

CAA-Net: Conditional Atrous CNNs With Attention for Explainable Device-Robust Acoustic Scene Classification

Zhao Ren¹, Student Member, IEEE, Qiuqiang Kong², Member, IEEE, Jing Han³, Member, IEEE, Mark D. Plumbley⁴, Fellow, IEEE, and Björn W. Schuller⁵, Fellow, IEEE

Abstract—Acoustic Scene Classification (ASC) aims to classify the environment in which the audio signals are recorded. Recently, Convolutional Neural Networks (CNNs) have been successfully applied to ASC. However, the data distributions of the audio signals recorded with multiple devices are different. There has been little research on the training of robust neural networks on acoustic scene datasets recorded with multiple devices, and on explaining the operation of the internal layers of the neural networks. In this article, we focus on training and explaining device-robust CNNs on multi-device acoustic scene data. We propose conditional atrous CNNs with attention for multi-device ASC. Our proposed system contains an ASC branch and a device classification branch, both modelled by CNNs. We visualise and analyse the intermediate layers of the atrous CNNs. A time-frequency attention mechanism is employed to analyse the contribution of each time-frequency bin of the feature maps in the CNNs. On the Detection and Classification of Acoustic Scenes and Events (DCASE) 2018 ASC dataset, recorded with three devices, our proposed model performs significantly better than CNNs trained on single-device data.

Index Terms—Acoustic scene classification, attention, conditional atrous convolutional neural networks, multi-device data, visualisation.

Manuscript received July 5, 2019; revised March 28, 2020 and August 12, 2020; accepted October 27, 2020. Date of publication November 12, 2020; date of current version November 18, 2021. This work was supported by European Union’s Horizon H2020 research and innovation programme under Marie Skłodowska-Curie grant agreement No. 766287 (TAPAS), EPSRC grant EP/N014111/1 “Making Sense of Sounds,” and a Research Scholarship from the China Scholarship Council (CSC) No. 201406150082. (Corresponding author: Zhao Ren.)

Zhao Ren is with the Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, 86159 Augsburg, Germany (e-mail: zhao.ren@informatik.uni-augsburg.de).

Qiuqiang Kong was with the Centre for Vision, Speech and Signal Processing, University of Surrey GU2 7XH Surrey, U.K. He is now with the ByteDance AI Laboratory, Shanghai 200233, China (e-mail: q.kong@surrey.ac.uk).

Jing Han was with the Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, 86159 Augsburg, Germany. She is now with the Department of Computer Science and Technology, University of Cambridge, CB3 0FD Cambridge, U.K. (e-mail: jh2298@cam.ac.uk).

Mark D. Plumbley is with the Centre for Vision, Speech and Signal Processing, University of Surrey, GU2 7XH Surrey, U.K. (e-mail: m.plumbley@surrey.ac.uk).

Björn W. Schuller is with the GLAM – Group on Language, Audio & Music, Imperial College London, SW7 2AZ London, U.K., and also with the Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, 86159 Augsburg, Germany (e-mail: schuller@ieee.org).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TMM.2020.3037534>.

I. INTRODUCTION

WITH the development of computer audition [1], *Acoustic Scene Classification* (ASC) has become a major research field, aiming to automatically recognise acoustic environments [2], [3]. The goal of ASC is to identify acoustic scenes in an audio stream, using computational approaches such as signal processing [4], [5], machine learning [6], and deep learning [7], [8]. ASC has been used in various applications, including context-aware services [9], intelligent wearable devices [10], and robot navigation systems [11].

In recent years, deep learning approaches have shown good performance in ASC [7], [12]. Deep learning can often learn high-level representations, which yield better performance than those using conventional machine learning methods [4], [13]. Many deep learning structures have been proposed for ASC, including *Convolutional Neural Networks* (CNNs) [7] and *Recurrent Neural Networks* (RNNs) [14]. In this article, we will use CNNs due to their ability to learn strong representations from spectrograms [15], [16]. Specifically, log mel spectrograms have been successfully utilised in ASC [7], [17]. In this regard, we extract log mel spectrograms as the inputs of the CNNs.

Varying characteristics of recording devices lead to various qualities of audio data, such as sampling rate, amplitude, and frequency response, and data distributions [18]. Audio data is often recorded with distinct devices, such as professional sound recording devices [19], and mobile devices [20]. However, it is challenging to optimise the training model on a multi-device audio dataset, due to the different data distributions of mismatched devices [21]. Supervised domain adaptation by transfer learning can adjust models from the source data to the target data [22]. Yet, it is time-consuming to train two deep models before and after transferring the learnt parameters. In this regard, jointly training a single model on multi-device data can reduce the computational complexity compared to training separate models. In previous studies [23], [24], researchers trained joint models on multi-device data. However, joint training cannot analyse the difference of data from mismatched devices, as the deep neural networks learn common parameters for multi-device data [25]. Inspired by the task of speaker verification [26], [27], we propose to condition the CNNs with device information, through feeding the device information into CNNs: we call this conditional training. Through conditional training, the CNN model

learns specific parameters for the data from each device. To further solve the problem caused by unknown device information in the test set, we show that our proposed multi-task conditional training can also condition the CNNs using predicted device information.

A CNN model generally contains convolutional layers, local pooling layers, and fully connected layers, to achieve a final classification [28]. To learn global information from the feature maps using the convolutional kernels, each convolutional layer applies either stride convolutions with strides of length $stride \geq 2$, where strides are the moving steps of convolutional kernels, or local pooling operations. Such CNNs can classify acoustic scenes with state-of-the-art performance [24]. Visualising the operation of the internal layers of CNNs can help us to understand what characteristics are important for classification. For example, the feature maps can be visualised according to the contribution of time-frequency bins to classification. This idea is inspired by image processing tasks, including object localisation [29], and saliency detection [30].

Two difficulties emerge when visualising ASC systems. Firstly, the size of the feature maps decreases due to either stride convolution or local pooling layers. Low-resolution feature maps lose time-frequency details compared to the log mel spectrograms. Secondly, the classification result depends on global pooling operations, which summarise the feature maps into fixed dimensional vectors for later classification. Global max and average pooling layers have been widely used in ASC [24]. However, if the value of each bin is viewed as the contribution, the contributions of the bins in feature maps are over- or underestimated in the output of these two kinds of global pooling layers [24]. To overcome these difficulties, we propose to employ atrous CNNs and an attention mechanism for visualisation. Atrous CNNs can retain the size of the feature maps to be the same as that of the input, by using dilated convolutions [31]. Furthermore, an attention mechanism can estimate the contribution of time-frequency bins better than global max or average pooling [32]. Our proposed conditional training framework, which is based on atrous CNNs with attention, is called *Conditional Atrous CNNs with Attention* (CAA-Net).

The remainder of this article is structured as follows. Related work will be introduced in Section II. In Section III, we will describe the proposed CAA-Net, including the framework overview, conditional training, atrous CNNs, and global pooling. The database description, experimental setup and evaluation metrics, experimental results, and the visualisation of the CNNs will then be presented in Section IV. Finally, the conclusions and future work will be given in Section V.

II. RELATED WORK

Since acoustic scene data from different devices has different qualities and distributions, CNNs trained on different single-device data will differ from each other. To reduce the computational burden caused by training a separate CNN model for each device, researchers proposed a series of learning strategies. In transfer learning, supervised domain adaptation [33] trained separate CNN models for the source and target data. In multi-task

learning, CNNs shared the parameters of the low-level convolutional layers, and trained the high-level convolutional layers separately [34]. However, more convolutional layers lead to more parameters, so that the training procedure tends to be sub-optimised [35]. Joint learning with shared context was proposed to process multi-source data using a selective loss function [25], but it is challenging to analyse the data difference among multiple devices using joint learning [25]. Conditional CNNs were proposed for multi-pose face recognition [36], where the convolutional kernels of the conditional CNNs were sparsely activated using a conditional activation function in each convolutional layer. Similar to the idea of conditional CNNs, a deep learning model was conditioned with the speaker information which was fed into the model as an extra input to achieve speech generation [37], [38]. Inspired by these approaches [37], [38], we propose to condition CNNs with device information. Furthermore, multi-task learning was employed for multi-modal emotion recognition, and the predicted difficulty level of the samples during the learning process was fed back into the deep neural networks [39]. Inspired by this work, we propose multi-task conditional training to learn CNNs, predicting acoustic scene classes and device information.

Many CNN models result in low-resolution feature maps, which cannot describe the time-frequency details [24]. A basic CNN model without pooling layers can retain the size of the feature maps to be the same as that of the input at each convolutional layer. However, more convolutional operations than those in CNNs with local pooling result in sub-optimised training procedures and lower performances [24]. Some image processing approaches have investigated increasing the size of feature maps at the final convolutional layer to be the same as the input. For example, encoder-decoder CNNs [40] learnt internal representations by employing encoders and decoders to respectively down-sample and up-sample the feature maps. *Fully Convolutional Networks* (FCNs) [41] utilised deconvolutional layers to up-sample the low-resolution feature maps. Deconvolutional layers were also applied in the architecture of *Generative Adversarial Networks* (GANs), generating synthetic images which can be used to augment the training data and further improve classification performance [42]. However, these image processing methods mostly use strongly labelled data, where each pixel of the input image is labelled. The data in ASC is weakly labelled as each audio clip has only one label, rather than each frame being annotated. It is challenging to apply these image processing methods to preserve the same size as the input on the feature maps at the final convolutional layer in ASC [24]. In a separate study of image segmentation [31], atrous CNNs employed dilated convolutions to retain the size of the feature maps to be the same of that of the input at each layer, instead of preserving the size of the feature maps at the final layer only. Inspired by the success of atrous CNNs in image segmentation, we apply atrous CNNs to achieve pixel-wise visualisation of the feature maps for audio classification.

For a classification task, convolutional layers are often followed by global pooling layers [43]. However, conventional global pooling methods, such as global max and average pooling, cannot estimate the contribution of each time-frequency bin

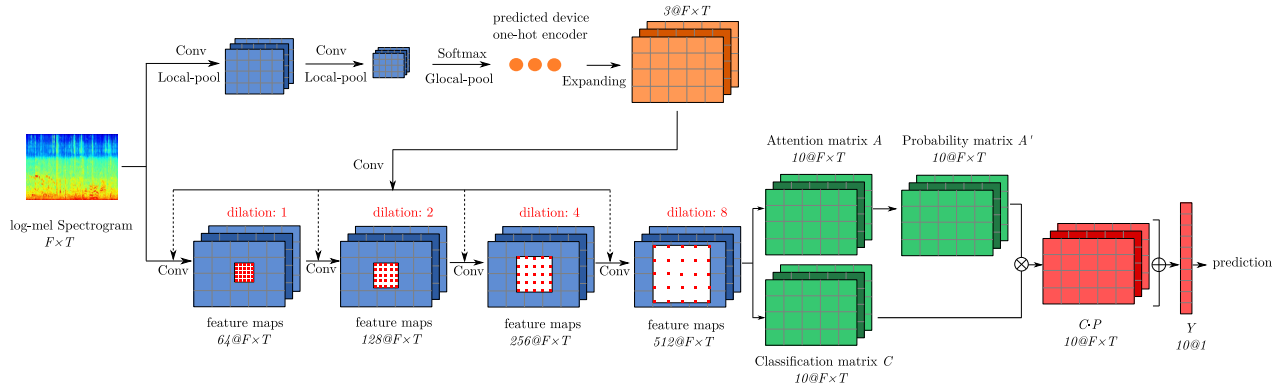


Fig. 1. The framework of our proposed multi-task CAA-Net. The log mel spectrograms are fed into two CNN models. The top CNN model predicts the device classes, and the bottom CNN model predicts the acoustic scenes. The predicted device information is transformed into a one-hot encoder, and is then expanded to a three dimensional tensor. Next, the convolutions of the tensors are fed into a convolutional layer of the bottom CNN model. Finally, a global attention layer is applied to the feature maps of the last convolutional layer. The dash lines indicate that the transformed device information is fed into one of the convolutional layers.

in the feature maps [24], [43]. Instead, weighted pooling was proposed to combine different convolutional layers by learning weighted masks, which solves this problem [44], [45]. Similarly, the attention mechanism was proposed for audio and speech-related tasks to estimate the contribution of the feature vectors at each time step [46]. Inspired by the attention mechanism idea, we propose to apply a global attention pooling layer to estimate the contribution of each time-frequency bin to the final classification. Finally, our proposed CAA-Net conditions the CNNs with device information, and applies dilated convolutions and an attention mechanism for visualising the intermediate layers of the CNNs.

III. PROPOSED METHODOLOGY

In this section, we will give an overview of the CAA-Net framework. Then, conditional CNNs will be proposed to adapt the CNNs to handle multi-device data. Finally, we will describe the atrous CNNs and the attention mechanism that we use to visualise the internal layers of the CNNs.

A. Framework Overview

The framework of our proposed multi-task CAA-Net is shown in Fig. 1. We extract log mel spectrograms from the multi-device audio waves. Then, the log mel spectrograms are used as input to train our proposed multi-task CAA-Net. CAA-Net contains two branches: one branch aims to predict acoustic scene classes, while the other branch predicts device classes. The first branch contains four convolutional layers and a global attention pooling layer. The four convolutional layers output feature maps with channels of dimensions 64, 128, 256, and 512. Notably, to preserve the same size of the feature maps as the log mel spectrograms, there is no local pooling following each convolutional layer. After the convolutional layers, a global attention pooling layer is applied to learn weight values for each time-frequency bin in the feature maps.

To process multi-device data, the CNN model in the first branch is conditioned by the predicted device information inferred from the CNN model in the second branch. The CNN

model in the second branch consists of two convolutional layers with channels of dimensions 64 and 128. Each convolutional layer is followed by a local pooling layer. A global max pooling layer is then used to summarise the feature maps to a fixed dimensional vector for classification. The predicted device classes are represented as one-hot encoders, quantifying the integer classes into binary vectors. The one-hot encoders are repeated along the time and frequency axes to form three-dimensional tensors, which are fed into the convolutional layers of the first branch. More details will be introduced in the next subsections.

B. Conditional Training

We now introduce four training strategies, including single-device training, joint training, teacher forcing conditional training, and multi-task conditional training.

1) *Single-Device Training*: Single-device training is an approach to train several separate ASC systems on the separate data from each device, as shown in Fig. 2(a). With single-device training, several CNN models can be learnt, on each single-device dataset as part of the multi-device dataset [43]. However, some datasets may be not big enough to train a deep learning model [47]. Additionally, single-device training requires many resources to train N systems for N -device data.

2) *Joint Training*: To overcome the shortcomings of single-device training, we jointly train the audio data from different devices, as shown in Fig. 2(b). We assume that data from multiple devices has common features that can be learnt in a CNN model [25], and that CNNs can learn complementary representations by jointly training on multi-device data. However, data from different devices share the same network structure and parameters with this approach, so that it is challenging to learn device-related features via joint learning [25].

3) *Teacher Forcing Conditional Training*: To learn device-related representations, the ground truth from a previous time step was fed into the current time step using conditional training in RNNs [48]. Conditional training has been applied to speech recognition [38], and speaker verification [49]. In our approach, the device information is forwarded to a convolutional layer, and

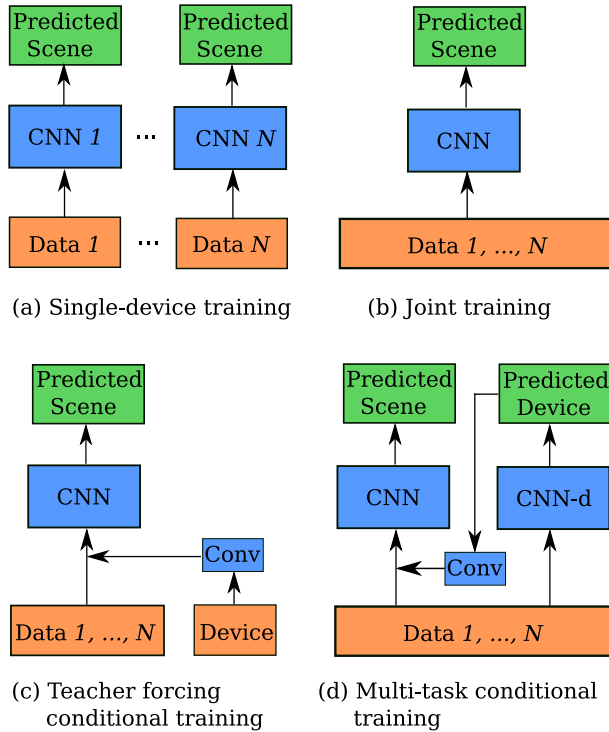


Fig. 2. Comparison of the four training strategies. In (d) ‘CNN-d’ means the CNN-device, which aims to predict the device classes.

then is fed into the classification branch, as shown in Fig. 2(c). In a similar way to a mask, the feature map learnt from the device information gives different weights to each time-frequency bin of the representations in the classification branch, filtering out the time-frequency bins of the feature maps which are not useful. In this regard, the representations computed from the device information are called as ‘masks’ in this article. We call this method teacher forcing conditional training.

The initial mask is generated by expanding the N -length device one-hot encoder into a matrix of size $N \times U_w \times U_l$, where $U_w \times U_l$ is the size of the feature maps in the CNN model of the classification branch. The expansion makes it possible to combine the device information with the feature maps. Referring to the linear transformation of the speaker one-hot encoder in a multi-speaker text-to-speech task [37], we then apply a convolutional layer with a kernel size of 1×1 , which is called a 1-by-1 convolution [43]. Notably, the output of the 1-by-1 convolution has the same number of channels as that of the convolutional layer in the classification branch, with which will be combined. The combined result is given by

$$M_{i+1} = \sigma(W_i * M_i + V_i * E_i), \quad (1)$$

where M_i is the input of the i -th convolutional layers, E_i is the expanded one-hot encoder, W_i and V_i are the weights of the convolutional procedures, $*$ denotes convolution, and σ is the ‘ReLU’ activation function.

4) *Multi-task Conditional Training*: In teacher forcing conditional training, the device information is required to be fed into the deep neural networks as prior knowledge. However, the device information is sometimes unknown, so we propose

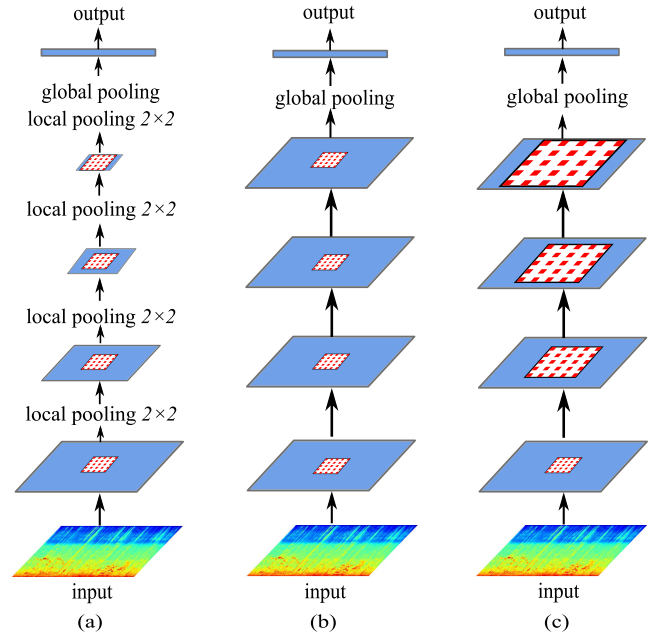


Fig. 3. Three CNN architectures: (a) CNNs with local pooling, (b) CNNs without local pooling, and (c) atrous CNNs.

to train a model to predict both the acoustic scenes and device classes simultaneously using multi-task learning [50], and we call this approach multi-task conditional training. As shown in Fig. 2(d), multi-task conditional training learns an additional CNN-device (CNN-d) model to predict the device classes, compared to teacher forcing conditional training. The structure of ‘CNN-d’ model was the same as that introduced in Section III-A. The local pooling following each convolutional layer has a kernel size of 2×2 to reduce the size of feature maps and accelerate the training procedure. The predicted device classes, represented as one-hot encoders, are expanded and processed using Eq. (1). In the training procedure, the loss function is computed as the weighted sum of the loss functions in the two CNN models, i. e.,

$$loss = loss_s + \lambda \times loss_d, \quad (2)$$

where $loss_s$, and $loss_d$ are the loss values of the CNN model for predicting scenes, and the CNN-d model for predicting devices, respectively. The weight factor λ can be used to balance the gradient difference between the two CNN models [51].

C. Atrous Convolutional Neural Networks

To classify the log mel spectrograms representing the audio waveforms, we use atrous CNNs, which can preserve the size of the feature maps to help visualise the internal layers of the CNNs [24].

1) *CNNs With Local Pooling*: In conventional CNNs, the convolutional layers are usually followed by local pooling layers. As shown in Fig. 3(a), each convolutional layer is followed by a local max pooling layer with a kernel size of 2×2 . Local max pooling is able to extract time-frequency shift-invariant features, and can accelerate the training procedure by reducing the size of the feature maps [52]. However, the local pooling means

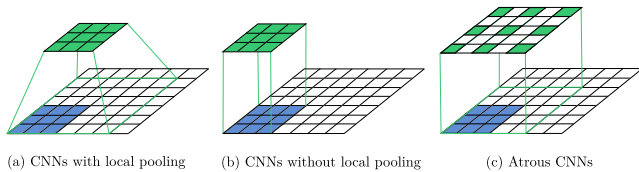


Fig. 4. Comparison of the receptive fields in the three CNN architectures. In each subfigure, the white grid at the bottom denotes the feature map in the first convolutional layer, and the blue grid means the convolutional kernel. The top green grid is the kernel in the second convolutional layer.

that the visualisation of the CNNs is limited to low-resolution feature maps.

2) *CNNs Without Local Pooling*: To solve the problem of low-resolution feature maps in CNNs with local pooling, the local max pooling layers are removed from the CNN model in Fig. 3(b). This is a basic approach to retain the size of the feature maps, that will lead to a performance reduction [24]. In our previous study [24], we found that the reason is related to the receptive field, the input area with a size of *number of frequency bins* \times *number of time frames* that can affect a single output pixel in the feature maps. For example, in Fig. 4(a-b), the two receptive fields in the first convolutional layer both have a size of $r \times r$, but in the second convolutional layer, the receptive field of the CNNs with local pooling has a size of $2r \times 2r$, while the receptive field of the CNNs without local pooling still has a size of $r \times r$. Further, at the l -th layer, $l \in \{1, \dots, L\}$ where L is the total number of convolutional layers, the size of the receptive field in the CNNs with local pooling will be $2^{l-1}r \times 2^{l-1}r$, while the CNNs without local pooling will have a fixed receptive field size of $r \times r$ at all layers. Hence, for CNNs with local pooling, the size of the receptive field increases exponentially with the number of layers; but for CNNs without local pooling, it does not change on the number of layers.

3) *Atrous CNNs*: To preserve the size of feature maps to be the same as that of the input, we employ atrous CNNs, increasing the size of the receptive field exponentially. In atrous CNNs, the convolutional kernel is replaced by a sparse kernel with holes, called a dilated kernel, the size of which is determined by the dilation rate. The dilation rate in each of the four convolutional layers is set to 1, 2, 4, and 8, so that the size of the receptive field increases exponentially, the same as in CNNs with local pooling. For example, in Fig. 4(c), the size of the receptive field in the second convolutional layer is $(2r - 1) \times (2r - 1)$, and $(2^{l-1}r - 1) \times (2^{l-1}r - 1)$ in the l -th layer. The convolutional operations in atrous CNNs are mostly less than those of the CNNs without local pooling.

D. Global Pooling Mechanism

We denote a log mel spectrogram as S of size $F \times T$, where F is number of the frequency bins and T is number of the time frames. The feature maps from the final convolutional layers of the CNNs are denoted as a three-dimensional tensor M of size $H \times P \times Q$, where H denotes the number of channels, and $P \times Q$ is the size of each feature map. The convolutional layers are then followed by a global pooling layer [53], which

summarises each of the H feature maps into a scalar value, which is stacked into an H -dimensional vector for classification. Global pooling performs better than flattening the feature maps into a vector, as global pooling can filter out the units from the feature maps which are not useful [43]. In this subsection, three conventional global pooling methods will be compared, and the proposed attention pooling will be introduced.

1) *Global Max Pooling*: With the assumption that the maximum value can represent the area of each kernel [54], global max pooling selects the maximum value from the kernel. Each kernel has the same size as the feature map. Global max pooling can be defined by

$$R_h = \max_{1 < q < Q} \max_{1 < p < P} M_{hpq}, \quad (3)$$

where R_h is the h -th value of the vector R which will be fed into a softmax layer for classification. Global max pooling has achieved success in many applications, such as image processing [55] and audio event detection [56]. However, global max pooling ignores the contribution of the units with smaller values than the maximum value. The lost information might lead to under-performance of the classification [43].

2) *Global Average Pooling*: Different from global max pooling, global average pooling treats every unit as having the same contribution for the classification [57]. Global average pooling computes the average value of each feature map by

$$R_h = \frac{1}{PQ} \sum_{q=1}^Q \sum_{p=1}^P M_{hpq}. \quad (4)$$

However, average pooling overestimates the contribution of some units which are not useful for the final classification, and underestimates the contribution of useful units [43]. Hence, global average pooling is also not suitable to accurately estimate the contribution of each unit.

3) *Global Region of Interest Pooling*: *Global Region of Interest (ROI) pooling* [58] first splits a feature map into several sub-areas, and then applies global max pooling on these sub-areas. The outputs from these sub-areas are then flattened into a vector for the final classification. Notably, while combining the ROI pooling and the global attention pooling, which will be introduced in the following paragraph, the outputs of the ROI pooling are fed into the attention mechanism instead of flattening, to learn the contribution of each time-frequency bin. To compare global ROI pooling methods to other pooling methods without considering the effect of hyper-parameters, the size of the sub-areas is set to 16×16 . Therefore, the feature maps after global ROI pooling have the same size as the feature maps from four local max pooling layers with a 2×2 kernel. Global ROI pooling can select more useful units from each feature map. However, like global max and average pooling, it cannot accurately estimate the contribution of each time-frequency bin [24].

4) *Global Attention Pooling*: Compared to the aforementioned pooling methods, the global attention pooling aims to reduce the dimensions of feature maps, by estimating the contribution of each unit [24]. Global attention pooling contains two branches: one is a classification branch, and the other is an attention branch. Global attention pooling can classify the feature

maps without a final linear transformation. To achieve the classification in attention pooling, we employ a 1-by-1 convolutional layer in each branch, and the channel number of the output is equal to the class number. Further, in the classification branch, a softmax activation function follows the 1-by-1 convolutional layer for classification. In the attention branch, a sigmoid function is employed. The sigmoid function transforms the values in the feature maps into the interval $[0, 1]$, representing the weight value of each time-frequency bin. In the attention branch, the obtained attention matrix A is normalised using

$$A'_{kpq} = A_{kpq} / \sum_{q=1}^Q \sum_{p=1}^P A_{kpq}, \quad (5)$$

where k means the k -th class, and A' denotes the probability matrix, holding the weight of each element in the H feature maps C from the classification branch. Finally, the prediction Y is computed by

$$Y_k = \sum_{q=1}^Q \sum_{p=1}^P A'_{kpq} \cdot C_{kpq}. \quad (6)$$

Notably, the matrices A' and C have the same size, and are multiplied using the element-wise product. In CNNs without local pooling and in atrous CNNs, the feature maps produced by the convolutional layers have the same size as the log mel spectrograms, i.e., $P = F$ and $Q = T$.

IV. EXPERIMENTS AND RESULTS

A. Database

We evaluate our proposed CAA-Net on the open dataset in the ASC task of the IEEE AASP Challenge on Detection and Classification of Acoustic Scenes and Events (DCASE) 2018 and 2019 [23]. We test the development set only, as the labels of the test set are not publicly released. The DCASE 2018 database was recorded in ten scene environments of six large European cities. The ten acoustic scenes include: *airport*, *shopping mall*, *metro station*, *street pedestrian*, *public square*, *street traffic*, *tram*, *bus*, *metro*, and *park*. Three devices were employed to record the audio waves simultaneously. The main recording device (called device A) comprises a Soundman OKM II Klassik/studio A3, an electret binaural in-ear microphone, and a Zoom F8 audio recorder. The other two recording devices are mobile phones: a Samsung Galaxy S7 (called device B), and an iPhone SE (called device C). During recording, device A was worn in the ears, device B was held in a hand, and device C was worn in a sleeve of the strap of a backpack. The sampling rate of device A was 48 kHz, and the sampling rates of device B and C were both 44.1 kHz. The original recordings were split into several audio clips with a length of 10 seconds. Finally, the development set consists of 8,640 audio files from device A, and 720 audio files from device B and C in parallel. The ten scene classes are balanced in the whole development set, so that each scene consists of 864 audio clips from device A, 72 clips from device B, and 72 clips from device C. The development set was officially split into a training subset and a test subset by the organisers of DCASE

2018. The training subset contains 6,122 audio clips from device A, 540 clips from device B, and 540 clips from device C. The test subset contains 2,518 audio clips from device A, 180 clips from device B, and 180 clips from device C.

The DCASE 2019 dataset was recorded in the same ten scene environments as the DCASE 2018 dataset, and extended from six cities to 12 cities. The three devices and data recording procedures were the same as those in DCASE 2018. Audio clips with a length of 10 seconds were obtained from the original recordings. The development set contains 16,560 audio clips in total (10,265 audio clips in the training subset, 5,265 clips in the test subset, and 1,030 clips from Milan). The training subset consists of 9,185 audio clips from device A, and 540 clips from devices B and C in parallel. The test subset contains 4,185 audio clips from device A, and 540 clips from devices B and C in parallel. As the data from Milan was not split into the training or test subsets due to data balance issue in the challenge, only the official training and test subsets are used in our experiments.

B. Experimental Setup and Evaluation Metrics

The audio clips are first resampled to 44.1 kHz. Then, the log mel spectrograms are extracted from the audio clips using a Hamming window of 2,048 samples length, an overlap of 672, and 64 mel bins. Hence, log mel spectrograms with a size of 64×320 are obtained, where 320 is the number of time frames. The log mel spectrograms are then fed into the CNN model for classification. The CNN model is optimised using an Adam optimiser with a learning rate of 0.001. During training, the learning rate is reduced by a factor of 0.9 at each 200-th iteration. Finally, the training procedure is stopped at the 15,000-th iteration. Notably, in the multi-task conditional training, the λ value in Eq. (2) is set to 1 and 0.0001 empirically before and after the accuracy of the ‘CNN-d’ on the test subset achieves 98%, respectively. The setting of λ aims to obtain a robust ‘CNN-d’ before training the CNN model for the ASC task. All of these training hyperparameters are chosen empirically. The source code of our work is publicly released.¹

For the evaluation of the performance, we use the unweighted average of class-wise accuracies, which are the number of correctly classified clips divided by the total number of clips in each class. Accuracy is the official DCASE evaluation metrics for ASC.

C. Results and Discussion

We tested the performance of the three CNN models introduced in Section III-C on the DCASE 2018 dataset, while conditioning the device information at four convolutional layers using teacher forcing conditional training. A set of global pooling methods introduced in Section III-D were applied for each CNN model. The performances of the CNN models on the data from the three devices are compared in Fig. 5. In Fig. 5, the performances of the CNN models on device A are better than those on devices B and C, which might be caused by the data imbalance

¹[Online]. Available: https://github.com/EIHW/CAANet_DCASE_ASC

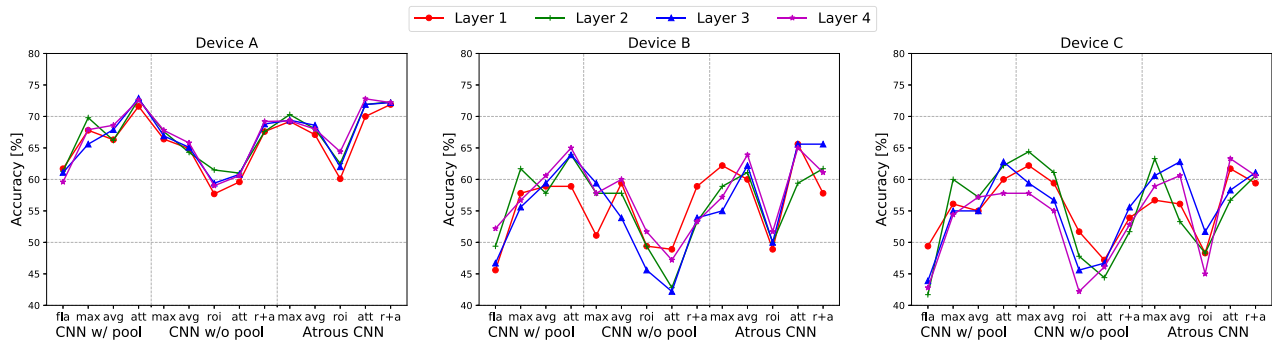


Fig. 5. Performance (accuracy) comparison of the CNN models evaluated on the DCASE 2018 dataset, while teacher forcing conditional training works at different convolutional layers. The three CNN topologies contain CNNs with local pooling, CNNs without local pooling, and atrous CNNs. The CNNs are followed by flattening ('fla') and five global pooling models, including max, average ('avg'), ROI, attention ('att'), and the combination of ROI and attention ('r+a'). The performance is evaluated on the data from the three devices A, B, and C.

TABLE I

THE AVERAGE PERFORMANCE (ACCURACY) COMPARISON OF THE CNN TOPOLOGIES EVALUATED ON THE DATA FROM THE THREE DEVICES (A, B, AND C) AVAILABLE IN THE DCASE 2018 DATASET, WHILE THE TEACHER FORCING CONDITIONAL TRAINING WORKS AT DIFFERENT CONVOLUTIONAL LAYERS. THE THREE TYPES OF CNN TOPOLOGIES CONTAIN CNNs WITH LOCAL POOLING, CNNs WITHOUT LOCAL POOLING, AND ATRous CNNs. THE CNNs ARE FOLLOWED BY FLATTENING AND FIVE GLOBAL POOLING MODELS, INCLUDING MAX, AVERAGE ('AVG'), ROI, ATTENTION ('ATT'), AND THE COMBINATION OF ROI AND ATTENTION ('ROI+ATT'). THE BEST RESULT CHOSEN FROM FOUR LAYERS IN EACH CNN MODEL IS HIGHLIGHTED

Accuracy [%]		Layer 1	Layer 2	Layer 3	Layer 4
CNN w/ pool	flatten	52.2	50.8	50.6	51.5
CNN w/ pool	max	60.6	63.8	58.7	59.7
CNN w/ pool	avg	60.1	60.4	60.8	62.1
CNN w/ pool	att	63.5	66.3	66.5	65.1
CNN w/o pool	max	59.9	63.3	61.9	61.1
CNN w/o pool	avg	61.2	61.1	58.6	60.3
CNN w/o pool	roi	52.9	52.9	50.2	51.0
CNN w/o pool	att	51.9	49.4	49.9	51.3
CNN w/o pool	roi+att	60.1	57.5	59.4	58.4
Atrous CNN	max	62.7	64.2	61.7	61.8
Atrous CNN	avg	61.1	60.8	64.5	64.2
Atrous CNN	roi	52.4	53.6	54.6	53.7
Atrous CNN	att	65.8	62.7	65.3	67.0
Atrous CNN	roi+att	63.0	64.9	66.3	64.6

among the data from devices A, B, and C, and the potential existence of more noise in the data from the mobile devices (B and C). While comparing the performances of all CNN models in each device, the atrous CNNs mostly perform better than CNNs with or without local pooling. The results from conditioning the device information at different convolutional layers are similar on each CNN model.

To better compare the performance of all CNN models and choose the best convolutional layer to achieve multi-task conditional training by each CNN model, the average performances on the data of the three devices are calculated in Table I. As shown in Table I, in the CNNs with local pooling and atrous CNNs, we can obtain better results from global attention pooling than from other global pooling methods. This indicates that attention pooling is suitable to estimate the contribution of each time-frequency bin in the feature maps to the final classification. Further, conditioning the

device information at different layers can affect the classification accuracy. Especially, in CNNs without local pooling, most models perform best when the transformed device information is fed into the first convolutional layer. In atrous CNNs, the models in which the device information was conditioned at the third convolutional layer produce better results than those conditioned in other layers. We select the best results from the four conditioned layers for later comparison of the four training strategies in Section III-B.

In Table II, teacher forcing conditional training and multi-task conditional training are compared to the other two training strategies, including single-device training and joint training. The results of multi-task conditional training were obtained while conditioning in the same convolutional layers as teacher forcing conditional training. The two conditional training strategies perform the best in most CNN models, except global attention and global ROI, and attention pooling in CNNs without local pooling. This may be caused by the large number of convolutional operations in those two global pooling methods in CNNs without local pooling. Teacher forcing conditional training performs slightly better than multi-task conditional training. This is reasonable, as the true device information is fed into the CNNs during teacher forcing conditional training. For devices B and C, the performance of joint training is comparable with that of single-device training. This indicates that joint training can learn complementary features for each single-device data. In contrast, conditional training can learn not only complementary features, but also device-related representations. Hence, in our results, conditional training can effectively improve the performance on devices B and C. Finally, our best result of 68.0% on average was obtained by multi-task conditional training. It is significantly higher than the result of 65.0% from the atrous CNNs with attention on a single device (in a one-tailed z-test, $p < 0.01$), and the result of 49.0% from the CNNs without local pooling with ROI pooling on a single device (in a one-tailed z-test, $p < 0.001$).

Finally, the confusion matrices of the best results obtained by our proposed multi-task CAA-Net are shown in Fig. 6. The proposed model performs the best on the data from device A, perhaps because the database contains more data for device A than B and C, and the data recorded with device A has a higher quality than the other data. The CAA-Net performs well in some

TABLE II

PERFORMANCE (ACCURACY) COMPARISON OF THE CNN TOPOLOGIES ASSESSED ON THE DCASE 2018 DATASET USING FOUR TRAINING STRATEGIES, INCLUDING SINGLE-DEVICE TRAINING, JOINT TRAINING, TEACHER FORCING CONDITIONAL TRAINING, AND MULTI-TASK CONDITIONAL TRAINING. THE POOLING METHODS ARE THE SAME AS THE METHODS IN TABLE I. THE PERFORMANCE IS EVALUATED ON THE DATA FROM THE THREE DEVICES A, B, AND C. THE BEST RESULT AMONG THE FOUR TRAINING STRATEGIES FOR EACH CNN MODEL IS HIGHLIGHTED

ACCURACY [%]		SINGLE-DEVICE TRAINING				JOINT TRAINING [32]				CONDITIONAL TRAINING							
Network	Pooling	A	B	C	Avg	A	B	C	Avg	TEACHER FORCING				MULTI-TASK			
										A	B	C	Avg	A	B	C	Avg
CNN w/ pool	flatten	61.1	49.4	40.6	50.4	61.6	49.4	46.7	52.6	61.7	45.6	49.4	52.2	63.5	52.8	50.6	55.6
CNN w/ pool	max	70.0	57.8	56.1	61.3	69.8	57.2	57.8	61.6	69.8	61.7	60.0	63.8	68.0	58.3	57.8	61.4
CNN w/ pool	avg	63.6	50.0	62.8	58.8	65.8	57.2	57.8	60.3	68.6	60.6	57.2	62.1	65.4	60.0	58.3	59.7
CNN w/ pool	att	73.1	51.1	58.3	60.8	72.6	62.2	56.1	63.6	72.9	63.9	62.8	66.5	72.6	60.6	59.4	64.2
CNN w/o pool	max	67.8	53.9	56.1	59.3	61.9	46.7	52.2	53.6	67.6	57.8	64.4	63.3	65.4	56.7	57.2	59.8
CNN w/o pool	avg	61.9	61.7	42.8	55.5	59.1	54.4	50.0	54.5	64.9	59.4	59.4	61.2	66.1	61.7	55.4	61.1
CNN w/o pool	roi	60.4	43.3	43.3	49.0	61.7	50.6	43.9	52.1	57.7	49.4	51.7	52.9	60.7	50.0	48.9	53.2
CNN w/o pool	att	62.7	56.7	60.0	59.8	59.6	45.0	43.3	49.3	59.6	48.9	47.2	51.9	59.2	46.1	49.4	51.6
CNN w/o pool	roi+att	67.8	61.1	60.0	63.0	69.2	56.1	50.6	58.6	67.6	58.9	53.9	60.1	67.5	58.3	53.3	59.7
Atrous CNN	max	66.2	60.6	42.8	56.5	69.7	60.0	59.4	63.0	70.3	58.9	63.3	64.2	67.0	61.1	58.9	62.3
Atrous CNN	avg	69.5	61.1	55.6	62.1	67.2	62.8	60.0	63.3	68.6	62.2	62.8	64.5	66.9	56.1	55.0	59.3
Atrous CNN	roi	61.2	45.6	40.6	49.1	62.6	48.3	43.9	51.6	62.0	50.0	51.7	54.6	63.5	47.2	45.0	51.9
Atrous CNN	att	70.7	61.1	63.3	65.0	73.2	64.4	62.2	66.6	72.8	65.0	63.3	67.0	72.4	68.9	62.8	68.0
Atrous CNN	roi+att	72.7	58.9	49.4	60.3	72.2	57.2	56.7	62.0	72.2	65.6	61.1	66.3	72.2	56.7	62.2	63.7

TABLE III

PERFORMANCE (ACCURACY) COMPARISON OF THE CNN TOPOLOGIES ASSESSED ON THE DCASE 2019 DATASET IN THE FOUR TRAINING STRATEGIES, INCLUDING SINGLE-DEVICE TRAINING, JOINT TRAINING, TEACHER FORCING CONDITIONAL TRAINING, AND MULTI-TASK CONDITIONAL TRAINING. THE POOLING METHODS ARE THE SAME AS THE METHODS IN TABLE I. THE PERFORMANCE IS EVALUATED ON THE DATA FROM THE THREE DEVICES A, B, AND C. THE BEST RESULT AMONG THE FOUR TRAINING STRATEGIES FOR EACH CNN MODEL IS HIGHLIGHTED

ACCURACY [%]		SINGLE-DEVICE TRAINING				JOINT TRAINING				CONDITIONAL TRAINING							
Network	Pooling	A	B	C	Avg	A	B	C	Avg	TEACHER FORCING				MULTI-TASK			
										A	B	C	Avg	A	B	C	Avg
CNN w/ pool	att	67.2	47.0	47.0	53.7	69.9	44.1	49.8	54.6	70.9	46.7	48.1	55.2	70.9	46.1	49.1	55.4
CNN w/o pool	att	60.9	48.9	48.3	52.7	63.5	36.9	36.5	45.6	65.0	50.6	56.3	57.3	63.3	43.3	46.3	51.0
Atrous CNN	att	72.2	45.4	50.2	55.9	71.4	44.3	50.9	55.5	72.0	43.7	51.1	55.6	71.3	44.4	52.6	56.1

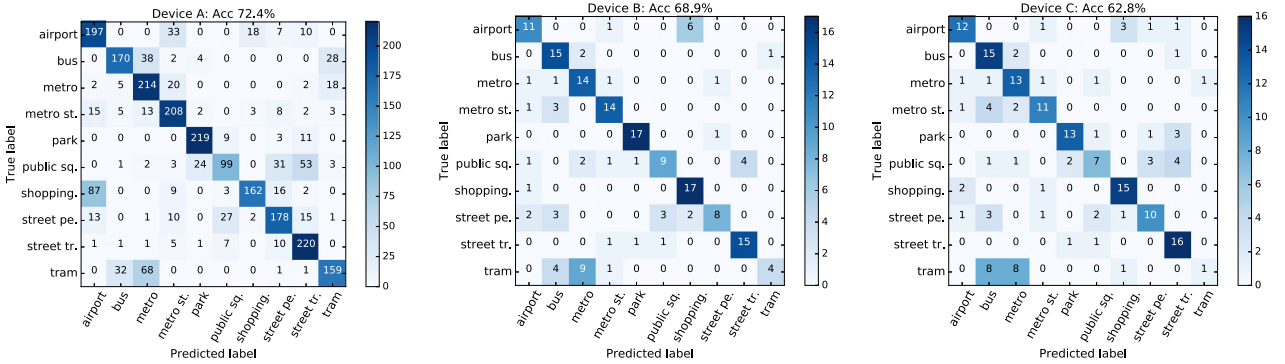


Fig. 6. Confusion matrices of the results on the data from devices A, B, and C computed by the proposed multi-task CAA-Net model, which achieves the best result on the DCASE 2018 dataset.

classes, such as *metro*, *park*, and *street traffic*, but for not well for other classes, including *public square*, and *street pedestrian*. The reason may be that the audio recordings in these two environments contain more noise. Furthermore, some classes of scenes are easy to be predicted as other similar scenes. For instance, the scene class *tram* is mostly predicted into the classes of *bus* and *metro*, especially on the data from devices B and C, perhaps because the audio recordings from these three scenes are all traffic sounds. This suggests that besides the types of devices, the type of acoustic scene might also be a factor to be considered in this classification task.

To verify the robustness of our proposed approach, the three CNN models with attention were trained on the DCASE 2019 dataset, as shown in Table III. For conditional training, the CNN layers corresponding to those highlighted in Table I were conditioned in the CNN models to predict the acoustic scenes. We can see that, in a similar way to the work on the DCASE 2018 dataset, conditional training achieves the best results in the three CNN models assessed on the DCASE 2019 dataset. The consistent results indicate that our proposed approach is more effective and robust than single-device training and joint training on multiple datasets.

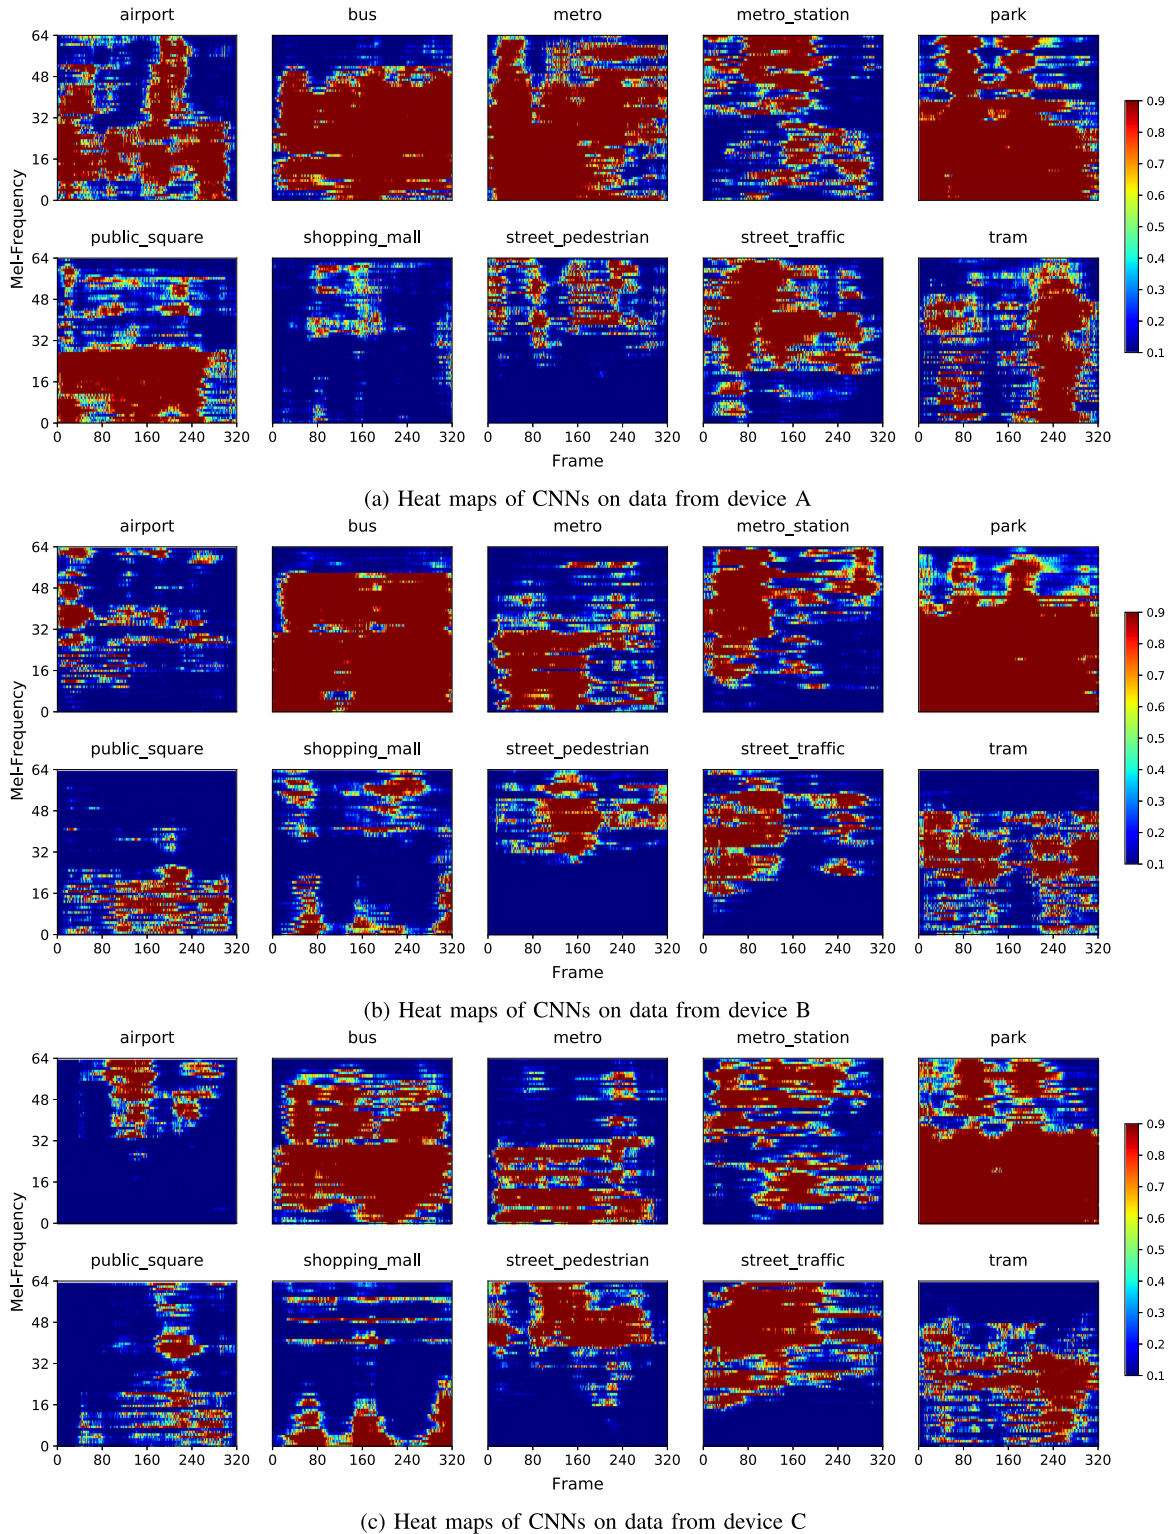


Fig. 7. Heat maps with a size of 64×320 are obtained by visualising the attention matrix A in our multi-task CAA-Net which works on DCASE 2018 dataset. The horizontal and vertical axes represent the time frames and frequency bins, respectively.

D. Visualisation of CNNs

To visualise the internal convolutional layers of the CNNs, heat maps of the attention matrices in our proposed multi-task CAA-Net are shown in Fig. 7. The heat maps are obtained from

the learnt attention matrices in each class of acoustic scene recorded at the same time and place by the three devices. We can see that, the learnt attention matrices are mostly similar across the three devices. For some classes, such as *airport*, *metro*, *public square*, and *tram*, more time-frequency bins in the heat maps

of device A have high values than in those of devices B and C. This may be caused by the fact that the data qualities differ from each other, as the recordings of the database were performed using three devices. Device A, as an in-ear device, has less noise than devices B and C, which both are mobile devices. The heat maps can reflect features of some special acoustic scenes. For instance, the heat maps of traffic scenes, such as *bus*, *metro*, and *tram*, contain time-continuity bins, which is consistent with the characteristics of traffic scenes. In *airport*, *metro station*, *park*, *shopping mall*, *street pedestrian*, and *street traffic* classes, there are some non-traffic sounds like speech, bird sound, etc. Hence, the time-frequency bins in these scenes concentrate on the high-frequency areas.

Visualisation of heat maps can provide a way to explain the decisions of CNNs in real-life applications, particularly in the security field. To ensure the machine is able to make an understandable and correct decision, it is essential to obtain an explanation of the decisions from the classifiers. Further, the visualisation technique offers the potential to improve the performance of other acoustic tasks such as in sound event detection, as the contribution weights of the heat maps might be related to the time and frequency of sound events.

V. CONCLUSIONS AND FUTURE WORK

We proposed multi-task conditional atrous *Convolutional Neural Networks* (CNNs) with attention, to visualise and understand acoustic scenes recorded with multiple devices. In the proposed model, log mel spectrograms were fed into two CNN models to predict device classes and acoustic scene classes. The predicted device classes were represented by one-hot encoders. The CNNs for acoustic scene classification were conditioned by feeding the transformation of device one-hot encoders into the convolutional layers. Four dilated convolutional layers, followed by a global attention pooling layer, were utilised to preserve the size of the feature maps for visualisation. Our proposed conditional training performed better than single-device training and joint training. Further, the intermediate layers of the CNNs were visualised and analysed in our work.

In future work, as the fixed weight of the loss function (the value of λ) is not friendly to optimise while training a neural network, and not flexible to adapt to a new dataset, we will explore learning the value of λ during the training procedure, to obtain a more robust and flexible model able to work on different datasets. Due to the effect of type of scenes as aforementioned, a hierarchical structure to classify both the type of scenes and device classes will also be considered. Next, CNNs trained on one dataset which is collected with three concrete devices are challenging to be applied to datasets recorded by other devices. Therefore, training of cross-database CNN models will be considered, so that the data from more devices can be included. The computation complexity of our conditional training approach will be a burden if the data is recorded by many devices (more than three) in practice. In this regard, efficient training strategies such as transfer learning [59], and lifelong learning [60] are promising to achieve an efficient training procedure. We will also attempt to utilise the learnt heat maps for other acoustic tasks, such as audio event detection.

REFERENCES

- [1] Q. Kun *et al.*, "Teaching machines on snoring: A benchmark on computer audition for snore sound excitation localisation," *Arch. Acoust.*, vol. 43, no. 3, pp. 465–475, Feb. 2018.
- [2] D. Barchiesi, D. Giannoulis, D. Stowell, and M. D. Plumbley, "Acoustic scene classification: Classifying environments from the sounds they produce," *IEEE Signal Process. Mag.*, vol. 32, no. 3, pp. 16–34, Apr. 2015.
- [3] D. Stowell, D. Giannoulis, E. Benetos, M. Lagrange, and M. D. Plumbley, "Detection and classification of acoustic scenes and events," *IEEE Trans. Multimedia*, vol. 17, no. 10, pp. 1733–1746, Oct. 2015.
- [4] K. Qian, Z. Ren, V. Pandit, Z. Yang, Z. Zhang, and B. Schuller, "Wavelets revisited for the classification of acoustic scenes," in *Proc. DCASE*, Munich, Germany, 2017, pp. 108–112.
- [5] Z. Yang, K. Qian, Z. Ren, A. Baird, Z. Zhang, and B. Schuller, "Learning multi-resolution representations for acoustic scene classification via neural networks," in *Proc. CSMT*, Harbin, P.R. China, 2019, pp. 133–143.
- [6] J. Geiger, B. Schuller, and G. Rigoll, "Large-scale audio feature extraction and SVM for acoustic scene classification," in *Proc. Workshop Appl. Signal Process. Audio Acoust.*, New Paltz, NY, 2013, pp. 1–4.
- [7] S. Hershey *et al.*, "CNN architectures for large-scale audio classification," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, New Orleans, LA, 2017, pp. 131–135.
- [8] Z. Ren, K. Qian, Z. Zhang, V. Pandit, A. Baird, and B. Schuller, "Deep scalogram representations for acoustic scene classification," *IEEE/CAA J. Automatica Sinica*, vol. 5, no. 3, pp. 662–669, May 2018.
- [9] C. Perera, A. Zaslavsky, P. Christen, and D. Georgakopoulos, "Context aware computing for the internet of things: A survey," *IEEE Commun. Surveys Tut.*, vol. 16, no. 1, pp. 414–454, May 2013.
- [10] A. Harma, J. Jakka, M. Tikander, M. Karjalainen, T. Lokki, and H. Nironen, "Techniques and applications of wearable augmented reality audio," in *Proc. Audio Eng. Soc. Conv.*, 2003, pp. 1–20.
- [11] E. Martinson and A. Schultz, "Robotic discovery of the auditory scene," in *Proc. ICRA*, Roma, Italy, 2007, pp. 435–440.
- [12] T. Hayashi, S. Watanabe, T. Toda, T. Hori, J. Roux, and K. Takeda, "Bidirectional LSTM-HMM hybrid system for polyphonic sound event detection," in *Proc. Detection Classification Acoust. Scenes Events Workshop*, Budapest, Hungary, 2016, pp. 35–39.
- [13] S. Amiriparian, M. Freitag, N. Cummins, and B. Schuller, "Sequence to sequence autoencoders for unsupervised representation learning from audio," in *Proc. Detection Classification Acoust. Scenes Events Workshop*, Munich, Germany, 2017, pp. 17–21.
- [14] H. Phan, P. Koch, F. Katzberg, M. Maass, R. Mazur, and A. Mertins, "Audio scene classification with deep recurrent neural networks," in *Proc. INTERSPEECH*, Stockholm, Sweden, 2017, pp. 3043–3047.
- [15] N. Koluguri, G. Meenakshi, and P. Ghosh, "Spectrogram enhancement using multiple window Savitzky-Golay MWSG filter for robust bird sound detection," *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 25, no. 6, pp. 1183–1192, May 2017.
- [16] Z. Zhao *et al.*, "Exploring deep spectrum representations via attention-based recurrent and convolutional neural networks for speech emotion recognition," *IEEE Access*, vol. 7, pp. 97 515–97 525, Jul. 2019.
- [17] J. Salamon and J. Bello, "Deep convolutional neural networks and data augmentation for environmental sound classification," *IEEE Signal Process. Lett.*, vol. 24, no. 3, pp. 279–283, Jan. 2017.
- [18] P. Primus, H. Eghbal-zadeh, D. Eitelsebner, K. Koutini, A. Arzt, and G. Widmer, "Exploiting parallel audio recordings to enforce device invariance in CNN-based acoustic scene classification," in *Proc. Detection Classification Acoust. Scenes Events Workshop*, 2019, pp. 204–208.
- [19] P. Brezina and S. Jeseničová, "Sound recording technologies and music education," *Ad Alta: J. Interdiscipl. Res.*, vol. 8, no. 2, pp. 13–18, 2018.
- [20] A. Goyal, S. Shukla, and R. Sarin, "Identification of source mobile hand sets using audio latency feature," *Forensic Sci. Int.*, vol. 298, pp. 332–335, May 2019.
- [21] T. Nguyen and F. Pernkopf, "Acoustic scene classification with mismatched recording devices using mixture of experts layer," in *Proc. ICME*, Shanghai, China, 2019, pp. 1666–1671.
- [22] A. Kumar, M. Khadkevich, and C. Fügen, "Knowledge transfer from weakly labeled audio using convolutional neural network for sound events and scenes," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Calgary, Canada, 2018, pp. 326–330.
- [23] A. Mesaros, T. Heittola, and T. Virtanen, "A multi-device dataset for urban acoustic scene classification," in *Proc. Detection Classification Acoust. Scenes Events Workshop*, Surrey, U.K., 2018, pp. 9–13.

- [24] Z. Ren, Q. Kong, K. Qian, M. D. Plumbley, and B. Schuller, "Attention-based convolutional neural networks for acoustic scene classification," in *Proc. Detection Classification Acoust. Scenes Events Workshop*, Surrey, U.K., 2018, pp. 39–43.
- [25] D. Fourure, R. Emonet, E. Fromont, D. Muselet, N. Neverova, A. Trémeau, and C. Wolf, "Multi-task, multi-domain learning: Application to semantic segmentation and pose regression," *Neurocomputing*, vol. 251, pp. 68–80, Aug. 2017.
- [26] D. Snyder, P. Ghahremani, D. Povey, D. Garcia-Romero, Y. Carmiel, and S. Khudanpur, "Deep neural network-based speaker embeddings for end-to-end speaker verification," in *Proc. SLT*, San Diego, CA, 2016, pp. 165–170.
- [27] N. Hojo, Y. Ijima, and H. Mizuno, "DNN-based speech synthesis using speaker codes," *IEICE Trans. Inf. Syst.*, vol. 101, no. 2, pp. 462–472, Feb. 2018.
- [28] Z. Ren, N. Cummins, V. Pandit, J. Han, K. Qian, and B. Schuller, "Learning image-based representations for heart sound classification," in *Proc. DH*, Lyon, France, 2018, pp. 143–147.
- [29] Z. Shi, T. Hospedales, and T. Xiang, "Bayesian joint modelling for object localisation in weakly labelled images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 10, pp. 1959–1972, Oct. 2015.
- [30] Y. Gao, M. Shi, D. Tao, and C. Xu, "Database saliency for fast image retrieval," *IEEE Trans. Multimedia*, vol. 17, no. 3, pp. 359–369, Jan. 2015.
- [31] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," 2017 *arXiv:1706.05587*, 2017.
- [32] Z. Ren, Q. Kong, J. Han, M. D. Plumbley, and B. Schuller, "Attention-based atrous convolutional neural networks: Visualisation and understanding perspectives of acoustic scenes," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Brighton, U.K., 2019, pp. 56–60.
- [33] S. Motiian, M. Piccirilli, D. Adjeroh, and G. Doretto, "Unified deep supervised domain adaptation and generalization," in *Proc. Int. Conf. Comput. Vis.*, Venice, Italy, 2017, pp. 5715–5725.
- [34] V. Sindagi and V. Patel, "CNN-based cascaded multi-task learning of high-level prior and density estimation for crowd counting," in *Proc. AVSS*, Lecce, Italy, 2017, pp. 1–6.
- [35] K. He and J. Sun, "Convolutional neural networks at constrained time cost," in *Proc. Conf. Comput. Vis. Pattern Recognit.*, Boston, MA, 2015, pp. 5353–5360.
- [36] C. Xiong, X. Zhao, D. Tang, K. Jayashree, S. Yan, and T.-K. Kim, "Conditional convolutional neural network for modality-aware face recognition," in *Proc. Int. Conf. Comput. Vis.*, Santiago, Chile, 2015, pp. 3667–3675.
- [37] A. Gibiansky *et al.*, "Deep voice 2: Multi-speaker neural text-to-speech," in *Proc. Conf. Neural Inf. Proc. Syst.*, Long Beach, CA, 2017, pp. 2962–2970.
- [38] A. Oord *et al.*, "WaveNet: A generative model for raw audio," in *ISCA Speech Synth. Workshop*, Sunnyvale, CA, 2016, 15 pages.
- [39] Z. Zhang, J. Han, E. Coutinho, and B. Schuller, "Dynamic difficulty awareness training for continuous emotion prediction," *IEEE Trans. Multimedia*, vol. 21, no. 5, pp. 1289–1301, May 2019.
- [40] V. Vukotić, S.-L. Pintea, C. Raymond, G. Gravier, and J. Gemert, "One-step time-dependent future video frame prediction with a convolutional encoder-decoder neural network," in *Proc. Int. Conf. Image Anal. Process.*, Catania, Italy, 2017, pp. 140–151.
- [41] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Med. Image Comput. Computer-Assisted Interv.*, Munich, Germany, 2015, pp. 234–241.
- [42] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification," *Neurocomputing*, vol. 321, pp. 321–331, 2018.
- [43] Z. Ren, V. Pandit, K. Qian, Z. Yang, Z. Zhang, and B. Schuller, "Deep sequential image features on acoustic scene classification," in *Proc. Detection Classification Acoust. Scenes Events Workshop*, Munich, Germany, 2017, pp. 113–117.
- [44] Q. Hu, H. Wang, T. Li, and C. Shen, "Deep CNNs with spatially weighted pooling for fine-grained car recognition," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 11, pp. 3147–3156, Nov. 2017.
- [45] D. Li, T. Yao, L.-Y. Duan, T. Mei, and Y. Rui, "Unified spatio-temporal attention networks for action recognition in videos," *IEEE Trans. Multimedia*, vol. 21, no. 2, pp. 416–428, Feb. 2018.
- [46] J. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Proc. Conf. Neural Inf. Proc. Syst.*, Montreal, Canada, 2015, pp. 577–585.
- [47] M. Valenti, S. Squartini, A. Diment, G. Parascandolo, and T. Virtanen, "A convolutional neural network approach for acoustic scene classification," in *Proc. Int. Joint Conf. Neural Networks*, 2017, pp. 1547–1554.
- [48] A. Goyal, A. Lamb, Y. Zhang, S. Zhang, A. Courville, and Y. Bengio, "Professor forcing: A new algorithm for training recurrent networks," in *Proc. NIPS*, Barcelona, Spain, 2016, pp. 4601–4609.
- [49] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification," in *Proc. INTERSPEECH*, Stockholm, Sweden, 2017, pp. 999–1003.
- [50] C. Hawthorne, E. Elsen, J. Song, A. Roberts, I. Simon, C. Raffel, J. Engel, S. Oore, and D. Eck, "Onsets and frames: Dual-objective piano transcription," in *Proc. Int. Soc. Music Inf. Retrieval*, Paris, France, 2018, pp. 50–57.
- [51] R. Cipolla, Y. Gal, and A. Kendall, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *Proc. CVPR*, Salt Lake City, UT, 2018, pp. 7482–7491.
- [52] H. Phan, F. Andreotti, N. Cooray, Y. O. Chèn, and M. De Vos, "DNN filter bank improves 1-max pooling CNN for single-channel EEG automatic sleep stage classification," in *Proc. Eng. Med. Biol. Conf.*, Honolulu, HI, 2018, pp. 453–456.
- [53] K. He, X. Zhang, S. Ren, and J. Sun, "Spatial pyramid pooling in deep convolutional networks for visual recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 9, pp. 1904–1916, Jan. 2015.
- [54] J. Amores, "Multiple instance classification: Review, taxonomy and comparative study," *Artif. Intell.*, vol. 201, pp. 81–105, Aug. 2013.
- [55] G. Tolia, R. Slicre, and H. Jégou, "Particular object retrieval with integral max-pooling of CNN activations," in *Proc. Int. Conf. Learn. Representations*, San Juan, Puerto Rico, 2015, pp. 1–12.
- [56] T.-W. Su, J.-Y. Liu, and Y.-H. Yang, "Weakly-supervised audio event detection using event-specific gaussian filters and fully convolutional networks," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, New Orleans, LA, 2017, pp. 791–795.
- [57] X. Xu, "Statistical learning in multiple instance problems," Department of Computer Science, M.S. Thesis, Dept. Compt. Sci., The University of Waikato, Jun. 2003.
- [58] Y.-L. Boureau, J. Ponce, and Y. LeCun, "A theoretical analysis of feature pooling in visual recognition," in *Proc. Int. Conf. Mach. Lear.*, Haifa, Israel, 2010, pp. 111–118.
- [59] X. Yang, T. Zhang, and C. Xu, "Cross-domain feature learning in multimedia," *IEEE Tran. Multimedia*, vol. 17, no. 1, pp. 64–78, Jan. 2015.
- [60] J. Qi, Y. Peng, and Y. Zhuo, "Life-long cross-media correlation learning," in *Proc. ACM Multimedia*, Seoul, Korea, 2018, pp. 528–536.



Zhao Ren (Student Member, IEEE) received the bachelor's and master's degrees in computer science and technology from Northwestern Polytechnical University (NWPU), Xi'an, China, in 2013 and 2017, respectively. She is currently working toward the Ph. D. degree with the Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Augsburg, Germany. She is an EU-Researcher, and she is involved with the H2020-MSCA-ITN-ETN Project TAPAS, the Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, for pathological speech analysis. She has coauthored publications in peer-reviewed journals and conference proceedings, including *IEEE JOURNAL OF BIOMEDICAL AND HEALTH INFORMATICS*, *IEEE International Conference on Acoustics, Speech, and Signal Processing*, *ISCA INTERSPEECH*, *DCASE*, etc. Her research interests include developing deep learning techniques for the applications in health care and wellbeing. She regularly reviews *IEEE TRANSACTIONS ON CYBERNETICS*, *IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS*, *IEEE TRANSACTIONS ON AFFECTIVE COMPUTING*, and *IEEE TRANSACTIONS ON MULTIMEDIA*.



Qiuqiang Kong (Member, IEEE) received the Ph.D. degree from the University of Surrey, Guildford, U.K., in 2020. Following his Ph.D., he joined ByteDance AI Laboratory as a research scientist. He has coauthored paper in journals and conferences including IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING (TASLP), ICASSP, INTERSPEECH, IJCAI, DCASE, EU-SIPCO, LVA-ICA, etc. His research interests include the classification, detection, and separation of general sounds and music. He is known for developing attention

neural networks for audio tagging, and winning the audio tagging task in the detection and classification of acoustic scenes and events challenge in 2017. He was nominated as the postgraduate research student of the year in University of Surrey, 2019. He is a frequent reviewer for over ten world well known journals and conferences including TASLP, TMM, SPL, TKDD, JASM, EURASIP, Neurocomputin, Neural Networks, ISMIR, CSMT, etc.



Jing Han (Member, IEEE) received the bachelor's degree in electronic and information engineering from Harbin Engineering University, Harbin, China, in 2011, the master's degree from Nanyang Technological University, Singapore, in 2014, and the doctoral degree in computer science from the University of Augsburg, Augsburg, Germany, in 2020. She is a Postdoctoral Researcher with the Department of Computer Science and Technology, University of Cambridge, U.K. She has coauthored more than 35 publications in peer-reviewed journals and

conference proceedings, such as IEEE TRANSACTIONS ON AFFECTIVE COMPUTING, IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, *IEEE Computational Intelligence Magazine*, ACM Multimedia, IEEE ICASSP, ISCA INTERSPEECH. Her research interests include deep learning for human-centric multimodal affective computing and health care. Besides, she Co-Chaired the 7th Audio/Visual Emotion Challenge and workshop in 2017, and was a program committee member of AVEC challenge and workshop in 2018 and TPC member of ACM Multimedia in 2019 and 2020. She was awarded student travel grants from IEEE SPS and ISCA to attend ICASSP and INTERSPEECH in 2018.



Mark D. Plumbley (Fellow, IEEE) received the B.A. (Hons.) degree in electrical sciences and the Ph.D. degree in neural networks from the University of Cambridge, Cambridge, U.K., in 1984 and 1991, respectively. Following his Ph.D., he became a Lecturer with Kings College London, before moving to Queen Mary University of London, in 2002, where he subsequently became Professor and Director of the Centre for Digital Music. He joined the University of Surrey in 2015 as Professor of Signal Processing, and in 2019 became the Head of School of Computer Science and

Electronic Engineering. He is a coeditor of a recent book on Computational Analysis of Sound Scenes and Events, and Co-Chair of the recent DCASE 2018 Workshop on Detection and Classifications of Acoustic Scenes and Events. His research interests include analysis and processing of audio and music, using a wide range of signal processing techniques, including matrix factorization, sparse representations, and deep learning. He is a member of the IEEE Signal Processing Society Technical Committee on Signal Processing Theory and Methods, and a fellow of the IET.



Björn W. Schuller (Fellow, IEEE) received the diploma, in 1999, the doctoral degree for his study on automatic speech and emotion recognition, in 2006, and the habilitation and Adjunct Teaching Professorship degree in the subject area of signal processing and machine intelligence, in 2012 all in electrical engineering and information technology from Technische Universität München, Munich, Germany. He is a tenured Full Professor heading the Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, Germany, and a Professor of

Artificial Intelligence heading GLAM – the Group on Language, Audio and Music, Department of Computing, the Imperial College London, London, U.K. He coauthored 5 books and more than 800 publications in peer reviewed books, journals, and conference proceedings leading to more than 33 000 citations (h-index 79). He was an elected member of the IEEE Speech and Language Processing Technical Committee, Editor in Chief of the IEEE TRANSACTIONS ON AFFECTIVE COMPUTING, and is President-emeritus of the AAC and Senior Member of the ACM.