

JOINT MODEL-BASED RECOGNITION AND LOCALIZATION OF OVERLAPPED ACOUSTIC EVENTS USING A SET OF DISTRIBUTED SMALL MICROPHONE ARRAYS

Rupayan Chakraborty and Climent Nadeu

TALP Research Center
Department of Signal Theory and Communications
Universitat Politècnica de Catalunya, Barcelona, Spain
{rupayan.chakraborty,climent.nadeu}@upc.edu

ABSTRACT

In the analysis of acoustic scenes, often the occurring sounds have to be detected in time, recognized, and localized in space. Usually, each of these tasks is done separately. In this paper, a model-based approach to jointly carry them out for the case of multiple simultaneous sources is presented and tested. The recognized event classes and their respective room positions are obtained with a single system that maximizes the combination of a large set of scores, each one resulting from a different acoustic event model and a different beamformer output signal, which comes from one of several arbitrarily-located small microphone arrays. By using a two-step method, the experimental work for a specific scenario consisting of meeting-room acoustic events, either isolated or overlapped with speech, is reported. Tests carried out with two datasets show the advantage of the proposed approach with respect to some usual techniques, and that the inclusion of estimated priors brings a further performance improvement.

Index Terms— Acoustic event detection, audio recognition, acoustic source localization, scene analysis, sound model, simultaneous sources, beamforming

1. INTRODUCTION

The automatic analysis of acoustic scenes requires several functionalities: detection, recognition, localization, separation, enhancement, etc. Usually, these functionalities are handled by different sub-systems. However, we can expect that carrying out some of them jointly, i.e. with a single system and sharing a given processing framework, an advantage is obtained in terms of effectiveness and efficiency.

Often recognition and localization are required for the same application scenario. This may happen with speech (e.g. [1, 2, 3]), or with other types of acoustic events (AE) (e.g. [4, 5, 6]). In this work, we have developed a system that can detect, recognize and localize acoustic events, i.e. it estimates the identities of the acoustic events and their positions along the time axis and the spatial axes. The system carries out all those tasks jointly, employing a single processing scheme

for all of them, a scheme that uses models of the involved acoustic events. Furthermore, the system tackles the problem of simultaneity of acoustic events.

In fact, time overlap of acoustic events is a major factor of complexity for both acoustic event recognition and acoustic source localization. Concerning the first, from the analysis of the results submitted by the participants in the meeting-room acoustic event detection (AED) task (which includes recognition) of the *CLEAR'2007* evaluation campaign [7], it was apparent that time overlapping of acoustic events caused more than two thirds of the detection errors. More recently, in the *D-CASE* evaluation [8] a big gap was observed between the AED accuracy of isolated and overlapped events. In terms of *F*-score, the best results were 45.17% and 8.45%, respectively. In this paper, we deal with the overlapping problem by taking into account the spatial diversity of the sound sources.

To estimate the position coordinates of the acoustic sources, most widely used acoustic source localization methods consider phase-based measures [9, 10, 11, 12, 13]. Conversely, in this work we use the probability or similarity measure delivered by a classifier that employs the whole time-spectral information of the signals as input. As the classifier uses models for the different sound classes, those models can be shared by both recognition and localization tasks.

In [14], a processing framework based on array beamformers and sound-model-based likelihood computations was introduced by these authors to recognize acoustic events that may overlap in time; there, recognition was actually reduced to a classification problem and only two microphone arrays were used. Posteriorly, a first trial for doing both recognition (actually, classification) and localization was reported in [15], but only the direction of arrival was estimated and once the sound identity was found. More recently, the same processing framework was applied to the source localization problem [16]. In all those works, knowledge of the time end-points of the events was assumed in order to compute event-level likelihoods. Conversely, that assumption is removed here. Furthermore, the acoustic events, which may overlap, are jointly detected, identified and localized with a single system. In

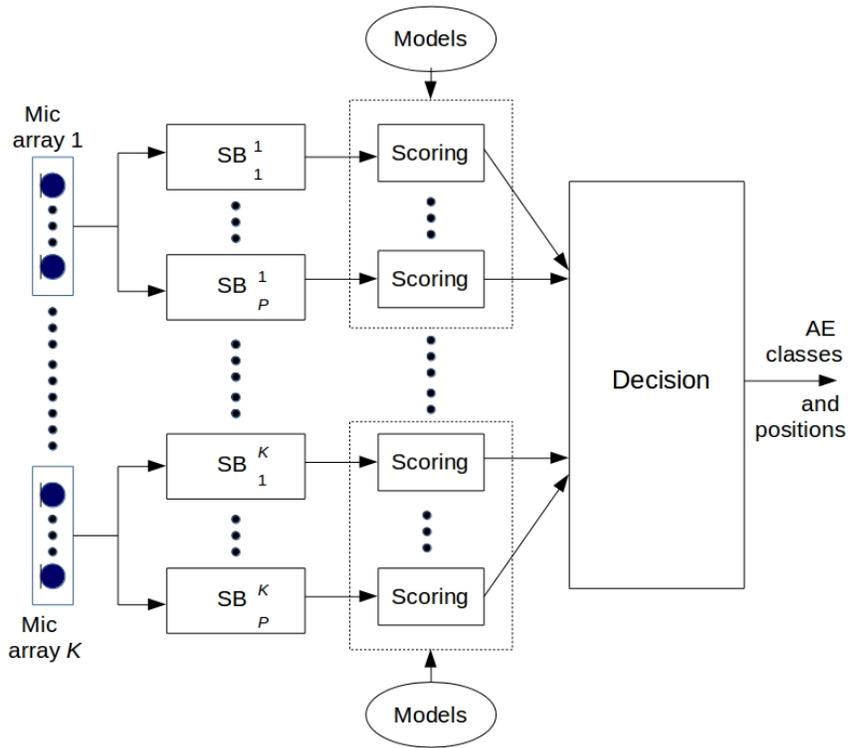


Fig. 1: Processing scheme for both classification and localization. For each of the K microphone arrays deployed in the room, a set of signals are produced at the output of P array processors (SB). Given an output signal, that results from targeting a specific room (cell) position, a score is computed for each of the C acoustic event classes using the set of models. The resulting $K \times P \times C$ scores are combined with a given criterion to estimate both the class identity and the room position.

other words, the presented system, which employs multiple arbitrarily-located small microphone arrays, is able to determine the classes of the acoustic events (that occur either in isolation or simultaneously) along with their time end-points and their respective positions in the room space.

Like in those previous works, the experiments reported in this paper are carried out for a concrete meeting-room scenario with one or two simultaneous sources, and using a database collected in the smart-room, where six microphone arrays are distributed on the walls. For jointly doing recognition and localization, a particular methodology is employed that follows a two-step approach. First the events are detected along time, and then their class identities and space positions are estimated using a maximum-a-posteriori criterion. The recognition and localization results are obtained with that two-step method using all the six three-microphone arrays in the room. Moreover, besides the experiments with the section of the database that was employed in previous works, which consists of isolated sound recordings, results with the other section of the database that includes actual two-source overlapping sounds are reported. The proposed technique shows better recognition

performance than a system which models all the combinations of classes that occur in the database, and it shows better event-based space localization accuracy than the well-known *SRP-PHAT* technique [11].

The rest of the paper is organized as follows. The conceptual approach is presented in Section 2, the developed system is described in Section 3, and in Section 4 we report the experimental work. Discussions and conclusions are given in Section 5.

2. JOINT CLASSIFICATION AND LOCALIZATION APPROACH

Let us assume a room with a set of K microphone arrays. Those arrays are arbitrarily located, so they do not follow a spatially-structured configuration; for system deployment, this is an advantage. The $2D$ room space is divided into a set of P pre-defined small-area cells. Note that the vertical coordinate is not considered in this study, though it could be easily managed by the proposed system.

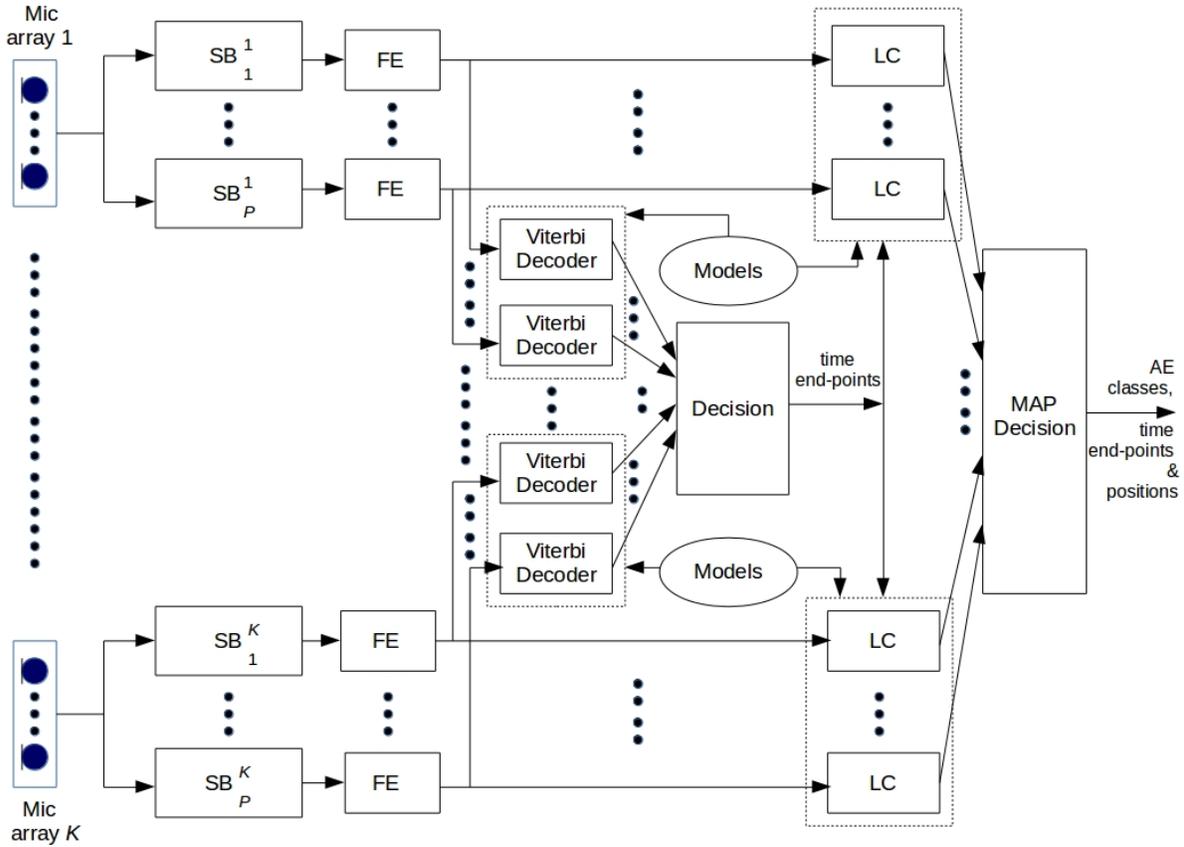


Fig. 2: The proposed two-step recognition and localization system. The *Viterbi* decoding of every channel has been included as a first step. In the intermediate decision block the sequence of time end-points that shows maximum likelihood is driven to the likelihood calculators from the second step. Note that steering beamformers (SB) and feature extractors (FE) are common to both steps.

Along time, N acoustic events take place, either in isolation or overlapped with other events. Those events belong to a set of C different acoustic classes. If we know the time end-points of every acoustic event, we can jointly assign a class to it and estimate its spatial position with the processing scheme depicted in Figure 1 already proposed by the authors in [14, 15]. In fact, for each microphone array, a set of P steering beamformers (SB) is designed, so that the j^{th} beamformer, $1 \leq j \leq P$, attenuates the signals coming from all directions except the one that corresponds to position (cell) s_j . Then, each beamformer output signal is separately fed into a classifier that computes a score for each pre-defined event class by using a set of C pre-trained acoustic models. The whole set of $K \times P \times C$ likelihood scores is used by the decision block to estimate both the identity and the source positions of the acoustic events using an optimality criterion.

The above processing scheme follows an event-based approach, i.e. to compute the likelihood with each class model, the time end-points of the (possibly overlapping) events must

be known. In order to remove that constraint, i.e. to do detection instead of just classification, several different alternatives could be devised. In this paper, a two-step approach is taken: (1) the events are firstly detected along time (but not recognized), and (2) the identity and the source position of each event are determined. The first step outputs an estimated time interval (two time end-points) for each detected event, which is subsequently used in the second step for classifying and localizing.

3. JOINT RECOGNITION AND LOCALIZATION SYSTEM

The whole scheme of the two-step technique implemented in this work is depicted in Figure 2. The two steps share the initial part of the processing scheme, that includes beamforming and feature extraction. In the first step, a sequence of events is detected for each beamformer output signal. In our experimental work, the acoustic events are modeled with hidden

Markov models (*HMM*), and the state emission probabilities are computed with continuous density Gaussian mixture models (*GMM*). Consequently, after feature extraction, the sequence of observation vectors is decoded with the *Viterbi* algorithm. Then, in the decision stage of this first step, the beamformer outputs are ranked according to the likelihood given by the *Viterbi* decoding. The time sequence of acoustic events corresponding to the beamformer output that shows the highest likelihood is then taken to the second step. In fact, as indicated in Figure 2, only the time end-points of the hypothesized events are taken to the likelihood calculators of the second step, which use the same *HMM-GMM* acoustic event models. If overlapped events are possible, the subsequent best sequences may also be taken to the second step.

In the second step, the last stage of the processing scheme from Figure 1 is applied for deciding the class identities and the corresponding source positions of the acoustic events. Taking the time end-points of the events from the first step, and using the same set of statistical models that are used in that first step, a $P \times C$ -dimensional matrix of likelihoods is build for each event and each microphone array. Then, both the class identities and the space positions of the events are estimated with a *MAP* criterion and combining the probabilities from all arrays.

3.1. Step 1: Time end-points estimation (detection)

Acoustic event recognition requires both segmentation of the audio stream and classification of the segments. For a single beamformer output, simultaneous segmentation and classification can be performed like it is usually done for continuous speech recognition [17, 18]. The goal for a single array can be formulated as follows: given X_k , the set of output vectors from the P feature extraction blocks corresponding to the k^{th} array, find the event sequence Ω that maximizes the posterior probability $p(\Omega|X_k)$:

$$\hat{\Omega} = \underset{\Omega}{\operatorname{argmax}} p(\Omega|X_k) = \underset{\Omega}{\operatorname{argmax}} p(X_k|\Omega)p(\Omega) \quad (1)$$

where the likelihood $p(X_k|\Omega)$ is computed using statistical models, and $p(\Omega)$ is the prior probability of the acoustic event sequence Ω . In order to avoid the dependence of that sequence to the particular room situation, all sequences of events are assumed equally probable, i.e. $p(\Omega)$ is constant and it does not affect the maximization in Equation 1.

As shown in Figure 2, given a set of P beamformers for each of the K microphone arrays, the decoding of every channel with the *Viterbi* algorithm [17, 18] yields a set of $L=PxK$ hypothesis, each consisting of a sequence of detected events with their corresponding end-points and its likelihood score. In the decision stage, the best channel is chosen by maximizing the posterior probability $p(\Omega_l|X_k)$, $1 \leq l \leq L$ across all channels.

The end-points of the acoustic events corresponding to the best channel are used in the second step for jointly doing classification and localization with the processing scheme depicted

in Figure 1. Note that, although the identities of the detected sounds are also hypothesized by this first-step algorithm, they are not used in the 2^{nd} step since a better accuracy is obtained using the *MAP* criterion in that step. Also, though an operationally simplified system would result from using just one sound class model that encompass all events, the end-points accuracy would suffer from the lack of specificity of that overall sound model, and model training would be more demanding as the specific *AE* models are also needed in 2^{nd} step but the single overall model would have to be additionally created.

3.2. Step 2: MAP-based classification and localization

To decide about the class identities and positions, all the channels (from all the arrays) are taken into account. As depicted in Figure 2, the feature vector sequence (already extracted in *Step 1*) from each beamformer output signal enters the classification system. Then, a likelihood score (*LC*) is computed for each of the event time intervals from *Step 1* by using the same *HMM-GMM* models that were used in that *Step 1*. Finally, a decision module carries out the classification and localization of the events by combining likelihood scores and priors using a *MAP* criterion and combining the probabilities from all arrays.

Given a room with K microphone arrays, let us assume we have a set of N events with labels c_i , $1 \leq i \leq N$, which belong to a set of C different classes. Given a grid of room positions s_j , $1 \leq j \leq P$, for each array, there is a set of P beamformers, so that the j^{th} beamformer attenuates the signals that come from all positions except position s_j . Thus, from that initial array signal processing block, we have a set of P output signals for each array, and after feature extraction and likelihood computations with the models of all classes, we have a $K \times P \times C$ -dimensional vector of likelihood scores. For identifying the sound class and localizing the sound sources, we first determine the posterior probability of a given class c_i and position s_j for each k^{th} array, i.e.

$$p(c_i, s_j|X_k) = p(X_k|c_i, s_j)p(c_i)p(s_j)/p(X_k) \quad (2)$$

where $p(X_k|c_i, s_j)$ is the likelihood of X_k given class i and position j , $p(c_i)$ is the (prior) probability of class i , $p(s_j)$ is the (prior) probability of position j , and $p(X_k)$ is the probability of the feature vector X_k .

In order to combine the posterior probabilities from the various microphone arrays, the product combination rule is used [19], and so the optimal class c_o and the optimal position s_o are chosen so as to maximize the product of posterior probabilities, i.e.

$$c_o, s_o = \underset{c_i, s_j}{\operatorname{argmax}} \prod_{k=1}^K p(c_i, s_j|X_k) \quad (3)$$

Note that for finding the maximum-a-posteriori (*MAP*) estimates c_o and s_o , the probability of the feature vector $p(X_k)$ is not required. Moreover, both the priors of classes and positions can be included in the optimization. In the experiments we will consider the latter.

3.3. Feature extraction and modeling

In the feature extraction block of the system, a set of audio spectro-temporal features is computed for each signal frame. Frames are 30 *ms* long, with a 20 *ms* shift, and a *Hanning* window is applied. We have used frequency-filtered logarithmic filter-bank energies (*FF-LFBE*) for the parametric representation of the spectral envelope of the audio signal, as it generally is a more robust alternative to *MFCC* features and requires a smaller amount of computations than them [20]. For computing the *FF-LFBE*s in a given frame, the *FIR* filter with transfer function $z-z^{-1}$ is applied to the sequence of log sub-band energies. This is equivalent to a convolution in the frequency domain that actually computes a smoothed derivative along frequency. Zeros are assumed beyond the vector ends and, after convolution, the end-points are also taken into account. In our experiments, 16 sub-bands are considered, so obtaining 16 *FF-LFBE* features per frame. Additionally, as it is usually done in speech recognition, temporal smoothed derivatives are also used to somewhat include the temporal evolution of the envelope. Therefore, the final feature vector, which has 32 elements, contains the first-order smoothed derivatives of the log sub-bank energies both along frequency and time.

Hidden Markov Models are used for modeling the *AE*s (including the speech event), and their state emission probabilities are modeled with continuous density Gaussian mixture Models. There is one left-to-right *HMM* with three emitting states for each *AE* and for each array. 32 Gaussian components with diagonal covariance matrix are used per state. Each *HMM-GMM* model is trained with the standard *Baum-Welch* algorithm [17], using all the signals from an array with the *HTK* toolkit [21]. The same models are used for both processing steps mentioned in Sub-section 3. To avoid mismatch between training and testing, the signals at the output of the beamformers are used for training the models. If each beamformer had its own set of models (for all acoustic events and silence), the total number of models would be too large. That is why we chose to have a common set of models for all the beamformers in a given array.

In the second step, where the optimal class identity and source position are obtained as indicated by Equation 3, flat priors are assigned to classes in the reported tests. However, regarding the positions, either flat or estimated prior probabilities are used. The latter are estimated by counting the event occurrences within the training dataset and computing the quotient between the number of occurrences in the cell and the total number of event occurrences in the dataset.

4. EXPERIMENTAL WORK

In the experimental work, a meeting-room scenario with a pre-defined set of eleven acoustic events plus speech is considered [7]. Like in [6], we assume that there may occur either 0, 1 or

2 simultaneous events.

4.1. Meeting-room acoustic scenario and database

Figure 3 shows the smart-room layout with the position of its six *T*-shaped four-microphone arrays on the walls [22]. The steering beamformers at the front-end of the system are designed to work with the horizontal row of three microphones available for each of the For system training, development and testing, we have used the audio part of a multi-modal database recorded in that smart-room with all microphones [6]¹. There are 11 meeting-room acoustic event (*AE*) classes: applause, spoon/cup jingle, chair moving, cough, door open/slam, key jingle, door knock, keyboard typing, phone ringing, paper work, and steps. The regions where the *AE* sources are located are shown in Figure 3. Signals were recorded with all the six *T*-shaped microphone arrays (24 microphones in total), but only the signals coming from the three horizontally-aligned microphones from every array are used in the experiments.

The database contains two different sections or datasets, and both will be used in these reported experiments. The first one, so called *S*-recordings, consists of 8 audio recording sessions which contain isolated acoustic events. For training, only these one-source signals are used in our experiments. The audio data corresponding to two simultaneous sources that are used for development and testing were obtained by adding those *AE* signals recorded in the room with a speech signal also recorded in the room from all 24 microphones. To do that, for each *AE* instance, a segment with the same length was extracted from the speech signal starting from a random position, and added to the *AE* signal. The mean power of speech was made equivalent to the mean power of the overlapping *AE*. That addition of signals produces an increment of the background noise level, since it is included twice in the overlapped signals; however, going from isolated to overlapped signals the *SNR* reduction is slight: from 18.7 *dB* to 17.5 *dB*. Although in our real meeting-room scenario the speaker may be placed at any point in the room, its position in the *S*-recordings dataset is fixed at a point at the left side (*'SP'* in Figure 3). However, this will not constrain the usefulness of the results, because the proposed system will not make use of that knowledge.

Additionally, a more realistic dataset was also used for testing: the *T*-recordings. In it, two participants (room positions *'SP 1'* and *'SP 2'* in Figure 3) interact with each other in a rather spontaneous way. While one is speaking, the other is randomly producing acoustic events, which belong to the above mentioned set of 11 meeting-room *AE* classes. Therefore, this dataset includes overlaps of an *AE* with speech that are more naturally produced than in the *S*-recordings dataset. All signals were recorded at 44.1 *kHz* sampling frequency, and they were further down-sampled to 16 *kHz* for our experiments.

¹That database is publicly available from the authors.

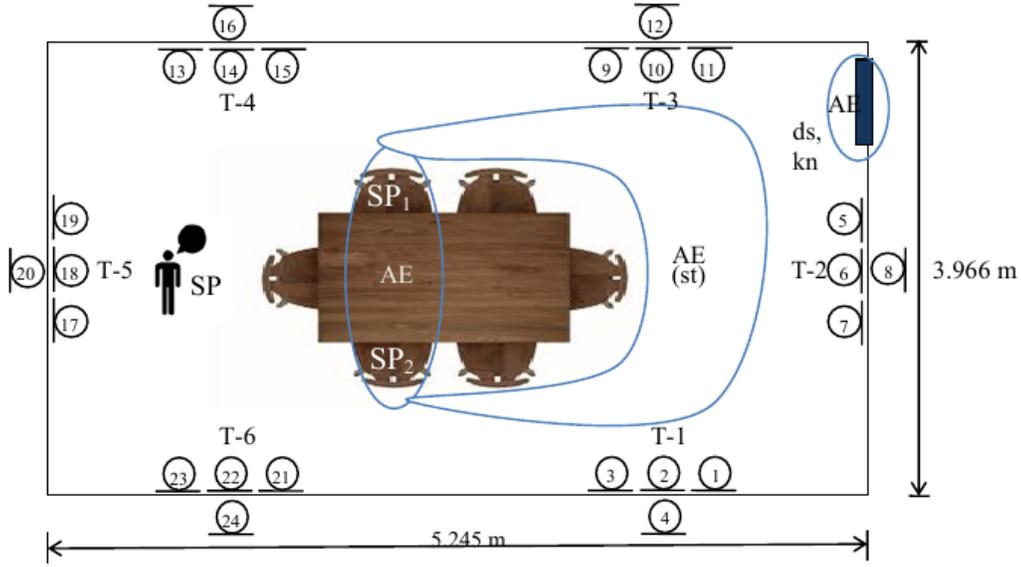


Fig. 3: Smart-room layout, with the positions of microphone arrays ($T-k$), acoustic events (AE), speaker (SP) in S -recordings, and speakers (SP_1 and SP_2) in T -recordings. Some place-specific acoustic events are indicated in the plot: steps (st), door-slam (ds) and door-knock (kn)

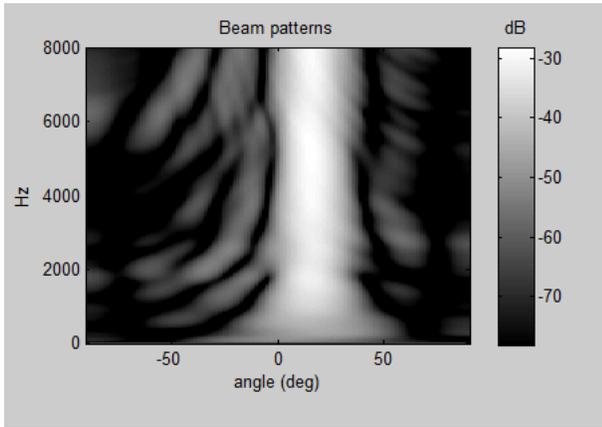


Fig. 4: Illustration of a frequency-invariant beam pattern

4.2. Experiments and results

In the reported experiments, the proposed system, depicted in Figure 2, is used to recognize the AE and to localize either one or two simultaneous acoustic event sources in the room environment. In order to deal with the broadband characteristics of the audio signal a frequency invariant beamforming used to determine the beamformer coefficients. The method, proposed in [23], uses a numerical approach to construct an optimal frequency invariant response for an arbitrary spacing a small number of microphones. It first decouples the spatial selectivity from the frequency selectivity by replacing the set

of real sensors by a set of virtual ones, which are frequency invariant. Then, the same array coefficients can be used for all frequencies.

All six microphone arrays ($T1$ to $T6$) available in the room are employed in the experiments. The steering beamformers at the front-end of the system are designed to work with the horizontal row of three microphones available for each of the arrays in the smart-room. An illustrative example of the beam pattern is shown in Figure 4 [22]. Note how the beams for the angle of interest are rather constant along frequency. Notice also that, due to the low number of microphones used in this particular experimental setup, the beamformer is steering to the target cell with a relatively broad beam. Regarding the design of the room grid considered in the experiments, as high localization precision is not required for the meeting-room application, a rather large cell size of $0.661\text{ m} \times 0.874\text{ m}$ has been chosen, and so there are 36 cells (6×6).

Testing results are obtained for both types of recordings and for all the sessions (i.e. eight sessions $S01$ to $S08$ for S -recordings, and nine sessions, $T01$ to $T09$, for T -recordings), using a leave-one-out criterion, i.e. recursively keeping one session for testing, while all the other sessions are used for training. The number of simultaneous sources is assumed known (1 or 2). For the two-source case, as mentioned above, one sound is always speech and the other is any one of the eleven acoustic events. However, the proposed system does not use that knowledge. In that two-source case, the identities and source positions of the two overlapping sounds are found

Table 1: Classification accuracies (%) for the various systems in the two-source scenario

			Proposed technique (1-step)		
	All-combinations modeling	BSS	Flat priors	Estimated priors	Known position
<i>S</i> -recordings	83.8	80.8	84.4	86.7	89.1
<i>T</i> -recordings	68.2	67	83.5	85.2	83.7

Table 2: Recognition accuracies (AED-ACC in %) in the two-source scenario

	All-combinations modeling	BSS	Proposed technique, only first step
<i>S</i> -recordings	80.1	77.6	80.4
<i>T</i> -recordings	65.6	63.8	75.6

by applying Equation 3 two consecutive times and leaving out the recognized class and the estimated position after the first time. Although speech is modeled as another acoustic event and recognized, only the system outputs corresponding to the eleven *AE*s are considered for computing the evaluation scores.

To evaluate both the recognition and the localization performance of our system, a *F*-score metric (harmonic mean between precision and recall) is used. That metric is defined as the one that was used for event detection in the *CLEAR'2007* [7] international evaluation.

4.2.1. Classification experiments

All the classification experiments are carried out for the two-source scenario. First of all, we present a test that was done assuming that the time end-points of the events are known, so it really consists in a classification (not detection) experiment. In that test, the *AE* classification performance of the proposed system from Figure 1 (one-step system) is compared with those from two reference systems, which use only one channel (microphone 18th). The first reference system uses a model for every possible combination of simultaneous events (all-combinations modeling). That is, the system has a model for each possible class whether it is an isolated event or a combination of events [24]. In our two-source scenario, this approach requires 11 models for the isolated *AE*s plus 11 models for the *AE*s overlapped with speech. Therefore, it does not need a prior separation of the two overlapped signals, but requires a number of models that may be too large in applications with a high number of classes or a high degree of simultaneity. Tests are carried out with a HMM-GMM based system [25] that was developed and tested with the same (*S*-recordings) database and accuracy metric used in this paper. The second reference system uses a blind source separation (*BSS*) approach that consists of a deflation-based iterative technique, where the source signals are extracted from the mixtures one by one [26, 27]. To reduce the time complexity for this *BSS* technique,

a quadratic contrast function with 4th order cumulants is used [27, 28]. In our experiments, we avoid the permutation problem of the algorithm by doing the right choice at the output of the separation block, before the classifier.

The acoustic event classification results are shown in Table 1. For this experiment, the performance of the various systems is measured in terms of classification accuracy, which is defined as the quotient between the number of correctly classified events and the total number of occurrences in the testing database. The results for *S*- and *T*-recordings are shown in two different rows in the table. Regarding the proposed joint system, two cases are considered using for the source positions either the flat or the estimated priors of source positions. The results for the *S*-recordings are also compared with those from a (virtual) version of the proposed system that assumes the source positions are known.

Notice from Table 1 that the proposed system shows a higher classification accuracy than the other two systems (all-combinations modeling and *BSS*-based), though a lower one than the system that assumes the source positions are known. As expected, the use of estimated position priors produces better classification accuracy for the proposed system than the use of flat priors. Logically, the classification accuracies are higher for the somewhat artificially-mixed signals from the *S*-recordings. When using the more realistic *T*-recordings, the performances from the two reference systems suffer a much significant degradation than the ones of the proposed technique. Concretely, the percentage of classification error reduction of the proposed technique with respect to the all-combinations modeling technique for the case of estimated priors is as high as 53.5%. Notice that, with estimated priors, the proposed system offers for *T*-recordings a higher accuracy than the system that uses knowledge of source positions. This apparent contradiction can be attributed to the fact that the experiment with known positions was carried out assuming that during recordings the two participants did not move from the room position initially assigned to each of them, but they actually

Table 3: Recognition accuracies (AED-ACC in %) for the two-source scenario

	Proposed technique (2-steps)		
	Flat priors	Estimated priors	Known positions
<i>S</i> -recordings	82.3	84.2	87.2
<i>T</i> -recordings	77.8	79.1	78.7

Table 4: Average localization accuracy (in %), assuming the time end-points of events are known

		<i>SRP – PHAT</i>	Proposed technique	
			Flat priors	Estimated priors
<i>S</i> -recordings	1 source	86.9	86.4	88.1
	2 source	70.9	83.9	85
<i>T</i> -recordings		82.3	85.6	86.2

moved due to the spontaneity requirement.

4.2.2. Recognition experiments

In the following, we report the recognition results, i.e. the ones obtained when the time end-points of the events are assumed unknown. As in Sub-section 4.2.1, all the experiments are carried out for the two-source scenario, but only the (non-speech) acoustic event is evaluated. First of all, only the first step of the proposed algorithm is tested to check how it is working. Its performance, which is measured in terms of the *F*-score (the *AED-ACC* metric [7]), computed from the *Viterbi*-decoding-based best channel is compared with those from the previously reported systems (all-combinations modeling and *BSS*-based). Notice from Table 2 that the proposed one-step recognition technique outperforms the *BSS*-based one and gets a similar result to that from the all-combinations modeling technique. Again, the degradation of the proposed algorithm is smaller than the ones from the other two techniques when moving to the more realistic *T*-recordings.

Table 3 shows the recognition results for the proposed two-step system, which can jointly estimate the *AE* identity and its cell-based position, using either flat or estimated priors of source positions. Like for the classification results in Table 1, here the performance of a reference system that assumes known source positions is also presented as reference. By comparing with the last column in Table 2, notice that the addition of the second step increases the recognition performance, especially when estimated priors are employed. That improvement is due to the fact that the two-step algorithm is effectively combining the likelihood scores from all the channels instead of choosing only the channel that gets the highest overall likelihood computed with the *Viterbi* algorithm. Again, the results from the *T*-recordings are worse than those from the *S*-recordings, but the proposed system shows a lower degradation. Actually, the error reduction percentage with respect to the all-combinations modeling technique for the *T*-recordings is 39.2% when using estimated priors. Again,

the lower score for the known positions case can be explained by the inadequacy of the assumption about steady participant position in those recordings. As it can be expected, a significant degree of degradation is observed going from known to unknown end-points, i.e. by including the automatic detection in the first step. In fact, by comparing Table 1 with Table 3 we notice that the loss in terms of relative increase of error from classification to recognition when estimated priors are used is 18.8% for *S*-recordings and 41.2% for *T*-recordings.

4.2.3. Localization experiments

Acoustic source localization is carried out at the event level for all the tested techniques, and using either one-source or two-source signals. As we did with the presentation of the recognition results, we firstly assume that the time end-points of the events are known. Consequently with that, the metric used is the localization accuracy, which is defined for a given class as the quotient between the number of correct localizations and the total number of event occurrences in the testing dataset. For a given event occurrence, a correct localization happens when the cell assigned to the true position is the same as the one estimated by the system. The true position for each event was obtained from visual inspection during signal recordings.

For the localization tests, the proposed system is compared with a reference *SRP-PHAT* system [11, 29, 30]. The *SRP – PHAT* localization technique explores the room space searching for the maximum of the global contribution of the *PHAT*-weighted cross-correlations from all the microphone pairs. To efficiently find the global maximum, the stochastic region contraction algorithm is used [30, 31] in our tests. *SRP-PHAT* works at the frame level so, in order to evaluate its results at the event level, we have averaged the estimated position coordinates along the whole event time interval (averaging gave us better results than the use of a voting procedure), which is assumed known.

The results for the case of known time end-points are pre-

Table 5: Average localization accuracy (in %) of the proposed joint approach

		Proposed technique	
		Flat Priors	Estimated priors
<i>S</i> -recordings	1 source	85.5	87.3
	2 source	83.6	84.6
<i>T</i> -recordings		84.7	85.4

sented in Table 4 in terms of average of accuracies along *AE* classes. As can be observed in that table, the joint recognition and localization system (the proposed technique) performs better in terms of localization accuracy than the *SRP-PHAT* system, both for *S*- and *T*-recordings. Notice again that the degradation from the one-source case to the two-source one is much smaller for the proposed technique. The error reduction in the two-source case with respect to *SRP-PHAT* is 48.4% for estimated priors, in spite that *SRP-PHAT* was implemented by looking at the two maxima of the sound map and counting a correct localization if anyone of the maxima was in the correct cell. Presumably, the inclusion of the time-spectral content of the *AEs* in the proposed system helps to better localize the sources, especially in the more complex two-source scenario. Table 5 shows the localization performance of the proposed system when the time end-points of the events are unknown and so they are estimated in the first step of the system. The performance is measured in terms of the *F*-score metric mentioned in Sub-section 4.2. The resulting scores are only slightly worse than those from the experiment reported in Table 4 with the system that uses knowledge of time end-points. Actually, the relative error increments are 6.7% for one source, 2.7% for two sources, and 5.8% for *T*-recordings. It is worth noticing that, in terms of localization performance, the system is less prone to end-point estimation errors than in terms of recognition performance.

5. CONCLUSIONS

A novel model-based approach for joint recognition and localization of simultaneously-occurring meeting-room acoustic events has been presented in this paper. Using multiple arbitrarily-located small microphone arrays and *HMM-GMM* models, a specific two-step system has been developed and tested. From the experiments, we observe that the system yields remarkably higher recognition scores than two common techniques, especially for the more realistic *T*-recordings. Although those two techniques do not require multichannel signals, their computational processing is not shared by a localization system as in the proposed technique, and furthermore, if all the classes can appear in all sources, the number of models required by the all-combinations modeling technique may be very high, since it is C^S , being *S* the number of simultaneous sources. Also, the event-based localization per-

formance of the proposed system is significantly better than that of the widely used *SRP-PHAT* method, especially in the two-source scenario, although the latter technique is able to work at the finer signal segment level. Actually, the proposed technique is more computationally expensive than the *SRP-PHAT* technique, but the required computations are shared with the recognition system. In conclusion, the proposed technique can be a proper alternative for localization in a multiple source scenario when it works together with an audio recognition system, since besides not requiring extra computations, it takes advantage of the sound characteristics in a multiple source scenario. Finally, the system is able to use prior probabilities of event class and room position, and we observed that both the recognition and the localization performance improve significantly by using the estimated priors of room positions. A further improvement can be expected also from using the event class probabilities whenever they are not flat. Notice that, since event sources will generally have different probabilities of occurrence at different room positions, the use of a joint class-position prior could further increase the system performance. In summary, the proposed approach shows the advantage of carrying out the two tasks, recognition and localization, with a single model-based system.

6. ACKNOWLEDGMENT

This work has been supported by the Spanish projects *SARAI* (*TEC2010 – 21040 – C02 – 01*) and *SPEECHTECH4ALL* (*TEC2012 – 38939 – C03 – 02*).

7. REFERENCES

- [1] DIRHA-Project, ,” <http://dirha.fbk.eu/>, 2014.
- [2] A. Brutti, M. Ravanelli, P. Svaizer, and M. Omologo, “A speech event detection and localization task for multiroom environments,” in *2014 4th Joint Workshop on Hands-free Speech Communication and Microphone Arrays (HSCMA)*, May 2014, pp. 157–161.
- [3] P. Giannoulis, A. Tsiami, I. Rodomagoulakis, A. Katsamanis, G. Potamianos, and P. Maragos, “The athena-rc system for speech activity detection and speaker localization in the dirha smart home,” in *Hands-free Speech Communication and Microphone Arrays (HSCMA)*, 2014 4th Joint Workshop on, May 2014, pp. 167–171.
- [4] Józef Kotus, Kuba Łopatka, and Andrzej Czyżewski, *Detection and Localization of Selected Acoustic Events in 3D Acoustic Field for Smart Surveillance Applications*, pp. 55–63, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [5] S. Goetze, J. Schroder, S. Gerlach, D. Hollosi, and J.-E. Appell, “Acoustic monitoring and localization for social care,” *Journal of Computing Science and Engineering*, vol. 6, no. 1, pp. 40–50, 2012.
- [6] Taras Butko, Cristian Canton-Ferrer, Carlos Segura, Xavier Giró, Climent Nadeu, Javier Hernando, and Josep R. Casas, “Acoustic event detection based on feature-level fusion of audio and video modalities,” *EURASIP J. Adv. Sig. Proc.*, 2011.
- [7] Andrey Temko, Climent Nadeu, Dušan Macho, Robert Malkin, Christian Zieger, and Maurizio Omologo, *Acoustic Event Detection and Classification*, pp. 61–73, Springer London, London, 2009.
- [8] Dimitrios Giannoulis, Emmanouil Benetos, Dan Stowell, Mathias Rossignol, Mathieu Lagrange, and Mark D Plumbley, “Detection and classification of acoustic scenes and events: an ieeea aasp challenge,” in *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. IEEE, 2013, pp. 1–4.
- [9] M. Brandstein and D. Ward, *Microphone Arrays: Signal Processing Techniques and Applications*, Springer, 1 edition, June 2001.
- [10] J. Dmochowski and J. Benesty, *Steered Beamforming Approaches for Acoustic Source Localization*, pp. 307–337, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [11] M. Omologo and P. Svaizer, “Use of the cross-power-spectrum phase in acoustic event location,” *IEEE Trans. on Speech and Audio Processing*, vol. 5, pp. 288–292, 1993.
- [12] C. Zhang, D. Florencio, D. E. Ba, and Z. Zhang, “Maximum likelihood sound source localization and beamforming for directional microphone arrays in distributed meetings,” *IEEE Transactions on Multimedia*, vol. 10, no. 3, pp. 538–548, April 2008.
- [13] H. Buchner, R. Aichner, J. Stenglein, H. Teutsch, and W. Kellennann, “Simultaneous localization of multiple sound sources using blind adaptive MIMO filtering,” in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2005, Philadelphia, USA*, 2005, vol. 3, pp. iii/97–iii/100 Vol. 3.
- [14] Rupayan Chakraborty and Climent Nadeu, “Real-time multi-microphone recognition of simultaneous sounds in a room environment,” in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2013, Vancouver, BC, Canada, May 26-31, 2013*, 2013, pp. 8672–8676.
- [15] Rupayan Chakraborty and Climent Nadeu, “Joint recognition and direction-of-arrival estimation of simultaneous meeting-room acoustic events,” in *INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France, August 25-29, 2013*, 2013, pp. 2948–2952.
- [16] Rupayan Chakraborty and Climent Nadeu, “Sound-model-based acoustic source localization using distributed microphone arrays,” in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2014, Florence, Italy, May 4-9, 2014*, 2014, pp. 619–623.
- [17] Lawrence Rabiner and Biing-Hwang Juang, *Fundamentals of Speech Recognition*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.
- [18] H. Ney and S. Ortmanms, “Dynamic programming search for continuous speech recognition,” *IEEE Signal Processing Magazine*, vol. 16, no. 5, pp. 64–83, Sep 1999.
- [19] Ludmila I. Kuncheva, *Combining Pattern Classifiers: Methods and Algorithms*, Wiley-Interscience, 2004.
- [20] Climent Nadeu, Dusan Macho, and Javier Hernando, “Time and frequency filtering of filter-bank energies for robust HMM speech recognition,” *Speech Communication*, vol. 34, no. 1-2, pp. 93–114, 2001.
- [21] S. J. Young, G. Evermann, M. J. F. Gales, T. Hain, D. Kershaw, G. Moore, J. Odell, D. Ollason, D. Povey, V. Valtchev, and P. C. Woodland, *The HTK Book, version 3.4*, Cambridge University Engineering Department, Cambridge, UK, 2006.
- [22] Rupayan Chakraborty, Climent Nadeu, and Taras Butko, “Source ambiguity resolution of overlapped sounds in

a multi-microphone room environment,” *EURASIP J. Audio, Speech and Music Processing*, vol. 2014, pp. 18, 2014.

- [23] Lucas C. Parra, “Steerable frequency-invariant beamforming for arbitrary arrays,” *The Journal of the Acoustical Society of America*, vol. 119, no. 6, pp. 3839–3847, 2006.
- [24] Andrey Temko and Climent Nadeu, “Acoustic event detection in meeting-room environments,” *Pattern Recogn. Lett.*, vol. 30, no. 14, pp. 1281–1288, oct 2009.
- [25] Taras Butko, Fran Gonzalez Pla, Carlos Segura, Climent Nadeu, and Javier Hernando, “Two-source acoustic event detection and localization: Online implementation in a smart-room,” in *Proceedings of the 19th European Signal Processing Conference, EUSIPCO 2011, Barcelona, Spain, August 29 - Sept. 2, 2011*, 2011, pp. 1317–1321.
- [26] C. Simon, Ph. Loubaton, and C. Jutten, “Separation of a class of convolutive mixtures: a contrast function approach,” *Signal Processing*, vol. 81, no. 4, pp. 883 – 887, 2001.
- [27] M. Castella, S. Rhioui, E. Moreau, and J. C. Pesquet, “Quadratic higher order criteria for iterative blind separation of a mimo convolutive mixture of sources,” *IEEE Transactions on Signal Processing*, vol. 55, no. 1, pp. 218–232, Jan 2007.
- [28] M. Castella and E. Moreau, “A new optimization method for reference-based quadratic contrast functions in a deflation scenario,” in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, April 2009, pp. 3161–3164.
- [29] Cha Zhang, Dinei Florencio, and Zhengyou Zhang, “Why does phat work well in low noise, reverberative environments?,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*. March 2008, IEEE.
- [30] Hoang Do, Harvey F. Silverman, and Ying Yu, “A real-time srp-phat source location implementation using stochastic region contraction(src) on a large-aperture microphone array,” in *ICASSP (1)*. 2007, pp. 121–124, IEEE.
- [31] M. F. Berger and H. F. Silverman, “Microphone array optimization by stochastic region contraction,” *IEEE Transactions on Signal Processing*, vol. 39, no. 11, pp. 2377–2386, Nov 1991.