

## Article

# Deep Learning for Audio Event Detection and Tagging on Low-Resource Datasets

Veronica Morfi \*  and Dan Stowell 

Machine Listening Lab, Centre for Digital Music (C4DM), Queen Mary University of London,  
London E1 4NS, UK; dan.stowell@qmul.ac.uk

\* Correspondence: g.v.morfi@qmul.ac.uk; Tel.: +44-751-923-2154

Received: 15 June 2018; Accepted: 14 August 2018; Published: 18 August 2018



**Abstract:** In training a deep learning system to perform audio transcription, two practical problems may arise. Firstly, most datasets are weakly labelled, having only a list of events present in each recording without any temporal information for training. Secondly, deep neural networks need a very large amount of labelled training data to achieve good quality performance, yet in practice it is difficult to collect enough samples for most classes of interest. In this paper, we propose factorising the final task of audio transcription into multiple intermediate tasks in order to improve the training performance when dealing with this kind of low-resource datasets. We evaluate three data-efficient approaches of training a stacked convolutional and recurrent neural network for the intermediate tasks. Our results show that different methods of training have different advantages and disadvantages.

**Keywords:** deep learning; multi-task learning; audio event detection; audio tagging; weak learning; low-resource data

## 1. Introduction

Machine learning has experienced a strong growth in recent years, due to increased dataset sizes and computational power, and to advances in deep learning methods that can learn to make predictions in extremely nonlinear problem settings [1]. However, a large amount of data is needed in order to train a neural network that can achieve a good quality performance. With the increased amount of audio datasets publicly available, there is also an increase of tagging labels available for them. We refer to these tagging labels, which only indicate the presence or not of a type of event in a recording and lack any temporal information about it, as weak labels.

A lot of research has been done in tagging of audio recordings. In [2], the authors proposed a content-based automatic music tagging algorithm using deep convolutional neural networks. In [3], the authors proposed to use a shrinking deep neural network incorporating unsupervised feature learning to handle the multi-label audio tagging. Furthermore, considering that only chunk level rather than frame-level labels are available, a large set of contextual frames of the chunk were fed into the network to perform this task. In [4,5], the authors use a stacked convolutional recurrent network to perform environmental audio tagging and tag the presence of birdsong, respectively. However, in [6], the authors explore two different models for end-to-end music audio tagging when there is a large amount of training data.

However, in recent decades, there has also been an increase in the demand for transcription predictions for a variety of audio recordings instead of just the tags of a recording. Transcription of audio recordings refers to audio event detection, which provides a list of audio events active in a recording along with temporal information about each of them, i.e., starting time and duration for each event [7–10]. Some potential applications where audio event transcription is necessary are

context awareness for cars, smartphones, etc., surveillance for dangerous events and crimes, analysis and monitoring of biodiversity, recognition of noise sources and machine faults and many more. Depending on the audio event to be detected and classified in each task, it may become difficult to collect enough samples for them. Furthermore, different tasks use task specific datasets, hence the amount of recordings available may be limited. Annotating data with strong labels, labels that contain temporal information about the events, to train transcription predictors is a time-consuming process involving a lot of manual labour. On the other hand, collecting weakly labelled data takes much less time, since the annotator only has to mark the active sound event classes and not their exact boundaries. We refer to datasets that only have these types of weak labels, may contain rare events and have limited amounts of training data as low-resource datasets.

In comparison to supervised techniques that are trained on strong labels, there has been relatively little work on learning to perform audio event transcription using weakly labelled data. In [11,12], the authors try to exploit weak labels in birdsong detection and bird species classification, while, in [13], the authors use deep networks to tag the location of bird vocalisations. In [14], singing voice is pinpointed from weakly labelled examples. In [15], the authors used a joint detection-classification network that slices the audio into blocks and an audio detector and classification on each block then uses the overall audio tag to train using the weak labels of a recording. In [10], the authors train a network that can do automatic scene transcription from weak labels and, in [16], audio from YouTube videos is used in order to train and compare different previously proposed convolutional neural network architectures for audio event detection and classification. Finally, in [17,18], the authors use weakly labelled data for audio event detection in order to move from the weak labels space to strong labels. Most of these methods formulate the provided weak labels of the recordings into a multi instance learning (MIL) problem. However, for the methods using neural networks, none of the datasets used could be considered low-resource. Most of the datasets used either come from transcription/detection challenges (e.g., the IEEE AASP Challenge on Detection and Classification of Acoustic Scenes and Events (DCASE)) or online sources, such as Youtube or Xeno-Canto, which contain a large number of training data.

Training a neural network to predict an audio transcription using a low-resource dataset can sometimes prove to be impossible. A network needs to have enough parameters to be able to predict all the different classes without ignoring any rare events, but also be small enough or have just the right amount of regularisation as to not overfit the limited amount of training data available. This becomes even harder when the task is a weak-to-strong prediction where the network needs to predict full transcriptions from weak labels. Unfortunately, there is no specific way of defining a network and type of training that ensures that a transcription will be predicted successfully. However, a full transcription task can be defined as multiple intermediate tasks of detection and classification that might be easier to train even when using a low-resource dataset. A similar approach is used to enhance the performance of automatic speech transcription [19] by using speaker diarisation [20,21] and speaker recognition [22] systems together in order to structure an audio stream into speaker turns and provide the speaker's true identity, respectively. However, these speech approaches are highly customised to characteristics of speech signals. Our method is focused on general audio with speech events considered just a single class amongst other audio events without distinguishing between individual speakers.

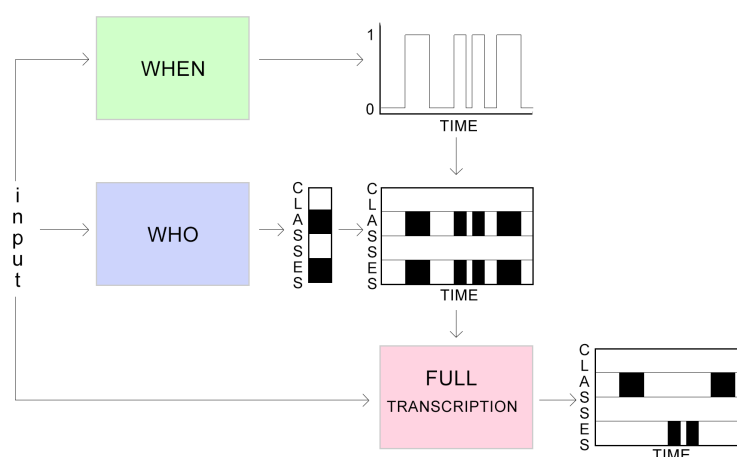
In this paper, we propose a factorisation of the final full transcription task into multiple simpler intermediate tasks of audio event detection and audio tagging in order to predict an intermediate transcription that can be used to boost the performance of the full transcription task. For each intermediate task, we propose a training setup to optimise their performance. Finally, we train the intermediate tasks independently and in two multi-task learning settings and compare their results.

The rest of the paper is structured as follows: Section 2 describes the way we factorise the transcription task into intermediate tasks and presents in detail our setup and network architectures. In Section 3, we propose three different training approaches for the intermediate tasks, two of which are implemented in a multi-task learning setting. In Section 4, we present our experiments and

compare the results of each training approach. Finally, in Section 5, we discuss our findings and future research directions.

## 2. Task Factorisation

A full audio transcription task can be described as audio event detection followed by event classification. In order to properly train a full transcription network, we need a large amount of data that is not available in a low-resource dataset. Since it is very hard to train a network to predict full transcription on a low-resource dataset, we factorise the final task of full transcription into intermediate tasks that can predict an intermediate transcription matrix that can later be used to boost the performance of a full transcription network. Figure 1 depicts the overall task factorisation into the intermediate tasks and how they interact with the final task of full transcription. We define a WHEN network that performs audio event detection considering all classes as one general class; in other words, it predicts when any event is present without taking into consideration the different event classes. We also define a WHO network that performs audio tagging without predicting any temporal information. By combining the two different predictions from these networks, we create an intermediate transcription that provides us with the events present in a recording and the times where any of these events could be present in a recording. This intermediate transcription is to be used as supplementary information when training the full transcription network in order to improve its performance by focusing its attention to the classes present in a recording and the time frames that may contain them.



**Figure 1.** Factorisation of the full transcription task. The WHEN network performs audio event detection considering all labels as one label. The WHO network performs audio tagging for all available labels. The predictions of WHEN and WHO produce an intermediate transcription that is used to boost the performance of the full transcription network.

When using a large enough dataset that provides satisfactory training data and has a good representation for each different class, many methods have been successful in performing both of the intermediate tasks. A few methods for audio event detection can be found in [13,14], while for audio tagging in [2–6,15]. These tasks are less challenging to train for than a full transcription task. However, using a low-resource dataset can degrade their performance. Hence, in order to achieve a satisfactory performance when training with a low-resource dataset, we propose a few training setups and techniques. The rest of this section describes in detail the task specific setups and techniques that we used.

### 2.1. Input Features

As input to all our intermediate networks, log mel-band energy is extracted from audio in 23 ms Hamming windows with 50% overlap. In order to do so, the librosa Python library is used. In total,

40 mel-bands are used in the 0–22,050 Hz range. For a given 5 s audio input, the feature extraction produces a  $T \times 40$  output ( $T = 432$ ).

## 2.2. Audio Event Detection (WHEN)

In our proposed task factorisation, the WHEN network performs a single class audio event detection as the first intermediate task towards full transcription. For a multi-class dataset, one would have to train a separate network for each class in order to perform single class event detection. However, in a low-resource dataset, training an audio event detector for each class can be nearly impossible. The number of classes might be too large, making it a time-consuming task. Furthermore, some of the classes might have very rare occurrences, limited to only a couple of recordings, hence making it infeasible to train a neural network for them. Nevertheless, many low-resource datasets are usually used for discriminating subclasses of a general class e.g., song of different bird species, sound of different car engines, barking of different dog breeds, and notes produced by an instrument. These subclasses usually share some common features and characteristics, hence, in order to achieve a good performance in the audio event detection task, we propose considering all subclasses as one general class and train a single WHEN network to perform single class event detection. This reduces the training time compared to training one network for each subclass and also solves any training issues caused by rare events.

### 2.2.1. Neural Network Architecture

For our audio event detector, we use a state-of-the-art stacked convolutional and recurrent neural network architecture. Table 1 describes the parameters of the proposed architecture.

**Table 1.** WHEN network architecture. Size refers to either kernel shape or number of units. #Fmaps is the number of feature maps in the layer. Activation denotes the activation used for the layer and l2\_regularisation the amount of l2 kernel regularisation used in the layer.

Layer	Size	#Fmaps	Activation	l2_regularisation
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Max Pooling	$1 \times 5$	-	-	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Max Pooling	$1 \times 4$	-	-	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Max Pooling	$1 \times 2$	-	-	-
Reshape	-	-	-	-
Bidirectional GRU	64	-	tanh	0.01
Bidirectional GRU	64	-	tanh	0.01
Time Distributed Dense	64	-	ReLU	0.01
Time Distributed Dense	1	-	Sigmoid	0.01
Flatten	-	-	-	-
Trainable parameters:	320,623			

The log mel-band energy feature extracted from the audio is fed to the neural network, which sequentially produces the predicted strong labels for each recording. The input to the proposed network is a  $T \times 40$  feature matrix. The convolutional layers in the beginning of the network are in charge of learning the local shift-invariant features of this input. We use a  $3 \times 3$  receptive field and the padding arguments set as 'same' in order to maintain the same size as the input in all our convolutional layers. The max-pooling operation is performed along the frequency axis after every convolutional layer to reduce the dimension for the feature matrix while preserving the number of frames  $T$ . The output of the convolutional part of the network is then fed to bi-directional gated recurrent units (GRUs) with tanh activation to learn the temporal structure of audio events. Next, we apply time distributed dense layers to reduce feature-length dimensionality. Note that the time resolution of  $T$  frames is maintained in both the GRU and dense layers. A sigmoid activation is used in the last time-distributed dense layer to produce a binary prediction of whether there is an event present in each time frame. This prediction layer outputs the strong labels for a recording. The dimensions of each prediction are  $T \times 1$ . Finally, we calculate the loss on this output as explained in Section 2.2.2.

### 2.2.2. Multi Instance Learning

When used for training audio event detectors, low-resource datasets present the issue of weak-to-strong prediction. Low-resource datasets only provide the user with weak labels, labels that don't include any temporal information about the events but only denote the presence or absence of a specific class in a recording. However, audio event detectors produce instance labels referred to as strong labels, hence providing full temporal information about the events in a recording.

The most common way to train a network for weak-to-strong prediction is the multi instance learning (MIL) setting. The concept of MIL was first properly developed in [23] for drug activity detection. MIL is described in terms of *bags*, with a bag being a collection of instances. The existing weak labels are attached to the bags, rather than the individual instances within them. Positive bags have at least one positive instance, an instance for which the target class is active. On the other hand, negative bags contain negative instances only, the target class is not active in them. A negative bag is thus pure while a positive bag is presumably impure, since the latter most likely contains both positive and negative instances. Hence, all instances in a negative bag can be uniquely assigned a negative label but for a positive bag this cannot be done. There is no direct knowledge of whether an instance in a positive bag is positive or negative. Thus, it is the bag-label pairs and not the instance-label pairs which form the training data, and from which a classifier that classifies individual instances must be learned.

Let the training data be composed of  $N$  bags, i.e.,  $\{B_1, B_2, \dots, B_N\}$ , the  $i$ -th bag is composed of  $M_i$  instances, i.e.,  $\{B_{i1}, B_{i2}, \dots, B_{iM_i}\}$ , where each instance is a  $p$ -dimensional feature vector, e.g., the  $j$ -th instance of the  $i$ -th bag is  $[B_{ij1}, B_{ij2}, \dots, B_{ijp}]^T$ . We represent the bag-label pairs as  $(B_i, Y_i)$ , where  $Y_i \in \{0, 1\}$  is the bag label for bag  $B_i$ .  $Y_i = 0$  denotes a negative bag and  $Y_i = 1$  denotes a positive bag.

One naïve but commonly used way of inferring the individual instances' labels from the bag labels is assigning the bag label to each instance of that bag: we refer to this method as *false strong labelling*. During training, a neural network in the MIL setting with false strong labels tries to minimise the average divergence between the network output for each instance and the false strong labels assigned to them, identically to an ordinary supervised learning scenario. However, it is evident that the false strong labelling approach is an approximation of the loss for a strong label prediction task, hence it has some disadvantages. When using false strong labels, some kind of early stopping is necessary since, when perfect accuracy is achieved, this would mean all positive instance predictions for a positive bag. However, there is no clear way of defining a specific point for early stopping. This is the same issue that all methods in the MIL setting face. As mentioned before, a positive bag might include both positive and negative instances; however, false strong labels will force the network towards positive predictions for both. Additionally, by using strong false labels, there is an imbalance of positive and negative instance labels compared to the true strong labels, since a substantial amount of negative

instances are considered as positive during training. Finally, a negative instance may appear in both a negative and positive recording; however, due to the false labelling of negative instances as positive in positive bags, the network may not learn the proper prediction for this kind of instance.

As an alternative to false strong labels, one can attempt to infer labels of individual instances in bag  $B_i$  by making a few educated assumptions. The most common ones are: if  $Y_i = 0$ , all instances of bag  $B_i$  are negative instances, hence  $y_{ij} = 0, \forall j$ , while, on the other hand, if  $Y_i = 1$ , at least one instance of bag  $B_i$  is equal to one. For all instances of bag  $B_i$ , this relation between the bag label and instance labels can be simply written as:

$$Y_i = \max_j y_{ij}. \quad (1)$$

The conventional way of training a neural network for strong labelling is providing instance specific (strong) labels for a collection of training instances. Training is performed by updating the network weights to minimize the average divergence between the network output in response to these instances and the desired output, the ground truth of the training instances. In the MIL setting using Equation (1) to define a characteristic of the strong labels, we must modify the manner in which the divergence to be minimized is computed, in order to utilize only weak labels, as proposed in [24].

Let  $o_{ij}$  represent the output of the network for input  $B_{ij}$ , the  $j$ -th instance in  $B_i$ , and the  $i$ -th bag of training instances. We define the bag-level divergence for bag  $B_i$  as:

$$E_i = \frac{1}{2} \left( \max_{1 \leq j \leq M_j} (o_{ij}) - Y_i \right)^2, \quad (2)$$

where  $Y_i$  is the label assigned to bag  $B_i$ .

The overall divergence on the training set is obtained by summing the divergences of all the bags in the set:

$$E = \sum_{i=1}^N E_i. \quad (3)$$

Equation (2) indicates that if at least one instance of a positive bag is perfectly predicted as positive, or all the instances of a negative bag are perfectly predicted as negative, then the error on the concerned bag is zero. Otherwise, the weights will be updated according to the error on the instance whose corresponding actual output is the maximal among all the instances in the bag. Note that such an instance is typically the most easy to be predicted as positive for a positive bag, while it is the most difficult to be predicted as negative for a negative bag. It seems that this sets a low burden on producing a positive output but a strong burden on producing a negative output. As indicated in [25], the value of a bag is fully determined by its instance with the maximal output, regardless of how many real positive or negative instances are in the bag. Therefore, in fact, the burden on producing a positive or negative output is not unbalanced, at least at bag-level. However, on an instance-level, when using max to compute the loss, only one instance per bag contributes to the gradient, which may lead to inefficient training. Additionally, as mentioned earlier, in positive bags, the network only has to accurately predict the label for the easiest positive instance to reach a perfect accuracy, thus not paying as much attention to the rest of the positive instances that might be harder to accurately detect.

In order to train our proposed WHEN network, we want all predictions to weigh in on the loss and not just the one with the maximum value, as is the case with MIL using max. In [26], the authors proposed the “noisy-or” pooling function to be used instead of max. However, noisy-or has been proven to not perform as well as max for audio event detection [27]. As discussed in [27], a significant problem with noisy-or is that the label of a bag is computed via the product of the instance labels as seen in Equation (4). This calculation relies heavily on the assumed conditional independence of instance labels, an assumption which is highly inaccurate in audio event detection. Furthermore, this can lead the system to believe a bag is positive even though all its instances are negative:



$$Y_i = 1 - \prod_{1 \leq j \leq M_j} (1 - y_{ij}). \quad (4)$$

Using all instances in a bag for computation of the loss and backpropagated gradient is important, since the network ideally should acquire some knowledge from every instance in every epoch. However, it is hard to find an elegant theoretical interpretation of the characteristics of the instances in a bag. On the other hand, we propose a couple simpler assumptions about these characteristics that can achieve a similar effect. One assumption is to consider the mean of the instance predictions of a bag. If a bag is negative, the mean should be zero, while if it is positive it should be greater than zero. The true mean is unknown in weakly labelled data. A naïve assumption is to presume that approximately half of the time a specific event will be present in a recording. Even though this is not true all of the time, it takes into consideration the predictions for all instances, and also inserts a bias to the loss that will keep producing gradient for training even after the max term has reached its perfect accuracy. However, this is indeed a naïve assumption that will guide the network to predict a balanced amount of positives and negatives that may make it more sensitive to all kinds of audio events, even when they are not the ones in question.

Another simple yet accurate assumption is that on both negative and positive recordings the minimum predictions at an instance-level should be zero. It is possible for a positive recording to have no negative frames; however, it is extremely rare in practice. This assumption could be used in synergy with max and mean to enforce the prediction of negative instances even on positive recordings and manage a certain level of the bias that is introduced with considering mean in the computation of the loss.

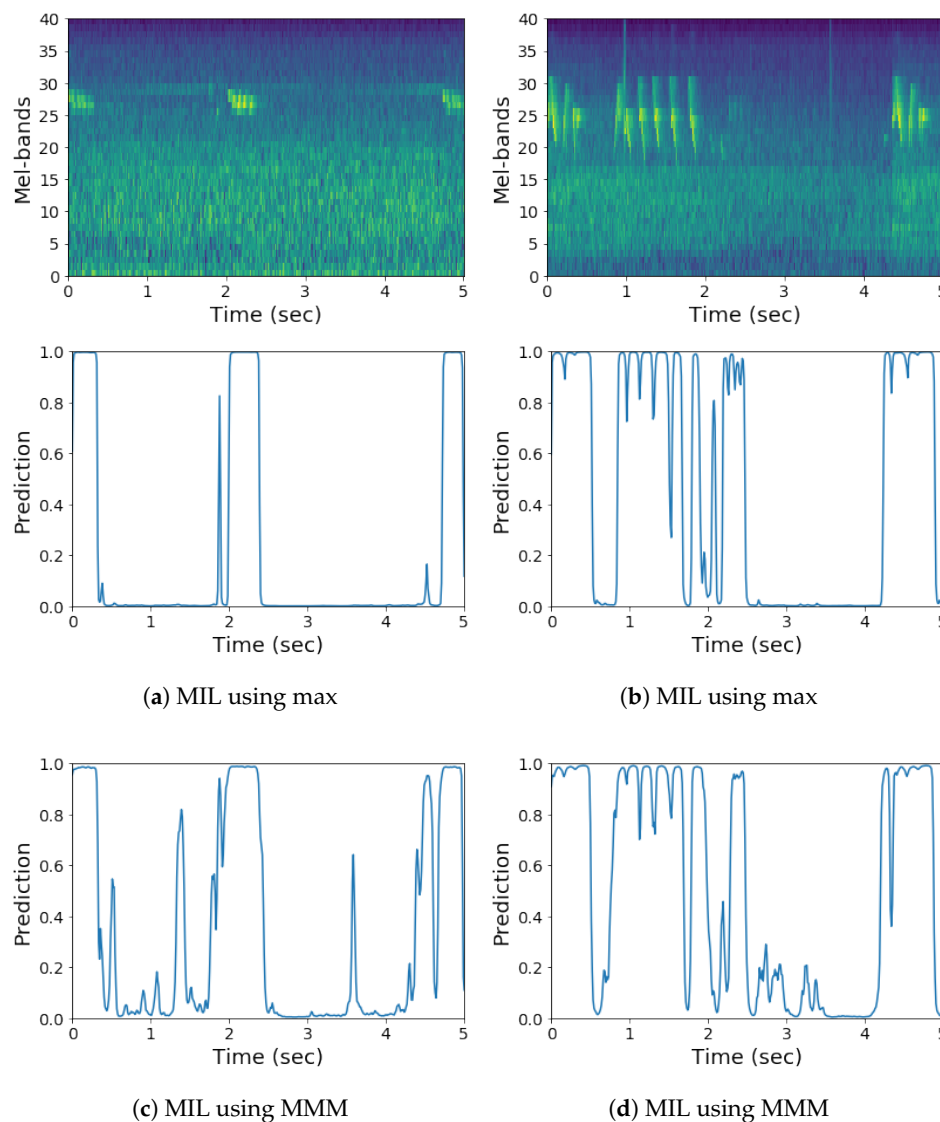
We train a network on a loss function that takes into account all the above-mentioned assumptions and compute the max, mean and min from the predictions of a recording and depending on whether a recording is positive or negative we predict their divergence from different conditions.

Our proposed loss function is computed as:

$$Loss = \frac{1}{3} (bin\_cr(max_j(o_{ij}), Y_i) + bin\_cr(mean_j(o_{ij}), \frac{Y_i}{2}) + bin\_cr(min_j(o_{ij}), 0)), \quad (5)$$

where  $bin\_cr(x, y)$  is a function that computes the binary cross-entropy between  $x$  and  $y$ ,  $o_{ij}$  are all the predicted strong labels of bag  $B_i$ , where  $j = 1 \dots M_i$  with  $M_i$  being the total number of instances in a bag, and  $Y_i$  is the label of the bag.

We refer to this as an MIL setting using max, mean and min (MMM). For negative recordings, Equation (5) will compute the binary cross-entropy between the max, mean and min of the predictions of the instances of a bag  $B_i$  and zero. This denotes that the predictions for all instances of a negative recording should be zero. On the other hand, for positive recordings, the predictions should span the full dynamic range from zero to one, biased towards a similar amount of positive and negative instances. Our proposed loss function is designed to balance the positive and negative predictions in a bag resulting in a network that has the flexibility of learning from harder-to-predict positive instances even after many epochs. This is due to the fact that there are no obvious local minima to get stuck in as in the max case. Some examples of the difference between the predictions produced by MIL using max and MIL using MMM when our proposed WHEN network is trained for birdsong detection are depicted in Figure 2. It becomes apparent that MIL using MMM can correctly classify harder to predict instances, especially when studying the difference between Figure 2a,c. In Figure 2c, one can notice that the network is able to correctly classify the harder to predict instances between the three main audio events.



**Figure 2.** Predicted transcription, of two recordings. Figure 2a,b depict the results of our WHEN network trained with max loss. Figure 2c,d depict the results of our WHEN network trained with MMM loss.

### 2.2.3. Half and Half Training

In the MIL setting for weak-to-strong labelling, it is of great importance to have a good balance between positive and negative bags, in order for the network to be able to distinguish what can be considered a positive instance and what can be considered a negative one. A simple approach to achieve this kind of balanced training is to have balanced minibatches. In our approach, we implement this by duplicating negative or positive recordings randomly during training depending on which ones are less in the whole dataset. Thus, each minibatch during training will consist of the same amount of positive and negative recordings, which in our case is four positive and four negative recordings. We call this kind of input Half and Half (HnH). Please note that balanced data for the WHEN task is not necessarily balanced data for the WHO task, an issue that we will return to.



### 2.3. Audio Tagging (WHO)

The second intermediate task of our approach is the WHO network that performs audio tagging using the provided weak labels of a low-resource dataset. This task follows supervised training since the weak labels provided are exactly the ones that the network will try to learn how to predict. Hence, our training techniques that we use for the WHO network follow standard approaches.

#### Neural Network Architecture

A similar network architecture to the one proposed for WHEN (see Table 1) is used for the first few layers of WHO in order to implement our proposed training approaches that we introduce in Section 3. Table 2 describes the structure of each individual layer used in the WHO network.

**Table 2.** WHO network architecture. Size refers to either kernel shape or number of units. #Fmaps is the number of feature maps in the layer. Activation denotes the activation used for the layer and l2\_regularisation the amount of l2 kernel regularisation used in the layer.

Layer	Size	#Fmaps	Activation	l2_regularisation
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Max Pooling	$1 \times 5$	-	-	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Max Pooling	$1 \times 4$	-	-	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Convolution 2D	$3 \times 3$	64	Linear	0.001
Batch Normalisation	-	-	-	-
Activation	-	-	ReLU	-
Max Pooling	$1 \times 2$	-	-	-
Global Average Pooling 2D	-	-	-	-
Dense	#labels	-	Sigmoid	0.001
Trainable parameters:	191,319			

Similar to the WHEN network, the log mel-band energy feature extracted from the audio is used as input with shape  $T \times 40$ , where  $T$  is the number of time frames in a recording. The convolutional layers in the beginning of the network are in charge of learning the local shift-invariant features of this input. We use a  $3 \times 3$  receptive field and the padding arguments set as 'same'. Max-pooling is performed along the frequency axis after every convolutional layer to reduce the dimension for the feature matrix. Global average pooling is finally applied to the output of the convolutional part of the network and the results are fed to a dense layer that has units equal to the number of labels for our tagging task with sigmoid activation that predict the probability of each class being present in a recording. The dimensions of each prediction are  $1 \times \text{\#labels}$ . Finally, we calculate the binary cross-entropy loss on this output and the ground truth extracted from the weak labels.

### 3. Training Methods

We investigate three different methods to train the two intermediate tasks. One of them is the simple and usual approach of training each network independently for each task. Additionally, two multi-task learning (MTL) methods were tested, namely joint training and tied weights training,


both of which follow a hard parameter sharing approach. All three different methods have advantages and disadvantages that will be compared in detail in Section 4.

MTL [28] aims to improve the performance of multiple learning tasks by sharing useful information among them. MTL can be very useful when using low-resource datasets since it can exploit useful information from other related learning tasks to help alleviate the issue of limited data. Based on the assumption that the multiple tasks are related, MTL is empirically and theoretically found to lead to better performance than independent learning. MTL is similar to transfer learning [29] which also transfers knowledge from one task to another. However, the focus of transfer learning is to help a single target task by initially training on one or multiple tasks while MTL uses multiple tasks to help each other. Furthermore, MTL can be viewed as a generalization of multi-label learning [30] when different tasks in multi-task learning share the same training data.

The motivation behind using MTL includes the implicit data augmentation, since a model that learns two tasks simultaneously is able to learn a more general representation. In addition, if data is limited, MTL can help the model focus its attention on those features that actually matter as other tasks will provide additional evidence for the relevance or irrelevance of those features. Finally, MTL acts as a regulariser by introducing an inductive bias that reduces the risk of overfitting. An overview of MTL can be found in [31].

### 3.1. Separate Training

First, we used separate training for the two tasks. As depicted in Figure 3, two independent networks are defined, namely WHEN and WHO with the architectures described in Sections 2.2 and 2.3, respectively. The WHEN network performs audio event detection considering all labels as a single general label, while the WHO network performs audio tagging. Different kinds of input can be used for each network. HnH input was used for WHEN and the normal (nonHnH) input for WHO. Thus, the minibatches used as input for the WHO network are randomly generated without taking into account the balance of positive and negative recordings in them. Different types of input were used for each task since they perform differently with different types of input even though the source of training data for each one is the same.



**Figure 3.** Separate training. Networks WHEN and WHO are defined and trained independently of one another, with different types of inputs.

The advantage of separate training is that each network can train with the type of input that works better for it. WHEN uses a balanced minibatch of positive and negative recordings (HnH) while WHO uses the conventional random type of minibatch (nonHnH). The main disadvantage of separate training is that each task trains independently of the other, which may mean wasted computation. However, these two tasks are somewhat related, hence they should be able to focus the attention of the network to important features and also regularise each other.

### 3.2. Joint Training

Joint training is one of the most common MTL approaches. In joint training, the same network is trained for more than one task. Usually, the network consists of a few shared layers in the beginning followed by task specific layers before the predictions for each task. For each task, a separate loss is computed and then combined into the general loss of the network, usually by weighting each loss. Joint training is a hard parameter sharing approach, since all tasks share the same initial layers and weights. Figure 4 depicts how our intermediate tasks are adapted to the joint training approach. The *Shared Convolutional Part* consists of the common convolutional and max pooling layers while the separate branches of the network consist of the task specific layers for WHEN and WHO as described in Tables 1 and 2, respectively.




Figure 4. Joint training.

The advantages of joint training are all the advantages presented by MTL. More specifically, information is shared between the tasks to help alleviate the issue of limited data. The model focuses its attention on features that are more relevant to all tasks. In addition, it reduces the risk of overfitting, since one task can act as the other's regulariser. One of the disadvantages of joint training is that both tasks train on the same input, which, depending on the type (HnH or nonHnH), degrades the performance of one of the tasks (WHO or WHEN, respectively), as we will show in Section 4.

### 3.3. Tied Weights Training

In order to achieve the advantages of both separate and joint training without any of their disadvantages, we propose a new approach of MTL. Tied weights training follows the hard parameter training convention, where layers and their weights are shared between tasks. However, in contrast to joint training, different types of input can be used to train each task. Figure 5 depicts the structure of tied weights training. *Shared Convolutional Part* refers to the common convolutional and max pooling layers of WHEN and WHO, and the weights between the two tasks are constrained to be identical in these layers. Each network is trained consecutively for one epoch, updating the weights of the shared layers. Using this approach, one can train each network with independent types of input as in separate training while keeping all the advantages of MTL learning.




Figure 5. Tied weights training.

#### 4. Evaluation

In order to test our approach in a low-resource dataset, we use the training dataset provided during the Neural Information Processing Scaled for Bioacoustics (NIPS4B) bird song competition of 2013 that is publicly available and contains 687 recordings of maximum length of 5 s each [32]. For the NIPS4B dataset the recordings have already been weakly labelled and the labels are provided by the organisers along with the dataset recordings. The dataset contains a total of 87 classes, with each being active in only 7 to 20 recordings. Each recording has 0 to 6 classes active in it. Such a dataset can be considered low-resource since the total amount of training time is less than one hour and also there are 87 possible labels that have very sparse activations, 7 to 20 positive recordings for each.

For our experiments, we split the NIPS4B 2013 training dataset into a training set and testing set. During the NIPS4B 2013 bird song competition, only the weak labels for the training dataset were released, hence we could only use these recordings and couldn't make any use of the NIPS4B 2013 testing dataset that consisted of more recordings. We enlisted an experienced bird watcher to manually annotate strong labels for most of the training dataset recordings [33]. For our training set, the first 499 recordings of the NIPS4B 2013 training dataset are used, while the rest are included in our testing set, excluding 14 recordings for which confident strong annotations could not be attained. Those 14 recordings were added to our training set for a total of 513 training recordings and 174 testing recordings.

In order to efficiently use the data provided by the NIPS4B 2013 training dataset for our WHEN task, we first consider all 87 unique labels as one general label 'bird' and train an audio event detection network for this class. Another limitation of this dataset is the imbalance of positive and negative recordings: out of the whole dataset (687 recordings), only 100 of them are labelled as negative (not having any bird present in them). We provide a balanced training set by using our Half and Half training approach. For this dataset, the training set consists of 450 recordings in total (385 positive, 65 negative), each recording being 432 time frames long, totalling less than 40 min of total training audio. During training, HnH will randomly duplicate the negative recordings in order to balance their amount with the positive recordings, hence creating a training set of 770 recordings half of which are unique positive recordings and the other half are randomly duplicated negative recordings.

#### Results

The same parameters are used for training both WHEN and WHO networks for all three different approaches. Our batchsize is equal to eight recordings. We use the Adam optimiser [34] with a learning rate scheduler that reduces the initial rate of  $1 \times 10^{-5}$  by half every 20 epochs until it reaches a minimum rate of  $1 \times 10^{-8}$ . The loss function used for the predictions of the WHEN network is the proposed MMM loss, while we use a binary cross-entropy loss for the multi-class predictions of the WHO network.

First, we trained WHEN and WHO independently. WHEN was trained with a HnH input, since not using HnH can cause the network to either ignore negative recordings or mix the negative and positive frames in a recording. On the other hand, WHO was trained with the conventional nonHnH input since using HnH for WHO made its performance worse. This is due to the fact that the active classes are already very sparse (0 to 6 active classes out of 87 per recording) and, for the NIPS4B dataset, the HnH input duplicates negative recordings, hence decreasing the activation rate for each class, making it even harder to predict.

Next, we trained two versions of the joint network, one of them uses a HnH input while the other a nonHnH input. When training the joint network with HnH, the WHO predictions tend to not have a satisfactory performance due to the increase in negative recordings. When training the joint network with the nonHnH input, the WHEN task performance is degraded. The loss value of the WHO task tends to be an order smaller than the one for WHEN, hence we trained with two different combination of weights for the task. For one of them, both task losses have the same weight of 0.5, while for the

other one the weight for the WHO task loss is an order larger than the WHEN; more specifically, we used weight 0.5 for WHEN loss and 5.0 for WHO loss.

Finally, we performed a tied weights training. This solved the issue of using only one type of input since it can train with both HnH and nonHnH input separately for each task as if the tasks are trained independently, while still sharing the weights of the shared layers like the joint training.

During the NIPS4B 2013 challenge, systems that performed audio tagging, similar to the task for our WHO network, were submitted. The winning solution [35] was trained on the whole training set (687 recordings) and tested on the NIPS4B testing set and reached 0.92 area under the receiver operating characteristic (ROC) curve (AUC) score. These results can be used as a performance baseline for this dataset. However, the winning NIPS4B method and our method cannot be explicitly compared since our network is trained in a subset of the NIPS4B training set and then evaluated in the rest of the recordings.

Table 3 shows the area under the ROC curve (AUC) results for each training approach. We can see that even though the tied weights training has a better overall performance compared to the joint training, separate training still has the best overall results. The best overall results for joint training were produced when using weights 0.5 and 5.0 for WHEN and WHO loss, respectively and also using nonHnH input. Hence, we can conclude that the WHO network is sharing important information with the WHEN network that can boost its performance when enough weight is given to its loss. As mentioned before, any type of joint training has so far been proven to outperform independent training, which is not the case in our experiments, when comparing results for both WHEN and WHO. We consider the two tasks to be closely related and use hard parameter sharing approaches. However, the tasks might be more loosely related than we originally considered and a soft parameter sharing approach [36–39] may increase performance.

**Table 3.** Area under the ROC curve (AUC) for the predictions of all training approaches. [WHEN: xx; WHO: yy] indicate the weights xx for WHEN task loss and yy for WHO task loss that were used during joint training. Best values are marked in bold.

Training	Input Type	WHEN	WHO
Method	WHEN   WHO	AUC	AUC
Separate	HnH   nonHnH	<b>0.90</b>	<b>0.94</b>
Joint [WHEN: 0.5; WHO: 0.5]	HnH	0.89	0.52
Joint [WHEN: 0.5; WHO: 0.5]	nonHnH	0.47	0.57
Joint [WHEN: 0.5; WHO: 5.0]	HnH	<b>0.90</b>	0.50
Joint [WHEN: 0.5; WHO: 5.0]	nonHnH	0.82	0.75
Tied Weights	HnH   nonHnH	0.87	0.77

## 5. Conclusions

In this paper, we present a way to factorise the task of full transcription into multiple intermediate tasks in order to improve performance for low-resource datasets. We propose two intermediate tasks of audio event detection on a single class and audio tagging, referred to as WHEN and WHO tasks, respectively. Additionally, we introduce a balanced input training and a new loss function in the multi instance learning (MIL) setting for the WHEN task. We train these tasks with three different approaches. Firstly, an independent training for each task and then two multi-task learning (MTL) approaches that use hard parameter sharing. One of them is the most commonly used joint training and the other one is our proposed tied weights training. In order to evaluate our approaches, we trained each network using a low-resource dataset for birdsong transcription. Our results show that, even though our proposed tied weights training outperforms joint training for these tasks, separate training still performs better than both.

For our future plans, we first intend to explore whether soft parameter sharing in MTL can further improve the performance of our intermediate tasks. Then, we plan to use the intermediate transcription to boost the performance of a full transcription network. To our current knowledge and

based on our latest experiments, attempting to train a network to perform full transcription without any intermediate tasks for this low-resource dataset is not feasible. Hence, we will attempt to achieve a satisfactory performance when using the intermediate transcription as a guide for the attention of the full transcription network.

**Author Contributions:** V.M. and D.S. conceived and designed the experiments; V.M. performed the experiments; V.M. and D.S. analysed the data; V.M. wrote the paper.

**Funding:** This research was funded by Engineering and Physical Sciences Research Council (EPSRC) Grant No. EP/L020505/1.

**Acknowledgments:** We would like to acknowledge the contributions of Hanna Pamula (AGH University of Science and Technology in Kraków, Poland) for acquiring the transcriptions of the NIPS4B 2013 training set by manual annotations.

**Conflicts of Interest:** The authors declare no conflict of interest.

## Abbreviations

The following abbreviations are used in this manuscript:

MIL	Multi instance learning
GRUs	Gated recurrent units
ReLU	Rectified linear unit
MMM	Max mean min
HnH	Half and half
nonHnH	Non half and half
MTL	Multi-task learning
NIPS4B	Neural Information Processing Scaled for Bioacoustics
ROC	Receiver operating characteristic
AUC	Area under the curve

## References

1. Lecun, Y.; Bengio, Y.; Hinton, G. Deep Learning. *Nature* **2015**, *521*, 436–444. [[CrossRef](#)] [[PubMed](#)]
2. Choi, K.; Fazekas, G.; Sandler, M.B. Automatic Tagging Using Deep Convolutional Neural Networks. In Proceedings of the 17th International Society for Music Information Retrieval Conference (ISMIR), New York, NY, USA, 7–11 August 2016; pp. 805–811.
3. Xu, Y.; Huang, Q.; Wang, W.; Foster, P.; Sigtia, S.; Jackson, P.J.B.; Plumbley, M.D. Unsupervised Feature Learning Based on Deep Models for Environmental Audio Tagging. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2017**, *25*, 1230–1241. [[CrossRef](#)]
4. Xu, Y.; Kong, Q.; Huang, Q.; Wang, W.; Plumbley, M.D. Convolutional gated recurrent neural network incorporating spatial features for audio tagging. In Proceedings of the 2017 International Joint Conference on Neural Networks (IJCNN), Anchorage, AK, USA, 14–19 May 2017; pp. 3461–3466.
5. Adavanne, S.; Drossos, K.; Çakir, E.; Virtanen, T. Stacked convolutional and recurrent neural networks for bird audio detection. In Proceedings of the 2017 25th European Signal Processing Conference (EUSIPCO), Kos, Greece, 28 August–2 September 2017; pp. 1729–1733. [[CrossRef](#)]
6. Pons, J.; Nieto, O.; Prockup, M.; Schmidt, E.M.; Ehmann, A.F.; Serra, X. End-to-end learning for music audio tagging at scale. Presented at the Workshop Machine Learning for Audio Signal Processing at NIPS (ML4Audio), Long Beach, CA, USA, 4–9 December 2017.
7. Cakir, E.; Heittola, T.; Huttunen, H.; Virtanen, T. Polyphonic sound event detection using multi label deep neural networks. In Proceedings of the 2015 International Joint Conference on Neural Networks (IJCNN), Killarney, Ireland, 12–17 July 2015; pp. 1–7. [[CrossRef](#)]
8. Parascandolo, G.; Huttunen, H.; Virtanen, T. Recurrent neural networks for polyphonic sound event detection in real life recordings. In Proceedings of the 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Shanghai, China, 20–25 March 2016; pp. 6440–6444. [[CrossRef](#)]
9. Lee, D.; Lee, S.; Han, Y.; Lee, K. Ensemble of Convolutional Neural Networks for Weakly-supervised Sound Event Detection Using Multiple Scale Input. In Proceedings of the Detection and Classification of Acoustic Scenes and Events 2017 Workshop (DCASE 2017), Munich, Germany, 16–17 November 2017; pp. 74–79.



10. Adavanne, S.; Virtanen, T. Sound Event Detection Using Weakly Labeled Dataset with Stacked Convolutional and Recurrent Neural Network. In Proceedings of the Detection and Classification of Acoustic Scenes and Events 2017 Workshop (DCASE 2017), Munich, Germany, 16–17 November 2017; pp. 12–16.
11. Briggs, F.; Lakshminarayanan, B.; Neal, L.; Fern, X.; Raich, R.; Hadley, S.J.K.; Hadley, A.S.; Betts, M.G. Acoustic classification of multiple simultaneous bird species: A multi-instance multi-label approach. *J. Acoust. Soc. Am.* **2014**, *131*, 4640–4650. [[CrossRef](#)] [[PubMed](#)]
12. Ruiz-Muñoz, J.F.; Orozco-Alzate, M.; Castellanos-Dominguez, G. Multiple Instance Learning-based Birdsong Classification Using Unsupervised Recording Segmentation. In Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI'15), Buenos Aires, Argentina, 25–31 July 2015; pp. 2632–2638.
13. Fanioudakis, L.; Potamitis, I. Deep Networks tag the location of bird vocalisations on audio spectrograms. *arXiv* **2017**, arXiv:1711.04347.
14. Schlüter, J. Learning to Pinpoint Singing Voice from Weakly Labeled Examples. In Proceedings of the 17th International Society for Music Information Retrieval Conference (ISMIR 2016), New York, NY, USA, 7–11 August 2016.
15. Kong, Q.; Xu, Y.; Wang, W.; Plumbley, M.D. A joint detection-classification model for audio tagging of weakly labelled data. In Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), New Orleans, LA, USA, 5–9 March 2017; pp. 641–645. [[CrossRef](#)]
16. Hershey, S.; Chaudhuri, S.; Ellis, D.P.W.; Gemmeke, J.F.; Jansen, A.; Moore, R.C.; Plakal, M.; Platt, D.; Saurous, R.A.; Seybold, B.; et al. CNN architectures for large-scale audio classification. In Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), New Orleans, LA, USA, 5–9 March 2017; pp. 131–135. [[CrossRef](#)]
17. Kumar, A.; Raj, B. Audio Event Detection Using Weakly Labeled Data. In Proceedings of the 2016 ACM on Multimedia Conference (MM'16), Amsterdam, The Netherlands, 15–19 October 2016; ACM: New York, NY, USA, 2016; pp. 1038–1047. [[CrossRef](#)]
18. Kumar, A.; Raj, B. Deep CNN Framework for Audio Event Recognition using Weakly Labeled Web Data. *arXiv* **2017**, arXiv:1707.02530.
19. Yu, D.; Deng, L. *Automatic Speech Recognition*; Springer: Berlin, Germany, 2016.
20. Anguera, X.; Bozonnet, S.; Evans, N.; Fredouille, C.; Friedland, G.; Vinyals, O. Speaker Diarization: A Review of Recent Research. *IEEE Trans. Audio Speech Lang. Process.* **2012**, *20*, 356–370. [[CrossRef](#)]
21. Garcia-Romero, D.; Snyder, D.; Sell, G.; Povey, D.; McCree, A. Speaker diarization using deep neural network embeddings. In Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), New Orleans, LA, USA, 5–9 March 2017; pp. 4930–4934. [[CrossRef](#)]
22. Tirumala, S.S.; Shahamiri, S.R. A Review on Deep Learning Approaches in Speaker Identification. In Proceedings of the 8th International Conference on Signal Processing Systems (ICSPS 2016), Auckland, New Zealand, 21–24 November 2016; ACM: New York, NY, USA, 2016; pp. 142–147. [[CrossRef](#)]
23. Dietterich, T.G.; Lathrop, R.H.; Lozano-Pérez, T. Solving the multiple instance problem with axis-parallel rectangles. *Artif. Intell.* **1997**, *89*, 31–71. [[CrossRef](#)]
24. Zhou, Z.H.; Zhang, M.L. *Neural Networks for Multi-Instance Learning*; Technical Report, AI Lab; Computer Science and Technology Department, Nanjing University: Nanjing, China, August 2002.
25. Amar, R.; Dooly, D.R.; Goldman, S.A.; Zhang, Q. Multiple-Instance Learning of Real-Valued Data. In Proceedings of the Eighteenth International Conference on Machine Learning (ICML '01), Williamstown, MA, USA, 28 June–1 July 2001; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 2001; pp. 3–10.
26. Liu, D.; Zhou, Y.; Sun, X.; Zha, Z.; Zeng, W. Adaptive Pooling in Multi-instance Learning for Web Video Annotation. In Proceedings of the 2017 IEEE International Conference on Computer Vision Workshops (ICCVW), Venice, Italy, 22–29 October 2017; pp. 318–327. [[CrossRef](#)]
27. Wang, Y.; Li, J.; Metze, F. Comparing the Max and Noisy-Or Pooling Functions in Multiple Instance Learning for Weakly Supervised Sequence Learning Tasks. *arXiv* **2018**, arXiv:1804.01146.
28. Caruana, R. Multitask Learning. *Mach. Learn.* **1997**, *28*, 41–75. [[CrossRef](#)]
29. Pan, S.J.; Yang, Q. A Survey on Transfer Learning. *IEEE Trans. Knowl. Data Eng.* **2010**, *22*, 1345–1359. [[CrossRef](#)]
30. Zhang, M.L.; Zhou, Z.H. A Review on Multi-Label Learning Algorithms. *IEEE Trans. Knowl. Data Eng.* **2014**, *26*, 1819–1837. [[CrossRef](#)]
31. Zhang, Y.; Yang, Q. An overview of multi-task learning. *Natl. Sci. Rev.* **2018**, *5*, 30–43. [[CrossRef](#)]

32. Neural Information Processing Scaled for Bioacoustics (NIPS4B) 2013 Bird Song Competition. Available online: <http://sabiody.univ-tln.fr/nips4b/challenge1.html> (accessed on 15 August 2018).
33. Transcriptions for the NIPS4B 2013 Bird Song Competition Training Set. Available online: [https://figshare.com/articles/Transcriptions\\_of\\_NIPS4B\\_2013\\_Bird\\_Challenge\\_Training\\_Dataset/6798548](https://figshare.com/articles/Transcriptions_of_NIPS4B_2013_Bird_Challenge_Training_Dataset/6798548) (accessed on 15 August 2018).
34. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. In Proceedings of the 3rd International Conference for Learning Representations, San Diego, CA, USA, 7–9 May 2015.
35. Lasseck, M. Bird song classification in field recordings: Winning solution for NIPS4B 2013 competition. In Proceedings of the International Symposium Neural Information Scaled for Bioacoustics, Lake Tahoe, NV, USA, 10 December 2013; pp. 176–181.
36. Duong, L.; Cohn, T.; Bird, S.; Cook, P. Low Resource Dependency Parsing: Cross-lingual Parameter Sharing in a Neural Network Parser. In Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing, Beijing, China, 27–29 July 2015.
37. Misra, I.; Shrivastava, A.; Gupta, A.; Hebert, M. Cross-Stitch Networks for Multi-task Learning. In Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Vegas, NV, USA, 27–30 June 2016; pp. 3994–4003. [CrossRef]
38. Yang, Y.; Hospedales, T. Trace Norm Regularised Deep Multi-Task Learning. In Proceedings of the 5th International Conference on Learning Representations Workshop, Toulon, France, 24–26 April 2017.
39. Ruder, S.; Bingel, J.; Augenstein, I.; Søgaard, A. Sluice networks: Learning what to share between loosely related tasks. *arXiv* **2017**, arXiv:1705.08142v2.



© 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).