Multiview approaches to event detection and scene analysis

Slim Essid, Sanjeel Parekh, Ngoc Q. K. Duong, Romain Serizel, Alexey Ozerov, Fabio Antonacci, Augusto Sarti

► To cite this version:

Slim Essid, Sanjeel Parekh, Ngoc Q. K. Duong, Romain Serizel, Alexey Ozerov, et al.. Multiview approaches to event detection and scene analysis. Tuomas Virtanen; Mark D. Plumbley; Dan Ellis. Computational Analysis of Sound Scenes and Events, Springer, pp.243-276, 2017, 978-3319634494. 10.1007/978-3-319-63450-0_9 . hal-01620341

HAL Id: hal-01620341

<https://hal.science/hal-01620341>

Submitted on 16 Nov 2017

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.

Chapter 9: Multiview approaches to event detection and scene analysis

Slim Essid, Sanjeel Parekh, Ngoc Q. K. Duong, Romain Serizel, Alexey Ozerov,
Fabio Antonacci and Augusto Sarti

1 Introduction

We now extend the study of sound scene and event classification to settings where the observations are obtained from multiple sensors, which we refer to as *multiview data*, each sensor contributing a particular *view* of the data. Instances of this include both *multi-channel* audio data, as acquired by microphone arrays, or more generally, *multimodal* data, i.e. heterogeneous data that involves two or more *modalities* such as the *audio* or *visual* modalities in video recordings.

Be it for applications in machine perception—at the heart of robots’ and virtual agents’ intelligence systems—or video description—as part of video surveillance or multimedia indexing systems—multiview approaches can lead to a significant boost in performance in challenging real world situations. Indeed, multiplying the sources

Slim Essid
Télécom ParisTech, Université Paris-Saclay, Paris, France, e-mail: `slim.essid@telecom-paristech.fr`

Quang-Khanh-Ngoc Duong
Technicolor, Rennes, France, e-mail: `Quang-Khanh-Ngoc.Duong@technicolor.com`

Sanjeel Parekh
Technicolor, Rennes, France, e-mail: `Sanjeel.Parekh@technicolor.com`

Alexey Ozerov
Technicolor, Rennes, France, e-mail: `Alexey.Ozerov@technicolor.com`

Romain Serizel
Université de Lorraine, LORIA, UMR 7503, Vandœuvre-lès-Nancy, France; Inria, Villers-lès-Nancy, France; CNRS, LORIA, UMR 7503, Vandœuvre-lès-Nancy, France, e-mail: `romain.serizel@loria.fr`

Fabio Antonacci
Politecnico di Milano, Italy e-mail: `fabio.antonacci`

Augusto Sarti
Politecnico di Milano, Italy `augusto.sarti@polimi.it`

of information, through different views, should result in a more robust overall “picture” of the scene being analyzed, where sensors, and consequently views, which are not reliable, e.g. noisy, at a particular time instant, are hopefully backed-up by others. This is for instance the case in video recordings where sound-emitting target events are not visible onscreen because of poor lighting conditions or occlusions.

Such an endeavor is unfortunately as promising as challenging, primarily because of the significant increase in the volume of the data to be analyzed, but also owing to the potential heterogeneity of the different streams of information (e.g. audio and visual streams), which additionally may not be perfectly synchronized. Another difficulty is that it is usually not possible to determine which streams are not reliable at every time instant. To see this, consider the scenario of scene analysis using a robot’s sensors. The data views available may then be composed of the multiple audio streams acquired by the robot’s microphone array, as well as RGB and depth-image streams captured by its cameras, possibly along with other signals recorded by inertial measurement units. As the cameras are pointed at an interactant, events of interest may appear only partially in their field of view, and be present in the audio recording only at a very low signal-to-noise ratio, both due to background noise (including the robot’s internal noise, the so-called *ego-noise*, typically produced by its cooling fans or its actuators) and the voice of the interactant, or the robot itself, in the foreground.

In this chapter, we briefly introduce some of the techniques that can be exploited to effectively combine the data conveyed by the different views under analysis for a better interpretation. Numerous good surveys have been written on the general topic of multimodal data fusion, notably the paper by Atrey *et al.* [10] which is quite comprehensive. Therefore, we hereafter merely provide a high-level presentation of generic methods that are particularly relevant in the context of multiview and multimodal sound scene analysis. We then more specifically present a selection of techniques used for audiovisual event detection and microphone array-based scene analysis.

1.1 Background and Overview

1.1.1 Multiview Architectures

Multiview data can be fused at the representation-level i.e. combining features from multiple views or at decision-level implying integration of partial classifier-outputs. Another important strategy, referred to as *joint subspace learning* in this chapter, does not fall within the stated categories. Herein the aim is to learn or non-trivially transform the representations based on inter-relationships across the views. Each of these methods have been discussed further in section 2.

1.1.2 Visual Features

As video modality is central to the content of this chapter, a short note on commonly employed visual features is in order. Features extracted from video streams can be broadly classified into two categories; appearance-based and motion-based features. For the former, several local and global descriptors representing appearance attributes namely, color, texture and shape are extracted. While some works utilize the raw pixel data or color histograms, others rely on standard features such as scale-invariant feature transform (SIFT) [100] and histograms of oriented gradients (HOG) [42]. Lately, features extracted from convolutional neural networks have dominated [92].

Motion-based features are typically computed using optical flow or tracking data. It is possible to represent temporal changes of segmented regions, objects and shapes by calculating velocity and acceleration i.e. optical flow and its derivative. Other popular features include histograms of optical flow (HOF) [158] and motion boundary histograms (MBH) [158]. As MBH is computed from optical flow derivatives, it is not affected by constant motion. This makes it robust to camera motion. The reader is referred to [80, 105] for an extensive review of visual representations used for multimodal analysis.

In multiview settings temporal synchronization across views is quite challenging. Notably, in the audiovisual case, since video frame rate, typically around 25-30 frames per second is significantly different from audio, features from both modalities must be appropriately sampled for temporal correspondence. Moreover, the natural asynchrony that exists between the two modalities must also be taken into account. This is to say that cues for an audiovisual event might not appear simultaneously in both modalities.

2 General techniques for multiview data analysis

Generally, such techniques operate either at the *representation-level* or the *decision-level* as further described in the next sections.

2.1 Representation and feature integration/fusion

This is the process of combining different types of features or low-level data representations from different views into a common representation (usually to be exploited by a prediction system).

In practice, this can be primarily achieved by concatenating the feature vectors $\mathbf{h}_{m,i}$, extracted from views m , $1 \leq m \leq M$, at the same time positions i , to build integrated feature vectors $\bar{\mathbf{h}}_i = [\mathbf{h}_{1,i}^\top, \dots, \mathbf{h}_{M,i}^\top]^\top$.¹

However, the dimensionality of the resulting representation is often too high, leading researchers to resort to dimensionality reduction methods. A common approach is then to use *feature transform* techniques (see Chapter 4, Section ??), possibly principal component analysis (PCA) [32], independent component analysis (ICA) [144], or linear discriminant analysis (LDA) [32]. An interesting alternative, is *feature selection* (see Chapter 4, Section ??). Indeed, when applied to the feature vectors $\bar{\mathbf{h}}_i$, the selection will hopefully retain a subset of the most “relevant” features across the various views (with respect to the selection criterion).

Nevertheless, in multimodal settings, the previous methods often turn out to be limited owing to the different physical nature of the features to be combined. In particular, the features do not necessarily live in the same metric spaces, and are not necessarily extracted from the same temporal segments. Consequently, there has been a number of proposals attempting to address these limitations.

An interesting approach, within the framework of multiple kernel learning, consists in considering separate kernels for different features, to build optimal convex combinations of these in order to use them for classification, as done for example in [31, 161].

Another approach that is worthy of note is the construction of joint multimodal representations, as done in video analysis applications, where various types of audiovisual representations have been envisaged. Examples include the creation of *audiovisual atoms* [78], or *audiovisual grouplets* [79], both exploiting audiovisual correlations. A joint audiovisual representation may in particular be built using one of the multimodal subspace methods described hereafter.

2.2 Joint subspace learning

2.2.1 Feature-space transformation

A number of techniques has been suggested to map the observed feature vectors from two modalities to a low dimensional space where a *measure of “dependency”* between them can be computed. Let us assume the N observed feature vectors from two modalities, $\mathbf{y}_{1,n} \in \mathbb{R}^{J_1}$ and $\mathbf{y}_{2,n} \in \mathbb{R}^{J_2}$ ($n = 1, \dots, N$), are assembled column-wise in matrices $\mathbf{Y}_1 \in \mathbb{R}^{J_1 \times N}$ and $\mathbf{Y}_2 \in \mathbb{R}^{J_2 \times N}$ respectively.² The methods we describe

¹ provided that the data analysis-rate and cross-view synchronization issues have been previously addressed.

² The underlying assumption is that the (synchronized) features from both modalities are extracted at the same rate. In the case of audio and visual modalities this is often obtained by downsampling the audio features or upsampling the video features, or by using temporal integration techniques [82].

here aim to find two mappings f_1 and f_2 (that reduce the dimensions of feature vectors in each modality), such that a dependency measure $S_{12}(f_1(\mathbf{Y}_1), f_2(\mathbf{Y}_2))$ is maximized. Various approaches can be described using this same formalism. The advantages of doing so are two-fold: (i) it appropriately modifies the feature spaces to uncover relationships between views specified by the measure of dependency and (ii) by projecting data into the same space, dimensionality difference between views is eliminated and direct comparison across views is made possible. Fisher et al. [53] choose the mutual information [37] as a dependency measure and seek single-layer perceptrons f_1 and f_2 projecting the audiovisual feature vectors to a 2-dimensional space. Other more popular approaches (for which closed-form solutions can be found) use linear mappings to project the feature streams:

- canonical correlation analysis (CCA), first introduced by Hotelling [69], aims at finding pairs of unit-norm vectors t_1 and t_2 such that

$$(t_1, t_2) = \arg \max_{(t_1, t_2) \in \mathbb{R}^{J_1} \times \mathbb{R}^{J_2}} \text{corr} \left(t_1^T \mathbf{Y}_1, t_2^T \mathbf{Y}_2 \right) \quad (1)$$

CCA can be considered equivalent to mutual information maximization for the particular case where the underlying distributions are elliptically symmetric [85]. Several variants have been proposed to incorporate sparsity and non-negativity into the optimization problem to resolve issues with interpretability and ill-posedness, respectively [86, 143]. In the context of multimodal neuronal data analysis, temporal kernel CCA [15] has been proposed to take into account the temporal dynamics.

- An alternative to the previous methods (expected to be more robust than CCA) is co-inertia analysis (CoIA). It consists in maximizing the covariance between the projected audio and visual features:

$$(t_1, t_2) = \arg \max_{(t_1, t_2) \in \mathbb{R}^{J_1} \times \mathbb{R}^{J_2}} \text{cov} \left(t_1^T \mathbf{Y}_1, t_2^T \mathbf{Y}_2 \right) \quad (2)$$

A possible reason for CoIA's stability is that it is a trade-off between CCA and PCA, thus it benefits from advantages of both [21].

- Yet another configuration known as cross-modal factor analysis (CFA), and found to be more robust than CCA in [96], seeks two matrices $\mathbf{T}_1$ and $\mathbf{T}_2$, such that

$$(\mathbf{T}_1, \mathbf{T}_2) = \arg \max_{(\mathbf{T}_1, \mathbf{T}_2)} \left(1 - \|\mathbf{T}_1 \mathbf{Y}_1 - \mathbf{T}_2 \mathbf{Y}_2\|_F^2 \right) = \arg \min_{(\mathbf{T}_1, \mathbf{T}_2)} \|\mathbf{T}_1 \mathbf{Y}_1 - \mathbf{T}_2 \mathbf{Y}_2\|_F^2 \quad (3)$$

with $\mathbf{T}_1 \mathbf{T}_1^T = I$ and $\mathbf{T}_2 \mathbf{T}_2^T = I$. $\|\mathbf{V}\|_F$ denotes the Frobenius norm of matrix $\mathbf{V}$.

Note that all the previous techniques can be kernelized to study non-linear coupling between the modalities considered (see for instance [66, 94]).

The interested reader is referred to [66, 69, 96] for further details on these techniques, and to [60] for a comparative study.

2.2.2 Multimodal dictionary learning

While previous approaches relied on modeling the association between the features across modalities, this class of techniques targets the extraction of meaningful multimodal structures to jointly represent all the modalities. This is useful because feature transformation techniques like CCA impose simplifying assumptions such as linearity and are adversely affected by lack of data. To this end, Monaci et al. [109] propose to learn multimodal dictionaries wherein the dictionary elements are learned using an algorithm that enforces synchrony between modalities and decorrelation between the learned dictionary elements. The learned templates can then be used for performing various tasks. Monaci et al. improve upon this foundational work by proposing a bimodal matching pursuit algorithm which integrates dictionary learning and coding [111]. The sparse shift-invariant generative model used for the audiovisual case can be given by defining multimodal dictionary elements $\{\phi_d\}_{d=1}^D = (\phi_d^a(t), \phi_d^v(x, y, t))$ consisting of audio, ϕ_d^a , and visual, ϕ_d^v , parts and a spatio-temporal shift operator $T_{(p,q,r)}\phi_d = (\phi_d^a(t-r), \phi_d^v(x-p, y-q, t-r))$ such that the multimodal signal s is approximated by the following equation:

$$s \approx \sum_{d=1}^D \sum_{i=1}^{n_d} c_{d_i} T_{(p,q,r)} \phi_d \quad (4)$$

where n_d is the number of instances of ϕ_d and c_{d_i} specifies the weights for AV components of ϕ_d at the i^{th} instance. Several limitations of this approach have been improved upon by proposing a new objective function and algorithm to balance the two modalities, reduce computational complexity and improve robustness [98].

2.2.3 Neural networks and deep learning

Lately, rapid progress in the application of deep learning methods to representation learning has motivated researchers to use them for fusing multiview data [3, 116, 145]. The primary advantage of neural networks is their ability to model very complex nonlinear correlations that exist between multiple views. Early insights into their use for multiview data were provided by Yuhas et al. [167] who trained a network to predict audio using visual input. Subsequently, Cutler et al. [41] proposed to learn audiovisual correlations for the task of speaker detection using a time-delayed neural network (TDNN). Recently, various multimodal autoencoder architectures for learning shared representations have been proposed, even for the case where only a single view is present at training and testing time [116]. Another interesting work extends CCA to learning two deep encodings, one for each view, such that their correlation is maximized [3]. Regularized deep neural networks [162] have also been proposed to construct shared representations taking into account the feature inter-relationships. Each of these methods has been developed independently in different settings. Their application to event analysis and detection still remains

to be explored. This is a rapidly growing area of research; we refer the interested reader to [58, 70] for recently proposed multimodal fusion architectures.

2.2.4 Co-factorization techniques

Matrix factorization techniques can be profitably used to extract meaningful representations for the data being analyzed.

When dealing with multiview or in case of audio multi-channel data, observations from multiple views or channels may be profitably assembled in *multi-way arrays*, aka *tensors*, before being modeled by tensor factorization methods. As for multi-channel audio data, a popular approach consists in collecting the spectrograms of signals from different channels (originating from different microphones) in a 3-way tensor, as illustrated in Figure 1, before processing it with the so-called PARAFAC (PARAllel FACtor analysis) decomposition method, possibly with non-negativity constraints. This can be interpreted as an attempt to explain audio spectra $\mathbf{v}_{im}$ observations as being linear combinations of elementary spectra $\mathbf{w}_k$, temporally weighted by activation coefficients h_{ki} up-to spatial modulation coefficients q_{mk} .

Fig. 1: PARAFAC decomposition of multi-channel audio spectra.

Such decompositions were found particularly useful in multichannel audio source separation [54, 122]. For more information about tensor factorization methods, we refer the reader to [34, 91, 163].

In contrast to the previous setting, data from different modalities usually live in feature spaces of completely different topology and dimensionality (think of audio as opposed to video), preventing the possibility of “naturally” representing them by the same tensor.

In this case, one may resort to so-called *co-factorization* techniques, that is techniques performing two (or more) factorizations in parallel, which are linked in a particular way. Because of the different nature of the modalities, this link has usually to be characterized through temporal dependencies between the temporal acti-

vations in cross-modal correspondence, and unlikely through dependencies between dictionary elements of different modalities.

Assuming that appropriate nonnegative features have been extracted at the same rate from the two modalities being analyzed³—say the audio and images of a video—so that two observation matrices $\mathbf{V}_1 \in \mathbb{R}_+^{J_1 \times N}$ and $\mathbf{V}_2 \in \mathbb{R}_+^{J_2 \times N}$ are available, for the audio and visual data. One may seek a model $(\mathbf{W}_1, \mathbf{W}_2, \mathbf{H})$ such that:

$$\begin{cases} \mathbf{V}_1 \approx \mathbf{W}_1 \mathbf{H} \\ \mathbf{V}_2 \approx \mathbf{W}_2 \mathbf{H} \\ \mathbf{W}_1 \geq 0, \mathbf{W}_2 \geq 0, \mathbf{H} \geq 0; \end{cases} \quad (5)$$

in such a way that the temporal activations be the same for both modalities. This is referred to as *hard co-factorization*, an approach that has been followed in a number of works (see e.g. [55, 164, 165]). Clearly, this approach is limited in that it does not account for possible local discrepancies across the modalities. This happens for example when there is a mismatch between the audio and the images information, say because of a visual occlusion in video analysis scenarios. This motivates the *soft co-factorization* model of Seichepine et al. [139], which merely encourages the temporal activations corresponding to each modality to be close, as opposed to equal, according to:

$$\begin{cases} \mathbf{V}_1 \approx \mathbf{W}_1 \mathbf{H}_1 \\ \mathbf{V}_2 \approx \mathbf{W}_2 \mathbf{H}_2 \\ \mathbf{H}_1 \approx \mathbf{H}_2 \\ \mathbf{W}_1 \geq 0, \mathbf{W}_2 \geq 0, \mathbf{H}_1 \geq 0, \mathbf{H}_2 \geq 0. \end{cases} \quad (6)$$

The model (6) is estimated by solving the following optimization problem:

$$\begin{cases} \min_{\theta} C_c(\theta) ; \theta \triangleq (\mathbf{W}_1, \mathbf{H}_1, \mathbf{W}_2, \mathbf{H}_2) \\ \mathbf{W}_1 \geq 0, \mathbf{W}_2 \geq 0, \mathbf{H}_1 \geq 0, \mathbf{H}_2 \geq 0; \end{cases} \quad (7)$$

$$C_c(\theta) \triangleq D_1(\mathbf{V}_1 | \mathbf{W}_1 \mathbf{H}_1) + \gamma D_2(\mathbf{V}_2 | \mathbf{W}_2 \mathbf{H}_2) + \delta P(\mathbf{H}_1, \mathbf{H}_2); \quad (8)$$

where:

- $D_1(\cdot | \cdot)$ and $D_2(\cdot | \cdot)$ are the measures of fit respectively relating to the first and second modalities; note that they may be chosen to be different divergences, each well suited to the corresponding feature space;
- $P(\cdot, \cdot)$ is a penalty on the difference between (properly rescaled) activation values occurring at the same instant; they can be either the ℓ_1 or ℓ_2 -norm of the difference between the rescaled activations;

³ To simplify, we consider the case of two modalities, but clearly the methods described here can be straightforwardly generalized to more than two data views by considering the relevant pairwise associations.

- γ and δ are regularization parameters controlling, respectively, the relative importance of each modality and the coupling penalty.

The interested reader is referred to [139] for more details on the algorithms.⁴ The soft co-factorization scheme has proven effective for multichannel [139] and multimodal audio source separation [124], as well as multimodal speaker diarization [138]. It is believed to be promising for audiovisual event detection tasks.

2.3 Decision-level integration/fusion

Decision-level fusion, aka *late integration* refers to the idea of combining intermediate decisions, i.e. partial classifier-outputs, in order to achieve a more accurate multimodal characterization of a content has been explored extensively, under various configurations. This can be seen as a particular case of *ensemble learning* [130] where the base classifiers (to be combined) operate on different views of the data.

Numerous works rely on *majority voting* procedures whereby final global decisions are made based on a weighted sum of individual voters, each typically corresponding to a decision taken on a particular view. The weights are often chosen using either heuristics or trial-and-error procedures (see for example [97]). This idea can be better formalized using a Bayesian framework that allows for taking into account the uncertainty about each classifier decisions [73, 107].

2.3.1 Probabilistic combination rules

When using classifiers providing local probabilistic outputs $p(G_c | \mathbf{h}_{m,i})$ for the i -th observation of the m -th view, $\mathbf{h}_{m,i}$, a simplistic decision strategy assumes feature-vector observations from different views to be independent, and the the decision rule consequently takes the form:

$$\hat{G} = \arg \max_c \log[p(G_c | \mathbf{h}_{m,i}, \dots, \mathbf{h}_{m,i})] = \arg \max_c \sum_{m=0}^{M-1} \log p(G_c | \mathbf{h}_{m,i}). \quad (9)$$

It is worth mentioning that alternative simple combination rules have also been employed that are discussed in Kittler et al. [88].

The previous approach does not allow for incorporating prior knowledge about the dependency structure in the data, in particular the cross-modal and temporal dependencies. To this end, sophisticated dynamic classifiers have been utilized, ranging from variants of (multi-stream) HMM [7, 62, 87, 115], through more general dynamic Bayesian networks [33, 61, 113], to even more general graphical models such as conditional random fields (CRF) [20].

⁴ Matlab implementations are available online at <http://plato.telecom-paristech.fr/publi/26108/>.

2.3.2 Neural networks

Neural networks can also be used for late integration. Some works have utilized them to adaptively learn the weights for fusing multiple classifiers or system outputs [76, 117]. This is typically carried out by training the network to minimize the error between estimated and oracle weights [76]. Besides, in order to take into account the temporal and multi-view dependencies, a frequently used strategy is to perform end-to-end training with “fusion” integrated as a layer (usually close to the output layer) into the architecture [51, 84]. Such methods cannot be termed as late fusion in the conventional sense as feature learning and decision fusion stages are not independent.

2.3.3 Other methods

Another widespread strategy consists in using the monomodal classifier-outputs as features, on the basis of which a new classifier, that is expected to optimally perform the desired multimodal fusion, is learned [13, 160].

Also, solutions to deal with the potential imprecision of some views have been proposed using the *Dempster-Shafer* theory [56].

Finally, it is important to note that the techniques described in this section are not mutually exclusive: in practice one may jointly consider different integration strategies for different features and views (possibly being driven by some expert knowledge), and different analysis time-horizons. This raises the difficult issue of effectively and efficiently exploiting, at the final prediction stage, heterogeneous representations: low-level instantaneous features, possibly over varying time-scales, intermediate prediction results—sometimes seen as outputs of *event* or *concept* detectors— bags-of-words or bags-of-systems extracted over longer texture-windows, *etc.*

3 Audiovisual event detection

3.1 Motivation

The target of audiovisual event detection (AVED) is to detect specific events that occur in an audiovisual recording or realtime stream, and to identify the class of those events. Though the task is more widely addressed through the analysis of the video images, information conveyed by the sound track may become key for a proper detection. Indeed, the visual information may not be sufficient since occlusions may occur and events may be localized in space, hence not visible in the images given that the camera field of view is necessarily restricted. Also the images may not be

usable because of poor lighting conditions, or fast camera motion. AVED then enables a more reliable detection of these events, by combining audio and visual cues.

3.1.1 Examples in video content analysis and indexing

Researchers continue to explore various techniques for improving video content analysis and indexing for better navigation and user experience. In this context, AV event analysis at various levels of granularity provides useful insights into the composition of such data in terms of objects, activities and scenes. This not only improves retrieval but also provides a representation closer to our understanding of the physical world. For example, a user could search a database for activity videos such as “dribbling a basketball” or “playing a violin”. Evidently, these are two very distinct tasks where the differences can be readily detected based on auditory and motion information. Moreover, joint analysis could reveal the presence of various objects (e.g. violin, basketball) and also the surroundings (e.g. concert hall, court).

Such an analysis makes object detection and segmentation [74], concept classification [77–79], scene segmentation and change detection [153], activity analysis and various other related tasks possible. Several systems submitted to TRECVID video content analysis tasks of multimedia event detection, story segmentation and search rely on AV analysis [2, 81].

3.1.2 Examples in AV Surveillance and robot perception

Video has recently become an increasingly important resource for forensics and surveillance [108, 129]. Video captured by CCTV systems or video recorded from mobile devices (and possibly shared on multimedia platforms) can provide essential clues in solving criminal cases. For example when considering an investigation about a missing person, video documents can help to localize the missing person or a suspect, providing crucial information about their whereabouts. The analysis of videos linked with a missing person or her/his social network can also help to understand the conditions of the disappearance (was it a kidnapping, a runaway...) and largely influence the investigation.

An investigator looking for a video in a large dataset may want to retrieve information based on the type of scene where the video was recorded or also, at a finer granularity level, based on specific events that occurred during the recording. In addition, the detection of specific events can help to confirm (or deny) the fact that a video was recorded in a particular scene. Some events are indeed representative of particular scenes for example train noise in all probability indicates the scene takes place in a train station or plates and cutlery noises indicate the scene is probably taking place in a restaurant [22]. On the other hand, some events are unlikely to happen in particular scenes. AVED can then help tracking anomalies to detect abnormal events (gunshots, crowd panic...) [101] or to identify a recording scene where information has voluntarily been concealed. This is the case, for exam-

ple, when a kidnapper sends a ransom video recorded from inside a building but a church bell or a train passing nearby can be heard during the video. This type of information that is not present visually can help to localize the place where the video was recorded [140].

3.2 AV event detection approaches

3.2.1 AV event detection and concept classification

Approaches to AV event detection have been very varied and data dependent. Many works for traditional event detection utilize Markov model variants such as the duration dependent input-output Markov model (DDIOMM) [114], multistream HMM or coupled HMM [71]. The former uses a decision level fusion strategy and the latter two do it at an intermediate level. These methods have been shown to perform better than single modality-based approaches with coupled-HMMs being particularly useful for modeling AV asynchrony.

Specifically, with regard to event detection in surveillance videos, Cristiani et al. [39] propose to use the AV concurrence matrix to identify salient events. The idea is to model the audio/video foreground and construct this matrix based on the assumption that simultaneously occurring AV foreground patterns are likely to be correlated. Joint AV analysis has also been employed extensively for sports video analysis and for broadcast analysis in general. In one approach, several feature detectors are built to encode various characteristics of field sports. Their decisions are then combined using a support vector machine (SVM) [132]. Several approaches for structuring TV news videos have also been proposed.

On the other hand, *joint codebook-based approaches* have been quite popular for the task of multimedia concept classification.⁵ In essence, each element of these multimodal codebooks captures some part of a salient AV event. Work on short-term audiovisual atoms (S-AVA) [78] aims to construct a codebook from multimodal atoms which are a concatenation of features extracted from tracked short-term visual-regions and audio. To tackle the problem of video concept classification, this codebook is built through multiple instance learning. Following this work, AV *grouplets* (AVG) [79] were proposed, where separate dictionaries are constructed from coarse audio and visual foreground/background separation. Subsequently, AVGs are formed based on the mixed-and-matched temporal correlations. For instance, an AVG could consist of frames where a basketball player is seen in the foreground with the audio of the crowd cheering in the background. As an alternative, Jhuo et al. [77] determine the relations between audio and visual modalities by constructing a bi-partite graph from their bag-of-words representation. Subsequently, spectral clustering is performed to partition and obtain *bi-modal* words. Unlike S-AVA and bimodal words, AVG has the advantage of explicitly tackling temporal interactions.

⁵ Here the term “concept classification” refers to generic categorization in terms of scene, event, object or location [78].

However, like S-AVA, it relies on video region tracking, which is quite difficult for unconstrained videos.

3.2.2 AV object localization and extraction

AV object localization and extraction refers to the problem of identifying sources visually and/or aurally. This section serves to show how objects responsible for audiovisual events can be extracted from either of the modalities through joint analysis. The general approach is to first associate the two modalities using methods discussed in section 2. The parameters learned during the former step can then be utilized for object localization and segmentation (visual part), audio source separation (audio counterpart) or unsupervised AV object extraction in both modalities. We now discuss approaches to each of these application scenarios.

Object localization and segmentation has been a popular research problem in the computer vision community. Various approaches have leveraged the audio modality to better perform this task with the central idea of associating visual motion and audio. Fisher et al. [53] proposed to use joint statistical modeling to perform this task using mutual information. Izadinia et al. [74] consider the problem of moving-sounding object segmentation, using CCA to correlate audio and visual features. The video features consisting of mean velocity and acceleration computed over spatio-temporal segments are correlated with audio. The magnitude of the learned video projection vector indicates the strength of association between corresponding video segments and the audio. Several other works have followed the same line of reasoning while using different video features to represent motion [86, 143]. Effectiveness of CCA can be illustrated with a simple example of a video with a person dribbling a basketball [74] (see Fig. 2). Simplifying Izadinia et al. 's [74] visual feature extraction methodology, we compute the optical flow and use mean velocity calculated over 40×40 blocks as the visual representation and mel-spectra as the audio representation. The heat map in Fig. 2 shows correlation between each image block and audio. Areas with high correlation correspond to regions with motion. If we instead use a soft co-factorization model [139], it is indeed possible to track the image blocks correlated with the audio in each frame.

Another approach worth mentioning is one that uses Gestalt principles for locating sound sources in videos [110]. Inspired by Gestalt principle of temporal proximity the authors propose to detect synchronous audiovisual events. A particularly different approach was taken by Casanovas et al. [30] who proposed an audio-visual diffusion coefficient to remove information from video image parts which are not correlated with the audio.

Audio source separation is the audio counterpart of the previously discussed problem. The aim is to extract sound produced by each source using video information. As done for videos, mutual information maximization has been used to perform source separation in a user-assisted fashion by identifying the source spatially. Recent methods perform this within the NMF-based source separation framework [124, 137].

Fig. 2: CCA Illustration: Heat map showing correlation between video image regions and audio. Black squares indicate highest correlation.

Several other approaches deal with both object segmentation and source separation *together* in a completely unsupervised manner. Work by Barzeley et al. [11] considers onset coincidence to identify AV objects and subsequently perform source separation. A particular limitation of this method is the requirement of setting multiple parameters for optimal performance on each example. Blind AV source separation work has also been attempted using nonnegative CCA [143] and sparse representations [29]. Independent component analysis over concatenated features from both modalities also extracts meaningful audio-visual objects [144]. However its application is limited to static scenes. Finally, multimodal dictionary learning has also been utilized in this context [98].

While the methods discussed in this section have been shown to work well in controlled environments, their performance is expected to degrade in dense audio-visual scenarios. Moreover, they make a simplifying assumption that all the objects are seen on-screen. It must be emphasized that most of these techniques can be considered symmetric, in the sense that they can be applied to tasks in either of the modalities with appropriate representations.

4 Microphone array based sound scene analysis

In complex sound scenes the sounds coming from different sources can be overlapping in both time and frequency. Single channel processing discriminates them mainly based on frequency. Therefore trying to separate sounds that are overlapping in frequency with single channel techniques will inevitably introduce degradations such as spectral distortion and randomly fluctuating noises (usually referred to as musical noise artefacts that can manifest themselves as, e.g., spurious bursts of energy or “tin can” like sound). Microphone arrays enable the usage of multichannel

techniques that exploit not only spectral diversity between sources but also spatial information about their location.

4.1 Spatial cues modeling

In order to exploit spatial information about the sound sources, audio scene analysis algorithms usually first model the spatial cues and then estimate the corresponding parameters. Both deterministic and probabilistic modeling of such spatial cues have been widely considered in the literature. The former case usually relies on (a) the *point source* assumption, where sound from a source is assumed to come from a single position, and (b) the *narrowband approximation*, where a mixing process from an audio source to the microphone array is characterized by a mixing frequency dependent vector [102]. Probabilistic modeling is usually applied for reverberated or diffuse sources, where sound from a source may come from many directions due to the reverberation, in e.g. source localization [18, 65, 120], separation [48, 75, 104], and beamforming systems [17, 50]. This section will discuss some typical spatial cue models, in both a deterministic and a probabilistic sense, for different audio scene analysis applications.

4.1.1 Binaural approach

Humans generally combine cues from several audiovisual streams to localize sound sources spatially. The main cues for localization in the horizontal hemisphere are related to binaural hearing (relying on the difference between the signal reaching the right ear and the signal reaching the left ear). All these cues are encoded in the so-called interaural transfer function (ITF) that includes the following:

- The **interaural time difference (ITD)** is the difference between the time-of-arrival of a signal at the left ear and the right ear. It is useful to localize sounds based on their onset and at low frequency (below 1.5 kHz) [93] (see also Fig. 3a).
- The **interaural phase difference (IPD)** is the phase difference between the signal at the left ear and the right ear. It is useful to localize on-going sound as long as the wavelength is larger than the diameter of the head (below 1.5 kHz) [166] (see also Fig. 3b);
- The **interaural intensity difference (IID)** is the difference in level between the signal at the left ear and the right ear due to the acoustic shadow produced by the head for sounds above 3 kHz (below the so-called head shadow effect is not present) [112] (see also Fig. 3c).

All the concepts mentioned above can be extended to general microphone array setups. These spatial cues are then extensively exploited to extract a signal of interest from the mixture using beamforming approaches described in Sect. 4.2 (for example the delay-and-sum beamformer directly relies on ITD and ILD). Spatial cues can

Fig. 3: Artificial representation of the binaural cues.

also be used directly for sound source localisation (see also Sec. 4.2.3) and, by proxy, for source separation (see also Sec. 4.2.1) and sound event detection (see also Sec. 4.2.2).

4.1.2 Beamforming methods

Fixed beamformers compose a first simple class of multichannel algorithms which can separate signals coming from different directions. A fixed beamformer tries to steer toward the direction from where the desired sound signal comes and to reject signals coming from other directions. The main categories of fixed beamformers include: delay-and-sum beamformers, filter-and-sum beamformers [68], superdirective microphone arrays [38] or the original formulation of the minimum variance distortion-less beamformer (MVDR) [26].

Adaptive beamformers try to steer toward the direction of the desired sound signal and to adaptively minimize the contributions from the undesired sources coming from other directions. This typically yields a constrained optimization problem. Frost introduced the linearly constrained minimum variance beamformer (LCMV) as an adaptive framework for MVDR [57].

The generalized side lobe canceler (GSC), also known as the Griffiths-Jim beamformer, is an alternative approach to the LCMV where the optimization problem is reformulated as an unconstrained problem [64]. The GSC can be decomposed as a fixed beamformer steering toward the desired source, a blocking matrix and a multichannel adaptive filter [67].

The multichannel Wiener filters (MWF) represent another class of multichannel signal extraction algorithms which are defined by an unconstrained optimization problem [47]. MWF-based algorithms can be implicitly decomposed into a spatial filter and a spectral filter; and can indeed be considered as beamformers [141]. Besides, a reformulation of MWF allows for explicitly controlling the spectral distortion introduced [47, 141].

4.1.3 Nonstationary Gaussian model

The nonstationary Gaussian framework has emerged in audio source separation [48, 52, 118, 123] as a probabilistic modeling of the reverberated sources. It was then also applied in e.g. multichannel acoustic echo cancellation [148] and multichannel speech enhancement [149]. In this paradigm, the short-time Fourier transform (STFT) coefficients of the *source images* $\mathbf{c}_j(t, f)$, i.e. the contribution of the j -th source ($1 \leq j \leq J$) at the microphone array, are modeled as a zero-mean Gaussian random vector whose covariance matrix $\hat{\mathbf{R}}_j(t, f) = \mathbb{E}(\mathbf{c}_j(t, f)\mathbf{c}_j^H(t, f))$ can be factorized as

$$\hat{\mathbf{R}}_j(t, f) = v_j(t, f)\mathbf{R}_j(t, f), \quad (10)$$

where $v_j(t, f)$ are scalar time-varying *variances* encoding the spectro-temporal power of the sources and $\mathbf{R}_j(t, f)$ are $I \times I$ *spatial covariance matrices* encoding their spatial position and spatial width. This model does not rely on the point source assumption nor on the narrowband assumption, hence it appears applicable to reverberated or diffuse sources. In the general situation where the sound source can be moving, the spatial cues encoded by $\mathbf{R}_j(t, f)$ are time-varying. However, in most cases where the source position is fixed and the reverberation is moderate, the spatial covariance matrices are time-invariant: $\mathbf{R}_j(t, f) = \mathbf{R}_j(f)$. Different possibilities of parameterizing $\mathbf{R}_j(f)$ have been considered in the literature resulting in either the rank-1 or the full-rank matrices, where the later case was shown to be more appropriate for modeling the reverberated and diffuse sources as it accounts directly for the interchannel correlation in the off-diagonal entries of $\mathbf{R}_j(f)$ [48].

4.2 Spatial cues based sound scene analysis

This section will discuss the use of spatial cue models presented in the previous section in some specific applications, namely sound source separation, acoustic event detection, and moving sound source localization and tracking.

4.2.1 Sound source separation

In daily life, recorded sound scenes often result from the superposition of multiple sound sources which prevent both human and machines from well localizing and perceiving the target sound sources. Thus, source separation plays a key role in sound scene analysis, and its goal is to extract the signals of individual sound sources from an observed mixture [102]. It offers many practical applications in e.g. communication, hearing-aid, robotics, and music information retrieval [6, 14, 103, 156].

Most source separation algorithms operate in the time-frequency (T-F) domain with the mixing process formulated as

$$\mathbf{x}(t, f) = \sum_{j=1}^J \mathbf{c}_j(t, f) \quad (11)$$

where $\mathbf{x}(t, f) \in \mathbb{C}^{I \times 1}$ denotes the STFT coefficients of the I -channel mixture at T-F point (t, f) , and $\mathbf{c}_j(t, f) \in \mathbb{C}^{I \times 1}$ is the j -th source image. As $\mathbf{c}_j(t, f)$ encodes both spectral information about the sound source itself and the spatial information about the source position, a range of spectral and spatial models has been considered in the literature resulting in various source separation approaches. In the determined case where $I \geq J$, non-Gaussian modeling such as frequency-domain independent component analysis (FDICA) has been well-studied [126, 133]. In the under-determined situation where $I < J$, sparse component analysis (SCA) has been largely investigated [19, 63, 83]. As a specific example of the nonstationary Gaussian modeling presented in Section 4.1.3, the parameters are usually estimated by the expectation maximization (EM) algorithm derived in either the maximum likelihood (ML) sense [48] or the maximum a posteriori (MAP) sense [49, 121, 123]. Then source separation is achieved by the multichannel Wiener filtering. Readers are referred to e.g. [99, 154] for the survey of recent advances on both *blind* scenarios, and *informed* scenarios which exploit some prior knowledge about the sources themselves [123] or the mixing process [49] to better guide the source separation.

4.2.2 Sound event detection

As different sound events usually occur at different spatial locations in the sound scene, spatial cues obtained from microphone array processing intrinsically offers important information for SED. As examples, information about the source directions inferred from the interchannel time differences of arrival (TDOA) was used to help partitioning home environments into several areas containing different types of sound events in [155]. The combination of these spatial features with the classic MFCC was reported to improve the event classification in the experiment. Motivated by binaural processing, in [1] the stereo log-mel-band energy is extracted from stereo recordings to train the neural networks in order to obtain a meaningful cue similarly to the IID.

4.2.3 Localization and Tracking of Acoustic Sources

Sound source localization and tracking are concerned with estimating and following the position of a target source within a sound scene. This active field of research in microphone array processing finds important applications e.g. in surveillance or video conferencing where the camera should be able to follow the moving speaker, and even can automatically switch the capture to an active sound source in multiple source environments [157]. Spatial cues offered by the multichannel audio capture play a key role in deriving the algorithms.

The problem of acoustic source localization has been a relevant topic in the audio processing literature for the past three decades because of its applicability to a wide range of applications [43, 150]. The most effective solutions rely on the use of spatial distributions of microphones in space, which sample the sound field at several locations. Spurious events, reverberation and environmental noise, however, can be a significant cause of localization error. In order to ease the problem, at least for those errors that are contained in a limited number of time frames, source tracking techniques can come in handy, as they are able to perform trajectory regularization, even on the fly. Typical approaches are based on Particle [5, 95, 159], Kalman [4, 59] or distributed Kalman [147] filtering.

Different methodologies have been developed for the localization of acoustic sources through microphone arrays. Those that gained in popularity are based on measurements of the time delay between array microphones. Working in the time domain is often a suitable choice for wideband signals, and most techniques tend to rely on the analysis of the Generalized Cross-Correlation (GCC) of the signals [28, 90] and variants thereof. Localization in the frequency domain, however, can be shown to attain good results for narrowband or harmonic sources immersed in a wideband noise and rely on the analysis of the covariance matrix of the array data. A taxonomy of the localization techniques is represented on Fig. 4.

Fig. 4: Taxonomy of source localization techniques.

Time-domain Localization

Steered Response Power (SRP, [44, 45, 106]) and Global Coherence Field (GCF, [119]) proceed through the computation of a coherence function that maps the GCC values at different microphone pairs on the hypothesized source location. A source location estimate is found as the point in space that maximizes the coherence function. In [23] the scenario of multiple sources is accommodated through a two-step procedure that, after localizing the most prominent source, deemphasizes its contribution in the GCC, so that other sources can be localized. These techniques are known for their high level of accuracy, and are suitable for networks of microphone arrays, where synchronization can only be guaranteed between microphones of the same array. One limitation of such solutions is their computational cost, which is proportional to the number of hypothesized source locations. This means that increasing the spatial resolution results in higher computational costs. Some solutions have been proposed in the literature to mitigate this problem. In [169] the authors propose a hierarchical method that begins with a coarser grid, and refines the estimate at different steps by computing the map for finer grids concentrated around the candidate locations estimated at the previous step. In [46] a similar approach

Fig. 5: Example of a coherent map using the Steered Response Power with Stochastic Region Contraction technique (SRP-SRC, from [46])

is adopted, but a Stochastic Region Contraction strategy is used for going from a coarser to a finer grid. An example of Steered Response Power with Stochastic Region Contraction map is shown on Figure 5. A coarser grid is also used in [35], where the coherence function is not mapped for the centroid of the grid element, but is integrated throughout all the points that fall within the considered element.

Less cumbersome are the solutions based on the Time Difference Of Arrival (TDOA), which is estimated as the time lag of the GCC that exhibits the maximum value. The TDOA is then converted into Range Difference (RD), which measures

the difference of the range between the source and the two microphones in the pair. The locus of candidate source locations corresponding to a given TDOA is a hyperbola whose foci are in the microphones locations, and whose aperture is proportional to the measured TDOA. The most straightforward technique for localization consists in intersecting hyperbolas corresponding to the TDOA measurements coming from different pairs of microphones. The cost function that is based on this procedure is strongly nonlinear, which makes the method sensitive to measurement errors. Least squares cost functions provide a good approximation [12, 36, 72, 134]. The main drawback of TDOA-based localization is its sensitivity to outlier measurements. In [24, 25, 135] techniques for removal of the outliers were presented. In particular, the DATEMM algorithm [135] is based on the observation that TDOAs over closed loop must sum to zero. An example of the improvement of the performance for a scenario of two sources (labelled as a and b) in a reverberant environment is shown in Table 1. In particular the average TDOA error for each of the two sources, and the resulting localization errors are shown.

Table 1: Improvement of the localization and TDOA measurements accuracy brought by the DATEMM algorithm (from [135]).

	$\Delta \hat{\mathbf{p}}$ [cm]		$\Delta \hat{\mathbf{h}}$ [samples]		
	$\hat{\mathbf{p}}_a$	$\hat{\mathbf{p}}_b$	$\hat{\mathbf{h}}_a$	$\hat{\mathbf{h}}_b$	$\hat{\mathbf{h}}_{II}$
DATEMM	4	3	5	28	232
GCC-PHAT	31	48	317	239	—
SRP-PHAT	14	12	—	—	—

Frequency-domain Localization

Techniques in the frequency domain are based on the observation that different microphones in the array will receive differently delayed replicas of the source signals. This, in the frequency domain, corresponds to a phase offset. For distant sources the phase offset between adjacent microphones is constant throughout the array. Delay-and-sum beamformers compensate the offsets so that the components related to a direction will sum up coherently and the others will not. The estimation of the direction of arrival (DOA) of the target source proceeds by searching for the direction that maximizes the output energy of the beamformer over a grid of directions [128, Chapter 6]. Delay and sum is known for its low resolution capabilities, which makes it difficult to distinguish sources that are seen under close angles from the array viewpoint. Minimum Variance Distortionless Response beamformer (MVDR, [27]), Multiple Signal Classification (MUSIC, [136]), and Estimation of the Signal Parameters through Rotational Invariance Techniques (ESPRIT, [131]) bring improvements in terms of resolution. However, they are known for their sensitivity to noise and reverberation, which tends to introduce spurious localizations. The superdirective data-independent beamformer [16] was shown to partially mit-

igate this problem. An interesting solution to the sensitivity to reverberation was proposed in [142] for the detection of gunshots using networks of sensors, each equipped with four or more microphones. For each sensor, both DOA and TDOA are measured. Source location is estimated by intersecting the loci of potential source locations (hyperbolas and direction of arrival) for the two kind of measurements from all the sensors. In reverberant conditions and in the presence of intereferers, some TDOAs and some DOAs could be related to spurious paths, thus providing multiple estimates of the gunshot location. The actual gunshot location is found as the one that maximizes the number of consistent TDOAs and DOAs.

It is important to notice that TDOA-based and frequency domain source localization techniques require the synchronization of the microphones within the array. This, in fact, becomes an issue when multiple independent small arrays are deployed in different locations. In [25] the authors propose a technique for the localization without requiring a preliminary synchronization of the arrays by including the time offsets between the arrays into the unknowns, along with the location of the source. Another important issue is the self-calibration of the array, i.e. the estimation of the mutual relative positions of the microphones [40, 151]. The widespread diffusion of mobile phones and devices equipped with one or more microphones enables the implementation of a wireless acoustic sensor network in seconds, for goals ranging from teleconferencing to security. In this context, however, both calibration and synchronization are needed before normal operation [127].

Acoustic Source Tracking

Independently of the adopted localization method, reverberation and interferers could introduce spurious localizations. The goal of source tracking is to alleviate the influence of outliers. The idea behind tracking is that measurements related to the actual source must follow a dynamical model whereas those related to spurious sources must not [159]. Another goal that can be pursued with tracking systems is that of fusing information coming from both audio and visual localization systems [9, 146]. Several solutions have been presented in the literature. The Kalman filter [59, 89] is a linear system characterized by two equations. The state equation models the evolution the state of a system (location and speed of the source) from one time frame to the next one. The observation equation links the state variables with the observable measurements. The goal of the Kalman filter is to estimate the current state from the knowledge of time series of the observations.

Recently, Distributed Kalman filters were used, which enable the tracking of acoustic sources also in the case of distributed array networks [147, 168], without requiring that all nodes communicate the whole state of the system.

Inherent assumptions that lie in the use of the Kalman filter are the linearity and Gaussianity of measurement and state vectors. In order to gain in robustness against the nonlinearity, the use of the Extended Kalman Filter has been proposed [146], which linearizes the nonlinear system around the working point. In order to gain in

robustness against non Gaussian conditions, however, one has to resort to a different modeling of the source dynamics. In recent years particle filter gained interest in the source localization community due to the fact that it is suitable also to perform tracking in nonlinear non-Gaussian systems and, more in general, for its higher performance [159]. Particle filter [8] assumes that both state and measurement vectors are known in a probabilistic form. Once a new measurement vector is available, the likelihood function of the current observation from a given state is sampled through particles. Each particle is assigned a weight, which determines its relevance in the likelihood function. Only relevant particles will be propagated to the next step. The source location is determined as the centroid of the set of particles. An example of tracking of one, two, or three acoustic sources on a given trajectory for DOA measurements is shown in Fig. 6.

Fig. 6: Example of tracking of one, two or three sources over a prescribed trajectory (from [152]).

In audio surveillance contexts, it is important to enable localization also when multiple sources are active at any time, with a small convergence time when acoustic sources alternate. This is important, for example, in events that involve multiple acoustic sources (brawls, people yelling, etc.). In recent years, swarm particle filtering has shown to address this scenario particularly well [125]. It is based on the idea that the propagation of each particle to the next step is determined not only by the previous history of the particle itself, but also by the particle that exhibits the best likelihood at the current time instant. Consequently, the overall behavior of the systems resembles that of a bird flock, rapidly moving towards the active source. An example of behavior of swarm particle filtering is shown in Fig. 7. Here two sets of particles at four consecutive time frames estimate the location of a source using Particle Filtering (PF) and Swarm Particle Filtering (Swarm). The two sets are initialized identically. It is possible to notice that after four steps the Swarm particles cluster around the source location, while the PF is still converging.

5 Conclusion and outlook

Multichannel and multimodal data settings represent opportunities to address complex real world scene and event classification problems in a more effective manner.

Fig. 7: Example of behavior of two sets of particles propagated using Particle Filtering (PF), and Swarm Particle Filtering (Swarm). The two sets of particles occupy the same location at the first time frame (from [125]).

The availability of concurrent, hence potentially complementary streams of data is amenable to a more robust analysis, by effectively combining them using appropriate techniques, be it at the input representation-level, the feature-level or the decision-level. Successful applications of such techniques have been realized in various multichannel audio and audiovisual scene analysis tasks.

Yet, a number of research questions remain open in these settings. Notably, it is still not clear how to generically detect when some of the data views are temporarily not reliable (typically noisy or out of focus, with respect to the classes of interest) and which strategies should be developed that can efficiently ignore such views and proceed with the classification (or any other similar data processing) using models which were perhaps trained assuming all views are available.

Also, given the complexity of accurately annotating all data views, especially for instantaneous multi-label event classification tasks, that is when multiple events may occur simultaneously, it is important to consider learning methods that can take advantage of very coarse ground-truth labels, which may have been obtained based on just one of the views, without necessarily being relevant for others. An example of this is the “blind” annotation of the audio track of a video (without considering the images) where sound events may not be visible onscreen at the same time stamps. Multiple instance learning and weakly supervised learning techniques may turn out to be effective learning paradigms to address these difficulties.

References

1. Adavanne, S., Parascandolo, G., Pertila, P., Heittola, T., Virtanen, T.: Sound event detection in multichannel audio using spatial and harmonic features. In: Proc IEEE AASP Chall Detect Classif Acoust Scenes Events (2016)
2. Amir, A., Berg, M., Chang, S.F., Hsu, W., Iyengar, G., Lin, C.Y., Naphade, M., Natsev, A., Neti, C., Nock, H., et al.: Ibm research trecvid-2003 video retrieval system. NIST TRECVID-2003 (2003)
3. Andrew, G., Arora, R., Bilmes, J.A., Livescu, K.: Deep canonical correlation analysis. In: Proc Int Conf Mach Learn (2013)
4. Antonacci, F., Lonoce, D., Motta, M., Sarti, A., Tubaro, S.: Efficient source localization and tracking in reverberant environments using microphone arrays. In: Proc IEEE Int Conf Acoust Speech Signal Process, vol. 4, pp. iv–1061. IEEE (2005)
5. Antonacci, F., Matteucci, M., Migliore, D., Riva, D., Sarti, A., Tagliasacchi, M., Tubaro, S.: Tracking multiple acoustic sources in reverberant environments using regularized particle filter. In: Proc Int Conf Digit Signal Process, pp. 99–102 (2007)
6. Arai, T., Hodoshima, H., Yasu, K.: Using steady-state suppression to improve speech intelligibility in reverberant environments for elderly listeners. IEEE Trans Audio Speech Lang Process **18**(7), 1775–1780 (2010)
7. Argones Rúa, E., Bredin, H.H., García Mateo, C., Chollet, G.G., González Jiménez, D.: Audio-visual speech asynchrony detection using co-inertia analysis and coupled hidden markov models. Pattern Anal Appl **12**(3), 271–284 (2008)
8. Arulampalam, M., Maskell, S., Gordon, N., Clapp, T.: A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. IEEE Trans Signal Process **50**(2), 174–188 (2002)
9. Asoh, H., Asano, F., Yoshimura, T., Yamamoto, K., Motomura, Y., Ichimura, N., Hara, I., Ogata, J.: An application of a particle filter to bayesian multiple sound source tracking with audio and video information fusion. In: Proc Fusion, pp. 805–812. Citeseer (2004)
10. Atrey, P.K., Hossain, M.A., El Saddik, A., Kankanhalli, M.S.: Multimodal fusion for multimedia analysis: a survey. Multimed Syst **16**(6), 345–379 (2010)
11. Barzelay, Z., Schechner, Y.Y.: Harmony in motion. In: Proc IEEE Int Conf Comput Vision Pattern Recognit, pp. 1–8 (2007)
12. Beck, A., Stoica, P., Li, J.: Exact and approximate solutions of source localization problems. IEEE Trans Signal Process **56**(5), 1770–1778 (2008)
13. Benmokhtar, R., Huet, B.: Neural network combining classifier based on Dempster-Shafer theory for semantic indexing in video content. Adv Multimed Model pp. 196–205 (2006)
14. Bertin, N., Badeau, R., Vincent, E.: Enforcing harmonicity and smoothness in Bayesian non-negative matrix factorization applied to polyphonic music transcription. IEEE Trans Audio Speech Lang Process **18**(3), 538–549 (2010)
15. Bießmann, F., Meinecke, F.C., Gretton, A., Rauch, A., Rainer, G., Logothetis, N.K., Müller, K.R.: Temporal kernel cca and its application in multimodal neuronal data analysis. Mach Learn **79**(1–2), 5–27 (2010)
16. Bitzer, J., Simmer, K.U.: Superdirective microphone arrays. In: Microphone arrays, pp. 19–38. Springer (2001)
17. Bitzer, J., Simmer, K.U., Kammeyer, K.D.: Theoretical noise reduction limits of the generalized sidelobe canceller (GSC) for speech enhancement. In: Proc IEEE Int Conf Acoust Speech Signal Process, pp. 2965–2968 (1999)
18. Blandin, C., Ozerov, A., Vincent, E.: Multi-source tdoa estimation in reverberant audio using angular spectra and clustering. Signal Process **92**(8), 1950–1960 (2012)
19. Bofill, P., Zibulevsky, M.: Underdetermined blind source separation using sparse representations. Signal Process **81**(11), 2353–2362 (2001)
20. Bousmalis, K., Morency, L.P.: Modeling hidden dynamics of multimodal cues for spontaneous agreement and disagreement recognition. In: Proc Int Conf Autom Face Gesture Recognit, pp. 746–752 (2011)

21. Bredin, H., Chollet, G.: Measuring audio and visual speech synchrony: methods and applications. *Proc IET Int Conf Vis Inf Eng* pp. 255–260 (2006)
22. Bregman, A.S.: Auditory scene analysis: The perceptual organization of sound. MIT press (1994)
23. Brutti, A., Omologo, M., Svaizer, P.: Localization of multiple speakers based on a two step acoustic map analysis. In: *Proc IEEE Int Conf Acoust Speech Signal Process*, pp. 4349–4352 (2008)
24. Canclini, A., Antonacci, F., Sarti, A., Tubaro, S.: Acoustic source localization with distributed asynchronous microphone networks. *IEEE Trans Audio Speech Lang Process* **21**(2), 439–443 (2013)
25. Canclini, A., Bestagini, P., Antonacci, F., Compagnoni, M., Sarti, A., Tubaro, S.: A robust and low-complexity source localization algorithm for asynchronous distributed microphone networks. *IEEE/ACM Trans Audio Speech Lang Process* **23**(10), 1563–1575 (2015)
26. Capon, J.: High-resolution frequency-wavenumber spectrum analysis. *Proc IEEE* **57**(8), 1408–1418 (1969)
27. Capon, J.: High-resolution frequency-wavenumber spectrum analysis. *Proc IEEE* **57**(8), 1408–1418 (1969)
28. Carter, G.C.: Coherence and time delay estimation. *Proc IEEE* **75**(2), 236–255 (1987)
29. Casanovas, A., Monaci, G., Vanderghenst, P., Gribonval, R.: Blind audiovisual source separation based on sparse redundant representations. *IEEE Trans Multimed* **12**(5), 358–371 (2010)
30. Casanovas, A.L., Vanderghenst, P.: Nonlinear video diffusion based on audio-video synchrony. *IEEE Trans Multimed* (2010)
31. Chang, S.F., Ellis, D., Jiang, W., Lee, K., Yanagawa, A., Loui, A.C., Luo, J.: Large-scale multimodal semantic concept detection for consumer video. In: *Proc Int Workshop Multimed Inf Retr, MIR '07*, pp. 255–264. ACM, New York, NY, USA (2007)
32. Chibelushi, C.C., Mason, J.S.D., Deravi, N.: Integrated person identification using voice and facial features. In: *Proc IEE Colloq Image Process Secur Appl*, pp. 4/1–4/5 (1997)
33. Choudhury, T., Reh, J.M., Pavlovic, V., Pentland, A.: Boosting and structure learning in dynamic Bayesian networks for audio-visual speaker detection. In: *Proc IEEE Int Conf Pattern Recognit*, vol. 3, pp. 789–794 (2002)
34. Cichocki, A., Zdunek, R., Amari, S.: Nonnegative Matrix and Tensor Factorization. *IEEE Signal Process Mag* **25**(1), 142–145 (2008)
35. Cobos, M., Marti, A., Lopez, J.J.: A modified srp-phat functional for robust real-time sound source localization with scalable spatial sampling. *IEEE Signal Process Lett* **18**(1), 71–74 (2011)
36. Compagnoni, M., Bestagini, P., Antonacci, F., Sarti, A., Tubaro, S.: Localization of acoustic sources through the fitting of propagation cones using multiple independent arrays. *IEEE Trans Audio Speech Lang Process* **20**(7), 1964–1975 (2012)
37. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. Wiley (2006)
38. Cox, H., Zeskind, R., Kooij, T.: Practical supergain. *IEEE Trans Acoust Speech Signal Process* **34**(3), 393–398 (1986)
39. Cristani, M., Bicego, M., Murino, V.: Audio-visual event recognition in surveillance video sequences. *IEEE Trans Multimed* **9**(2), 257–267 (2007)
40. Crocco, M., Bue, A.D., Murino, V.: A bilinear approach to the position self-calibration of multiple sensors. *IEEE Trans Signal Process* **60**(2), 660–673 (2012)
41. Cutler, R., Davis, L.: Look who's talking: Speaker detection using video and audio correlation. In: *Proc IEEE Int Conf Multimed Expo*, vol. 3, pp. 1589–1592. IEEE (2000)
42. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *Proc IEEE Int Conf Comput Vision Pattern Recognit*, vol. 1, pp. 886–893. IEEE (2005)
43. D'Arca, E., Robertson, N., Hopgood, J.: Look who's talking: Detecting the dominant speaker in a cluttered scenario. In: *Proc IEEE Int Conf Acoust Speech Signal Process* (2014)
44. DiBiase, J., Silverman, H., Brandstein, M.: Robust localization in reverberant rooms. In: *Microphone Arrays*, pp. 157–180. Springer (2001)

45. Dmochowski, J., Benesty, J., Affes, S.: A generalized steered response power method for computationally viable source localization. *IEEE Trans Audio Speech Language Process* **15**(8), 2510–2526 (2007)
46. Do, H., Silverman, H., Yu, Y.: A real-time srp-phat source location implementation using stochastic region contraction (src) on a large-aperture microphone array. In: *Proc Int Conf Acoust Speech Signal Process*, vol. 1, pp. 1121 – 1124. IEEE (2007)
47. Doclo, S., Moonen, M.: Gsvd-based optimal filtering for single and multimicrophone speech enhancement. *IEEE Trans Signal Process* **50**(9), 2230–2244 (2002)
48. Duong, N.Q.K., Vincent, E., Gribonval, R.: Under-determined reverberant audio source separation using a full-rank spatial covariance model. *IEEE Trans Audio Speech Lang Process* **18**(7), 1830–1840 (2010)
49. Duong, N.Q.K., Vincent, E., Gribonval, R.: Spatial location priors for gaussian model based reverberant audio source separation. *EURASIP J Adv Signal Process* **2013**(1), 1–11 (2013)
50. Elko, G.W.: Spatial coherence functions for differential microphones in isotropic noise fields. In: *Microphone Arrays: Signal Processing Techniques and Applications*, pp. 61–85. Springer (2001)
51. Feichtenhofer, C., Pinz, A., Zisserman, A.: Convolutional two-stream network fusion for video action recognition. *arXiv preprint arXiv:1604.06573* (2016)
52. Févotte, C., Cardoso, J.F.: Maximum likelihood approach for blind audio source separation using time-frequency Gaussian models. In: *Proc IEEE Workshop Appl Signal Process Audio Acoust*, pp. 78–81 (2005)
53. Fisher, J., Darrell, T., Freeman, W.T., Viola, P., Fisher III, J.W.: Learning Joint Statistical Models for Audio-Visual Fusion and Segregation. In: *Proc Adv Neural Inf Process Syst, ML*, pp. 772–778 (2001)
54. FitzGerald, D., Cranitch, M., Coyle, E.: Extended Nonnegative Tensor Factorisation Models for Musical Sound Source Separation. *Comput Intell Neurosci* **2008** (2008)
55. Fitzgerald, D., Cranitch, M., Coyle, E.: Using tensor factorisation models to separate drums from polyphonic music. In: *Proc Int Conf Digit Audio Eff* (2009)
56. Foucher, S., Lalibert, F., Boulianne, G., Gagnon, L.: A Dempster-Shafer Based Fusion Approach for Audio-Visual Speech Recognition with Application to Large Vocabulary French Speech. In: *Proc IEEE Int Conf Acoust Speech Signal Process* (2006)
57. Frost, O.L.: An algorithm for linearly constrained adaptive array processing. *Proc IEEE* **60**(8), 926–935 (1972)
58. Gandhi, A., Sharma, A., Biswas, A., Deshmukh, O.: Gethr-net: A generalized temporally hybrid recurrent neural network for multimodal information fusion. *arXiv preprint arXiv:1609.05281* (2016)
59. Gehrig, T., Nickel, K., Ekenel, H., Klee, U., McDonough, J.: Kalman filters for audio-video source localization. In: *Proc IEEE Workshop Appl Signal Process Audio Acoust*, pp. 118–121. IEEE (2005)
60. Goecke, R., Millar, J.B.: Statistical Analysis of the Relationship between Audio and Video Speech Parameters for Australian English. In: *Proc ISCA Tutor Res Workshop Audit-Vis Speech Process*, pp. 133–138 (2003)
61. Gowdy, J.N., Subramanya, A., Bartels, C., Bilmes, J.A.: DBN based multi-stream models for audio-visual speech recognition. In: *Proc IEEE Int Conf Acoust Speech Signal Process* (2004)
62. Gravier, G., Potamianos, G., Neti, C.: Asynchrony modeling for audio-visual speech recognition. In: *Proc Int Conf Hum Lang Technol Res*, pp. 1–6. Morgan Kaufmann Publishers Inc., San Diego, California (2002)
63. Gribonval, R., Zibulevsky, M.: Sparse component analysis. In: *Handbook of Blind Source Separation, Independent Component Analysis and Applications*, pp. 367–420. Academic Press (2010)
64. Griffiths, L., Jim, C.: An alternative approach to linearly constrained adaptive beamforming. *IEEE Trans Antennas Propag* **30**(1), 27–34 (1982)
65. Gustafsson, T., Rao, B.D., Trivedi, M.: Source localization in reverberant environments: Modeling and statistical analysis. *IEEE Trans Speech Audio Process* **11**, 791–803 (2003)

66. Hardoon, D.R., Szedmak, S., Shawe-Taylor, J.: Canonical correlation analysis: an overview with application to learning methods. *Neural Comput* **16**(12), 2639–2664 (2004)
67. Haykin, S.: *Adaptive Filter Theory*, 5 edn. Pearson Education (2014)
68. Haykin, S., Justice, J.H., Owsley, N.L., Yen, J., Kak, A.C.: *Array signal processing*. Prentice-Hall, Inc., Englewood Cliffs, NJ (1985)
69. Hotelling, H.: Relations between two sets of variates. *Biometrika* **28**(3 - 4), 321–377 (1936)
70. Hu, D., Li, X., lu, X.: Temporal multimodal learning in audiovisual speech recognition. In: *Proc IEEE Conf Comput Vis Pattern Recognit* (2016)
71. Huang, P.S., Zhuang, X., Hasegawa-Johnson, M.: Improving acoustic event detection using generalizable visual features and multi-modality modeling. In: *Proc IEEE Int Conf Acoust Speech Signal Process*, pp. 349–352. IEEE (2011)
72. Huang, Y., Benesty, J., Elko, G., Mersereau, R.: Real-time passive source localization: A practical linear-correction least-squares approach. *IEEE Trans Speech Audio Process* **9**(8), 943–956 (2001)
73. Ivanov, Y., Serre, T., Bouvrie, J.: Error weighted classifier combination for multi-modal human identification. Tech. Rep. MIT-CSAIL-TR-2005-081, MIT (2005)
74. Izadinia, H., Saleemi, I., Shah, M.: Multimodal analysis for identification and segmentation of moving-sounding objects. *IEEE Trans Multimed* **15**(2), 378–390 (2013)
75. Izumi, Y., Ono, N., Sagayama, S.: Sparseness-based 2CH BSS using the EM algorithm in reverberant environment. In: *Proc IEEE Workshop Appl Signal Process Audio Acoust*, pp. 147–150 (2007)
76. Jaureguiberry, X., Vincent, E., Richard, G.: Fusion methods for speech enhancement and audio source separation. *IEEE/ACM Trans Audio Speech Lang Process* **24**(7), 1266–1279 (2016)
77. Jhuo, I.H., Ye, G., Gao, S., Liu, D., Jiang, Y.G., Lee, D., Chang, S.F.: Discovering joint audio–visual codewords for video event detection. *Mach Vision Appl* **25**(1), 33–47 (2014)
78. Jiang, W., Cotton, C., Chang, S.F., Ellis, D., Loui, A.: Short-term audiovisual atoms for generic video concept classification. In: *Proc ACM Int Conf Multimed*, pp. 5–14. ACM (2009)
79. Jiang, W., Loui, A.C.: Audio-visual grouplet: temporal audio-visual interactions for general video concept classification. In: *Proc ACM Int Conf Multimed*, pp. 123–132. Scottsdale, USA (2011)
80. Jiang, Y.G., Bhattacharya, S., Chang, S.F., Shah, M.: High-level event recognition in unconstrained videos. *Int J Multimed Inf Retr* **2**(2), 73–101 (2013)
81. Jiang, Y.G., Zeng, X., Ye, G.: Columbia-ucf trecvid2010 multimedia event detection: Combining multiple modalities, contextual concepts, and temporal matching. In: *Proc NIST TRECVID-2003* (2003)
82. Joder, C., Essid, S., Richard, G.: Temporal Integration for Audio Classification with Application to Musical Instrument Classification. *IEEE Trans Audio Speech Lang Process* (2008)
83. Jourjine, A., Rickard, S., Yilmaz, O.: Blind separation of disjoint orthogonal signals: Demixing N sources from 2 mixtures. In: *Proc IEEE Int Conf Acoust Speech Signal Process*, pp. 2985–2988 (2000)
84. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: *Proc IEEE Conf Comput Vision Pattern Recognit*, pp. 1725–1732 (2014)
85. Kay, J.: Feature discovery under contextual supervision using mutual information. In: *Proc Int Jt Conf Neural Netw*, vol. 4, pp. 79–84 vol.4 (1992)
86. Kidron, E., Schechner, Y., Elad, M.: Pixels that sound. In: *Proc IEEE Conf Comput Vision Pattern Recognit*, vol. 1, pp. 88–95 vol. 1 (2005)
87. Kijak, E., Gravier, G., Gros, P., Oisel, L., Bimbot, F.: HMM based structuring of tennis videos using visual and audio cues. In: *Proc IEEE Int Conf Multimed Expo*, pp. 309–312. IEEE Computer Society (2003)
88. Kittler, J., Hatef, M., Duin, R.P.W., Matas, J.: On Combining Classifiers. *IEEE Trans Pattern Anal Mach Intell* **20**(3), 226–239 (1998)

89. Klee, U., T. G., McDonough, J.: Kalman filters for time delay of arrival-based source localization. *EURASIP J Adv Signal Process* **2006**(1), 1–15 (2006)
90. Knapp, C., Carter, G.: The generalized correlation method for estimation of time delay. *IEEE Trans Acoust Speech Signal Process* **24**(4), 320–327 (1976)
91. Kolda, T.G., Bader, B.W.: Tensor decompositions and applications. *SIAM review* **51**(3), 455–500 (2009)
92. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Proc Adv Neural Inf Process Syst*, pp. 1097–1105 (2012)
93. Kuhn, G.F.: Model for the interaural time differences in the azimuthal plane. *J Acoust Soc Am* **62**(1), 157–167 (1977)
94. Lai, P.L., Fyfe, C.: Kernel and nonlinear canonical correlation analysis. *Int J Neural Syst* **10**(5), 365–378 (2000)
95. Levy, A., Gannot, S., Habets, E.: Multiple-hypothesis extended particle filter for acoustic source localization in reverberant environments. *IEEE Trans Audio Speech Lang Process* **19**(6), 1540–1555 (2011)
96. Li, D., Dimitrova, N., Li, M., Sethi, I.: Multimedia content processing through cross-modal association. In: *Proc ACM Int Conf Multimed*. Berkeley, CA, USA (2003)
97. Lim, A., Nakamura, K., Nakadai, K., Ogata, T., Okuno, H.G.: Audio-visual musical instrument recognition. In: *Proc Natl Conv Audio-V Inf Process Soc* (2011)
98. Liu, Q., Wang, W., Jackson, P.J., Barnard, M., Kittler, J., Chambers, J.: Source separation of convolutive and noisy mixtures using audio-visual dictionary learning and probabilistic time-frequency masking. *IEEE Trans Signal Process* **61**(22), 5520–5535 (2013)
99. Liutkus, A., Durrieu, J.L., Daudet, L., Richard, G.: An overview of informed audio source separation. In: *Proc Int Workshop Image Anal Multimed Interact Serv*, pp. 1–4. IEEE (2013)
100. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int J Comput Vis* **60**(2), 91–110 (2004)
101. Mahadevan, V., Li, W., Bhalodia, V., Vasconcelos, N.: Anomaly detection in crowded scenes. In: *Proc IEEE Conf Comput Vision Pattern Recognit*, vol. 249, p. 250 (2010)
102. Makino, S., Lee, T.W., Sawada, H.: *Blind Speech Separation*. Springer (2007)
103. Mandel, M., Bressler, S., Shinn-Cunningham, B., Ellis, D.: Evaluating source separation algorithms with reverberant speech. *IEEE Trans Audio Speech Lang Process* **18**(7), 1872–1883 (2010)
104. Mandel, M., Ellis, D.: EM localization and separation using interaural level and phase cues. In: *Proc IEEE Workshop Appl Signal Process Audio Acoust*, pp. 275–278 (2007)
105. Maragos, P., Gros, P., Katsamanis, A., Papandreou, G.: Cross-modal integration for performance improving in multimedia: A review. In: *Multimodal processing and interaction*, pp. 1–46. Springer (2008)
106. Marti, A., Cobos, M., Lopez, J., Escolano, J.: A steered response power iterative method for high-accuracy acoustic source localization. *The Journal of the Acoustical Society of America* **134**(4), 2627–2630 (2013)
107. Metallinou, A., Lee, S., Narayanan, S.: Decision level combination of multiple modalities for recognition and analysis of emotional expression. In: *Proc IEEE Int Conf Acoust Speech Signal Process*, pp. 2462–2465 (2010)
108. Milani, S., Fontani, M., Bestagini, P., Barni, M., Piva, A., Tagliasacchi, M., Tubaro, S.: An overview on video forensics. *APSIPA Trans Signal Inf Process* **1**, e2 (2012)
109. Monaci, G., Jost, P., Vanderghenst, P., Mailhé, B., Lesage, S., Gribonval, R.: Learning multimodal dictionaries. *IEEE Trans Image Process* **16**(9), 2272–2283 (2007)
110. Monaci, G., Vanderghenst, P.: Audiovisual gestalts. In: *Proc IEEE Conf Comput Vision Pattern Recognit*, pp. 200–200 (2006)
111. Monaci, G., Vanderghenst, P., Sommer, F.T.: Learning bimodal structure in audio-visual data. *IEEE Trans Neural Netw* **20**(12), 1898–1910 (2009)
112. Moore, B.C.J.: *Introduction to the psychology of hearing* / Brian C. J. Moore. Macmillan London (1977)
113. Murphy, K.P.: *Dynamic Bayesian Networks: Representation, Inference and Learning*. Ph.D. thesis, University of California, Berkeley (2002)

114. Naphade, M.R., Garg, A., Huang, T.S.: Audio-visual event detection using duration dependent input output markov models. In: Proc IEEE Workshop Content-Based Access Image Video Libr, pp. 39–43. IEEE (2001)
115. Nefian, A.V., Liang, L., Pi, X., Xiaoxiang, L., Mao, C., Murphy, K.P.: A coupled {HMM} for audiovisual speech recognition. In: Proc IEEE Int Conf Acoust Speech Signal Process, vol. 2. IEEE (2002)
116. Ngiam, J., Khosla, A., Kim, M., Nam, J., Lee, H., Ng, A.Y.: Multimodal deep learning. In: Proc Int Conf Mach Learn, pp. 689–696 (2011)
117. Nguyen, V.T., Nguyen, D.L., Tran, M.T., Le, D.D., Duong, D.A., Satoh, S.: Query-adaptive late fusion with neural network for instance search. In: Proc IEEE Int Workshop Multimed Signal Process, pp. 1–6. IEEE (2015)
118. Nikunen, J., Virtanen, T.: Direction of arrival based spatial covariance model for blind sound source separation. *IEEE/ACM Trans Audio Speech Lang Process* **22**(3), 727–739 (2014)
119. Omologo, M., Svaizer, P.: Acoustic event localization using a crosspower-spectrum phase based technique. In: Proc IEEE Int Conf Acoust Speech Signal Process, vol. 2. IEEE (1994)
120. Otsuka, T., Ishiguro, K., Sawada, H., Okuno, H.G.: Bayesian nonparametrics for microphone array processing. *IEEE/ACM Trans Audio Speech Lang Proc* **22**(2), 493–504 (2014)
121. Ozerov, A., Févotte, C.: Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation. *IEEE Trans Audio Speech Lang Process* **18**(3), 550–563 (2010)
122. Ozerov, A., Févotte, C., Blouet, R., Durrieu, J.L.: Multichannel nonnegative tensor factorization with structured constraints for user-guided audio source separation. In: Proc IEEE Int Conf Acoust Speech Signal Process. Prague, Czech Republic (2011)
123. Ozerov, A., Vincent, E., Bimbot, F.: A general flexible framework for the handling of prior information in audio source separation. *IEEE Trans Audio Speech Lang Process* **20**(4), 1118–1133 (2012)
124. Parekh, S., Essid, S., Ozerov, A., Duong, N.Q.K., Pérez, P., Richard, G.: Motion informed audio source separation. In: Under Review (2017)
125. Parisi, R., Croene, P., Uncini, A.: Particle swarm localization of acoustic sources in the presence of reverberation. In: Proc IEEE Int Symp Circuits Syst, pp. 4–pp. IEEE (2006)
126. Parra, L., Spence, C.: Convolutional blind separation of non-stationary sources. *IEEE Trans Speech Audio Process* **8**(3), 320–327 (2000)
127. Pertilä, P., Mieskolainen, M., Hämäläinen, M.: Closed-form self-localization of asynchronous microphone arrays. In: Proc Joint Workshop Hands-free Speech Commun Microphone Arrays, pp. 139–144. IEEE (2011)
128. P. Stoica, Moses, R.: Spectral analysis of signals. Pearson Prentice Hall Upper Saddle River, NJ (2005)
129. Rocha, A., Scheirer, W., Boulton, T., Goldenstein, S.: Vision of the unseen: Current trends and challenges in digital image and video forensics. *ACM Comput Surv* **43**(4), 26 (2011)
130. Rokach, L.: Ensemble-based classifiers. *Artif Intell Rev* (2010)
131. Roy, R., Kailath, T.: Esprit-estimation of signal parameters via rotational invariance techniques. *IEEE Trans Acoust Speech Signal Process* **37**(7), 984–995 (1989)
132. Sadlier, D.A., O'Connor, N.E.: Event detection in field sports video using audio-visual features and a support vector machine. *IEEE Trans Circuits Syst Video Technol* **15**(10), 1225–1233 (2005)
133. Sawada, H., Mukai, R., Araki, S., Makino, S.: A robust and precise method for solving the permutation problem of frequency-domain blind source separation. *IEEE Trans Speech Audio Process* **12**(5), 530–538 (2004)
134. Schau, H., Robinson, A.: Passive source localization employing intersecting spherical surfaces from time-of-arrival differences. *IEEE Trans Acoust Speech Signal Process* **35**(8), 1223–1225 (1987)
135. Scheuing, J., Yang, B.: Disambiguation of tdoa estimation for multiple sources in reverberant environments. *IEEE Trans Audio Speech Lang Process* **16**(8), 1479–1489 (2008)
136. Schmidt, R.: Multiple emitter location and signal parameter estimation. *IEEE Trans Antennas Propag* **34**(3), 276–280 (1986)

137. Sedighin, F., Babaie-Zadeh, M., Rivet, B., Jutten, C.: Two multimodal approaches for single microphone source separation. In: *Proc Eur Signal Process Conf* (2016)
138. Seichepine, N., Essid, S., Févotte, C., Cappe, O.: Soft nonnegative matrix co-factorization with application to multimodal speaker diarization. In: *Proc IEEE Int Conf Acoust Speech Signal Process*. Vancouver (2013)
139. Seichepine, N., Essid, S., Févotte, C., Cappe, O.: Soft nonnegative matrix co-factorization. *IEEE Trans Signal Process* **PP**(99) (2014)
140. Serizel, R., Bisot, V., Essid, S., Richard, G.: Machine listening techniques as a complement to video image analysis in forensics. In: *Proc IEEE Int Conf Image Process*, pp. 948–952. IEEE (2016)
141. Serizel, R., Moonen, M., van Dijk, B., Wouters, J.: Low-rank approximation based multi-channel wiener filter algorithms for noise reduction with application in cochlear implants. *IEEE/ACM Trans Audio Speech Lang Process* **22**(4), 785–799 (2014)
142. Showen, R., Calhoun, R., Dunham, J.: Acoustic location of gunshots using combined angle of arrival and time of arrival measurements (2009). US Patent 7,474,589
143. Sigg, C., Fischer, B., Ommer, B., Roth, V., Buhmann, J.: Nonnegative CCA for Audiovisual Source Separation. In: *Proc IEEE Workshop Mach Learn Signal Process*, pp. 253–258. IEEE (2007)
144. Smaragdis, P., Casey, M.: Audio Visual Independent Components. In: *Proc Int Symp Indep Compon Anal Blind Signal Sep*, pp. 709–714 (2003)
145. Srivastava, N., Salakhutdinov, R.R.: Multimodal learning with deep boltzmann machines. In: *Proc Adv Neural Inf Process Syst*, pp. 2222–2230 (2012)
146. Strobil, N., Spors, S., Rabenstein, R.: Joint audio-video object localization and tracking. *IEEE Signal Processing Mag* **18**(1), 22–31 (2001)
147. Tian, Y., Chen, Z., Yin, F.: Distributed kalman filter-based speaker tracking in microphone array networks. *Appl Acoust* **89**, 71–77 (2015)
148. Togami, M., Hori, K.: Multichannel semi-blind source separation via local Gaussian modeling for acoustic echo reduction. In: *Proc Eur Signal Process Conf* (2011)
149. Togami, M., Kawaguchi, Y.: Simultaneous optimization of acoustic echo reduction, speech dereverberation, and noise reduction against mutual interference. *IEEE/ACM Trans Audio Speech Lang Process* **22**(11), 1612–1623 (2014)
150. Trifa, V., Koene, A., Moren, J., Cheng, G.: Real-time acoustic source localization in noisy environments for human-robot multimodal interaction. In: *Proc IEEE Int Symp Robot Hum Interact Commun* (2007)
151. Valente, S., Tagliasacchi, M., Antonacci, F., Bestagini, P., Sarti, A., Tubaro, S.: Geometric calibration of distributed microphone arrays from acoustic source correspondences. In: *Proc IEEE Int Workshop Multimed Signal Process*, pp. 13–18 (2010)
152. Valin, J., Michaud, F., Rouat, J.: Robust 3d localization and tracking of sound sources using beamforming and particle filtering. In: *Proc IEEE Int Conf Acoust Speech Signal Process*, vol. 4. IEEE (2006)
153. Velivelli, A., Ngo, C.W., Huang, T.S.: Detection of documentary scene changes by audio-visual fusion. In: *Proc Int Conf Image Video Retr*, pp. 227–238. Springer (2003)
154. Vincent, E., Bertin, N., Gribonval, R., Bimbot, F.: From blind to guided audio source separation: How models and side information can improve the separation of sound. *IEEE Signal Process Mag* **31**(3), 107–115 (2014)
155. Vuegen, L., Broeck, B.V.D., Karsmakers, P., hamme, H.V., Vanrumste, B.: Automatic monitoring of activities of daily living based on real-life acoustic sensor data: a preliminary study. In: *Proc Int Workshop Speech Lang Process Assist Technol*, pp. 113–118 (2013)
156. Wang, D.L.: Time-frequency masking for speech separation and its potential for hearing aid design. *Trends Amplif* pp. 332–352 (2008)
157. Wang, H., Chu, P.: Voice source localization for automatic camera pointing system in video-conferencing. In: *Proc IEEE Int Conf Acoust Speech Signal Process* (1997)
158. Wang, H., Kläser, A., Schmid, C., Liu, C.L.: Dense trajectories and motion boundary descriptors for action recognition. *Int J Comput Vis* **103**(1), 60–79 (2013)

159. Ward, D.B., Lehmann, E.A., Williamson, R.C.: Particle filtering algorithms for tracking an acoustic source in a reverberant environment. *IEEE Trans Speech Audio Process* **11**(6), 826–836 (2003)
160. Wilkins, P., Adamek, T., Byrne, D., Jones, G., Lee, H., Keenan, G., McGuinness, K., O'Connor, N.E., Smeaton, A.F., Amin, A., Obrenovic, Z., Benmokhtar, R., Galmar, E., Huet, B., Essid, S., Landais, R., Vallet, F., Papadopoulos, G.T., Vrochidis, S., Mezaris, V., Kompatsiaris, I., Spyrou, E., Avrithis, Y., Morzinger, R., Schallauer, P., Bailer, W., Piatrik, T., Chandramouli, K., Izquierdo, E., Haller, M., Goldmann, L., Samour, A., Cobet, A., Sikora, T., Praks, P., Kompatsiaris, Y.: K-Space at TRECVID 2007. In: *TRECVID 2007* (2007)
161. Wu, Y., Lin, C.Y.Y., Chang, E.Y., Smith, J.R.: Multimodal information fusion for video concept detection. In: *Proc IEEE Int Conf Image Process*, vol. 4, pp. 2391–2394. IEEE, Singapore (2004)
162. Wu, Z., Jiang, Y.G., Wang, J., Pu, J., Xue, X.: Exploring inter-feature and inter-class relationships with deep neural networks for video classification. In: *Proc ACM Int Conf Multimed*, pp. 167–176. ACM (2014)
163. Yilmaz, K., Cemgil, A.T.: Probabilistic latent tensor factorisation. In: *Proc Int Conf Latent Var Anal Signal Sep*, pp. 346–353 (2010)
164. Yokoya, N., Yairi, T., Iwasaki, A.: Coupled Nonnegative Matrix Factorization Unmixing for Hyperspectral and Multispectral Data Fusion. *IEEE Trans Geosci Remote Sens* **50**(2), 528–537 (2012)
165. Yoo, J., Choi, S.: Matrix co-factorization on compressed sensing. In: *Proc Int Joint Conf Artif Intell* (2011)
166. Yost, W.A.: Discriminations of interaural phase differences. *J Acoust Soc Am* **55**(6), 1299–1303 (1974)
167. Yuhas, B.P., Goldstein, M.H., Sejnowski, T.J.: Integration of acoustic and visual speech signals using neural networks. *IEEE Commun Mag* **27**(11), 65–71 (1989)
168. Zhang, Q., Chen, Z., Yin, F.: Distributed marginalized auxiliary particle filter for speaker tracking in distributed microphone networks. *IEEE/ACM Trans Audio Speech Lang Process* **24**(11), 1921–1934 (2016)
169. Zotkin, D.N., Duraiswami, R.: Accelerated speech source localization via a hierarchical search of steered response power. *IEEE Trans Speech Audio Process* **12**(5), 499–508 (2004)