



---

# Audio Engineering Society Conference Paper

Presented at the Conference on  
Audio Forensics  
2019 June 18 – 20, Porto, Portugal

*This paper was peer-reviewed as a complete manuscript for presentation at this conference. This paper is available in the AES E-Library (<http://www.aes.org/e-lib>) all rights reserved. Reproduction of this paper, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.*

---

## Beyond Equal-Length Snippets: How Long is Sufficient to Recognize an Audio Scene?

Huy Phan<sup>1,2</sup>, Oliver Y. Chén<sup>2</sup>, Philipp Koch<sup>3</sup>, Lam Pham<sup>1</sup>, Ian McLoughlin<sup>1</sup>, Alfred Mertins<sup>3</sup>, and Maarten De Vos<sup>2</sup>

<sup>1</sup>University of Kent, School of Computing, UK

<sup>2</sup>University of Oxford, Department of Engineering Science, UK

<sup>3</sup>University of Lübeck, Institute for Signal Processing, Germany

Correspondence should be addressed to Huy Phan ([h.phan@kent.ac.uk](mailto:h.phan@kent.ac.uk))

### ABSTRACT

Due to the variability in characteristics of audio scenes, some scenes can naturally be recognized earlier than others. In this work, rather than using equal-length snippets for all scene categories, as is common in the literature, we study to which temporal extent an audio scene can be reliably recognized given state-of-the-art models. Moreover, as model fusion with deep network ensemble is prevalent in audio scene classification, we further study whether, and if so, when model fusion is necessary for this task. To achieve these goals, we employ two single-network systems relying on a convolutional neural network and a recurrent neural network for classification as well as early fusion and late fusion of these networks. Experimental results on the LITIS-Rouen dataset show that some scenes can be reliably recognized with a few seconds while other scenes require significantly longer durations. In addition, model fusion is shown to be the most beneficial when the signal length is short.

### 1 Introduction

Audio scene recognition (ASC) [1, 2] is an important task in machine hearing [3, 4]. A common goal in ASC is to use machines to recognize environments based on acoustic signals, as in human hearing. Over the past few years, research on this problem has advanced rapidly, both in performance [5, 6, 7, 8, 9] and in the number of available datasets [10, 11, 12, 13].

Audio scenes vary significantly in their characteristics, i.e. their background noise and foreground sound events, and, as a result, the duration required for human ears to perceive and recognize them is different

from one to another. For instance, one might recognize with relative ease an “in airplane” acoustic environment within a few seconds due to its loud and distinguishable background noise. However, it may take minutes of listening to accumulate sufficient acoustic cues to differentiate a “restaurant” from a “cafe” or even a “busy street” with similar babble noise. As one of the goals of machine hearing is to achieve human hearing intelligence, this variability should also be generalizable to a machine hearing system. However, most, if not all, available ASC datasets assume a fixed signal length for every scene category and instance. Typical lengths of 30 seconds [11, 13] or 10 seconds [10, 12] have been commonly adopted when designing such

datasets. There also exist studies with shorter signal duration, such as six seconds [14]. Although a fixed common length makes dataset design and experimentation easier, it does not reflect the reality of the ASC task. Hence, this paper aims to investigate into the perspective, namely how long is sufficient to reliably recognize different acoustic scenes using state-of-the-art models. Findings of this study may also suggest how an ASC should be implemented in real applications to improve the quality of services. Our results show that, for such a system, the listening duration should be adapted to different kinds of scenes in order to reach a certain level of recognition certainty.

Methodology-wise, approaches based on convolutional neural networks (CNNs) [6, 15, 2, 16, 7] and recurrent neural networks (RNNs) [5, 17, 9] have been demonstrated to be most efficient for ASC. In addition, state-of-the-art performance has often been achieved with ensembles of multiple networks [8, 18, 19]. Although ensemble methods have been well-established to improve performance of machine learning systems, and have been applied to ASC as a rule of thumb, no prior work has studied whether and to which temporal extent model fusion is necessary and most useful for the ASC task. We therefore aim to examine this question in this work. To accomplish these goals, we use two state-of-the-art single-network models based on CNN and RNN as well as two fusion schemes of these models for classification purpose. An early fusion scheme constructs a two-stream convolutional-recurrent neural network (C-RNN) and explores in-network fusion of the two streams' learned features before classification. A late fusion scheme trains two standalone networks independently and probabilistically aggregates their classification results afterwards. The experiments are based on the LITIS-Rouen dataset [13].

## 2 CNN and RNN for ASC

The CNN model proposed here improves on its counterpart in [2] in such a way that the convolutional filters are trained to take into account invariance across multiple input feature channels. The RNN model is based on [5], which reported state-of-the-art performance on the LITIS-Rouen dataset [13]. Both the CNN and RNN models share common input features.

### 2.1 Input features

An audio signal, e.g. a segment of 4 seconds long (see more detail in Section 5.1), is first decomposed

into small frames of 250 ms with 50% overlap. We employed the label tree embedding (LTE) features proposed in [20] to represent an audio frame. To this end, low-level feature vectors are first extracted for the audio frames and used to construct a label tree. Given  $C$  scene categories, the constructed label tree consists of  $C - 1$  split nodes which index  $2(C - 1)$  meta-classes in their left and right child nodes. A low-level feature vector of an audio frame is then mapped into an LTE feature vector  $\mathbf{x} \in [0, 1]^F$ ,  $F = 2(C - 1)$ , whose entries encode the posterior probabilities of the audio frame belonging to the meta-classes. As a result, we obtain  $T$  such LTE feature vectors for the audio snippet consisting of  $T$  frames.

We employ three different low-level feature sets: (1) Gammatone spectral coefficients [21], (2) MFCCs, and (3) log-frequency filter bank coefficients. We extract LTE features with the presence/absence of the background noise as in [5, 2]. In total, we obtain  $D = 6$  feature channels so the audio snippet is eventually represented by a multi-channel image feature  $\mathbf{X} \in [0, 1]^{F \times T \times D}$ .

### 2.2 CNN

Audio scene recognition using CNNs on LTE features have been previously explored in [2]. In that work, over-time convolution with 3-dimensional convolutional filters, which fully cover feature and channel dimensions [2], did not explore the feature invariance across LTE channels. As an improvement, the CNN proposed here is designed to have 2-dimensional convolutional filters to perform convolution over time as well as across input feature channels.

Let us denote such a 2-dimensional filter as  $\mathbf{w} \in \mathbb{R}^{F \times w}$  where  $w < T$  represents the temporal width of the filter. Convolution of the filter  $\mathbf{w}$  with the multi-channel image input  $\mathbf{X}$  over-time and across-channel results in a 2-dimensional feature map  $\mathbf{O} \in \mathbb{R}^{(T-w+1) \times D}$  whose entries are given by

$$o_{ij} = (\mathbf{X} * \mathbf{w})_{ij} = \sum_{m,n} (\mathbf{X}_j[i : i + w - 1] \odot \mathbf{w})_{m,n}. \quad (1)$$

Here,  $*$  and  $\odot$  indicate the convolution and element-wise multiplication operations, respectively.  $\mathbf{X}_j[i : i + w - 1]$  denotes an image slice from time index  $i$  to  $i + w - 1$  on the channel index  $j$ . *Rectified Linear Units* (ReLU) activation [22] is then applied, followed by

1-max pooling [23, 24] on the 2D feature map to retain the most prominent feature:

$$z^{\text{conv}} = \max_{i,j} o_{ij}. \quad (2)$$

Similar to [2], we design the CNN to have  $R = 3$  filter sets corresponding to three temporal widths  $w \in \{3, 5, 7\}$  with each filter set consisting of  $Q$  convolutional filters with the same width. We adopt  $Q = 1000$  as a large number of convolutional filters was shown to be more efficient for this “shallow” and simple CNN [2]. The total number of filters is, therefore,  $R \times Q$ ; this leads to a convolutional feature vector  $\mathbf{z}^{\text{conv}} \in \mathbb{R}^{RQ}$  after the pooling layer.

During network training, the convolutional feature vector  $\mathbf{z}^{\text{conv}}$  is presented to a softmax layer for classification. The CNN is trained to minimize the cross-entropy error over the training examples:

$$E(\boldsymbol{\theta}) = - \sum_n \mathbf{y}_n \log(\hat{\mathbf{y}}_n(\mathbf{X}_n, \boldsymbol{\theta})) + \frac{\lambda}{2} \|\boldsymbol{\theta}\|_2^2. \quad (3)$$

Here,  $\boldsymbol{\theta}$  denotes the network parameters and the  $\lambda$  denotes the hyper-parameter of the  $\ell_2$ -norm regularization term. For further regularization, dropout [25] is also applied to the CNN feature map.

After network training, we extract the convolutional feature vector  $\mathbf{z}^{\text{conv}}$  and train a linear SVM for classification, in replacement of the softmax as in [2, 5].

### 2.3 RNN

For the RNN, we stack different channels of the multi-channel LTE image input  $\mathbf{X}$  in the feature dimension and treat it as a temporal sequence of feature vectors  $(\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \dots, \tilde{\mathbf{x}}_T)$  where each  $\tilde{\mathbf{x}}_t \in \mathbb{R}^{FD}$ ,  $1 \leq t \leq T$ . The RNN then reads the input sequence into the sequence of recurrent output vectors  $(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_T)$ , where

$$\mathbf{z}_t = \mathbf{h}_t \mathbf{W}_z + \mathbf{b}_z, \quad (4)$$

$$\mathbf{h}_t = \mathcal{H}(\tilde{\mathbf{x}}_t, \mathbf{h}_{t-1}). \quad (5)$$

$\mathbf{h}_t \in \mathbb{R}^H$  denotes the hidden state vector of size  $H$  at time step  $t$ ,  $\mathbf{W}_z \in \mathbb{R}^{H \times H}$  is a weight matrix and  $\mathbf{b}_z \in \mathbb{R}^H$  denotes a bias term.  $\mathcal{H}$  represents the hidden layer function of the recurrent layer and is realized using a Gated Recurrent Unit (GRU) [26]. The RNN is designed to have two recurrent layers with the hidden state vector size of  $H = 256$ . The recurrent layers are

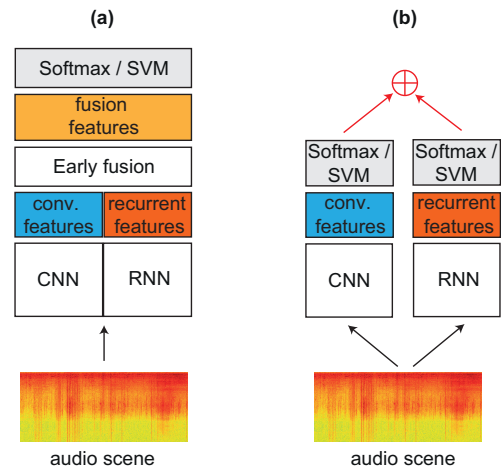
stacked on top of each other to construct a deep RNN similar to [5, 27].

We eventually retain the output vector at the last time index  $T$  as the recurrent feature vector, i.e.  $\mathbf{z}^{\text{rec}} \equiv \mathbf{z}_T \in \mathbb{R}^H$ , as it is expected to have encoded information of the entire input sequence. Dropout is also applied to the recurrent feature vector for regularization. Similarly to the CNN case, softmax and the cross-entropy loss given in (3) are employed for classification during network training. After training, the recurrent features  $\mathbf{z}^{\text{rec}}$  are also extracted and used to train a linear SVM classifier for classification.

### 3 Early Fusion with Two-Stream C-RNN

For the early fusion scheme, we construct a two-stream C-RNN network as illustrated in Fig. 1(a) and fuse the features learned by the CNN and RNN streams before the classification taken place. The idea is to allow the network to explore combinations of two feature types to optimize the classification task. The CNN and RNN streams in the C-RNN network have the same body architectures as the CNN and RNN described in Sections 2.2 and 2.3.

At the fusion layer of the C-RNN network, a fusion function  $f : (\mathbf{z}^{\text{conv}}, \mathbf{z}^{\text{rec}}) \mapsto \mathbf{z}^f$  fuses the convolutional feature vector  $\mathbf{z}^{\text{conv}}$  of the CNN stream and the recurrent feature vector  $\mathbf{z}^{\text{rec}}$  of the RNN stream to produce an output map  $\mathbf{z}^f$ . We investigate the following fusion



**Fig. 1:** (a) Early fusion with two-stream C-RNN and (b) late fusion with two networks, CNN and RNN.

functions  $f$ : *sum fusion*, *max fusion*, and *concatenation fusion*. In addition, as  $\mathbf{z}^{\text{conv}}$  and  $\mathbf{z}^{\text{rec}}$  differ in size, we transform them via a fully-connected layer with *sigmoid* activation beforehand to make their sizes compatible (i.e. with sum fusion and max fusion) or to equalize their contributions to the output map (i.e. with concatenation fusion);

$$\tilde{\mathbf{z}}^{\text{conv}} = \mathbf{z}^{\text{conv}} \mathbf{W}^{\text{conv}} + \mathbf{b}^{\text{conv}}, \quad (6)$$

$$\tilde{\mathbf{z}}^{\text{rec}} = \mathbf{z}^{\text{rec}} \mathbf{W}^{\text{rec}} + \mathbf{b}^{\text{rec}}. \quad (7)$$

Here,  $\mathbf{W}^{\text{conv}} \in \mathbb{R}^{RQ \times M}$  and  $\mathbf{W}^{\text{rec}} \in \mathbb{R}^{H \times M}$  denote weight matrices while  $\mathbf{b}^{\text{conv}} \in \mathbb{R}^M$  and  $\mathbf{b}^{\text{rec}} \in \mathbb{R}^M$  denotes bias terms of the fully-connected layers.  $M$  is the desired size of the transformed feature vectors  $\tilde{\mathbf{z}}^{\text{conv}}$  and  $\tilde{\mathbf{z}}^{\text{rec}}$ .

**Sum fusion.**  $\mathbf{z}^{\text{sum}} = f_{\text{sum}}(\tilde{\mathbf{z}}^{\text{conv}}, \tilde{\mathbf{z}}^{\text{rec}})$  computes the sum of the two feature vectors  $\mathbf{z}^{\text{sum}} = \tilde{\mathbf{z}}^{\text{conv}} + \tilde{\mathbf{z}}^{\text{rec}}$ .

**Max fusion.**  $\mathbf{z}^{\text{max}} = f_{\text{max}}(\tilde{\mathbf{z}}^{\text{conv}}, \tilde{\mathbf{z}}^{\text{rec}})$  takes the maximum of the two feature vectors  $z_i^{\text{max}} = \max(z_i^{\text{conv}}, z_i^{\text{rec}})$ ,  $1 \leq i \leq M$ .

**Concatenation fusion.**  $\mathbf{z}^{\text{cat}} = f_{\text{cat}}(\tilde{\mathbf{z}}^{\text{conv}}, \tilde{\mathbf{z}}^{\text{rec}})$  simply concatenates the two feature vectors to make a larger one  $\mathbf{z}^{\text{cat}} = [\tilde{\mathbf{z}}^{\text{conv}}, \tilde{\mathbf{z}}^{\text{rec}}]$ .

Network training and evaluation are performed similarly to the standalone CNN and RNN with a softmax layer used for classification and minimization of the cross-entropy loss given in (3). In particular, dropout is further applied to the fusion features for regularization purpose. Again, the fusion features are extracted to train a linear SVM for classification after network training.

## 4 Late Fusion of CNN and RNN

For the late fusion scheme, we probabilistically combine the classification results of the standalone CNN and RNN in Sections 2.2 and 2.3, respectively, as illustrated in Fig. 1(b). We study three fusion methods: *max fusion*, *mean fusion*, and *multiplication fusion*. Let  $\mathbf{P}^{\text{conv}} = (P_1^{\text{conv}}, P_2^{\text{conv}}, \dots, P_C^{\text{conv}})$  and  $\mathbf{P}^{\text{rec}} = (P_1^{\text{rec}}, P_2^{\text{rec}}, \dots, P_C^{\text{rec}})$  denote the posterior probabilities obtained by the CNN and RNN, respectively. The classification likelihood  $\mathbf{P} = (P_1, P_2, \dots, P_C)$  after fusion with the three methods is then given by

$$P_c = \max(P_c^{\text{conv}}, P_c^{\text{rec}}) \quad \text{for } 1 \leq c \leq C, \quad (8)$$

$$P_c = \frac{1}{2}(P_c^{\text{conv}} + P_c^{\text{rec}}) \quad \text{for } 1 \leq c \leq C, \quad (9)$$

$$P_c = \frac{1}{2}(P_c^{\text{conv}} \times P_c^{\text{rec}}) \quad \text{for } 1 \leq c \leq C, \quad (10)$$

respectively. The final output label is then determined by likelihood maximization on the classification likelihood  $\mathbf{P}$ .

## 5 Experiments

### 5.1 LITIS-Rouen dataset and modification for this study

We conducted experiments using the LITIS-Rouen dataset [13]. The dataset consists of 19 scene categories with 3026 examples in total. All instances have the same length of 30 seconds and were recorded with a sampling rate of 22050 Hz.

For this study, we did not use the full 30-second snippets. Instead, we decomposed each 30-second snippet into segments of 4 seconds length without overlap (except the last segment). Thus  $S = 8$  such segments were obtained from each 30-second snippet. The classification models were trained with 4-second segments extracted from the training data. To understand how the classification performance varies with different test signal lengths, we sequentially evaluated a classification model on the 4-second segments of a 30-second snippet and aggregated the classification results over  $\{1, 2, \dots, S\}$  segments which are equivalent to  $\{4, 8, \dots, 28, 30\}$  seconds. In order to aggregate classification results over multiple segments in the late-fusion system, we used the same fusion method that we used for model fusion. For the standalone CNN, RNN, and the two-stream C-RNN, we employed the probabilistical multiplication fusion of segment-wise results, as it has been shown to be efficient for this purpose [28, 5].

For an experiment with a specific test signal length, we follow the training/testing splits in the seminal work [13] with the reported performances averaged over 20 splits.

### 5.2 Network training and parameters

A dropout rate of 0.5 and 0.1 was used for the CNN and the RNN, respectively. Particularly, for the early fusion C-RNN network, a dropout rate of 0.5 was further applied to the fusion features. For all networks, the regularization parameter  $\lambda$  was commonly set to  $10^{-3}$  and the network training was accomplished using the *Adam* optimizer [29] with a learning rate of  $10^{-4}$ . In addition, the hyper-parameter  $C$  of the SVMs used for classification after network training was fixed to 1.0.

**Table 1:** Overall performance of the classification models with different test signal lengths of  $\{4, 8, \dots, 28, 30\}$  seconds. Early fusion and late fusion are abbreviated by EF and LF, respectively.

Method	Accuracy (%)									F1-score (%)								
	4s	8s	12s	16s	20s	24s	28s	30s	Avg.	4s	8s	12s	16s	20s	24s	28s	30s	Avg.
CNN	91.7	94.3	95.6	96.0	96.3	96.9	97.4	97.5	95.7	91.4	93.9	95.3	95.8	96.1	96.8	97.3	97.4	95.5
RNN	91.2	94.5	95.6	96.2	96.6	<b>97.2</b>	<b>97.7</b>	<b>97.8</b>	95.8	90.8	94.2	95.4	96.0	96.5	<b>97.2</b>	<b>97.6</b>	<b>97.7</b>	95.7
EF - Sum	90.1	93.1	94.5	95.1	95.7	96.2	96.5	96.9	94.8	89.6	92.8	94.2	94.9	95.6	96.1	96.4	96.8	94.5
EF - Max	90.2	93.1	94.3	94.9	95.3	95.9	96.4	96.8	94.6	89.5	92.7	94.0	94.6	95.1	95.7	96.3	96.7	94.3
EF - Concat	90.4	93.1	94.5	95.2	95.5	96.1	96.6	96.8	94.8	90.0	92.7	94.1	94.8	95.3	95.9	96.5	96.7	94.5
LF - Max	92.0	94.6	95.4	95.9	95.9	96.6	96.9	97.0	95.5	91.6	94.2	95.1	95.7	95.8	96.4	96.8	96.9	95.3
LF - Mean	<b>92.1</b>	94.5	95.6	96.1	96.5	97.0	97.5	97.6	95.9	<b>91.7</b>	94.5	95.3	95.9	96.3	96.9	97.4	97.6	95.7
LF - Mult	<b>92.1</b>	<b>94.9</b>	<b>95.9</b>	<b>96.3</b>	<b>96.7</b>	<b>97.2</b>	<b>97.7</b>	<b>97.8</b>	<b>96.1</b>	<b>91.7</b>	<b>94.6</b>	<b>95.6</b>	<b>96.2</b>	<b>96.6</b>	97.1	<b>97.6</b>	<b>97.7</b>	<b>95.9</b>

## 5.3 Experimental results

### 5.3.1 Overall performance

Table 1 shows the overall performance in terms of accuracy and F1-score obtained by the studied classification systems over different test signal lengths  $\{4, 8, \dots, 28, 30\}$ . Fig. 2 provides a visual snapshot of variation of four systems, including the CNN, RNN, the best early-fusion system (i.e. EF-Sum) and the best late-fusion system (i.e. LF-Mult), over the spectrum of the test signal length.

On average, the standalone CNN marginally underperforms its RNN counterpart. Their performance gap is particularly noticeable with large test signal lengths ( $\geq 20$  seconds). These results highlight the importance of sequential modelling for temporal data. It may also be worthwhile noticing that the CNN proposed here achieves an F1-score of 97.4%, improving that of the over-time-convolution CNN in [2] (i.e. 96.5%) by 0.9% absolute. Regarding the fusion systems, sum fusion and concatenation fusion perform comparably in early fusion and appear to be better than max fusion. However, among the late fusion methods, multiplication fusion performs the best. Compared to other methods, multiplication fusion favours likelihood for categories that have consistent classification results, and more aggressively suppresses those with diverged classification results [28, 5].

More importantly, the results in Table 1 and Fig. 2 show that early fusion consistently worsens the classification performance. On average, the best early-fusion system (i.e. EF-Sum) reduces the F1-score by 1.0% and 1.2% absolute compared to the standalone CNN and RNN,


respectively. Late fusion is more efficient as LF-Mult yields performance gains of 0.4% and 0.2% absolute on the average F1-score compared to the standalone CNN and RNN, which was previously reported with state-of-the-art performance at 30-second test signal length on the experimental dataset [5].

### 5.3.2 When is model fusion useful?

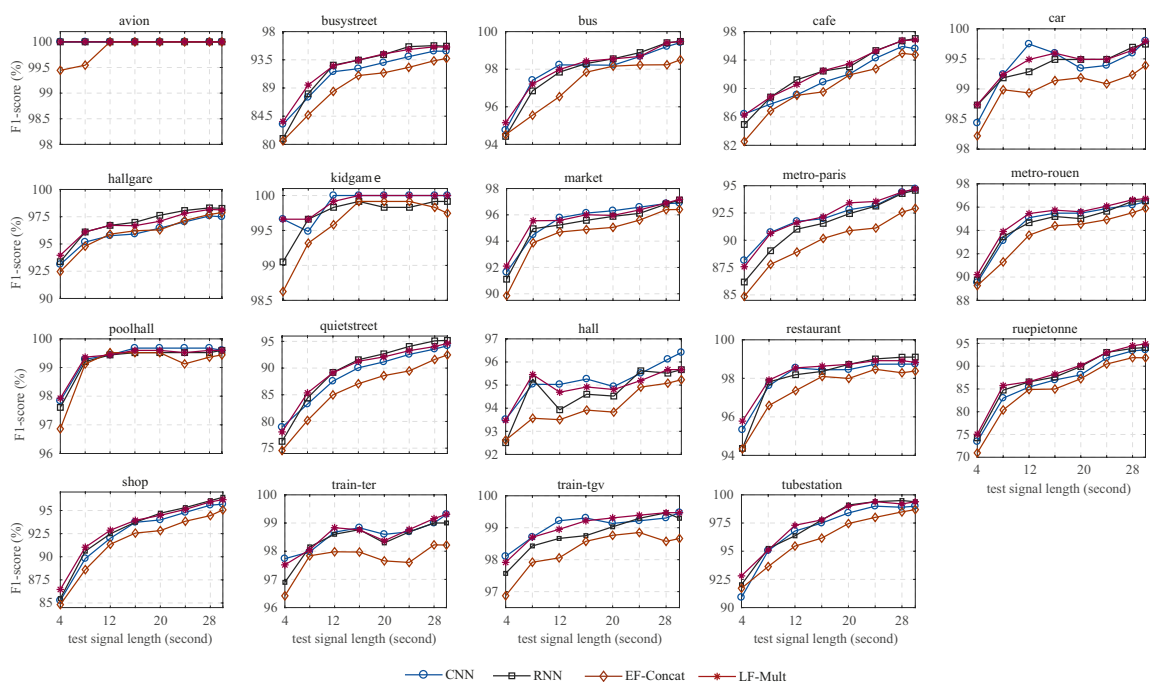
Inspection on Fig. 2 further reveals that model fusion is the most productive when the test signal length is small (i.e.  $< 20$  seconds). When the test signal length is larger than 20 seconds, the performance gain (if any) achieved by LF-Mult is very marginal. Intuitively, with short signal lengths, multiple views on a short duration scene can compensate each other in the ensemble. However, when listening longer, the best standalone system (namely RNN) has accumulated more information about a scene and can recognize it quite reliably. Meanwhile, the weaker standalone model (namely the CNN) does not appear to bring sufficient new information about the scene into the ensemble.

### 5.3.3 How long is sufficient to recognize a scene?

Fig. 3 shows the F1-score variation patterns for several different scene categories over the spectrum of test signal length. It is unsurprising to find that the patterns are very category-specific. Several categories can be reliably recognized even with a very short signal length. For example, we achieve perfect or near perfect recognition accuracy rate on “avion”, “kidgame”, and “poolhall” within 4, 12, and 16 seconds, respectively. Some other scenes, such as “car”, “metro-rouen”, “restaurant”, “train-ter”, and “train-tgv”, can



**Fig. 2:** Variation in overall performance of CNN, RNN, EF-Sum (i.e. the best early fusion system) and LF-Mult (i.e. the best late fusion system) over different test signal lengths. (a) Accuracy and (b) F1-score.



**Fig. 3:** Variation in category-specific F1-scores of CNN, RNN, EF-Sum (i.e. the best early fusion system) and LF-Mult (i.e. the best late fusion system) over the range of test signal lengths.

also be recognized reliably within 16 seconds, and accumulating more information about these scenes further improve performance. In contrast, scenes such as “cafe”, “quietstreet”, “ruepietonne”, and “shop” require much longer test signal lengths (e.g. 30 seconds) to

achieve good performance. It would be expected that a system should listen even longer (i.e. > 30 seconds) to improve the detection certainty for those scenes.

## 6 Future research

In this study, we made use of high-level LTE features that were learned to encode the hierarchy of the scene categories [20] to train the networks. Future research should look for an alternative method to encode this hierarchy with a neural network. Such a network should be trained jointly with the classification task in an end-to-end fashion. In addition, the networks were trained with audio segments of 4 seconds long and a coarse temporal resolution for signal decomposition and feature extraction (i.e. 250 ms frames) may overlook the fine structure of the signal. Using low-level features, such as log Mel-scale spectrograms [28], and employing smaller frames for signal decomposition may avoid this potential issue. They would also allow us to train a network with shorter audio segments, enabling studying audio scene recognition with variable lengths at a finer temporal resolution.

## 7 Conclusion

This work studied an important aspect of the audio scene recognition task, namely how much time is sufficient to recognize an audio scene. We additionally investigated which kind of model fusion and to which temporal extent model fusion is most useful for the task. These studies were accomplished by exploring ASC using a CNN, an RNN and their fusion using various early and late fusion strategies. Experimental results on the LITIS-Rouen dataset showed that the duration required to recognize a scene is category-specific, just as it is for humans. Some scenes were reliably recognized within a few seconds while others required significantly longer durations. Furthermore, while late fusion was shown to improve the classification performance, it is most beneficial when the test signal length is limited. These findings suggest that practical implementation of audio scene recognition systems should adapt their recognition strategies to different kinds of target scenes to improve their service quality.

## References

- [1] Stowell, D., Giannoulis, D., Benetos, E., Lagrange, M., and Plumbley, M. D., "Detection and Classification of Acoustic Scenes and Events," *IEEE Trans. Multimedia*, 17(10), pp. 1733–1746, 2015.
- [2] Phan, H., Hertel, L., Maass, M., Koch, P., Mazur, R., and Mertins, A., "Improved Audio Scene Classification based on Label-Tree Embeddings and Convolutional Neural Networks," *IEEE/ACM Trans. on Audio, Speech and Language Processing (TASLP)*, 25(6), pp. 1278–1290, 2017.
- [3] Lyon, R. F., "Machine Hearing: An Emerging Field," *IEEE Signal Processing Magazine*, 27(5), pp. 131–139, 2010.
- [4] Wang, D. and Brown, G. J., *Computational Auditory Scene Analysis: Principles, Algorithms, and Applications*, Wiley-IEEE Press, 2006.
- [5] Phan, H., Koch, P., Katzberg, F., Maass, M., Mazur, R., and Mertins, A., "Audio Scene Classification with Deep Recurrent Neural Networks," in *Proc. Interspeech*, pp. 3043–3047, 2017.
- [6] Eghbal-Zadeh, H., Lehner, B., Dorfer, M., and Widmer, G., "CP-JKU submissions for DCASE-2016: a hybrid approach using binaural I-vectors and deep convolutional neural networks," Technical report, Detection and Classification of Acoustic Scenes and Events 2016, 2016.
- [7] Mun, S., Park, S., Han, D., and Ko, H., "Generative Adversarial Network Based Acoustic Scene Training Set Augmentation and Selection Using SVM Hyper-Plane," in *Proc. Workshop on Detection and Classification of Acoustic Scenes and Events*, 2017.
- [8] Sakashita, Y. and Aono, M., "Acoustic Scene Classification by Ensemble of Spectrograms Based on Adaptive Temporal Divisions," Technical report, DCASE2018 Challenge, 2018.
- [9] Zhang, T., Zhang, K., and Wu, J., "Temporal transformer networks for acoustic scene classification," in *Proc. Interspeech 2018*, pp. 1349–1353, 2018.
- [10] Mesaros, A., Heittola, T., and Virtanen, T., "A multi-device dataset for urban acoustic scene classification," *arXiv preprint arXiv:1807.09840*, 2018.
- [11] Mesaros, A., Heittola, T., Benetos, E., Foster, P., Lagrange, M., Virtanen, T., and Plumbley, M. D., "Detection and Classification of Acoustic Scenes and Events: Outcome of the DCASE

- 2016 Challenge,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*, 26(2), pp. 379–393, 2018.
- [12] Mesaros, A., Heittola, T., Diment, A., Elizalde, B., Shah, A., Vincent, E., Raj, B., and Virtanen, T., “DCASE 2017 Challenge setup: Tasks, datasets and baseline system,” in *Proc. Workshop on Detection and Classification of Acoustic Scenes and Events*, 2017.
- [13] Rakotomamonjy, A. and Gasso, G., “Histogram of gradients of time-frequency representations for audio scene classification,” *IEEE/ACM Trans. Audio, Speech, and Language Processing*, 23(1), pp. 142–153, 2015.
- [14] Guo, J., Xu, N., Li, L.-J., and Alwan, A., “Attention based CLDNNs for short-duration acoustic scene classification,” in *Proc. Interspeech*, 2017.
- [15] Valenti, M., Diment, A., Parascandolo, G., Squarini, S., and Virtanen, T., “DCASE 2016 Acoustic Scene Classification Using Convolutional Neural Networks,” Technical report, Detection and Classification of Acoustic Scenes and Events 2016, 2016.
- [16] Phan, H., Koch, P., Hertel, L., Maass, M., Mazur, R., and Mertins, A., “CNN-LTE: A Class of 1-X Pooling Convolutional Neural Networks on Label Tree Embeddings for Audio Scene Classification,” in *Proc. ICASSP*, 2017.
- [17] Vu, T. H. and Wang, J.-C., “Acoustic Scene and Event Recognition Using Recurrent Neural Networks,” Technical report, Detection and Classification of Acoustic Scenes and Events 2016, 2016.
- [18] Bae, S. H., Choi, I., and Kim, N. S., “Acoustic Scene Classification Using Parallel Combination of LSTM and CNN,” Technical report, Detection and Classification of Acoustic Scenes and Events 2016, 2016.
- [19] Yin, Y., Shah, R. R., and Zimmermann, R., “Learning and Fusing Multimodal Deep Features for Acoustic Scene Categorization,” in *Proc. ACMMM*, pp. 1892–1900, 2018.
- [20] Phan, H., Hertel, L., Maass, M., Koch, P., and Mertins, A., “Label Tree Embeddings for Acoustic Scene Classification,” in *Proc. ACMMM*, pp. 486–490, 2016.
- [21] Ellis, D. P. W., “Gammatone-like spectrograms,” 2009.
- [22] Glorot, X., Bordes, A., and Bengio, Y., “Deep Sparse Rectifier Neural Networks,” in *Proc. 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 315–323, 2011.
- [23] Kim, Y., “Convolutional neural networks for sentence classification,” in *Proc. EMNLP*, pp. 1746–1751, 2014.
- [24] Phan, H., Hertel, L., Maass, M., and Mertins, A., “Robust audio event recognition with 1-max pooling convolutional neural networks,” in *Proc. Interspeech*, pp. 3653–3657, 2016.
- [25] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R., “Dropout: A Simple Way to Prevent Neural Networks from Overfitting,” *Journal of Machine Learning Research (JMLR)*, 15, pp. 1929–1958, 2014.
- [26] Cho, K., van Merriënboer, B., Gulcehre, C., Bougares, F., Schwenk, H., and Bengio, Y., “Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation,” in *Proc. EMNLP*, pp. 1724–1734, 2014.
- [27] Graves, A., Mohamed, A., and Hinton, G., “Speech recognition with deep recurrent neural networks,” in *Proc. ICASSP*, pp. 6645–6649, 2013.
- [28] Han, Y. and Lee, K., “Convolutional Neural Network with Multiple-Width Frequency-Delta Data Augmentation for Acoustic Scene Classification,” Technical report, DCASE2016 Challenge, 2016.
- [29] Kingma, D. P. and Ba, J. L., “Adam: a Method for Stochastic Optimization,” in *Proc. ICLR*, pp. 1–13, 2015.