance worsens, however, it always provides a much better localization accuracy and much lower error compared to both MUSIC and SRP-PHAT.

Considering same noise level, performance of the proposed method in both rooms is relatively similar compared to the signal processing based methods which have a considerably better performance in the less reverberant room (Room 1). One of the main reasons for this difference is the assumption of free-field sound propagation in the formulation of the signal processing based methods which leads to considerable deterioration in their performance in more reverberant conditions. The proposed

TABLE III  
RESULTS FOR TWO DIFFERENT ROOMS WITH VARYING LEVELS OF SPATIALLY WHITE NOISE COMPUTED OVER 3150 SPEECH SEGMENTS OF 0.8 s FOR EACH ARRAY POSITION. FOR EACH SNR, THE RESULT IS AVERAGED OVER THE FOUR DIFFERENT ARRAY POSITIONS IN THE ROOM

Test Room	Room 1						Room 2					
	10 dB		20 dB		30 dB		10 dB		20 dB		30 dB	
Measure	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.
SRP-PHAT	26.7	37.7	22.7	49.3	18.9	60.4	29.2	28.2	27.6	36.9	21.6	48.3
MUSIC	23.3	43.4	16.2	63.4	13.4	71.3	27.1	37.7	18.6	51.7	16.6	59.6
Proposed	14.5	73.5	3.5	93.2	1.5	98.1	16.8	63.4	4.3	88.9	2.7	96.3

TABLE IV  
RESULTS FOR TWO DIFFERENT ROOMS WITH VARYING LEVELS OF BABBLE NOISE COMPUTED OVER 3150 SPEECH SEGMENTS OF 0.8 s FOR EACH ARRAY POSITION. FOR EACH SNR, THE RESULT IS AVERAGED OVER THE FOUR DIFFERENT ARRAY POSITIONS IN THE ROOM

Test Room	Room 1						Room 2					
	-5 dB		0 dB		5 dB		-5 dB		0 dB		5 dB	
Measure	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.
SRP-PHAT	22.4	40.8	21.8	46.1	19.9	57.8	23.7	40.2	20.8	46.6	20.1	48.3
MUSIC	23.9	39.2	18.8	49.4	16.3	59.9	25.9	36.3	19.2	49.9	18.1	52.1
Proposed	5.0	91.9	2.1	96.8	1.1	98.7	7.1	82.9	3.4	94.3	2.0	97.5

supervised learning based method is trained in a diverse set of acoustic conditions, leading to a much better robustness to adverse acoustic environments.

Overall, it can be seen that the proposed method has a superior performance, in terms of both MAE and localization accuracy, compared to the traditional signal processing based methods for all the different levels of spatially white noise in both rooms. Among the two signal processing based methods, MUSIC performs much better since the averaged broadband MUSIC pseudo-spectrum contains clearer peaks compared to SRP-PHAT which tends to exhibit a flatter distribution over the DOAs.

2) *Generalization to Unseen Noise Type*: In the previous experiment, the performance of the proposed method was evaluated for different levels of spatially white noise, which is a noise type seen by the network during training. In this section, we consider the presence of diffuse babble noise in the acoustic environment, which has different spatial as well as spectral characteristics compared to white noise, and is a noise type with which the CNN was not trained. A 40 s long sample of multi-channel diffuse babble noise was generated using the acoustic noise field generator [36], assuming an isotropic spherically diffuse noise field. The generated babble noise was divided into 20 segments of 2 s each and randomly chosen segments were added to each mixture.

The performance of the methods was evaluated for three different input SNRs of babble noise: -5 dB, 0 dB and 5 dB. Along with diffuse babble noise, spatially white noise with an input SNR of 40 dB was also added and results for the two different rooms are shown in Table IV. Similar to previous experiment, results for each input SNR of babble noise was averaged over the four different array positions considered in each room.

Though the proposed method is not trained with diffuse babble noise, it can be seen from the results that even at the lowest input SNR of -5 dB, the proposed method is able to perform

accurate localization of the two speakers in both rooms for approximately 90% of the speech mixtures. Since we consider an isotropic spherically diffuse noise field, the spatial coherence of the babble noise is frequency dependent whereas white noise is incoherent for all frequencies. Despite this difference, since the proposed method is trained to localize directional sources and due to multi-condition training, as long as the noise source is not directional the proposed method can provide very good performance.

If the results from Table III are compared to Table IV, it can be seen that the deterioration in performance of the proposed method, in terms of the objective measures, as the noise levels increase is more prominent when white noise is considered compared to diffuse babble noise. The main reason for this difference is the spectral characteristics of the two different types of noises. On the one hand, spatially white noise is present across the spectrum, therefore the input features at all frequency sub-bands are equally affected. On the other hand, babble noise is mostly dominant at low frequencies, therefore since each filter kernel in the convolution layers of the CNN learns from the complete input feature space, the filters are able to extract the relevant features for localization from the high SNR regions of the input to compensate for the lack of information in the low SNR regions.

Overall, the proposed method provides a much better localization accuracy and lower error than the signal processing based methods, with the difference in performance being especially significant at low input SNRs of diffuse babble noise.

3) *Influence of Source-Array Distance*: The CNN used for the earlier evaluations was trained for each room and array position for two specific source-array distances of 1 m and 2 m. To investigate the influence of source-array distance, in this part, the localization performance of the proposed method is evaluated for varying source-array distances.



For this experiment, we simulated a room with dimensions  $10 \times 11 \times 3 \text{ m}^3$  and a reverberation time of 0.38 s. The test data was generated for three different array positions. For each of these array positions, the sound sources were placed at distances varying from 0.4 m to 3 m. It should be noted that both the speakers were placed at the same distance for each setup. A single two speaker mixture was used and spatially white noise was added resulting in input SNR of 20 dB.

The results for this experiment, in terms of both MAE and localization accuracy, is shown in Fig. 5. Each point in the plot corresponds to a specific source-array distance. For each of these points, the measures were averaged over all possible angular combinations for the two speakers at each of the different array positions in the room.

From the result plots, it can be seen that when the sources are very close to the microphone array the error in localization is higher, since the CNN was trained considering a far-field scenario, however, for very small source-array distances, the sources are essentially in the near-field of the array. The minimum error as well as maximum accuracy in localization can be observed for the two specific distances of 1 m and 2 m, which were part of the training setup. Additionally, for distances close to these training distances, the errors are also relatively lower. When the sources are between the two training distances, the errors are slightly higher, however, if we observe the absolute value of the MAE as well as the accuracy, the degradation in performance is not significant. Similarly for distances larger than 2 m, it can be seen that the localization performance deteriorates slightly.

Overall, observing the absolute value of the objective measures, it can be seen that though the network is trained with two specific source-array distances, there is small deterioration in performance for other distances, except when the sources are very close to the microphone array.

4) *Influence of Number of Convolution Layers:* In the previous experiments, we considered a ULA with  $M = 4$  microphones and the CNN architecture used was the same architecture that was proposed in [22], [23] which consisted of three convolution layers followed by two fully connected layers. In this section we empirically demonstrate that given the choice of small filters of size  $2 \times 1$  for all the convolution layers, with the aim to learn the relevant features for localization from the phase correlations at neighboring microphones, a CNN architecture with three convolution layers is not always the best performing architecture. Here we show that the number of convolution layers need to be  $M - 1$  to obtain the best localization performance.

For this experiment we consider a ULA with 8 microphones with an inter-microphone distance of 2 cm. From this array, we select two sub-arrays, one with 6 microphones and the other with 4 microphones that are formed by selecting the respective number of middle microphones from the main eight element array, as shown in Fig. 4, to get a ULA with  $M = 6$  and another ULA with  $M = 4$ , respectively. All the arrays have the same inter-microphone distance and array center.

Using the same training data configuration from previous experiments (Table I), multiple CNNs with number of convolution layers varying from 2 to  $M - 1$  are trained for each of the arrays. The number of convolution layers is restricted to  $M - 1$




Fig. 4. Array setup for experiment presented in Section VI-B4.

since further 2D convolution layers are not possible as the microphone dimension of the phase map input is reduced to 1 after the  $M - 1$ -th layer. For the eight microphone array, 6 CNNs are trained, whereas for the six microphones and the four microphone array, 4 and 2 CNNs are trained, respectively. All the networks were trained with the same amount of data. To analyze the performance of the 12 different trained networks, test data corresponding to the Room 1 configuration in Table II is generated for each of the arrays. Spatially white noise is added for an input SNR of 30 dB.

The results for this experiment, in terms of both MAE and localization accuracy, is shown in Fig. 6. In the figures, the center of the circle markers correspond to the value of the objective measure and the area of the markers denote the number of trainable/free parameters for that specific network.

The first trend that can be noticed from the figures is that for each of the arrays, as the number of convolution layers is decreased from  $M - 1$  the performance of the networks degrades in terms of both MAE and localization accuracy. This shows that with small filters of size  $2 \times 1$ , to aggregate the phase correlation features from all the microphone pairs in an array,  $M - 1$  convolution layers are required. When lesser number of convolution layers are used, as the same filter size is used in each of these layers, phase correlation information from all microphone pairs are not incorporated into the learned features leading to deterioration in performance.

It can also be seen from the figures that the best localization performances of the three arrays is different and it is better for the array with higher number of microphones. This difference in performance comes from the different apertures of the considered arrays, and similar to signal processing based localization methods, here also we observe better performance for a ULA with a larger aperture. A similar observation can be made by comparing the localization results from Section VI-B1, where the ULA consists of only four microphones, however, the inter-microphone distance is 8 cm compared to 2 cm used in this experiment leading to an array with a larger aperture.

In Fig. 6, we also observe that as the number of convolution layers is decreased the number of trainable/free parameters increases, as depicted by the area of the markers for each network. From Fig. 2, it can be seen that when  $M - 1$  convolution layers are used, the size of each feature map at the end of the convolution layers is always  $1 \times K$ . As the number of convolution layers is decreased the size of each feature map at the end of the convolution layers actually becomes larger leading to a larger number of trainable/free parameters for the complete network. This further demonstrates the need of  $M - 1$  convolution layers, as very large number of free parameters can lead to problems of over fitting, if the amount of available training data is not sufficient.

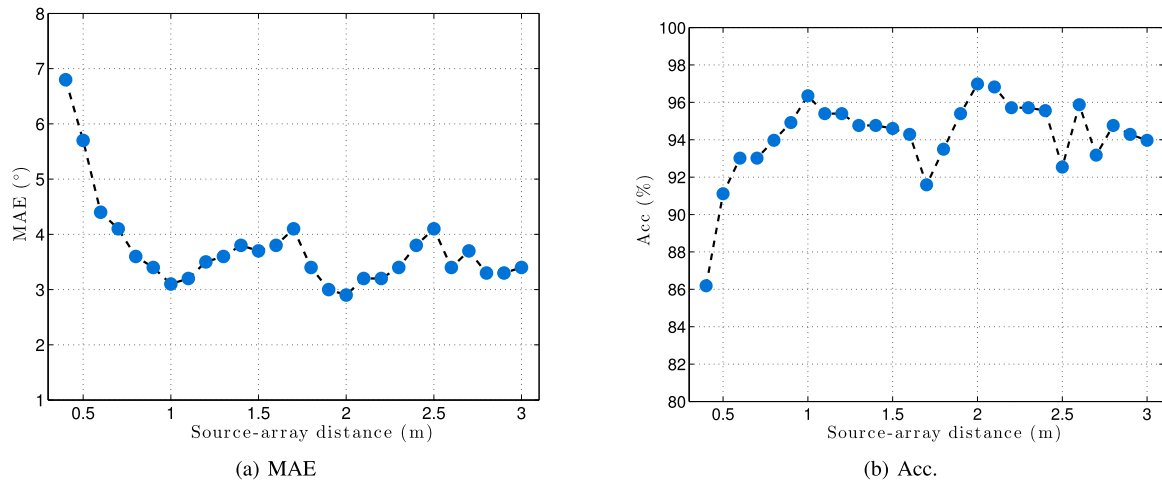


Fig. 5. Results for the experiment showing the performance of the proposed method for increasing source-array distances presented in Section VI-B3.

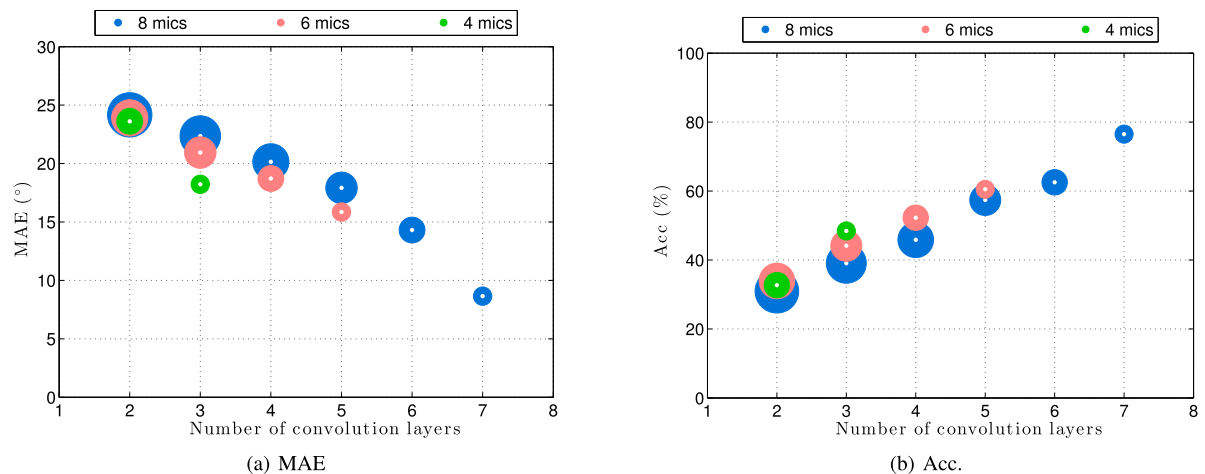


Fig. 6. Results for the experiment on the influence of convolution layers on the proposed method presented in Section VI-B4.

Since the requirement of  $M - 1$  convolution layers is mainly related to the aggregation of information in the feature space by the slowly growing receptive field of the small filters used in our framework, techniques for a more aggressive expansion of the receptive field of the filters can also be employed. This is however beyond the scope of this paper and is a topic for future research.

### C. Experiments With Measured RIRs

For the experiments with measured RIRs, we used the Multichannel Impulse Response Database from Bar-Ilan University [37]. The database consists of RIRs measured at Bar-Ilan University's acoustics lab, of size  $6 \times 6 \times 2.4 \text{ m}^3$ , for three different reverberation times of  $\text{RT}_{60} = 0.160, 0.360, \text{ and } 0.610$  s. The recordings were done for several source positions placed on a spatial grid of semi-circular shape covering the whole angular range for a linear array, i.e.,  $[0^\circ, 180^\circ]$ , in steps of  $15^\circ$  at distances of 1 m and 2 m from the center of the microphone array.

The recordings were done with a linear microphone array with three different microphone spacings. For our experiment, we chose the  $[8, 8, 8, 8, 8, 8, 8, 8]$  cm setup [37], which consists of eight microphones where the distance between the microphones

is 8 cm. We selected a sub-array of the four middle microphones out of the total eight microphones used in the original setup, to have a ULA with  $M = 4$  elements with an inter-microphone distance of 8 cm, which corresponds to the array setup used in experiments with simulated RIRs. Therefore, the CNN trained with simulated data used for the earlier evaluations in Section VI-B1 and VI-B2 was also used for these experiments. We used the same five mixtures from earlier, with the total number of mixtures for evaluation being  $C = 76 * 5 = 380$ , where 76 is the number of all possible angular combinations with discretization of the complete DOA space of a ULA with  $15^\circ$  resolution.

The results for all the different reverberation times and source-array distances are shown in Table V. For this experiment, spatially white noise was added to each mixture resulting in an input SNR of 30 dB.

Even when trained with simulated data only, the results show that the proposed method is able to provide good localization performance in real conditions, even when the sources are placed far from the array in reverberant conditions. The performance of all the compared methods is better when the sources are close to the array, however, the difference in performance, for different distances, for the signal processing based methods is

TABLE V  
RESULTS WITH MEASURED RIRs

RT <sub>60</sub>	0.160 s				0.360 s				0.610 s			
	1 m		2 m		1 m		2 m		1 m		2 m	
Measure	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.	MAE	Acc.
SRP-PHAT	12.8	75.0	15.33	64.2	15.8	61.8	19.8	49.2	15.3	57.4	21.5	42.9
MUSIC	4.9	87.0	9.33	78.2	10.4	72.8	15.2	54.2	11.3	70.7	18.5	47.3
Proposed	1.9	89.7	3.4	86.1	3.27	88.2	4.35	79.9	3.14	85.5	4.43	80.2

considerable since the effect of reverberation is more significant when the sources are further away from the array.

Overall, the proposed method provides significantly better performance compared to both MUSIC and SRP-PHAT, and the difference is more prominent as the acoustic environment becomes more reverberant.

1) *Dynamic Acoustic Scenario*: In all the previous experiments, we considered the two speaker scenario for the evaluation of the performance of the proposed method. In this experiment we show that even though the CNN is trained to estimate the frame-level probabilities of a maximum of two sources, with the proposed method it is possible to estimate the DOA of more than two sources for a short segment. Simultaneously, it is also shown that since the input to the CNN is the phase map for a single STFT time frame, the proposed method is also able to handle dynamic acoustic scenarios where the number of speakers changes over time.

For this experiment, we consider the reverberation time of 0.36 s and source-array distance of 2 m from the measured RIR database used in the previous experiment. A 6 s speech mixture segment is created where for the first 1 s only one source from 60° is active. For the next 2 s, an additional source is active from 105°. A third source from 135° is active for the next 2 s along with the first two sources. For the final 1 s duration, only the third source is active. The source activities for each segment and the corresponding ground truth DOAs of the sources are shown in the bottom figure of Fig. 7(a). Spatially white noise and diffuse babble noise are added to the speech mixture resulting in input SNRs of 40 dB and 5 dB, respectively.

The estimated frame-level probabilities for the proposed method and broadband MUSIC are depicted in the top and middle figures of Fig. 7(a), respectively. Since from the previous experiments, it was found that MUSIC is the better performing method out of the two considered signal processing based techniques, the results for SRP-PHAT are not presented. It can be seen that the estimated frame-level probabilities for the proposed method is much more concentrated towards the actual source DOAs compared to MUSIC.

In Fig. 7(b), the frame level probabilities are averaged over the time frames in each segment and then normalized to a maximum value of 1. This specific normalization is done for the purpose of visualization only. From these figures, it can be seen that the proposed method exhibits much clearer peaks at the true source DOAs compared to MUSIC which lead to the superior performance of the proposed method in previously presented evaluations even with the simple post-processing method considered in this work for obtaining the final DOA estimates. It

can also be seen that in the segment S3, where three sources are simultaneously active, though the network is trained to estimate frame level probabilities of two speakers, clear peaks are visible at all the three true source DOAs. Also, when only one source is active (S1 and S4), the highest peaks correspond to the true DOA.

#### D. Experiments With Real Recordings

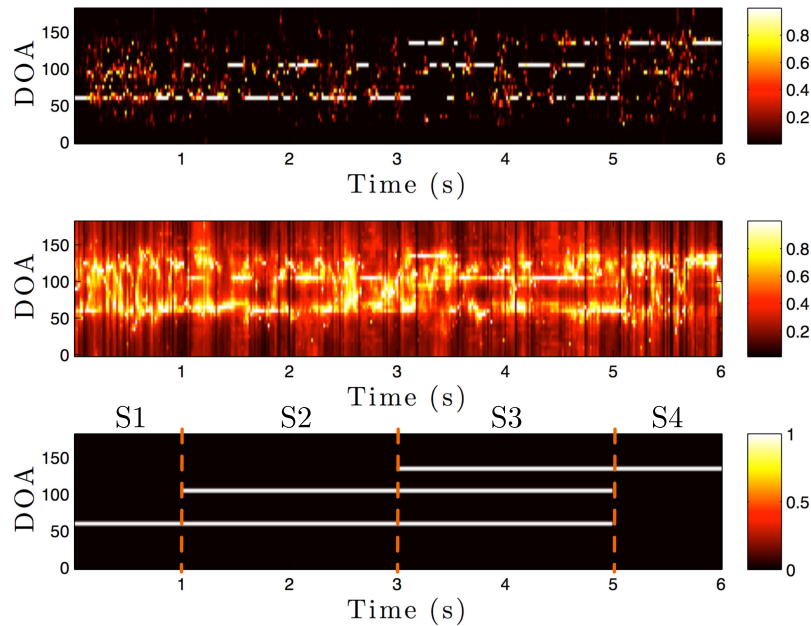
To further analyze the practical applicability of the proposed approach, we conducted experiments with real speech recordings. The recordings were done in a seminar room at Fraunhofer IIS, of size  $9.8 \times 5 \times 2.9 \text{ m}^3$  with a reverberation time of  $RT_{60} \approx 0.43 \text{ s}$ . The room consists of multiple tables and chairs along with a white-board at the front and a projector mounted on the ceiling.

For the recordings, we used a ULA with  $M = 4$  AKG cardioid microphones and an inter-microphone distance of 8 cm, similar to the array considered in Sections VI-B and VI-C. The array was mounted on a stand at a height of 1.45 m from the floor. The CNN trained with simulated data used for the earlier evaluations in Section VI-B1, VI-B2 and VI-C was also used for these experiments. Please note that the CNN was trained assuming omni-directional microphones. Also, no phase or position calibration was performed prior to recording.

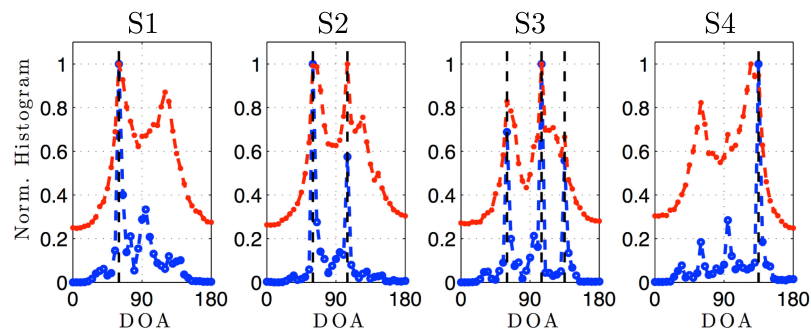
We considered two recording setups with different speaker combinations at different spatial positions. The recording setups and the corresponding results are described in the following.

1) *Recording Setup 1*: In the first setup we consider a female and a male speaker positioned at 115° and 70°, respectively. The female speaker was at a distance 2.5 m from the array center whereas the male speaker was at a distance of 2.2 m. This recording was 27 s long, where the female speaker first spoke in Spanish for about 6 s followed by short period of silence. It was then followed by the male speaker speaking in English for few seconds followed by another short period of silence. In the final segment, both speakers spoke in their respective languages at the same time for about 6 s. The speakers read excerpts from their mobile phones, therefore were not always facing the array and head movements were unavoidable.

The complete recording was transformed to the STFT domain using the same parameters as earlier and the phase map for each STFT time frame was given as input to the CNN. The localization result for this recording setup with the proposed method is shown in Fig. 8(a). In the upper figure, from the top, the estimated frame-level probabilities using the proposed method, estimated probabilities using MUSIC, the ground truth



(a) Frame level DOA probabilities for the proposed method (top) and MUSIC (middle). The ground truth DOAs and source activities for each segment are shown in the bottom figure.



(b) Normalized histogram computed from the frame level probabilities for each segment. Blue lines correspond to the proposed method while the Red lines correspond to MUSIC.

Fig. 7. Results for experiment presented in Section VI-C1 with measured RIR and a four microphone ULA. The reverberation time of the room is 0.36 s with the source placed 2 m away from the array center. Spatially uncorrelated noise and diffuse babble noise were added to the mixture signal with input SNRs of 40 dB and 5 dB, respectively.

and the spectrogram of the recording are shown. In the bottom figure, the averaged probabilities for the three distinct speech activity segments are shown. For the averaging,  $N = 50$  time frames from the different segments were selected.

From the results, it can be seen that even though the CNN was trained with simulated data, it is able to accurately localize simultaneously active human speakers in a real scenario. Also, the mismatch between the microphones considered for training and those used for the recording does not hamper the localization ability of the proposed method. From the figures showing averaged probabilities for each DOA class, it can be seen that in the single speaker segments the highest peaks for the proposed method correspond to the true DOA of the speaker. In the segment with simultaneous activity of the speakers also, the two highest peaks for the proposed method correspond to the true speaker DOAs. The averaged probabilities for MUSIC exhibit clear peaks around the true DOAs for the single speaker case. However, for the segment with simultaneous activity of the speakers, the peak for one of the speakers is not clear.

The silence periods in the recordings were introduced to demonstrate the behavior of the proposed method when there is no speech activity. From the estimated frame-level probabilities for these segments, it can be seen that the proposed method does not assign a high probability to any specific DOA class, rather all the DOA classes get assigned low probabilities leading to an almost uniform distribution across all the DOA classes. In contrast, MUSIC tends to assign high likelihoods in the silence periods to all directions. This makes the proposed approach more suitable than MUSIC for practical purposes as these silence or no speech activity periods can be explicitly detected via simple thresholding.

2) *Recording Setup 2*: In this setup, we considered a more challenging scenario with two male speakers positioned at  $115^\circ$  and  $100^\circ$ , at distances of 2.5 m and 2 m, respectively. The female speaker from the previous setup was substituted by a German speaking male speaker at the same position while the English speaking male speaker was moved closer to the other speaker. This recording was 25 s long with the speech activity pattern similar to the previous setup.

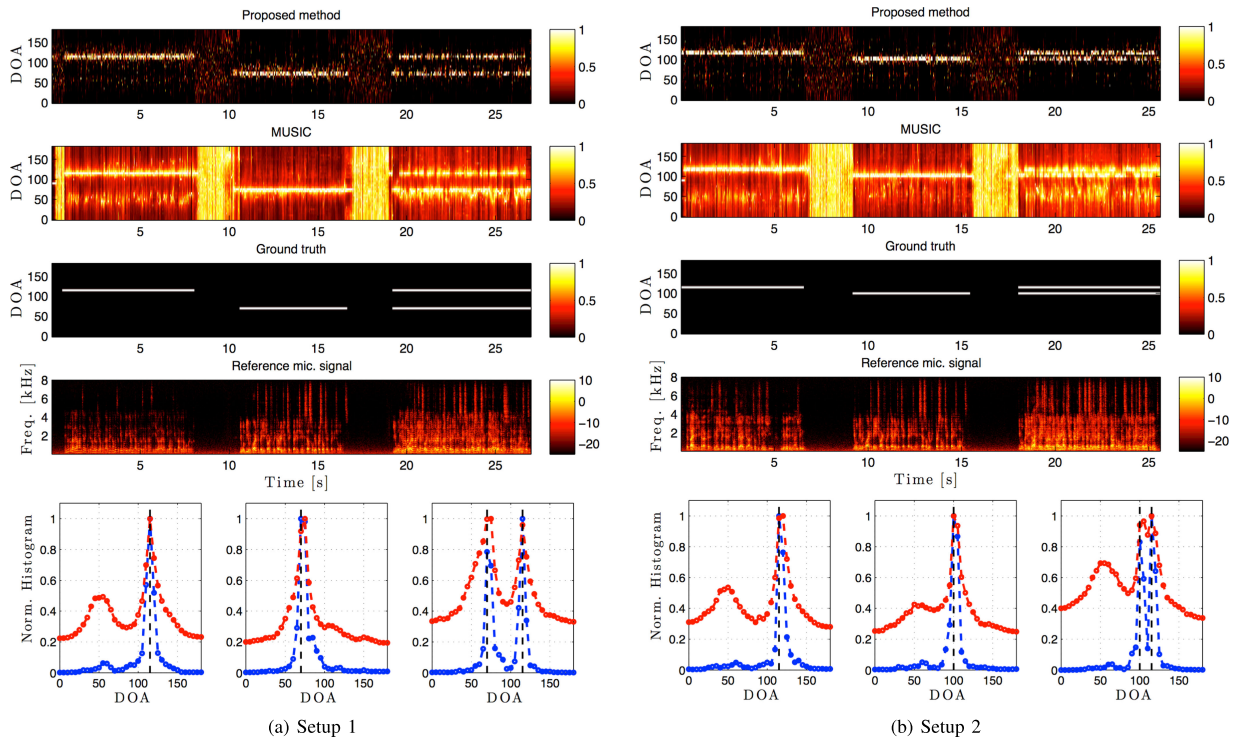


Fig. 8. Results for the experiments with real recordings. For each setup, the estimated frame-level probabilities using the proposed method (top) and MUSIC (second from the top) are shown, followed by the ground truth and the spectrogram of the recording. The averaged probabilities for the three speech activity segments are shown in the bottom figure. Blue lines correspond to the proposed method while the Red lines correspond to MUSIC.

The results for this setup are shown in Fig. 8(b). From the estimated frame-level probabilities, it can be seen that the proposed method is able to accurately localize simultaneously active human speakers positioned  $15^\circ$  apart in a real scenario. From the averaged probabilities for the corresponding segment, it can be seen that the even for this challenging scenario, the two highest peaks correspond to the true speaker DOAs. In contrast, it can be seen that in this challenging setup, it becomes difficult for MUSIC to localize the two speakers in the segment with simultaneous activity. A similar behavior to the previous setup can be observed during the silence periods for both methods.

## VII. CONCLUSION

A convolutional neural network based supervised learning method for DOA estimation of multiple speakers was presented that is trained using synthesized noise signals. Through experimental evaluation, it was shown that the proposed method provides excellent localization performance in unseen acoustic environments as well as in the presence of unseen noise types. It was also shown to exhibit a far superior performance compared to the signal processing based localization methods, SRP-PHAT and MUSIC, for the tested conditions. The ability of the proposed method to deal with acoustic scenarios with varying number of sources was shown. Using experiments with real recordings, we demonstrated the practical applicability of the proposed approach.

For the design choice of the number of convolution layers in the proposed architecture, it was empirically shown that for a microphone array with  $M$  microphones,  $M - 1$  convolution layers are required for the best localization performance. It was

also shown that such a choice leads to lesser number of trainable parameters. The choice of  $M - 1$  convolution layers is required for the aggregation of the phase correlation information from all microphone pairs in the extracted features, when using contiguous convolution operations, as done in this work.

## ACKNOWLEDGMENT

The authors would like to thank Moritz Wild for his help with the setup for the real recordings.

## REFERENCES

- [1] R. O. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propag.*, vol. AP-34, no. 3, pp. 276–280, Mar. 1986.
- [2] J. P. Dmochowski, J. Benesty, and S. Affes, "Broadband MUSIC: Opportunities and challenges for multiple source localization," in *Proc. IEEE Workshop Appl. Signal Process. Audio Acoust.*, Oct. 2007, pp. 18–21.
- [3] C. Knapp and G. Carter, "The generalized correlation method for estimation of time delay," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-24, no. 4, pp. 320–327, Aug. 1976.
- [4] J. H. Dibiase, "A high-accuracy, low-latency technique for talker localization in reverberant environments using microphone arrays," Ph.D. dissertation, Brown University, Providence, RI, USA, 2000.
- [5] J. Benesty, J. Chen, and Y. Huang, *Microphone Array Signal Processing*. Berlin, Germany: Springer-Verlag, 2008.
- [6] J. Benesty and Y. Huang, Eds., *Adaptive Signal Processing: Application to Real-World Problems* (Series Signals and Communication Technology). Berlin, Germany: Springer-Verlag, 2003.
- [7] P. Stoica and K. C. Sharman, "Maximum likelihood methods for direction-of-arrival estimation," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 7, pp. 1132–1143, Jul. 1990.
- [8] S. Delikaris-Manias, D. Pavlidis, A. Mouchtaris, and V. Pulkki, "DOA estimation with histogram analysis of spatially constrained active intensity vectors," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Mar. 2017, pp. 526–530.

- [9] A. Moore, C. Evers, and P. Naylor, "2D direction of arrival estimation of multiple moving sources using a spherical microphone array," in *Proc. Eur. Signal Process. Conf.*, 2016, pp. 1217–1221.
- [10] M. Mandel, D. Ellis, and T. Jebara, "An EM algorithm for localizing multiple sound sources in reverberant environments," in *Proc. Adv. Neural Inf. Process. Syst.*, 2006, vol. 19, pp. 953–960.
- [11] C. Blandin, A. Ozerov, and E. Vincent, "Multi-source TDOA estimation in reverberant audio using angular spectra and clustering," *Signal Process.*, vol. 92, no. 8, pp. 1950–1960, 2012. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0165168411003525>
- [12] O. Nadiri and B. Rafaely, "Localization of multiple speakers under high reverberation using a spherical microphone array and the direct-path dominance test," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 22, no. 10, pp. 1494–1505, Oct. 2014.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. 25th Int. Conf. Adv. Neural Inf. Process. Syst., 26th Annu. Conf. Neural Inf. Process. Syst.*, 2012, pp. 1106–1114.
- [14] G. Hinton *et al.*, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 82–97, Nov. 2012.
- [15] T. Hirvonen, "Classification of spatial audio location and content using convolutional neural networks," in *Proc. 138th Audio Eng. Soc. Conv.*, 2015, vol. 2.
- [16] 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 Trans. Audio, Speech, Lang. Process.*, vol. 25, no. 12, pp. 2444–2453, Dec. 2017.
- [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 *Proc. IEEE 26th Int. Workshop Mach. Learn. Signal Process.*, 2016, pp. 1–6.
- [18] R. Takeda and K. Komatani, "Sound source localization based on deep neural networks with directional activate function exploiting phase information," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Mar. 2016, pp. 405–409.
- [19] 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 *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Apr. 2015, pp. 2814–2818.
- [20] R. Takeda and K. Komatani, "Discriminative multiple sound source localization based on deep neural networks using independent location model," in *Proc. IEEE Spoken Lang. Technol. Workshop*, Dec. 2016, pp. 603–609.
- [21] N. Yalta, K. Nakadai, and T. Ogata, "Sound source localization using deep learning models," *J. Robot. Mechatronics*, vol. 29, pp. 37–48, 2017.
- [22] S. Chakrabarty and E. A. P. Habets, "Broadband DOA estimation using convolutional neural networks trained with noise signals," in *Proc. IEEE Workshop Appl. Signal Process. Audio Acoust.*, Oct. 2017, pp. 136–140.
- [23] S. Chakrabarty and E. A. P. Habets, "Multi-speaker localization using convolutional neural network trained with noise," in *Proc. Mach. Learn. Audio Process. (MLAudio) Workshop Neural Inf. Process.*, 2017.
- [24] W. He, P. Motlicek, and J. Odobez, "Deep neural networks for multiple speaker detection and localization," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2018, pp. 74–79.
- [25] S. Adavanne, A. Politis, and T. Virtanen, "Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network," in *Proc. Eur. Signal Process. Conf.*, 2018, pp. 1462–1466.
- [26] S. Rickard and O. Yilmaz, "On the approximate W-disjoint orthogonality of speech," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, May 2002, pp. 529–532.
- [27] J. Read, B. Pfahringer, G. Holmes, and E. Frank, "Classifier chains for multi-label classification," *Mach. Learn.*, vol. 85, no. 3, pp. 333–359, Jun. 2011. [Online]. Available: <https://doi.org/10.1007/s10994-011-5256-5>
- [28] F. Stoeter, S. Chakrabarty, B. Edler, and E. A. P. Habets, "Classification vs. regression in supervised learning for single channel speaker count estimation," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2018, pp. 436–440.
- [29] Y. LeCun and Y. Bengio, *The Handbook of Brain Theory and Neural Networks*, (Convolutional Networks for Images, Speech, and Time Series), M. A. Arbib, Ed., Cambridge, MA, USA: MIT Press, 1998, pp. 255–258. [Online]. Available: <http://dl.acm.org/citation.cfm?id=303568.303704>
- [30] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. Cambridge, MA, USA: MIT Press, 2009.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, arXiv:1412.6980.
- [32] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, Jan. 2014.
- [33] F. Chollet *et al.*, "Keras," 2015. [Online]. Available: <https://github.com/fchollet/keras>
- [34] E. A. P. Habets, "Room impulse response (RIR) generator," 2016. [Online]. Available: <https://github.com/ehabets/RIR-Generator>
- [35] M. S. Brandstein and H. F. Silverman, "A robust method for speech signal time-delay estimation in reverberant rooms," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Apr. 1997, vol. 1, pp. 375–378.
- [36] E. A. P. Habets and S. Gannot, "Generating sensor signals in isotropic noise fields," *J. Acoustical Soc. Amer.*, vol. 122, no. 6, pp. 3464–3470, Dec. 2007.
- [37] E. Hadad, F. Heese, P. Vary, and S. Gannot, "Multichannel audio database in various acoustic environments," in *Proc. Int. Workshop Acoust. Echo Noise Control*, Sep. 2014, pp. 313–317.



**Soumitro Chakrabarty** (S'10) received the B.Eng. degree from Manipal University, Manipal, India, in 2010 and the M.Sc. degree from Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland, in 2013. He is currently working toward the Ph.D. degree with International Audio Laboratories Erlangen, Germany on the topic of robust estimation of spatial information for microphone array processing.

His current research interests include spatial filtering, sound source localization, multi-microphone source extraction, and machine learning for array processing.



**Emanuël A. P. Habets** (S'02–M'07–SM'11) received the B.Sc. degree in electrical engineering from the Hogeschool Limburg, The Netherlands, in 1999, and the M.Sc. and Ph.D. degrees in electrical engineering from the Technische Universiteit Eindhoven, Eindhoven, The Netherlands, in 2002 and 2007, respectively.

He is currently an Associate Professor with the International Audio Laboratories Erlangen (a joint institution of the Friedrich-Alexander-Universität Erlangen-Nürnberg and Fraunhofer IIS), and the

Head of the Spatial Audio Research Group, Fraunhofer IIS, Erlangen, Germany.

From 2007 to 2009, he was a Postdoctoral Fellow with the Technion—Israel Institute of Technology and Bar-Ilan University, Israel. From 2009 to 2010, he was a Research Fellow in the Communication and Signal Processing Group, Imperial College London, U.K.

His research activities center around audio and acoustic signal processing, and include spatial audio signal processing, spatial sound recording and reproduction, speech enhancement (dereverberation, noise reduction, and echo reduction), and sound localization and tracking.

Dr. Habets was a member of the organization committee of the 2005 International Workshop on Acoustic Echo and Noise Control in Eindhoven, The Netherlands, a General Co-Chair of the 2013 International Workshop on Applications of Signal Processing to Audio and Acoustics in New Paltz, NY, USA, and General Co-Chair of the 2014 International Conference on Spatial Audio in Erlangen, Germany. He was a member of the IEEE Signal Processing Society Standing Committee on Industry Digital Signal Processing Technology from 2013 to 2015, a Guest Editor for the IEEE JOURNAL OF SELECTED TOPICS IN SIGNAL PROCESSING and the *EURASIP Journal on Advances in Signal Processing*, an Associate Editor for the IEEE SIGNAL PROCESSING LETTERS from 2013 to 2017, and the Editor-in-Chief for the *EURASIP Journal on Audio, Speech, and Music Processing* from 2016 to 2018. He is the co-recipient, with S. Gannot and I. Cohen, of the 2014 IEEE Signal Processing Letters Best Paper Award. He is currently a member of the IEEE Signal Processing Society Technical Committee on Audio and Acoustic Signal Processing, and the Chair of the EURASIP Special Area Team on Acoustic, Sound, and Music Signal Processing.