



HAL
open science

An evaluation framework for event detection using a morphological model of acoustic scenes

Mathieu Lagrange, Grégoire Lafay, Mathias Rossignol, Emmanouil Benetos,
Axel Roebel

► **To cite this version:**

Mathieu Lagrange, Grégoire Lafay, Mathias Rossignol, Emmanouil Benetos, Axel Roebel. An evaluation framework for event detection using a morphological model of acoustic scenes. 2015. hal-01111381v1

HAL Id: hal-01111381

<https://hal.science/hal-01111381v1>

Preprint submitted on 30 Jan 2015 (v1), last revised 28 Sep 2016 (v2)

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.

An evaluation framework for event detection using a morphological model of acoustic scenes

Mathieu Lagrange, *IEEE memberships*, Grégoire Lafay, Mathias Rossignol, *IEEE memberships*, Emmanouil Benetos, *Member, IEEE*, Axel Roebel, *IEEE memberships*,

Abstract—This paper introduces a model of environmental acoustic scenes which adopts a morphological approach by abstracting temporal structures of acoustic scenes. To demonstrate its potential, this model is employed to evaluate the performance of a large set of acoustic events detection systems. This model allows us to explicitly control key morphological aspects of the acoustic scene and isolate their impact on the performance of the system under evaluation. Thus, more information can be gained on the behavior of evaluated systems, providing guidance for further improvements. The proposed model is validated using submitted systems from the IEEE DCASE Challenge; results indicate that the proposed scheme is able to successfully build datasets useful for evaluating some aspects the performance of event detection systems, more particularly their robustness to new listening conditions and the increasing level of background sounds.

Index Terms—acoustic event detection, auditory scene analysis, experimental validation.

I. INTRODUCTION

OVER the past decades, the amount of audio data in our sonic environment have considerably grown. Recent research fields such as eco-acoustics [1], [2] start to massively record environmental sounds around the world in order to measure potential animal biodiversity modification over large temporal scales due to human activity or climate change [3]–[5]. Other research fields focus on human activities for context inference and surveillance [6]–[9].

If research on Automatic Speech Recognition (ASR) [10] and Music Information Retrieval (MIR) [11] are now well established, research addressing automatic analysis of complex environmental acoustic scenes remains relatively young. In particular, those open research avenues suggest a large range of experimentation in order to 1) gain knowledge about the important characteristics of those acoustic scenes and how they can be modeled, 2) propose new algorithmic approaches to contribute to the above cited applications area: eco-acoustics and urban sensing. Being relatively new research fields, only few data sets are available, though this number may grow as the interest of scientific and engineering communities for such tasks increases, see [12] and [13] for research effort in human environments related tasks.

This paper focuses on building an evaluation framework for the task of detecting events of interest in acoustic scenes using simulated data. It builds upon the IEEE AASP Challenge on

Detection and Classification of Acoustic Scenes and Events (DCASE), which was organised by the Centre for Digital Music of Queen Mary University of London and by the Institute for Research and Coordination in Acoustics/Music (IRCAM), under the auspices of the Audio and Acoustic Signal Processing (AASP) technical committee of the IEEE Signal Processing Society in 2013 [14]. The DCASE Challenge is the second challenge dedicated to this task after the CLEAR Challenge [15].

During the formal definition of the event detection task of this challenge, besides important questions about evaluation metrics, an interest rose about the potential benefit of considering simulated data to enlarge the scope of evaluation of the submitted systems. Varying the power level of the background, the density of the events, their intra class diversity, all seemed important aspects to would be desirable to study though costly to tackle with recorded and annotated data. To this end, a simulation protocol was needed, which would be based on a morphological model of environmental acoustic scenes. As discussed in details in Section VI, we acknowledge that the use of simulated data shall not be considered as sufficient for the final evaluation of engineering systems. That being said, the above described potential benefits are still sufficient to justify pursuing that avenue of research.

The aim of the morphological model proposed in this paper is to generate acoustic scenes as a “skeleton of events on a bed of texture” [16]. As the final use of the simulated scenes are to be analyzed by event recognizers trained on recorded data, one shall minimize both the discrepancy between the simulated scenes and recorded ones and its potential impact. Thus, we do not consider approaches based on actual synthesis of sounds. It thus departs significantly from models used in research fields such as wave field synthesis [17], binaural or spatial scene synthesis [18], acoustic event synthesis [19] and texture synthesis [20]–[22].

The proposed model is based on several sequences of sound events issued from the same source, where each sound event is drawn from a collection of carefully chosen sound samples. The morphological aspects of the scene, *i.e.* which sound sample is played at what time and which level, are then modeled in an abstract manner, allowing us to control high level properties of the scene. The contribution of this paper is threefold: 1) propose a computational model for the generation of simulated data sets, 2) motivate important morphological aspects of the model based on perceptual considerations, and 3) consider this simulation paradigm to gain knowledge about the behavior of several event detection systems developed by

different research teams worldwide initially submitted to the DCASE challenge [12].

To this end, Section II motivates some design choices, and details the structure of the so called “sound collections”, that is, the input data of the simulation process. Section III presents the proposed model of acoustic scenes which underlines the simulation process. Sections IV and V present the evaluation framework for event detection systems using simulated acoustic scenes. Then, the use of simulated data to evaluate detection algorithms is discussed in Section VI.

II. THE NOTION OF SOUND COLLECTION

A. Auditory scene as a sum of sound sources

As a simulation process couldn’t practically deal with each acoustic event that may occur in an acoustic scene separately, the proposed model adopts a “source-driven” approach by considering an acoustic scene as a sum of sound sources. This approach is consistent with the way humans perceive their sonic environment. Studies addressing the Auditory Scenes Analysis (ASA) [23] problem, and more specifically the sound segregation process [24] [25] [26] [27], show that humans make sense from their sonic world by isolating information related to individual sound sources. Considering a bottom-up approach, the segregation process relies on generic rules involving Gestalt-like principles [27] to group sounds with similar acoustic indicators (common onset, spectral regularity and harmonicity), as well as similar perceptual attributes (timbre, loudness, perceived location and pitch) into perceptual entities called “auditory streams”. Recently, several neurophysiological studies have shown evidence of the existence of auditory streams [28].

Besides ASA studies which mostly consider pure tones or simple complex sounds [25], more recent studies adopting a psycho-linguistic approach to describe recorded sounds, have also demonstrated the existence of top-down source-driven grouping processes involved in sound perception. Investigating the qualitative evaluation of urban acoustic scenes using categorization tasks and linguistic analysis, studies of Dubois and colleagues [29] [30] have shown that listeners categorize sound environments on the basis of semantic features, that is the meaning attributed to the recalled sound sources.

Considering both the ASA and the psycho-linguistic approach, it seems intuitive for the simulation process to consider separately the sound activity of each sound source of the scene. In practical terms, to materialize these sound activities, each sound source has to be related to a collection of sound recordings. But this approach introduces fundamental questions about the very nature of such a collection. It first questions the existence of a standardized taxonomy of sounds. Such taxonomy must be a hierarchical classification system putting together sounds according to their shared characteristics. Each group must be labeled in a way that a specific name may describe its corresponding class, an instance of it, but also at which level of the classification it fits. Unfortunately, if such systems exist for plants, animals or colors, it is not the case for sounds [31]. Main reasons are:

- Sound description and identification are highly subjective. In other words, a same sound may be described quite differently according to the subject. This is due to the relative lack of basic lexicalized terms to describe acoustic phenomena [32]
- Sound description and identification are highly context dependent, that is, sound source identification depends on the nature of the other co-occurring sound sources [27], [33], [34]

Even if there is no systematic way to build a sound collection, one may take into account some perceptual considerations to guarantee a certain level of ecological validity. Those considerations are addressed in the next sections.

B. Action and Sources

Event detection tasks evaluate if an algorithm is able to detect a specific set of sound classes. Ideally, to prevent from low generalization capability, the training set of a given class shall be consistent, that is, class exemplars should be representative of the diversity of the sounds suggested by the class label that may occur in the real world. In our case, the class exemplars are the recordings of a sound collection.

Some perceptual considerations can be taken into account to guide the collection building process. First, one may look at the way humans classify / categorize sounds. As explained by [35], “Categorization is a cognitive process that unites different entities of an equivalent status”. Among other categorization strategies, several studies show that humans categorize sounds according to 1) the type of source (agent, object, functions) and / or 2) the action / movement causing the sounds [27], [29], [36]–[39].

Human categorization occurs at several levels. Rosch [40] proposed three levels of categorization for real-world objects namely superordinate, basic, and subordinate. The higher the level, the higher is the abstraction degree of the categories. Considering sound perception, Guyot et al. [36] proposed a framework where listeners identified sound categories of abstract concepts at a superordinate level (noise generated by a mechanical excitation), action at the basic level (grating, scratching, rubbing) and source at the subordinate level (Dishes, Pen sharpening, Door). Although Houix et al. [35] found some differences by showing that sounds seem “to be categorized as sound sources first and only second as actions”, it appears that source and action are adequate verbal descriptors for category.

One way to make a sound collection consistent is to consider low-level categories as the intra-category diversity decreases with the level. Considering that, one may label a sound collection using a couple “source-action” (*passing-car*), or at least one of the two, in order to minimize the expected diversity of its recordings. Any name referring to higher category levels may lead to sound collections comprising a too large variety of objects. Such definition of collection then raise two issues:

- building such collections would suppose the availability of a large number of recorded sounds to be representative of the diversity suggested by the collection label;

- adopting a data-centered approach, such collections may lead to a misinterpretation of the results of a detection task for someone who did not build them, as the nature of the entities suggested by the collection labels are ambiguous (ex: a sound collection of *traffic sounds* vs. a sound collection of *passing-car*).

Considering the source-action couple is not sufficient. The experimenters must also choose generic labels for the couple. To do so, one may refer to the work of Gaver [41] who proposed a phenomenological taxonomy of everyday sounds, the work of Niessen et al. [31] who assessed the consensus of categories mentioned in 166 papers of different research domain using linguistic analysis, and recently, the work of Salamon et al. [13] who built a taxonomy of urban sounds based on the work of Brown et al. [42].

This section shows evidence that labeling a class using the source-action nomenclature helps us to reduce the expected intra-class diversity. However, it does not address the issue of inter-class diversity. Indeed, a naive source-driven approach supposes to record in a source-wise way all the sound activities that may occur in an environment. Considering dense environments such as cities or forest, this may raise important practical issues. To circumvent this problem, one may assume that all the sources do not carry the same potential information, and are not required to be recorded separately.

C. Texture vs. Event

The human brain may easily distinguish between a voice sound and a background of other competing sounds [26]. Considering the example of an urban acoustic scene, global traffic hubbub sounds are typically uninformative, compared with closer human sounds [43].

Maffiolo [44] showed the existence of two distinct cognitive processes depending on the listener's ability to identify separate sound events. By asking subjects to categorize recordings of urban environments and using linguistic analysis of the verbal descriptions of the categories, she found two cognitive categories of sound environments called respectively "event sequences" and "amorphous sequences". Event sequences (sound environments in which distinct events or sequences of events can be identified) are processed analytically, that is, based on the meaning of the identified sound sources, whereas amorphous sequences (sound environments in which no event can be isolated) are processed holistically using global acoustical indicators (intensity, spectral content). The distinction observed by Maffiolo was validated by Guastavino [32]. Using semantic analysis of verbal descriptions of specific sounds populating the urban environment, Guastavino showed that verbal descriptions of low pitched sounds may be divided into two categories called "source events" (sound events which can be attributed to a sound source), and "background noise" (where no identifiable event can be isolated).

What comes out from these studies is that sound perception highly depends on semantic features (source identification), but also on the informativeness of the isolated source. Sound sources that carry information of interest are processed separately, whereas the other are processed together in a single stream.

Based on this notion of informativeness, another common distinction is made between two perceptual objects called "sound events" and "sound textures". Based on previous studies on vision, McDermott and Simoncelli [21], [45] showed that the perception of sound textures may derive from simple statistics of early auditory representations. These summary statistics would be sufficient to recognize sounds having some temporal homogeneity.

That said, there are few formal definitions concerning the texture object [20]. The most notable attempt has been made by Saint-Arnaud [46] and Saint-Arnaud and Popat [47]. From their experiment, they derived the following properties (quoted from [46]):

- *Sound textures are formed of basic sound elements, or atoms;*
- *atoms occur according to a higher-level pattern, that can be periodic or random, or both;*
- *the high-level characteristics must remain the same over long time periods (which implies that there can be no complex message);*
- *the high-level pattern must be completely exposed within a few seconds ("attention span");*
- *high level randomness is also acceptable, as long as there are enough occurrences within the attention span to make a good example of the random properties.*

Considering these properties, a texture may be understood as a composite object with two hierarchical levels, the top level being the high level pattern, and the leaf level being the atom. The nature of an atom remains adaptable as the latter may be considered at several time scales. Thus and to some extent, texture may be considered as a concatenation of recordings, each of them being a sequence of atoms. In this case, these recordings must comprise at least the high level pattern of the texture, that is, if we consider a texture of 'gallop', recordings of atom sequences must be at least composed of the first three sounds of hoofs.

To summarize the previous statements, it appears that all sounds are not processed as a sum of distinct events:

- amorphous sequences that convey low semantic information are processed holistically;
- sound textures with stable acoustic properties over long period are processed using summary statistics of these acoustics properties.

To circumvent the issue of recording a representative number of sound collections to simulate an acoustic scene, one can take into account those considerations, and use recordings of mixed sound sources, provided that they are amorphous sequences or textures. We believe that there exist some links between the notions of amorphous sequences and textures. Both trigger holistic processing based on global acoustical properties for amorphous sequences [29], [44] and summary statistics for textures [45], and both convey a low potential information content [46]). Yet, amorphous sequences are described as "background sounds" with no identifiable events,

whereas the texture definition comprises sequences of events such as “gallop” that do not meet this last criterion. Considering that, one can consider an amorphous sequence to be a texture, as the physical characteristics of an amorphous sequence remain stable over time, but the reverse is not systematic.

D. Definition of a sound collection

From the considerations discussed above, we derive two types of sound collections to be used as basic elements by the simulation process: the “event collections” and the “texture collections”. For both collections, a stream is modeled as being a temporal sequence of sound recordings coming for the same sound collection. For the texture collection, each recording is an atom sequence, or more precisely, a sequence of sound events which follow a periodic or a stochastic pattern. The nature of the sequence to be recorded depends on the type of texture considered. For a texture with a periodic pattern such as gallop, recordings are event sequences comprising at least the first three sounds of hooves. And for a texture with a stochastic pattern such as “rain”, the recordings are simply samples of rain sounds. This method offers a certain flexibility, as it makes it possible to quickly generate various versions of a same texture with few recorded samples, by varying the apparition order of the sequence. Obviously for a texture to be realistic, sequences have to come from the same recording session. Moreover, as the human brain is very sensitive to repetition of identical sounds, even when they are individual chunks of white noise [48], a sequence shall not be concatenated with itself.

To summarize, the proposed source-driven model uses collections as basic element for the simulation process:

- Each collection is a group a sound recordings.
- Insofar as possible, the label of the collection should be of the form “source + action” and labels of source and action must be generic.
- There are two types of collections called respectively the event collections, and the texture collections.
- Sound recordings of a same event collection come from the same sound source.
- Sound recordings of a same texture collection are atomic sequences emitted by one or a mixture of sound sources.
- Sound recordings of a texture collection must at least comprise the high-level pattern of the texture (*e.g.* three sounds of hooves for the gallop texture).
- A texture built from the concatenation of recordings must convey a low semantic information and / or have stable acoustic properties over time.

III. PROPOSED MODEL

A. Design choices

Building on the above discussed matters, the proposed simulation process considers an acoustic scene as a sum of sound sources. Each sound source activity is symbolized as a semantic sound track, which is a sequence of acoustic samples all emitted by the considered sound source (see Figure 2).

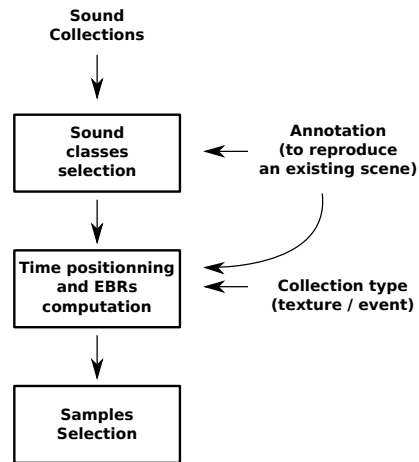


Fig. 1. Schematic of the simulation process.

To generate each semantic sound track, the model takes into account a set of four parameters being respectively:

- 1) the mean / variance of the Event to Background power Ratio (EBR) between acoustic samples
- 2) the mean / variance time interval between consecutive onsets of acoustic samples
- 3) the mean / variance duration between acoustic samples
- 4) the start / end times of the track

As motivated in Section II-C, the model distinguishes between sound events and texture. A track of events is made of discrete sound samples, whereas a track of texture consists of one continuous sound, or a concatenation of samples (see Figure 2). Thus, for texture track, the mean/variance time interval between samples as well as the variance EBR are set to 0.

Each semantic track, texture or event, is related to a specific sound collection. As discussed in Section II a sound collection may be seen as a group of similar recordings, each of which comprising sound signals that are emitted by the same sound source. For the purpose of this study, the notion of sound collection greatly overlaps the notion of sound class, as this term is understood when tackling automatic detection tasks.

The resulting simulation model is depicted on Figure 1. First, the experimenter selects a number of sound sources or class to be used, each of which being related to a specific sound collection. Second, the experimenter sets the simulation parameters depending on the nature of the track (event or texture). Those parameters can also be estimated from pre-existing annotated recorded sound scenes. According to those parameters, the simulation process computes the number of samples used in each semantic track. Lastly, samples are randomly drawn from the corresponding sound collection using a discrete uniform distribution.

B. Formal Definition

The proposed model is source-driven as it uses as basic elements semantic sound tracks, gathering sounds coming from the same collection of either sound events or sound textures (see Figure 1). The nature of the recordings depends

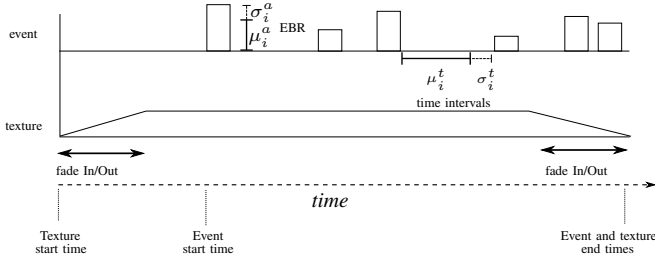


Fig. 2. Two semantic sound tracks (event and texture) and their controlling parameters.

on the type of collection (event of texture) which is considered. After selecting the classes to be used, the putative recorded samples are sequenced to generate the sound environment. The sequencing process depends on the type of collection. Ideally, the sound collection design has to fulfill some perceptual constraints for the simulation to be ecologically valid, *e.g.* to produce realistic scenes, as described in Section II.

Considering that $s(n)$ is a given acoustic scene composed of C sound classes c_i , the proposed model is such that:

$$s(n) = \sum_{i=1}^C t_i(n) \quad (1)$$

Where each n is the time index, t_i is a semantic sound track. For the sake of simplicity, we only detail here the model of an event track, then explain the adaptation of the model to texture tracks.

t_i is defined as a sequence of n_i sound events $e_i^k(n)$ randomly chosen among the $|c_i|$ samples in class c_i : for each k in $[1..n_i]$, $e_i^k = c_{i,\mathcal{U}(1,|c_i|)}$, where $\mathcal{U}(x,y)$ represents an uniformly distributed integer random value between x and y included. Each event is scaled by an amplitude factor sampled from a real normal distribution with average μ_i^a and variance σ_i^a . The interval separating the onset times of consecutive samples for track i is, similarly, randomly chosen following a normal distribution with average μ_i^t and variance σ_i^t . Formally, each sequence t_i is thus expressed as:

$$t_i(n) = \sum_{j=1}^{n_i} \mathcal{N}(\mu_i^a, \sigma_i^a) c_{i,\mathcal{U}(1,|c_i|)}(n - n_i^j) \quad (2)$$

$$n_i^j = n_i^{j-1} + \mathcal{N}(\mu_i^t, \sigma_i^t) \quad (3)$$

where n_i^0 is set to 0 by convention. The signal of an event is defined in such a way that $e(n) = 0$ if $n < 0$ or beyond the signal's duration.

In the case of a texture track, two implementation differences must be observed to maintain a perceptually acceptable output: first, signal amplitude is only drawn at random once, and that value is applied to all samples; second, sample start times are not randomized but chosen so that the texture recordings chosen from class c_i will be played back-to-back with sufficient overlap to create an equal-power cross-fade between them, thus generating a continuous, seamless track.

While implementing this model for the generation of acoustical sound scenes, additional treatments and constraints are

applied in order to improve perceptual quality. Namely, the fading of the onset and offsets of events and the whole track, and the fact that the same sound sample cannot be sequenced consecutively.

IV. CORPUS SIMULATION

This section describes the different corpora of simulated acoustic scenes considered in the experiments described in Section V. All the scenes are simulated using the DCASE challenge test set annotations for the ‘Office Live’ (OL) task [14]. We run the same automatic event detection algorithms used for the DCASE challenge, and compare the results obtained with the simulated scenes to those obtained with the real scenes of the DCASE test set.

The root corpus is the test set considered in the DCASE challenge. It is called “test-QMUL”. This corpus is composed of 11 recordings of office live scenes roughly one minute long. Scenes have been recorded in 5 different acoustic environments. The audio events have been divided into 16 sound event classes to be annotated: door knock, door slam, speech, human laughter, clearing throat, coughing, drawer, printer, keyboard click, mouse click, object (specifically pen, pencil or marker) put on table surfaces, switch, keys (put on table), phone ringing, short alert (beep) sound and page turning. Two different annotations coming from two distinct individuals have been used to measure the algorithm performances, thus leaving us with 22 scene-annotator couples. There is no time overlap between events.

Four corpora of simulated scenes are generated as depicted in Figure 3. They are respectively called “instance-QMUL”, “abstract-QMUL”, “instance-IRCCYN” and “abstract-IRCCYN”. The labels “IRCCYN” and “QMUL” refer to the two different datasets of event recordings used to generate the corpora which have been recorded in different offices, the ones of Queen Mary University of London (QMUL) for the former and the ones of the Institute of Research on Communications and Cybernetics of Nantes (IRCCYN). The labels “instance” and “abstract” correspond to two distinct simulation processes.

To generate the two QMUL corpora, we use recordings of audio events that have been extracted from recordings done during the preparation of the DCASE challenge, but unused during the challenge, see [12] for further information on recording conditions. The extracted samples were therefore recorded in the same conditions than the test-QMUL corpus.

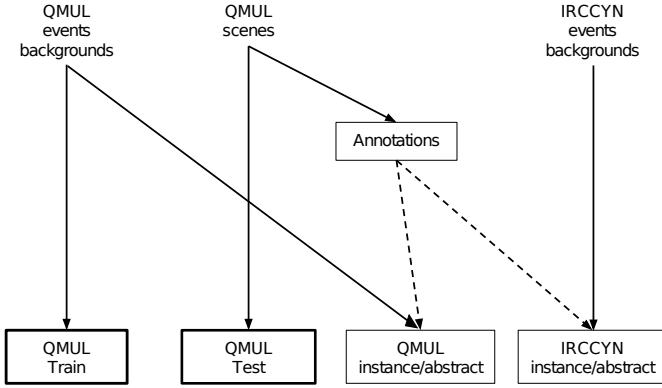


Fig. 3. Generation process of the corpora considered in this evaluation. As part of the DCASE challenge, systems were trained on QMUL Train and tested on QMUL Test during the DCASE challenge.

Depending on the sound class considered, 3 to 23 events per class are extracted. We also use event-free background recordings (texture) coming from the same acoustic environments than those of the test-QMUL corpus. These background recordings are used to generate the background noise (texture) of the instance-QMUL and abstract-QMUL corpora.

The two IRCCYN corpora are generated using new recordings of sound events with respect to the sound classes of the DCASE challenge. All recordings were performed at IRCCYN in a calm environment using the shotgun microphone AT8035 connected to a ZOOM H4n recorder. 20 samples of each class are used to generate the instance-IRCCYN and abstract-IRCCYN corpora, which corresponds to the cardinality of the DCASE challenge train set in terms of event classes [12].

A. “Instance” simulation process

The instance simulation process simulates acoustic scenes with the same temporal structure and Event to Background Ratios (EBRs) than the annotation of the test-QMUL corpus. The EBR of an event of N sample length is obtained by computing the ratio in decibel between the event E_{rms} and the background B_{rms} root means square measures.

$$EBR = 20 \log_{10} \left(\frac{E_{rms}}{B_{rms}} \right) \quad (4)$$

with

$$X_{rms} = \left(\frac{1}{N} \sum_{n=1}^N x(n)^2 \right)^{1/2}$$

$x(n)$ may be replaced by $e(n)$ and $b(n)$, the sound pressures at sample n of respectively the sound event and the background noise.

For each event of each scene-annotator couple of the test-QMUL corpus, the onset-offset times and an approximation of the EBR are considered. As it is not possible to isolate the background under the events, the background level needed to

compute the EBR is obtained using a event-free sequence of each real scene. These onsets-offsets and EBR are then used to generate the simulated scenes. For each simulated scene, at each onset of the corresponding annotator-couple scene, we randomly place an audio event belonging to the same audio class. To ensure that samples of recorded audio events are not too long comparing to the annotated ones, recordings are cut off to the annotation length if the recording duration is larger than the annotation duration of at least 0.5 seconds.

Each event has its amplitude scaled to the same EBR than the test-QMUL corpus. Instance simulation process provides us with simulated scenes with temporal structures and sound levels that are close as possible as those of the real corpus test-QMUL.

B. “Abstract” simulation process

For the abstract simulation process, the goal is to abstract temporal structures and EBRs of the real scenes. To do so, the model described in III-B is instantiated using estimations of the μ_i^a , σ_i^a , μ_i^t and σ_i^t parameters (see eq. 2 and 3). Estimation is done for each annotator-scene couple, using both the sound signals and the annotations of the test-QMUL corpus. To generate the simulated scenes, EBRs and time intervals between events are respectively obtained from the Normal distributions $\mathcal{N}(\mu_i^a, \sigma_i^a)$ and $\mathcal{N}(\mu_i^t, \sigma_i^t)$.

Similarly to the instance simulation process, event recordings are chosen randomly. For practical considerations, the start and termination times of the class sequence (semantic sound track) are the same as the ones of the test-QMUL corpus. To ensure that the recorded samples are not significantly longer compared to the annotation times, the sample duration of a considered sound class i has its duration D thresholded as follows: $D - \mu_i - \sigma_i > 5$, with μ_i and σ_i being respectively the average and standard deviation of the duration of the events belonging to the class i in a given annotation. Setting the lower bound to 5 seconds allows us to minimize the impact of such operation on short impulsive sounds.

V. EXPERIMENTS

A. Evaluation Metric

The performance of event detection systems can be evaluated following several metrics. In order to improve legibility of the following, we shall retain one evaluation metric that is considered to be the most informative for our study.

Among the four metrics considered in the DCASE Challenge [12], namely the Acoustic Event Error Rate (AEER) [15], the Precision, Recall, and F-measure, the F-measure is selected as the most common and interpretable one.

Another variation is that those metrics can be computed over each frame or on event boundaries. In the latter case, the detection of the onset boundary can be considered solely or together with the offset. As annotating and consequently detecting the duration and the offset of events is notoriously difficult, we focus on the detection of the onset as the main objective. Furthermore, in order to achieve more comparable results across datasets and to ensure that repetitive events do

TABLE I
SUMMARY OF SUBMITTED EVENT DETECTION SYSTEMS.

System	Method
CPS [51]	Segmentation - Likelihood ratio test classification
DHV [52]	MFCCs (features) - HMMs (detection)
GVV [53]	NMF (detection) - HMMs (postprocessing)
NVM [54]	Hierarchical HMMs + Random Forests (classification)
NR [55]	MFCCs (features) - SVMs (classification)
SCS [56]	Gabor filterbank (features) - HMMs (classification)
VVK [57]	MFCCs (features) - GMMs (detection)
Baseline [12]	NMF with learned bases (detection)

not dominate the accuracy of an algorithm, the metric shall be class normalized. That is:

$$f = \frac{1}{C} \sum_{i=1}^C f_i \quad (5)$$

where f_i is the F-measure achieved by the system while detecting event i . Thus, by considering the Class-Wise Event onset based F-measure (CWEBF), performance evaluation is more invariant to event duration and distribution. We thus select this metric that was also collectively agreed upon by DCASE participants through the challenge mailing list.

B. Detection systems

Together with a Baseline system provided by the organizers, 8 detection systems have been evaluated at the DCASE challenge. Those systems roughly follow the processing chain shown on Figure 4 with some variety on the implementation of the different nodes. Features are most commonly Mel-Frequency Cepstral Coefficients (MFCC)s [49] but other sets of spectral features are also considered, with or without pre-processing such as denoising. The classifier of choice is the 2 layer Hidden Markov Model (HMM) [50] where the second layer models the transition between events but other classifiers are also considered such as Random Forests (RF), Support Vector Machines (SVM) or Non-negative Matrix Factorization (NMF).

All those algorithmic differences as well as their specific tuning result in specific behaviors that are interesting to evaluate in different testing conditions, especially those which evaluate their generalization capabilities.

C. Datasets

Five corpora called respectively test-QMUL (described in Section IV), instance-QMUL, abstract-QMUL, instance-IRCCYN and abstract-IRCCYN are used for the evaluation. This section details the corpus specificities. Figure 3 illustrates the corpora generation process.

All corpora with the labels “QMUL” or “IRCCYN” are obtained using respectively the QMUL or IRCCYN event datasets (see section IV). Similarly, all corpora with the labels “instance” or “abstract” are obtained using respectively the instance or abstract simulation process (see Sections IV-A and IV-B resp.).

The instance-QMUL corpus is composed of 4 sub-corpora called respectively “insQ-EBR_6”, “insQ-EBR_0”, “insQ-EBR_-6” and “insQ-EBR_-12”. For the insQ-EBR_0 sub-corpora, EBRs of test-QMUL scenes are preserved. To measure

dataset	testQ	insQ	absQ
Baseline	9.0±4.8	10.5±3.0*	9.9±3.5
CPS	0.7±0.8	0.8±1.3	0.8±1.4*
DHV	30.7±8.4	34.5±7.5*	34.0±7.9
GVV	13.2±8.0	15.0±6.4*	14.6±6.2
NR	21.5±6.5*	6.8±5.7	7.4±5.8
NVM	28.2±5.9*	9.7±9.6	10.8±9.9
SCS2	41.5±7.6*	39.3±8.2	39.4±8.2
VVK	24.6±6.8*	19.7±8.7	19.2±9.2

TABLE II
RESULTS OF THE EVALUATED SYSTEMS ON THE QMUL DATASETS. RESULTS IN BOLD ARE EQUIVALENT PER ROW (PAIRED T-TEST AT 0.05 SIGNIFICANCE LEVEL) TO THE BEST PERFORMANCE PER ROW (DEPICTED WITH A *).

the impact of different EBRs on the algorithm performances, we generate three other sub-corpora (insQ-EBR_6, insQ-EBR_-6 and insQ-EBR_-12) by adding three offsets, one for each sub-corpora, to the EBRs of the test-QMUL scenes. The offsets are of +6dB, -6dB and -12dB.

For all the sub-corpora of the instance-QMUL corpus (insQ-EBR_6, insQ-EBR_0, insQ-EBR_-6 and insQ-EBR_-12) as well as the other corpora (abstract-QMUL, instance-IRCCYN, abstract-IRCCYN), each scene is simulated 10 times, each time using different recording instances. Each of these corpora / sub-corpora is composed of 220 simulated scenes (22 * 10) corresponding to the 22 scene-annotator couples of the test-Q scenes replicated 10 times.

D. Results on QMUL datasets

Granted with permission of the authors of the submitted systems, we ran the above described systems on the simulated datasets on the same computing servers as the ones used for the challenge with the same computing environment. Also, a rerun of the systems over the QMUL Test set has been done in order to ensure replication of the published results.

Table II shows the class-wise event based F-measure in percent achieved by the evaluated systems over the QMUL Test set and the simulated sets QMUL Instance and QMUL Abstract. The baseline, CPS, GVV and SCS systems performed equivalently across the 2 datasets. The DHV system performed better, but not by a significant margin. The VVK, NVM, and NR2 systems have their performance decreased, by a significant margin for the latter two. The CPS system submitted to the DCASE Challenge had an implementation issue that prevent it to run correctly at the time of the challenge, giving poor results that are consistently replicated over the simulated datasets. For this reason, the CPS system will not be discussed further in the remaining of the paper. Leaving aside the NR2 and NVM systems, the ranking of the systems are same for the 3 datasets. This result comfort us with the use of such simulation scheme to replicate and extend evaluation results achieved on recorded and annotated datasets.

We shall now investigate further the reasons explaining the behavior of the NVM and NR2 systems. In test mode, both systems first compute features and then run a classifier on them. Therefore, features were first checked for inconsistent values. The minimal and maximal values did not change




Fig. 4. Schematic of event detection systems (nodes with a * are not systematically used). Below, state-of-the-art design choices are given as examples.

System	testQ	insQ	absQ
Baseline	3.14 (drawer)	8.63 (drawer)	7.40 (drawer)
CPS	2.66 (knock)	9.04 (doorslam)	7.84 (doorslam)
DHV	8.44 (drawer)	6.88 (drawer)	8.01 (keyboard)
GVV	3.08 (pageturn)	3.78 (pageturn)	3.55 (pageturn)
NR	4.33 (keyboard)	25.35 (doorslam)	20.68 (doorslam)
NVM	1.26 (laughter)	22.48 (cough)	19.22 (cough)
SCS	1.18 (alert)	2.70 (drawer)	1.72 (doorslam)
VVK	1.81 (alert)	8.73 (doorslam)	8.20 (doorslam)

TABLE III

MAXIMAL NUMBER OF FALSE POSITIVE AVERAGED OVER SCENES AND CORRESPONDING EVENT.

across datasets, and the distribution of the features are indeed different across datasets but not by a large margin.

Close inspection of the inter-class confusion matrices for each systems revealed that for the two systems the classification node may be responsible for this performance. Indeed, one event is triggered almost all the time which drastically increase the false alarm rate. This explains for a large part the decrease of performance of the NR and NVM systems. This behavior can easily be seen on Table III which displays the maximal number of false positive averaged over scenes of each datasets and their corresponding event; on the simulated datasets, the doorslam event for NR and cough event for NVM are falsely triggered very often.

We conclude that this decrease of performance is most probably not due some potential synthesis inconsistencies produced by the simulation process but more due to an over fit of the classification node. Considering that both systems are the only submissions based on discriminative approaches, SVMs and RFs for the NR2 and NVM respectively, we may conjecture that the training framework of the DCASE challenge is not well suited for such classification schemes.

System performance when the EBR is varied is now studied. Figure 5 shows their performance, where 0 dB of EBR roughly corresponds to the EBR level of the QMUL Test set. As expected, most of the systems have their performance decreasing with respect the decrease of the EBR. The ranking is preserved, and the spread between the 3 lowest performing systems greatly reduces with respect the decrease of the EBR. The only system that does not follow this trend is the SCS system, which maintains a stable performance across all EBR ranges. This may be due to an effective signal enhancement which is an important pre processing node of this system [56].

E. Results on IRCCYN datasets

When tackling a classification task, an important issue is whether the classification system under evaluation is able to generalize to unseen data whose annotation is consistent with the one used for training and tuning.

To evaluate this generalization capability, it is useful to consider results achieved by the systems on the IRCCYN datasets, where the background and events are recorded in




Fig. 5. Class wise event based F-measure in percent achieved by the systems on the QMUL instance datasets with varying EBR.

dataset	testQ	insI	absI
Baseline	9.0±4.8*	5.9±2.9	5.6±2.9
DHV	30.7±8.4*	10.0±5.8	9.5±5.6
GVV	13.2±8.0*	5.6±3.7	5.5±3.6
NR	21.5±6.5*	4.6±3.4	5.4±4.5
NVM	28.2±5.9*	3.1±3.1	3.2±3.0
SCS	41.5±7.6*	35.4±7.2	34.0±6.7
VVK	24.6±6.8*	6.6±5.7	7.3±6.3

TABLE IV

RESULTS OF THE EVALUATED SYSTEMS ON THE IRCCYN DATASETS. RESULTS IN BOLD ARE EQUIVALENT (T-TEST PER ROW AT 0.05 SIGNIFICANCE LEVEL) TO THE BEST PERFORMANCE (DEPICTED WITH A *).

a different environment than the one used for recording the training data.

Whereas the expected behavior with the use of the QMUL instance and abstract datasets was a equivalent performance compared to the ones achieved on test QMUL, the expected behavior with the IRCCYN dataset is a drop in performance. As can be seen on Table IV, this drop is significant for all the systems. More importantly, all the systems except the SCS one




Fig. 6. Class-wise event based F-measure achieved by the different systems on the QMUL and IRCCYN datasets.

achieve similar performance when compared to the baseline on the IRCCYN datasets, meaning that for most systems, the performance gain may solely be due to an over adaptation of the system to the training data. Figure 6 summarizes the results, where the good behavior of the SCS system can be clearly seen.

VI. DISCUSSION

In the light of the results discussed above, we believe that considering carefully designed simulated data is useful for gaining knowledge about the properties and behaviors of the systems under evaluation, thus helping the designer in his algorithmic choices and their evaluation. Important factors influencing the performance such as the noise level, the level of polyphony, the intra-class diversity (acoustical difference between training and testing data) can be evaluated independently, without the burden of experimentally recording data which the desired properties and manually annotating them.

Even though the sole use of synthetic data for validating a computational approach is clearly not sufficient, we believe that the sole use of real data may not be sufficient either, should one wish to gain deep knowledge about the impact of some design and parametrization issues involved in the implementation of an engineering system.

Indeed, real data which is well annotated is most of the time a scarce resource as the careful design of a large evaluation dataset is a demanding task. Moreover, depending on the task at hand, which may not be always well posed, the annotation can be a critical issue leading to some compromise that will greatly contribute to the difficulty of evaluating the performance of the algorithms.

We thus believe that considering simulated data is an in between approach, that together with final validation using real data may be very useful in order to produce more knowledge about the engineering systems under evaluation. We shall stress that such approach is taken in more mature fields, for

example robust ASR where challenges are conducted using simulated data such as the CHIME challenges [58], [59]. The simulated acoustic scenes datasets have been generated using a dedicated set of Matlab functions publicly available¹.

VII. CONCLUSION

A morphological model of acoustic scenes has been presented. Following a collection based approach, it generates a set of sound tracks which are sequences of event realizations drawn from specifically tailored sound sample collections. Its potential for generating simulated corpuses of office events scenes is evaluated, by building upon the results obtained thanks to the IEEE AASP DCASE challenge on the detection of events in an office environment.

We believe that considering those simulated corpora allows us to gain important knowledge about the behavior of the systems under evaluation. As most of the systems under evaluation were built for monophonic inputs (one event occurring at a given time), this paper focused on modifying the acoustical properties of the background or the events. Future research will focus on the influence of the degree of overlap when facing polyphonic scenes, potentially with temporal interactions between events, both for single events (e.g. repetitions for a single event) as well as interactions between event classes.

ACKNOWLEDGMENT

The authors would like to thank Mark Plumbley for his support. Research project partly funded by ANR-11-JS03-005-01.

REFERENCES

- [1] Ecology and acoustics: emergent properties from community to landscape. page 94, Paris, France, June 2014. Sueur, J. and Farina, A. and Bobryk, C. and Llusia, D. and McWilliam, J. and Pieretti, N., Musum national d’Histoire naturelle.
- [2] Bryan C. Pijanowski, Almo Farina, Stuart H. Gage, Sarah L. Dumyahn, and Bernie L. Krause. What is soundscape ecology? an introduction and overview of an emerging new science. *Landscape Ecology*, 26(9):1213–1232, 2011.
- [3] Steven R. Ness, Helena Symonds, Paul Spong, and George Tzanetakis. The archive : Data mining a massive bioacoustic archive. *International Workshop on Machine Learning for Bioacoustics*, 2013.
- [4] Dan Stowell and Mark D. Plumbley. Segregating event streams and noise with a markov renewal process model. *Journal of Machine Learning Research*, 14:2213–2238, 2013.
- [5] Dan Stowell and Mark D. Plumbley. Large-scale analysis of frequency modulation in birdsong databases. *Methods in Ecology and Evolution*, 11, 2013.
- [6] Toni Heittola, Annamaria Mesaros, Antti Eronen, and Tuomas Virtanen. Context-dependent sound event detection. *EURASIP Journal on Audio, Speech, and Music Processing*, 2013(1), 2013.
- [7] R. Radhakrishnan, A. Divakaran, and P. Smaragdus. Audio analysis for surveillance applications. In *Applications of Signal Processing to Audio and Acoustics, 2005. IEEE Workshop on*, pages 158–161, Oct 2005.
- [8] T Heittola, A Mesaros, A Eronen, and T Virtanen. Context-dependent sound event detection. *EURASIP Journal on Audio, Speech, and Music Processing*, 2013.
- [9] Johnathan Turner Tae Hong Park, Michael Musick, Jun Hee Lee, Christopher Jacoby, Charlie Mydlarz, and Justin Salamon. Sensing urban soundscapes. In *EDBT/ICDT Workshops*, pages 375–382, 2014.
- [10] Lawrence Rabiner and Biing-Hwang Juang. *Fundamentals of Speech Recognition*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.

¹<https://bitbucket.org/mlagrange/simscene/downloads>

- [11] Meinard Müller. *Information Retrieval for Music and Motion*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2007.
- [12] D. Giannoulis, D. Stowell, E. Benetos, M. Rossignol, M. Lagrange, and M. D. Plumbley. A database and challenge for acoustic scene classification and event detection. In *Proceedings of the European Signal Processing Conference (EUSIPCO)*, 2013.
- [13] C. Jacoby J. Salamon and J. P. Bello. A dataset and taxonomy for urban sound research. In *Proc. 22nd ACM International Conference on Multimedia*, pages 158–161, Nov 2014.
- [14] 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 (WASPAA)*, pages 1–4. IEEE, 2013.
- [15] Rainer Stiefelhagen, Keni Bernardin, Rachel Bowers, John Garofolo, Djamel Mostefa, and Padmanabhan Soundararajan. The clear 2006 evaluation. In Rainer Stiefelhagen and John Garofolo, editors, *Multimodal Technologies for Perception of Humans*, volume 4122 of *Lecture Notes in Computer Science*, pages 1–44. Springer Berlin Heidelberg, 2007.
- [16] Israel Nelken and Alain de Cheveign. An ear for statistics. *Nature neuroscience*, 16(4):381382, 2013.
- [17] Sascha Spors, Heinz Teutsch, Achim Kuntz, and Rudolf Rabenstein. Sound field synthesis. In Yiteng Huang and Jacob Benesty, editors, *Audio Signal Processing for Next-Generation Multimedia Communication Systems*, pages 323–344. Springer US, 2004.
- [18] D.N. Zotkin, R. Duraiswami, and L.S. Davis. Rendering localized spatial audio in a virtual auditory space. *Multimedia, IEEE Transactions on*, 6(4):553–564, Aug 2004.
- [19] Charles Verron, Mitsuko Aramaki, Richard Kronland-Martinet, and Grégory Pallone. A 3-d immersive synthesizer for environmental sounds. *Audio, Speech, and Language Processing, IEEE Transactions on*, 18(6):1550–1561, 2010.
- [20] Diemo Schwarz. State of the art in sound texture synthesis. In *Proc. Digital Audio Effects (DAFx)*, pages 221–231, 2011.
- [21] Josh H. McDermott and Eero P. Simoncelli. Sound texture perception via statistics of the auditory periphery: Evidence from sound synthesis. *Neuron*, 71(5):926–940, September 2011.
- [22] Richard Turner and Maneesh Sahani. Modeling natural sounds with modulation cascade processes. In *Advances in neural information processing systems*, pages 1545–1552, 2008.
- [23] Albert S Bregman. *Auditory scene analysis: The perceptual organization of sound*. MIT press, 1994.
- [24] Joel S Snyder and Claude Alain. Toward a neurophysiological theory of auditory stream segregation. *Psychological bulletin*, 133(5):780, 2007.
- [25] Valter Ciocca. The auditory organization of complex sounds. *Frontiers in bioscience: a journal and virtual library*, 13:148–169, 2007.
- [26] Robert P Carlyon. How the brain separates sounds. *Trends in cognitive sciences*, 8(10):465–471, 2004.
- [27] James A Ballas and James H Howard. Interpreting the language of environmental sounds. *Environment and behavior*, 19(1):91–114, 1987.
- [28] Israel Nelken and Omer Bar-Yosef. Neurons and objects: the case of auditory cortex. *Frontiers in neuroscience*, 2(1):107, 2008.
- [29] Danièle Dubois, Catherine Guastavino, and Manon Raimbault. A cognitive approach to urban soundscapes: Using verbal data to access everyday life auditory categories. *Acta Acustica united with Acustica*, 92(6):865–874, 2006.
- [30] Manon Raimbault and Danile Dubois. Urban soundscapes: Experiences and knowledge. *Cities*, 22(5):339–350, October 2005.
- [31] Maria Niessen, Caroline Cance, and Danile Dubois. Categories for soundscape: toward a hybrid classification. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, volume 2010, page 58165829, 2010.
- [32] Catherine Guastavino. The ideal urban soundscape: Investigating the sound quality of french cities. *Acta Acustica United with Acustica*, 92:945–951, 2006.
- [33] Brian Gygi and Valeriy Shafiro. The incongruity advantage for environmental sounds presented in natural auditory scenes. *Journal of Experimental Psychology: Human Perception and Performance*, 37(2):551, 2011.
- [34] Maria E Niessen, Leendert Van Maanen, and Tjeerd C Andringa. Disambiguating sound through context. *International Journal of Semantic Computing*, 2(03):327–341, 2008.
- [35] Olivier Houix, Guillaume Lemaitre, Nicolas Misdariis, Patrick Susini, and Isabel Urdapilleta. A lexical analysis of environmental sound categories. *Journal of Experimental Psychology: Applied*, 18(1):52–80, 2012.
- [36] F. Guyot, M. Castellengo, and B. Fabre. *Catégorisation et Cognition: De la Perception au Discours*, chapter Étude de la catégorisation d’un corpus de bruits domestiques (A study of the categorization of an everyday sound set), pages 41–58. Édition Kimé, Paris, France, 1997.
- [37] Brian Gygi, Gary R Kidd, and Charles S Watson. Similarity and categorization of environmental sounds. *Perception & psychophysics*, 69(6):839–855, 2007.
- [38] Michael M Marcell, Diane Borella, Michael Greene, Elizabeth Kerr, and Summer Rogers. Confrontation naming of environmental sounds. *Journal of clinical and experimental neuropsychology*, 22(6):830–864, 2000.
- [39] N. J. Vanderveer. *Ecological acoustics: Human perception of environmental sounds*. Cornell University, Ithaca, NY, 1979.
- [40] Eleanor Rosch and Barbara B Lloyd. Cognition and categorization. *Hillsdale, New Jersey*, pages 27–48, 1978.
- [41] William W Gaver. What in the world do we hear?: An ecological approach to auditory event perception. *Ecological psychology*, 5(1):1–29, 1993.
- [42] A.L. Brown, Jian Kang, and Truls Gjestland. Towards standardization in soundscape preference assessment. *Applied Acoustics*, 72(6):387–392, May 2011.
- [43] Michael Southworth. The sonic environment of cities. *Environment and behavior*, 1969.
- [44] Valérie Maffiolo. De la caractérisation sémantique et acoustique de la qualité sonore de l’environnement urbain. *Semantic and acoustic characterization of urban environmental sound quality* (Ph. D. dissertation, Université du Maine, France), 1999.
- [45] Josh H McDermott, Michael Schemitsch, and Eero P Simoncelli. Summary statistics in auditory perception. *Nature neuroscience*, 16(4):493–498, 2013.
- [46] Nicolas Saint-Arnaud. *Classification of sound textures*. PhD thesis, Massachusetts Institute of Technology, 1995.
- [47] Nicolas Saint-Arnaud and Kris Popat. Analysis and synthesis of sound textures. In *in Readings in Computational Auditory Scene Analysis*. Citeseer, 1995.
- [48] Trevor R. Agus, Simon J. Thorpe, and Daniel Pressnitzer. Rapid formation of robust auditory memories: Insights from noise. *Neuron*, 66(4):610–618, May 2010.
- [49] Steven Davis and Paul Mermelstein. Comparison of Parametric Representations for Monosyllabic Word Recognition in Continuously Spoken Sentences. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 28(4):357–366, 1980.
- [50] L. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, Feb 1989.
- [51] S. Chauhan, S. Phadke, and C. Sherland. Event detection and classification. Technical report, 2013. <http://c4dm.eecs.qmul.ac.uk/sceneseventschallenge/abstracts/OL/CPS.pdf>.
- [52] A. Diment, T. Heittola, and T. Virtanen. Sound event detection for office live and office synthetic AASP challenge. Technical report, 2013. <http://c4dm.eecs.qmul.ac.uk/sceneseventschallenge/abstracts/OL/DHV.pdf>.
- [53] J. F. Gemmeke, L. Vuegen, B. Vanrumste, and H. Van hamme. An exemplar-based NMF approach for audio event detection. Technical report, 2013. <http://c4dm.eecs.qmul.ac.uk/sceneseventschallenge/abstracts/OL/GVV.pdf>.
- [54] W. Nogueira, G. Roma, and P. Herrera. Automatic event classification using front end single channel noise reduction, MFCC features and a support vector machine classifier. Technical report, 2013. <http://c4dm.eecs.qmul.ac.uk/sceneseventschallenge/abstracts/OL/NVM.pdf>.
- [55] M. E. Niessen, T. L. M. Van Kasteren, and A. Merentitis. Hierarchical sound event detection. Technical report, 2013. <http://c4dm.eecs.qmul.ac.uk/sceneseventschallenge/abstracts/OL/NR2.pdf>.
- [56] J. Schröder, B. Cauchi, M. R. Schädler, N. Moritz, K. Adiloglu, J. Anemüller, S. Doelo, B. Kollmeier, and S. Goetze. Acoustic event detection using signal enhancement and spectro-temporal feature extraction. Technical report, 2013. <http://c4dm.eecs.qmul.ac.uk/sceneseventschallenge/abstracts/OL/SCS.pdf>.
- [57] L. Vuegen, B. Van Den Broeck, P. Karsmakers, J. F. Gemmeke, B. Vanrumste, and H. Van hamme. An MFCC-GMM approach for event detection and classification. Technical report, 2013. <http://c4dm.eecs.qmul.ac.uk/sceneseventschallenge/abstracts/OL/VVK.pdf>.
- [58] J. Vincent Barker, E. Ma, N., Christensen, C., Green, and P. The pascal chime speech separation and recognition challenge. *Computer Speech and Language*, 27(3), 2013.
- [59] E. Vincent, J. Barker, S. Watanabe, J. Le Roux, F. Nesta, and M. Matasoni. The second ‘chime’ speech separation and recognition challenge: Datasets and tasks and baselines. In *ICASSP*, 2013.