

Device-Robust Acoustic Scene Classification via Impulse Response Augmentation

Tobias Morocutti², Florian Schmid¹, Khaled Koutini², Gerhard Widmer^{1,2}

¹*Institute of Computational Perception*, ²*LIT Artificial Intelligence Lab*

Johannes Kepler University, Linz, Austria

{tobias.morocutti, florian.schmid}@jku.at

Abstract—The ability to generalize to a wide range of recording devices is a crucial performance factor for audio classification models. The characteristics of different types of microphones introduce distributional shifts in the digitized audio signals due to their varying frequency responses. If this domain shift is not taken into account during training, the model’s performance could degrade severely when it is applied to signals recorded by unseen devices. In particular, training a model on audio signals recorded with a small number of different microphones can make generalization to unseen devices difficult. To tackle this problem, we convolve audio signals in the training set with pre-recorded device impulse responses (DIRs) to artificially increase the diversity of recording devices. We systematically study the effect of DIR augmentation on the task of Acoustic Scene Classification using CNNs and Audio Spectrogram Transformers. The results show that DIR augmentation in isolation performs similarly to the state-of-the-art method Freq-MixStyle. However, we also show that DIR augmentation and Freq-MixStyle are complementary, achieving a new state-of-the-art performance on signals recorded by devices unseen during training.

Index Terms—Recording Device Generalization, Impulse Response Augmentation, Freq-MixStyle, Acoustic Scene Classification

I. INTRODUCTION

Common machine learning theory builds on the assumption that training and test data are drawn from the same underlying distribution [1]. This assumption is often violated in practical applications, causing a performance drop at application time [2]. A well-known instance of this problem in the audio domain is the distributional shift caused by the physical characteristics of recording devices [3]. Digitizing an audio signal using a specific type of microphone results in the device-specific characteristics being encoded into the digitized signal. The problem is particularly severe if the training dataset is recorded using a limited number of recording devices while the application requires the model to generalize to devices unseen during training.

This problem has been extensively studied in the context of Acoustic Scene Classification (ASC) as part of the annual DCASE challenges [3], [4]. ASC is the task of predicting an acoustic scene (e.g. traveling by bus, traveling by tram, indoor shopping mall, urban park) given an audio clip. The TAU

Urban Acoustic Scenes 2022 Mobile development dataset [3] used in Task 1 of the DCASE’22 challenge [4] uses a Soundman OKM II Classic/Studio A3 as the main recording device referred to as *device A*. The training set further involves two real devices (*device B*: Samsung Galaxy S7, *device C*: GoPro Hero5 Session) and three simulated devices (*S1-S3*). The training subset is imbalanced containing 102,150 samples of device A while other devices are represented by only 7,500 samples. The test subset is balanced across all devices and involves three devices (*S4-S6*) unseen at training time. Previous challenge results [3], [4] have underlined the difficulty to generalize across different recording devices, showing test accuracy gaps of up to 15 percentage points between device A and unseen devices.

Prior efforts to improve device generalization involve suppressing device-specific information of audio signals [5], [6] or increasing the device diversity in the training set by device balancing [7] or augmentation [8]. As part of the latter, we propose to augment audio signals using a dataset consisting of 66 pre-recorded Impulse Responses (IR) from different microphones. Convolution of the audio signals with these device IRs (DIRs) artificially increases the recording device variety in the training dataset.

We validate the performance impact of our proposed DIR augmentation on different architectures including Convolutional Neural Networks (CNNs) and Transformers. With the receptive-field regularized CNN *CP-ResNet* [9], [10] and the Patchout faSt Spectrogram Transformer (*PaSST*) [11], we choose two state-of-the-art ASC models and evaluate them on the TAU Urban Acoustic Scenes 2020 Mobile development (*TAU20*) [3] and the TAU Urban Acoustic Scenes 2022 Mobile development (*TAU22*) [3] datasets using the official train-test split.

The main contribution of this work is a systematic analysis of the effect of DIR augmentation on the outlined models and datasets. We study the impact of DIR augmentation on the generalization abilities to unseen recording devices and on the overall performance. We compare DIR augmentation to Freq-MixStyle, the device generalization technique that lead to the best accuracies on unseen devices in the ASC task of the DCASE’22 challenge [4]. The results show that DIR augmentation consistently improves the generalization performance of models, especially when tested on short clips and in combination with the state-of-the-art method

The LIT AI Lab is supported by the Federal State of Upper Austria. GW’s work is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement No 101019375 (Whither Music?).

Freq-MixStyle [8].¹

II. RELATED WORK

A. Recording Device Generalization

Many different directions and methods have been explored to increase a model’s robustness against recording devices. In this regard, Koutini et al. [5] apply the classical Domain Adaptation objectives Maximum Mean Discrepancy [12] and Sliced Wasserstein Distance [13] to match recording device distributions and force the model to learn device-invariant high-level embeddings. Primus et al. [14] additionally exploit parallel recordings to create domain-invariant hidden layer representations. Lee et al. [7] balance the device distribution by randomly dropping samples of the dominant *device A*. Device translators are built by Kim et al. [15] to artificially generate samples for underrepresented recording devices. Relaxed instance frequency-normalization is introduced by the same authors [6] to eliminate device-specific information. Freq-MixStyle [6], [8], a method that mixes frequency statistics of spectrograms, led to the best generalization performance on unseen devices in the ASC task of the DCASE’22 challenge [16].

B. Impulse Response Augmentation

Impulse response augmentation is a very common technique to augment audio signals and increase a model’s robustness against different acoustic environments by simulating their reverberation. In this regard, Szöke et al. [17] publish a dataset of room impulse responses (RIRs) and show that it can be used to improve automatic speech recognition (ASR) when carefully selecting a set of RIRs that match the target environment. Ko et al. [18] use simulated and real RIRs to create robust models for far-field ASR. Similarly, [19] and [20] use RIRs to create reverberation robust speech recognition systems. Compared to RIRs, device impulse responses (DIRs) are less commonly used. Ferràs et al. [21] improve speaker-recognition performance by simulating degraded recordings with a variety of augmentation techniques including convolution with DIRs. TADA [22], an audio augmentation pipeline including background noise addition and convolution with RIRs and DIRs, shows to improve a model’s robustness against environmental distortions for audio classification tasks. According to our knowledge, our work is the first detailed study to show that convolution with DIRs drastically improves the generalization to unseen devices for state-of-the-art ASC models.

III. DEVICE IR AUGMENTATION

As a source of impulse responses, we use 66 freely available DIRs from MicIRP². The DIRs are recorded using a swept-sine method with the microphones placed at a distance of about 20 to 30 cm from the source. MicIRP consists of DIRs of vintage microphones, such as the Toshiba Type G. As

the Frequency Magnitude Response in Figure 1 shows, the devices used have a very specific sonic signature, making them a perfect source for simulating a dataset with high device variability.

During training, we augment samples from *device A*, the main recording device of the datasets *TAU20* [3] and *TAU22* [3], with a probability of p_{dir} using DIR augmentation. We choose one of the 66 DIRs at random and convolve the waveform with the DIR, cutting off trailing samples to not change the overall length of the waveform.

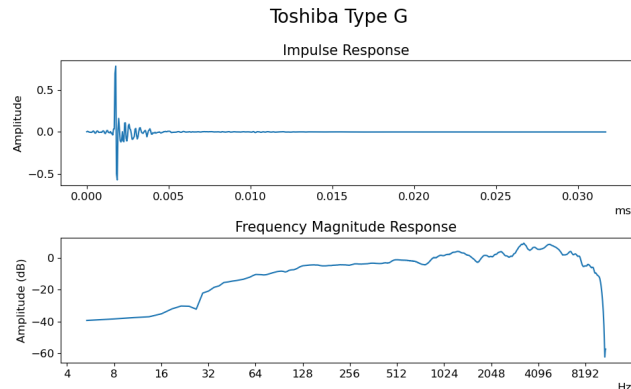


Fig. 1. Impulse response and frequency magnitude response of a *Toshiba Type G* recording device.

IV. EXPERIMENTAL SETUP

We evaluate our proposed method based on two different datasets and three different models. The two datasets *TAU20* [3] and *TAU22* [3] are identical except for the fact that the samples of *TAU20* have a duration of ten seconds, while *TAU22* is split into one-second clips, which results in a more difficult classification task. *TAU20* and *TAU22* are recorded with a limited amount of devices and the official train-test split strongly encourages models to have cross-device generalization capabilities. In the following, we describe the architectures used in the experiments, the audio preprocessing, the training setup and the device generalization techniques we compare our proposed method to.

A. Architectures

As a first architecture, we use *CP-ResNet*, a receptive-field regularized CNN (RFR-CNN) optimized for ASC [9], [10]. *CP-ResNet* has proven its effectiveness in previous editions of the DCASE ASC challenge [5], [16], [23]. We use *CP-ResNet* with the number of base channels set to 128, resulting in a model with approximately 4M parameters. Secondly, we build a low-complexity version of the *CP-ResNet* by reducing the number of blocks and the number of channels per layer. The resulting model has less than 128K parameters and less than 30 million multiply-accumulate operations (MACs), satisfying the complexity requirements for the ASC task of the DCASE’22 challenge [16]. We denote this model *LC CP-ResNet*. Finally, we use the Audio Spectrogram Transformer

¹Code available on GitHub: <https://github.com/theMoro/DIRAugmentation>

²<http://micirp.blogspot.com>

PaSST [11], which is a complex, self-attention-based model, pre-trained on AudioSet [24], that recently has been shown to achieve state-of-the-art performances on multiple downstream tasks, including ASC [11]. *PaSST* consists of 85M parameters.

B. Audio Preprocessing

For *CP-ResNet* and *LC CP-ResNet*, we use mono audio with a sampling rate of 22.05kHz. For computing spectrograms, a window size of 2048 samples (≈ 93 ms) and a hop size of 512 samples (≈ 23 ms) is applied. A Mel-scaled filterbank is used to create mel spectrograms with 256 mel bins. Regarding *PaSST*, we match the AudioSet [24] pre-training and use audio at a sampling rate of 32kHz, a hop size of 320 samples (10 ms), a window size of 800 samples (25 ms) and 128 mel bins to compute the mel spectrograms.

C. Training Setup

We use AdamW [25] as an optimizer with a learning rate schedule including an exponential warm-up, a linear decrease, and a fine-tuning phase to train all models. We use peak learning rates of 2×10^{-4} and 3×10^{-5} for *CP-ResNet* and *LC CP-ResNet*, respectively, and train both models for 50 epochs using a batch size of 64. Fine-tuning of *PaSST* on ASC requires only 25 epochs to converge and uses a peak learning rate of 1×10^{-5} . We use batch sizes of 80 and 8 for the datasets *TAU22* and *TAU20*, respectively. To avoid overfitting on the downstream task, we apply a structured patchout of 6 on the frequency dimension, similar to [8].

D. Device Generalization Methods

We compare the device generalization abilities of our proposed DIR augmentation method (referenced as *DIR* in short) to Mixup [26] and Freq-MixStyle (*FMS*) [6], [8]. Mixup creates virtual training samples by linearly mixing two spectrograms and their labels, which has been shown to improve generalization on ASC tasks before [27]. *FMS* showed the best generalization capabilities to unseen devices in the latest edition of the ASC task in the DCASE'22 challenge [4]. Similar to the original MixStyle [28] which mixes channel-wise statistics to improve domain generalization, the goal of *FMS* is to enhance device generalization by mixing frequency-wise statistics, as the device-style in spectrograms primarily resides in them [6]. Similar to *DIR*, *FMS* is guided by a probability p_{fms} that determines the augmentation strength by specifying the proportion of batches to augment. Additionally, we experiment with using *DIR* and *FMS* in combination (*DIR + FMS*). In this case, we first augment the waveforms using the *DIRs* and apply *FMS* afterward.

E. Device Generalization Score

To quantify device generalization abilities, we introduce the device generalization score (*DG-Score*). We define this measure as the standard deviation of the individual device accuracies. Therefore, a low *DG-Score* indicates a balanced performance of the model across different recording devices.

V. RESULTS

Table I presents the main results of this paper showing the performance of *CP-ResNet*, *LC CP-ResNet* and *PaSST* on the datasets *TAU20* [3] and *TAU22* [3] resulting in six different setups. Each model is evaluated without device generalization methods (*baseline*), with *Mixup*, with *DIR*, with *FMS* and with the combination of *DIR* and *FMS* (*DIR + FMS*). Each setup is presented in terms of individual device accuracies, accuracy on all real, simulated and unseen devices, overall accuracy and the *DG-Score* introduced in Section IV-E. We refer to the accuracy on the unseen devices as the *unseen accuracy*. The results presented are averages over three independent runs and the last five epochs of training.

A. Comparison to other Device Generalization Methods

Table I shows that experiments including *DIR* augmentation achieve the highest overall accuracy on five out of the six setups. The results indicate that *DIR* and *FMS* are complementary. The combination leads to the highest overall accuracy on three setups and is especially powerful for generalizing to unseen devices, giving the highest unseen accuracy on five out of the six setups and always leading to the lowest *DG-Score*. *Mixup* performs best for *PaSST* on *TAU20* both in terms of unseen and overall accuracy, as analyzed in more detail in Section V-B3.

B. Effect on different Models

1) *CP-ResNet*: On both datasets, the best-performing single method for *CP-ResNet* in terms of overall accuracy is *FMS*. *DIR* cannot catch up with the performance improvement achieved by *FMS* but outperforms *Mixup* and substantially improves upon the baseline increasing the overall accuracy by approximately 2.5 percentage points on both datasets. However, the combination of *DIR* and *FMS* improves the overall accuracy slightly compared to using only *FMS* and increases the unseen accuracy by more than one percentage point on both datasets, leading to the best *DG-scores*. In comparison to the baseline, *DIR + FMS* increases the accuracy on the unseen devices by approximately 11 and 7.5 percentage points for the datasets *TAU22* and *TAU20*, respectively. Additionally, using *FMS + DIR* improves upon the much more complex *PaSST* model on *TAU22* and outperforms the best *PaSST* results reported in [8], the approach with the best performance in the DCASE'22 challenge.

2) *Low-Complexity CP-ResNet*: As for the *CP-ResNet*, the low-complexity network achieves the best performance on both datasets in the setups including *DIR*. Combining *DIR* and *FMS* improves the results on *TAU22* further while *DIR* alone achieves the best overall accuracy on *TAU20*. *DIR + FMS* achieves the best results on the unseen devices boosting the accuracy by multiple percentage points on both datasets compared to the baseline. However, we also note that the gain in accuracy on unseen devices comes at the cost of decreasing accuracy on *device A* for some setups and methods and is particularly severe when using the *DIR + FMS* combination on the *LC CP-ResNet*.

	Method	Real Devices			Simulated Devices				Unseen Devices				Overall	DG	
		A	B	C	Real	S1	S2	S3	Sim	S4	S5	S6			Unseen
TAU22	CP-ResNet	70.76	63.20	63.25	65.74	56.47	52.62	57.57	55.55	51.52	48.67	44.66	48.28	56.52	7.73
	+ Mixup	70.76	62.36	62.97	65.37	58.43	54.34	59.90	57.56	53.09	50.78	48.31	50.73	57.88	6.62
	+ DIR	70.94	63.01	64.09	66.02	59.42	56.54	59.29	58.42	55.70	53.46	49.47	52.88	59.10	6.00
	+ FMS	71.67	61.34	63.97	65.66	62.05	58.62	63.32	61.33	60.02	57.85	55.91	57.93	61.64	4.34
	+ DIR + FMS	69.22	61.38	64.09	64.90	62.49	58.89	62.65	61.35	59.46	61.20	57.23	59.30	61.84	3.30
	LC CP-ResNet	67.51	55.15	57.36	60.01	51.73	48.30	53.19	51.07	46.86	43.25	39.28	43.13	51.40	7.86
	+ Mixup	66.21	53.92	57.17	59.11	50.71	49.39	52.93	51.01	48.02	44.70	40.27	44.33	51.48	7.04
	+ DIR	66.04	53.95	60.65	60.22	54.30	49.72	54.97	53.00	48.08	49.99	43.44	47.17	53.46	6.40
	+ FMS	66.53	54.59	57.41	59.52	50.78	49.23	54.26	51.42	48.19	47.59	41.87	45.89	52.27	6.65
	+ DIR + FMS	64.98	53.76	59.95	59.57	53.55	49.98	55.71	53.08	52.01	52.53	43.44	49.33	53.99	5.71
	PaSST	71.75	63.38	66.95	67.36	57.6	56.34	58.44	57.46	56.95	57.42	53.37	55.91	60.24	5.60
	+ Mixup	72.00	63.36	67.69	67.69	58.3	55.91	58.45	57.55	57.05	57.45	54.50	56.33	60.52	5.60
	+ DIR	71.56	62.58	67.70	67.29	60.35	58.62	60.16	59.71	59.39	59.45	56.53	58.46	61.82	4.55
	+ FMS	71.77	62.85	66.64	67.09	60.08	57.11	59.55	58.91	58.81	58.51	56.60	57.97	61.32	4.71
	+ DIR + FMS	70.21	62.03	66.66	66.30	60.96	58.30	59.78	59.68	60.51	59.89	57.25	59.22	61.73	3.93
TAU20	CP-ResNet	81.92	69.71	75.66	75.77	68.63	65.19	70.04	67.95	66.91	60.91	56.08	61.30	68.34	7.18
	+ Mixup	80.83	72.71	75.04	76.19	69.41	66.89	71.70	69.33	65.56	63.90	57.66	62.37	69.29	6.42
	+ DIR	79.98	72.16	75.85	76.00	72.46	68.14	72.61	71.07	68.57	67.72	60.48	65.59	70.88	5.26
	+ FMS	80.20	71.75	73.80	75.26	71.33	67.39	73.68	70.80	69.17	68.89	65.17	67.74	71.26	4.18
	+ DIR + FMS	78.18	74.55	72.68	75.14	72.69	66.79	71.07	70.18	69.37	70.75	66.10	68.74	71.35	3.61
	LC CP-ResNet	72.30	57.00	59.98	63.10	56.73	51.92	58.42	55.69	53.27	49.90	46.34	49.84	56.21	7.02
	+ Mixup	71.64	57.14	58.18	62.33	54.81	55.43	60.99	57.08	55.33	52.71	43.01	50.35	56.58	7.10
	+ DIR	70.99	58.76	66.95	65.57	59.29	52.40	62.81	58.17	54.85	53.88	45.05	51.26	58.33	7.45
	+ FMS	73.43	56.86	60.69	63.67	59.49	52.99	60.40	57.63	53.92	54.97	45.66	51.52	57.60	7.12
	+ DIR + FMS	67.11	58.38	63.81	63.10	57.86	51.43	60.12	56.47	55.45	58.83	44.83	53.04	57.53	6.18
	PaSST	80.12	78.32	78.56	79.00	76.04	71.58	73.37	73.66	73.72	74.59	74.73	74.34	75.67	2.83
	+ Mixup	82.30	77.95	81.36	80.54	76.57	71.68	74.34	74.20	75.43	74.83	75.03	75.10	76.61	3.33
	+ DIR	80.65	77.99	79.51	79.39	76.08	71.82	74.57	74.15	74.85	75.45	74.75	75.02	76.18	2.71
	+ FMS	79.90	77.22	77.08	78.07	74.95	71.11	72.85	72.97	74.67	73.96	74.14	74.26	75.10	2.56
	+ DIR + FMS	78.53	76.62	78.03	77.73	75.19	71.35	74.91	73.82	74.91	73.98	74.08	74.32	75.29	2.42

TABLE I

PERFORMANCE ON THE *TAU Urban Acoustic Scenes 2020 / 2022 Mobile development dataset (TAU20 / TAU22)* [3] FOLLOWING THE OFFICIAL TEST SPLIT: DEVICE-WISE COMPARISON BETWEEN BASELINE, MIXUP, FREQ-MIXSTYLE (FMS), DEVICE IMPULSE RESPONSE AUGMENTATION (DIR) AND A COMBINATION OF DIR AND FMS (DIR + FMS). THE PROVIDED ACCURACIES (%) AND DEVICE GENERALIZATION SCORES (DG) ARE AVERAGED OVER THREE RUNS AND THE LAST FIVE EPOCHS OF TRAINING. THE DEVICES ARE GROUPED ACCORDING TO REAL DEVICES (**REAL**: A, B, C), THE SEEN, SIMULATED DEVICES (**SIM**: S1, S2, S3) AND THE UNSEEN, SIMULATED DEVICES (**UNSEEN**: S4, S5, S6). FOR EACH SETUP, WE MARK THE BEST-PERFORMING METHOD IN GRAY.

3) *PaSST*: AudioSet [24] is collected from YouTube videos and consists of audio clips recorded by a large variety of different recording devices of diverse quality. Since *PaSST* is pre-trained on AudioSet, the baseline is already robust against different devices, performing much better than the baselines of the other models that are trained from scratch. In comparison to CP-ResNet and LC CP-ResNet, the performance gap between the baseline and the device generalization methods is much lower for both the overall and unseen accuracies. In the case of *TAU20*, pre-training and fine-tuning clip length (10 seconds) match exactly, allowing *PaSST* to fully exploit its robustness gained from the pre-training stage. In this scenario, using DIR or FMS can not improve the overall or unseen accuracy. On *TAU22*, consisting of one-second samples, DIR, FMS and the combination of them all enhance the performance with DIR achieving the best overall accuracy.

C. Hyperparameter Sensitivity

Relevant hyperparameters for the device generalization methods include the parameters α to determine the shape of the Beta distributions for Mixup and FMS, and p_{dir} and p_{fms}




Fig. 2. Parallel coordinate plot visualizing the relationship between the Freq-MixStyle probability (p_{fms}), the DIR augmentation probability (p_{dir}), accuracy on unseen devices and overall accuracy. Each line is an average over three experiments of running *PaSST* on *TAU22* using DIR + FMS as the device generalization method.

to specify the probabilities to augment a sample and a batch for DIR and FMS, respectively. As suggested in [8], α is set to 0.3 and 0.4 for Mixup and Freq-MixStyle, respectively, and p_{fms}

takes a value of 0.4. We found that values in the range from 0.4 to 0.6 for p_{dir} yield robust results across models and datasets and decided to use $p_{dir} = 0.6$. We choose these values for every setup except for training CP-ResNet on TAU22 where we experimentally found that values of 0.8 and 0.4 for p_{fms} and p_{dir} give superior results. The hyperparameter configuration is more sensitive when using the combination of DIR and FMS. Figure 2 visualizes the performance impact of p_{fms} and p_{dir} on the unseen and overall accuracy for PaSST on TAU22. The plot reveals that suitable values for p_{dir} are in the range 0.2 to 0.6 and for p_{fms} between 0.2 and 0.4. Choosing either of p_{dir} or p_{fms} too high or setting them to zero decreases performance. This observation holds across all models and datasets.

VI. CONCLUSION

In this paper, we address the issue of the distributional shift in digitized acoustic signals caused by different recording device characteristics, which can lead to severe performance drops for audio classification models at application time. We propose device impulse response augmentation as a means to artificially increase the recording device diversity in the training set. Our proposed method drastically improves the device generalization across multiple models and datasets for the task of Acoustic Scene Classification. We compare to other device generalization techniques and conclude that device impulse response augmentation performs similarly to the state-of-the-art method Freq-MixStyle when tested in isolation. However, we also show that device impulse response augmentation and Freq-MixStyle are complementary. The combination is especially powerful to generalize to devices unseen during training, giving the highest accuracy on unseen devices on five out of the six setups tested.

REFERENCES

- [1] V. Vapnik, *The Nature of Statistical Learning Theory*, ser. Statistics for Engineering and Information Science. Springer, 2000.
- [2] G. Wilson and D. J. Cook, "A survey of unsupervised deep domain adaptation," *ACM Trans. Intell. Syst. Technol.*, 2020.
- [3] T. Heittola, A. Mesaros, and T. Virtanen, "Acoustic scene classification in DCASE 2020 challenge: Generalization across devices and low complexity solutions," in *Proceedings of the 5th Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2020.
- [4] I. Martín-Morató, F. Paissan, A. Ancilotto, T. Heittola, A. Mesaros, E. Farella, A. Brutti, and T. Virtanen, "Low-complexity acoustic scene classification in DCASE 2022 challenge," in *Proceedings of the 7th Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2022.
- [5] K. Koutini, F. Henkel, H. Eghbal-zadeh, and G. Widmer, "CP-JKU Submissions to DCASE'20: Low-Complexity Cross-Device Acoustic Scene Classification with RF-Regularized CNNs," *Tech. Rep., DCASE Challenge*, 2020.
- [6] B. Kim, S. Yang, J. Kim, H. Park, J. Lee, and S. Chang, "Domain generalization with relaxed instance frequency-wise normalization for multi-device acoustic scene classification," in *23rd Annual Conference of the International Speech Communication Association (Interspeech)*. ISCA, 2022.
- [7] J.-H. Lee, J.-H. Choi, P. M. Byun, and J.-H. Chang, "Multi-scale architecture and device-aware data-random-drop based fine-tuning method for acoustic scene classification," in *Proceedings of the 7th Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2022.

- [8] F. Schmid, S. Masoudian, K. Koutini, and G. Widmer, "Knowledge distillation from transformers for low-complexity acoustic scene classification," in *Proceedings of the 7th Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2022.
- [9] K. Koutini, H. Eghbal-zadeh, M. Dorfer, and G. Widmer, "The receptive field as a regularizer in deep convolutional neural networks for acoustic scene classification," in *27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019.
- [10] K. Koutini, H. Eghbal-zadeh, and G. Widmer, "Receptive field regularization techniques for audio classification and tagging with deep convolutional neural networks," *IEEE ACM Trans. Audio Speech Lang. Process.*, 2021.
- [11] K. Koutini, J. Schlüter, H. Eghbal-zadeh, and G. Widmer, "Efficient training of audio transformers with patchout," in *23rd Annual Conference of the International Speech Communication Association (Interspeech)*. ISCA, 2022.
- [12] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. J. Smola, "A kernel two-sample test," *J. Mach. Learn. Res.*, 2012.
- [13] S. Kolouri, K. Nadjahi, U. Simsekli, R. Badeau, and G. K. Rohde, "Generalized sliced wasserstein distances," in *Advances in Neural Information Processing Systems 32 (NeurIPS)*, 2019.
- [14] P. Primus, H. Eghbal-zadeh, D. Eitelsebner, K. Koutini, A. Arzt, and G. Widmer, "Exploiting parallel audio recordings to enforce device invariance in cnn-based acoustic scene classification," in *Proceedings of the 4th Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2019.
- [15] B. Kim, S. Yang, J. Kim, and S. Chang, "QTI submission to DCASE 2021: residual normalization for device-imbalanced acoustic scene classification with efficient design," *Tech. Rep., DCASE Challenge*, 2021.
- [16] I. Martín-Morató, T. Heittola, A. Mesaros, and T. Virtanen, "Low-complexity acoustic scene classification for multi-device audio: Analysis of DCASE 2021 challenge systems," in *Proceedings of the 6th Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2021.
- [17] I. Szöke, M. Skácel, L. Mosner, J. Paliesek, and J. H. Cernocký, "Building and evaluation of a real room impulse response dataset," *IEEE J. Sel. Top. Signal Process.*, 2019.
- [18] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017.
- [19] A. Sriram, H. Jun, Y. Gaur, and S. Satheesh, "Robust speech recognition using generative adversarial networks," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018.
- [20] M. Ritter, M. Müller, S. Stüker, F. Metzke, and A. Waibel, "Training deep neural networks for reverberation robust speech recognition," in *Proceedings of the 12th ITG Symposium on Speech Communication*. IEEE, 2016.
- [21] M. Ferras, S. R. Madikeri, P. Motlíček, S. Dey, and H. Bourlard, "A large-scale open-source acoustic simulator for speaker recognition," *IEEE Signal Process. Lett.*, 2016.
- [22] V.-V. Eklund, "Data augmentation techniques for robust audio analysis," *Master's Thesis, Tampere University*, 2019.
- [23] K. Koutini, S. Jan, and G. Widmer, "CPJKU Submission to DCASE21: Cross-Device Audio Scene Classification with Wide Sparse Frequency-Damped CNNs," *Tech. Rep., DCASE Challenge*, 2021.
- [24] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017.
- [25] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *7th International Conference on Learning Representations (ICLR)*. OpenReview.net, 2019.
- [26] H. Zhang, M. Cissé, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *6th International Conference on Learning Representations (ICLR)*. OpenReview.net, 2018.
- [27] K. Koutini, H. Eghbal-zadeh, and G. Widmer, "CP-JKU submissions to DCASE'19: Acoustic scene classification and audio tagging with receptive-field-regularized CNNs," *Tech. Rep., DCASE Challenge*, 2019.
- [28] J. Fu, Y. Zhong, and F. Yang, "Adversarial domain generalization with mixstyle," in *International Conference on Advanced Robotics and Mechatronics (ICARM)*. IEEE, 2022.