

# Acoustic Features Fusion using Attentive Multi-channel Deep Architecture

Gaurav Bhatt<sup>1</sup>, Akshita Gupta<sup>1</sup>, Aditya Arora<sup>1,2</sup>, Balasubramanian Raman<sup>1</sup>

<sup>1</sup>Indian Institute of Technology Roorkee <sup>2</sup>Inception Institute of Artificial Intelligence, Abu Dhabi

gauravbhatt.deeplearn@gmail.com, akshitadvlp@gmail.com, adityadvlp@gmail.com,  
balarfma@iitr.ac.in

## Abstract

In this paper, we present a novel deep fusion architecture for audio classification tasks. The multi-channel model presented is formed using deep convolution layers where different acoustic features are passed through each channel. To enable dissemination of information across the channels, we introduce attention feature maps that aid in the alignment of frames. The output of each channel is merged using interaction parameters that non-linearly aggregate the representative features. Finally, we evaluate the performance of the proposed architecture on three benchmark datasets :- DCASE-2016 and LITIS Rouen (acoustic scene recognition), and CHiME-Home (tagging). Our experimental results suggest that the architecture presented outperforms the standard baselines and achieves outstanding performance on the task of acoustic scene recognition and audio tagging.

**Index Terms:** acoustic scene recognition, audio tagging, deep learning, feature fusion.

## 1. Introduction

While deep architectures have been shown to achieve top performance on simple audio classifications tasks like speech recognition and music genre detection, their application to complex acoustic problems has a significant room for improvements. Two challenging audio classification tasks that have been recently introduced are acoustic scene recognition (ASR) [1] and audio tagging [2]. ASR is defined as the identification of environments in which an audio is captured, while audio tagging is a multi-label classification task. To solve these challenge, a majority of current research has shown the effectiveness of feature fusion with deep architectures such as deep neural networks (DNN) [1], convolution neural networks (CNN) [3], and recurrent neural models [4, 5]. A general problem with multi-channel deep networks is the limited memory and low interaction between subsequent layers. To the problem of limited memory, an attention mechanism has been introduced. In this paper, we demonstrate an attention mechanism that can be used to guide the information flow across multiple channels enabling a smoother convergence that results in better performance.

A popular approach for acoustic scene recognition (ASR) and the tagging task is to use the low-level or high-level acoustic features such as Mel-frequency cepstral coefficients (MFCCs), Mel-spectrogram, Mel-bank, log Mel-bank features, etc., with the state-of-the-art deep models [6, 7]. Some of these acoustic features possess complementary qualities, that is, for two given features, one is apt in identifying certain specific classes, while the other is suitable for the rest. This complementarity property may depend upon the spectrum range in which these features operate. Hence, it is possible to obtain a boost in performance when multiple complementary features are combined together as the overall range over which the learning models can operate is increased. For an instance, the augmentation of the delta

and acceleration coefficients with MFCC proves to be more effective for acoustic scene classification [8, 9]. Similarly, Mel frequency components and the log of Mel components are examples of one such complementary pair that we use in our work.

In this paper, we combine acoustic features using a multi-channel approach, where we add subsequent convolution and pooling layers to the input low-level complementary features. To effectively amalgamate the properties of several acoustic features we introduce three feature fusion techniques: early fusion, late fusion, and hybrid fusion, depending on the position where acoustic features are fused together (here, position refers to the intermediate neural layers). The early fusion strategy comprises of stacked attention layers that introduce a flow of information between the channels to facilitates better convergence. In the late fusion, we introduce trainable parameters to the model which enables better generalization for the audio classification and tagging tasks. Finally, we demonstrate the performance evaluation of the proposed model on DCASE-2016 (ASR), LITIS-Rouen (ASR) and CHiME-Home (audio tagging) datasets.

## 2. Related Work

In this section, we discuss previous work related to audio classification and audio tagging. These domains have recently gained popularity because of open challenges such as Dcase2013 [10], Dcase2016 [11], and Chime2018 [12].

For audio classification and tagging, the Mel frequency cepstrum coefficient (MFCC) and the Gaussian mixture model (GMM) are widely used as a baseline [13, 14]. Most published works in this domain uses Mel-spectrograms as features with deep parallel convolution architectures [3, 6, 15]. Traditional techniques focus on using hand-crafted audio features as the input to various machine learning classifiers. Some of the recent research has been focused on passing the term-frequency representation of the waveform through convolution neural networks [16, 15, 6, 5, 3], or deep neural networks [17, 18, 1] However, deep networks have not yet outperformed feature-based approaches.

Currently, the main challenges to audio tagging datasets are the uneven distribution of samples, along with uneven labels [19]. For audio tagging tasks, attention models have been introduced, which have shown more accurate results than other hybrid combinations of deep models [20].

## 3. Complementary Acoustic Features

The Mel and log-Mel are a set of complementary acoustic features (CAF). The Mel frequencies capture classes which lie in the higher frequency domain and log-Mel frequencies capture classes that lie in the lower frequency domain. We conjecture that passing the features via a multi-channel model it is possible to efficiently combine the complementary properties inhibited

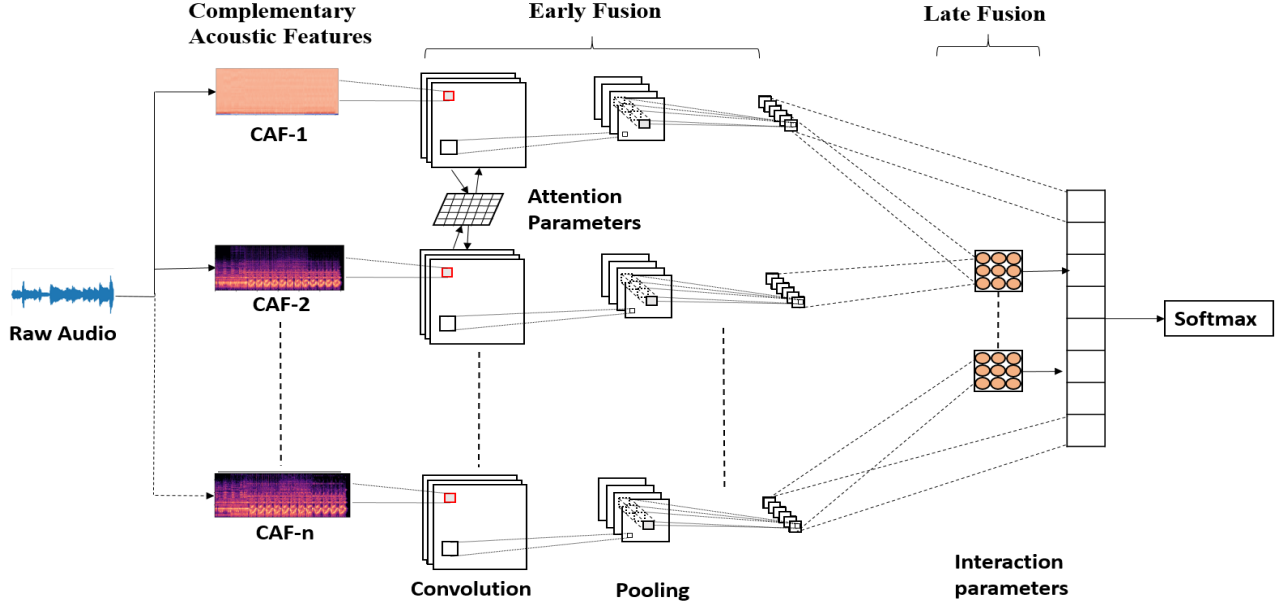


Figure 1: The proposed multi-channel deep fusion architecture with complementary acoustic features (CAF) as input.

by these features. We calculate the Mel spectrum by taking the spectrogram of the raw audio and combining it with the transposition of the spectrogram. The Mel frequencies are kept at 40, resulting in 40-dim Mel features with non-overlapping frames and 50% hop size.

We also compute the Mel-Frequency components (MFCCs) and select 13 mel frequency ceptral coefficients (including the 0<sup>th</sup> order coefficient) with a window size of 1024 frames with 50% overlapping. With the time-varying information, we combine the first and second derivatives (i.e. the delta and acceleration coefficients).

For the Constant Q Transformation (CQT) features, we select 80 bins per octave with 50% hop size.

## 4. Multi-Channel Deep Fusion

In the model presented, each complementary acoustic feature is passed on to a separate CNN, thereby forming a multiple channel architecture. Each channel is formed using 128 kernels in the first layer, with a receptive filter of size  $3 \times 3$ . This gives us the convolved features which are then sub-sampled using a max/global pooling with filter size  $2 \times 2$ . In the second convolution layer, we use a large number of kernels (256) for exploring higher-level representations. The activation function that we use is the rectified linear units (Relu) in the subsequent convolution layers. All the parameters are shared across the layers.

Some problems related to long audio recordings are that the channels can be noisy and the number of foreground events may not be sufficient. To improve the performance of the underlying deep architectures, we follow the work of [5, 3] and divide of the audio features into segments. In place of using the whole audio feature as the training sample, we decompose the snippet into T segments having  $length = 1024$  frames with a hop size of 512 frames. This is done to ensure that the underlying learning model is able to capture the important foreground events in the long recordings.

### 4.1. Early fusion

In the multi-channel architecture, each layer of the channel  $i$  captures higher level representation based on the acoustic feature passed through. These representative features differ from the corresponding layer of channel  $j$ , where  $j \in \{N - i\}$  (shown in Figure 1). The temporal sequences input can be aligned together with the help of the attention mechanism [21], so that the properties of the audio sequences important in one channel can be reflected in the others. We compute an attention matrix to align two audio representative feature maps, which is followed by the addition of trainable parameters to transform the matrix into convolution feature maps (Equation 3 and 4). This is essentially attentive convolution and helps the model to assign a higher score important events than the rest (shown in Figure 2).

Inspired by [21], we introduce a similarity feature matrix  $S_{ij}$  that can influence convolutions across multiple channels, where  $S_{ij}$  is shared across channels  $i$  and  $j$ . The similarity features assign a higher score to those frames in channel  $c_i$  that are

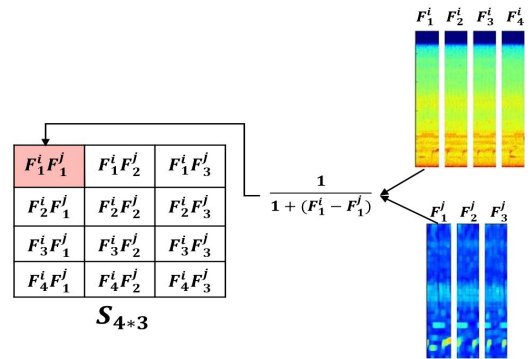


Figure 2: Computing similarity matrix  $S_{4 \times 3}$  for two acoustic features,  $F^i$  and  $F^j$ , with frame sizes of 4 and 3 respectively.

relevant to frames in channel  $c_j$ , where  $j \in \{N - i\}$ . The values in the row of  $S_{ij}$  denote the distribution of the similarity weights of the  $t^{th}$  frame of  $c_i$  with respect to  $c_j$ , and the columns of  $S_{ij}$  represents the distribution of similarity weights of the  $T^{th}$  frame of  $c_j$  with respect to  $c_i$ . As shown in Figure 2, the rows of the similarity matrix  $S_{4 \times 3}$  represent the distribution of audio frames of feature  $i$  and columns represent the distribution of feature  $j$ .

Let the representative feature for channel  $i$  and  $j$  be given by  $F_i^r$  and  $F_j^r$  respectively. The similarity feature matrix  $S_{ij}$  is then computed as

$$S_{ij}^{x,y} = \text{similarity}_{score}(F_i^r[:, x], F_j^r[:, y]) \quad (1)$$

where  $\text{similarity}_{score}(a, b)$  is given by

$$\text{similarity}_{score}(a, b) = \frac{1}{1 + |a - b|} \quad (2)$$

Using the similarity matrix  $S_{ij}$ , the attentive feature representations  $F_{k \in \{i,j\}}^s$  are computed as

$$F_i^s = W_i \cdot S_{ij}^T \quad (3)$$

$$F_j^s = W_j \cdot S_{ij} \quad (4)$$

The channel representative features  $F_i^r$  and the attentive representative features  $F_i^s$  are stacked together as an order-3 tensor matrix before passing to further convolution layers.

The idea of attentive convolution works in the early stages and we call it early fusion technique.

## 4.2. Late Fusion

Once all the features maps are obtained using the multi-channel architecture, we use an interaction matrix  $W$  to compute the interaction between them (shown as interaction parameters in Figure 1). Given the feature maps for channel  $i$  and  $j$  are  $F_i^p$ ,  $F_j^p$  respectively, the interaction score is computed as

$$\text{Score}_I(F_i^p, F_j^p) = (F_i^p)^T \cdot W \cdot F_j^p \quad (5)$$

We share a common weight matrix  $W$  across all the possible pairs. This is done to ensure that all the channels interact with each other and the computation burden over the network is reduced.

## 4.3. Parameters Sharing

The proposed multi-channel architecture, along with early and late fusion, increases the number of trainable parameters, which results in optimization issues. We solve this problem by sharing the parameters across layer  $i$ . That is, the convolution layer shown in Figure 1 is shared across all the channels. We also share the attention weights  $W_i$  and  $W_j$ , as computed by Equation 3 and 4. The interaction matrix  $W$  used in the late fusion is also shared across all the channels. Finally, we combine the representative features computed at each pooling layer with the similarity scores (shown in Figure 1 as late fusion). This ensures that the model can capture representations at each level of abstraction.

# 5. Experiments and Results

## 5.1. Dataset Description

**DCASE-2016.** The dataset consists of 1560 samples which are divided into the development dataset (1170) and evaluation

Table 1: *Hyper-parameters used in feature fusion layer for training.*

Hyper-parameter	Value
Depth	6
Number of neurons	600
Regularization	L2 ; dropout
L2	0.005
Dropout	0.3
Optimizer	adam
Loss	Binary crossentropy
Learning rate	0.001
Batch size	100
Number of Epochs	10
filter Length	(60,3)
Number of filters	256

dataset (390). Each audio class in the development set consists of 78 samples (39 minutes of audio) while the evaluation dataset is comprised of 26 samples (13 minutes of audio) for each class. The organizers of DCASE-2016 have provided the four cross-fold validation meta-data which is used for tuning the parameters of the network.

**LITIS Rouen.** This dataset consists of audio recordings of 30 second duration, which are divided into 19 classes. The total number of audio samples is 3026. We divide the dataset into 10 cross-validation sets with 80:20 random splits each. The final performance of the proposed techniques is computed by averaging the accuracy on all 10 test sets.

**CHiME-Home.** This dataset consists of audio recordings of 4-second duration in two sampling frequencies: 48Khz in stereo and 16Khz in mono. We use mono audio data with the 16Khz sampling frequency, which is further divided into 7 classes. The total number of audio samples is 2792, which are divided into 1946 development sets and 846 evaluation sets. Each piece is annotated with a single or multiple labels. [6].

## 5.2. Hyperparameters

The architecture details of the individual channels are described in Table 1. The output of the global average pooling layers is concatenated and then passed on to the intermediate matrix, which computes the interaction between them. Finally, we use *adam* as the optimizer for binary cross-entropy loss.

## 5.3. Baselines

For DCASE-2016 we use the Gaussian Mixture Model (GMM) with MFCC (including acceleration and delta coefficients) as the baseline system. This baseline is provided by DCASE-2016 organizers. The other baseline used is DNN with mel-components [1]. For LITIS-Rouen we use the HOG+CQA and DNN + MFCC results as the baseline. These results are taken from [7]. For the CHiME-Home dataset, we use the standard baseline of the MFCC+GMM system [14] and mel+DNN [17].

## 5.4. Results

The results for the task of ASR on DCASE-2016 and LITIS-Rouen dataset are shown in Table 2 (a) and (b), respectively. The use of deep fusion achieves the highest precision and F1 measure when compared to current state-of-the-art techniques on LITIS Rouen, while we achieve an accuracy of **88.7** on

Table 2: Performance of various architectures on ASR and Audio Tagging. The first half of each table shows the baseline results, the middle section is the performance of state-of-the-art systems and the bottom section is the performance of the proposed model.

(a) DCASE-2016		(b) LITIS-Rouen			(c) CHiME-Home	
Techniques	Accuracy %	Techniques	P %	F1 %	Techniques	EER
MFCC + GMM	72.50	DNN + MFCC	92.20	-	MFCC + GMM [14]	21.0
Mel + DNN [1]	81.00	CNN+Fusion [5]	96.30	96.50	Mel + DNN [17]	20.9
MFCC + fusion [8]	<b>89.70</b>	RNN+Fusion [7]	97.50	97.70	CQT + CNN [16]	16.6
Spectrogram + NMF [22]	87.70	LTE+Fusion [24]	95.90	96.20	MFCC + GMM [9]	17.4
LTE-fusion + CNN [5]	81.20	Proposed	<b>98.00</b>	<b>98.25</b>	MFCC + DNN [18]	17.85
BiMFCC + I-vector [23]	81.70				DAE + DNN [2]	14.8
Mel + CNN [3]	86.20				Mel + IMD [15]	<b>12.3</b>
BiMFCC + RNN [23]	80.20				log-mel + CRNN [6]	<b>11.3±0.6</b>
BiMFCC + Fusion [23]	88.10				Proposed	<b>14.0</b>
Proposed	88.70					

Table 3: Ablation results of the proposed model. Eight system configurations for the proposed deep fusion architecture.

Techniques	DCASE	ROUEN (F1)	CHiME
Vanilla	85.50	96.85	15.6
EF	86.1	96.36	15.0
LF	87.00	96.80	14.6
EF+LF	<b>88.70</b>	<b>98.25</b>	<b>14.0</b>

DCASE-2016, which is comparable to current top methods.. A similar technique of feature fusion, CNN+Fusion, was used by [5], where they used label-based embeddings. However, our proposed method of fusing complementary features results in better performance than all of their architectures.

For the task of audio tagging (Table 2 (c)), the proposed model achieved an equal error rate (EER) of **14.0**, which is better than the baselines and all DCASE-2016 challenge submissions. These submissions all omitted the *silence* class. In contrast, we keep the *silence* class and omit the *others* class, since it has high variance due to the introduction of random samples. We keep the *silence* class instead. In the task of audio tagging, the temporal models outperformed the non-temporal techniques and the reason being the dependence of temporal sequences on multi-label classification. As shown in Table 2 (b), most of the successful systems consist of CNN and RNN based models.

Finally, the ablation results are presented in Table 3, which demonstrates the performance of the deep fusion techniques introduced in our work. Here, Vanilla represents the basic feature fusion model without attention and interaction matrix. Here, we present four systems describing the in-depth analysis of the proposed architecture - vanilla (no fusion), early fusion (EF), late fusion (LF) and hybrid (EF+LF).

### 5.5. Discussion

The model combining early and late fusion (EF+EL) had the highest performance amongst all the introduced models. This is due to the enhanced interactivity across the channels. As shown in Table 3, the vanilla system is an amalgamation of multiple features without any fusion and acts as a baseline model. The early and late fusion models achieve a better performance than the baseline model. The LF system constitutes the fusion-by-

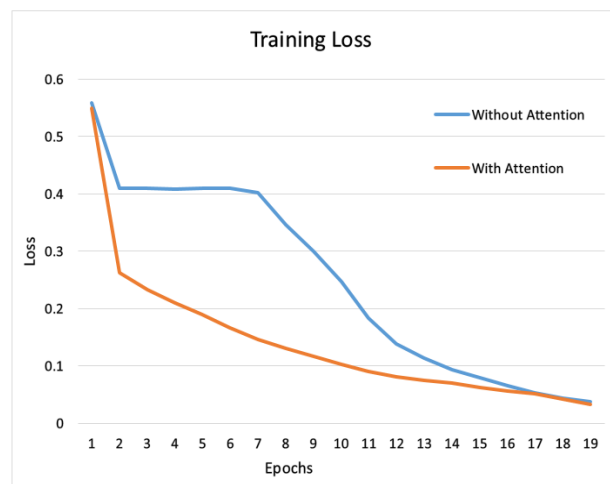


Figure 3: Training loss curves for the vanilla system along with the proposed attentive network.

multi-channel architecture, where the interaction parameters are responsible for non-linear feature augmentation. The similarity features are accompanied by additional trainable parameters, which result in higher performance but are computationally expensive to train.

Finally, we present the training curves for the vanilla model and the proposed architecture (Figure 3 and 4). We keep track of the mean square error (MSE) for each iteration along with the binary cross-entropy loss. This is done for the vanilla model and the proposed attention-based model. The training loss and MSE for the attention-based systems show a steep decrease in loss as compared to the vanilla model. Not only are the decreases in loss quick but the overall losses for attention-based models are also lower than that of the vanilla system. This demonstrates that the introduced attention and similarity parameters are responsible for a smoother convergence.

## 6. Conclusion

In this paper, we present a multi-channel architecture for the fusion of complementary acoustic features. Our idea is based on the fact that the introduction of attention parameters between



Figure 4: Training MSE curves for the vanilla system along with the proposed attentive network.

the channels results in better convergence. The proposed technique is general and can be applied to any audio classification task. A possible extension to our work would be to use the pairs or triplets of audio samples of similar classes and pass them through the multi-channel architecture. This could help to align the diverse audio samples of similar classes, making the model robust to audio samples that are difficult to classify.

## 7. Acknowledgement

We would like to express our gratitude towards Institute Computer Center (ICC), and the Indian Institute of Technology Roorkee, for providing us with the necessary resources for this work.

## 8. References

- [1] Q. Kong, I. Sobieraj, W. Wang, and M. D. Plumbley, "Deep neural network baseline for dcase challenge 2016," *Proceedings of DCASE 2016*, 2016.
- [2] Y. Xu, Q. Huang, W. Wang, P. Foster, S. Sigtia, P. J. Jackson, and M. D. Plumbley, "Unsupervised feature learning based on deep models for environmental audio tagging," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1230–1241, 2017.
- [3] M. Valenti, A. Diment, G. Parascandolo, S. Squartini, and T. Virtanen, "DCASE 2016 acoustic scene classification using convolutional neural networks," DCASE2016 Challenge, Tech. Rep., September 2016.
- [4] S. H. Bae, I. Choi, and N. S. Kim, "Acoustic scene classification using parallel combination of LSTM and CNN," Tech. Rep., September 2016.
- [5] H. Phan, L. Hertel, M. Maass, P. Koch, R. Mazur, and A. Mertins, "Improved audio scene classification based on label-tree embeddings and convolutional neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1278–1290, 2017.
- [6] G. Parascandolo, T. Heittola, H. Huttunen, T. Virtanen *et al.*, "Convolutional recurrent neural networks for polyphonic sound event detection," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1291–1303, 2017.
- [7] H. Phan, P. Koch, F. Katzberg, M. Maass, R. Mazur, and A. Mertins, "Audio scene classification with deep recurrent neural networks," *arXiv preprint arXiv:1703.04770*, 2017.
- [8] H. Eghbal-Zadeh, B. Lehner, M. Dorfer, and G. Widmer, "CP-JKU submissions for DCASE-2016: a hybrid approach using binaural i-vectors and deep convolutional neural networks," Tech. Rep., September 2016.
- [9] S. Yun, S. Kim, S. Moon, J. Cho, and T. Kim, "Discriminative training of gmm parameters for audio scene classification and audio tagging," *IEEE AASP Challenge Detect. Classification Acoust. Scenes Events*, 2016.
- [10] D. Stowell, D. Giannoulis, E. Benetos, M. Lagrange, and M. D. Plumbley, "Detection and classification of acoustic scenes and events," *IEEE Transactions on Multimedia*, vol. 17, no. 10, pp. 1733–1746, 2015.
- [11] A. Mesaros, T. Heittola, and T. Virtanen, "Tut database for acoustic scene classification and sound event detection," in *Signal Processing Conference (EUSIPCO), 2016 24th European*. IEEE, 2016, pp. 1128–1132.
- [12] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth 'chime' speech separation and recognition challenge: Dataset, task and baselines," *arXiv preprint arXiv:1803.10609*, 2018.
- [13] T. Heittola, A. Mesaros, and T. Virtanen, "DCASE2016 baseline system," DCASE2016 Challenge, Tech. Rep., September 2016.
- [14] P. Foster and T. Heittola, "DCASE2016 baseline system," DCASE2016 Challenge, Tech. Rep., September 2016.
- [15] Y. Xu, Q. Kong, Q. Huang, W. Wang, and M. D. Plumbley, "Convolutional gated recurrent neural network incorporating spatial features for audio tagging," in *Neural Networks (IJCNN), 2017 International Joint Conference on*. IEEE, 2017, pp. 3461–3466.
- [16] T. Lidy and A. Schindler, "CQT-based convolutional neural networks for audio scene classification and domestic audio tagging," DCASE2016 Challenge, Tech. Rep., September 2016.
- [17] Q. Kong, I. Sobieraj, W. Wang, and M. Plumbley, "Deep neural network baseline for DCASE challenge 2016," Tech. Rep.
- [18] Y. Xu, Q. Huang, W. Wang, and M. D. Plumbley, "Fully DNN-based multi-label regression for audio tagging," DCASE2016 Challenge, Tech. Rep., September 2016.
- [19] P. Foster, S. Sigtia, S. Krstulovic, J. Barker, and M. D. Plumbley, "Chime-home: A dataset for sound source recognition in a domestic environment," in *WASPAA*, 2015, pp. 1–5.
- [20] Y. Xu, Q. Kong, W. Wang, and M. D. Plumbley, "Large-scale weakly supervised audio classification using gated convolutional neural network," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 121–125.
- [21] W. Yin, H. Schütze, B. Xiang, and B. Zhou, "Abcnn: Attention-based convolutional neural network for modeling sentence pairs," *Transactions of the Association for Computational Linguistics*, vol. 4, pp. 259–272, 2016.
- [22] V. Bisot, R. Serizel, S. Essid, and G. Richard, "Supervised non-negative matrix factorization for acoustic scene classification," Tech. Rep., September 2016.
- [23] J. Li, W. Dai, F. Metze, S. Qu, and S. Das, "A comparison of deep learning methods for environmental sound detection," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 126–130.
- [24] H. Phan, L. Hertel, M. Maass, P. Koch, and A. Mertins, "Label tree embeddings for acoustic scene classification," in *Proceedings of the 2016 ACM on Multimedia Conference*. ACM, 2016, pp. 486–490.