

A JOINT DETECTION-CLASSIFICATION MODEL FOR AUDIO TAGGING OF WEAKLY LABELLED DATA

Qiuqiang Kong, Yong Xu, Wenwu Wang, Mark D. Plumbley

Center for Vision, Speech and Signal Processing (CVSSP)
University of Surrey
{q.kong, yong.xu, w.wang, m.plumbley}@surrey.ac.uk

ABSTRACT

Audio tagging aims to assign one or several tags to an audio clip. Most of the datasets are *weakly labelled*, which means only the tags of the clip are known, without knowing the occurrence time of the tags. The labeling of an audio clip is often based on the audio events in the clip and no event level label is provided to the user. Previous works have used the *bag of frames* model assume the tags occur all the time, which is not the case in practice. We propose a *joint detection-classification* (JDC) model to detect and classify the audio clip simultaneously. The JDC model has the ability to attend to informative and ignore uninformative sounds. Then only informative regions are used for classification. Experimental results on the “CHiME Home” dataset show that the JDC model reduces the equal error rate (EER) from 19.0% to 16.9%. More interestingly, the audio event detector is trained successfully without needing the event level label.

Index Terms— audio tagging, weakly labelled data, *joint detection-classification* model, acoustic event detection

1. INTRODUCTION

Audio tagging aims to assign an audio clip with one or several tags. The clips are typically short segments such as 10 seconds of a long recording. Audio tagging has applications such as audio retrieval [1] and audio classification [2]. Accurate audio tagging relies on the amount of labelled audio data, including *clip level* labelled audio data and *event level* labelled audio data. In clip level labelling, each audio clip is labelled with one or several tags without indicating their occurrence time, while in event level labelling, each audio clip is labelled with both tags and their occurrence time. We refer to the clip level labelled data as *weakly labelled data* [3]. Figure 1 shows an audio clip and its event level label. The clip¹ level label is “children speech”, “percussion” and “other

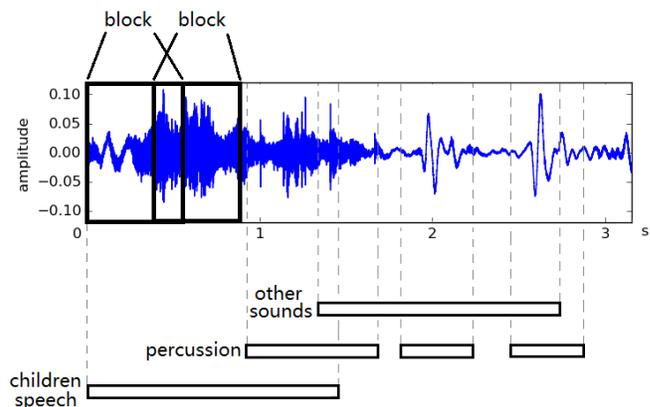


Fig. 1. Waveform of an audio clip and its event level label. Clip level label: “children speech”, “percussion”, “other sounds”.

sounds”.

Event level labelled data is scarce and insufficient for industrial use [4]. On the other hand, the amount of weakly labelled data on the internet is exploding in recent years. Every minute, over 300 hours videos are uploaded to YouTube². Each of the clip has a few tags roughly tagged by the uploaders describing the content such as “music”, “dogs”, “Obama”. These weakly labelled data have huge potential values but are currently difficult to use. First, the tags only indicate the existence but not the occurrence times of the events. Second, the duration of audio events varies from events to events. Some events such as “music” may last for a few minutes, while other events such as “gun shot” may only last for hundreds of milliseconds. Third, overlapping events are difficult to classify. These difficulties limit the use of weakly labelled data.

Previous work in audio tagging includes the *bag of blocks* (BOB) model [2][5] and the *global model* [6]. The BOB model is a generalization of the *bag of frames model* [2]. In the BOB model, each audio clip is divided into overlapping blocks (Figure 1). Each block contains one or several frames. In the training phase, for a machine learning system,

This research is supported by EPSRC grante EP/N014111/1 “Making Sense of Sounds” and research scholarship from the China Scholarship Council (CSC).

¹The clip is CR_lounge_220110_0731.s0_chunk39 from the “CHiME Home” dataset.

²<http://tubularinsights.com/youtube-300-hours/>

each block is assigned all tags in the clip as ground truth for training. In the recognition phase, the predicted tags from all blocks in a clip are used to estimate a predicted probability of the tag in that clip. The BOB model is based on an assumption that tags occur in all the blocks, which is however not the case in practice.

In contrast with the bag of blocks model, the global model feeds the whole clip into the model, without dividing the clip into blocks (Figure 2(b)). The advantage of the global model over the BOB model is that there is no “tags occur all the time” assumption. In addition, long-term context is considered. Inspired by image classification [7], the global model can be for example, a convolutional neural network (CNN) on the clip’s spectrogram [6].

However, neither of the BOB model nor the global model looks for the audio events likes humans do before tagging an audio clip. In this paper, we propose a *joint detection-classification* (JDC) model inspired by humans perception [8]. Humans use two steps to label an audio, a detection step and a classification step [8]. In the detection step, humans listen when to *attend* to sounds and when to *ignore* sounds. They attend to informative audio events and ignore uninformative sounds or silences. In a classification task, humans only tag informative events. Our JDC model is trained on weakly labelled data, without needing event level labelled data.

Our work is organized as follows, Section 2 introduces related works. Section 3 describes the JDC model. Section 4 shows experimental results. Section 5 gives conclusions.

2. RELATED WORK

Audio classification and detection have obtained increasing attention in recent years. such as the open challenges for audio classification, audio event detection and audio tagging including CLEAR 2007 [9], DCASE2013 [2] and DCASE2016 [5].

For audio classification and audio tagging, Mel frequency cepstrum coefficient (MFCC) and Gaussian mixture model (GMM) is widely used as baseline [5]. Recent methods include deep neural networks (DNNs) and convolutional neural networks (CNNs), with inputs varying from Mel energy [10], spectrogram [11] to constant Q transform (CQT) [12]. For non neural network models, nonnegative matrix factorization (NMF) on the spectrogram [13] was used and obtained competitive result in the DCASE2016 Challenge. Methods combining the BOB model and the global model are proposed in [4].

For audio event detection, generative models such as NMF [14], probabilistic latent component analysis (PLCA) [15] and hidden Markov model (HMM) [16] are used with inspiration from music transcription. Neural network methods including recurrent neural networks (RNNs) [17] and bidirectional long-short term memory (BLSTM) [18] are also

used for event detection.

3. JOINT DETECTION–CLASSIFICATION MODEL

For the audio tagging task, assume there are K different audio tags. We denote the labelled target of an audio clip as $t \in \{0, 1\}^K$ where $t_k \in \{0, 1\}$ represents the existence of the k -th tag. For the bag of blocks model (Figure 2(a)), the audio clip is divided into overlapped blocks $x_m, m = 1, 2, \dots, M$. We denote y_{mk} as the predicted occurrence probability of the k -th tag on the m -th block. In the training phase, the “tags occur in all blocks” assumption is applied and each block is assigned with the clip’s tags as ground truth $t_{mk} = t_k$. The model is trained by minimizing the loss function

$$loss = \frac{1}{M} \sum_{k=1}^K \sum_{m=1}^M d(y_{km}, t_{mk}) \quad (1)$$

where $d(\cdot, \cdot)$ is the binary cross-entropy error.

In the recognition phase, we obtain y_{km} on each k -th tag and each m -th block by feeding the m -th block to the model. The occurrence probability of the k -th tag in the clip p_k is obtained by averaging y_{km} over the blocks

$$p_k = \frac{1}{M} \sum_{m=1}^M y_{km} \quad (2)$$

However, the “tags occur in all blocks” assumption in the BOB model can not always be satisfied in practice, such as the “gun shot” will not occur all the time in an audio clip. The JDC model is able to overcome this problem. Similar to the BOB model, the audio clip is divided into overlapped blocks. The JDC model consists of a detector and a classifier acting on each tag and each block with output value between 0 and 1. The output of the classifier y_{km} on the k -th tag and the m -th block indicates how possible the m -th block has tag k . The output of the detector w_{km} on the k -th tag and the m -th block indicates how informative the m -th block is when classifying the k -th tag. If w_{km} is close to 1 it means the m -th block is informative and should be attended when classifying the k -th tag. If w_{km} is close to 0 it means the m -th block is uninformative and should be ignored when classifying the k -th tag.

Compared with the BOB model in Figure 2(a), the JDC model in Figure 2(b) applies a merging operation to calculate the occurrence probability of the k -th tag p_k directly, without calculating the probability over the blocks, hence avoids the “tags occur in all blocks” assumption. It is easy to define $p_k = \sum_{m=1}^M w_{km} y_{km}$, but this will cause problem. First, the detector w_{km} and the classifier y_{km} are no longer distinguishable because they are symmetric. Second, p_k may exceed 1 which is no longer a probability. One way to solve this problem is to enforce the constraint $\sum_{m=1}^M w_{km} = 1$, or equally,

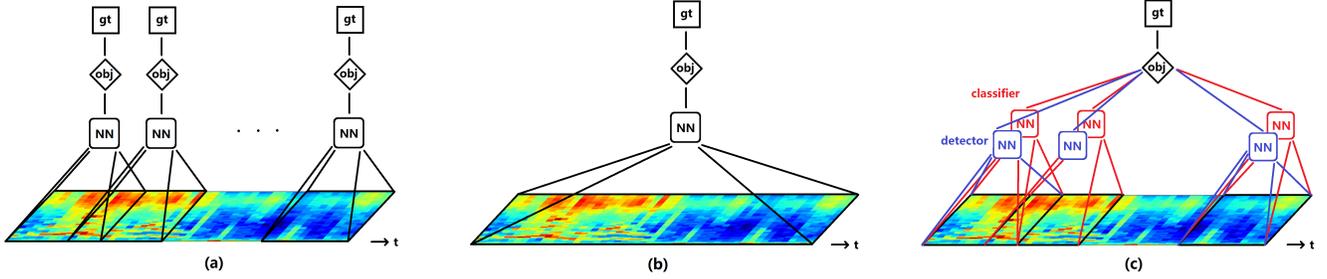


Fig. 2. (a) The BOB model. (b) The global model. (c) The JDC model

use normalized w_{km} over blocks to turn the constrained problem to an unconstrained problem. We define \tilde{w}_{km} as the normalized detector

$$\tilde{w}_{km} = \frac{w_{km}}{\sum_{m=1}^M w_{km}} \quad (3)$$

Define p_k as the occurrence probability of the k -th tag in the clip

$$p_k = \sum_{m=1}^M \tilde{w}_{km} y_{km} \quad (4)$$

Define the loss function of the JDC model as

$$loss = \sum_{k=1}^K d(p_k, t_k) \quad (5)$$

where $d(\cdot, \cdot)$ is the binary cross-entropy error.

In the recognition phase, the occurrence probability of the k -th tag in the clip is calculated by equation (4).

Figure 2(a-c) demonstrates the BOB model, the global model and the JDC model, respectively. The rounded rectangle, diamond and square represents the NN model, objective function and ground truth respectively. The NN can be any kinds of DNNs including CNNs, RNNs and LSTMs. In the training phase, the error back propagates from loss function to both the detector and the classifier in the JDC model (Figure 2(c)). Both the detector and the classifier can be updated according to the gradient.

The JDC model has the following features:

- Weakly labelled data can be fed into the model directly without the event level label.
- The detector’s attend and ignore mechanism simulate the humans perception [8].
- The detector’s attend and ignore mechanism facilitate the recognition of short audio events.
- Audio event detector is trained without needing the event level label.

4. EXPERIMENTS

We apply both the BOB model and the JDC model on the CHiME home dataset [19]. The development set consists of 1946 4 second chunks divided into 5 folds. Tags include “child speech”, “male speech”, “female speech”, “TV”, “percussive sounds”, “broadband noise” and “other identifiable sounds” and ‘silence’, denoted by ‘c’, ‘m’, ‘f’, ‘v’, ‘p’, ‘b’, ‘o’ and ‘S’, respectively. Another 816 chunks are kept for evaluation. The dataset is weakly labelled without event level labels.

4.1. Features

We apply fast Fourier transform (FFT) with non overlapping Hamming window with size of 1024 on the 16kHz audio files. Mel filter bank features with 40 bins are extracted using librosa toolbox³. Each block contains 11 frames and the hop size between the neighboring blocks is 1.

4.2. Models

To compare the BOB model and the JDC model, we simply use the fully connected (FC) neural network to model both the classifier and the detector. This can be extended to any neural network (Section 3).

For the BOB model, the classifier takes one block as input followed by 3 FC layer with 500 units per hidden layer. ReLU activation function and dropout probability of 0.2 are applied to each hidden layer. Sigmoid function is applied on the classifier’s output. For the JDC model, the classifier is the same as the BOB model, the detector takes one block as input followed by a mean pool layer along time axis and a single layer linear transform with sigmoid output. Experiments show that the detector described by a shallow layer NN is sufficient on the “CHiME Home” dataset. Adam optimizer is used and the learning rate is tuned to 0.0002. The system is built on *Hat* deep learning framework⁴. The source code of the JDC model is available on github⁵.

³<https://github.com/librosa/librosa>

⁴<https://github.com/qiuqiangkong/Hat>

⁵https://github.com/qiuqiangkong/audio_tagging_jdc

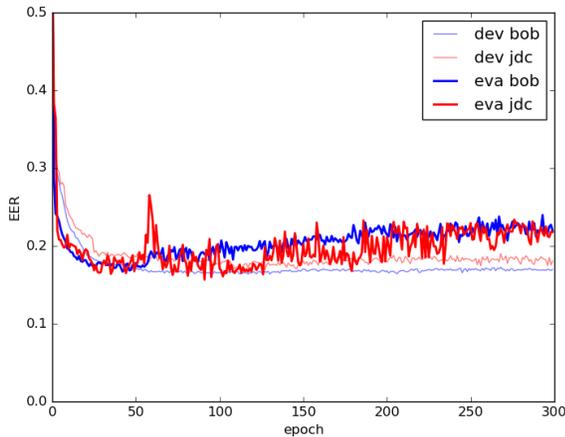


Fig. 3. EER of the BOB model and the JDC model on development set and evaluation set.

4.3. Results

We use 5-fold cross validation on development set to select the best model (65 epochs for BOB model and 110 epochs for JDC model). Then the best model is applied to the evaluation set. The EERs of the BOB model and the JDC model on development set and evaluation set are shown in Figure 3 and Table 1. It is noticeable the JDC model performs slightly worse than the BOB model on development set (equal error rate (EER) of 17.3% against 16.7%) but performs better on the evaluation set (EER of 16.9% against 19.0%). This indicates with more training data, the JDC is more effective and suffers less from overfitting.

Table 1. EER on development set and evaluation set.

	EER (dev.)	EER (eva.)
MFCC + GMM (baseline)	21.3%	20.9%
BOB	16.7%	19.0%
JDC	17.3%	16.9%

4.4. Visualize Learned Representation

It is interesting to visualize what the BOB model and the JDC model learn. Figure 4(a) shows the Mel energy spectrogram corresponds to the audio clip in Figure 1. The tags for this clip include children speech, percussion and other sounds, denoted by ‘c’, ‘p’, ‘o’, respectively. Figure 4(b) shows the detector output \tilde{w}_{km} , where the bright regions indicate informative regions. Figure 4(c) shows the classifier output y_{km} , where the bright regions indicate the probability of the occurrence of the events regardless of the detector. Figure 4(d) is the detection-classification output $\tilde{w}_{km}y_{km}$, where bright re-

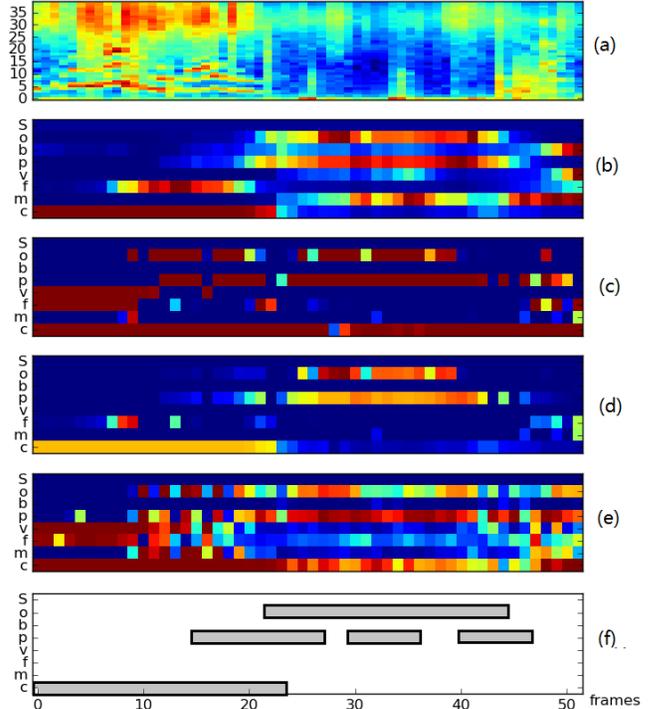


Fig. 4. (a) Mel energy spectrogram of the audio clip in Figure 1. (b) Detector output of the JDC model. (c) Classifier output of the JDC model. (d) Detector-classifier output of the JDC model. (e) Classifier output of the BOB model. (f) Manually labelled event level label.

gions indicates the occurrence probability of the events. Figure 4(e) is the classifier output of the BOB. Figure 4(f) is the event level label of the clip. Compared Figure 4(d) and Figure 4(f), the JDC detects the audio events successfully, even if no event level label data is used in the training phase. However the classifier output of the BOB model in Figure 4(e) does not contain the occurrence times of the events.

5. CONCLUSION

We proposed a *joint detection-classification* (JDC) model to use *weakly labelled data*. By introducing the detector, the model simulates the humans’ attend and ignore ability. In addition, the JDC is able to perform event detection without needing event level label. For the future work, CNNs, RNNs, LSTMs can be used to replace the fully connected NN. In Figure 3 the EER curve of the JDC model oscillate considerably, which may be mitigated by applying different learning rates on the detector and classifier. Expectation maximization (EM) algorithm could be used to replace the joint training method in the JDC model in the future.

6. REFERENCES

- [1] Guodong Guo and Stan Z Li, “Content-based audio classification and retrieval by support vector machines,” *IEEE transactions on Neural Networks*, vol. 14, no. 1, pp. 209–215, 2003.
- [2] Dan Stowell, Dimitrios Giannoulis, Emmanouil Benetos, Mathieu Lagrange, and Mark D Plumbley, “Detection and classification of acoustic scenes and events,” *IEEE Transactions on Multimedia*, vol. 17, no. 10, pp. 1733–1746, 2015.
- [3] Gideon S Mann and Andrew McCallum, “Generalized expectation criteria for semi-supervised learning with weakly labeled data,” *Journal of machine learning research*, vol. 11, no. Feb, pp. 955–984, 2010.
- [4] Anurag Kumar and Bhiksha Raj, “Audio event detection using weakly labeled data,” *arXiv preprint arXiv:1605.02401*, 2016.
- [5] Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen, “Tut database for acoustic scene classification and sound event detection,” in *24th European Signal Processing Conference*, 2016, vol. 2016.
- [6] Keunwoo Choi, George Fazekas, and Mark Sandler, “Automatic tagging using deep convolutional neural networks,” *arXiv preprint arXiv:1606.00298*, 2016.
- [7] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [8] Albert S Bregman, *Auditory Scene Analysis: The Perceptual Organization of Sound*, MIT press, 1994.
- [9] Mong How Ooi, Tom Solomon, Yuwana Podin, Anand Mohan, Winnie Akin, Mohd Apandi Yusuf, Sylvia del Sel, Kamsiah Mohd Kontol, Boon Fu Lai, Daniela Clear, et al., “Evaluation of different clinical sample types in diagnosis of human enterovirus 71-associated hand-foot-and-mouth disease,” *Journal of Clinical Microbiology*, vol. 45, no. 6, pp. 1858–1866, 2007.
- [10] Michele Valenti, Aleksandr Diment, Giambattista Parascandolo, Stefano Squartini, and Tuomas Virtanen, “DCASE 2016 acoustic scene classification using convolutional neural networks,” 2016.
- [11] Yoonchang Han and Kyogu Lee, “Acoustic scene classification using convolutional neural network and multiple-width frequency-delta data augmentation,” *arXiv preprint arXiv:1607.02383*, 2016.
- [12] Thomas Lidy and Alexander Schindler, “CQT-based convolutional neural networks for audio scene classification,” in *Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE 2016)*, Budapest, Hungary, 2016.
- [13] Victor Bisot, Romain Serizel, Slim Essid, et al., “Acoustic scene classification with matrix factorization for unsupervised feature learning,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 6445–6449.
- [14] Arnaud Dessen, Arshia Cont, and Guillaume Lemaitre, “Real-time detection of overlapping sound events with non-negative matrix factorization,” in *Matrix Information Geometry*, pp. 341–371. Springer, 2013.
- [15] Emmanouil Benetos, Mathieu Lagrange, Mark D Plumbley, et al., “Detection of overlapping acoustic events using a temporally-constrained probabilistic model,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 6450–6454.
- [16] Toni Heittola, Annamaria Mesaros, Antti Eronen, and Tuomas Virtanen, “Context-dependent sound event detection,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2013, no. 1, pp. 1–13, 2013.
- [17] Simon Leglaive, Romain Hennequin, and Roland Badeau, “Singing voice detection with deep recurrent neural networks,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 121–125.
- [18] Giambattista Parascandolo, Heikki Huttunen, and Tuomas Virtanen, “Recurrent neural networks for polyphonic sound event detection in real life recordings,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 6440–6444.
- [19] Peter Foster, Siddharth Sigtia, Sacha Krstulovic, Jon Barker, and Mark D Plumbley, “CHiME-home: A dataset for sound source recognition in a domestic environment,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2015, pp. 1–5.