

Acoustic Scene Classification Using Multichannel Observation with Partially Missing Channels

Keisuke Imoto
Doshisha University, Japan
keisuke.imoto@ieee.org

Abstract—Sounds recorded with smartphones or IoT devices often have completely missing parts due to microphone failure and packet loss in data transmission over the network, and partially unreliable observations caused by clipping, wind noise. In this paper, we investigate the impact of the partially missing channels on the performance of acoustic scene classification using multichannel audio recordings, especially for a distributed microphone array. Missing observations cause not only losses of time-frequency and spatial information on sound sources but also a mismatch between a trained model and evaluation data. We thus investigate how a missing channel caused by the microphone failure and packet loss affects the performance of acoustic scene classification in detail. We also propose simple data augmentation methods for scene classification using multichannel observations with partially missing channels and evaluate the scene classification performance using the data augmentation methods.

Index Terms—Acoustic scene classification, multichannel processing, missing observation, data augmentation

I. INTRODUCTION

Acoustic scene classification (ASC), which classifies sound recordings into the predefined class such as recording environments, places, and daily activities, is one of the core search problems in environmental sound analysis [1]–[3]. ASC has significant potential for various applications such as monitoring infants/elderly people [4], automatic surveillance [5], automatic life-logging [6], and media retrieval [7].

Many methods for ASC utilizing spectral information have been proposed. For instance, Eronen *et al.* [8] and Mesaros *et al.* [9] have proposed methods based on mel-frequency cepstral coefficients (MFCCs) and Gaussian mixture models (GMMs). Valenti *et al.* [10], Han *et al.* [11], and Jallet *et al.* [12] have proposed methods using mel-spectrograms and a convolutional neural network (CNN). Liping *et al.* [13], Tanabe *et al.* [14], and Raveh and Amar [15] have proposed Xception-based, VGG-based, and ResNet-based ASC methods, respectively.

More recently, environmental sound analysis utilizing spatial information, which is extracted from time differences or sound power ratios between channels, has also been studied [14], [16]–[20]. Conventional microphone array processing requires that microphones are synchronized between channels and/or microphone locations or array geometry is known. However, spatial information based on accurate time differences or sound power ratios between channels cannot be extracted using a combination of unsynchronized distributed

microphones such as smartphones, IoT devices, and surveillance cameras. To utilize unsynchronized distributed microphones whose locations or array geometry is unknown for multichannel ASC, Kürby *et al.* [21] have proposed a method based on the late fusion of scene classification results obtained with each microphone. Many conventional methods for multichannel ASC also apply this strategy [22], [23]. Imoto *et al.* have proposed ASC methods using the spatial cepstrum and graph cepstrum that can be applied under an unsynchronized condition [18], [20].

On the other hand, sounds recorded with smartphones or IoT devices often have missing parts caused by microphone failure, packet loss in data transmission over the network, or unreliable observations caused by clipping and wind noise. To analyze acoustic scenes from intermittently missing observations with a single-channel microphone, Imoto and Ono have proposed a method of simultaneously analyzing acoustic scenes and estimating missing observations [24]. However, the conventional method is not for multichannel audio recordings, and the impact of partially missing channels on the ASC performance using multichannel audio recordings has not been investigated in the conventional works. In this paper, we thus investigate the impact of partially missing channels on the performance of multichannel ASC under the condition that several microphones do not completely work.

In machine-learning-based multichannel ASC, missing channels cause not only losses of time-frequency and spatial information on sound sources but also a mismatch between a trained model and evaluation data. Therefore, to realize a robust ASC system, it is important to investigate how a missing channel affects the ASC performance. We then apply simple data augmentation methods for multichannel ASC with partially missing channels and evaluate the scene classification performance using the data augmentation methods.

The remainder of this paper is organized as follows. In section 2, we discuss conventional acoustic scene classification using multichannel observation. In section 3, we introduce three simple data augmentation methods for multichannel ASC with missing channels. In section 4, we report the results of experiments carried out to evaluate the performance of ASC with partially missing channels and the impact of missing channels on the ASC performance. Finally, we conclude this paper in section 5.

II. CONVENTIONAL METHODS FOR SCENE CLASSIFICATION

Let us consider a model f and model parameter θ . The purpose of ASC is to estimate an acoustic scene label \hat{z} in an evaluated sound as

$$\hat{z} = \arg \max_z f(\mathbf{X}, \theta), \quad (1)$$

where z and \mathbf{X} ($\in \mathbb{R}^{F \times T \times C}$) are the acoustic scene class and acoustic feature, respectively. F , T , and C are the numbers of frequency bins, time frames, and channels, respectively. The model parameter θ is preliminarily determined using the training dataset $\mathcal{D} = \{(\mathbf{X}_1, z_1), \dots, (\mathbf{X}_l, z_l), \dots, (\mathbf{X}_L, z_L)\}$. Here, \mathbf{X}_l is the acoustic feature of the l th sound clip and z_l indicates an acoustic scene label in the l th sound clip. For the acoustic feature \mathbf{X}_l , the mel-band energy and MFCCs are often used. As the model f , GMMs, a CNN, a ResNet-based, or a VGG-based method has often been applied. In the neural-network-based methods, the model parameter θ is estimated using the softmax cross-entropy loss function and the backpropagation technique.

Most conventional methods assume that there is no missing channel in a multichannel observation. However, in the scenario of a distributed microphone array, we may have partially missing channels caused by microphone failure, in which some acoustic feature \mathbf{X}_c in the c th channel cannot be utilized in the evaluation data.

III. DATA AUGMENTATION FOR MULTICHANNEL SCENE CLASSIFICATION

In this work, we apply three data augmentation methods for multichannel scene classification with partially missing channels. These data augmentation methods are reasonably simple to implement and enable us to investigate how partially missing channels affect the ASC performance.

A. Channel Mask

The data missing in the evaluation stage causes a mismatch between the trained model and evaluation data. To avoid this mismatch, we apply simple binary masking throughout the input time-frequency features for the random channels in the model training stage as follows:

$$\mathbf{X}_{l,c} = O, \quad (2)$$

where $\mathbf{X}_{l,c}$ is the acoustic feature of the l th sound clip in the c th channel. O is the zero matrix when the acoustic feature is the linear spectrum, whereas it is the matrix that has negative infinity values in its element when the acoustic feature is the log spectrum.

TABLE I
NUMBER OF RECORDED AUDIO SEGMENTS

Acoustic scene	# segments
Absence	4,874
Cooking	1,281
Dishwashing	356
Eating	577
Other	515
Social activity	1,236
Vacuum cleaning	243
Watching TV	4,662
Working	4,661
Total	18,405

TABLE II
DETAILED NETWORK ARCHITECTURE USED FOR EVALUATION EXPERIMENTS

Layer	Input size	Output size
Conv. ($7 \times 1 \times 64$) + BN + ReLU	$40 \times 501 \times 16$	$40 \times 501 \times 64$
Max pooling (4×1) + Dropout (rate=0.2)	$40 \times 501 \times 64$	$10 \times 501 \times 64$
Conv. ($10 \times 1 \times 128$) + BN + ReLU	$10 \times 501 \times 64$	$10 \times 501 \times 128$
Conv. ($1 \times 7 \times 256$) + BN + ReLU	$10 \times 501 \times 128$	$10 \times 501 \times 256$
Global max pooling + Dropout (rate=0.5)	$10 \times 501 \times 256$	256
Dense	256	128
Softmax	128	9

TABLE III
SCENE CLASSIFICATION PERFORMANCE WITH MISSING CHANNELS IN EVALUATION DATASET

	w/o missing	1ch	2ch	4ch	8ch	12ch
Micro-Fscore	92.26%	83.37%	72.82%	57.54%	43.30%	28.62%
Macro-Fscore	88.59%	64.27%	42.81%	34.79%	29.92%	26.10%

TABLE IV
SCENE CLASSIFICATION PERFORMANCE WITH SAME MISSING CHANNELS IN TRAINING AND EVALUATION DATASETS

	1ch	2ch	4ch	8ch	12ch
Micro-Fscore	91.47%	90.98%	91.20%	89.26%	84.76%
Macro-Fscore	86.86%	84.49%	84.60%	81.25%	77.02%

B. Channel Overwrite and Random Copy

When applying *channel mask*, a large difference in acoustic features may remain between the unmasked channel in the training data and the missing channel in the evaluation data. This may not fully mitigate the mismatch between the trained model and evaluation data. To bridge this gap, we apply a data augmentation method using *channel overwrite* in the model training stage and a random copy in the evaluation stage. *Channel overwrite* mandatorily overwrites the time-frequency features between channels in model training as follows:

$$\mathbf{X}_{l,c} = \mathbf{X}_{l,c'}, \quad (3)$$

where c and c' are different channel indices. In the evaluation stage, we randomly copy the acoustic features from non-missing channels to missing channels.

TABLE V
MICRO-FSCORE FOR CLASSIFICATION PERFORMANCE WITH DATA AUGMENTATION.
“*n*CH MISSING” DENOTES THE NUMBER OF MISSING CHANNELS IN THE EVALUATION DATASET.

	w/o missing	1ch missing	2ch missing	4ch missing	8ch missing	12ch missing
w/o augmentation	92.26%	83.37%	72.82%	57.54%	43.30%	28.62%
<i>Channel mask</i>	91.65%	91.46%	91.39%	90.94%	87.70%	78.13%
<i>Channel overwrite</i> + Random copy	91.89%	91.70%	91.56%	90.03%	87.82%	84.32%
<i>Channel swap</i> + Random copy	91.93%	91.86%	91.75%	90.29%	88.72%	87.59%

TABLE VI
MACRO FSCORE FOR CLASSIFICATION PERFORMANCE WITH DATA AUGMENTATION
“*n*CH MISSING” DENOTES THE NUMBER OF MISSING CHANNELS IN THE EVALUATION DATASET.

	w/o missing	1ch missing	2ch missing	4ch missing	8ch missing	12ch missing
w/o augmentation	88.59%	64.27%	42.81%	34.79%	29.92%	26.10%
<i>Channel mask</i>	86.40%	86.83%	86.22%	85.18%	79.91%	64.61%
<i>Channel overwrite</i> + Random copy	86.58%	86.86%	86.37%	86.77%	80.21%	79.88%
<i>Channel swap</i> + Random copy	86.72%	86.90%	86.61%	86.13%	82.27%	81.12%

C. Channel Swap and Random Copy

Channel mask and *channel overwrite* lose time-frequency information since we discard time-frequency features in the training stage. To train the scene classification model without wasting time-frequency information, we apply a data augmentation method using *channel swap*. *Channel swap* simply swaps the time-frequency features between channels in model training as follows:

$$\begin{cases} \mathbf{X}_{l,c} = \mathbf{X}_{l,c'} \\ \mathbf{X}_{l,c'} = \mathbf{X}_{l,c}. \end{cases} \quad (4)$$

In the evaluation stage, we randomly copy the acoustic features from non-missing channels to missing channels as with the data augmentation in *channel overwrite*.

IV. EXPERIMENTS

A. Experimental Conditions

We evaluate the impact of a missing channel on the performance of scene classification using various data augmentation methods. To evaluate the performance, we use the SINS dataset [26] and the experimental setting in the DCASE2018 Challenge Task 5 [25]. The SINS is a database for the detection of daily activities in a home environment from sounds recorded by a distributed microphone array. From the SINS dataset, we first select four microphone arrays (array 1, 4, 6, 8 in the SINS dataset), each of which consists of four linearly arranged microphones; that is, each audio segment contains 16 channels. Then, we calculate cross-correlation coefficients between *Node1_xx_audio.wav* in the SINS dataset and *DevNode1_xx.wav* in the development dataset of DCASE2018 Challenge Task 5 to determine the reference time positions of audio segments. After that, we split the continuous recordings of the array 1, 4, 6, 8 into 16 channel/10 s audio segments based on the reference time positions. As shown in Table I, the dataset contains 18,405 audio segments with nine acoustic

scene classes, and we split them into a similar 4-fold cross-validation setup as in DCASE2018 Challenge Task 5, in which we use the metadata of *DevNode1_xx.wav*.

For the acoustic features, we use the 40-dimensional log mel-band energy, which has a frame length of 40 ms with hop size of a 20 ms. In this paper, we select missing channels randomly and regard the missing channels as silent with zeros filled in the time domain. As the classification model, we apply the same network proposed by Inoue *et al.* [22], which achieved the best score in DCASE2018 Challenge Task 5, except for the input channel size of the network. The detailed network structure is shown in Table II. We utilize the RAdam optimizer [27] with a learning rate of 0.001. For each method, we conduct the evaluation experiment 16 (random combinations of missing channels) \times 4 (fold) times.

B. Experimental Results

1) *Impact of Missing Channel on Classification Performance*: We evaluate the performance degradation caused by missing channels in the evaluation data. Table III shows the scene classification performance in terms of micro- and macro-Fscore. The result shows that the missing channels cause severe performance degradation in multichannel ASC.

To investigate how the missing channels affect the ASC performance, we also evaluate the ASC performance with the same channels missing in the training and evaluation datasets. Table IV shows the scene classification performance in terms of micro- and macro-Fscore. The results show that the performance degradation is significantly smaller than the corresponding results in Table III, even though some channels are missing in the training dataset. This indicates that, in multichannel ASC, the mismatch between the trained model and evaluation data is a much more severe problem than missing spectral and spatial information. Thus, the mismatch between the trained model and evaluation data must be addressed preferentially in multichannel ASC with partially missing channels.

Absence	98.6	0.0	0.0	0.0	0.2	0.1	0.0	0.0	1.1
Cooking	0.0	97.6	0.0	1.5	0.3	0.6	0.0	0.0	0.0
Dishwashing	0.0	6.6	80.2	11.0	0.0	2.2	0.0	0.0	0.0
Eating	1.4	0.0	1.4	94.4	0.7	2.1	0.0	0.0	0.0
Other	30.3	0.8	0.0	3.9	60.2	1.6	0.0	0.0	3.2
Social activity	1.0	0.3	0.0	0.0	1.0	97.7	0.0	0.0	0.0
Vacuum clearing	0.0	0.0	0.0	0.0	0.0	0.0	100.0	0.0	0.0
Watching TV	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0	0.0
Working	1.9	0.0	0.0	0.1	0.1	0.2	0.0	0.0	97.7

Fig. 1. Example of scene classification result without missing channels (recall, %)

Absence	67.2	0.0	0.0	0.0	32.8	0.0	0.0	0.0	0.0
Cooking	0.0	0.0	0.0	0.0	98.2	0.0	0.0	0.0	1.8
Dishwashing	0.0	0.0	13.8	0.0	77.4	0.0	0.0	0.0	8.8
Eating	4.2	0.0	0.0	0.0	93.8	0.0	0.0	0.0	0.0
Other	17.7	0.0	0.0	0.0	74.7	0.0	0.0	0.0	7.6
Social activity	2.8	0.0	0.0	0.0	97.2	0.0	0.0	0.0	0.0
Vacuum clearing	0.0	0.0	0.0	0.0	100.0	0.0	0.0	0.0	0.0
Watching TV	0.0	0.0	0.0	0.0	39.8	0.0	0.0	54.6	5.6
Working	29.2	0.0	0.0	0.0	32.6	0.0	0.0	2.2	36.0

Fig. 2. Example of scene classification result with four missing channels in evaluation data (recall, %)

Absence	98.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.9
Cooking	0.3	98.2	0.0	0.0	1.2	0.0	0.0	0.0	0.3
Dishwashing	0.0	15.1	73.9	2.2	6.6	0.0	0.0	0.0	2.2
Eating	2.1	0.0	1.4	95.8	0.0	0.0	0.0	0.0	0.7
Other	15.0	0.0	0.0	0.8	67.7	3.9	0.0	0.0	12.6
Social activity	1.5	0.0	1.1	2.7	5.9	76.2	0.0	0.4	4.2
Vacuum clearing	0.0	0.0	0.0	0.0	0.0	0.0	100.0	0.0	0.0
Watching TV	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0	0.0
Working	4.1	0.0	0.0	0.0	0.3	0.0	0.0	0.0	95.6

Fig. 3. Example of scene classification result based on *channel swap* with four missing channels in evaluation data (recall, %)

2) *Evaluation of Data Augmentation Technique for Multichannel ASC*: We next evaluate the ASC performance with the proposed data augmentation methods. In this experiment, we randomly select a number of channels from 0 to 8 for data augmentation in each iteration of model training. Tables V and VI show the scene classification performance in terms of micro- and macro-Fscore with the proposed data augmentation methods. The results show that the three data augmentation methods achieve reasonable performances. In particular, *channel overwrite* and *channel swap* achieve comparable ASC performance to the result without missing channels. Comparing these results with Table IV indicates that *channel overwrite* and *channel swap* can almost completely avoid the mismatch between the trained model and evaluation data. On the other hand, the scene classification performance when using *channel mask* is lower than that of *channel overwrite* and *channel swap*. This is because the correlation between acoustic features of non-missing channels is higher than that between missing and non-missing channels; thus, copying the acoustic features from non-missing channels to missing channels is a better complement than using the zero-padding data.

3) *Detailed Scene Classification Performance*: Figs. 1–3 show the detailed scene classification performance using no data augmentation method and *channel swap* with four channels missing. The results show that most of the audio segments are predicted as “other.” On the other hand, the classification result using *channel swap* achieves comparable performance to that without missing channels. From these results, we conclude that, in multichannel ASC, partially missing channels may cause a severe degradation of the ASC performance, and avoiding the mismatch between the trained model and evaluation data is important to achieve a robust ASC system in a realistic situation.

V. CONCLUSION

In this paper, we investigated the impact of partially missing channels caused by microphone failure or packet loss on the ASC performance using multichannel audio recordings. We also proposed three data augmentation methods for multichannel ASC: *channel mask*, *channel overwrite*, and *channel swap*. The experimental results showed that, in multichannel ASC, the mismatch between the trained model and evaluation data is a much more severe problem than missing spectral and spatial information. To avoid this negative impact on the ASC performance, the data augmentation based on *channel overwrite* and *channel swap* is effective and can avoid the performance degradation caused by the model mismatch. In more practical situations, there are also unreliable observations such as heavily clipped, distorted, and noise-polluted sounds, and scene classification methods with these unreliable observations need to be addressed in future works.

ACKNOWLEDGMENT

This work was supported by JSPS KAKENHI Grant Number JP20H00613.

REFERENCES

- [1] T. Virtanen, M. Plumbley, and D. Ellis, Eds., *Computational Analysis of Sound Scenes and Events*. Springer, 2017.
- [2] K. Imoto, “Introduction to acoustic event and scene analysis,” *Acoustical Science and Technology*, vol. 39, no. 3, pp. 182–188, 2018.
- [3] T. Heittola, A. Mesaros, and T. Virtanen, “Acoustic scene classification in DCASE 2020 challenge: Generalization across devices and low complexity solutions,” *Proc. Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, pp. 56–60, 2020.
- [4] A. Mesaros, T. Heittola, A. Diment, B. Elizalde, A. S. E. Vincent, B. Raj, and T. Virtanen, “DCASE 2017 challenge setup: Tasks, datasets and baseline system,” *Proc. Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, pp. 85–92, 2017.
- [5] S. Ntalampiras, I. Potamitis, and N. Fakotakis, “On acoustic surveillance of hazardous situations,” *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 165–168, 2009.

- [6] K. Imoto and S. Shimauchi, "Acoustic scene analysis based on hierarchical generative model of acoustic event sequence," *IEICE Trans. Inf. Syst.*, vol. E99-D, no. 10, pp. 2539–2549, 2016.
- [7] Q. Jin, P. F. Schulam, S. Rawat, S. Burger, D. Ding, and F. Metzger, "Event-based video retrieval using audio," *Proc. INTERSPEECH*, pp. 2085–2088, 2012.
- [8] A. Eronen, V. T. Peltonen, J. T. Tuomi, A. P. Klapuri, S. Fagerlund, T. Sorsa, G. Lorho, and J. Huopaniemi, "Audio-based context recognition," *IEEE Trans. Audio Speech Lang. Process.*, vol. 14, no. 1, pp. 321–329, 2006.
- [9] A. Mesaros, T. Heittola, A. Eronen, and T. Virtanen, "Acoustic event detection in real life recordings," *Proc. European Signal Processing Conference (EUSIPCO)*, pp. 1267–1271, 2010.
- [10] M. Valenti, S. Squartini, A. Diment, G. Parascandolo, and T. Virtanen, "A convolutional neural network approach for acoustic scene classification," *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 1547–1554, 2017.
- [11] Y. Han, J. Park, and K. Lee, "Convolutional neural networks with binaural representations and background subtraction for acoustic scene classification," *Proc. Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, pp. 1–5, 2017.
- [12] H. Jallet, E. Çakır, and T. Virtanen, "Acoustic scene classification using convolutional recurrent neural networks," *Proc. Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, pp. 1–5, 2017.
- [13] Y. Liping, C. Xinping, and T. Lianjie, "Acoustic scene classification using multi-scale features," *Tech. Rep. DCASE Challenge Task 1*, 2018.
- [14] R. Tanabe, T. Endo, Y. Nikaido, T. Ichige, P. Nguyen, Y. Kawaguchi, and K. Hamada, "Multichannel acoustic scene classification by blind dereverberation, blind source separation, data augmentation, and model ensembling," *Tech. Rep. DCASE Challenge Task 5*, 2018.
- [15] A. Raveh and A. Amar, "Multi-channel audio classification with neural network using scattering transform," *Tech. Rep. DCASE Challenge Task 5*, 2018.
- [16] H. Kwon, H. Krishnamoorthi, V. Berisha, and A. Spanias, "A sensor network for real-time acoustic scene analysis," *Proc. IEEE International Symposium on Circuits and Systems*, pp. 169–172, 2009.
- [17] P. Giannoulis, A. Brutti, M. Matassoni, A. Abad, A. Katsamanis, M. Matos, G. Potamianos, and P. Maragos, "Multi-room speech activity detection using a distributed microphone network in domestic environments," *Proc. 23rd European Signal Processing Conference (EUSIPCO)*, pp. 1271–1275, 2015.
- [18] K. Imoto and N. Ono, "Spatial cepstrum as a spatial feature using distributed microphone array for acoustic scene analysis," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 25, no. 6, pp. 1335–1343, 2017.
- [19] K. Nakadai and D. R. Onishi, "Partially-shared convolutional neural network for classification of multi-channel recorded audio signals," *Tech. Rep. DCASE Challenge*, 2018.
- [20] K. Imoto, "Graph cepstrum: Spatial feature extracted from partially connected microphones," *IEICE Trans. Inf. Syst.*, vol. E103-D, no. 03, pp. 631–638, 2020.
- [21] J. Kürby, R. Grzeszick, A. Plinge, and G. A. Fink, "Bag-of-features acoustic event detection for sensor networks," *Proc. Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, pp. 55–59, 2016.
- [22] T. Inoue, P. Vinayavekhin, S. Wang, D. Wood, N. Greco, and R. Tachibana, "Domestic activities classification based on CNN using shuffling and mixing data augmentation," *Tech. Rep. DCASE Challenge Task 5*, 2018.
- [23] H. Liu, F. Wang, X. Liu, and D. Guo, "An ensemble system for domestic activity recognition," *Tech. Rep. DCASE Challenge Task 5*, 2018.
- [24] K. Imoto and N. Ono, "Acoustic topic model for scene analysis with intermittently missing observations," *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 27, no. 2, pp. 367–382, 2019.
- [25] G. Dekkers, L. Vuegen, T. Waterschoot, B. Vanrumste, M. Verhelst, and P. Karsmakers, "DCASE 2018 challenge - task 5: Monitoring of domestic activities based on multi-channel acoustics," *Tech. Rep. DCASE Challenge Task 5*, 2018.
- [26] G. Dekkers, S. Lauwereins, B. Thoen, M. W. Adhana, H. Brouckxon, T. Waterschoot, B. Vanrumste, M. Verhelst, and P. Karsmakers, "The SINS database for detection of daily activities in a home environment using an acoustic sensor network," *Proc. Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, pp. 32–36, 2017.
- [27] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," *Proc. International Conference on Learning Representations (ICLR)*, pp. 1–13, 2020.