Multiscale CNN based Deep Metric Learning for Bioacoustic Classification: Overcoming Training Data Scarcity Using Dynamic Triplet Loss

Anshul Thakur, Daksh Thapar, Padmanabhan Rajan, and Aditya Nigam

School of Computing and Electrical Engineering, IIT Mandi, Mandi, Himachal Pradesh-175005, India[a)]

(Dated: 28 March 2019)

This paper proposes multiscale convolutional neural network (CNN)-based deep metric learning for bioacoustic classification, under low training data conditions. The proposed CNN is characterized by the utilization of four different filter sizes at each level to analyze input feature maps. This multiscale nature helps in describing different bioacoustic events effectively: smaller filters help in learning the finer details of bioacoustic events, whereas, larger filters help in analyzing a larger context leading to global details. A dynamic triplet loss is employed in the proposed CNN architecture to learn a transformation from the input space to the embedding space, where classification is performed. The triplet loss helps in learning this transformation by analyzing three examples, referred to as triplets, at a time where intra-class distance is minimized while maximizing the inter-class separation by a dynamically increasing margin. The number of possible triplets increases cubically with the dataset size, making triplet loss more suitable than the softmax cross-entropy loss in low training data conditions. Experiments on three different publicly available datasets show that the proposed framework performs better than existing bioacoustic classification frameworks. Experimental results also confirm the superiority of the triplet loss over the cross-entropy loss in low training data conditions.

[http://dx.doi.org(DOI number)]

I. INTRODUCTION

Habitat destruction induced by global warming and human activities has pushed many avian and amphibian species to the brink of extinction. With this looming threat of population decline and species extinction, a huge escalation can be witnessed in the conservation efforts[1,2]. The surveying and monitoring are the principal steps in any conservation effort. Manual surveying and monitoring are difficult, time-consuming and require experienced personnel[3,4]. Owing to the rich acoustic communication in birds and frogs, automated acoustic monitoring provides an appropriate way to survey different species of interest in their natural habitat and alleviates the requirement of manual monitoring[5]. The bioacoustic signal classification module is the mainstay of such acoustic monitoring systems[6] and often includes tasks such as bird and frog species classification. The major impediment in many bioacoustic classification tasks is the scarcity of the labeled training data. Moreover, the target species and hence, the training data requirements often vary from one ecosystem to other. This makes it unfeasible to collect and label a large amount of bioacoustic data for all the possible species. Thus, there is a requirement for bioacoustic classification frameworks that could provide effective classification in low training data conditions.

In recent times, deep convolution neural networks (CNN) have become the cornerstone for achieving state-of-art performances in various audio classification tasks[7–9]. In comparison to

---

[a)]anshul_thakur@students.iitmandi.ac.in

the shallow learning techniques, CNNs often require large amount of training data (subject to the task in hand) to generalize and provide effective classification. However, the scarcity of the labeled data for many bioacoustic tasks makes it undesirable to utilize these data-intensive CNNs. The lesser amount of training data often leads to over-fitting in CNNs. This over-fitting can be avoided by using regularizers and early stopping, which can restrict the modeling capabilities of the CNNs. Many studies on CNN based audio classification have used data augmentation techniques to overcome the training data scarcity[10,11]. These methods augment the training data with synthetic examples that is generated by deforming the original data. Some common deformations used for data augmentation include pitch alterations and time stretching. However, these augmentation techniques are not always useful and can affect the classification performances[11]. As a result, the effectiveness of these techniques are data dependent and often require a trial-error approach. In the case of bioacoustics, coming up the effective augmentation requires domain-expert knowledge about the nature of vocalizations of each target species. Apart from augmentation, many studies have explored transfer learning for overcoming training data scarcity[12,13]. In the case of CNNs, existing (or pre-trained) networks trained for any audio classification tasks can be fine-tuned for achieving effective performance[14]. Fine-tuning helps in transferring the knowledge from the pre-trained network to the domain and task of interest[15]. In low training data conditions, fine-tuning an existing network is easier and effective than training the network from scratch. Thus, CNN based transfer learning presents an effective way to overcome the training data scarcity in bioacoustic applications.

In literature, CNNs have mildly been explored for different bioacoustic classification tasks. Due to the recently conducted *bird activity detection (BAD) challenges*[16], large datasets have been publicly released for the task of bird activity detection. This led to an influx of CNN-based frameworks that provide state-of-art performance for the aforementioned task[17–19]. However, due to the scarcity of the labeled data, only a few studies have addressed the task of vocalization segmentation and species identification using deep learning approaches. Lostanlen *et*

al.[20] released a bird flight call detection dataset along with a CNN-based benchmark. Salamon *et al.*[21] experimentally showed that the late fusion of scores obtained from CNN (deep learning) and a random forest classifier (shallow learning) results in better performance for the task of bird species classification from flight calls. The same CNN architecture, consisting of three convolution layers and two dense layers, is used in both the aforementioned studies. Ibrahim and Zhuang[22] proposed to use a recurrent neural network (RNN) and CNN to classify grouper species. Tóth[23], Sprengel *et al.*[24] and Piczak[25] utilized spectrogram enhancement methods before applying CNNs to identify bird species from their songs or calls. This spectrogram enhancement helps in removing the effect of overwhelming background disturbances on the classification procedure.

Apart from deep learning, many classical machine learning techniques have been successfully utilized for bioacoustic classification. Stowell and Plumbley[26] proposed spherical K-means based unsupervised feature learning for large scale bird species classification. Building on their work, Thakur *et al.*[27,28] proposed to use archetypal analysis[29] and deep archetypal analysis for obtaining supervised convex representations for bioacoustic classification. Kernel-based extreme learning machines are used by Qian *et el.*[30] for bird species classification. This study utilizes active learning to alleviate the problem of unlabelled bioacoustic data. Many studies have used dynamic kernels (such as the intermediate matching kernel and probabilistic sequence kernels) based support vector machines (SVM) for different bioacoustic classification tasks such as bird activity detection and bird species classification[31–33].

In this work, the authors propose to use CNN-based deep metric learning (DML)[34] for bioacoustic classification. DML deals with learning a mapping from the input space to a compact Euclidean space where similarity among examples is in direct correspondence with distance among them. As a result, DML directly provides class-specific clustering in this space. Thus, a classifier trained in this space can provide better classification than the one trained in the input feature space. This study utilizes CNN powered by the triplet loss[35] to map the input examples to the 128-dimensional embeddings in the de-

sired transformation space. The triplet loss processes three examples, called a triplet, at a time. Triplet consists of an anchor, a positive example and a negative example. The anchor and the positive examples belong to the same class whereas the negative example can be from any other class. A CNN with triplet loss tries to learn a transformation where a triplet constraint is imposed on all the training examples. This constraint states that the distance between the negative-anchor pair should be greater than the positive-anchor by a fixed margin in the transformation space. Only triplets that violate this constraint are chosen for training. More details about the triplet loss and its implementation are in later sections. Although triplet loss has been successfully utilized for many applications such as face recognition[35] and person re-identification[34], it has also received some criticism for its slow convergence[36] on large datasets. This slow convergence can be attributed to the fact that as the size of dataset increases, the possible number of triplets increase cubically[37]. This drawback of triplet loss can prove advantageous in the low training data conditions as the possible number of triplets even in small datasets can be large enough to effectively train a CNN. Moreover, the nature of each triplet (in terms of distance between positive-anchor and negative-anchor pairs) is usually unique. As a result, no redundant information is utilized in the training procedure.

CNN used in the proposed DML framework is characterized by the utilization of different filter sizes in the convolution layers. Each filter size helps in analyzing the input bioacoustic events at a different scale. The smaller filters help in extracting the minute local details whereas the large filters analyze a larger receptive field and help in obtaining the global details from the input bioacoustic event. This notion of multiscalar analysis is inspired by the Inception[38] model that was proposed for large scale image classification. This multiscale CNN is empowered by a dynamic variant of the classical triplet loss[35] to learn the desired transformation space. During training, the margin of the loss function is slowly increased based on a pre-defined heuristic (see section 2). This dynamically varying triplet loss has a twofold advantage:

1. Starting with a smaller margin and slowly increasing the margin can be seen as the
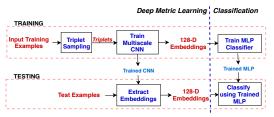

FIG. 1. Proposed DML framework for bioacoustic classification.

warm start. First, the CNN is taught to learn a relatively simpler task of separating the examples of one class from the others in the embedding space by a smaller margin. Then, the complexity of this task is slowly increased by increasing the margin. This warm start helps in better convergence even when the number of classes is very large.

2. Dynamically varying margin increases the number of triplets used for training. It can be attributed to the fact that triplets which satisfy the triplet constraint at a lower margin can violate the constraint as the margin is increased.

The main contribution of this study are as follows:

- To the best of authors' knowledge, this is the first study that utilizes deep metric learning for bioacoustics.

- A simple multiscale CNN is proposed for bioacoustic classification.

- This study experimentally shows that the utilization of triplet loss helps in overcoming the training data scarcity without utilizing any data augmentation and transfer learning.

- A dynamic variant of triplet loss is proposed.

The rest of this paper is organized as follows. In Section II, the proposed DML based classification framework and the dynamic triplet loss is described. Experimental setup is explained in Section III. Experimental results are in discussed in Section IV. Section V concludes this paper.

II. PROPOSED DML FRAMEWORK

In this section, the proposed DML framework for bioacoustic classification is described. The

overall design of the framework is depicted in Fig. 1. The proposed framework is composed of two neural networks: a multiscale CNN and a multi-layer perceptron (MLP). The multiscale CNN, equipped with dynamic triplet loss, is used to learn a transformation from input to the embedding space. The embeddings generated by CNN are given as input to the MLP for learning the discrimination between classes in the embedding space.

This section starts with the feature extraction procedure. Then, the architectures of the proposed multiscale CNN and MLP are described. Later, dynamic triplet loss and other details regarding the training of neural networks are highlighted. Finally, the procedure to classify the bioacoustic signals using the trained DML framework is explained.

A. Feature Extraction

Most bird species such as Passerines are known for producing the harmonically rich sounds. However, there are many species such as woodpeckers, snipes and storks that are characterized by drumming, winnowing, clattering and other mechanically produced sounds. These sounds are more or less percussive in nature. Thus, the difference in harmonic and percussive components of a bioacoustic sound has some class-specific characteristics. This difference in harmonic and percussive components of sounds produced by white-bellied woodpecker and Indian peafowl is evident in Fig. 2. Inspired by this observation, Mel-spectrogram along with its harmonic and percussive components[39] are given as a three-channel input to the proposed framework. The spectrogram of the input audio recording is decomposed into its harmonic and percussive components using the method proposed in Ref. 39. The original, harmonic and percussive spectrograms are multiplied by Mel filterbank to obtain the respective Mel spectrograms that form the three channels of an input example. All three channels are converted to decibel scale and are normalized with respect to the maximum value.

B. Neural Network Designs

1. Multiscale CNN Architecture

The proposed CNN consists of five convolution ($CONV$) layers, four Inception[38] inspired multiscale analysis modules (MAM), three dense layers and has 1,286,410 trainable parameters. Each multiscale analysis module consists of seven convolution layers having different filter sizes that enable the network to analyze each input at different scales. The overall network design is illustrated in Fig. 3(a). The main components of the network are:

Input: As discussed earlier, audio examples, represented by Mel-spectrograms and their harmonic and percussive components ($40 \times M \times 3$, 40 Mel-filters and M frames) are given as input.

Multiscale Analysis Module: Multiscale analysis modules utilize kernels of different sizes (1×1, 3×3. 5×5 and 7×7). This multiscale analysis helps in better feature extraction from short duration vocalizations (such as birdsong syllables or flight calls) as well as from the longer vocalizations such as birdsong phrases. The shorter vocalizations occupy smaller spatial space on Mel-spectrograms as compared to the longer vocalizations. Hence, a smaller kernel size is more appropriate for the shorter bioacoustic events and vice-versa. The smaller 3×3 filter helps in learning minute details from an input Mel-spectrogram where as the larger filters (5×5 and 7×7) help in capturing more global traits due to the larger receptive fields. In bioacoustics, these minute details can be low energy harmonics or vocalizations recorded in a far-field setting. The global traits can include the information about the frequency contents or bandwidth and the coarser time-frequency modulations of bioacoustic events.

Each MAM receives an input of 64 feature maps ($N \times M \times 64$) that are processed by seven convolution layers arranged in four parallel strands as shown in Fig. 3(b). First strand contains one convolution layer that has 64 filters of 1×1 kernel size. These filters are mainly concerned with selecting the inter feature map patterns rather than the spatial analysis of feature maps. It is a known fact that each feature map has some complementary information[40]. Hence, learning these inter feature map patterns can be helpful in distinguishing one class from another. The second, third and fourth strands consist of two convolution layers. The first convolution layer in all these strand consists of 32 filters of 1×1, whereas, the second convolution layers have 64 filters of sizes 3×3, 5×5 and 7×7 respectively. Here the 1×1 convolution layers
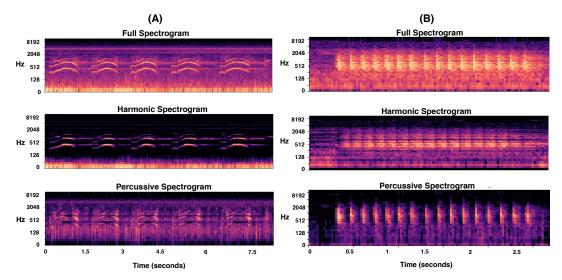
FIG. 2. Difference in harmonic and percussive components of sounds produced by (A) Indian peafowl and (B) white-bellied woodpecker.

serve two purposes: 1). It decrease the number of input channels from 64 to 32 and reduce the computational requirements for the following convolutional layer in each strand. 2). As discussed earlier, 1×1 filters are used for selecting the discriminative feature patterns from the input feature maps. It must be noted that each strand utilizes a separate 1×1 convolution layer. As the appropriate feature patterns may be scale dependent, a separate 1×1 convolution layer provides independence in the feature selection for different scales in each parallel strand. The output of these layers are processed by convolution layers having filter size of 3×3, 5×5 and 7×7 in second, third and fourth strand respectively for multiscale analysis. The difference in responses of filters of different strands is illustrated in Fig. 4. The feature maps obtained from all the four strands are concatenated in a channel-wise manner to output 256 feature maps from each module. It must be noted that zero-padding is used to make sure that each feature map is of same dimension before concatenation.

Bottleneck, Global Pooling and Dense Layers: The network consists of five convolution layers having 64 filters of 3×3. Apart from the first convolution layer, all other convolution layers act as the bottleneck layers. They use strided convolutions to down-sample the feature maps by a factor of 2 (or 5 in case of the last convolution layer) along the Mel-energy axis. Apart from down-sampling, they also help in selecting the relevant features from the feature maps, obtained from multiscalar analysis modules, by analyzing the inter feature map correlations[40]. Due to this feature selection, the number of channels are decreased from 256 (generated by multiscale analysis module) to 64 which also helps in decreasing the computation requirements for the corresponding layers. After all the convolution layers and multiscale analysis modules, global average pooling (GAP) is applied to obtain a 64-D vector. This averaging operation helps in making the framework invariant towards the time differences in onsets-offsets of the bioacoustic events in audio recordings or their Mel-spectrograms. Then, this 64-D vector is passed to the dense layers. The network has three dense layers having 256, 128 and 128 hidden units.

Activation, Regularization and Optimizer: Each convolution layer (whether standalone or in multiscale analysis module) and first two dense layers are followed by rectified linear unit activation. The output of the last dense layer is normalized to have the unit norm such that embeddings produced by the proposed CNN lies on a unit hypersphere. A dropout of 0.5 is added before each dense layer. Along with dropout, exponential weight decay of 0.0001 is also used to avoid over-fitting and improve generalization[41]. Adagrad with a learning rate 0.001 is used as optimizer.
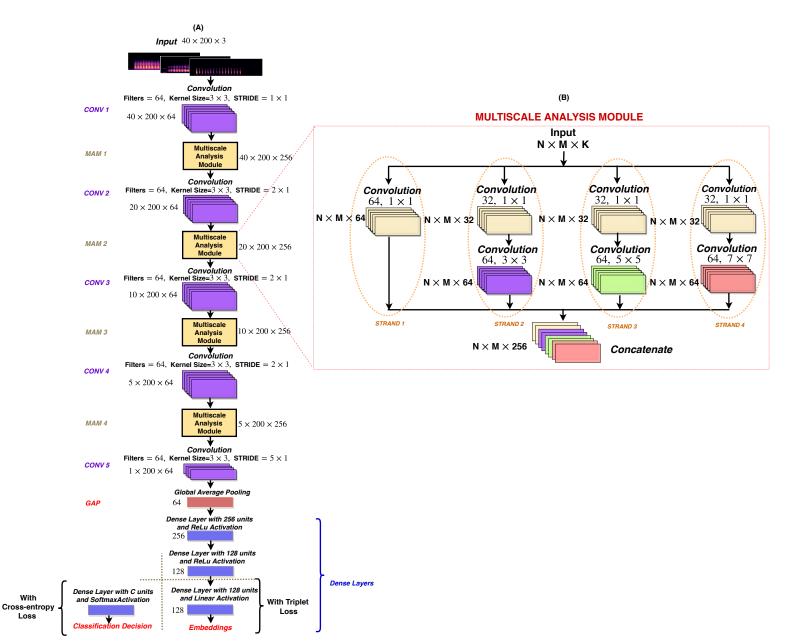
FIG. 3. Illustration of (A) the Proposed multiscale CNN architecture and (B) a multiscale scale analysis module. To use the proposed multiscale CNN architecture for classification (with cross-entropy loss), the last layer of the architecture is replaced by a fully connected layer having C (number of classes) units and softmax activation.

2. MLP Architecture

The MLP used in the proposed framework consists of three layers: an input layer with 128 units, a hidden layer with 256 units and an output layer with C units (C is the number of classes). The rectified linear unit function is used as activation for the hidden layer neurons. The weight optimization is performed by utilizing Adam solver. A constant learning rate of 0.001 and L2 regularization term of 0.0001 is used for optimization.

C. Multiscale CNN Training: Dynamic Triplet Loss

Dynamic triplet loss is utilized for training the proposed multiscale CNN. The aim of the triplet loss is to learn an embedding space where all possible triplets satisfy the triplet constraint. Let $(\mathbf{x}_i^a, \mathbf{x}_i^p, \mathbf{x}_i^n)$ be ith triplet of embeddings belonging to the possible set of triplets present in the training data embeddings. $\mathbf{x}_i^a$, $\mathbf{x}_i^p$ and $\mathbf{x}_i^n$ be
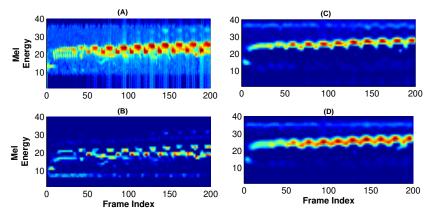
FIG. 4. (A) Mel-spectrogram given as input to the proposed CNN. (B), (C) and (D) depicts the filter responses obtained for 6^{th} 3×3, 32^{nd} 5×5 and 62^{nd} 7×7 filters of the first multiscale analysis module (*MAM 1*) of the trained multiscale CNN. These particular filters are chosen for their expressivity.

an anchor, a positive and a negative example. Then, the triplet constraint can be defined as:

$$\|\mathbf{x}_i^a - \mathbf{x}_i^n\|_2^2 - \|\mathbf{x}_i^a - \mathbf{x}_i^p\|_2^2 \geq \alpha. \quad (1)$$

Here α is the enforced margin or distance between the positive examples and the negative examples. Thus, based on the triplet constraint, the loss function to be minimized is[35]:

$$\mathcal{L} = \sum_{i=1}^{N} \max(\|\mathbf{x}_i^a - \mathbf{x}_i^p\|_2^2 - \|\mathbf{x}_i^a - \mathbf{x}_i^n\|_2^2 + \alpha, 0), \quad (2)$$

where N is the possible number of triplets to be used for training.

For training, triplets are sampled from a mini-batch in an online fashion. A forward pass is performed on the CNN to obtain embeddings for a mini-batch of input examples. These embeddings are processed to select triplets that are later used for optimizing the current state of CNN. The performance of triplet loss is directly dependent on the choice of triplets to be used for training. Choosing triplets that satisfy triplet constraint (equation 1) will lead to no change in the state of CNN. Hence, triplets violating the triplet constraint are of interest. These triplets are of two types: hard triplets and semi-hard triplets. A triplet $\mathcal{X} = (\mathbf{x}^a, \mathbf{x}^p, \mathbf{x}^n)$ is classified as hard or semi-hard according to the following criteria[35]:

$$\mathcal{X} \text{ is } \begin{cases} hard & d(\mathbf{x}^a, \mathbf{x}^n) < d(\mathbf{x}^a, \mathbf{x}^p) \\ semi-hard & d(\mathbf{x}^a, \mathbf{x}^p) < d(\mathbf{x}^a, \mathbf{x}^n) \text{ and} \\ & d(\mathbf{x}^a, \mathbf{x}^n) < d(\mathbf{x}^a, \mathbf{x}^p) + \alpha. \end{cases}$$

Here $d()$ refers to the Euclidean distance.

Although hard triplets appear to be more informative for training, they result in higher loss values, leading to the larger weight updates. These larger weight updates result in significant change to the current state of network, hence, undoing the optimization work done by the previous weight updates. Thus, utilizing these hard triplets may lead to instability during training. It has been shown in Ref. 35 that semi-hard triplets often leads to faster convergence and effective training than the hard triplets. Building on this information, the semi-hard triplets are used for training the proposed CNN. In semi-hard triplet, the distance between anchor-negative pairs is greater than the anchor-positive pairs as desired. However, this distance is not greater than the desired separation margin α. Thus, the weight updates obtained in case of semi-hard triplets are not as large as the hard triplets, leading to a stable training.

As discussed in Section I, a dynamic variant of the triplet loss is used in this work. In the proposed implementation of triplet loss, α or the margin is considered as a dynamic variable whose value is changed over the course of training. The overall procedure to calculate dynamic triplet loss is depicted in Algorithm 1.

We start with a small margin, $\alpha = 0.2$, and force the network to learn the embedding space

where examples of each class are separated from other by a distance of 0.2. As the network is trained, the number of semi-hard triplets mined from training dataset decreases. If this number of mined triplets is less than a pre-defined threshold for three consecutive iterations, the value of α is incremented by 0.05. Again, the network is trained to satisfy the new triplet loss induced by new value of α. This process is continued till α reaches a pre-defined maximum value of 0.6. As discussed in Section I, this dynamic triplet loss can lead to faster convergence. Along with threshold, values of α and margin updates are determined experimentally.

Implementation Details: Each mini-batch is composed of at least 5 examples per class. Hence, all classes are represented in a mini-batch. The semi-hard triplets are sampled from this mini-batch and are used for training the CNN. The number of triplets that can be processed simultaneously (let us say triplet batch size) is often limited by the available GPU memory. In our implementation, we set this triplet batch size to be 150 input examples or 50 semi-hard triplets. Thus, the semi-hard triplets sampled from a mini-batch are presented to the CNN in an iterative manner where during each iteration, the triplet batch (having 50 or less triplets) is used to calculate triplet loss and update the weights. The number of mini-batches in an epoch are limited to 1000. The CNN is trained till it reaches the desired value of margin (α) and the triplet loss is converged to zero. Figure 5 depicts the number of triplets and triplet loss calculated from each triplet batch during the training of multiscale CNN on CLO-43SD dataset (see Section III for dataset details).

D. Classification

As illustrated in Fig. 1, first the multiscale CNN is trained to learn the transformation or embedding space using dynamic triplet loss. This trained CNN is used to extract 128-D embeddings from all the training examples. These embeddings show high class-specific signatures as embeddings of each class occupy a compact Euclidean space which is separated from every other classes by a significant margin. This behaviour is highlighted in Fig. 6. Then, a multilayer perceptron (MLP) with Adam optimizer (described in Section II B 2) is trained on these embeddings for classification. During testing

phase, embeddings are extracted from the test examples using the trained CNN. These embeddings are classified using the trained MLP classifier.

III. EXPERIMENTAL SETUP

In this section, we describe datasets, comparative methods, and parameter setting for the performance evaluation.

A. Datasets Used

The performance of the proposed DML framework is evaluated on three different datasets:

- **Birdcalls71**: This dataset consists of audio recordings of 71 different bird species that are obtained from three different sources. The recordings of 38 bird species were provided by the Macaulay Library[42] on an academic license. The recordings of 7 bird species were downloaded from bird database maintained by Art & Science Centre, UCLA[43]. The recordings of 26 bird species were obtained from the Great Himalayan national park (GHNP) dataset[31] and were provided on request. All these recordings are sampled at 44.1 kHz and vary in duration from 0.5 to 320 seconds. Due to licensing issues, the authors can not make this dataset public. However, the processed Mel-spectrograms extracted from these recordings are hosted on a public platform for analysis[44].

- **Anuran dataset**: The publicly available Anuran dataset contains audio recordings of 10 different frog species found in the Amazon rainforest. These recordings are sampled at 44.1 kHz and are of variable durations (from 3 to 360 seconds). More details about the frog species and dataset can in found in Ref. 6.

- **CLO-43SD**: This public dataset is provided by Salamon *et al.*[45] and contains processed Mel-spectrograms of flight calls of 43 different North American wood-warblers. These Mel-spectrograms were extracted from audio recordings sampled at 22.05 kHz using 11.6 ms frame size with an overlap of 1.25 ms and 40 Mel bands. It must be noted that 11.6 ms frame size is optimum to analyze flight calls[21,45].
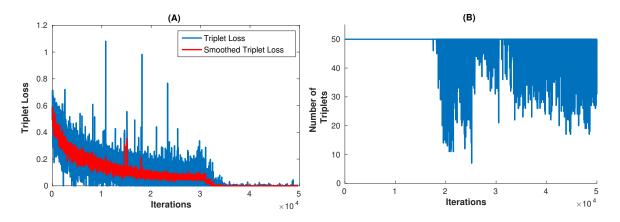
FIG. 5. Illustration of (A) triplet loss and (B) number of semi-hard triplets involved during during each iteration of training the multiscale CNN with CLO-43SD dataset. The red plot in (A) represents the smoothed triplet loss over a window of 15.
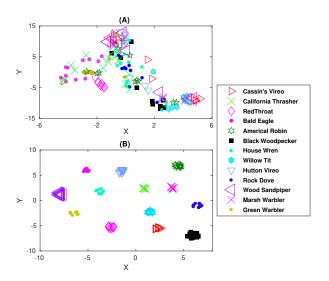

FIG. 6. Two dimensional t-SNE visualization of 128-d embeddings extracted from audio examples of 12 different bird species using (A) untrained multiscale CNN and (B) fully trained multiscale CNN.

- **Combined**: To analyze the scalability of the proposed framework, all three datasets are combined together to create a larger dataset having 124 classes.

B. Comparative studies

The classification performance of the proposed DML framework is compared with six existing bioacoustic and audio classification frameworks. Table I lists different methods used for the performance comparison. These com-

parative methods include both shallow and deep learning based methods. The three mentioned shallow learning baselines include polynomial kernel based extreme learning machines (KELM)[30] and random forest classifiers, wherein (1) the KELM is trained on low-level audio descriptors while the later one is trained on (2) the unsupervised and (3) supervised feature representations. The unsupervised feature representations are obtained from spherical K-means[26] (SKM) while the supervised representations[28] are acquired by deep convex matrix factorization (DCR). The input feature representations (Mel-spectrogram and compressed spectral frames) used in the respective studies are also used here. In SKM, frame-wise feature representations for each input example are aggregated using mean and standard deviation to obtain a fixed dimensional representation. In DCR, a random forest classifier is trained on a frame-wise deep convex representation and a voting rule is used on these frame-wise decisions to classify the input example.

Apart from these shallow techniques, CNN proposed by Salamon *et al.*[21] (SAL) and VGG (used in Ref. 9 for audio classification) are used as deep learning baselines. The total number of trainable parameters in SAL and VGG are 1,226,554 and 9,653,831 respectively (if last dense layer has 71 units). To evaluate the proposed DML framework against transfer learning, we fine-tuned VGG network for bioacoustic classification. This VGG network is pre-trainied on AudioSet database (see Ref. 9 for more details).

The dense layers of VGG are changed with three dense layers having 256, 128 and C (number of classes) hidden units. A dropout of 0.5 is used before each dense layer. The first two dense layers have Relu activation where as the last dense layer is followed by softmax activation. The final baseline is the proposed multiscale CNN with the cross-entropy loss (MS-CNN). The last dense layer of the propose CNN is replaced with a dense layer having C units and softmax activation (C is the number of target classes). The Keras implementations of these CNN baselines are publicly available along with datasets[44].

C. Parameter Setting

All the audio recordings used for experimentation are divided into fixed length segments of 2 seconds. These segments are used for training and performance comparison. If the duration of any recording is less than 2 seconds, then the signal is repeated (from the beginning) to force the fixed duration of 2 seconds. The short-term spectral analysis is done using a 20 ms frame size with 50% overlap to obtain respective feature representations (from all databases except CLO-43SD) used in all the considered studies. Thus, each input example consists of 200 frames. In case of CLO-43SD dataset, frames of a pre-computed Mel-spectrogram are repeated to obtain a fixed number of frames i.e. 200 per example (for maintaining uniformity between datasets). It must be noted that there is a difference in frame sizes used for the short-term analysis done for CLO-43SD and other datasets to extract Mel-spectrograms. However, this difference is ignored to create a combined dataset for the sake of analyzing the scalability of the proposed DML framework.

For implementing SKM (for all datasets), spherical K-means with 128 clusters and random forest with 100 trees is used. For implementing DCR, a three-level archetypal analysis based matrix factorization (with an order of 128, 64 and 32) is used to learn the class-specific dictionaries. A random forest with 100 trees is used for classification. For implementing polynomial kernel and extreme learning machine in KELM, the parameter values used in Qian et al[30] were found to be optimal for all the datasets used in this study. All the aforementioned parameters are empirically determined on a validation dataset.

For VGG, VGG with fine tuning (VGG-FT) and MS-CNN, a batch size of 32 and Adam optimizer with a learning rate of 0.001 is used. These model are trained for 200 epochs and checkpoints are used after each epoch to determine the setting that provides least validation loss. For implementing the network proposed in Salamon *et al.* (SAL), the details and parameter setting described in the respective study is also used here. The implementation details of the proposed DML framework (MS-CNN-TL) are already discussed in Section II. The multiscale CNN is trained using triplet loss for 150 epochs (each epoch consists of 1000 iterations). After 150 epochs, the MS-CNN-TL shows full convergence. In the current context, the convergence simply means that all semi-hard triplets in the training dataset satisfy the triplet constraint across all dynamically chosen margins.

D. Train-Test Data distribution and Performance Metric

Birdcalls71, CLO-43SD and Anuran datasets have 1982, 5428 and 2445 examples respectively. All datasets considered in this study exhibit significant imbalance in the number of examples per class. In Birdcalls71, the number of examples per class vary from 9 to 75. Similarly, in CLO-43SD, this number varies from 10 to 1256. Thus, a single class of CLO-43SD forms about 23% of the dataset. To avoid this huge class imbalance, this class is not included in the study. The number of examples per class varies from 50 to 705 in Anuran dataset. The 50%, 15% and 35% of the examples per class (in all datasets) are randomly chosen for training, validation and testing. All these files are publicly available at *Figshare*[44].
Performance Metric: The imbalance between classes is large, however, we desire to give equal weightage to the classification performance obtained for each class. Hence, macro F1-score[46] is used as a metric for the performance comparison. The macro F1-score is the average of class-specific F1-scores where F1-score is the harmonic mean of precision and recall.

IV. RESULTS AND DISCUSSION

In this section, the classification performances of the proposed DML framework and different baselines are presented. Apart from that, the significance of using harmonic-percussive components and generalization of the proposed

framework are also discussed. Finally, the extension of the proposed framework for open-set classification is described.

A. Classification Performance

Fig. 7 shows the classification performances of different methods on all four datasets. Following can be inferred from the analysis of this figure:

- Shallow learning techniques (DCR, SKM and KELM) are significantly outperformed by CNN based frameworks including the proposed MS-CNN and MS-CNN-TL.

- MS-CNN performs better than VGG and SAL on all but Anuran dataset, highlighting the superiority of the proposed multiscale CNN. On Anuran dataset, all these baselines show comparable performances.

- VGG-FT outperforms VGG and SAL while showing similar performance to that of MS-CNN. This shows that the utilization of transfer learning or fine-tuning a pre-trained model improves the classification performance.

- MS-CNN-TL (proposed) outperforms all baselines including MS-CNN on all datasets. This confirms the claim that utilizing triplet loss based DML framework leads to a better classification than the cross-entropy loss based CNN. Moreover, the better performance of MS-CNN-TL than VGG-FT shows that the proposed framework provides effective classification without using any fine-tuning or pre-trained model.

- The class unbalance has no significant effect on the performance of MS-CNN-TL as evident from the high macro F1-score across all datasets. Since the target of triplet loss based metric learning is to separate examples of one class from other by a fixed margin, the difference in number of examples per class may not have a significant effect the learning procedure. Moreover, as discussed earlier, during mini-batch creation, at least 5 examples per class are always present in each batch and each class is represented in the training procedure.

B. Generalization of the proposed DML framework

In bioacoustic classification tasks, the training examples often do not contain the whole repertoire of vocalizations that a species can produce. In field conditions, the test examples often contain vocalizations that are not used for training. Thus, an effective classification framework must be able to generalize on these unseen examples. To study the generalization ability of the proposed framework, the unseen vocalizations (not included in training train the model) of Cassin's vireo (one of the species in Birdcalls71 dataset) are used. The ten unseen song phrases are extracted from the audio recordings available at https://goo.gl/x17fYf and are provided for analysis along with Birdcalls71 dataset. The t-SNE representation of the embeddings extracted from these unseen song phrases using the trained MS-CNN-TN are shown in Fig. 8. The analysis of this figure makes it clear that embeddings generated from the unseen song phrases of Cassin's vireo exhibit more similarity to the training Cassin's vireo examples than embeddings of other species. This is attributed to the fact that the triplet loss used in the proposed framework deals with grouping the similar vocalizations together and separating them from the dissimilar examples. Generally, vocalizations of a species are more similar to each other than the sounds produced by other species (though exceptions are always present in natural systems). As a result, the embeddings extracted from seen or unseen vocalizations of a species are bound to be grouped together. Thus, the utilization of deep metric learning (DML) helps in overcoming small variations in the nature of vocalizations as well as differences in the recording environment during training and testing.

C. Effect of using harmonic and percussive components

To analyze the effect of using harmonic and percussive components along with mel-spectrograms, MS-CNN (with cross-entropy loss) and MS-CNN-TL (with triplet loss) are trained on Birdscalls71 dataset using both three-channeled input and Mel-spectrograms only. In each scenario, a model is trained 10 times to compensate for the effects of random weight initialization. The violin plots depicting the classification performances achieved in both these scenarios are illustrated in Fig. 9. The anal-
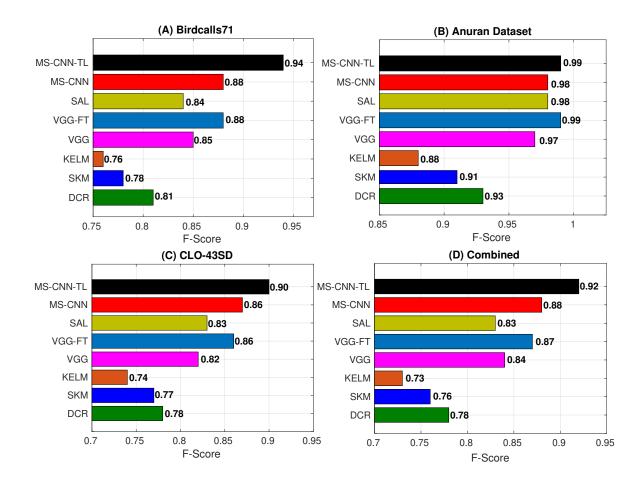
FIG. 7. Classification performances of the proposed framework along with various baselines on (A) Birdcalls71, (B) Anuran, (C) CLO-43SD and (D) Combined datasets.
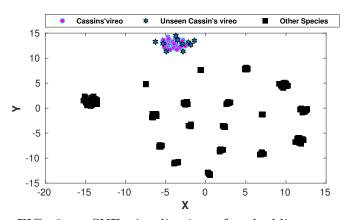

FIG. 8. t-SNE visualization of embeddings generated from seen and unseen Cassin's vireo song phrases using MS-CNN-TL.

ysis of these violin plots makes it clear that utilizing percussive and harmonic components along with Mel-spectrograms helps in achieving a better classification performance. This demonstrates that the differences in harmonic and per-

cussive components of animal sounds can be exploited for species classification.

D. Exploring MS-CNN-TL for open set classification

Open set classification is a challenging issue in designing bioacoustic classification frameworks for field conditions. Current frameworks including the proposed MS-CNN-TL assign any input test example to a class that exhibits maximum similarity, even if this test example does not belong to any of the classes involved in training. Thus, there must be a way to reject such examples without affecting the classification performance. To tackle open set classification, we propose to use the metric learning module of MS-CNN-TL. First an input test example is fed to MS-CNN-TL and an output label is obtained. Now, if the embedding of test example lies at a significant distance from training embeddings of
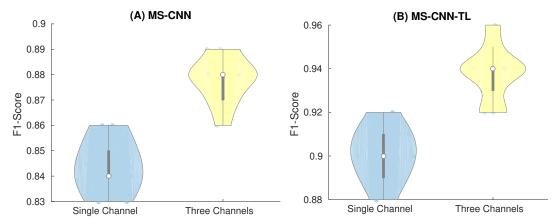
FIG. 9. Violin plots depicting the effect of utilizing the harmonic and percussive components along with Mel-spectrogram on the classification performances of (A) MS-CNN and (B) MS-CNN-TL.

the predicted class, it can be considered as an outlier and must be rejected.

To model the distance from training embeddings, an unimodel Gaussian distribution is utilized. The embeddings of training and validation examples of each class are obtained from the trained MS-CNN-TL. The training embeddings are averaged to obtain a mean vector. A Gaussian distribution is fitted over the distance of validation embeddings from this mean vector. The maximum likelihood estimation is used to calculate the mean and variance of these Gaussian distributions. During testing, first an input test example is classified by MS-CNN-TL. Then, the distance of the test embedding from the mean vector of the predicted class is calculated. Finally, the likelihood of this distance is computed with respect to the Gaussian distribution of the predicted class. If this likelihood is less than 0.5, the test example is considered as outlier and is rejected.

A small experiment is designed to analyze the rejection accuracy of the above mentioned framework. First, MS-CNN-TL is trained on Birdcalls71 dataset and Anuran dataset is used for outlier rejection. Then, the framework is trained on Anuran dataset and Birdcalls71 is used for outlier rejection. The results of this experiment are documented in Table II. The analysis of this table makes it clear that in both setups, MS-CNN-TL with the aforementioned rejection mechanism is able to reject the outliers with good accuracy of 97% and 95%. However, a small relative drop in macro F1-scores is observed. This shows that as expected, incorporating outlier rejection mechanism in MS-CNN-

TL leads to a small drop in classification performance. However, this observed classification performance is still competitive in comparison to the other methods considered in this study (see Fig. 7).

V. CONCLUSION

In this paper, the authors presented a deep metric learning based framework for bioacoustic classification. The authors proposed a multiscale CNN and a dynamic triplet loss to achieve effective deep metric learning even in scarcity of the training data. The proposed multiscale CNN utilizes different kernel sizes to extract features at different granularities. Whereas, the nature of dynamic triplet loss significantly increases the amount of triplets during course of training. The embeddings extracted from multiscale CNN based DML are used as a feature representation for classifying an input example. The experimental results on four different datasets show that the proposed DML based classification framework performs better than existing bioacoustic classification frameworks and various CNN architectures trained using cross-entropy loss. The authors also presented a simple augmentation that enables the proposed framework to perform open-set classification.

A major drawback of the proposed framework (and most of the existing metric learning frameworks) is that it can not handle multi-label classification. Future work may involve developing metric learning frameworks to overcome this drawback.

[1] F. van Bommel, "Birds in Europe: Population estimates, trends and conservation status," British Birds **98**, 269–271 (2005).

[2] S. A. Cushman, "Effects of habitat loss and fragmentation on amphibians: a review and prospectus," Biological conservation **128**(2), 231–240 (2006).

[3] A. L. Borker, M. W. McKown, J. T. Ackerman, C. A. Eagles-Smith, B. R. Tershy, and D. A. Croll, "Vocal activity as a low cost and scalable index of seabird colony size," Conservation biology **28**(4), 1100–1108 (2014).

[4] B. J. Furnas and R. L. Callas, "Using automated recorders and occupancy models to monitor common forest birds across a large geographic region," The Journal of Wildlife Management **79**(2), 325–337 (2015).

[5] T. S. Brandes, "Automated sound recording and analysis techniques for bird surveys and conservation," Bird Conservation International **18**(S1), S163–S173 (2008).

[6] B. Gatto, J. Colonna, E. M. dos Santos, and E. F. Nakamura, "Mutual singular spectrum analysis for bioacoustics classification," in *Proceedings of Mach. Learn. Sig. Process. (MLSP)* (Sept., 2017).

[7] E. Cakir, G. Parascandolo, T. Heittola, H. Huttunen, T. Virtanen, E. Cakir, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen, "Convolutional recurrent neural networks for polyphonic sound event detection," IEEE Trans. Audio, Speech, Language Process **25**(6), 1291–1303 (2017).

[8] Y. Xu, Q. Kong, W. Wang, and M. D. Plumbley, "Large-scale weakly supervised audio classification using gated convolutional neural network," in *Proceedings of Int. Conf. Acoust. Speech, Signal Process. (ICASSP)* (2018), pp. 121–125.

[9] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, *et al.*, "CNN architectures for large-scale audio classification," in *Proceedings of Int. Conf. Acoust. Speech, Signal Process. (ICASSP)* (2017), pp. 131–135.

[10] J. Salamon and J. P. Bello, "Deep convolutional neural networks and data augmentation for environmental sound classification," IEEE Signal Processing Letters **24**(3), 279–283 (2017).

[11] R. Lu, Z. Duan, and C. Zhang, "Metric learning based data augmentation for environmental sound classification," in *Proc. of IEEE Workshop on Applications of Signal Processing to Audio and Acoustics* (2017), pp. 1–5.

[12] P. Hamel, M. E. P. Davies, K. Yoshii, and M. Goto, "Transfer learning in MIR: Sharing learned latent representations for music audio classification and similarity," in *Proceedings of Int. Conf. Music Info. Retrieval* (2013).

[13] S. Ntalampiras, "Bird species identification via transfer learning from music genres," Ecological Informatics **44**, 76–81 (2018).

[14] N. Tajbakhsh, J. Y. Shin, S. R. Gurudu, R. T. Hurst, C. B. Kendall, M. B. Gotway, and J. Liang, "Convolutional neural networks for medical image analysis: Full training or fine tuning?," Trans. on Medical Imaging **35**(5), 1299–1312 (2016).

[15] S. Poria, I. Chaturvedi, E. Cambria, and A. Hussain, "Convolutional mkl based multimodal emotion recognition and sentiment analysis," in *Proceedings of Int. Conf. Data Mining* (2016), pp. 439–448.

[16] D. Stowell, M. D. Wood, H. Pamuła, Y. Stylianou, and H. Glotin, "Automatic acoustic detection of birds through deep learning: the first bird audio detection challenge," Methods in Ecology and Evolution (2018).

[17] E. Cakir, S. Adavanne, G. Parascandolo, K. Drossos, and T. Virtanen, "Convolutional recurrent neural networks for bird audio detection," in *Proceedings of European Sig. Process. Conf. (EUSIPCO)* (2017), pp. 1744–1748.

[18] T. Grill and J. Schlüter, "Two convolutional neural networks for bird detection in audio signals," in *Proceedings of European Sig. Process. Conf. (EUSIPCO)* (2017), pp. 1764–1768.

[19] T. Pellegrini, "Densely connected cnns for bird audio detection," in *Proceedings of Eusipco* (2017), pp. 1734–1738.

[20] V. Lostanlen, J. Salamon, A. Farnsworth, S. Kelling, and J. P. Bello, "Birdvox-full-night: A dataset and benchmark for avian flight call detection," in *Proceedings of Int. Conf. Acoust. Speech, Signal Process. (ICASSP)* (2018), pp. 266–270.

[21] J. Salamon, J. P. Bello, A. Farnsworth, and S. Kelling, "Fusing shallow and deep learning for bioacoustic bird species classification," in *Proceedings of Int. Conf. Acoust. Speech, Signal Process. (ICASSP)* (2017), pp. 141–145.

[22] A. K. Ibrahim, H. Zhuang, L. M. Chérubin, M. T. Schärer-Umpierre, and N. Erdol, "Automatic classification of grouper species by their sounds using deep neural networks," The Journal of the Acoustical Society of America **144**(3), EL196–EL202 (2018).

23 B. P. Tóth and B. Czeba, "Convolutional neural networks for large-scale bird song classification in noisy environment.," in *CLEF (Working Notes)* (2016), pp. 560–568.

24 E. Sprengel, M. Jaggi, Y. Kilcher, and T. Hofmann, "Audio based bird species identification using deep learning techniques.," in *CLEF (Working Notes)* (2016), pp. 547–559.

25 K. J. Piczak, "Recognizing bird species in audio recordings using deep convolutional neural networks.," in *CLEF (Working Notes)* (2016), pp. 534–543.

26 D. Stowell and M. D. Plumbley, "Automatic large-scale classification of bird sounds is strongly improved by unsupervised feature learning," PeerJ **2**, e488 (2014).

27 A. Thakur, V. Abrol, P. Sharma, and P. Rajan, "Compressed convex spectral embedding for bird species classification," in *Proceedings of Int. Conf. Acoust. Speech, Signal Process. (ICASSP)* (April, 2018).

28 A. Thakur, V. Abrol, P. Sharma, and P. Rajan, "Deep convex representations: Feature representations for bioacoustics classification," in *Proceedings of Interspeech* (2018).

29 Y. Chen, J. Mairal, and Z. Harchaoui, "Fast and robust archtypal analysis for representation learning," in *Proceedings of Comp. Vision Pattern Recog. (CVPR)* (2014), pp. 1478–1485.

30 K. Qian, Z. Zhang, A. Baird, and B. Schuller, "Active learning for bird sound classification via a kernel-based extreme learning machine," J. Acoust. Soc. Am. **142**(4), 1796–1804 (2017).

31 D. Chakraborty, P. Mukker, P. Rajan, and A. Dileep, "Bird call identification using dynamic kernel based support vector machines and deep neural networks," in *Proceedings of Int. Conf. Mach. Learn. App.* (2016).

32 A. Thakur, R. Jyothi, P. Rajan, and A. Dileep, "Rapid bird activity detection using probabilistic sequence kernels," in *Proceedings of European Sig. Process. Conf. (EUSIPCO)* (2017), pp. 1754–1758.

33 V. Abrol, P. Sharma, A. Thakur, P. Rajan, A. Dileep, and A. K. Sao, "Archetypal analysis based sparse convex sequence kernel for bird activity detection," in *Signal Processing Conference (EUSIPCO), 2017 25th European*, IEEE (2017), pp. 1774–1778.

34 D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Deep metric learning for person re-identification," in *Proceedings of Int. Conf. Pattern Recognition (ICPR)* (2014), pp. 34–39.

35 F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of Comp. Vis. Pattern Recog.*, (2015), pp. 815–823.

36 O. Rippel, M. Paluri, P. Dollar, and L. Bourdev, "Metric learning with adaptive density discrimination," in *Proceedings of Int. Conf. Learn. Represent.* (2016).

37 G. Pahariya, B. Ravindran, and S. Das, "Dynamic class learning approach for smart cbir," in *National Conference on Computer Vision, Pattern Recognition, Image Processing, and Graphics* (2018), pp. 327–337.

38 C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning.," in *Proceedings of AAAI* (2017), Vol. 4, p. 12.

39 J. Driedger, M. Müller, and S. Disch, "Extending harmonic-percussive separation of audio signals.," in *Proceedings of ISMIR* (2014), pp. 611–616.

40 A. Pankajakshan, A. Thakur, D. Thapar, P. Rajan, and A. Nigam, "All-conv net for bird activity detection: Significance of learned pooling," in *Proc. Interspeech* (2018).

41 A. Krogh and J. A. Hertz, "A simple weight decay can improve generalization," in *Proceedings of Advances in neural information processing systems* (1992), pp. 950–957.

42 http://www.macaulaylibrary.org.

43 http://artsci.ucla.edu/birds/database.html.

44 https://figshare.com/s/4af71d71d94e04afcd5f.

45 J. Salamon, J. P. Bello, A. Farnsworth, M. Robbins, S. Keen, H. Klinck, and S. Kelling, "Towards the automatic classification of avian flight calls for bioacoustic monitoring," PloS one **11**(11), e0166866 (2016).

46 V. Van Asch, "Macro and micro-averaged evaluation measures," Belgium: CLiPS (2013).

Algorithm 1: Training CNN using Dynamic Triplet Loss

---

input : $f()$: CNN (randomly initialized)
X: Training dataset

output: $f()$: Trained CNN for metric learning

---

1   $\alpha = 0.2$ // Initial value of margin
2   $count\_list = [\ ]$ // List to store number of triplets sampled in each iteration
3   $thresh = 15$ // Threshold for margin update (determined empirically)
4   $K = 100$ // Number of epochs
5   **for** $J \leftarrow 1$ **to** K **do**
6      $\mathcal{I} = \text{createBatches}(X, n)$ // Returns n batches stored in List $\mathcal{I}$
7   **for** $i \leftarrow 1$ **to** n **do**
8          $\mathcal{E} = f(\mathcal{I}[i])$ // Forward pass to get embeddings for ith batch
9          $\mathcal{T}, t = \text{getTriplets}(\mathcal{E}, \alpha)$ // Returns $\mathcal{T}$, a set containing t semi-hard triplets, sampled from ith batch
10         $count\_list.\text{append}(t)$ // Store the number of semi-hard triplets sampled from ith batch
11         $num = len(count\_list)$ // Number of elements in $count\_list$ currently
12   **if** $num \geq 3$ AND $\alpha \leq 0.6$ **then**
13   **if** $(count\_list[num] < thresh$ AND $count\_list[num-1] < thresh$ AND $count\_list[num-2] < thresh)$ **then**
14               $\alpha = \alpha + 0.05$ // Margin update
15   **end**
16   **end**
17         $L = \text{calculateTripletLoss}(f(), \mathcal{T}, \alpha)$ // Calculate Triplet Loss using Equation 2
18         $f() = \text{UpdateWeights}(f(), L)$ // Back-propagate L through $f()$ to get the updated $f()$
19   **end**
20 **end**

TABLE I. Comparative methods used for the performance evaluation.

Method	Nature
Spherical K-means and Random forest (SKM)[26]	Unsupervised Feature Learning
Deep Convex and Random Forest (DCR)[28]	Supervised Dictionary Learning
Kernel Based Extreme Learning Machines (KELM)[30]	Shallow Learning
VGG[9]	CNN
Fine-tuned VGG (VGG-FT)[9], Pre-trained on AudioSet	Transfer Learning/CNN /Deep Learning
CNN proposed by Salamon et al.[10] (SAL)	CNN
Multiscale CNN with Cross Entropy Loss (MS-CNN)	CNN/Deep Learning
Multiscale CNN with Triplet Loss and MLP (MS-CNN-TL)	CNN based Deep Metric Learning

TABLE II. Classification and outlier rejection performances of the proposed MS-CNN-TL framework in different training-testing setup. In each setup, a dataset is used for training and classification evaluation whereas a different dataset is used for evaluating the outlier rejection mechanism.

		Classification Setup	Outlier Rejection Setup	
Training Dataset	Testing Dataset	Classification Performance (Macro F1-score)	Outlier Dataset	Rejection Accuracy (%)
Birdcalls71	Birdcalls71	0.91	Anuran	97
Anuran	Anuran	0.95	Birdcalls91	93