Applied Acoustics 151 (2019) 73-86

Contents lists available at ScienceDirect

**Applied Acoustics** 

journal homepage: www.elsevier.com/locate/apacoust

# Directional embedding based semi-supervised framework for bird vocalization segmentation

# Anshul Thakur\*, Padmanabhan Rajan

School of Computing and Electrical Engineering, Indian Institute of Technology Mandi, Himachal Pradesh, India

#### ARTICLE INFO

Article history: Received 16 October 2018 Accepted 23 February 2019

Keywords: Bird vocalization segmentation Bioacoustics Directional embedding

# ABSTRACT

This paper proposes a data-efficient, semi-supervised, two-pass framework for segmenting bird vocalizations. The framework utilizes a binary classification model to categorize frames of an input audio recording into the background or bird vocalization. The first pass of the framework automatically generates training labels from the input recording itself, while model training and classification is done during the second pass. The proposed framework utilizes a reference directional model for obtaining a feature representation called directional embeddings (DE). This reference directional model acts as an acoustic model for bird vocalizations and is obtained using the mixtures of Von-Mises Fisher distribution (moVMF). The proposed DE space only contains information about bird vocalizations, while no information about the background disturbances is reflected. The framework employs supervised information only for obtaining the reference directional model and avoids the background modeling. Hence, it can be regarded as semi-supervised in nature. The proposed framework is tested on approximately 79,000 vocalizations of seven different SNRs. Experimental results convey that the proposed framework performs better than the existing bird vocalization segmentation methods.

© 2019 Elsevier Ltd. All rights reserved.

# 1. Introduction

Birds are of significant importance in maintaining the equilibrium of various ecosystems. However, many bird species are facing the threat of population decline due to habitat destruction and global climate change [1]. As a result, there is a pressing need to study the effects of human activities and climate change on avian diversity. Surveying and monitoring of bird species in their natural habitat are generally the first steps in analyzing these effects [2]. Many bird species are highly vocal, hence, acoustic monitoring provides a passive and convenient way to monitor birds. The development of sophisticated programmable recording devices has made acoustic monitoring feasible. Acoustic data recorded by these devices can be analyzed to perform various tasks such as species identification, the tracking of migrant species or in examining the avian biodiversity of a given region. Manual processing of this data is tedious and time-consuming, and requires considerable knowledge of bird vocalizations. Hence, it is essential to develop automatic methods for processing the recorded acoustic data.

Segmenting bird vocalizations and discarding the background is the principal step in many bioacoustic tasks such as studying the semantics of bird songs or species identification. The performance of this segmentation step directly influences the overall performance of other subsequent tasks [3]. Bird vocalization segmentation appears to be a specialized version of sound event detection (SED). These two tasks differ from each other in nature of the target acoustic events. The typical target events, such as gun-shot, babycry, glass-break and dog-bark, in any SED task (such as the DCASE challenges [4]) are well-defined i.e. the knowledge about the temporal-frequency structure of the target events is present as prior information. However, in bird vocalization segmentation, the nature of the target event (bird vocalizations) is not well defined. This can be attributed to the large amount of variations present in the frequency profiles of vocalizations of different species. This task of bird vocalization segmentation is often challenging in the field conditions due to reasons listed below:

• *Background modeling*: The presence of various biotic (e.g. vocalizations of other animals) and abiotic (such as rain, wind or human-induced sounds) background sounds make it difficult to isolate the bird vocalizations from the background. Irrespec-







<sup>\*</sup> Corresponding author. E-mail addresses: anshul\_thakur@students.iitmandi.ac.in (A. Thakur), padman@iitmandi.ac.in (P. Rajan).

tive of the source, these background disturbances are unpredictable and highly volatile. This makes the acoustic modeling of the background non-trivial.

- *Bird vocalization modeling*: The vocalizations of different bird species can be significantly different from each other. For example, the temporal-frequency modulations present in Cassin's vireo (a North American song bird) vocalizations are completely different from the temporal-frequency modulations present in vocalizations of Greater sooty owl and Hooded crow. Apart from that, many species exhibit a large repertoire of vocalization types, which results in high intra-class variations. Hence, coming up with one universal acoustic model that can be used for segmenting the vocalizations of many bird species is not straight-forward.
- Lack of labeled training data: Like other bioacoustic tasks, bird vocalization segmentation also suffers from the scarcity of labeled training data. This limits the use of data-intensive state-of-art audio classification frameworks such as convolutional neural networks (CNN). In recent times, large datasets with weak labels (i.e. presence or absence of birdcalls) have been released for the audio tagging tasks (also known as *bird activity detection*) [5]. These datasets lack vocalization-level time-stamps or frame-level labels and can be used for training powerful neural networks in a weakly supervised setup [6] for segmenting bird vocalizations. However, the segmentation performance of these weakly supervised neural networks is not on-par with their supervised counterparts.

In this work, we propose a data-efficient semi-supervised twopass framework which overcomes the aforementioned challenges associated with the task of bird vocalization segmentation. The proposed segmentation framework does not require any explicit background acoustic model and utilizes a reference bird vocalization acoustic model to discriminate between bird and non-bird sounds. This reference vocalization model is obtained using a very small amount of the labeled training data. Since only bird vocalizations are modeled (not the background), the proposed framework can be considered as a semi-supervised framework. This notion of semi-supervision is borrowed from non-negative matrix factorization (NMF) based speech enhancement literature where either speech samples or noise is used for training [7,8].

Given a reference model for bird vocalizations, the proposed framework is a two-pass process. During first pass, the reference vocalization model is utilized to convert the input audio recording into a frame-wise feature representation called directional embeddings (DE). The feature space corresponding to DE retains information only about the bird vocalizations and most of the background disturbances are inherently not reflected in the feature space (more details are in Section 3). Mutual information (MI) is obtained between consecutive pairs of DE frames to get an initial estimate of the regions containing the background and vocalizations respectively. MI exhibits high values for the background and low values for vocalization regions (see Section 4). A fixed number of audio frames exhibiting highest and lowest MI are then automatically labeled as background and vocalizations respectively. These labels and DE features are given as input to the second pass of the framework. In the second pass, a classification model is trained on the DE features using the generated labels. This trained classification model is utilized to classify each input audio frame as background or bird vocalization. Fig. 1 illustrates the proposed two-pass semi-supervised framework for bird vocalization segmentation, where a support vector machine (SVM) is used as the classification model in the second pass.

DE forms the nucleus of the proposed framework. As discussed earlier, DE provides ability to discriminate bird sounds from nonbird background disturbances. Hence, the utilization of DE in the proposed framework helps in decreasing false alarm rate significantly. DE are obtained by projecting the time-frequency representations of an input audio on a dictionary (which forms the reference vocalization model). The atoms of this dictionary are unit vectors pointing in the dominant directions of bird vocalizations. These dominant directions are learned by clustering a small amount of audio frames (henceforth termed as training data) containing bird vocalizations on the unit hypersphere. Clustering is achieved using mixtures of Von-Mises Fisher distribution (moVMF) [9] and the mean directional vector of each cluster forms an atom of the dictionary. Also, moVMF can be seen as equivalent of Gaussian mixture models (GMM) on the unit hypersphere. Clustering the spectral frames on the unit hypersphere helps in capturing the information about the presence/absence of dominant frequency bins rather than the frequency magnitudes. This mitigates the effect of near-field and far-field conditions on the clustering process to some extent. Hence, two similar vocalizations of a bird species, recorded at different distances from the recording device. are clustered together. This behaviour is not present in other clustering techniques such as K-means where magnitude information influences the assignment of a data point to any cluster. Instead of moVMF, earlier studies [10,11] have used spherical K-means to cluster the data for audio feature learning. However, moVMF has been reported to provide better clustering than spherical Kmeans [12]. The other advantage of moVMF is the presence of the concentration parameter (see Section 3) which helps in discarding mean vectors of clusters having very low concentration around the mean. This helps in overcoming the effects of outliers which can be introduced in the learning process of moVMF due to incorrect labels in the training data.

The main contributions of this work are listed below:

- A semi-supervised two pass framework which can be used for segmenting vocalizations of many bird species in various recording environments. The framework does not require any background modeling and utilizes a very small amount of training data to learn a bird vocalization acoustic model. Also, the framework avoids the train-test conditions mismatch by using a classifier trained on the DE features extracted from the input recording itself.
- Treating the audio frames as directional data to mitigate the effect of near-field and far-field conditions. This directional data is obtained by projecting the audio frames on the unit hypersphere.
- Utilizing moVMF as an acoustic model to capture the behaviour of a particular class (in this case, bird vocalizations). To the best of authors' knowledge, this is the first study to treat moVMF as an acoustic model.
- An intermediate feature representation (directional embeddings) suitable to facilitate bird vocalization segmentation in a semi-supervised manner. This feature representation equips the proposed framework with the ability to discriminate between bird and non-bird sounds.

The rest of the paper is organized as follows: In Section 2, various existing studies targeting bird vocalization segmentation and acoustic event segmentation are discussed. The proposed directional embeddings based feature learning is described in Section 3. The proposed framework is discussed in Section 4. Later, we discuss experimental details, results and conclusion in Sections 5–7 respectively.

# 2. Previous studies

The literature contains many studies addressing the task of bird vocalization segmentation. These studies can be categorized into the following:



Fig. 1. Proposed semi-supervised framework for bird vocalization segmentation.

*Manual segmentation:* A few early bioacoustic studies dealing with bioacoustics [13,2] have used manually segmented bird vocalizations. Manual segmentation can be tedious and unfeasible if large amount of audio data has to be processed.

Segmentation using time-domain features: Segmentation methods based on time-domain features such as energy, entropy and inverse spectral flatness (ISF) have been used for segmenting bird vocalizations. Many studies[14–16] threshold the energy envelope to extract syllables from bird song recordings. Spectral entropy [17] and KL-divergence [18] have also been used to extract bird vocalizations from audio recordings. The regions containing bird vocalizations tend to exhibit lower entropy in comparison to the background regions. Also, KL-divergence between normalized power spectral density (PSD) of an audio frame and the uniform density is an indicator of entropy; lower KL-divergence corresponds to high entropy and vice versa. Lakshminarayanan et al. [18] exploited this behavior for segmenting bird vocalizations. In one of our earlier studies, we used short term energy (STE) and inverse spectral flatness (ISF) for segmentation [19].

All these methods are unsupervised in nature which is desirable. However, these methods are not able to discriminate bird vocalizations from any other sound event. Also, the performance of energy/ISF based methods rapidly deteriorate as the SNR decreases.

Methods based on spectrogram processing: Various spectrogram processing based methods [20-23] have been proposed in the literature for bird vocalization segmentation. These methods have a common theme. First, morphological processing is applied on spectrograms to remove the low-energy background acoustic events. Then, thresholding is applied on the sum of frequency magnitudes of each frame of the morphologically processed spectrogram to segment the vocalizations. Apart from these methods, Koluguri et al. [24] proposed to use directional filtering on a spectrogram enhanced using a multiple window Savitzky-Golay filter. This enhancement procedure can enhance far-field sounds which are often ignored by many other methods. The directional filtering is applied in four different directions to capture the time-frequency modulations of bird vocalizations. Like time-domain based methods, these spectrogram based methods are also not able to distinguish bird vocalizations from other background sounds and output any high-energy acoustic event as bird vocalization.

*Other Methods:* A noise-robust template matching based method is proposed by Kaewtip et al. [25]. This method performs well in challenging conditions. However, the major disadvantage

of template matching is that we must know beforehand what vocalizations we wish to segment. Hence, this method may not scale in real-world scenarios. Neal et al. [26] proposed a noise robust bird vocalization segmentation method which uses a random forest classifier to yield a probability mask from the spectrograms. A hard threshold is applied on this probability mask to generate the binary decision i.e. bird/non-bird at pixel level of the spectrogram. The main disadvantage of this method is that it requires large amount of labeled training data which is usually scarce.

In our earlier studies [19,27], we proposed model-based unsupervised frameworks for segmentation. These frameworks bypasses the need of thresholding by using a classification model. However, like other existing methods, these frameworks are also not able to distinguish between bird and non-bird sounds. We also proposed a semi-supervised framework [28] that utilizes a singular value decomposition (SVD) based feature representation to discriminate between bird/non-bird sounds. The ideal number of singular vectors required to obtain maximum discrimination between bird and non-bird sounds depend on the species under study. This limits the effectiveness of this method in field conditions.

## 3. Feature learning: directional embedding

Feature learning, in the context of bird vocalization segmentation, deals with finding a representation, which not only highlights the bird vocalizations but also removes background sounds. In this section, we describe the proposed feature representations called directional embedding (DE) for segmenting bird vocalizations. Von Mises-Fisher mixture model (moVMF), which is based on directional statistics, is used to learn a dictionary whose atoms correspond to the dominant directions of bird vocalizations. The mean directions of mixtures of vMF are regarded as dominant directions. As discusses in Section 1, the DE representation is obtained by projecting the spectrogram of an audio recording on this dictionary. Here, first we explain the pre-processing step to convert the input audio recordings into a super-frame representation. Then, we describe moVMF to obtain the dictionary of dominant directions (the reference vocalization model). Later, the method to obtain DE from an audio recording using the trained moVMF is explained. A discussion comparing DE with features obtained from NMF is also included in this section. The overall process to obtain DE is documented in Algorithm 1.

**Algorithm 1**: Procedure to obtain the proposed DE representation

Training: Obtaining dominant directions of bird vocalizations **Input:** Training audio recordings. **Output:** A dictionary of the mean directions.

- Obtain magnitude spectrograms of all the training audio recordings.
- Convert each spectrogram into the super-frame based representation and project each super-frame on the unit hypersphere as explained in Section 3.
- Use the available ground truth of the training audio recordings to extract and pool the unit norm super-frames corresponding to the bird vocalization regions. The background super-frames are simply discarded.
- Learn a moVMF having *Z* mixtures on the selected unit norm super-frames.
- •Form a dictionary, **M** having *Z* atoms, whose atoms are the mean directions of the mixtures of VMF (moVMF).
- Feature learning: Obtaining DE for any audio recording

**Input:** An audio recording and the dictionary of the mean directions.

**Output:** DE features.

- Calculate the spectrogram from any input audio recording.
- Process this spectrogram to obtain super-frames.
- Project these super-frames on the dictionary **M** (learned during training) to obtain the proposed DE representation.

3.1. Pre-processing audio recordings

Most bird vocalizations exhibit temporal and frequency (T-F) modulations [29]. These modulations give specific character to a bird vocalization. However, when the short-term Fourier transform is applied to obtain the spectrogram of an audio recording, the information about T-F modulations is smeared out as one frame is not capable of capturing these modulations effectively. These T-F modulations are not only helpful in classifying bird sounds but also provides important cues to discriminate them from other background sounds. Hence, we embed the temporal context around each frame to capture these modulations. A moving window of *w* frames with a stride of a single frame is used to embed this temporal context. The *w* frames within the moving window are concatenated, one below the other, to form a super-frame. If the T-F representation has dimensions of  $d \times n$  (d frequency bins and n frames), then after applying concatenation process, a  $wd \times n$  dimensional super-frame based representation is obtained. This representation is used for both learning the dictionary and for segmenting the input recordings.

For learning the dictionary, we are interested in clustering the unit norm representations of super-frames on the unit hypersphere. Therefore, we normalize each training super-frame to get a unit norm vector that lies on the surface of the unit hypersphere. The magnitude information is lost due to this transformation. However, the desired information about dominant frequencies is still preserved. Fig. 2 illustrates the process of converting any input spectrogram into a unit norm based representation.

#### 3.2. Mixtures of von Mises-Fisher distribution (moVMF)

The super-frames obtained by concatenation of *w d*-dimensional frames lies in the Euclidean space,  $\mathbb{R}^{wd}$  and the unit norm variants of these super-frames lie on surface of the unit hypersphere,  $\mathbb{S}^{wd-1}$ . Mixtures of von Mises-Fisher distribution [9] can be used as clusters of these unit norm vectors. moVMF has found its application across many fields in various applications such as speaker clustering [30], medical imaging [31], text mining [32] and estimation of shadow and illumination in computer vision [33]. Taking a cue from these studies, we propose to use vMF mixture modeling to cluster unit norm representations of the superframes. Originally two variants of moVMF were proposed: soft assignments and hard assignments [9]. In this work, we utilize moVMF with soft assignments only.

To learn moVMF as the bird vocalization model, the training audio recordings are manually segmented, using the available ground truth, to extract the bird vocalization regions. The superframes corresponding to these bird vocalization regions are pooled together. Suppose  $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$  is a set of unit norm representations of these pooled super-frames where  $\mathbf{x}_i \in \mathbb{R}^{wd}, \|\mathbf{x}_i\| = 1$  and thus  $\mathbf{x}_i \in \mathbb{S}^{wd-1}$ . The  $\mathbf{x}_i$  has *wd*-variate vMF distribution if its probability density is [9,34]:

$$\rho(\mathbf{x}_i; \boldsymbol{\mu}, \kappa) = C_{wd}(\kappa) \exp^{\kappa \boldsymbol{\mu}^* \mathbf{x}_i},\tag{1}$$



Fig. 2. Schematic representation of the process to convert a spectrogram into the unit norm based representation.

where  $\|\boldsymbol{\mu}\| = 1, \kappa > 0$  and  $wd \ge 2$ .  $C_{wd}(\kappa)$  is a normalizing constant defined as:

$$C_{wd}(\kappa) = \frac{\kappa^{(wd/2)-1}}{(2\pi)^{wd/2} I_{wd/2-1}(\kappa)}.$$
(2)

Here  $I_{wd/2-1}(.)$  is a modified Bessel's function of first kind and of order wd/2 - 1. The density function,  $\rho(\mathbf{x}_i; \boldsymbol{\mu}, \kappa)$ , is parameterized by mean direction,  $\boldsymbol{\mu}$ , and concentration parameter,  $\kappa$  which characterizes the concentration of unit norm vectors around the mean direction. As  $\kappa \to 0$ ,  $\rho(\mathbf{x}_i; \boldsymbol{\mu}, \kappa)$  becomes the uniform density on the surface of the unit hypersphere. On the other hand, when  $\kappa \to \infty$ , it reduces to a point density.

Building on this, the density of any unit norm vector,  $\mathbf{x}_i$  drawn from a mixture of *Z* different vMF distributions is defined as:

$$\rho\left(\mathbf{x}_{i}; \left\{\boldsymbol{\mu}_{z}\right\}_{z=1}^{Z}, \left\{\kappa_{z}\right\}_{z=1}^{Z}\right) = \sum_{z=1}^{K} \pi_{z} \rho\left(\mathbf{x}_{i}; \boldsymbol{\mu}_{z}, \kappa_{z}\right)$$
(3)

where  $\pi_z \ge 0$  and  $\sum_{z=1}^{Z} \pi_z = 1$ . An Expectation-Maximization (EM) algorithm is used for estimating the mixture parameters. The mixture log-likelihood is defined as:

$$l\left(\mathcal{X}; \left\{\pi_{z}, \boldsymbol{\mu}_{z}, \kappa_{z}\right\}_{z=1}^{Z}\right) = \sum_{i=1}^{N} \left(\log \sum_{z=1}^{Z} \rho\left(\mathbf{x}_{i}; \boldsymbol{\mu}_{z}, \kappa_{z}\right)\right)$$
  
subject to  $||\boldsymbol{\mu}_{z}|| = 1$  and  $\kappa_{z} \ge 0.$  (4)

In the E-step, the probability of  $\mathbf{x}_i$  being sampled from *z*th mixture ( $\gamma_{iz}$ ) is calculated as:

$$\gamma_{iz} = \frac{\pi_z \rho(\mathbf{x}_i; \boldsymbol{\mu}_z, \kappa_z)}{\sum_{l=1}^{Z} \pi_l \rho(\mathbf{x}_i; \boldsymbol{\mu}_l, \kappa_l)}.$$
(5)

During the M-step of EM, this  $\gamma_{iz}$  is used to calculate parameter updates which maximizes the log-likelihood defined in Eq. (4). These parameter updates are defined below:

$$\pi_{z} = \frac{1}{N} \sum_{i=1}^{N} \gamma_{iz}, \quad \boldsymbol{\mu}_{z} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{x}_{i} \gamma_{iz}, \quad \overline{r} = \frac{\|\boldsymbol{\mu}_{z}\|}{N\pi_{z}},$$

$$\boldsymbol{\mu}_{z} = \boldsymbol{\mu}_{z} / \|\boldsymbol{\mu}_{z}\|, \quad \kappa_{z} = \overline{r} (wd - \overline{r}^{2}) / (1 - \overline{r}^{2}).$$
(6)

The complete EM algorithm and detailed derivation of EM updates is documented in Banerjee et al. [9].

# 3.3. Obtaining directional embeddings (DE) using moVMF

The trained moVMF model is used to obtain directional embeddings (DE) of an input test audio recording. An under-complete dictionary,  $\mathbf{M} \in \mathbb{R}^{wd \times Z}$  (*wd* is dimensions of a super-frame and *Z* is the number of unit mean vectors), is formed using the mean directions of the mixtures ( $\boldsymbol{\mu}_z$ ) as its atoms. This dictionary forms the reference vocalization model. The mean directions of the mixtures having very low concentration are not included in the dictionary. To obtain the DE for an audio recording, its super-frame representation,  $\mathbf{P} \in \mathbb{R}^{wd \times K}$ , (*wd* is dimensions of a super-frame and *K* is the number of super-frames in the test audio recording) is projected on  $\mathbf{M}$  as:  $\mathbf{F}_{proj} = \mathbf{M}^T \times \mathbf{P}, \mathbf{F}_{proj} \in \mathbb{R}^{Z \times K}$  whose columns contains the *Z*-dimensional DE representation for the input super-frames.

*Nature of DE coefficients:* The magnitude of DE coefficients (**F**<sub>proj</sub>) obtained from a super-frame corresponding to the bird vocalizations is significantly higher than the magnitude of DE coefficients calculated from the background regions. This can be attributed to the differences in alignments of bird and background superframes with the unit mean directions (which form the atoms) accumulated in the dictionary. A bird vocalization super-frame is aligned in the directions of unit mean vectors, which leads to the high magnitude of DE coefficients. On the other hand, the background super-frames show no correlations with the dictionary atoms. Thus, the DE coefficients of background regions are characterized by their relatively low magnitude. This difference in magnitude of the DE coefficients leads to the distinction between bird vocalizations and background. This behavior is illustrated in Fig. 3. A dictionary (containing 10 mean directions) learned from the vocalizations of Cassin's vireo, a North American songbird, is used to obtain DE (shown in Fig. 3(b)) from the spectrogram shown in Fig. 3(a). A context window of 5 i.e. w = 5 is used for obtaining the super-frames. By analyzing Fig. 3(b), it is clear that only bird vocalization information is reflected in DE space and most of the other background disturbances are implicitly removed. It must be noted that no explicit filtering process is applied in DE space to remove the background sounds. This background filtering is a



Fig. 3. (a) Spectrogram containing two vocalizations of Cassin's Vireo (b) DE obtained using the proposed approach from (a).

result of non-alignment of background super-frames with the dictionary atoms.

Directional embeddings can also distinguish bird vocalizations from various types of biotic and abiotic background sounds such as the sound of rain, waterfall, human speech and insects. This behavior is highlighted in Figs. 4-6. Fig. 4(a) shows the spectrogram of two Cassin's vireo vocalizations with overwhelming waterfall sound present in the background (SNR 0 dB). The DE obtained for this spectrogram, depicted in Fig. 4(b), indicates that the proposed representation is able to filter-out most of the background. Fig. 4(c) and 4(d) depict the spectrogram of an audio recording containing human speech with a Cassin's vireo vocalization and its respective DE. It is clear from Fig. 4(d) that, as desired, the speech information is not reflected in the DE space. Also, Fig. 5 depicts the robustness of DE against a commonly occurring biotic background disturbance in the field conditions, namely cicadas. These insects are highly vocal and often overwhelm the bird vocalizations, making it difficult to correctly segment them. Fig. 5(a)shows the spectrogram of an audio containing Cassin's vireo song phrases with cicadas constantly vocalizing in the background. DE representation obtained for this spectrogram is shown in Fig. 5 (b). The analysis of this figure highlights that DE is correctly able to identify the bird vocalizations even in the presence of overwhelming noise produced by cicadas. The same dictionary as described earlier is also used here.

To further analyze this behavior, we created an artificial audio recording containing a Cassin's vireo vocalization, an emerald dove vocalization (which has relatively low frequency vocalizations) and the sound of a siren. The spectrogram of this audio recording is shown in Fig. 6(a). Again, the aforementioned dictionary learned from Cassin's vireo is used. The analysis of Fig. 6(b) shows that the DE space has most of the information about the Cassin's vireo vocalization. However, no dominant information about emerald dove vocalization or the siren is present. This is due to the fact that mean directions are learned from the vocalizations of Cassin's vireo, whose frequency profile is significantly different from the frequency profiles of dove vocalizations and siren sounds. In many applications, the siren is undesired (which is achieved in the DE space). Other applications might require that both the dove and Cassin's vireo vocalizations are captured properly (which is not achieved in this illustration.) This exhibits the trade-off between the generalization ability and background removal ability of DE. The DE can be generic as long as the frequency profile of target vocalizations is similar to the ones used for learning the dictionary. This behavior may not an issue in many field conditions, as a dictionary learned from one species can be used to segment vocalizations of many target species, which exhibit similar frequency profiles. For example, a Cassin's vireo dictionary can be used to segment vocalizations of many song bird species. Fig. 7 depicts the DE obtained for eight different species using Cassin's vireo dictionary. The analysis of this figure highlights the generalization claim (under restriction of the similarity between the reference and testing vocalizations). This claim is further corroborated in the experimentation (see Sections 5 and 6).

# 3.4. Comparison with Non-negative matrix factorization

Non-negative matrix factorization (NMF) has been successfully used in various tasks such as acoustic scene classification [35], speech enhancement [36] and voice activity detection [37]. For the task of bird vocalization segmentation, a dictionary or reference model of vocalizations can be learned using supervised or semi-supervised NMF [7]. Since the proposed framework does not use the background modeling, we are only considering the semi-supervised NMF where the vocalization dictionary is kept constant and the background atoms are updated in each iteration. The feature representation obtained using NMF can distinguish between bird sounds and silence regions effectively. However, this representation also reflects the presence of most of the background sounds. Both of these behaviors are evident in Fig. 8. The semisupervised NMF features (see Fig. 8) obtained for spectrograms shown in Fig. 4(c) and 6(a) show the significant presence of background acoustic events in the feature domain. This can be attributed to the fact that NMF models the source spectra by a convex cone which often contains the regions where source spectra does not exist [8]. Moreover, these regions can contain other undesirable background acoustic events. This may hinder the ability of NMF to distinguish bird vocalizations from other background sounds in a semi-supervised setup.

In addition to this, obtaining NMF representation is computationally expensive (iterative and requires multiple matrix updates) in comparison to the calculation of DE from a pre-trained dictionary.

# 4. Method

In this section, first, we formulate the problem statement and then, we describe the proposed framework for bird vocalization segmentation.

#### 4.1. Problem formulation

The objective of the bird vocalization segmentation is to temporally locate any bird vocalization irrespective of its nature (call or song) and the species producing the vocalization. There are two steps in formulating this problem: feature representation and classification. First, a short-term feature representation is obtained from an input recording, where a feature vector  $\mathbf{x}_i \in \mathbb{R}^y$  (*y* is the number of features) is obtained for the *i*th frame. In classification stage,  $\mathbf{x}_i$  is classified as the bird vocalization or the background.

#### 4.2. Proposed framework

The proposed two-pass framework (Fig. 1) for bird vocalization segmentation is described in this sub-section. Here, first we describe pass 1 of the framework: the process to generate automatic training labels from an input test recording using DE and mutual information. Then, we discuss pass 2 of the framework, where labels and the DE features, generated during pass 1 from the input recording itself, are used for training a classification model which in turn is used for making bird/background decisions.



Fig. 4. (a) and (c) Depicts the spectrograms of recordings having vocalizations of Cassin's Vireo with overwhelming waterfall noise and human speech respectively. (b) and (d) Illustrates their respective DE representation.



**Fig. 5.** (a) The spectrogram of an audio recording having Cassin's Vireo vocalizations with significant background noise produced by cicadas (b) DE representation obtained from the spectrogram shown in (a).



**Fig. 6.** (a) Spectrogram containing a Cassin's vireo vocalization, a low frequency emerald dove vocalization and an air horn (b) DE obtained from the spectrogram shown above in (a).

# 4.3. Pass 1: Generating automatic training labels

Obtaining DE from input audio recording: A test audio recording and a dictionary containing unit mean directions of a trained moVMF are given as input to pass 1 of the framework. The input audio recording is converted into the super-frame based representation as described in Section 3.1. This super-frame based representation is projected on the dictionary to obtain the DE representation. This process is discussed in detail in Section 3.3.

Using mutual information to automatically generate labels: As discussed in Section 3.3, the magnitude of DE coefficients obtained from non-bird super-frames is significantly less than DE coefficients of bird vocalization super-frames. Also, there is a large variation in magnitudes of DE coefficients of the bird super-frames. This is because a bird super-frame exhibits varying degree of alignments with different unit mean directions of the dictionary. However, there is very little variation in magnitude of DE coefficients of a non-bird super-frame. This can be attributed to the fact that a non-bird super-frame has no or very little similarity with any of the dictionary atoms. This dissimilarity leads to low magnitudes (approaching zero) of DE coefficients. Fig. 9 highlights this behavior.

The difference in magnitudes of DE coefficients between bird and non-bird regions can be captured using mutual information (MI) between DE features of two consecutive super-frames. MI of a vector with itself is highest, therefore MI between two almost similar DE feature vectors will be higher than the two which are different. Since the DE feature vectors for any background sound are almost similar (all coefficients approach zero), MI calculated for background regions is almost constant and is higher than the MI calculated for the bird vocalization regions. Fig. 10 illustrates this behavior.

To calculate MI, DE vectors  $(\mathbf{F}_{proj})$  obtained from an input recording are normalized using the softmax function [28] as:

$$\hat{\mathbf{f}}_{n_j} = \frac{e_{\mathbf{n_j}}^{\mathbf{f}}}{\sum_{z=1}^{Z} e_{\mathbf{n_z}}^{\mathbf{f}}}, \quad \text{for } j = 1, 2, \dots, Z.$$

$$(7)$$

Here  $\mathbf{f}_n$  is the *n*th column of  $\mathbf{F}_{\text{proj}}$ .  $\mathbf{f}_{n_j}$  represents the *j*th element of  $\mathbf{f}_n$  and *Z* represents the dimensionality of  $\mathbf{f}_n$ . Since softmax function calculates the categorical distribution of an event over *Z* different events,  $\hat{\mathbf{f}}_{n_j}$  represents the probability estimate of the degree of the alignment of  $\mathbf{f}_n$  with *j*th mean direction of the learned dictionary **M**. As discussed earlier, the DE vectors of non-bird regions exhibit no or very low alignment with any of the dictionary atoms. As a result, these non-bird DE vectors (softmax normalized) become almost equiprobable vectors i.e. each element of the normalized DE vector exhibits almost same value. Similarly, elements of a nor-



**Fig. 7.** (a), (c), (e), (j), (k), (m) and (o) depict the spectrograms containing vocalizations of Cassin's vireo, Hutton vireo, Black-headed grosbeak, Greater sooty owl, California thrasher, Western tanager, Black phoebe and Indian peafowl respectively. Their respective feature representations/DE are shown in (b), (d), (f), (h), (j), (l), (n) and (p).





malized DE vector, corresponding to a bird super-frame, exhibits different probabilities of the degree of alignment of super-frame with different mean directions.

MI between the normalized DE vectors of each pair of the consecutive super-frames (i.e. between  $n_{th}$  and  $(n-1)_{th}$  columns of normalized **F**<sub>proj</sub>) is calculated as:



**Fig. 9.** Box plots of the DE coefficients calculated for (a) 100 background frames (b) 100 speech frames and (c) 100 Cassin's vireo (CV) vocalization frames (d) 100 air-horn frames and (e) 100 Califorina Thrasher (CT), a bird species frames. In all cases, the dictionary/reference model learned from Cassin's vireo vocalizations is used.

$$MI(\hat{\mathbf{f}}_{n},\hat{\mathbf{f}}_{n-1}) = H(\hat{\mathbf{f}}_{n}) + H(\hat{\mathbf{f}}_{n-1}) - H(\hat{\mathbf{f}}_{n},\hat{\mathbf{f}}_{n-1}).$$
(8)

Here H() represents the entropy and can be calculated as:

$$H(\hat{\mathbf{f}}_n) = -\sum_{j=1}^{Z} \hat{\mathbf{f}}_{n_j} \log_2 \hat{\mathbf{f}}_{n_j}, \text{ where Z is the dimensions of DE vector.}$$

(9)



**Fig. 10.** (a) DE representation obtained for spectrogram shown in Fig. 4(a). (b) Normalized mutual information calculated for the DE representation shown in (a).

Also,  $\hat{\mathbf{f}}_{n_j}$  is a multinoulli probability estimate (as discussed earlier) and Z is the dimensions of  $\hat{\mathbf{f}}_n$ .

The joint entropy  $H(\hat{\mathbf{f}}_n, \hat{\mathbf{f}}_{n-1})$  can be calculated as:

$$H\left(\hat{\mathbf{f}}_{n-1}, \hat{\mathbf{f}}_{n}\right) = -\sum_{j=1}^{Z} P\left(\hat{\mathbf{f}}_{n-1_{j}}, \hat{\mathbf{f}}_{n_{j}}\right) \log_{2} P\left(\hat{\mathbf{f}}_{n-1_{j}}, \hat{\mathbf{f}}_{n_{j}}\right), \tag{10}$$

where the joint probability estimate  $P(\hat{\mathbf{f}}_{n-1_j}, \hat{\mathbf{f}}_{n_j})$  is calculated as:

$$P(\hat{\mathbf{f}}_{n-1_{j}}, \hat{\mathbf{f}}_{n_{j}}) = \frac{\text{number of times } (\hat{\mathbf{f}}_{n-1_{j}}, \hat{\mathbf{f}}_{n_{j}}) \text{ jointly occurs in } \hat{\mathbf{f}}_{n-1} \text{ and } \hat{\mathbf{f}}_{n}}{Z}$$
(11)

MI, calculated using Eq. (7), is used to generate automatic training labels in the proposed framework. Q super-frames having lowest and Q super-frames having highest MI are labeled as bird vocalizations and background respectively. The rest of the superframes remain unused. These generated labels and DE feature vectors are given as input to the second pass.

#### 4.4. Pass 2: Training classification model and extracting segments

During the second pass, a classification model is trained using the input labels (generated during pass 1) and DE feature vectors. Once the model is trained, each DE vector of the input recording is classified as either the background or bird vocalization. In this work, we have used Support vector machines (SVM) for classification. Since the features generated from the input recording is itself used for training the classification model, the proposed framework does not suffer from the classical problem of mismatch in trainingtesting conditions. In bioacoustics, this mismatch is very common due to the volatile/unpredictable nature of the field conditions and differences in the audio recording devices.

Figs. 11(c) and 11(d) show the output of the proposed framework for the input test audio recordings whose spectrograms are shown in Figs. 11(a) and 11(b) respectively. The '0' represents the background while the '1' represents the bird vocalization. The analysis of this figure makes it clear that the proposed framework can identify the bird vocalizations in an audio recording successfully.

# 5. Experiments

This section describes the datasets and experimental setup used for the performance evaluation of the proposed segmentation framework.

#### 5.1. Datasets used

Experimental validation of the proposed framework is performed on the audio recordings of seven different species. These species include Cassin's Vireo (CV1), California Thrasher (CT), Hutton Vireo (HV), Black-headed Grosbeak (BG), Grey Shrike-Thrush (GS), Redthroat (RT) and Western Tanager (WT). The recordings of these 7 bird species were obtained from the bird audio database maintained by the Art & Science center, UCLA [38]. Apart from this, a different set of Cassin's Vireo recordings (CV2) [39] is also used for the performance evaluation. All the audio recordings available at both these sources are 16-bit mono WAV files having a sampling rate of 44.1 kHz. The details about the number of recordings and the number of vocalizations available for each species are tabulated in Table 1. The choice of species used in this study is restricted by the public availability of the strongly labeled datasets (having segment level time stamps). The datasets used in this study have both labeled and unlabeled audio recordings. However, only the labeled audio recordings from all datasets are used here. A list of audio file used in this study and the process to retrieve these files is available athttps://bit.ly/20th03R. The ground truth provided with datasets is in terms of time-stamps (onset/offset times) of the bird vocalizations. These time-stamps are processed to obtain frame-level ground truths, which are used for evaluating the segmentation performance.

To test the proposed framework in noisy conditions, background noise are artificially added to the recordings of the CV2 dataset. Four different types of background sounds i.e. rain, waterfall, river and cicadas at 0 dB, 5 dB, 10 dB, 15 dB and 20 dB SNR are added using Filtering and Noise Adding Tool (FaNt) [40]. These background sound files are downloaded from FreeSound [41]. The bird vocalization recordings are most likely to be affected by similar sounds in the field conditions.

# 5.2. Performance evaluation setup

Three experiments are designed to evaluate the segmentation performance of the proposed framework. In the first experiment, the segmentation performance of the proposed framework is compared with various existing algorithms over the datasets mentioned earlier. Also, using the noisy versions of CV2, the robustness of the proposed framework and other comparative algorithms is evaluated in different SNR conditions. The second experiment is conducted to highlight the generic nature of the proposed framework. In this experiment, the dictionary learned from one species is used to segment the vocalizations of another species. The performance of this setup is compared with the segmentation performance obtained when same species is involved in both training and testing. For performance comparison, frame-level F1-score is used as a performance metric.

For comparison, six different existing methods (discussed in Section 2) are used as baselines in this study. Apart from these algorithms, a variant of the proposed framework is also used for comparative studies. This variant targets the feature representation module of the proposed framework by replacing the DE with a semi-supervised non-negative matrix factorization (NMF) (discussed in Section 3). Table 2 summarizes all the algorithms used for comparative study.

Third experiment is conducted to compare the performances of the weakly supervised neural networks with the proposed framework. In Section 1, the authors explained that the problem of the scarcity of strong training labels for the task in hand can also be addressed by the data-intensive, weakly supervised neural networks that are trained for audio tagging tasks. These networks only require the recording level labels during training and bypass the need of strong or frame-level labels. The frame-level predictions



**Fig. 11.** (a) and (b) depict the spectrograms of two input audio recordings containing the vocalizations of Black-headed grosbeak and California thrasher respectively. (c) and (d) show the decisions outputted by the proposed framework on these audio recordings.

#### Table 1

Total audio recordings and number of vocalizations present in each dataset considered for the performance evaluation.

Dataset	CV1	CV2	CT	HV	BG	GS	RT	WT	Total	
No. of Recordings	497	13	97	5	108	12	7	3	744	
Vocalizations (Approx.)	40 K	900	20 K	278	16 K	120	1.2 K	500	78998	

#### Table 2

Methods used for comparative studies.

Algorithm	Classification Method	Nature
Two-pass segmentation (TP) [19]	Gaussian Mixture Models (GMM)	Unsupervised
SVD + MI (SVD) [28]	Thresholding	Semi-supervised
Morphological noise removal on spectrograms + Energy (FOD) [21]	Thresholding	Unsupervised
Morphological noise removal and finding connected segments + Energy (	(LAS) [22] Thresholding	Unsupervised
Morphological Opened Spectrogram + Energy (MOS) [20]	Thresholding	Unsupervised
Spectrogram Enhancement and Directional Filtering + Energy (DF) [24]	Thresholding	Unsupervised
DE + SVM Based Framework (Proposed)	SVM	Semi-supervised
NMF + SVM (NMF), variant of the proposed framework	SVM	Semi-supervised

generated by these networks are the by-product of audio-tagging tasks. Although, to the best of our knowledge, no existing study has used weakly supervised neural networks for bird vocalization segmentation. However, for the proper performance evaluation of the proposed framework, it is essential to incorporate such weakly supervised neural networks in the performance comparison. Thus, we utilize off-the-shelf weakly supervised convolutional neural networks, proposed in [6,42] for the performance comparison.<sup>1</sup> In [42], the authors used attention and localization mechanism in a CNN-RNN (convolutional-recurrent neural network) architecture (*Att\_loc\_CNN*) to detect the audio events. In [6], a gated convolution neural network (GCNN) is proposed for detecting rare sound events. The gated convolutional layers control the flow of information from layer to other and act as attention mechanism. Both these networks can be easily extended for bird activity detection (audio tagging task) and hence, bird vocalization segmentation. Since the nature of methods listed in Table 2 is completely different from the weakly supervised neural networks, it is not possible to use the same type and amount of data for training. Hence, we train these networks on audio recordings obtained from *Wablr* and *Freefield* datasets. These datasets are provided for training as a part of *Bird Activity Detection Challenge*, 2017 [5]. A total 16000 recordings (8000 bird and 8000 non-bird approx.) are used for training both the networks. Each recording is 10 s long and is sampled at 44.1 kHz. The testing is done on all the datasets listed in Table 1.

#### 5.3. Parameter settings used

#### 5.3.1. Parameters used in the proposed framework

For obtaining the spectrograms from audio recordings, STFT is performed on 20 ms frames with an overlap of 50%, using 1024 FFT points. The super-frames are obtained by concatenating 5 frames i.e. w = 5. The dictionary, **M**, with 10 atoms is learned from the training vocalizations. moVMF with 15 mixtures is applied to obtain the dominant directions of bird vocalizations. Out of these 15 mixtures, five having lowest concentration were discarded (Z = 10). While auto-labeling super-frames in pass 1, 2000 superframes having lowest MI are labeled as bird vocalizations while 2000 super-frames having highest MI are labeled as background

<sup>&</sup>lt;sup>1</sup> Python codes available at: https://github.com/yongxuUSTC/dcase2017\_task4\_ cvssp https://github.com/yongxuUSTC/att\_loc\_cgrnn.

(i.e Q = 2000.) All these values are experimentally determined on a validation dataset. Support vector machines with cubic polynomial kernel is used for classification in pass 2.

#### 5.3.2. Parameters used in the other comparative algorithms

For SVD, the number of dictionary atoms used in the respective study [28] i.e. 5, is also used here. For implementing directional filtering (*DF*), an order of 3 and a matrix length of 21 is used to apply Savitzky-Golay filtering. For *NMF*, 20 dictionary atoms are used. A classical NMF algorithm [43] with Itakura-Saito distance is used in this work. All these parameters including thresholds are tuned empirically on a validation dataset to provide the optimal segmentation performance. For weakly supervised neural networks, only a small change is done in architectures to incorporate the required inputs. Mel-spectrograms are given as input to these networks. The input size is changed to  $40 \times 1000$  to incorporate 10 s of audio recordings (1000 20 ms frames with 50% overlap and 40 is the number of Mel bands). The rest of parameters and architecture details proposed in the corresponding studies [6,42] are also used here.

#### 5.3.3. Data distribution for training, testing and validation

For first experiment: From each dataset, two recordings are used as validation set for tuning the threshold and parameters. These recordings are not used for training or testing purposes. The remaining audio recordings are used for the evaluation in a five fold cross-validation manner and the results presented here are averaged across these five folds. For each fold, 20% of the recordings are chosen from the dataset for learning the dictionary and rest of the recordings are used for testing. The same test recordings are given to the unsupervised methods for a fair performance comparison. It must be noted that the test recordings contain many vocalization types, which are not present in the training recording used for learning the dictionary.

For second experiment (cross-species segmentation): In this experiment, one dataset (corresponding to a particular species) is used for training i.e. for learning the dictionary while all the other datasets are used for testing. From the training dataset, two recordings are randomly chosen as the validation set and out of the remaining recordings, 20% of recordings are randomly chosen for learning the dictionary. All the recordings present in the test datasets are used for evaluating the segmentation performance.

For third experiment (comparison against weakly supervised neural networks): For weakly supervised neural networks, 80% of the recordings in *Freefield* and *Warblr* are used for training while remaining recordings are used for validation. The testing is done on all the datasets listed in Table 1. For a fair comparison, only the files used for testing the proposed framework in experiment 1 are used i.e. audio files used for validation are not used here. Also, the segmentation performances of these networks are compared against the average F1-score achieved across five-folds by the proposed framework.

#### 6. Results and discussion

In this section, we describe the results obtained for three experiments mentioned in the previous section. In addition to that, we also discuss the various aspects of the proposed framework such as the effect of increasing the size of context window, increasing the number of mixtures in moVMF and using MI over energy/entropy.

#### 6.1. First experiment: comparison with existing methods

The segmentation performance of the comparative algorithms including the proposed framework is depicted in Table 3. From the analysis of Table 3, the following can be inferred:

- The proposed framework outperforms most of the existing algorithms across all the datasets considered.
- SVD and the proposed framework significantly outperform the unsupervised methods. This is due to their ability to discriminate bird vocalizations from most of the background disturbances.
- *DF* outperforms all the other spectrogram-processing based methods. This can be attributed to the spectrogram enhancement and directional filtering steps before applying energy-based thresholding. This processing on spectrogram significantly increases the energy contrast between vocalizations and background leading to a better segmentation. However, like other unsupervised methods, *DF* is also not able to discriminate between bird and non-bird sounds.
- NMF variant of the proposed framework also outperforms most of the existing algorithms and its segmentation performance is second only to the proposed framework and *SVD*. Better segmentation performance can be attributed to the reference modeling of the bird vocalizations using NMF. However, as discussed earlier, NMF may not be able to discriminate between birds and other acoustic events effectively in a semi-supervised setup, which leads to a poorer segmentation performance than the proposed framework.

The segmentation performance of these methods is also evaluated on four noisy versions of CV2, each having SNR of 0 dB, 5 dB, 10 dB, 15 dB and 20 dB. The results are shown in Fig. 12. These results highlight the noise robustness of the proposed framework, which shows no significant drop in segmentation performance as the SNR deteriorates (20 dB–0 dB). Apart from this, all the semisupervised methods perform well at low SNR conditions. Again, the feature extraction step in all these methods mitigate the effect of the background noise to a large extent which results in better segmentation performance. On the other hand, all the unsupervised methods show a significant drop in performance on moving from high to low SNR conditions. An average relative drop of 12%, 4.5%, 16%, 16.3%, 15.5%, 17%, 3.6% and 6.4% is observed in performance of *TP*, *SVD*, *FOD*, *LAS*, *MOS*, *DF*, *Proposed* and *NMF* respectively on moving from 20 dB SNR to 0 dB SNR across all noise types.

Table 3

Segmentation performance of different algorithms over different datasets (*Results of first experiment*). F-Score is used as the metric for comparison. The dictionary is learned from the vocalizations of the species, which are targeted for segmentation. Entries in bold represent the best F-score obtained for the corresponding dataset.

Method	Datasets								
	CV1	CV2	CT	HV	BG	GS	RT	WT	
TP	0.69	0.72	0.68	0.7	0.64	0.64	0.66	0.69	
SVD	0.79	0.82	0.77	0.8	0.76	0.78	0.79	0.8	
FOD	0.61	0.62	0.58	0.67	0.59	0.6	0.63	0.65	
LAS	0.6	0.6	0.57	0.68	0.57	0.57	0.61	0.64	
MOS	0.52	0.57	0.53	0.65	0.57	0.58	0.6	0.62	
DF	0.66	0.67	0.66	0.7	0.64	0.65	0.7	0.71	
Proposed	0.81	0.83	0.8	0.8	0.77	0.79	0.81	0.81	
NMF	0.76	0.78	0.76	0.76	0.74	0.72	0.72	0.76	



Fig. 12. Segmentation performance of different algorithms on CV2 corrupted by (a) Rain, (b) River, (c) Waterfall and (d) Cicadas noise at different SNR. Here 'Prop.' represents the proposed framework.

#### Table 4

Segmentation performance for cross-species train-test setup (*Results of the second experiment*). F-Score is used as the performance metric. The dictionary learned from one species is used for segmenting all the other species considered in the study. Entries in bold represent the best F-score obtained for the corresponding dataset.

Training Dataset		Testing Datasets							
	СТ	HV	BG	GS	RT	WT	CV1		
СТ	0.79	0.81	0.78	0.76	0.79	0.79	0.78		
HV	0.77	0.81	0.76	0.76	0.78	0.79	0.77		
BG	0.78	0.8	0.78	0.77	0.79	0.8	0.79		
GS	0.76	0.8	0.77	0.77	0.78	0.78	0.78		
RT	0.78	0.79	0.76	0.75	0.8	0.78	0.79		
WT	0.77	0.8	0.77	0.75	0.79	0.81	0.78		
CV1	0.78	0.8	0.77	0.76	0.8	0.8	0.8		

#### 6.2. Second experiment: generic nature of the proposed framework

In this experiment, the generic nature of the proposed framework is evaluated using different species for training (learning the dictionary) and testing. Table 4 tabulates the segmentation performance observed for this cross-species setup. The analysis of Table 4 shows that the best segmentation performance is obtained when the training and testing species are same. However, there is no dramatic drop in the performance when the dictionary is learned from the vocalization of a species that is not being targeted for the segmentation. Moreover, this drop in F1-score is minute and segmentation performance across all the cases is near optimum. This upholds the generalization claim of the proposed framework.

# 6.3. Third experiment: comparison against weakly supervised neural networks

The performances achieved by the proposed framework and the weakly supervised neural networks on different datasets are listed in Table 5. The analysis of Table 5 shows that the segmentation performance of the proposed framework is better than the weakly

supervised neural networks considered in this study. The performance of GCNN[6] and Att\_loc\_CNN[42] is not up to the expected standards as these networks were designed for rare sound event detection<sup>2</sup> and not for the bird vocalization segmentation. Moreover, the differences in training and test datasets may have also amounted to the average performance of these neural networks. It must be noted that both these networks were able to identify the presence and absence of bird vocalizations. However, the onset-offset boundaries detected by Att\_loc\_CNN are much better than boundaries detected by GCNN. This can be attributed to the fact that Att\_loc\_CNN utilizes time-distributed convolutional operations on each input frame, while GCNN utilizes two-dimensional convolutional operations at each layer. Due to these two-dimensional convolutional operations, the *field-of-view* observed at each frame in the last convolution layer spans multiple input frames [44]. This elevated span deters the ability of network to detect the correct onset-offset times. Also, the results obtained from this experiment confirm that the weak supervision provides a convenient way to tap into the massive potential of huge datasets with weak labels.

<sup>&</sup>lt;sup>2</sup> http://www.cs.tut.fi/sgn/arg/dcase2017/challenge/task-rare-sound-event-detection.

Table 5	
Segmentation performance of the proposed framework against	t weakly supervised neural networks.
Mathad	Detecete

Method	Datasets							
	CV1	CV2	СТ	HV	BG	GS	RT	WT
GCNN [6]	0.59	0.62	0.63	0.65	0.6	0.58	0.59	0.61
Att_loc_CNN [42]	0.72	0.73	0.75	0.74	0.71	0.75	0.73	0.76
Proposed	0.8	0.83	0.8	0.8	0.77	0.79	0.81	0.81

# 6.4. Effect of w and Z on segmentation performance of the proposed framework

The context window (of size w) controls the amount of spectraltemporal information to be incorporated in a super-frame. This helps in capturing frequency and temporal modulations present in most of the bird vocalizations. The smaller value of w gives rise to a super-frame representation having less context information while using a large value of w increases the dimensionality of super-frames which is often undesirable. The large values of wcan also affect the generic nature of the proposed algorithm. The mean directions learned from the super-frames having very high context information can get highly biased towards the vocalizations used for training. Therefore, an appropriate value of w must be used. As discussed previously, we have used w = 5 for all the experiments. This value of w is determined experimentally and it balances the trade-off between preserving the context information and high dimensionality of the super-frames.

The number of mixtures (*Z*) in moVMF directly controls the dimensionality of the feature representation. We intend to train moVMF with less amount of training data (usually a single audio recording). Hence, the number of mixtures should not be very large. To choose an appropriate value of *Z*, we experimented with Z = 15, 20 and 40. The behavior of DE for all these cases was same i.e. no improvement was observed by increasing *Z* from 15 to 40. This shows that the proposed DE representations can be learned effectively using few mean directions only. Considering the dimensionality of DE representation and usage of less training data, we chose to use Z = 15 for all the experiments.

#### 6.5. Using MI over energy/entropy

Most of the information in DE space is about the bird vocalizations. However, DE space can also exhibit background information to some extent (as can be seen in Fig. 4(b)). These background disturbances are slowly varying and for two consecutive superframes, DE representations are almost same. Hence, MI for background regions will still be almost constant. However, metrics like energy and entropy are known to be susceptible to these disturbances, leading to a varying energy/entropy for background regions. This may lead to the poor automatic labeling in the proposed framework. Therefore, MI is used in the proposed framework over energy and entropy.

# 7. Conclusion

In this paper, we proposed a two pass semi-supervised framework for bird vocalization segmentation. A feature representation called directional embedding, tailored for this segmentation task, has also been proposed. The proposed framework gets its semisupervised nature from the feature learning module which utilizes a small amount of training data to model the dominant directions of the vocalizations of many bird species. The utilization of this feature representation in the proposed framework helps in discriminating the bird vocalizations from other background acoustic events. The segmentation performance of the proposed framework has been evaluated on the datasets of seven different species and the experimental results establish the superiority of the proposed framework over various existing bird vocalization segmentation methods.

One possible shortcoming of the proposed framework is that when bird vocalizations used for learning the dictionary are similar to the undesired background events, the framework cannot discriminate between bird and non-bird sounds. However, it must be noted that these cases rarely occur in field conditions. Future work will deal with developing segmentation techniques which could tackle such situations. Also, the proposed DE representations will be explored in a supervised setup for species identification and acoustic event detection.

## Acknowledgements

This work is partially supported by IIT Mandi under the project IITM/SG/PR/39 and Science and Engineering Research Board, Government of India under the project SERB/F/7229/2016-2017.

#### References

- Brandes TS. Automated sound recording and analysis techniques for bird surveys and conservation. Bird Conserv Int 2008;18(S1):S163–73.
- [2] Lee C-H, Han C-C, Chuang C-C. Automatic classification of bird species from their sounds using two-dimensional cepstral coefficients. IEEE Trans Audio Speech Language Process 2008;16(8):1541–50.
- [3] Fagerlund S. Automatic recognition of bird species by their sounds (Masters' thesis). Helsinki University Of Technology; 2004.
- [4] DCASE 2017.http://www.cs.tut.fi/sgn/arg/dcase2017/index [accessed 05.04.2018].
- [5] Bird Audio Detection Challenge.http://machine-listening.eecs.qmul.ac.uk/birdaudio-detection-challenge/ [accessed 10.05.2018].
- [6] Xu Y, Kong Q, Wang W, Plumbley M. Large-scale weakly supervised audio classification using gated convolutional neural network. In: Proceedings of Int. Conf. Acoust., Speech, Signal Process..
- [7] Mohammed S, Tashev I. A statistical approach to semi-supervised speech enhancement with low-order non-negative matrix factorization. In: Proceedings of Int. Conf. Acoust., Speech, Signal Process., p. 546–50.
- [8] Kim M, Smaragdis P. Mixtures of local dictionaries for unsupervised speech enhancement. IEEE Signal Process Lett 2015;22(3):293–7.
- [9] Banerjee A, Dhillon IS, Ghosh J, Sra S. Clustering on the unit hypersphere using von mises-fisher distributions. J Mach Learn Res 2005;6:1345–82.
- [10] Stowell D, Plumbley MD. Automatic large-scale classification of bird sounds is strongly improved by unsupervised feature learning. PeerJ 2014;2:e488.
- [11] Dieleman S, Schrauwen B. Multiscale approaches to music audio feature learning. In: Proceedings of Int. Soc. Music Info. Retrieval. p. 116–21.
- [12] Zhong S, Ghosh J. A comparative study of generative models for document clustering. In: Proceedings of Workshop on clustering high dimensional data and its applications in SIAM data mining conference.
- [13] Trifa VM, Kirschel ANG, Taylor CE, Vallejo EE. Automated species recognition of antbirds in a Mexican rainforest using hidden Markov models. J Acoust Soc Am 2008;123(4):2424–31. <u>https://doi.org/10.1121/1.2839017</u>.
- [14] Harma A, Somervuo P. Classification of the harmonic structure in bird vocalization. In: Proceedings of Int. Conf. Acoust. Speech, Signal Process. p. 701-4. <u>https://doi.org/10.1109/ICASSP.2004.1327207</u>.
- [15] Somervuo P, Harma A, Fagerlund S. Parametric representations of bird sounds for automatic species recognition. IEEE Trans Audio Speech Language Process 2006;14(6):2252-63. <u>https://doi.org/10.1109/TASL.2006.872624</u>.
- [16] Fagerlund S. Bird species recognition using support vector machines. EURASIP J. Appl Signal Process 2007;2007(1). <u>https://doi.org/10.1155/2007/38637</u>. 64– 64.
- [17] Wang NC, Hudson RE, Tan LN, Taylor CE, Alwan A, Yao K. Bird phrase segmentation by entropy-driven change point detection. In: Proceedings of Int. Conf. Acoust. Speech, Signal Process. p. 773–7. <u>https://doi.org/10.1109/ ICASSP.2013.6637753</u>.

- [18] Lakshminarayanan B, Raich R, Fern X. A syllable-level probabilistic framework for bird species identification. In: Proceedings of Int. Conf. Mach. Learn. Applicat., p. 53–9. <u>https://doi.org/10.1109/ICMLA.2009.79</u>.
- [19] Thakur A, Rajan P. Model-based unsupervised segmentation of birdcalls from field recordings. In: Proceedings of Int. Conf. Signal Process. Commun. Syst..
- [20] de Oliveira AG, Ventura TM, Ganchev TD, de Figueiredo JM, Jahn O, Marques MI, Schuchmann K-L. Bird acoustic activity detection based on morphological filtering of the spectrogram. Appl Acoust 2015;98:34–42.
- [21] Fodor G. The ninth annual mlsp competition: first place. In: Proceedings of MLSP. p. 1–2.
- [22] Lasseck M. Large-scale identification of birds in audio recordings. CLEF (Working Notes) 2014:643–53.
- [23] Glotin H, LeCun Y, Artieres T, Mallat S, Tchernichovski O, Halkias X. The NIPS4B 2013 competition. In: Proceedings of neural information processing scaled for bioacoustics: from neurons to big data.
- [24] Koluguri NR, Meenakshi GN, Ghosh PK. Spectrogram enhancement using multiple window savitzky-golay (mwsg) filter for robust bird sound detection. IEEE/ACM IEEE Trans Audio Speech Language Process 2017;25(6):1183–92.
- [25] Kaewtip K, Tan LN, Taylor CE, Alwan A. Bird-phrase segmentation and verification: a noise-robust template-based approach. In: Proceedings of Int. Conf. Acoust. Speech, Signal Process. p. 758–62.
- [26] Neal L, Briggs F, Raich R, Fern XZ. Time-frequency segmentation of bird song in noisy acoustic environments. In: Proceedings of Int. Conf. Acoust. Speech, Signal Process. p. 2012–5. <u>https://doi.org/10.1109/ICASSP.2011.5946906</u>.
- [27] Thakur A, Rajan P. Unsupervised birdcall activity detection using source and system features. In: Proceedings of NCC. IEEE; 2017.
- [28] Thakur A, Rajan P. Rényi entropy based mutual information for semisupervised bird vocalization segmentation. Proc MLSP 2017.
- [29] Marler P. Bird calls: their potential for behavioral neurobiology. Ann N Y Acad Sci 2004;1016(1):31–44.
- [30] Tang H, Chu SM, Huang TS. Generative model-based speaker clustering via mixture of von mises-fisher distributions. In: Proceedings of Int. Conf. Acoust. Speech, Signal Process. p. 4101–4.

- [31] Bangert M, Hennig P, Oelfke U. Using an infinite von mises-fisher mixture model to cluster treatment beam directions in external radiation therapy. In: Proceedings of Int. Conf. Mach. Learn. App., p. 746–51.
- [32] Srivastava AN, Sahami M. Text mining: classification, clustering, and applications. CRC Press; 2009.
- [33] Panagopoulos A, Samaras D, Paragios N. Robust shadow and illumination estimation using a mixture model. In: Proceedings of Comp. Vision and Pattern Recog., IEEE; 2009. p. 651–8.
- [34] Gopal S, Yang Y. Von mises-fisher clustering models. In: Proceedings of Int. Conf. on Machine Learn., p. 154–62.
- [35] Bisot V, Serizel R, Essid S, Richard G. Acoustic scene classification with matrix factorization for unsupervised feature learning. In: Proceedings of Int. Conf. Acoust., Speech, Signal Process., p. 6445–9.
- [36] Mohammadiha N, Smaragdis P, Leijon A. Supervised and unsupervised speech enhancement using nonnegative matrix factorization. IEEE Trans Audio Speech Language Process 2013;21(10):2140–51.
- [37] Germain FG, Sun DL, Mysore GJ. Speaker and noise independent voice activity detection. Proc Interspeech 2013:732–6.
- [38] Art-sci center, University of California.http://artsci.ucla.edu/birds/database. html/ [accessed 10.07.2018].
- [39] Cassin's vireo recordings.http://taylor0.biology.ucla.edu/al/bioacoustics/ [accessed 20.03.2016].
- [40] Filtering and noise adding tool.http://dnt.kr.hs-niederrhein.de/ [accessed 14.11.2016].
- [41] Freesound.http://freesound.org/ [accessed 13.03.2017].
- [42] Xu Y, Kong Q, Huang Q, Wang W, Plumbley MD. Attention and localization based on a deep convolutional recurrent model for weakly supervised audio tagging. Interspeech; 2017. p. 1267–71.
- [43] Lee DD, Seung HS. Algorithms for non-negative matrix factorization. Adv Neural Inf Process Syst 2001:556–62.
- [44] Goodfellow I, Bengio Y, Courville A. Deep learning. MIT Press; 2016. URL: http://www.deeplearningbook.org.