

# End-to-End Classification of Reverberant Rooms using DNNs

Constantinos Papayiannis, *Member, IEEE*, Christine Evers, *Senior Member, IEEE*,  
and Patrick A. Naylor, *Fellow, IEEE*

**Abstract**—Reverberation is present in our workplaces, our homes, concert halls and theatres. This paper investigates how deep learning can use the effect of reverberation on speech to classify a recording in terms of the room in which it was recorded. Existing approaches in the literature rely on domain expertise to manually select acoustic parameters as inputs to classifiers. Estimation of these parameters from reverberant speech is adversely affected by estimation errors, impacting the classification accuracy. In order to overcome the limitations of previously proposed methods, this paper shows how DNNs can perform the classification by operating directly on reverberant speech spectra and a CRNN with an attention-mechanism is proposed for the task. The relationship is investigated between the reverberant speech representations learned by the DNNs and acoustic parameters. For evaluation, AIRs are used from the ACE-challenge dataset that were measured in 7 real rooms. The classification accuracy of the CRNN classifier in the experiments is 78% when using 5 hours of training data and 90% when using 10 hours.

**Index Terms**—Room Classification, Attention Mechanisms, Deep Neural Networks, Convolutional Recurrent Neural Networks, Reverberation, Reverberant Speech Classification, Room Acoustics.

## I. INTRODUCTION

Acoustic environments shape and define aspects of the sounds we hear and through this process we experience the world around us from an audible perspective. At the same time, audio recordings provide listeners with cues that enable the understanding of properties of the environments [1]. For example, as human listeners, we are able to tell whether we are sitting in a large concert hall compared to a tiled bathroom as the two would have very different acoustics. Discriminative models can be used to allow machines to make similar distinctions between different types of acoustic environments [2]. The ability to classify a recording in terms of the room in which it was recorded is becoming very important as the adoption of smart-home devices is increasing. For instance, knowing the room in which a speaker is located is critical to a machine that will take action to the command "Turn on the lights!".

Constantinos Papayiannis is currently with Amazon Alexa, Cambridge MA, USA (e-mail: papayiac@amazon.com). This work was done prior to that when he was with the Department of Electrical and Electronic Engineering, Imperial College London. Christine Evers was also with the Department of Electrical and Electronic Engineering, Imperial College London. She is now with the School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, U.K. (e-mail: c.evers@soton.ac.uk). P. A. Naylor is with the Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2AZ, U.K. (e-mail: p.naylor@imperial.ac.uk).

This work received support from the UK EPSRC Fellowship Grant EP/P001017/1, awarded to C. Evers while at Imperial College London.

Classifying the reverberation effect has been addressed in the literature in the past. The work in [3] proposed the use of a Gaussian Mixture Model (GMM)-Universal Background Model (UBM) classifier using Mel-frequency Cepstral Coefficients (MFCCs). The approach in [4] proposed the use of the Negative-Side Variance (NSV) for the classification using a Gaussian Naive Bayes Classifier (NBC). In [2] Frequency-Dependent Reverberation Times (FDRTs) were used as input features to a NBC. The work in [1] showed that spectral and energy-decay features, such as the FDRTs can be used for room classification and can be reliably derived from Acoustic Impulse Responses (AIRs). However, their estimation from reverberant speech is known to be challenging [5]. Errors in the estimation of acoustic parameters impact the classification accuracy of methods that use them [6]. In contrast to room classification, deep learning has been used in the past to learn other properties of reverberant environments, such as the reverberation time [7]–[9], the Early-To-Late Reverberation Ratio (ELRR) [8] and the room volume [10].

The aim of this work is to overcome the limitations of previously proposed room classification methods by proposing state-of-the-art classifiers for the task. Deep Neural Networks (DNNs) are used to perform the classification directly from spectrograms extracted from reverberant speech, avoiding the challenging estimation of parameters such as the FDRTs. Working only with reverberant speech spectra also avoids time-consuming AIR measurements [11]. The performance of a set of candidate DNN architectures is evaluated. The best performing model is a Convolutional Recurrent Neural Network (CRNN) that incorporates an attention-mechanism. CRNNs have been successfully used for Automatic Speech Recognition (ASR) [12], Sound Event Detection (SED) [13], and media-presence detection [14]. Attention-mechanisms have provided substantial performance improvements for ASR [15], translation [16] and image-description [17]. Motivated by their success in other areas, CRNNs and attention-mechanisms are now applied to the task of room classification. The proposed combination shows benefits compared to the other approaches considered. The performance of the proposed model is compared to that of previous classifiers in the literature. Important contributions of this work include the use of deep learning for room classification and the insight derived from the analysis of the attention-vectors and representations learned by DNNs. This reveals what information is identified as important to the task during training and the relationship of this information with acoustic parameters.

The structure of the remainder of this paper is as follows:

Section II presents the model of reverberant speech. Section III discusses the candidate DNN architectures for the task. Section IV presents the method used to train DNNs and the data used. Section V presents the experiments that evaluate the performance of DNNs on the task of room classification. Section VI analyzes the representations learned by the DNNs and their relationship with acoustic parameters. Conclusions are drawn in Section VII.

## II. SIGNAL MODEL

In a dataset of size  $M$ , the  $m^{\text{th}}$  reverberant speech signal  $\mathbf{x}_m \triangleq [x_m(0), \dots, x_m(N-1)]$ , is defined as

$$x_m(n) = h_m(n) * s_m(n), \quad (1)$$

where  $*$  describes the convolution operation,  $\mathbf{s}_m \triangleq [s_m(0), \dots, s_m(N-1)]$  the anechoic speech signal for sample index  $n \in \{0, \dots, N-1\}$  and  $\mathbf{h}_m \triangleq [h_m(0), \dots, h_m(N_h-1)]$  denotes the AIR of length  $N_h$ .

Vectors  $\mathbf{x}_m$ ,  $\mathbf{s}_m$  and  $\mathbf{h}_m$  are transformed into the log-power Short Time Fourier Transform (STFT) domain to obtain  $\mathbf{X}_m$ ,  $\mathbf{S}_m$  and  $\mathbf{H}_m$  respectively, similar to [14]. The size of the STFT frames  $N_s$  is selected for relevant experiments in later Sections. An overlap of  $\frac{N_s}{2}$  samples exists between each STFT frame.

## III. CANDIDATE DNN ARCHITECTURES

### A. Architectures' Description

The first candidate model is a Convolutional Neural Network (CNN). CNNs have been previously used for ASR using audio spectrograms [18] and for Reverberation Time ( $T_{60}$ ) estimation [9]. Inputs to the CNNs are processed by a stack of convolutional layers [19] that apply small filters to the input data. Convolutional layers are typically separated by Max Pooling layers [20] to reduce dimensionality while maintaining the most relevant information. Dropout [21] is used at the input of fully-connected layers to avoid overfitting. The Rectified Linear Unit (ReLU) function [22] is typically used as an activation. A softmax activation is applied to the output of the last fully connected layer. To incorporate the sequential nature of speech signals, the second candidate model is a Recurrent Neural Network (RNN). The inputs to the RNNs are first processed by Time Distributed (TD) layers. A TD layer with output  $\mathbf{Y} \in \mathbb{R}^{N_f \times D_y}$  performs the operation

$$\mathbf{y}_i = \mathbf{W}\mathbf{x}_i + \mathbf{b} \quad \forall i \in \{0, \dots, N_f - 1\}, \quad (2)$$

where  $\mathbf{y}_i$  are rows of  $\mathbf{Y}$  and  $\mathbf{x}_i$  are rows of the input  $\mathbf{X} \in \mathbb{R}^{N_f \times D_x}$ .  $\mathbf{W} \in \mathbb{R}^{D_y \times D_x}$  and  $\mathbf{b} \in \mathbb{R}^{D_y}$  are learnable parameters. The 3<sup>rd</sup> candidate model combines convolutional and recurrent layers to form a CRNN [23]. All the architectures are given in Fig. 1. The RNN and CRNN networks are studied with and without an attention-mechanism, which is described in the following section.

### B. Attention Mechanism

The RNN and CRNN architectures for room classification shown in Fig. 1 incorporate an attention-mechanism. The mechanism is used to allow each model to emphasize or de-emphasize the importance of individual time-frames. The mechanism ‘‘guides’’ the classifier to make more accurate predictions by treating RNN layers as encoders and computing an attention-vector  $\alpha_i$  for their output sequence  $\zeta_i$ , where  $i \in \{0, \dots, N_i - 1\}$  a time-step of the recurrent layer and  $N_i$  the total number of time-steps.

The attention-mechanism is implemented as [16], [17]

$$\alpha_i = \frac{\exp\{a_f(\zeta_i)\}}{\sum_{i=0}^{N_i-1} \exp\{a_f(\zeta_i)\}}, \quad (3)$$

where  $\alpha_i$  the attention applied to frame  $i$  and  $a_f$  a function parametrised by a 1-layer feed-forward network. The mechanism operates on the sequence of the last recurrent layer and the resulting context vector

$$\mathbf{c} = \sum_{i=0}^{N_i-1} \alpha_i \zeta_i \quad (4)$$

is used as the input to fully-connected layers, as shown in Figs. 1b and 1c.

## IV. NETWORK TRAINING AND GENERALIZATION

### A. Dataset Generation

The candidate models are trained using spectrograms of reverberant speech signals, obtained by convolving anechoic speech signals from TIMIT [24] with AIRs from the Acoustic Characterization of Environments (ACE) challenge [25] and Queen Mary, University of London (QMUL) datasets [26], downsampled to 16 kHz. The ACE dataset contains AIRs measured in 7 rooms (see Table I) using 5 microphone-arrays. The arrays contain between 2 to 32 microphones. 100 AIRs are provided for each room at 10 different microphone-array positions, giving a total of 700 AIRs at 70 array positions. Using the QMUL dataset allows the evaluation of the generalisation of the presented training methods and model architectures. ACE AIRs were predominantly recorded in small offices and meeting rooms whereas QMUL AIRs were recorded in much larger environments such as halls (see Table I). The Reverberation Time ( $T_{60}$ ) and Direct-to-Reverberant Ratio (DRR) for each AIR from each room in [25] and [26] are shown in Fig. 3.

Anechoic speech utterances are created by concatenating sentences from a single speaker from the TIMIT dataset. Each resulting utterance is convolved with one AIR to produce reverberant speech. Each AIR is convolved with  $N_u$  speech utterances. Prior to using a speech utterance, it is randomly time-shifted and truncated to 5 s. The selection of which utterance to convolve with an AIR is random with replacement, which means that the speech data are oversampled. It is important to note that speech from the same speaker is never used for both training and testing and the training and test splits of the TIMIT distribution are maintained. The resulting reverberant speech samples are used to create a matrix  $\mathbf{X}$  of




Fig. 1: Candidate architectures for room classification from reverberant speech.

dimensions  $[N_u \times M, N_f, \frac{N_s}{2} + 1]$ . The framesize  $N_s$  for the analysis of speech is set to 20 ms (320 samples) and a Hanning window is applied to each frame. For the utterance length of 5 s and a sampling rate of 16 kHz,  $N_f = 500$ . Values for the number of utterances convolved per AIR,  $N_u$ , and the total number of AIRs,  $M$ , are varied for each experiment as detailed in Section IV-D.

### B. Training data batches

The ACE AIR data [25] is evenly distributed across rooms, however, it is significantly imbalanced in terms of the number of AIRs available for each measurement position. This proves problematic when constructing training batches by randomly sampling the dataset as performance is biased towards the modes that contributed most of the training data. Therefore, batches are constructed with a balanced number of samples from each of the measurement positions. The method includes in each batch a fixed and equal number of AIRs from each of the measurement positions. This creates a balance across the positions and, for an equal number of positions per microphone-array and the same set of microphone-arrays used per room, this simultaneously creates a balance across rooms. The batch construction method is shown in Fig. 2. The “Selection” block is a mechanism that iterates through the AIRs contained in the previous block. Each block selects and contributes two AIRs to each batch, which are convolved with anechoic speech. For the case of the ACE challenge shown in Fig. 2, the resulting batch size is  $M_b = 140$ .

During training, an epoch is defined by  $\lceil MN_u \rceil$  weight updates, with one batch constructed for each. In the method described above for creating batches, AIRs are reused within the same epoch. As stated in Section IV-A, for each AIR and each batch, reverberant speech is created using randomly selected and shifted samples from the TIMIT dataset. These random

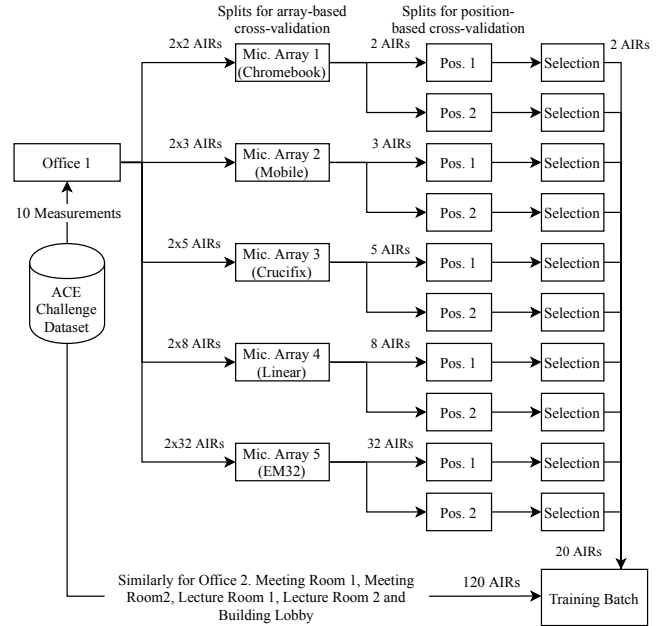


Fig. 2: Organization of the data from the ACE dataset according to rooms, microphone-arrays and recording positions for the purpose of training DNN room classifiers. The “Selection” block maintains a counter looping through the AIRs contained in the previous block.

permutations of the data avoid the repeated presentation of identical training samples to the model during training and the subsequent bias and overfitting to specific examples.

### C. Optimization

The Adam optimizer is used to train the networks by minimizing the categorical cross-entropy between the model’s

Dataset	Room	L (m)	W (m)	H (m)	Vol. (m <sup>3</sup> )
ACE [25]	Office 1	3.32	4.83	2.95	47.30
	Office 2	3.22	5.10	2.94	48.30
	Meeting Room 1	6.61	5.11	2.95	99.60
	Meeting Room 2	10.30	9.07	2.63	246.00
	Lecture Room 1	6.93	9.73	3.00	202.00
	Lecture Room 2	13.60	9.29	2.94	370.00
	Building Lobby	4.47	5.13	3.18	72.90
QMUL [26]	Classroom	7.50	9.00	3.50	236.00
	Great Hall*	16.00	23.00	**	**
	Octagon*	23.00	23.00	21.00	9500.00

\*Room is not “shoe-box” shaped., \*\*Information not provided by authors.

TABLE I: Datasets of AIRs used in this work and the corresponding room information.




Fig. 3: DRR and  $T_{60}$  values for the AIRs used in this work grouped by room.

output and the label for the room associated with each sample. To avoid overfitting, early stopping is used to terminate the training when the validation loss has not improved for 10 epochs. Models are trained for a maximum of 50 epochs. The validation set is constructed by stratified selection across rooms and amounts to 15% of the available data for training.

#### D. Evaluation

Cross-validation is used for the evaluation of classifiers. For the ACE dataset, two types of partitioning of the dataset into folds are studied. The first type, similar to [1], partitions the dataset into 70 folds, each one corresponding to a recording position (see Fig. 2). The second partitions the data into 5 folds, each one corresponding to the microphone-array used for each measurement. The second case is more challenging as more data is held-out during training and the microphone-array used for recording the held-out samples is unseen during training. For the QMUL dataset the data is split based on positions on the densely populated two-dimensional measurement grid [26]. In order to ensure dissimilarity of data points between the splits, the receiver-position grid is split in three uniformly sized parts along the x-axis value of the microphone’s position. The two extreme ends are kept and each one forms one fold.

In total, 288 AIRs are used from the dataset. Two uniformly sized folds are therefore used to split 288 AIRs from the three rooms that form the QMUL dataset. Each fold is tested by predicting the class of the held-out samples by training a network using the remaining samples. As the evaluation measure, the accuracy of predictions is used.

The baselines considered for the task are the following: Baseline 1 is the method proposed in [2], which is based on the use of FDRTs as “Roomprints”, combined with Gaussian Naive Bayes Classifiers (NBCs). Baseline 2 uses the same inputs with a 3-layer Feed Forward (FF) network as the classifier with 32 units for each hidden layer and dropout of probability of 10%. This is equivalent to replacing the NBC classifier of [2] with a DNN classifier. Baseline 3 is the SVM classifier and features proposed in [1]. The features are a collection of acoustic parameters that include FDRTs, MFCCs and Frequency-Dependent DRRs (FDDRRs). All of the above baselines take acoustic parameters as inputs, which are estimated with access to the AIRs. The performance of the DNNs is also compared with the method of [3]. It is based on a GMM-UBM setup, training a base GMM model with 128 mixtures, that is later copied and adapted to form one GMM for each of the rooms. The inputs to the model are reverberant speech MFCCs, augmented with  $\Delta s$  and  $\Delta\Delta s$  of the parameters. This is an alternative to the use of DNNs.

## V. EXPERIMENTAL EVALUATION

The classification accuracy of the DNNs and the baselines described in Section IV-D are summarized in Table II. The results show that for the ACE challenge dataset, the best performing DNN architecture for the task is consistently the Attention-CRNN, which achieves accuracies of 90% for both cross-validation types. The superior performance of the specific architecture is attributed to the fact that it features the most diverse set of layers, first using convolutional layers to extract feature-maps of reverberant speech, which are then processed as a sequence by bidirectional recurrent layers before the attention mechanism is applied. The confusion matrix associated with the microphone-array based cross-validation is shown in Figure 4. For the case of the QMUL dataset, the task is much simpler as 3 rooms are involved, compared to the case of the ACE challenge, which involves 7 rooms. All classifiers perform better in this case. The best accuracy is provided by the Attention-RNN at 94%. The Attention-CRNN achieves an accuracy of 90%. While its accuracy is almost identical to that on the ACE dataset, the remaining models show significant differences between the two cases. The generalization of the Attention-CRNN to both datasets is illustrated by this observation. When considering the joint case of ACE and QMUL, the best performing model is the CRNN model, with an accuracy of 82%. The Attention-RNN is a close second with an accuracy of 81%. The Attention-CRNN gives a comparable accuracy of 79%. The remaining models show significantly lower performance on this most challenging test-case.

Comparing the performance of the DNNs with Baseline 4 shows that all architectures outperform the baseline in all

Datasets included	ACE folds cross-validation	AIRs			Reverberant Speech					
		Baseline 1 [2] FDET NBC	Baseline 2 FDRT FF	Baseline 3 [1] SVM Mult.	Baseline 4 [3] GMM-UBM	RNN	Att.-RNN	CNN	CRNN	Att.-CRNN
QMUL	–	<b>1.000</b>	0.851	<b>1.000</b>	0.610	0.894	<b>0.936</b>	0.896	0.852	0.899
ACE	Position	0.981	0.966	<b>1.000</b>	–*	0.793	0.770	0.732	0.865	<b>0.896</b>
ACE	Mic. Array	0.568	0.752	<b>0.874</b>	0.693	0.730	0.798	0.714	0.777	<b>0.904</b>
ACE & QMUL	Mic. Array	0.680	0.769	<b>0.844</b>	0.620	0.718	0.814	0.690	<b>0.821</b>	0.787

\* Not evaluated due to large number of folds and long training times required by the baseline.

TABLE II: Results of room classification as the accuracy of each classifier. Baselines 1–3 rely on the availability of AIRs whereas Baseline 4 uses reverberant speech MFCCs. DNN classifiers operate directly on reverberant speech spectra.

cases. The remaining baselines assume access to the AIRs to estimate acoustic parameters as features. This makes the task significantly easier as the acoustic channel does not have to be inferred from reverberant speech. Estimation of such parameters from reverberant speech is a significantly challenging task [5] and estimation errors impact the performance of the classifiers [6]. Nevertheless, the DNN classifiers provide comparable or even greater accuracy than these baselines in Table II. For the case of the ACE challenge dataset and with microphone-arrays as the cross-validation splits, the Attention-CRNN provides an accuracy of 90%, compared to 84% provided by [1]. The method of [1] is consistently the best performing baseline. However, it uses FDDRR values, which are notoriously difficult to estimate without access to AIRs [5]. The easiest task for the baselines is the case of the QMUL dataset. This is attributed to the well separated  $T_{60}$  values of the AIRs in the dataset (see Fig. 3). As all classifiers use FDRTs derived from AIRs, the task becomes trivial. The same does not apply for the case of the ACE challenge, which is reflected in the accuracy numbers. Notably, classifying AIRs with the microphone-arrays as splits makes the task much more challenging over the case of having only the positions as splits, both for DNNs and the baselines. For baselines with FDRTs as the features this is particularly the case with the accuracy of [2] falling to 57% from 98%. Using a fully-connected layer instead of the NBC provides respective scores of 97% and 75%. This indicates that information about the microphone-array used during inference is important for FDRT based room classifiers.

Considering the potential applications of the classifiers, an important dimension to consider is the amount of audio that is needed to train them at a satisfactory level of accuracy. Such an analysis for each of the architectures studied is performed and the results are shown in Fig. 5. The figure shows the accuracy of each network when trained with an increasing amount of training hours. All models trained are tested on a static test set of 15 hours. It can be seen that with less than 2 hours of available training hours the performance of the networks is significantly worse than having 15 hours available. The most significantly impacted is the CNN model that shows a drop of 49% relative accuracy from the highest of 73%. With five hours available, all networks achieve more than 80% of relative performance compared to having 15 hours of audio available for training. The Attention-CRNN results in an accuracy of 78%. It is also interesting to note that, while for other networks providing more than 7 hours of training data

Actual \ Predicted	BL	L1	L2	M1	M2	O1	O2
BL	92.4% 1386	5.5% 83	1.8% 27	0.1% 2			0.1% 2
L1	2.5% 38	96.4% 1446		0.3% 4	0.3% 5	0.4% 6	0.1% 1
L2	0.7% 10	0.1% 1	99.3% 1489				
M1		0.9% 13		92.9% 1394	2.5% 37	1.2% 18	2.5% 38
M2		1.1% 16		13.3% 199	78.4% 1176	6.7% 101	0.5% 8
O1		4.4% 66		6.2% 93	8.8% 132	79.4% 1191	1.2% 18
O2				5.6% 84	0.1% 2	0.5% 7	93.8% 1407

Legend: MR1 - Meeting Room 1, MR2 - Meeting Room 2, BL - Building Lobby, O1 - Office 1, O2 - Office 2, LR1 - Lecture Room 1, LR2 - Lecture Room 2.

Fig. 4: Confusion matrix for room classification based on reverberant speech by the Attention-CRNN classifier using data from the ACE and TIMIT datasets.

offered little improvements, the Attention-CRNN continued to improve beyond that point, reaching an accuracy of 91% with 15 hours used for training. This behaviour is attributed to the more diverse set of mechanisms that are part of the model, allowing it to more effectively leverage spectral and temporal patterns in the input, which are enhanced by the ability to learn to attend to the most informative parts of its input. As shown in (3), information from all time-steps are weighted and effectively used for the classification, compared to the case of not using the attention mechanism that forces the recurrent layer to retain all important information in the last time-frame, which is driving the final classification. The attention-mechanism also offers a benefit to the RNN model but with an overall lower accuracy than the CRNN equivalent.

## VI. ANALYSIS OF LEARNED REPRESENTATIONS

As the accuracy of different DNNs and baselines differs across datasets, an additional experiment is performed that provides insight into the relationship between acoustic parameters,




Fig. 5: Accuracy of DNN classifiers versus available training audio hours for the ACE dataset. The test set is fixed to 15 hours of reverberant speech, equally distributed across rooms.




Fig. 6: Reverberant speech spectrogram, waveform and attention vector applied by the Attention-CRNN.

namely the DRR and  $T_{60}$ , and the representation of reverberant speech at the outputs of deeper layers of the networks. This experiment studies the representations at the output of the last convolutional stage of the CNN and CRNN of Fig. 1. The representation is formed by a set of output channels, each one being a higher level and lower dimensional representation of the reverberant speech spectrogram. To investigate the correlation between acoustic parameters and the representations, the values of the  $T_{30}$ ,  $T_{60}$  and DRR are linearly predicted directly from the representations. To this end, a weight vector  $\mathbf{w} \in \mathbb{R}^{D_c}$  and a scalar  $b$  are trained to minimize

$$l = (y_m - \hat{y}_m)^2 \quad (5)$$

where  $y_m$  the value of the acoustic parameter being estimated, extracted from the AIR  $\mathbf{h}_m$ , and  $\hat{y}_m$  its estimate as

$$\hat{y}_m = \mathbf{w} f_{\text{conv}}(\mathbf{X}_m) + b. \quad (6)$$

	CNN	CRNN	Rand. Init.
DRR $\rho_{ccc}$	0.257	0.260	-0.039
T30 $\rho_{ccc}$	0.038	0.022	0.036
T60 $\rho_{ccc}$	-0.004	0.046	-0.030

TABLE III: Accuracy of prediction of acoustic parameters using DNN embeddings. Accuracy is measured by the concordance correlation coefficient  $\rho_{ccc}$  between the predicted and estimated values of the parameters from the AIR.

Spectral Feature	Attention-CRNN	Rand. Init.
Centroid $z_c(i)$	-0.115	0.002
Bandwidth $z_b(i)$	-0.188	-0.005
Roll-off $z_r(i)$	-0.004	0.007

TABLE IV: Pearson correlation coefficient between attention-vectors applied to reverberant speech by the Attention-CRNN and spectral features. A randomly initialized and untrained network of identical architecture is used as a reference.

The function  $f_{\text{conv}}$  corresponds to the convolutional layers of the models in Figs. 1a and 1c, which outputs a representation-vector of length  $D_c$ . The matrix  $\mathbf{X}_m$  is the reverberant speech spectrogram. The values of  $\mathbf{w}$  and  $b$  are learned using the Adam optimizer [27] using batches of 16 AIRs. The accuracy of the predictions is measured by the concordance correlation coefficient [28] defined as

$$\rho_{ccc} = \frac{2\rho_{y\hat{y}}\sigma_y\sigma_{\hat{y}}}{\sigma_y^2 + \sigma_{\hat{y}}^2 + (\mu_y - \mu_{\hat{y}})^2} \quad (7)$$

where  $\rho_{y\hat{y}}$  is the Pearson's correlation coefficient between the estimates and values of parameters,  $\sigma_y$  and  $\sigma_{\hat{y}}$  the standard deviation for each and similarly for the means  $\mu_y$  and  $\mu_{\hat{y}}$ . The results of this experiment are shown in Table III and show that there is a non-negligible correlation  $\rho_{ccc}$  of 0.3 between linear estimates of DRR from the representations and the values of the DRR measured from the AIRs.

A similar analysis is performed on the attention-vectors calculated by the Attention-CRNNs. The correlation between the values of the vectors defined in (3) and spectral features is measured. The attention values  $\alpha_i$ , where  $i \in \{0, \dots, N_f - 1\}$  is the frame index, are calculated for a reverberant speech signal  $\mathbf{x}_m$  with STFT matrix  $\mathbf{F} \in \mathbb{R}^{N_f \times K}$ . The correlation is evaluated between the attention-vector and the following spectral features:

1. The spectral centroid [29], defined as

$$z_c(i) = \frac{\sum_{k=0}^{K-1} F_k(i) f_k}{\sum_{k=0}^{K-1} F_k(i)}, \quad (8)$$

where  $f_k$  the centre frequency of STFT bin with index  $k$  and  $F_k(i)$  the magnitude of the STFT at frame  $i$ .

2. The spectral bandwidth [29], defined as

$$z_b(i) = \left\{ \sum_{k=0}^{K-1} F_k(i) \{f_k - z_c(i)\}^2 \right\}^{\frac{1}{2}}, \quad (9)$$

which is a measure of the energy away from the centroid defined by (8).

3. The spectral roll-off  $z_r(i)$  [30], which indicates the frequency  $f_k$  that corresponds to the bin under (and



including) which 85% of the energy of frame  $i$  is captured.

The three features are complimentary, in the sense that the centroid indicates the frequency region where energy is most concentrated, the bandwidth measures the spread of the energy across the spectrum and the roll-off indicates the frequency above which little energy exists. The results in Table IV show the Pearson correlation coefficient between the attention-vector and each feature. The results show that all correlations are negative. The strongest negative correlation of -0.19 exists between the bandwidth of the reverberant speech signal and the attention-vector, indicating that more attention is given to frames that contain energy concentrated closer to the centroid  $z_c(i)$ . Smaller correlation exists between the location of the centroid and the attention-vector and no correlation is shown between the roll-off frequency  $z_r(i)$  and the attention-vector. As an example, Fig. 6 shows a reverberant speech signal and its spectrum, along with the corresponding attention-vector. What is characteristic in the figure is that the attention vector has higher values near the end of the utterance where energy is decaying, indicating the importance of this region for the task.

## VII. DISCUSSION AND CONCLUSION

This paper proposed the use of neural networks for room classification. Five model architectures were investigated as classifiers for the task: 1) CNN 2) RNN 3) Attention-RNN 4) CRNN 5) Attention-CRNN. The performance of the trained classifiers was compared to the GMM-UBM of [3] using MFCCs. In the experiments presented, the Attention-CRNN architecture provided the highest classification accuracy in most of the test cases. The model also provided the highest accuracy when very limited amounts of training hours were presented during training. Since the models are trained using spectrograms of reverberant speech, estimation errors due to acoustic parameter estimation are avoided. The experiments presented on the ACE challenge dataset have shown that the Attention-CRNN achieves a classification accuracy of 78% using 5 hours and 90% using 10 hours of training data.

Convolutional layers in the networks learn representations of reverberant speech spectrograms. An investigation of the predictive power of the representations regarding a set of acoustic parameters indicated that the DRR can be linearly predicted from the representation with a concordance correlation coefficient of 0.3 between the DRR values and their estimates. Furthermore, a correlation of -0.19 was identified between the attention-vectors produced by the Attention-CRNN and the spectral bandwidth of the input reverberant speech. A notable finding of the above analysis is the stronger relationship between the DRR and the learned representations compared to the  $T_{60}$ . This is notable because the  $T_{60}$  is generally considered to be a room-dependent parameter in contrast to the DRR that relates to the source-receiver distances [31]. This finding is in line with the observations in [1], which show that sub-band DRR measurements are useful for the task of room classification when using the ACE dataset. A relevant aspect of the ACE dataset is the fact that the rooms are of the

scale typically expected in residential and office buildings, which makes their  $T_{60}$  relatively short compared to larger rooms, such as the ones found in the QMUL dataset (see Fig. 3). One interpretation of the above is that, for spaces with similar characteristics to the ones in the ACE dataset, the ratio between the energy of the strong early reflections to the remaining energy can be more informative for room classification than features that only capture the energy decay. Of course the learned representations are not bound to follow a strict definition of the DRR or the  $T_{60}$ , which generally makes DNNs able to learn the most useful features for the task. This benefit sets DNN classifiers apart from methods such as [2], where the representation of the input is fixed prior to training. Another important point to note here is that because of this benefit the learned representations are to some degree “tailored” to the training dataset and using different datasets will lead to different representations. This is not a concern in settings where the training and testing rooms are fixed and known at the training time but might impact the usefulness of the learned representations with regards to transfer-learning to other tasks or to incremental-training settings, where rooms are incrementally added to the existing set.

The code implementation of the work discussed in this paper can be found at: [https://github.com/papayiannis/reverberation\\_learning\\_python](https://github.com/papayiannis/reverberation_learning_python).

## REFERENCES

- [1] C. Papayiannis, C. Evers, and P. A. Naylor, “Discriminative feature domains for reverberant acoustic environments,” in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, Louisiana, USA, Mar. 2017, pp. 756–760.
- [2] A. H. Moore, M. Brookes, and P. A. Naylor, “Room identification using roomprints,” in *Proc. Audio Eng. Soc. (AES) Conf. on Audio Forensics*, June 2014.
- [3] N. Peters, H. Lei, and G. Friedland, “Name that room: Room identification using acoustic features in a recording,” in *Proceedings of the 20th ACM International Conference on Multimedia*, 2012, pp. 841–844.
- [4] A. H. Moore, P. A. Naylor, and M. Brookes, “Room identification using frequency dependence of spectral decay statistics,” in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 6902–6906.
- [5] J. Eaton, N. D. Gaubitch, A. H. Moore, and P. A. Naylor, “Proceeding of the ACE Challenge,” *Proceedings, New Paltz, NY, USA*, Oct. 2015.
- [6] C. Papayiannis, *Models for learning reverberant environments*, Ph.D. thesis, Imperial College London, 2019.
- [7] T. J. Cox, F. Li, and P. Darlington, “Extracting room reverberation time from speech using artificial neural networks,” *J. Audio Eng. Soc. (AES)*, vol. 49, no. 4, pp. 219–230, 2001.
- [8] F. Xiong, S. Goetze, B. Kollmeier, and B. T. Meyer, “Joint estimation of reverberation time and early-to-late reverberation ratio from single-channel speech signals,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 2, pp. 255–267, 2019.
- [9] H. Gamper and I. Tashev, “Blind reverberation time estimation using a convolutional neural network,” in *Proc. International Workshop on Acoustic Signal Enhancement (IWAENC)*, September 2018, pp. 1–5, IEEE, Nominated for best paper award.
- [10] A. F. Genovese, H. Gamper, V. Pulkki, N. Raghuvanshi, and I. J. Tashev, “Blind room volume estimation from single-channel noisy speech,” in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 231–235.
- [11] A. Farina, “Simultaneous measurement of impulse response and distortion with a swept-sine technique,” in *Proc. Audio Eng. Soc. (AES) Convention*, Feb. 2000, pp. 1–23.
- [12] T. Sainath, R. J. Weiss, K. Wilson, B. Li, A. Narayanan, E. Variiani, M. Bacchiani, I. Shafran, A. Senior, K. Chin, A. Misra, and C. Kim, “Multichannel Signal Processing with Deep Neural Networks for Automatic Speech Recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, pp. 965 – 979, 2017.

- [13] J. Yan and Y. Song, “Weakly labeled sound event detection with residual CRNN using semi-supervised method,” Tech. Rep., University of Science and Technology of China, National Engineering Laboratory for Speech and Language Information Processing, Hefei, China, June 2019.
- [14] C. Papayiannis, J. Amoh, V. Rozgic, S. Sundaram, and C. Wang, “Detecting Media Sound Presence in Acoustic Scenes,” in *Proc. Conf. of Intl. Speech Commun. Assoc. (INTERSPEECH)*, Hyderabad, India, Sept. 2018, pp. 1363–1367.
- [15] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 4960–4964.
- [16] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” in *Int. Conf. on Learning Representations (ICLR)*, 2015.
- [17] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, “Show, attend and tell: Neural image caption generation with visual attention,” in *Int. Conf. on Machine Learning*, 2015, pp. 2048–2057.
- [18] O. Abdel-Hamid, A. Mohamed, H. Jiang, L. Deng, G. Penn, and D. Yu, “Convolutional neural networks for speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 10, pp. 1533–1545, 2014.
- [19] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *Int. Conf. on Learning Representations (ICLR)*, 2015.
- [20] Y.-T. Zhou, R. Chellappa, A. Vaid, and B. K. Jenkins, “Image restoration using a neural network,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 36, no. 7, pp. 1141–1151, July 1988.
- [21] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. Salakhudinov, “Improving neural networks by preventing co-adaptation of feature detectors,” *Computing Research Repository*, vol. abs/1207.0580, 2012.
- [22] R. H. R. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung, “Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit,” *Nature*, vol. 405, pp. 947, June 2000.
- [23] T. N. Sainath, O. Vinyals, A. Senior, and H. Sak, “Convolutional, Long Short-Term Memory, fully connected Deep Neural Networks,” in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Brisbane, Australia, Apr. 2015, pp. 4580–4584.
- [24] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, N. L. Dahlgren, and V. Zue, “TIMIT acoustic-phonetic continuous speech corpus,” Corpus, Linguistic Data Consortium (LDC), Philadelphia, 1993.
- [25] J. Eaton, N. D. Gaubitch, A. H. Moore, and P. A. Naylor, “The ACE Challenge - corpus description and performance evaluation,” in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, New Paltz, NY, USA, 2015.
- [26] R. Stewart and M. Sandler, “Database of omnidirectional and B-format room impulse responses,” in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Dallas, Texas, USA, Mar. 2010, pp. 165–168.
- [27] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Int. Conf. on Learning Representations (ICLR)*, 2015.
- [28] L. I.-K. Lin, “A concordance correlation coefficient to evaluate reproducibility,” *Biometrics*, vol. 45, no. 1, pp. 255–268, 1989.
- [29] P. Herrera-Boyer, A. Klapuri, and M. Davy, *Automatic Classification of Pitched Musical Instrument Sounds*, pp. 163–200, Springer US, Boston, MA, 2006.
- [30] G. Peeters, “A large set of audio features for sound description (similarity and classification) in the cuidado project,” *CUIDADO IST Project Report*, vol. 54, no. 0, pp. 1–25, 2004.
- [31] P. A. Naylor and N. D. Gaubitch, Eds., *Speech Dereverberation*, Springer, 2010.