

Impact of Acoustic Event Tagging on Scene Classification in a Multi-Task Learning Framework

Rahil Parikh¹, Harshvardhan Sundar², Ming Sun², Chao Wang², Spyros Matsoukas²

¹ University of Maryland, College Park, MD, USA

² Alexa AI, Amazon.com Inc, Cambridge, MA, USA

rahil@umd.edu, sundarhs@amazon.com

Abstract

Acoustic events are sounds with well-defined spectro-temporal characteristics which can be associated with the physical objects generating them. Acoustic scenes are collections of such acoustic events in no specific temporal order. Given this natural linkage between events and scenes, a common belief is that the ability to classify events must help in the classification of scenes. This has led to several efforts attempting to do well on Acoustic Event Tagging (AET) and Acoustic Scene Classification (ASC) using a multi-task network. However, in these efforts, improvement in one task does not guarantee an improvement in the other, suggesting a tension between ASC and AET. It is unclear if improvements in AET translates to improvements in ASC. We explore this conundrum through an extensive empirical study and show that under certain conditions, using AET as an auxiliary task in the multi-task network consistently improves ASC performance. Additionally, ASC performance further improves with the AET data-set size and is not sensitive to the choice of events or the number of events in the AET data-set. We conclude that this improvement in ASC performance comes from the regularization effect of using AET and not from the network's improved ability to discern between acoustic events.

Index Terms: acoustic scene classification, acoustic event tagging, multi-task learning, joint-training, auxiliary task

1. Introduction

Computational Environmental Audio Analysis [1] aims to extract and interpret information related to the environment from which an acoustic signal is recorded. Two major sub-classes of problems in this field are Acoustic Event Tagging (AET) [2–5] and Acoustic Scene Classification (ASC) [6–9]. AET is a multi-label classification problem which involves classifying an audio sample into one or more predefined events. ASC is a multi-class classification problem in which an audio sample is classified into one of the predefined acoustic scenes that represents the environment from which the audio sample was recorded. Acoustic events represent information at lower levels of abstraction with clear time-frequency patterns such as ‘car engine’, ‘dog-bark’, etc., while scenes are collection of acoustic events in no specific temporal order and represent information at higher levels of abstraction such as ‘street traffic’ and ‘urban park’.

Studies in auditory perception and cognition suggest that humans leverage event information for scene classification [10]. For instance, knowledge of the event ‘jet-engine’ helps classify a given acoustic scene as ‘airport’ instead of ‘shopping mall’. This has inspired several researchers to integrate ASC with AET or Acoustic Event Detection (AED) [11–15]. These works are grounded on the assumption that a model’s capability to classify or detect events (i.e. to perform AET or AED) must help in its ability to classify scenes (i.e. to perform ASC). Various

frameworks and model architectures have been implemented in an attempt to use the AET/AED task to improve ASC performance. In [15] the authors explore pre-training a network on the AET task and then fine-tuning it to perform ASC. Jointly training for AED/AET and ASC using a multi-task learning framework has been explored in [11–14]. These works employ a common encoder and task-specific decoders for ASC and AET. Here, the network is trained on a linear combination of the individual task losses, thereby presenting multiple hyper-parameters that need to be tuned. In the absence of an existing data set with both event and scene annotations, most works use task-specific individual data sets from the DCASE Challenges to train each task. However, in each of these works, researchers have observed that jointly training with AET does not guarantee an improvement in ASC performance. In [14], various multi-task learning model configurations are explored to leverage event and scene cross-information sharing for ASC, AET and AED tasks, despite which, the improvements in performance are often marginal.

The marginal improvements obtained from jointly-training on AET and ASC tasks suggest a ‘tension’ between the two tasks which prevents them from simultaneous improvements. A likely explanation for this tension arises from the limitations of the human annotation process and data collection [16, 17]. To elaborate, data sets are commonly annotated with ‘foreground’ sounds, such as *speech*, *distinct musical instruments* etc. These sounds are informative of the acoustic events. In contrast, ‘background’ sounds such as *wind*, *distant home appliances* etc. which are more informative of the scene are seldom annotated. In a model jointly trained to optimize ASC and AET performance, an improvement in the AET performance may come at the cost of the ASC performance since low SNR regions are ignored. Similarly, the ASC data set contains several un-annotated events for each scene which may or may not be present in the AET data set. Jointly-training a network on such data may degrade the AET performance. This tension questions the improvement of the scene classification task when jointly-trained with AET, and suggests the presence of an optimal weighing of the objective function for the AET and ASC tasks to obtain this improvement, if any.

The works in [11–14] are aimed towards improving ASC and AET performance by jointly-training the network on fixed, pre-defined AET and ASC tasks. Instead, in this paper, we probe the tension between AET and ASC tasks in a multi-task learning framework by fixing the ASC task and analyzing the impact on this fixed ASC task by changing various attributes of the AET task. To facilitate this, we create multiple AET data sets by varying the distribution of event labels, the size of the data set and the number of the event labels and examine its effect on the fixed ASC task.

Through such empirical studies we show that:

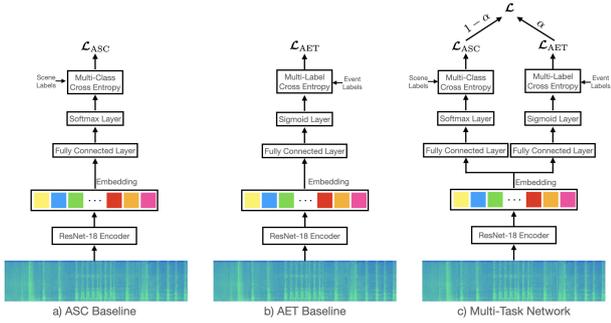


Figure 1: *The ASC and AET baselines along with the multi-task network used in this work.*

- Performance improvements in AET task are not in general correlated with improvements in ASC task. With AET and ASC tasks fixed, the performance improvement in AET **does not** guarantee performance improvement in ASC.
- Consistent improvements in ASC (up to 5% accuracy compared to a competitive baseline) are observed when the multi-task loss function is balanced w.r.t. the contribution of the individual AET and ASC losses.
- Improvement in ASC depends on the AET training data set size and is less sensitive to the number of events being classified or the choice of events used in the AET task.

Thus, our empirical studies reveal that improvements in ASC can be attributed to the regularization effect of using AET as an auxiliary task. Interestingly, the regularization effect of using AET to improve ASC is more effective than some of the well known regularization strategies of Mix-up [18], Dropout [19], and SpecAugment [20].

The factors like AET training data size, choice and the number of events used can be used as parameters to control the regularization effect. We emphasize that the motivation of our work is not to improve scene classification performance using a multi-task learning framework, but instead to investigate the impact of AET as an auxiliary task on scene classification with empirical evidence to analyze what controls on the AET task improve ASC performance.

2. Multi-task Learning for ASC and AET

We jointly train our network to perform ASC and AET using a multi-task learning approach. The two tasks share an encoder that generates an embedding from the input features. This is followed by the ASC head— a fully-connected layer with softmax activation and an AET head— a fully-connected layer with sigmoid activation. A ResNet-18 [21] architecture is used as our common encoder. We use log-mel-spectrograms as input features of 10s long audios. Figure 1 shows the individual task specific baselines along with the multi-task network used for analysis.

The network is trained on the overall loss (\mathcal{L}) which is a convex combination of the individual AET and ASC losses parameterized by α as:

$$\mathcal{L} = \alpha \cdot \mathcal{L}_{\text{AET}} + (1 - \alpha) \cdot \mathcal{L}_{\text{ASC}}; \alpha \in [0, 1] \quad (1)$$

where \mathcal{L}_{AET} and \mathcal{L}_{ASC} are the multi-label and multi-class classification losses of AET and ASC respectively. Higher the value

assigned to α , higher will be the contribution of the AET loss to the overall loss. During training, each mini-batch contains a fixed number of samples from the AET and ASC data sets. We experimented with different number of samples in the set $\{8, 16, 32, 64, 128, 256\}$ per mini-batch for the ASC data set keeping the number of samples from the AET data set fixed at 128. Best results were obtained for a mini-batch size of 192 samples with 64 samples from ASC and 128 samples from AET data sets. We used this setting for all other experiments reported in this paper. We use SpecAugment [20] and MixUp [18] to augment our training data for regularization.

3. Experimental Setup

3.1. Datasets

We use the TUT Urban Acoustic Scenes data set from the DCASE 2018 Task-1 challenge [22] for the ASC task. This balanced data set consists of 24 hours of audio from 10 urban scenes. Examples of these scenes are: ‘airport’, ‘street traffic’, ‘metro’, etc. Each sample in the data set has a unique scene corresponding to it. We keep this data set fixed throughout our analysis for comparability of ASC performance.

We sample AudioSet [23] to create multiple AET data sets by varying the choice of events in the data set, the size of the data set (T) and number of events in the data set (k). In its entirety, AudioSet is annotated for 530 event-classes and is 5600 hours long. Each sample in AudioSet may contain multiple event-classes.

Data from the AET data set is used to train the common encoder and the AET head while data from the ASC data set is used to train the encoder and the ASC head.

3.2. Evaluation Metrics

We evaluate our models for scene-classification and event-tagging using the metrics followed by the DCASE ASC (Task-1, Editions: 2018-’21) and DCASE AET (Task 2, Editions: 2018-’19) Challenge: **Macro-Average Accuracy (Accuracy)** and **Label-Weighted Label-Ranking Average Precision (lwl-rap)** respectively. A higher value suggests better model performance for both metrics.

3.3. Common Encoder Architecture and Baseline System

The task specific baselines for ASC and AET are shown in Figure 1 a) and 1 b) respectively. The ASC baseline was trained using only the TUT Urban Acoustic Scenes data set with SpecAugment [20] and MixUp [18] techniques for regularization and \mathcal{L}_{ASC} as the loss function. We obtain a scene classification accuracy of 68% with this approach which is comparable to the state-of-art single-channel, single model accuracy for the same data set [24]. It is worthwhile to note that, we obtain similar performance with the multi-task network shown in Figure 1 c) with $\alpha = 0$. While we tried other common architectures such as VGG [25] and CRNN variants [26], we observe that the ResNet-18 outperforms them. Therefore, we use ResNet-18 as the common encoder for all our subsequent experiments as discussed in Section 2.

The AET baseline also employs a ReNet-18 encoder and is trained on several different subsets of AudioSet. The performance of this baseline will be specified based on the specific subset of AudioSet in the following sections.

4. Analysing the Impact of AET on ASC

In this section we analyze the impact of changing the AET task in various ways on the ASC task. We run 4 trials for each experiment, initialized with 4 different random-seeds.

4.1. Impact of α on ASC Performance

We investigate the impact of changing the relative contribution of the two individual task-specific losses \mathcal{L}_{AET} and \mathcal{L}_{ASC} to the overall loss given in (1) by varying α . We perform this analysis on the multi-task network shown in Figure 1 c), using different values of α incremented by 0.1 in the range $[0, 0.9]$ and by 0.01 in $[0.90, 1]$. The AET data set is chosen to comprise of 35 randomly selected event classes with a total size of 170 hours. Figure 2 shows the scene classification accuracy and event tagging performance (lwlap) of the multi-task network along with the performance of the individual task specific baselines. At $\alpha = 0$, the model is only trained on \mathcal{L}_{ASC} . The scene accuracy of the multi-task network is thus equivalent to the ASC baseline model while the multi-task network’s event tagging performance is equivalent to *random chance*. Similarly, at $\alpha = 1.0$, the multi-task network is only trained on \mathcal{L}_{AET} and its the event tagging performance is similar to that of the AET baseline while its performance on scene classification is at chance. As α increases from 0 to 1, the event tagging performance (lwlap) of the multi-task network increases. Strangely, the multi-task network outperforms the ASC baseline in terms of scene classification accuracy for higher values of $\alpha \in [0.8, 1.0]$. We repeat this experiment on 6 separate AET data sets sampled from AudioSet each with 35 randomly chosen events with the total number of hours in the range of 11 to 360 hours. In each of these cases, we observe that the performance of the multi-task network improves over the ASC baseline for higher values of α . On closer examination, we observe that upon convergence $\mathcal{L}_{\text{AET}}^* \approx \mathcal{L}_{\text{ASC}}^*/20$, where $\mathcal{L}_{\text{AET}}^*$ and $\mathcal{L}_{\text{ASC}}^*$ represent the individual task losses upon convergence. Additionally, the range of α over which we observe consistent improvements of the multi-task network over the ASC baseline, the individual loss terms are nearly equal, i.e.,

$$\alpha \cdot \mathcal{L}_{\text{AET}}^* \approx (1 - \alpha) \cdot \mathcal{L}_{\text{ASC}}^* \quad (2)$$

$$\therefore \alpha \approx \frac{\mathcal{L}_{\text{ASC}}^*}{\mathcal{L}_{\text{ASC}}^* + \mathcal{L}_{\text{AET}}^*} \quad (3)$$

As a simplifier, we use $\mathcal{L}_{\text{ASC}}^*$ and $\mathcal{L}_{\text{AET}}^*$ as converged losses from the training of the respective task specific baselines and fix $\alpha = 0.95$ for the remainder of the experiments.

Interestingly, from Figure 2, we observe that for a multi-task network trained with $\alpha = 0.5$ the event lwlap is around 76% whereas for a multi-task network trained with $\alpha = 0.$, the event lwlap is around 10%. Thus, compared to the multi-task network trained with $\alpha = 0$, the multi-task network trained with $\alpha = 0.5$ is a much better event tagger. However, this improved event tagging ability does not translate to improved scene classification accuracy of the multi-task network. Thus, contrary to the assumptions made in [12, 14], our experiments show strong evidence that such multi-task network’s improved scene classification does not come for the ability of the network to do better event tagging but rather comes about from the regularization effect of using AET as a challenging auxiliary task. The value of α chosen to balance the two losses contributes to making AET a challenging auxiliary task while the same network is required to also do well on the scene classification task. Additionally, we believe AET is an even more effective regularizer

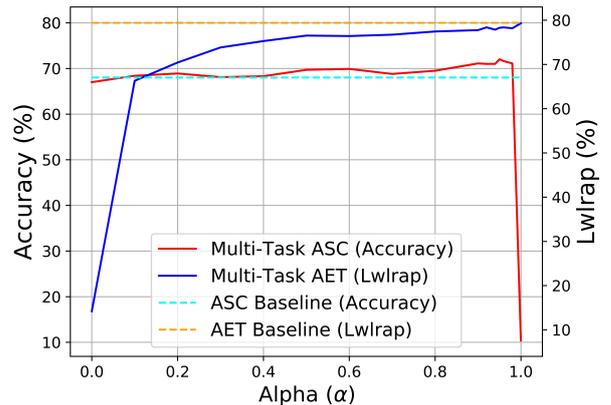


Figure 2: Joint-Training with Different α using the Correlated Event-Set with $k = 35$, $T = 170$ hours.

for the ASC tasks than the generally useful SpecAugment [20] and MixUp [18] techniques.

4.2. Impact of Choice of Events in the AET Task on ASC Performance

To analyze the impact of the choice of event classes in the AET task on ASC, we sample k event-classes with a data set size of 170 hours from AudioSet. The following event-class sampling strategies are explored to evaluate the impact of choosing events used to train the AET task:

1. Random Selection: k events are randomly sampled from the set of 530 event-classes present in AudioSet.
2. Correlated Event Selection: Each clip from the TUT Urban Acoustic Scenes data set used for the ASC task may contain multiple events. However, these events are not annotated as these clips are only associated with scene labels. We wish to train the AET task on the k events that are most likely to be present in this data set. To predict these events in the ASC data set, we use a trained fine-grained event tagger. This tagger is a DenseNet [27] trained on over 3.58M samples from the Amazon Instant Video (AIV) data set and Amazon Alexa’s internal production data set for 203 classes. On AIV and the internal production data set, this classifier has a Top-1 accuracy of 34% and 99.98% respectively, and a Top-5 accuracy of 65% and 99.99% respectively. We choose the k event-classes in AudioSet with the highest prediction scores since they are most-likely to be present in ASC data. Examples of such events are ‘speech’, ‘traffic noises’, ‘footsteps’, ‘jet-engine’ etc.
3. Uncorrelated Event Selection: To sample the k events least likely to be present in the TUT Urban Acoustic Scenes data set, we use the above-mentioned event tagger and choose the k events in AudioSet with the lowest prediction scores on the ASC data set. These events are least-likely to appear in ASC data and include events such as ‘roaring cats’, ‘orchestra’, ‘gunshots’, etc.

Our AET data sets thus consists of a subset of AudioSet data corresponding to k classes, sampled using either of the above-mentioned strategies.

Figure 3(a) shows that multi-task learning model using the AET event-classes obtained from either of the sampling strate-

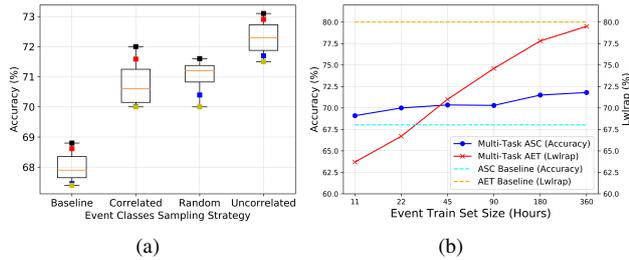


Figure 3: Studying ASC performance by modifying the AET task: (a) Events sampled using three different sampling strategies for $k = 35, T = 170$ hours. (b) AET training data sets of different sizes (T) for $k = 35$. Events are randomly-sampled.

gies outperforms the baseline ASC system. It is noteworthy that this improvement in performance is uncorrelated with the choice of the events used to train the AET task.

4.3. Impact of AET Dataset Size on ASC Performance

To compare the multi-task learning model performance when trained on AET data sets of different sizes (T), we randomly sample $k = 35$ events and vary the amount of training data corresponding to these events, ranging from 11 to 360 hours. The AET validation data set is kept fixed across all models for fair comparison between the AET performance. Figure 3(b) illustrates that the multi-task learning model outperforms the baseline ASC system for all values T . Interestingly, ASC accuracy continues to increase with an increase in the event training data size. The AET data set size and ASC accuracy have a Pearson Correlation Coefficient [28] $\rho = 0.89$, suggesting a strong positive correlation between the ASC performance and the size of the AET data set. As expected, the lwlrup also increases with an increase in training data for AET.

4.4. Impact of Number of Evens in AET Dataset on ASC Performance

To analyze the impact of the number of events (k) in the AET data set on ASC performance, we prepare AET data sets with k ranging from 2 to 115, as illustrated in Figure 4 for $T=180$ hours. The event-classes are randomly sampled for $k > 2$. For $k = 2$, we train the AET task using the event class ‘speech’ and ‘music’, which are most abundant in AudioSet. We observe that even when the AET task is trained to classify only between ‘music’ and ‘speech’ events, the multi-task learning scene classifier outperforms the baseline model. We also observe that after $k = 35$, increasing the number of event classes further has marginal influence on the scene accuracy. In general, the ASC performance has a weak positive correlation with the number of events in the AET data sets, $\rho = 0.24$.

4.5. Comparing Pre-Training with Joint-Training

An alternate framework to leverage AET to improve ASC is to pre-train a network to perform AET and then fine-tune it for ASC [15]. In Figure 5, we observe that both frameworks outperform the baseline and joint-training is marginally better than fine tuning. Additionally, joint-training is more scalable and can be simultaneously used for both tasks, while the pre-training is sequential and eventually can be used only for ASC.

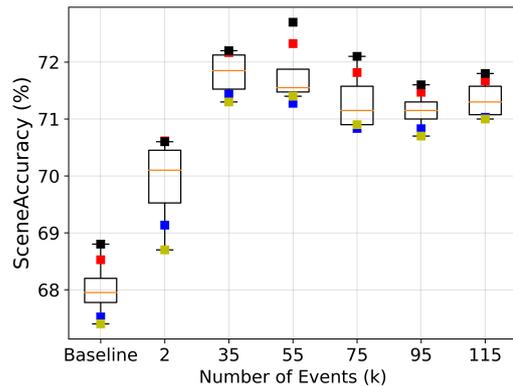


Figure 4: Joint-Training with different number of event-classes (k) in the AET data-set. T is 180 hours in all experiments.

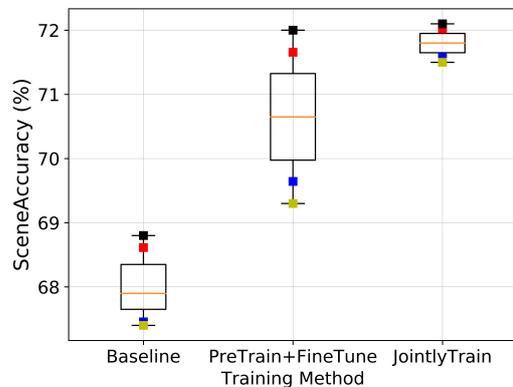


Figure 5: Comparison Between Pre-Training and Multi-Task Learning Frameworks

5. Conclusion

We demonstrate empirically that scene classification performance can be consistently improved using event tagging as an auxiliary task in a multi-task learning framework when the individual task losses contribute equally to the overall loss. Through ablation studies on carefully handcrafting the data used for the auxiliary AET task, we show that the performance of the main ASC task improves with the size of the AET data set while being agnostic to the specific choice of events or the number of events used in the AET task. Contrary to previous works, we show empirically that this improvement in ASC performance cannot be attributed to the network gaining knowledge of events or its ability to discern between events. Instead, it is most likely due to the regularization effect of using AET as an auxiliary task. This work opens the door to deeper analysis on the interplay between ASC and AET.

6. Acknowledgements

We thank Viktor Rozgic, Chieh-Chi Kao and Qingming Tang at Amazon Inc. for their feedback and suggestions on the experiment design.

7. References

- [1] T. Virtanen, M. D. Plumbley, and D. P. W. Ellis, *Computational analysis of sound scenes and events*. Springer, 2018.
- [2] Y. Wang, J. Li, and F. Metzger, “A comparison of five multiple instance learning pooling functions for sound event detection with weak labeling,” in *Proc. Int’l. Conf. on Acoust., Speech and Sig. Process. (ICASSP)*. IEEE, 2019, pp. 31–35.
- [3] C. C. Kao, B. Shi, M. Sun, and C. Wang, “A joint framework for audio tagging and weakly supervised acoustic event detection using densenet with global average pooling,” *arXiv:2008.03350 preprint*, 2020.
- [4] E. Fonseca, M. Plakal, F. Font, D. P. W. Ellis, X. Favory, J. Pons, and X. Serra, “General-purpose tagging of Freesound audio with Audioset labels: Task description, dataset, and baseline,” *arXiv preprint arXiv:1807.09902*, 2018.
- [5] Q. Kong, Y. Xu, W. Wang, and M. D. Plumbley, “A joint detection-classification model for audio tagging of weakly labelled data,” in *Proc. Int’l. Conf. on Acoust., Speech and Sig. Process. (ICASSP)*. IEEE, 2017, pp. 641–645.
- [6] D. Barchiesi, D. Giannoulis, D. Stowell, and M. D. Plumbley, “Acoustic scene classification: Classifying environments from the sounds they produce,” *IEEE Sig. Process. Magazine*, vol. 32, no. 3, pp. 16–34, 2015.
- [7] S. H. Bae, I. Choi, and N. S. Kim, “Acoustic scene classification using parallel combination of lstm and cnn,” in *Proc. Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2016, pp. 11–15.
- [8] D. Battaglino, L. Lepauloux, N. Evans, F. Mougins, and F. Biot, “Acoustic scene classification using convolutional neural networks,” *IEEE AASP Challenge on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2016.
- [9] J. T. Geiger, B. Schuller, and G. Rigoll, “Large-scale audio feature extraction and svm for acoustic scene classification,” in *Workshop on Applications of Signal Processing to Audio and Acoustics*. IEEE, 2013, pp. 1–4.
- [10] C. Guastavino, “Current trends in urban soundscape research,” in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, vol. 261, no. 1. Institute of Noise Control Engineering, 2020, pp. 5857–5868.
- [11] J. Jung, H. Shim, J. Kim, S. Kim, and H. Yu, “Acoustic scene classification using audio tagging,” *arXiv:2003.09164 preprint*, 2020.
- [12] N. Tonami, K. Imoto, M. Niitsuma, R. Yamanishi, and Y. Yamashita, “Joint analysis of acoustic events and scenes based on multitask learning,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2019, pp. 338–342.
- [13] H. L. Bear, I. Nolasco, and E. Benetos, “Towards joint sound scene and polyphonic sound event recognition,” *arXiv:1904.10408 preprint*, 2019.
- [14] J. Jung, H. Shim, J. Kim, and H. Yu, “DCASENet: An integrated pretrained deep neural network for detecting and classifying acoustic scenes and events,” in *Proc. Int’l. Conf. on Acoust., Speech and Sig. Process. (ICASSP)*. IEEE, 2021, pp. 621–625.
- [15] R. Zhang, W. Zou, X. Li, and D. Chuxing, “Cross-task pre-training for acoustic scene classification,” *arXiv:1910.09935v1 e-prints*, 2019.
- [16] J. Abeßer, “A review of deep learning based methods for acoustic scene classification,” *Applied Sciences*, vol. 10, no. 6, 2020.
- [17] Y. Wu and T. Lee, “Enhancing sound texture in cnn-based acoustic scene classification,” *arXiv:1901.01502 e-prints*, 2019.
- [18] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” *arXiv:1710.09412 preprint*, 2017.
- [19] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, no. 56, pp. 1929–1958, 2014. [Online]. Available: <http://jmlr.org/papers/v15/srivastava14a.html>
- [20] D. S. Park, W. Chan, Y. Zhang, C. C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *arXiv preprint arXiv:1904.08779*, 2019.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comp. Vision and Patt. Recog.*, 2016, pp. 770–778.
- [22] A. Mesaros, T. Heittola, and T. Virtanen, “A multi-device dataset for urban acoustic scene classification,” in *Proc. Detection and Classification of Acoustic Scenes and Events (DCASE) Workshop*, November 2018, pp. 9–13. [Online]. Available: <https://arxiv.org/abs/1807.09840>
- [23] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *Proc. Int’l. Conf. on Acoust., Speech and Sig. Process. (ICASSP)*. IEEE, 2017, pp. 776–780.
- [24] O. Mariotti, M. Cord, and O. Schwander, “Exploring deep vision models for acoustic scene classification,” *Proc. Detection and Classification of Acoustic Scenes and Events (DCASE) Workshop*, pp. 103–107, 2018.
- [25] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv:1409.1556 preprint*, 2014.
- [26] B. Shi, X. Bai, and C. Yao, “An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition,” *IEEE Trans. Patt. Analysis and Machine Intelligence*, vol. 39, no. 11, pp. 2298–2304, 2016.
- [27] G. Huang, Z. Liu, L. V. D. Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” *arXiv:1608.06993*, 2018.
- [28] A. V. Oppenheim, R. W. Schaffer, and J. R. Buck, “Discrete-time signal processing,” *Prince Hall, Sec.*, vol. 11, 1999.