



HAL
open science

BeamLearning: an end-to-end Deep Learning approach for the angular localization of sound sources using raw multichannel acoustic pressure data

Hadrien Pujol, Eric Bavu, Alexandre Garcia

► To cite this version:

Hadrien Pujol, Eric Bavu, Alexandre Garcia. BeamLearning: an end-to-end Deep Learning approach for the angular localization of sound sources using raw multichannel acoustic pressure data. *Journal of the Acoustical Society of America*, 2021, Special issue on Machine Learning in Acoustics, 149 (6), pp.4248-4263. 10.1121/10.0005046 . hal-03219058

HAL Id: hal-03219058

<https://hal.science/hal-03219058v1>

Submitted on 6 May 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

BeamLearning: an end-to-end Deep Learning approach for the angular localization of sound sources using raw multichannel acoustic pressure data

Hadrien Pujol
LMSSC, Cnam,
Paris, France
hadrien.pujol@lecnam.net

Éric Bavu
LMSSC, Cnam,
Paris, France
eric.bavu@lecnam.net

Alexandre Garcia
LMSSC, Cnam,
Paris, France
alexandre.garcia@lecnam.net

Abstract—

The following article has been submitted to the special issue on Machine Learning in Acoustics in JASA. After it is published, it will be found at <http://asa.scitation.org/journal/jas>.

Sound sources localization using multichannel signal processing has been a subject of active research for decades. In recent years, the use of deep learning in audio signal processing has allowed to drastically improve performances for machine hearing. This has motivated the scientific community to also develop machine learning strategies for source localization applications. In this paper, we present BeamLearning, a multi-resolution deep learning approach that allows to encode relevant information contained in unprocessed time domain acoustic signals captured by microphone arrays. The use of raw data aims at avoiding simplifying hypothesis that most traditional model-based localization methods rely on. Benefits of its use are shown for realtime sound source 2D-localization tasks in reverberating and noisy environments. Since supervised machine learning approaches require large-sized, physically realistic, precisely labelled datasets, we also developed a fast GPU-based computation of room impulse responses using fractional delays for image source models. A thorough analysis of the network representation and extensive performance tests are carried out using the BeamLearning network with synthetic and experimental datasets. Obtained results demonstrate that the BeamLearning approach significantly outperforms the wideband MUSIC and SRP-PHAT methods in terms of localization accuracy and computational efficiency in presence of heavy measurement noise and reverberation.

*Index Terms—*Sound source localization - direction of arrival - deep end-to-end learning - reverberating environments - time domain - atrous convolution

I. INTRODUCTION

Sound sources localization (SSL) is a research field in acoustics that has only recently made use of machine learning approaches. Even if research works on the subject can be found as early as in the 1990s [1] – and sporadically until 2015 [2]–[5] – it is mainly from 2017 onwards, in particular with the development of deep learning, that the scientific community proposed to use advanced machine learning techniques for SSL tasks. The present paper proposes a supervised deep learning

approach for the localization of acoustic sources in aerial environment, using raw, unprocessed time-domain acoustic pressure measurements on a compact array of microphones. Although not aiming at giving a systematic and exhaustive bibliographical review on this subject, this section presents an overview of the methods recently proposed in the scientific literature on this subject in which motivates the development of the BeamLearning approach.

The growing interest in using deep learning techniques for SSL has been recently illustrated by the publication of the work of several authors [6]–[9] in a special issue on acoustic source localization [10]. The LOCATA challenge [11], which brings together a corpus of data recorded by different microphone arrays for different scenarios (single or multiple, moving or static sound sources combinations recorded on static or moving arrays), has also motivated the proposal of the use of a deep neural network [12], which has been shown to outperform both the baseline method and the conventional MUSIC method. While most publications on the subject use supervised learning strategies [6]–[9], [13]–[39] for audible acoustic sources in air and for underwater localization [18], [19], unsupervised or semi-supervised learning approaches have also recently been successfully applied for SSL tasks in reverberant and noisy environments [40], [41]. Among these studies, the majority of authors apply the SSL task to a single active source – this is also the case in this paper – while others evaluate their methods for multiple sources localization [8], [20]–[24] or propose specific deep learning approaches to count the number of active sources [42], [43].

The rise of deep learning algorithms based on convolutional neural networks for machine hearing and acoustics [44] – along with their ability to improve the robustness of direction of arrival (DoA) estimation techniques in presence of noise and reverberation – led to the use of various kinds of inputs to the proposed neural network architectures. Among them, the most common choices are based on time-frequency representations based on short-time Fourier transforms [6], [8], [20], [23]. Using this bidimensional representation, some

authors proposed to exploit the phase information only [20], others used both amplitude and phase [8], [21], and some used the power computed from the STFT matrix [22], [25], [32]. Still, other kinds of input features have been proposed in the literature, such as the GCC features [35], the eigenvectors of the spatial covariance matrix [33], [34], or ambisonics signals [7], [8], [31], [39]. However, there is still no consensus on the best representation to use in order to better encode the information needed to localize sound sources.

Since the unprocessed, time-domain audio signals contain all the information to be extracted, the scientific community has recently put some efforts to directly use the raw waveforms as inputs for deep learning models, either for machine hearing tasks [45]–[47], [47]–[49], [51] or for SSL tasks [24], [36]–[38]. Joint acoustic model learning from the raw waveform has therefore emerged as an active area of research in the last few years, and recent works have shown that this approach, also known as *end-to-end learning*, allows to successfully learn the temporal and spatial dynamics scales of the waveforms. These studies, along with recent advances in machine learning architectures for one-dimensional signals [52]–[54] and the work of the authors on a similar approach for speech and environmental sounds recognition [49] has motivated the present work, which aims at showing the benefit of a deep-learning multi-resolution approach for the SSL task, that allows to avoid the need to pre-process the multichannel waveforms in order to encode the relevant information contained in raw acoustic measurements.

Among the already published studies on SSL using deep learning techniques, the problem is the most commonly treated in a classification framework, where the space is divided into distinct zones. These zones can be angular sectors, and thus represent portions of azimuthal angles [25]–[28], or combine azimuthal and elevation angles to define portions of a sphere [21]. Perotin et al. [7] proposed to use a classification task to also estimate the distance between the source and the microphone array. However, for a SSL task, a regression approach seems as legitimate as classification [55]. When a high angular precision is required, a classification approach may also not be longer sufficient, and a regression approach, giving a continuous numerical value of the source’s angular position [29]–[31] even appears to be more accurate in scenarios with diffuse interference [55]. In a recent study [24], the authors propose to combine the two approaches in order to determine the position of several speakers in a room. The position is first roughly inferred using a classification approach, and then refined using a regression approach. With regard to the BeamLearning network detailed in the present paper, both approaches will be studied in order to compare the observed results.

Machine learning approaches for SSL localization require a large amount of training data. Since signals captured by microphone arrays with different geometries are radically

different, individual experimental data collection is required for each specific type of microphone array. This is why most SSL studies use simulated data, which therefore need to be as realistic as possible in order to adapt to real recording conditions. Some solutions have been proposed in the literature using domain adaptation [56] when the mismatch between training and testing conditions is too high. We therefore also developed an efficient tensor GPU-based computation of synthetic room impulse responses using fractional delays for image source models.

In this article, we present the Beamlearning approach, which relies on the use of a deep neural network (DNN) DoA estimation system, using raw multichannel measurements. The BeamLearning neural network architecture has been specifically developed in order to be insensitive to the microphone array geometry and to the number of sensors involved in the measurement array, and to be compatible with both classification and regression approaches. The obtained results are evaluated for both experimental and synthetic data using speech signals, in anechoic and reverberant environments, with strong noisy conditions. These results are compared on the same data with two state of the art model-based SSL methods (MUSIC and SRP-PHAT). Section II provides details on the method used to build the synthetic datasets, and the BeamLearning neural network architecture is extensively detailed. The different datasets used in this paper are described in Section III, along with the evaluation methods and the training procedure. The ensemble of three multichannel, multiresolution learnable stacked filterbanks, which constitutes the core of the first subnet of the BeamLearning network is analyzed in detail in Section IV. Finally, several SSL experiments are presented and analyzed in Section V, before drawing our conclusions in Section VI.

II. METHODS

A. Synthetic and experimental datasets generation

For both the dataset generation and the neural network architecture, the BeamLearning approach proposed in the present paper has been tailored and extensively tested in order to accommodate to any microphone array configurations in arbitrary reverberant rooms. However, in this paper, for sake of clarity and concision, we chose to illustrate the results using only a particular microphone array. For the experimental dataset, we used the MiniDSP[®] UMA-8 compact microphone array (<https://www.minidsp.com/products/usb-audio-interface/uma-8-microphone-array>). This compact array is composed of $N_c = 7$ digital MEMS microphones. The first one is placed at the center of the circular array, and the 6 others microphones are evenly distributed over the perimeter of a 8 cm diameter circle. This geometry is mostly similar to the microphone arrays used in personal assistants such as Amazon Echo[®], or Apple HomePod[®], where the estimation of a speaker’s direction could represent an important pre-requisite in the processing chain for hands-free communication during voice interaction with these devices in noisy and reverberating

environments.

The experimental datasets have been generated using the UMA-8 microphone array, placed at the "sweet spot" of the laboratory's 5-th order ambisonic playback system [57], which allows a great flexibility and reproducibility for the generation and labelling task of experimental datasets. This sound field synthesis system consists in a 1.07 m radius spherical loudspeaker array following a 50-node Lebedev grid. The loudspeakers are made with a tubular plastic cabinets and 2" drivers (AuraSound[®] NSW2-380-8A). Datasets of 72000 sources positions (80% for learning, 10% evaluation and 10% for testing) have been encoded in the ambisonic domain in order to drive the 50 loudspeakers [58]. These sources positions are randomly generated using the procedure detailed in III-A.

In this paper, the BeamLearning network has also been trained and tested using synthetic datasets corresponding to different room configurations. For these datasets, the microphone array geometry matches the one used in the experimental datasets. Since there is a need to model accurately a large number of sound sources positions and signals in various simulated environments, we specifically developed for this purpose a GPU-based computation of synthetic room impulse responses (RIRs) using fractional delays [59]. This computation of realistic RIRs relies on the use of the image source method [60]. In this process, all the reflections on the room's boundaries are computed, for any time of arrival on the array smaller than the reverberation time (RT) of the room. As an example, for a typical classroom size of $7 \times 10 \times 3.70$ m with a RT of 0.5 s, the whole number of image sources requires a high order of reflections (up to the 60th order), and represents more than 80000 image sources for each of RIR of the synthetic dataset. For each simulated environment, 48000 primary source positions are used in the RIR computations. These sources positions are randomly generated using the procedure detailed in III-A. Using the UMA-8 microphone array, the dataset therefore corresponds to the computation of 336000 RIRs for each room. For each of these RIRs, the whole number of image sources positions and corresponding attenuations that contribute to the acoustic field for a time interval of $[0; T_R]$ are computed using the Pyroomacoustics library [61]. For such a large number of RIRs and image sources, the existing frameworks did not allow to perform the RIR computation efficiently and accurately enough. This is the reason why we developed for this specific application a fast parallel batch RIR computation performed on the GPU using the Tensorflow APIs [62]. This implementation is achieved using sparse tensors computations and an efficient fractional delay filtering implementation [59]. For a compact microphone such as the UMA-8, there is indeed a critical need to keep the precision of the time delays corresponding to the distance between each of the image sources involved in the RIR computation and the microphones. In order to implement the RIR with non integer sample delays, we therefore used a 7th

order Lagrange interpolation [63], that allows to perform a better accuracy than standard truncated sinc interpolations, with a much lower computation cost. The integer part of the sample delay is represented as a sparse array, and the 8 coefficients of each fractional delay filters along with the individual dampings corresponding to the cumulative effect of each wall of the room and the distance between the image source and the microphones are then used to compute the RIRs using every source image contributions [59]. Using this large number of RIRs, the complete dataset is built using batch convolutions with real-life recordings for each simulated environments.

B. Proposed neural network architecture

The proposed BeamLearning neural network is composed of three subnetworks, which are further detailed in the following subsections. This DNN architecture is fed with raw multichannel audio data measured by a microphone array (see Fig. 1), without any further preprocessing. This allows to perform a joint feature learning along with the SSL task, and to avoid manual input features choices that may not be well adapted to particular microphone array geometries. We also specifically developed this architecture in order to allow real-time SSL inference after convergence of the training process. This led to the use of input frames of length $N_t = 1024$ samples. Since most conventional beamforming methods can be thought as filter and sum approaches [64], the global neural architecture depicted on Fig. 1 has been developed in a similar way. The DNN architecture is mainly based on the use of a learnable filterbank, with specific operations allowing the neural network to be expressive enough to achieve the SSL task in reverberant and noisy environments using raw multichannel audio input data of dimension $N_c \times N_t$, with $N_c = 7$ for the particular array used in this paper.

The first subnetwork can be thought as a set of finite impulse-response (FIR) filters with learnable coefficients, that self-adapts to the multichannel audio input data for the SSL task. The convolutional neural cells, which are widely used in DNN architectures, can indeed be thought as a strict equivalent to learnable finite impulse response (FIR) filters in standard digital signal processing (DSP). For this purpose, we propose the use of residual networks of one-dimensional depthwise separable atrous convolutions, which allow to operate in a computationally efficient way, with a similar strategy to the one used in the FrameNet subnet that we proposed in a previous published study for sound recognition [49]. A depthwise separable convolution consists in a convolution performed independently over each channel of an input, followed by a pointwise convolution, *i.e.* a 1x1 convolution, projecting the channels output by the depthwise convolution onto a new channel space. This process allows to operate a convolution on data, faster, with much less parameters than standard convolutions [50]. This residual subnetwork of depthwise separable atrous convolutions is

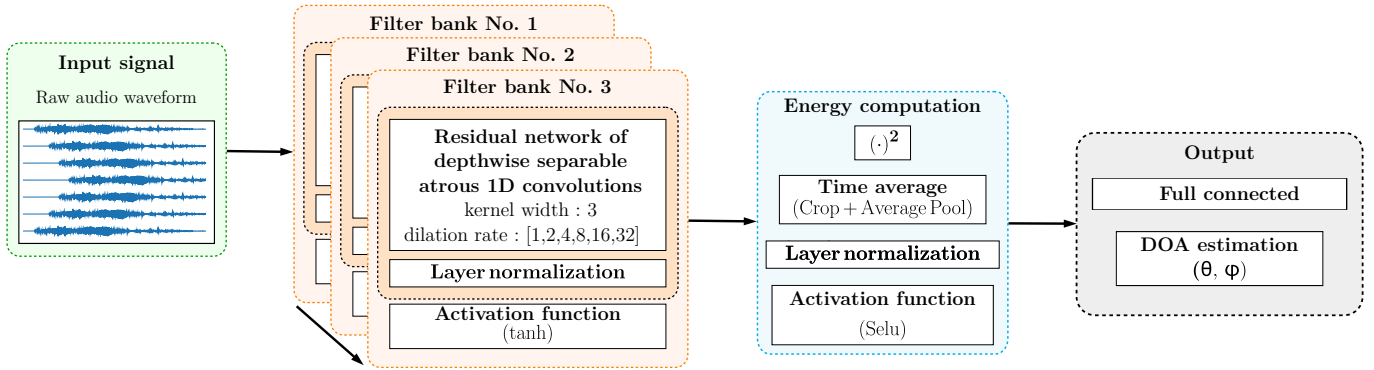


Fig. 1. Schematical representation of the global architecture of the BeamLearning network. This neural network takes a raw multichannel waveform measured by a microphone array, and outputs the DOA estimation of the emitting source.

described in detail in II-B1. The overall output of these three filterbanks – which are composed of 3×768 cascaded learnable filters, each of them having a channel multiplier of 4 – is a two-dimensional map, where the first dimension is N_t , and the second dimension represents the N_f outputs of the cascaded learnable filters. In the present paper, for a 2D-DoA finding task, $N_f = 128$. This hyperparameter value has been determined after extensive preliminary training and testing phases for the proposed DOA task, in order to optimize the neural network performances while limiting the impact on the computational cost on a single GPU. The obtained data representation is then fed to the second subnetwork depicted on Fig. 1, which aims at computing an energy-like representation in N_f dimensions. This subnet is inspired by Beamforming approaches, where the sound source position is determined using a maximization of the Steered Response Power (SRP) of a beam former [64]. The last part of the BeamLearning network aggregates the energy of the N_f channels using a full connected subnetwork in order to infer the sound source position. In the following, the architecture of each of these sub-networks is presented and discussed in relation to the physics of the associated SSL problem.

1) Learnable filterbanks : raw waveform processing:

As introduced in the previous subsection, from a machine-learning point of view, the first subnetwork in the BeamLearning DNN consists of three stacked modules, which are residual networks of one-dimensional depthwise separable atrous convolutions. The proposed filterbank module architecture is detailed on Fig. 2. The use of convolutional layers is a common operation in deep learning for audio applications, that allows to extract features from the data. From a signal processing point of view, the kernel of the convolutional cells can be seen as the Finite Impulse Response (FIR) of a learnable filter. However, acoustic model learning from the raw waveform using dense kernels can lead to large FIR sizes in order to be efficient at filtering at low frequency [49], [51]. As a consequence, rather than making convolutions with large and dense kernels, the proposed filterbank is implemented using one-dimensional atrous

convolutional kernels, which have a long history in classical signal processing, and have originally been developed for the efficient computation of the undecimated wavelet transform [65]. The use of this kind of convolution with upsampled filters has been revisited in the context of Deep Learning and have already been shown as an efficient architecture for audio generation [52], denoising [54], neural machine translation [53], and word recognition [49]. Similarly to previous published studies [49], [52], [54], we use dilation rates which are multiplied by a factor of two for each successive layers and very short kernel widths of 3 samples (early experiments with larger filters of length 5, 7 and 9 showed inferior performances). This allows to achieve a large receptive field (127 samples for a single residual subnetwork of depthwise separable atrous convolutions, and a total receptive field of 379 samples for the three stacked learnable filterbanks), with only 3×6 sets of one-dimensional convolutions with kernels of size 1×3 , corresponding to dilation rates ranging from 1 to 32.

The proposed network topology has been developed in order to optimize the SSL performances in the frequency range of [100 Hz ; 4000 Hz]. For such a large frequency range, using raw audio data, it is useful to offer different timescales of analysis in order to extract the relevant information for the SSL task. The use of three stacked residual atrous convolutional filterbanks depicted on Fig. 1 and Fig. 2 – which share the same architecture but do not operate on the same data, being in cascade – allow the network to operate on multiple time scales in the range of [68 μ s ; 8.6 ms]. The corresponding filter impulse response lengths can thus take values of one period for frequencies ranging from 115 Hz to 14.7 kHz without impacting the computational efficiency. Since the frequency response of a FIR filter is highly dependent on its impulse response length, the proposed network topology – along with the skip connections offered by the residual network – allows for a wide range of filtering possibilities in the targeted frequency range.

In our approach, we use non-causal depthwise separable

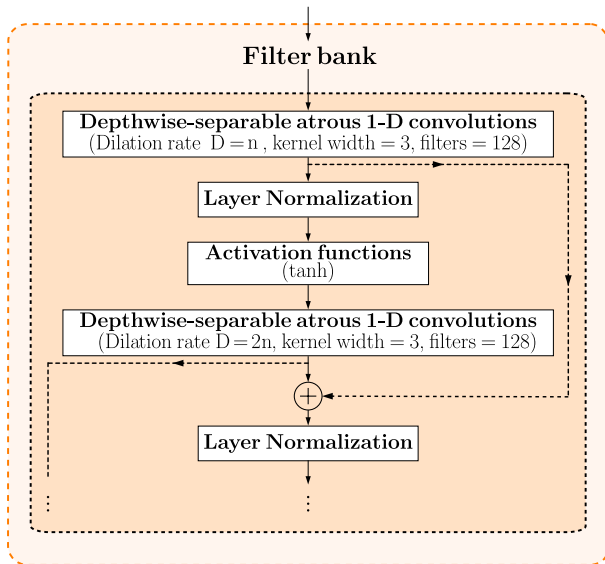


Fig. 2. Inner architecture of a learnable filterbank module, showing the residual connections (dotted lines) between the successive layers. Each layer corresponds to a different dilation rate D , taking values 1,2,4,8,16, and 32. In each depthwise-separable convolutional layer, the depthwise convolution is achieved using a channel multiplier $m = 4$, and a pointwise convolution mapping the outputs of the depthwise convolution into $N_f = 128$ pooled channels.

convolutions, which present the considerable advantage of making a much more efficient use of the parameters available for representation learning than standard convolutions [53]. The convolutions are performed independently over channels (depthwise separable convolutions), with a channel multiplier m . These computed depthwise convolutions are then projected onto a new channel space of dimension N_f for each layer using a pointwise convolution, which is a type of convolution that uses a kernel of size 1×1 : this kernel iterates through every single time sample of the input tensor. This kernel has a depth of however many channels its input tensor has. From a signal processing point of view, this approach aims at pooling together the contents in the filtered multichannel soundwave that share similar spatial and spectral features, in order to ease the sound localization task: the pointwise convolution therefore helps combining the pooled channels in order to enhance the expressivity of the network.

As shown on Fig. 1, three of these filterbank modules are stacked, and residual connections are added between each layers of the three successive filterbanks (see Fig. 2), in order to allow shortcut connections between layers and offer increased representation power by circumventing some of the learning difficulties introduced by deep layers [66]. These skip connections allow the information to flow across the layers easier by bypassing the activations from one layer to the next. This limits the probability of saturation or deterioration phenomenons of the learning process, both for forward and backward computations in deep neural networks [66]–[68].

As depicted on Fig. 2, each atrous convolutional layer is followed by a layer-normalization layer (LN) [69], which allows to compute layer-wise statistics and to normalize the nonlinear tanh activation across all summed inputs within the layer, instead of within the batch with a standard batch-normalization [70], [71] process. The layer-normalization approach has the great advantage of being insensitive to the mini-batchsize [69], and to be usable even if the batch dimension is reduced to one, when the DNN is used for realtime inference. Each LN layer is then followed in this module by a tanh nonlinear activation function, which was selected after extensive comparison with other standard activation functions for the SSL task using the BeamLearning architecture.

2) Pseudo-energy computation:

In the following, the set of N_f time-domain outputs $s^{(i)}[n]$ of the filterbank subnetwork will be denoted as $\mathbf{S}_{i,n}$ - where i stands for the filter channel index, and n for the time sample - the bold notation signifying that this a two-dimensional tensor. $\mathbf{S}_{i,n}$ is fed to the second subnetwork of the BeamLearning DNN, which allows to compute a pseudo-energy for each of the N_f pooled channels. As shown in Fig. 1, from a machine learning point of view, $\mathbf{S}_{i,n}$ is squared, and cropped in the time dimension - therefore only keeping the time frames corresponding to the valid part for all the atrous convolutional layers used in the filterbanks subnetwork - and undergoes an averagepool operation. This pooling operation acts as a dimensionality reduction and has been preferred to standard maxpool operation, since the proposed pseudo-energy can be seen as an equivalent to a mean quadratic pressure, which is similar to the value that is maximized by standard model-based Beamforming algorithms [64]. The output of this deterministic pseudo-energy computation is then fed to a LN layer [69] in order to normalize the Selu [72] nonlinear activations, which in this case helps to speed up the convergence process.

3) Matching the pseudo-energy features to the source angular output:

The last layer of the BeamLearning network is a standard full connected layer, which aims at combining the previously computed pseudo energy features in order to infer the source DoA. In our experiments, the BeamLearning approach is evaluated for 2D DoA determination tasks using either a classification framework or a regression framework. For the classification problem experiments, the BeamLearning DNN output is a n -dimensional vector representing the probabilities of the source belonging to n different angular partitions. For a regression problem, and in agreement with Adavanne et al. [8] proposal, the output corresponds to the projection of the source position on the unit circle (in 2D) or on the unit sphere (in 3D). This output vector therefore corresponds to the unit vector pointing towards the source DOA. Because of the angular periodicity of the DoA determination problem,

this allows an easier retro-propagation process of errors from these projections, than directly from the error of a periodic angle. The angular output DoA θ (and if needed ϕ in 3D) is then deduced from this projected output.

4) Training procedure:

Using a classification framework, the BeamLearning DNN is trained with one-hot encoded labels corresponding to the angular area where the source is located, therefore allowing to compute the cross-entropy loss function between estimated labels and ground truth labels. For the regression problem experiments, the BeamLearning DNN is trained with labels corresponding to the projection of the source position on the unit circle, therefore allowing to compute a standard L2-norm loss function between estimated labels and ground truth labels, which represents a geometrical distance between the prediction and the true DOA using cartesian coordinates.

The DNN variables updates and the back propagation of errors through the neural network are optimized using the Adaptive Moment Estimation (Adam [73]) algorithm, which performs an exponential moving average of the gradient and the squared gradient, and allows to control the decay rates of these moving averages. In addition to the natural decay of the learning rate that Adam performs during the learning process, we set an exponential learning rate decay. The learning rate λ for each step is set to decrease exponentially at each learning iteration k , from a maximum value $\lambda_{Max} = 10^{-3}$ to a minimum value $\lambda_{min} = 10^{-6}$:

$$\lambda = \lambda_{min} + (\lambda_{Max} - \lambda_{min}) * e^{-\frac{k}{N_{iteration}}}$$

The models have been implemented and tested using the Python Tensorflow [62] open source software library, and computations were carried out on a Nvidia[®] GTX 1080Ti GPU card, using mini-batches of 100 raw multichannel waveforms for each training steps. All the network variables were initialized using a centered truncated normal distribution with a standard deviation of 0.1. Using the proposed datasets, the convergence of the training process requires approximately 150 epochs using a Nvidia[®] GTX 1080Ti GPU card.

III. EVALUATION

A. Datasets

In the present paper, we evaluate the performances of the proposed BeamLearning DNN for a 2D DoA SSL task, using experimental and synthetic datasets. Both dataset generation methods have been detailed in subsection II-A. In the present section, the source positions around the microphone array used in the datasets are detailed, and the different propagating environments are also presented.

In the following, the source positions are denoted using spherical coordinates (R, θ, ϕ) centered on the microphone array, where R is the distance to the central microphone, and

(θ, ϕ) are the azimuthal and elevation angles defined in Fig. 3. In the present study, since we evaluate the BeamLearning approach for a 2D DoA determination task, the θ coordinate is the only DNN global output. However, in order to increase the robustness to spatial variations, the source positions are randomly picked from a uniform distribution in a torus of section $R \times (2\Delta R) \times \sin(2\Delta\phi)$ and of central radius R , as depicted in Fig. 3.

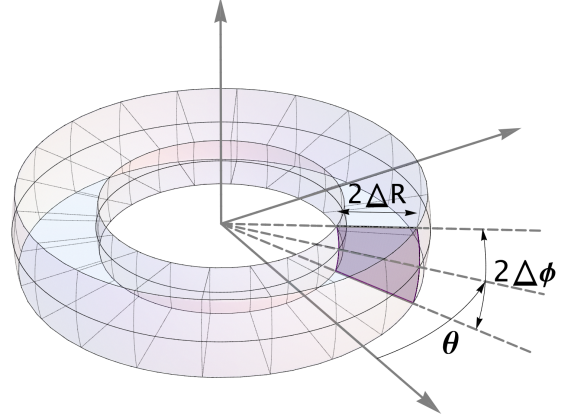


Fig. 3. Sound sources positioning strategy around the microphone array for the generation of the experimental and synthetic datasets.

For each synthetic dataset, a set of 48000 sources positions is randomly generated using this procedure. For the experimental dataset, a set of 72000 sources positions has also been randomly generated using the same geometry. The datasets are split into training, validation and test sets in the standard ratios of 80:10:10.

In this following, the BeamLearning DNN is trained and tested using synthetic data in two different reverberating environments, and using experimental data in a real room. The first synthetic dataset is generated in a simulated room corresponds to a typical classroom size of $7 \times 10 \times 3.7$ m and a reverberation time of 0.5 s. In this case, the microphone array is positioned at coordinates $[4, 6, 1.5]$ m, the origin of the Cartesian coordinate system corresponding to a corner of the room. The second simulated room exhibits an almost cubic shape, with dimensions $5 \times 5 \times 4$ m, and a reverberation time $T_r = 0.5$ s. For both simulated rooms, the room boundaries exhibit homogeneous absorbing coefficients computed using the method proposed by Lehmann et al. [74], [75] in order to ensure a simulated reverberation time of 0.5 s. In order to also study the localization performances offered by the proposed BeamLearning approach in presence of measurement noise independantly of reverberation, a free field dataset has also been generated for this particular experiment.

For the synthetic datasets, a dataset of multichannel RIRs is first computed for each propagating environment using the method detailed in II-A. These sets of multichannel RIRs are then convoluted with several kinds of signals

emitted from the position of the simulated sound sources. This common method [7], [21], [37] allows a great flexibility in terms of emitted signals in a memory-efficient way. In the present study, the signals used to generate the synthetic datasets using the computed RIRs are composed of anechoic recordings of symphonic music [76] (<https://users.aalto.fi/~ktlokki/Sinfrec/sinfrec.html>), outside recordings of multiple women talking in Danish (<https://odeon.dk/downloads/anechoic-recordings/>), and car horn honks from the UrbanSound8K dataset [77].

During the training phase, each multichannel recordings of the training mini batches are added to a random multichannel Gaussian noise with randomly picked signal-to-noise ratio (SNR) values in the range of $[x, +\infty[$ dB. This process can be interpreted as a data augmentation strategy, which allows to ease the generalization capabilities of the proposed method during the learning process. In the following, the presented results will include data augmentation strategies during the training phase with $x = 5$ dB (section V), $x = 20$ dB (section IV), $x = 25$ dB (section V-A), and $x = +\infty$ dB (without any data augmentation, section V-A). The proposed data augmentation strategy will be shown to offer an increase of robustness to measurement noise in comparison to state of the art SSL algorithms.

In order to evaluate the 2D-DoA SSL performances of the BeamLearning approach and to compare it in a reproducible way to model-based algorithms (MUSIC [78] and SRP-PHAT [79]), each method is tested in section V using the exact same testing datasets. These testing datasets consist in 3600 sources, randomly distributed in the torus depicted on Fig 3, emitting unseen signals during the training phase. In order to evaluate the robustness of each methods to noisy measurements in reverberating environments, the statistics of the DoA mismatch are derived from these 3600 positions, for each proposed propagating environments, and for different deterministic (SNR) values, ranging from +40 dB to -1 dB.

B. Evaluation metrics

In the present paper, the BeamLearning approach is evaluated both in a classification framework and in a regression framework. As a consequence, in order to analyze precisely the performances of the proposed DNN for the task of 2D-DoA SSL, several evaluation metrics will be used in the following. For the task of supervised multi-class angular classification, the overall accuracy is computed during the training phase and for testing using the number of correctly recognized class examples (true positives, t_{p_i}), the number of correctly recognized examples that do not belong to the class (true negatives, t_{n_i}), and examples that either were incorrectly assigned to the class (false positives, f_{p_i}) or that were not recognized as class examples (false negatives, f_{n_i}) [80]. Other angular mismatch estimators will also be detailed in section IV-C. For the task of supervised angular regression, the DoA

mismatch between the estimated angular positions $\theta(\tilde{k})$ and the groundtruth angular positions $\theta(k)$ is evaluated for each sources using statistics derived from the absolute angular mismatch $|\theta(\tilde{k}) - \theta(k)|$ over all the sources in the testing dataset.

IV. BEAMLEARNING NETWORK ANALYSIS

A. Learnable filterbanks inner working analysis

In this subsection, we analyse the variables learnt in the first subnetwork described in II-B1, in order to give further insight on the learning process involved. The architecture of this first subnetwork of the BeamLearning DNN has been specifically developed to automatically build a 2D map $\mathbf{S}_{i,n}$ that can be interpreted as the filtering of raw input multichannel measurements by 128 different tunable filters. On contrary to conventional Beamforming, these 128 time domain outputs cannot be interpreted as beams, but rather as 128 optimized channels which aim at achieving a joint filtering in space and time. Since this subnetwork is composed of 3 filterbanks constituted by 6 depthwise separable atrous convolutions with a channel multiplier of 4 and a number of output channels of 128, it is important to note that the global output of this first subnet results from the combination of nonlinear filtering achieved by the equivalent $3 \times 6 \times 4 \times 128 = 9216$ learnable filters followed by mathematical operations operated by the LN layer, and their nonlinear activation functions. Rather than looking at each of these filters independantly, this analysis aims at giving insightful representations of the spatio-temporal filtering achieved when the multichannel input data has partially passed through the n first layers of the learnable filterbank subnetwork. This analysis will allow to better understand the multiresolution offered by the residual networks of one-dimensional depthwise separable atrous convolutions that constitute the main ingredient of this subnetwork.

In order to perform this analysis, the BeamLearning network is trained for a regression problem using the synthetic dataset described in III-A. This training dataset, will be referred in the following as $\mathcal{D}_{\text{train}}^{\text{Room}^1}$. In this dataset, each of the sources located in a room of dimensions $7 \times 10 \times 3.7$ m and a reverberation time of 0.5 s. emit ambient speech noise signals, car horn signals, and symphonic music. The data augmentation procedure detailed with added measurement noise detailed in III-A is used with a $SNR \geq 20$ dB. Following the procedure detailed in III-A, the 48000 sources positions included in $\mathcal{D}_{\text{train}}^{\text{Room}^1}$ are randomly distributed in a torus of section $R \times (2\Delta R) \times \sin(2\Delta\phi)$ and of central radius $R = 2$ m, and with $\Delta R = 0.5$ m and $\Delta\phi = 7$ degrees (see Fig 3). All the parameters of the BeamLearning network are then frozen after convergence of the learning process. This frozen BeamLearning network is then used to build 2 dimensional (θ, f) magnitude maps of the 128×18 spatio-temporal filterings offered at the output of the 18 successive depthwise-separable atrous convolutional layers used in the learnable filterbank subnetwork. Among these




Fig. 4. Directivity maps (in dB) of one of the 128 equivalent filters obtained by passing through the first layers of the network, for the 1st (a), 4th (b), and 6th (c) depthwise-separable atrous convolutional layers of the filterbank subnetwork. These 3 layers correspond respectively to a dilatation factor of 1, 8 and 32.

2304 possible (θ, f) transfer function magnitude maps offered at increasing depths of the subnetwork, Fig. 4 shows three of them. These three maps are selected from the first filterbank, and corresponding to dilatation factors of 1, 8, and 32 in order to illustrate the multiresolution capabilities offered by the BeamLearning approach.

The representations in Fig. 4 are obtained using unseen input data during the DNN training. These unseen testing input data correspond to raw measurements on the microphone array using 360 sources uniformly sampled on a circle of radius 2 m at $\phi = \pi/2$, emitting monochromatic signals ranging from 100 Hz to 4000 Hz. This allows to compute the magnitude of the filtering process for each frequency and for each azimuthal angle at any of the layer composing the learnable filterbank subnetwork (including nonlinear activations and normalization layers), which is a strict equivalent of representing the directivity maps obtained at different depths of the subnetwork.

When analysing the obtained directivity maps on Fig. 4, one can observe that the frequency and spatial selectivity of the learnt filters increase with the depth through which the input data passes into the DNN. For the very first layer of the filterbank subnetwork corresponding to Fig. 4a, the obtained directivity obtained for this particular filter corresponds to a dipolar response above 1000 Hz, with directivity notches at $\theta = 150$ degree and $\theta = 330$ degree. At low frequencies, this particular filter (out of the 128 possible filters for this layer) exhibits a global attenuation of -20 dB, without any strong directivity pattern. It is also interesting to note that among the 128 filters obtained for this particular first convolutional layer (data not shown), most filters exhibit such a dipolar angular response at high frequencies, with various spatial notches values. This behaviour can be explained by the fact that for this layer, the filter kernels have a very small size of 3 samples, with a dilatation factor of 1.

However, when the data passes through more depthwise-separable atrous convolutional layers in the filterbank (Fig. 4b and Fig. 4c), the combination of optimized filters begin to be more and more selective, both in the frequency domain and in

the angular domain. This is even observed at low frequencies, thanks to the kernel widening offered by the increased dilatation factors and the combination of filters at previous layers offered by the residual connections. This observation suggests that the stacking of the layers composing the three successive filterbanks converge to a global nonlinear filtering process into 128 channels, where each output channel specializes in several frequency bands and several angle incidences. When the source signal exhibits a particular spectrum in a reverberating environment, the energies carried by these 128 channels are then combined by the full connected layer in order to infer the source position.

B. BeamLearning output analysis for a classification problem

Most published studies on SSL using a deep learning approach treat the localization problem as an angular classification problem, where the objective is to determine to which discrete portion of the angular space the emitting source belongs to [7], [20], [23], [33]. When seeking a better angular resolution for the SSL problem, these approaches require to increase the number of classes.

In this n-class classification approach, the source angular position $\theta \in [-180, 180]$ is first converted to an integer i representing the belonging of the source to the i^{th} angular partition among the n possibilities :

$$i = \lfloor (\theta + 180) \frac{n}{360} \rfloor$$

Using this approach, θ is then approximated to a corresponding discrete θ_i (the center of the i^{th} angular partition), which is equivalent at discretizing the space on an angular grid :

$$\theta_i = (i + \frac{1}{2}) \times \frac{360}{n} - 180$$

The label used for classification can then be converted using a one-hot encoded target strategy using the integer i [20], [55]. Some authors also proposed the use of a soft Gibbs target [23], [55] exploiting the absolute angular distance between the true angular position θ and the attributed discrete value θ_i on the grid.

In the following, we study the behaviour of the DNN output neurons for a 8-class classification problem. The analysis of the BeamLearning DNN outputs for a regression problem will be carried out similarly in IV-D. For the classification problem, the BeamLearning DNN is trained with a cross-entropy loss function preceded by a softmax activation function and with one-hot encoded labels. Using the same training dataset $\mathcal{D}_{\text{train}}^{\text{Room}1}$ as in IV-A with data augmentation procedure with $SNR \geq 20$ dB and after convergence of the BeamLearning network, the frozen DNN is fed with unseen input data during the DNN training. These unseen testing inputs corresponds to 4800 sources randomly distributed over a torus of mean radius $R = 2$ m, emitting unseen speech signals. This testing dataset, which will be used several times throughout the present paper, will be referred in the following as $\mathcal{D}_{\text{test speech}}^{\text{Room}1}$. The use of such a testing dataset allows to represent quantitatively the probability value obtained using the softmax function applied to the estimated multiclass label for each of the 8 directions as a polar representation, which can be interpreted as a directivity diagram (see Fig.5).




Fig. 5. Directivity diagrams of 3 of the 8 output neurons for the SSL task seen as a classification problem. The diagrams are obtained by plotting the output probability (solid blue) obtained for each output neuron when testing the BeamLearning network with the 4800 sources positions in $\mathcal{D}_{\text{test speech}}^{\text{Room}1}$. The sources positions where the predicted output is the i^{th} angular partition are depicted using light blue shaded areas.

The analysis of the obtained directivity diagrams on Fig. 5 allows to highlight the fact that during learning, each output neuron specializes to its corresponding angular partition, with a sharp directivity pattern. However, a precise analysis of these directivity pattern shows that some of these directivity patterns slightly overlap adjacent angular partitions that define the angular classes used for the SSL problem. As a consequence, a loss of accuracy can be observed for sound sources located in the vicinity of angular partitions. This also clearly shows that the misclassified sources mostly belong to contiguous angular sectors to the estimated angular sector. This suggest that for a rather coarse angular grid, the use of a soft Gibbs target or a Gibbs-weighted cross-entropy loss such as proposed in [23], [55] may only offer small accuracy improvements for the proposed BeamLearning approach over the simple one-hot encoded labelling strategy we use. However, for finer angular grids and with a higher number of angular partitions, the number of boundaries between classes increase, which could lead to a loss of global classification accuracy, which may

justify the use of a regression approach.

C. Influence of the number of angular classes

Since an angular classification problem is directly linked to the choice of a discrete angular grid to define angular partitions, the present subsection aims at giving further insight on the influence of the number of classes when aiming at obtaining a better angular resolution for the SSL problem. As observed in the previous section, for a small number of classes, most misclassified source positions are located in the vicinity of the chosen angular partitions. In order to verify if this assumption still stands for a high count of classes, we conducted several n-class classification experiments, with n ranging from 8 to 128. Those experiments have been conducted both with the experimental dataset \mathcal{D}^{exp} detailed in III-A with $\Delta R = 0.5$ m and $\Delta \phi = 7^\circ$ partitionned in $\mathcal{D}_{\text{train}}^{\text{exp}}$ and $\mathcal{D}_{\text{test}}^{\text{exp}}$, and the same training and testing synthetic datasets $\mathcal{D}_{\text{train}}^{\text{Room}1}$ and $\mathcal{D}_{\text{test speech}}^{\text{Room}1}$ than those used in IV-A. Though, it is interesting to note that for the experimental dataset \mathcal{D}^{exp} , the BeamLearning network has been trained using 72000 acoustic sources generated using the laboratory's higher order ambisonics sound field synthesis system [57], each source emitting a simple set of 6 pure sinusoidal signals, whose frequencies correspond to the central frequencies of the octave bands between 125 Hz and 4000 Hz. Since the sound field synthesis system is located in a rather moderately reverberating environment ($T_R \approx 0.2$ s) and since the frequency content and the dynamics of those emitted signals are much less diverse than in $\mathcal{D}_{\text{train}}^{\text{Room}1}$, this represents an easier learning task than the one achieved using the synthetic dataset.

For both training datasets and for each number of classes n , after convergence of the learning process, the frozen Beamlearning DNNs are fed with unseen input data during the training phase. These unseen testing inputs corresponds to 4800 sources randomly distributed over a torus of mean radius $R = 2$ m around the microphone array emitting unseen signals during the DNN training phase. Table I presents the obtained performances in terms of multi-class accuracy and in terms of angular accuracy. Since the SSL determination is treated here as a classification problem, two kinds of angular accuracies can be calculated, which both carry some insightful information about the localization performances of the algorithm :

$$\Delta \theta_{\text{class}} = \overline{|\tilde{\theta}_i - \theta_i|}$$

$$\Delta \theta = \overline{|\tilde{\theta}_i - \theta|}$$

, where θ is the set of continuous ground truth positions of the 4800 testing sources, θ_i is the set of corresponding groundtruth grid positions defined in IV-B, and $\tilde{\theta}_i$ is the set of discrete grid positions estimated by the BeamLearning network.

The analysis of the results in Table I allow to show that even with a rather low classification accuracy for a high

TABLE I

MULTICLASS CLASSIFICATION PERFORMANCES (ACCURACY, AND MEAN ABSOLUTE ANGULAR ERRORS) WITH INCREASING NUMBER n OF CLASSES. RESULTS ARE PRESENTED BOTH FOR AN EXPERIMENTAL DATASET AND A SYNTHETIC DATASET IN A REVERBERATING ENVIRONMENT, WITH A $SNR = 20$ DB. THE TESTING DATASET IS COMPOSED OF 4800 POSITIONS RANDOMLY DISTRIBUTED AROUND THE MICROPHONE ARRAY. RESULTS WITH BOTH $\Delta\theta \leq 5^\circ$ AND $\frac{\Delta\theta}{\alpha} \leq 1$ ARE IN BOLD.

n		8	16	32	64	128
Half angular width α		22.5°	11.25°	5.63°	2.81°	1.4°
Experim. dataset	Accuracy	99.2%	97.1%	94.3%	91.3%	87.7%
	$\Delta\theta_{class}$	0.3°	0.6°	0.6°	0.5°	0.4°
	$\Delta\theta$	11.1°	5.8°	3.0°	1.5°	0.9°
Synthetic dataset	Accuracy	92.1%	87.6%	79.5%	61.4%	37.9%
	$\Delta\theta_{class}$	5.9°	5.1°	3.0°	4.1°	4.7°
	$\Delta\theta$	14.9°	8.8°	4.3°	4.5°	4.8°

count of angular classes, the mean absolute angular errors on the grid $\Delta\theta_{class}$ remain very low and almost constant across the number of classes. This suggests that for a high count of angular classes, the growing number misclassified sources are attributed to grid points that remain close to their ground truth positions. Since the angular grid refines when n increases, the observed decrease of classification accuracy is therefore compensated by the increasing of the angular resolution offered by a higher count of angular classes. However, there is a clear influence of this grid resolution when estimating $\Delta\theta$, which also includes the uncertainty due to the grid resolution. Indeed, the computation of $\Delta\theta$ also includes the absolute angular error for sources that are classified in the right angular partition. One can observe that for the synthetic dataset in a reverberating environment, $\Delta\theta$ remains smaller than the half angular width α of the angular partitions for $n \leq 32$. However, for $n = 64$ and $n = 128$, the mean absolute angular error $\Delta\theta$ exceed α , which signifies that for such a angular resolution, a regression approach could be more appropriate.

It is also interesting to study the spatial repartition of the misclassified sources. Table II presents two insightful statistics over the population Ω_m of misclassified sources. P_{adj} represents the percentage of Ω_m elements whose groundtruth position is located in an contiguous angular partition to the estimated angular partition. β represents the mean absolute angular distance between the ground truth position of misclassified sources and the boundary of the mis-estimated angular partition.

The analysis of these metrics – and in particular of P_{adj} – allows to conclude that even with a high count of classes and a refined discrete grid of possible angles, the great majority of misclassified sources mostly belong to contiguous angular sectors to the estimated angular sector.

TABLE II

STATISTICS OVER THE MISCLASSIFIED POPULATIONS Ω_m WITH INCREASING NUMBER OF CLASSES n (PERCENTAGE OF MISCLASSIFIED SOURCES, AND MEAN ABSOLUTE ANGULAR ERRORS) WITH INCREASING NUMBER n OF CLASSES. RESULTS ARE PRESENTED BOTH FOR AN EXPERIMENTAL DATASET AND A SYNTHETIC DATASET IN A REVERBERATING ENVIRONMENT, WITH A $SNR = 20$ DB. THE TESTING DATASET IS COMPOSED OF 4800 POSITIONS RANDOMLY DISTRIBUTED AROUND THE MICROPHONE ARRAY. RESULTS WITH BOTH $P_{adj} \geq 85\%$ AND $\frac{\beta}{\alpha} \leq 1$ ARE IN BOLD.

n		8	16	32	64	128
Half angular width α		22.5°	11.25°	5.63°	2.81°	1.4°
Experim. dataset	Card(Ω_m)	34	139	273	417	591
	P_{adj}	100%	100%	100%	100%	100%
	β	0.01°	0.03°	0.09°	0.17°	0.37°
Synthetic dataset	Card(Ω_m)	379	595	984	1853	2981
	P_{adj}	88.9%	82.2%	94.6%	88.1%	57.8%
	β	1.8°	2.0°	1.0°	2.4°	3.0°

The only situation where misclassified sources begin to spread on non contiguous angular partitions corresponds to a 128-class classification problem in a reverberating and noisy environment. Using the proposed BeamLearning network, this observation again advocates for the limited interest of the use of a proximity-aware Gibbs weighing strategy proposed in [23], [55] for labelling the sound sources positions. The proposed BeamLearning network is only trained for classification with one-hot encoded labels that do not take into account any proximity weighting. This result suggests that the obtained representation at the output of the learnable filterbanks and the pseudo-energy network already encodes efficiently this information.

Furthermore, the analysis of β shows that even with a high count of classes, the misclassified sources even remain very close to the boundary of the estimated angular section (less than 3° in all cases). Except for $n = 64$ and $n = 128$ classification problems in a reverberating and noisy environment, the mean distance to the boundary remains smaller than the quarter width of the angular sectors. However, when seeking an angular resolution of less than 5° (corresponding to $n > 32$), these results suggest that a regression approach could be more appropriate and less dependent to the choice of the angular partitions, which confirms the results obtained in [55] for a similar problem in 3D, but with a completely different DNN architecture and different kind of input signals representation.

D. BeamLearning output analysis for a regression problem

In this subsection, the BeamLearning DNN is used for DoA determination, in a regression scenario. As explained in sections II-B3 and II-B4, in this case, the sound source angular position is inferred from the output vector of the network, which represents the estimated cartesian coordinates of the unit vector pointing to the source DoA. The whole

BeamLearning DNN architecture remains unchanged with respect to the classification problem, except for the last layer and the cost function described in section II-B4. In a regression framework though, the DNN output cannot be used to derive directivity patterns as it was in IV-B, since the output is directly linked to the continuous DoA space. However, a polar representation of the absolute angular errors committed by the DNN over a large testing dataset (see Fig. 6) allows to represent the SSL performances of the proposed approach in a compact way.




Fig. 6. Angular diagram of statistical performance metrics derived from the absolute angular mismatches obtained by the BeamLearning DNN using an unseen testing dataset $\mathcal{D}_{\text{test speech}}^{\text{Room1}}$ with a fixed $SNR = 20$ dB. Global median over the whole dataset (dashed black line), local sliding median over an angular window of 5 degrees (solid blue line), and local interquartile range (light blue shaded area).

The polar diagram on Fig. 6 has been obtained by freezing the BeamLearning network after convergence of the training process for a DoA regression task in a reverberating environment, trained with a variable $SNR \geq 20$ dB. The training dataset used for this task is the same synthetic dataset $\mathcal{D}_{\text{train}}^{\text{Room1}}$ as in IV-A, IV-B, and IV-C. The frozen BeamLearning DNN is then fed with input data $\mathcal{D}_{\text{test speech}}^{\text{Room1}}$ (unseen during the training phase), with a fixed $SNR = 20$ dB. For each of the 4800 testing positions, 50 different draws of random noise have been added to the microphone measurements. Using the 4800×50 obtained DoA estimations, statistics of the absolute angular errors are computed, including the global median over $\mathcal{D}_{\text{test speech}}^{\text{Room1}}$ (1.9°), the local median computed using a sliding average over a window of 5 degrees (ranging from 1.5° to 2.2°), and the local interquartile range (ranging from 0.7° to 3.9°) which represents the confidence interval of the DoA mismatch estimations. Since this experiment uses the same training and testing datasets as those used for classification experiments in previous subsections, this allows to compare the obtained performances for this particular 2D DoA task. The obtained results show that the proposed regression approach with the BeamLearning DNN is more accurate than

any of the classification test cases, even with a high count of classes. As a consequence, in the following, all the results and experiments will be conducted using a regression framework.

V. RESULTS AND DISCUSSION USING A REGRESSION APPROACH

In this section, the SSL performances of the BeamLearning approach will be extensively studied and systematically compared with the performances of two state-of-the-art model-based SSL algorithms (MUSIC and SRP-PHAT) for a 2D DoA determination problem. This choice has been made using extensive preliminary comparisons between the TOPS, CSSM, wideband MUSIC, SRP-PHAT, and FRIDA algorithms included in the Pyroomacoustics library [61] and by selecting the most robust methods among them.

SRP-PHAT [79] is a widely used algorithm, which is often considered by the scientific community as one of the ideal strategies for SSL systems [81]–[83]. This algorithm aims at increasing robustness to signal and ambient conditions by combining the robustness of the steered response power [84] (SRP) approach with phase transformation filtering [85] (PHAT).

The MUSIC algorithm [78] and its wideband extensions [86] is also a widely used and tested algorithm [86]–[89] for DoA estimation in noisy and reverberant environments, which also makes it an excellent choice for comparative tests. This algorithm, which is able to handle arbitrary geometries and multiple simultaneous sources, belongs to the family of subspace approaches, and depends on the eigen-decomposition of the covariance matrix of the observation vectors. In the present paper, we used the broadband implementation of the MUSIC algorithm proposed in [61], [86], which is obtained by decomposing the wideband signal into multiple frequency bins, and then applying the narrowband algorithm for each frequency bin in order to get the source direction. In the following, for each methods, frequency content within $[100 \text{ Hz}; 4 \text{ kHz}]$ is used for DOA estimations. This frequency range corresponds approximately to 90 frequency bins for the wideband MUSIC and SRP-PHAT methods.

It is also important to note that the MUSIC and SRP-PHAT methods are both based on an angular grid search. Therefore, we used a 1-degree angular grid resolution, which is smaller than the observed interquartile ranges of the three compared SSL methods. Although this affects the computational cost of the two traditional DOA estimation algorithms, it ensures a fair comparison of DOA estimation performance between BeamLearning, MUSIC, and SRP-PHAT.

In the following, a particular attention will be paid to the robustness to measurement noise in reverberating environments, to changes in the propagation environment or in the characteristics of the emitted signals. The computational efficiency

of the three localization methods will also be compared for realtime 2D-DoA determination.

A. Robustness to noisy measurements

In this subsection, we evaluate the SSL performances of the BeamLearning DNN along with those of the MUSIC and SRP-PHAT algorithms, which both have been developed to accommodate to noisy measurement conditions, ranging from sensor ego-noise and acquisition boards electronic noise, to residual acoustic noise due to other interference sources. As explained in III-A, during the training phase, each multichannel recordings of the training mini batches are added to a random multichannel Gaussian noise with randomly picked signal-to-noise ratio (SNR) values in the range of $[x, +\infty[$ dB. This data augmentation procedure intends at offering an improved robustness to noisy measurements for SSL.

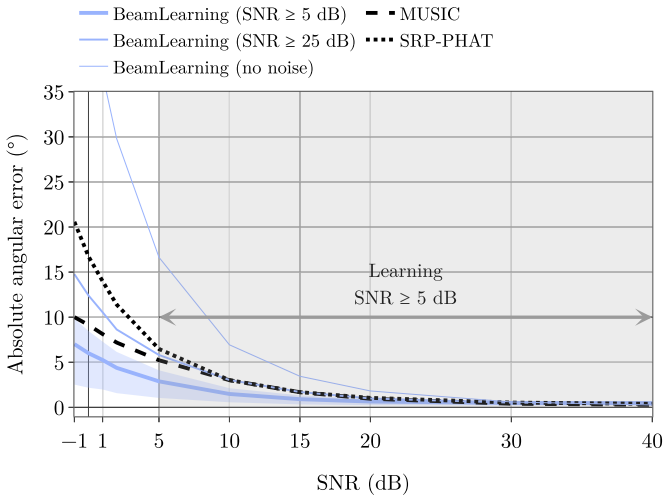


Fig. 7. (Color online) Mean absolute angular mismatch obtained in an anechoic environment for different SNRs, using an unseen speech testing dataset (3600 positions for each SNR). The DoA errors curves are plotted for wideband MUSIC (dashed), SRP-PHAT (dotted), and the proposed BeamLearning approach (solid blue lines). In order to emphasize the influence of the proposed data augmentation strategy, the BeamLearning DNN has been pre-trained three different times using noisy measurements, with random $\text{SNR} \geq 5$ dB (thick blue line), with random $\text{SNR} \geq 25$ dB (medium blue line), and without any noise added (thin blue line). The interquartile range of absolute angular mismatch is also represented for BeamLearning trained with $\text{SNR} \geq 5$ dB (light blue shaded area).

The BeamLearning DNN has been trained and frozen after convergence for two different environments: a perfectly anechoic environment, and the same reverberating room than in previous subsections. In the anechoic environment, the training procedure has been performed with three different minimum SNR values ($x = 5$ dB, $x = 25$ dB, and $x = +\infty$ dB, *i.e.* noiseless data) for data augmentation. In the reverberating environment, results are presented with a data augmentation procedure with $x = 5$ dB. This allows to observe the robustness to measurement noise only (see Fig. 7), and to demonstrate the capabilities of the proposed

approach in both noisy and reverberating environments (see Tab. III).

For both environments, the testing datasets used for the evaluation of the MUSIC and SRP-PHAT algorithms and the frozen BeamLearning DNNs are composed of noisy measurements of 3600 sources positions emitting speech signals at a distance of 2 ± 0.5 m from the microphone array, with a varying SNR ranging from -1 dB to $+\infty$ dB. Both the sources positions and the emitted signals were unseen during the training phases of the BeamLearning DNN. The analysis of the obtained results presented on Fig. 7 and Tab. III shows that the proposed BeamLearning approach is significantly more insensitive to measurement noise than MUSIC and SRP-PHAT, both in a free field situation and in presence of reverberation.

TABLE III
MEAN ABSOLUTE ANGULAR MISMATCH OBTAINED IN A REVERBERATING ENVIRONMENT (RT = 0.5 s), FOR DIFFERENT SNRS, USING AN UNSEEN SPEECH TESTING DATASET (3600 POSITIONS). THE DNN HAS BEEN PRE-TRAINED USING NOISY MEASUREMENTS, WITH RANDOM $\text{SNR} \geq 5$ dB. THE BEST RESULTS ARE IN BOLD.

SNR (dB)	-1	2	5	10	15	20
SRP-PHAT	41.6°	28.7°	17.4°	7.9°	4.6°	3.4°
MUSIC	27.5°	21.3°	17.8°	13.1°	11.0°	10.9°
BeamLearning	21.1°	13.0°	8.2°	4.8°	3.4°	2.7°

A comparison of the results obtained with the BeamLearning approach on Fig. 7 allows to conclude that the proposed data augmentation process offers a very efficient way to improve SSL estimations. Over all the proposed DOA estimation methods in Fig. 7, the best results are obtained using the BeamLearning approach trained using a data augmentation process with a random $\text{SNR} \geq 5$ dB. This result still stands when the testing dataset is used with SNR as small as -1 dB. For a moderate noise data augmentation (a random SNR training larger than 25 dB), the BeamLearning approach offer very similar SSL performances than both MUSIC and SRP-PHAT for unseen testing data with a $\text{SNR} \geq 10$ dB. Without any data augmentation, the BeamLearning approach only outperforms MUSIC and SRP-PHAT for unseen testing data with a $\text{SNR} \geq 30$ dB. It is also interesting to note that the use of the proposed data augmentation mechanism does not degrade the localization accuracy at high SNR, even if the data augmentation uses very noisy data. This property is directly linked to the fact that the training process is performed with randomly picked signal-to-noise ratio (SNR) values in the range of $[x, +\infty[$ dB, which does not give more importance to noisy data than to noise-free data when training the BeamLearning DNN.

In free-field (Fig. 7) and for low-noise measurements, the three methods offer very similar SSL performances. However,

as expected, the MUSIC algorithm performs significantly better than SRP-PHAT for heavily noisy measurements ($\text{SNR} \leq 5$ dB). It is although interesting to note that the proposed BeamLearning approach performs even better than MUSIC at low SNRs in free field. Since the mean absolute angular errors obtained using wideband MUSIC and SRP-PHAT stand outside the interquartile range of the BeamLearning approach trained with a random $\text{SNR} \geq 5$ dB, one can also conclude that the improvement offered by the proposed BeamLearning approach is statistically significant over those two methods at low SNR : the DOA estimation precision is improved over MUSIC by 30% (from 10.0° to 7.0° at -1 dB).

The same trend is also observed in a reverberating environment (Tab. III), where the BeamLearning approach offers better angular estimations for all tested SNRs. The improvements in terms of SSL accuracies range from 23% at -1 dB to 75% at 20 dB when compared to MUSIC, and from 49% at -1dB to 21% at 20 dB when compared to SRP-PHAT. It is also interesting to note that for low-noise measurements ($\text{SNR} \geq 15$ dB) in this reverberating environment, the SSL performances offered by MUSIC are degraded, since this subspace-based algorithm relies on the assumption that signal paths are uncorrelated or can be fully decorrelated, which does not hold true with early reflections in room acoustics, leaving DoA estimation a difficult problem in these situations for the wideband MUSIC method.

B. Robustness to mismatch between the propagating environments in the learning and testing phases

One of the main common pitfalls when using machine learning approaches resides in overfitting training data [90]. In order to avoid this phenomenon, the authors in [7] proposed the use of a training dataset composed by a large number of randomly sized rooms, which is an excellent way of circumventing these potential problems.

Still, it is interesting to study the robustness of the BeamLearning DNN, when tested in a totally different propagating environment than the one used for training. Using the same frozen network as in the previous subsection, trained using $\mathcal{D}_{\text{train}}^{\text{Room1}}$, the frozen BeamLearning DNN is tested using 3600 sources positions in $\mathcal{D}_{\text{test speech}}^{\text{Room2}}$ in a $5 \times 5 \times 4$ m room, and the results are compared to those obtained using MUSIC and SRP-PHAT (see Tab. IV), for different SNR situations.

This testing room exhibits a volume reduced by half, when compared to the one used during the training process. However, with a same reverberation time ($T_r = 0.5$ s), the testing room walls are therefore much more reflective, and the positions of the sources are much closer to the room boundaries than during the training process, which is commonly recognized as a difficult task for DoA estimation [91]. The degraded performances obtained using MUSIC in

TABLE IV
MEAN ABSOLUTE ANGULAR MISMATCH OBTAINED FOR 3600 SOURCES POSITIONS IN $\mathcal{D}_{\text{TEST SPEECH}}^{\text{ROOM2}}$, FOR DIFFERENT SNRS. THE DNN HAS BEEN PRE-TRAINED USING $\mathcal{D}_{\text{TRAIN}}^{\text{ROOM1}}$, WITH RANDOM $\text{SNR} \geq 5$ DB. THE BEST RESULTS ARE IN BOLD.

SNR (dB)	-1	2	5	10	15	20
SRP-PHAT	53.7°	42.0°	30.5°	15.5°	9.2°	6.6°
MUSIC	41.9°	36.7°	30.8°	26.9°	24.6°	25.0°
BeamLearning	35.7°	25.9°	18.7°	11.1°	8.0°	6.8°

this testing room illustrates very well this observation. On the other hand, despite this significant change in the properties of the room between the training dataset and the testing dataset, the BeamLearning approach still offers significantly enhanced performances when compared to MUSIC and SRP-PHAT, except for a high SNR of 20 dB, where SRP-PHAT and BeamLearning allow equivalent SSL accuracies.

C. Robustness to change of signal characteristics

Similarly, we also investigate the influence of a mismatch between test signals emitted by the sources and the signals used during the training procedure. We use the same frozen network as in the previous subsection. As detailed in III-A, the signals emitted by the labelled sources in $\mathcal{D}_{\text{train}}^{\text{Room1}}$ during the training process are composed of speech signals, symphonic music, and car horn honks. Even if the speech signals used for testing in the previous subsections were unseen during training, it is important to check if the performances remain satisfying when using a signal whose temporal and spectral dynamics are dissimilar to the ones used during the training process. This is the reason why Tab. V shows the localization performances obtained for SNRs of 5 dB and 15 dB using sound sources emitting dog bark excerpts from the UrbanSound8K dataset [77], along with the SSL results obtained with unseen voice signals.

The wideband MUSIC method allows a better localization of dog barks when compared to voice signals, mainly due to a smaller bandwidth occupied by the signal subspace when compared to the noise subspace. Still, the BeamLearning approach offers again better localization performances than both SRP-PHAT and wideband MUSIC in this situation. The analysis of these results allows to conclude that in each case, the BeamLearning approach outperforms MUSIC and SRP-PHAT in terms of SSL accuracy, and that the SSL performances observed with the frozen DNN do not degrade significantly when a signal with a different spectral and temporal characteristics is emitted by the sources to be located.

D. Computational efficiency for realtime inference

Since the proposed SSL method is a learning based method, the optimization of the 1.1×10^6 variables of the BeamLearning DNN represents a great amount of computational time (approx. 7 days on a Nvidia[®] GTX 1080TI GPU card using the Tensorflow framework). This

TABLE V

MEAN ABSOLUTE ANGULAR MISMATCH OBTAINED USING SEVERAL UNSEEN SIGNALS IN AN REVERBERANT ENVIRONMENT (RT = 0.5 s), FOR DIFFERENT SNRS FOR 3600 SOUND SOURCES POSITIONS. THE DNN HAS BEEN PRE-TRAINED USING $\mathcal{D}_{\text{TRAIN}}^{\text{ROOM1}}$, WITH RANDOM SNR ≥ 5 DB. THE BEST RESULTS ARE IN BOLD.

SNR	Localization approach	Voice signals	Dog bark
5 dB	MUSIC	17.8°	8.9°
	SRP-PHAT	17.4°	16.9°
	BeamLearning	8.2°	6.7°
15 dB	MUSIC	11.0°	6.4°
	SRP-PHAT	4.6°	6.1°
	BeamLearning	3.4°	4.2°

corresponds to a total of 10^6 successive iterations for learning (approx. 434 epochs). Each of these learning iterations includes the calculation of the feed forward propagation, cross entropy loss, back-propagation, gradients computations, and variables updates using Adam, for each mini-batch. On this GPU architecture, and for each of the audio excerpts of 23.2 ms (1024 samples), the mean computation time for the learning process is therefore only 6 ms for the whole learning process involved. Given that some amount of these 6 ms are dedicated to the gradient optimizer operations, which are not needed for the inference with a frozen model, this gives confidence for a realtime inference.

TABLE VI
COMPUTATION TIMES FOR 3600 SOURCES LOCALIZATIONS

	Mean time (CPU)	Mean time (GPU)
MUSIC	45 min 40 s	<i>not implemented</i>
SRP-PHAT	3 min 01 s	<i>not implemented</i>
BeamLearning	2 min 47 s	16.2 s

In order to compare the computational time required by the three algorithms for an inference task, we benchmarked the SSL tasks detailed in section V (see Table VI). The BeamLearning DNN has been tested both on a Nvidia[®] GTX 1080TI GPU and a i7-6900K CPU (3.20GHz) with a mini-batch of size 1. On the other hand, MUSIC and SRP-PHAT have only been tested on the same CPU, since this is the only possible implementation proposed by [61]. In both cases, the BeamLearning approach allows to estimate the position of the sources in a faster way than the model approaches tested under the same conditions, with drastic improvements using a GPU computation.

VI. CONCLUSION

In this paper, we have presented BeamLearning, a new deep learning approach for the localization of acoustic sources. The proposed DNN allows to achieve a 2D-DoA task in real-time

from raw multichannel measurements on a microphone array. The proposed DNN architecture is partly inspired by the principle of filter and sum beamforming. In particular, the use of depthwise separable atrous convolutions combined with pointwise convolutions allows to process the input audio data at multiple timescales. The BeamLearning network acts as a joint-feature learning process, that efficiently encodes the spatio-temporal information contained in raw measurements, and accomodates heavy measurement noise situations and reverberation. An extensive analysis of the BeamLearning approach using both a classification framework and a regression framework also allowed to show that the regression approach seems better suited when a fine angular resolution is sought, with computation times that open up the possibility of precise and real-time localization using a deep learning approach, without any pre-processing of microphone array measurements. The analysis of the BeamLearning approach for SSL tasks in noisy and reverberating environments also proves that BeamLearning outperforms state of the art SRP-PHAT and MUSIC in almost each test situations.

In addition to the proposed network architecture, we also proposed the use of a fast and accurate RIR generator on GPU in order to build large realistic training datasets in an efficient way, using an image source method. Using these datasets, the proposed data augmentation process allowed to obtain increased robustness to measurement noise when compared to state of the art model-based SSL algorithms that were specifically crafted to handle this kind of degraded SNR situations. We also demonstrate that even if the propagation environment in which the localization is carried out significantly differs from the one used during the training phase, the BeamLearning approach remains quite robust.

The good SSL performances obtained for a 2D-DoA task motivates under development extensions of this work, which include realtime 3D-DoA and simultaneous localization / sound source recognition using the BeamLearning approach. Further developments may include the use of multi-room datasets and multi-source localization.

ACKNOWLEDGMENTS

This work has been partially supported through the Deeplo-matics project funded by DGA/AID (ANR-18-ASTR-0008).

REFERENCES

- [1] B. Z. Steinberg, M. J. Beran, S. H. Chin, and J. H. Howard Jr, "A neural network approach to source localization," *The Journal of the Acoustical Society of America* **90**(4), 2081–2090 (1991).
- [2] K. Guentchev and J. Weng, "Learning-based three dimensional sound localization using a compact non-coplanar array of microphones," in *Proc. AAAI Spring Symposium on Intelligent Environments* (1998).
- [3] J. Weng and K. Y. Guentchev, "Three-dimensional sound localization from a compact non-coplanar array of microphones using tree-based learning," *The Journal of the Acoustical Society of America* **110**(1), 310–323 (2001).
- [4] R. Talmon, I. Cohen, and S. Gannot, "Supervised source localization using diffusion kernels," in *2011 IEEE workshop on applications of signal processing to audio and acoustics (WASPAA)*, IEEE (2011), pp. 245–248.

- [5] A. Deleforge, F. Forbes, and R. Horaud, "Acoustic space learning for sound-source separation and localization on binaural manifolds," *International journal of neural systems* **25**(1), 1440003 (2015).
- [6] S. Chakrabarty and E. A. Habets, "Multi-speaker DOA estimation using deep convolutional networks trained with noise signals," *IEEE Journal of Selected Topics in Signal Processing* **13**(1), 8–21 (2019).
- [7] L. Perotin, R. Serizel, E. Vincent, and A. Guerin, "CRNN-based multiple DoA estimation using acoustic intensity features for Ambisonics recordings," *IEEE Journal of Selected Topics in Signal Processing* **13**(1), 22–33 (2019).
- [8] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, "Sound Event Localization and Detection of Overlapping Sources Using Convolutional Recurrent Neural Networks," *IEEE Journ. of Selected Topics in Signal Processing* **13**(1), 34–48 (2019).
- [9] A. Brendel and W. Kellermann, "Distributed source localization in acoustic sensor networks using the coherent-to-diffuse power ratio," *IEEE Journal of Selected Topics in Signal Processing* **13**(1), 61–75 (2019).
- [10] S. Gannot, M. Haardt, W. Kellermann, and P. Willett, "Introduction to the issue on acoustic source localization and tracking in dynamic real-life scenes," *IEEE Journal of Selected Topics in Signal Processing* **13**(1), 3–7 (2019).
- [11] H. W. Löllmann, C. Evers, A. Schmidt, H. Mellmann, H. Barfuss, P. A. Naylor, and W. Kellermann, "The LOCATA challenge data corpus for acoustic source localization and tracking," in *2018 IEEE 10th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, IEEE (2018), pp. 410–414.
- [12] J. Pak and J. W. Shin, "LOCATA Challenge: A Deep Neural Networks-Based Regression Approach for Direction-Of-Arrival Estimation," in *Proc. of LOCATA Challenge Workshop-a satellite event of IWAENC* (2018).
- [13] W. Liu, Y. Yang, M. Xu, L. Lü, Z. Liu, and Y. Shi, "Source localization in the deep ocean using a convolutional neural network," *The Journal of the Acoustical Society of America* **147**(4), EL314–EL319 (2020).
- [14] J. Pak and J. W. Shin, "Sound localization based on phase difference enhancement using deep neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **27**(8), 1335–1345 (2019).
- [15] L. Comanducci, F. Borra, P. Bestagini, F. Antonacci, S. Tubaro, and A. Sarti, "Source localization using distributed microphones in reverberant environments based on deep learning and ray space transform," *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **28**, 2238–2251 (2020).
- [16] M. Yasuda, Y. Koizumi, S. Saito, H. Uematsu, and K. Imoto, "Sound Event Localization Based on Sound Intensity Vector Refined by Dnn-Based Denoising and Source Separation," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2020), pp. 651–655.
- [17] R. Varzandeh, K. Adiloğlu, S. Doclo, and V. Hohmann, "Exploiting Periodicity Features for Joint Detection and DOA Estimation of Speech Sources Using Convolutional Neural Networks," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2020), pp. 566–570.
- [18] K. L. Gemba, S. Nannuru, and P. Gerstoft, "Robust ocean acoustic localization with sparse bayesian learning," *IEEE Journal of Selected Topics in Signal Processing* **13**(1), 49–60 (2019).
- [19] Y. Liu, H. Niu, and Z. Li, "A multi-task learning convolutional neural network for source localization in deep ocean," *The Journal of the Acoustical Society of America* **148**(2), 873–883 (2020).
- [20] S. Chakrabarty and E. A. Habets, "Broadband DOA estimation using convolutional neural networks trained with noise signals," in *Applications of Signal Processing to Audio and Acoustics (WASPAA), 2017 IEEE Workshop on applications* (2017), pp. 136–140.
- [21] S. Adavanne, A. Politis, and T. Virtanen, "A Multi-room Reverberant Dataset for Sound Event Localization and Detection," in *Submitted to Detection and Classification of Acoustic Scenes and Events 2019 Workshop (DCASE2019)*, Munich, Germany (2019), <https://arxiv.org/abs/1905.08546>.
- [22] T. Hirvonen, "Classification of spatial audio location and content using convolutional neural networks," in *Audio Engineering Society Convention 138*, Audio Engineering Society (2015).
- [23] W. He, P. Motlicek, and J.-M. Odobez, "Deep neural networks for multiple speaker detection and localization," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE (2018), pp. 74–79.
- [24] H. Sundar, W. Wang, M. Sun, and C. Wang, "Raw waveform based end-to-end deep convolutional network for spatial localization of multiple acoustic sources," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2020), pp. 4642–4646.
- [25] N. Yalta, K. Nakadai, and T. Ogata, "Sound source localization using deep learning models," *Journal of Robotics and Mechatronics* **29**(1), 37–48 (2017).
- [26] 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* **25**(12), 2444–2453 (2017).
- [27] Q. Nguyen, L. Girin, G. Bailly, F. Elisei, and D.-C. Nguyen, "Autonomous sensorimotor learning for sound source localization by a humanoid robot," in *Workshop on Crossmodal Learning for Intelligent Robotics in conjunction with IEEE/RJSJ IROS*.
- [28] S. Sivasankaran, E. Vincent, and C. T. Jin, "Keyword-based speaker localization: Localizing a target speaker in a multi-speaker environment," in *Interspeech 2018 - 19th Annual Conference of the International Speech Communication Association*.
- [29] 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 *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*, IEEE (2016), pp. 1–6.
- [30] E. L. Ferguson, S. B. Williams, and C. T. Jin, "Sound source localization in a multipath environment using convolutional neural networks," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2018), pp. 2386–2390.
- [31] Z. Tang, J. D. Kanu, K. Hogan, and D. Manocha, "Regression and classification for direction-of-arrival estimation with convolutional recurrent neural networks," *arXiv preprint arXiv:1904.08452* (2019).
- [32] D. Salvati, C. Drioli, and G. L. Foresti, "Exploiting cnns for improving acoustic source localization in noisy and reverberant conditions," *IEEE Transactions on Emerging Topics in Computational Intelligence* **2**(2), 103–116 (2018).
- [33] R. Takeda and K. Komatani, "Sound source localization based on deep neural networks with directional activate function exploiting phase information," in *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on Acoustics, speech and Signal Processing (ICASSP)*, IEEE (2016), pp. 405–409.
- [34] M. Yiwere and E. J. Rhee, "Distance estimation and localization of sound sources in reverberant conditions using deep neural networks," *International Journal of Applied Engineering Research* **12**(22), 12384–12389 (2017).
- [35] X. Xiao, S. Zhao, X. Zhong, D. L. Jones, E. S. Chng, and H. Li, "A learning-based approach to direction of arrival estimation in noisy and reverberant environments," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2015), pp. 2814–2818.
- [36] D. Suvorov, G. Dong, and R. Zhukov, "Deep residual network for sound source localization in the time domain," *arXiv preprint arXiv:1808.06429* (2018).
- [37] J. Vera-Diaz, D. Pizarro, and J. Macias-Guarasa, "Towards End-to-End Acoustic Localization Using Deep Learning: From Audio Signals to Source Position Coordinates," *Sensors* **18**(10), 3418 (2018).
- [38] Y. Huang, X. Wu, and T. Qu, "A time-domain unsupervised learning based sound source localization method," in *2020 IEEE 3rd International Conference on Information Communication and Signal Processing (ICICSP)*, IEEE (2020), pp. 26–32.
- [39] D. Comminiello, M. Lella, S. Scardapane, and A. Uncini, "Quaternion convolutional neural networks for detection and localization of 3d sound events," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2019), pp. 8533–8537.
- [40] Y. Hu, P. Samarasinghe, S. Gannot, and T. Abhayapala, "Semi-supervised multiple source localization using relative harmonic coefficients under noisy and reverberant environments," *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **28**, 3108–3123 (2020).
- [41] Y. Hu, P. N. Samarasinghe, T. D. Abhayapala, and S. Gannot, "Unsupervised multiple source localization using relative harmonic coefficients," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2020), pp. 571–575.
- [42] F.-R. Stöter, S. Chakrabarty, B. Edler, and E. A. Habets, "Countnet: Estimating the number of concurrent speakers using supervised learning,"

- IEEE/ACM Transactions on Audio, Speech, and Language Processing **27**(2), 268–282 (2018).
- [43] P. Grumiaux, S. Kitić, L. Girin, and A. Guérin, “High-resolution speaker counting in reverberant rooms using CRNN with ambisonics features,” in *28th European Signal Processing Conference, EUSIPCO 2020, Amsterdam, Netherlands, January 18-21, 2021*, IEEE (2020), pp. 71–75.
- [44] M. J. Bianco, P. Gerstoft, J. Traer, E. Ozanich, M. A. Roch, S. Ganot, and C.-A. Deledalle, “Machine learning in acoustics: Theory and applications,” *The Journal of the Acoustical Society of America* **146**(5), 3590–3628 (2019).
- [45] S. Dieleman and B. Schrauwen, “End-to-end learning for music audio,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2014), pp. 6964–6968.
- [46] W. Dai, C. Dai, S. Qu, J. Li, and S. Das, “Very deep convolutional neural networks for raw waveforms,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2017), pp. 421–425.
- [47] T. N. Sainath, R. J. Weiss, A. Senior, K. W. Wilson, and O. Vinyals, “Learning the speech front-end with raw waveform CLDNNs,” in *Sixteenth Annual Conference of the International Speech Communication Association* (2015).
- [48] J. Lee, J. Park, K. L. Kim, and J. Nam, “Samplecnn: End-to-end deep convolutional neural networks using very small filters for music classification,” *Applied Sciences* **8**(1), 150 (2018).
- [49] É. Bavu, A. Ramamonjy, H. Pujol, and A. Garcia, “TimeScaleNet : a Multiresolution Approach for Raw Audio Recognition using Learnable Biquadratic IIR Filters and Residual Networks of Depthwise-Separable One-Dimensional Atrous Convolutions,” *IEEE Journ. of Selected Topics in Signal Processing* **13**(2), 220–235 (2019).
- [50] F. Chollet, “Deep learning with depthwise separable convolutions,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, IEEE (2017), pp. 1251–1258.
- [51] M. Ravanello and Y. Bengio, “Speaker recognition from raw waveform with sincnet,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*, IEEE (2018), pp. 1021–1028.
- [52] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” arXiv preprint arXiv:1609.03499 (2016).
- [53] L. Kaiser, A. N. Gomez, and F. Chollet, “Depthwise separable convolutions for neural machine translation,” arXiv preprint arXiv:1706.03059 (2017).
- [54] D. Rethage, J. Pons, and X. Serra, “A wavenet for speech denoising,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2018), pp. 5069–5073.
- [55] L. Perotin, A. Défossez, E. Vincent, R. Serizel, and A. Guérin, “Regression versus classification for neural network based audio source localization,” in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, IEEE (2019), pp. 343–347.
- [56] W. He, P. Motlicek, and J. Odobez, “Adaptation of multiple sound source localization neural networks with weak supervision and domain-adversarial training,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2019), pp. 770–774.
- [57] P. Lecomte, P.-A. Gauthier, C. Langrenne, A. Berry, and A. Garcia, “A fifty-node lebedev grid and its applications to ambisonics,” *Journal of the Audio Engineering Society* **64**(11), 868–881 (2016).
- [58] P. Lecomte, “Ambitools: Tools for sound field synthesis with higher order ambisonics-v1. 0,” in *1st international Faust Conference (IFC-18)*.
- [59] H. Pujol, E. Bavu, and A. Garcia, “Source localization in reverberant rooms using Deep Learning and microphone arrays,” in *23rd International Congress on Acoustics (ICA 2019 Aachen)*, Aachen, Germany (2019), pp. 6929–6936.
- [60] J. B. Allen and D. A. Berkley, “Image method for efficiently simulating small-room acoustics,” *The Journal of the Acoustical Society of America* **65**(4), 943–950 (1979).
- [61] R. Scheibler, E. Bezzam, and I. Dokmanić, “Pyroomacoustics: A python package for audio room simulation and array processing algorithms,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE (2018), pp. 351–355.
- [62] M. Abadi, A. Agarwal, *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems” (2015), <http://tensorflow.org/> software available from tensorflow.org.
- [63] J. O. Smith, *Introduction to Digital Filters with Audio Applications* (W3K Publishing, 2007).
- [64] M. Brandstein and D. Ward, *Microphone arrays: signal processing techniques and applications* (Springer-Verlag, Berlin, Germany, 2001).
- [65] M. Holschneider, R. Kronland-Martinet, J. Morlet, and P. Tchamitchian, “A real-time algorithm for signal analysis with the help of the wavelet transform,” in *Wavelets* (Springer, 1990), pp. 286–297.
- [66] K. He, X. Zhang, S. Ren, and J. Sun, “Identity mappings in deep residual networks,” in *European conference on computer vision ECCV’16*, Springer (2016), Vol. 9908, pp. 630–645.
- [67] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)* (2016), pp. 770–778.
- [68] R. K. Srivastava, K. Greff, and J. Schmidhuber, “Training very deep networks,” in *NIPS’15: Proceedings of the 28th International Conference on Neuronal information processing systems* (2015), pp. 2377–2385.
- [69] J. L. Ba, J. R. Kiros, and G. E. Hinton, “Layer normalization,” arXiv preprint arXiv:1607.06450 (2016).
- [70] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” arXiv preprint arXiv:1502.03167 (2015).
- [71] S. Ioffe, “Batch renormalization: Towards reducing minibatch dependence in batch-normalized models,” in *Advances in neural information processing systems* (2017), pp. 1945–1953.
- [72] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, “Self-normalizing neural networks,” in *Advances in Neural Information Processing Systems 30* (2017), pp. 972–981.
- [73] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” arXiv preprint arXiv:1412.6980 (2014).
- [74] E. A. Lehmann, A. M. Johansson, and S. Nordholm, “Reverberation-time prediction method for room impulse responses simulated with the image-source model,” in *2007 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, IEEE (2007), pp. 159–162.
- [75] E. A. Lehmann and A. M. Johansson, “Prediction of energy decay in room impulse responses simulated with an image-source model,” *The Journal of the Acoustical Society of America* **124**(1), 269–277 (2008).
- [76] T. Lokki, J. Patynen, and V. Pulkki, “Recording of anechoic symphony music,” *Journal of the Acoustical Society of America* **123**(5), 3936–3936 (2008).
- [77] J. Salamon, C. Jacoby, and J. P. Bello, “A dataset and taxonomy for urban sound research,” in *22nd ACM International Conference on Multimedia (ACM-MM’14)*, Orlando, FL, USA, pp. 1041–1044.
- [78] R. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE transactions on antennas and propagation* **34**(3), 276–280 (1986).
- [79] J. H. DiBiase, H. F. Silverman, and M. S. Brandstein, “Robust localization in reverberant rooms,” in *Microphone Arrays*, Chap. 8, pp. 157–180.
- [80] M. Sokolova and G. Lapalme, “A systematic analysis of performance measures for classification tasks,” *Information processing & management* **45**(4), 427–437 (2009).
- [81] H. Do and H. F. Silverman, “SRP-PHAT methods of locating simultaneous multiple talkers using a frame of microphone array data,” in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, IEEE (2010), pp. 125–128.
- [82] A. Badali, J.-M. Valin, F. Michaud, and P. Aarabi, “Evaluating real-time audio localization algorithms for artificial audition in robotics,” in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE (2009), pp. 2033–2038.
- [83] M. Cobos, A. Martí, and J. J. Lopez, “A modified SRP-PHAT functional for robust real-time sound source localization with scalable spatial sampling,” *IEEE Signal Processing Letters* **18**(1), 71–74 (2010).
- [84] J. P. Dmochowski and J. Benesty, “Steered beamforming approaches for acoustic source localization,” in *Speech processing in modern communication* (Springer, 2010), pp. 307–337.
- [85] C. Zhang, D. Florêncio, and Z. Zhang, “Why does PHAT work well in lownoise, reverberative environments?,” in *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, IEEE (2008), pp. 2565–2568.
- [86] S. Argentieri and P. Danes, “Broadband variations of the MUSIC high-resolution method for sound source localization in robotics,” in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE (2007), pp. 2009–2014.

- [87] J. P. Dmochowski, J. Benesty, and S. Affes, "Broadband MUSIC: Opportunities and challenges for multiple source localization," in *2007 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, IEEE (2007), pp. 18–21.
- [88] C. T. Ishi, O. Chatot, H. Ishiguro, and N. Hagita, "Evaluation of a MUSIC-based real-time sound localization of multiple sound sources in real noisy environments," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE (2009), pp. 2027–2032.
- [89] J. X. Zhang, M. G. Christensen, J. Dahl, S. H. Jensen, and M. Moonen, "Robust implementation of the MUSIC algorithm," in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, IEEE (2009), pp. 3037–3040.
- [90] R. Roelofs, V. Shankar, B. Recht, S. Fridovich-Keil, M. Hardt, J. Miller, and L. Schmidt, "A meta-analysis of overfitting in machine learning," *Advances in Neural Information Processing Systems* **32**, 9179–9189 (2019).
- [91] S. Shen, D. Chen, Y.-L. Wei, Z. Yang, and R. R. Choudhury, "Voice localization using nearby wall reflections," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking* (2020), pp. 1–14.