

# UC San Diego

## UC San Diego Previously Published Works

### Title

Microphone Conversion: Mitigating Device Variability in Sound Event Classification

### Permalink

<https://escholarship.org/uc/item/16455736>

### Authors

Ryu, Myeonghoon

Oh, Hongseok

Lee, Suji

et al.

### Publication Date

2024-04-14

### DOI

10.1109/icassp48485.2024.10447432

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# MICROPHONE CONVERSION: MITIGATING DEVICE VARIABILITY IN SOUND EVENT CLASSIFICATION

\*Myeonghoon Ryu<sup>1,2</sup> \*Hongseok Oh<sup>1,3</sup> Suji Lee<sup>1</sup> Han Park<sup>1</sup>

<sup>1</sup> Deeply Inc.

<sup>2</sup> Seoul National University

<sup>3</sup> University of California, San Diego

## ABSTRACT

In this study, we introduce a new augmentation technique to enhance the resilience of sound event classification (SEC) systems against device variability through the use of CycleGAN. We also present a unique dataset to evaluate this method. As SEC systems become increasingly common, it is crucial that they work well with audio from diverse recording devices. Our method addresses limited device diversity in training data by enabling unpaired training to transform input spectrograms as if they are recorded on a different device. Our experiments show that our approach outperforms existing methods in generalization by 5.2% - 11.5% in weighted f1 score. Additionally, it surpasses the current methods in adaptability across diverse recording devices by achieving a 6.5% - 12.8% improvement in weighted f1 score.

*Index Terms*— Sound event classification, device mismatch, generative adversarial network, deep learning

## 1. INTRODUCTION

Sound event classification (SEC) aims to automatically identify various types of sounds, like speech, music, and environmental sound, using signal processing and machine learning techniques. Despite recent progress that has led to practical applications, existing models still struggle with distortions caused by different recording devices. These distortions, although often subtle to human ears as shown in Figure 1, can significantly hamper SEC system performance[1].

Past strategies to tackle this domain shift have largely centered on data augmentation and normalization. Initiatives like the DCASE Challenge[1, 2] have tackled this problem but have limitations due to their focus on synthetic evaluation data, reducing the thoroughness of its evaluations.

Previous studies have applied CycleGAN to tasks like speaker verification and emotion recognition[3, 4]. However, they depended on datasets lacking detailed device information and did not use optimal recording environments. This oversight can introduce biases and challenges.

To bridge these gaps, we created a dataset featuring sound events recorded in an anechoic chamber using real-world devices, simultaneously. This dataset facilitates a deeper understanding of device-induced performance variations.

Our key contribution is the introduction of the unique dataset and Microphone Conversion, a plug-and-play augmentation technique for SEC systems. With the CycleGAN[5] framework, our method enables more effective deployment on heterogeneous devices without requiring paired or labeled data. This plug-and-play solution generates spectrograms resembling recordings from desired devices. It empowers SEC systems to either generalize across devices or optimize for a specific one, showing enhanced performance compared to existing state-of-the-art methods.

## 2. DATASET

In this section, we outline the process of creating the unique dataset and provide an overview of its contents.

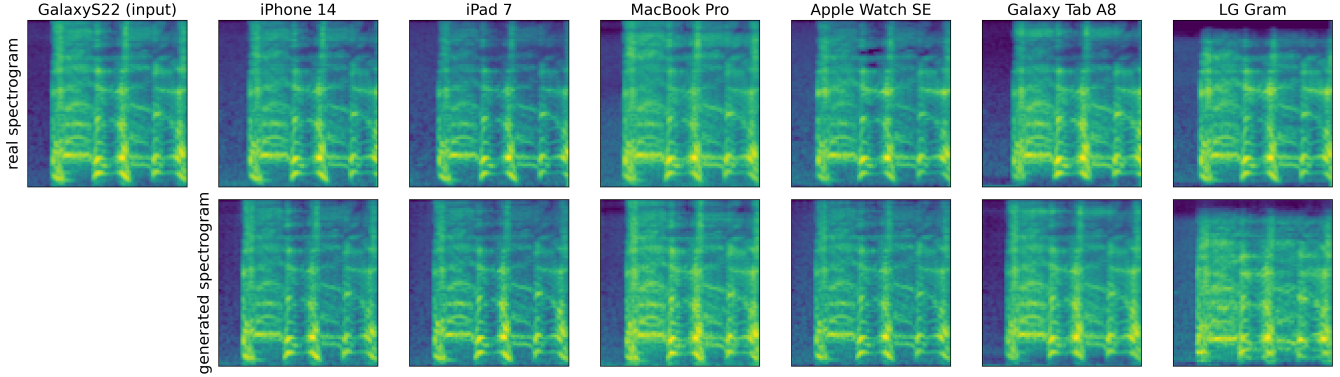
### 2.1. Sound Events

The dataset contains 75 types of sound. Of these, 25 were directly produced and recorded, while the remaining 50 are recordings of the playback of anechoic recordings from the RWCP Sound Scene Database (RWCP-SSD)[6]. The former includes sound of things, human sound, and musical instruments. The RWCP-SSD contains non-speech sounds recorded in an anechoic room. These RWCP-SSD samples were played using a KRK VXT6 loudspeaker. We manually annotated the start and end times of each sound event, resulting in a paired dataset intended for fair evaluation.

### 2.2. Devices

We have constructed our dataset using 18 distinct recording devices. Of these, 11 end-user devices come with built-in microphones, chosen specifically to capture distortions caused by varying microphones and digital signal processing units. Additionally, we have included 7 microphones to investigate the distortions due to transducer types and polar patterns.

\*These authors contributed equally to this work.



**Fig. 1:** Spectrograms of the real and generated coughing sounds are presented in the top and bottom rows, respectively. The generated ones are produced by corresponding Microphone Conversion networks using a real spectrogram of Galaxy S22.

### 2.3. Environment

All recordings were made in an anechoic chamber meeting ISO 3745 standards<sup>1</sup>. Sound sources were located on a side of the chamber, while the recording devices were set 1.5 meters away, with their microphones aimed at the sources.

## 3. OUR METHOD

The objective of Microphone Conversion is to train a mapping function  $F$  that transforms source device spectrograms  $X_A$  into those of a target device. In essence, the generator produces output spectrograms  $F(X_A)$  that are statistically similar to the data distribution of the target device B.

### 3.1. CycleGAN


Our approach employs CycleGAN, an unsupervised framework for image-to-image translation that can learn mappings between two domains with unpaired data. Built on the Generative Adversarial Networks (GANs), CycleGAN involves two generators,  $F$  and  $G$ , and two discriminators,  $D_A$  and  $D_B$ .

$$\mathcal{L}_{adv}(F, D_B, X_A, X_B) = \mathbb{E}_{X_A \sim p(x_a)}[\log(1 - D_B(F(X_A)))] + \mathbb{E}_{X_B \sim p(x_b)}[\log D_B(X_B)] \quad (1)$$

$$\mathcal{L}_{cycle}(F, G, X_A, X_B) = \mathbb{E}_{X_A \sim p(x_a)}[\|G(F(X_A)) - X_A\|_1] + \mathbb{E}_{X_B \sim p(x_b)}[\|F(G(X_B)) - X_B\|_1] \quad (2)$$

Each generator performs bijective mapping between domain A and B, serving as inverse functions of one another. The discriminators aim to differentiate real images from converted counterparts. The adversarial loss (Eq. 1) encourages the generators to produce outputs indistinguishable from the target domain, while the cycle-consistency loss (Eq. 2)

<sup>1</sup><https://www.iso.org/standard/45362.html>



**Fig. 2:** A difference spectra between a real spectrum ( $x_{source}$ , iPhone 14) and a real/generated spectrum ( $x_{target}$ ,  $\hat{x}_{target}$ , Other devices). Spectra are calculated using Welch’s method.

ensures that the generated images, when reverse-mapped, closely resemble their originals. This dual-loss setup allows CycleGAN to preserve the content while translating the style.

### 3.2. Development Data

To build a development set for our experiment, we select data from seven end-user devices from  $data_{full}$ . We create 315,966 audio segments using a 930-ms window and a 50% overlap. They are filtered based on a 10% threshold for sparse classes like ‘coughing’ and ‘clap,’ and a 50% threshold for the others, resulting in a 246,544-segment development set.

The entire development set is then divided into three subsets:  $data_{train,mc}$ ,  $data_{train,sec}$ , and  $data_{val}$  are 45%, 45%, 10%, respectively. We use stratification to balance sound events across subsets and to guarantee that each audio segment has counterparts from different devices in each subset.

## 4. EXPERIMENTS AND RESULTS

In this section, we explore the performance decline in SEC systems when faced with heterogeneous recording devices. We demonstrate that incorporating Microphone Conversion can significantly improve the SEC systems’ resilience to device variability. Additionally, we show our approach is effective in adapting SEC systems for specific devices.

### 4.1. Implementation Details

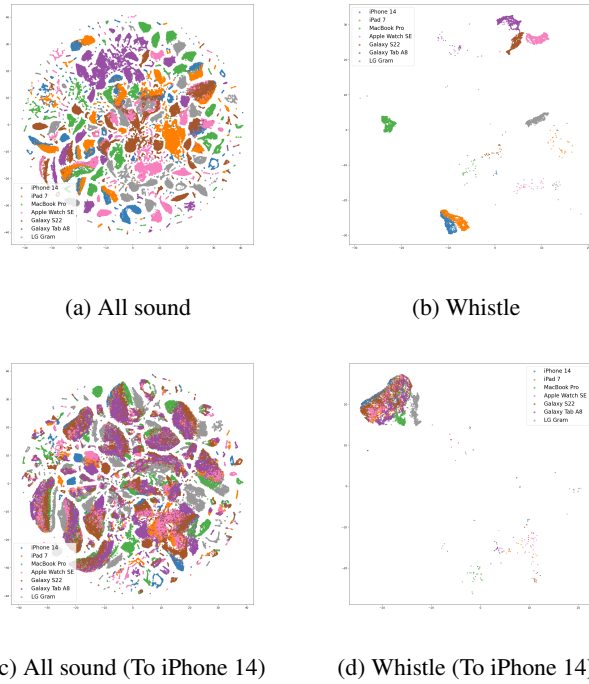
The audio data, down-sampled to 22,050 Hz, is converted into log Mel spectrograms using a 1,024-sample Hanning window, a 256-sample hop length, and 80 Mel bands. We replace the negative log likelihood loss in Eq. (1) with the least square loss[9] for Microphone Conversion models. we randomly draw an initial learning rate from [0.00002, 0.002] and halved at random intervals between [10, 50], optimized by the Adam optimizer with  $\beta = (0.5, 0.999)$ , with a batch size of 100. Discriminator updates employ a generated image buffer[10]. These models are trained for 100 and 200 epochs, as MC-100 and MC-200, using  $data_{train,mc}$ , respectively. We set  $\lambda$  to 10 for  $L_{cycle}$  in the total loss function to control its relative importance. Parameter tuning is done using 10 iterations of Bayesian search, minimizing  $L_{cycle}$  for final evaluation.

For SEC systems, we employ ResNet50[11] as our baseline network and use iPhone 14 segments of  $data_{train,sec}$  for training. We specifically choose the iPhone 14 as our source device because smartphones are among the most commonly used devices for sound reception in the modern world, making them ideal for our study. Model optimization is carried out using AdamW[12] with  $\beta = (0.9, 0.999)$  and an initial learning rate of 0.001, reduced by 0.1 every 25 epochs. All implementations are executed on NVIDIA T4 and RTX6000 GPUs. During training, the source data, iPhone 14, is randomly converted to six other devices through Microphone Conversion networks to assess generalization capabilities. MC-100-Gen and MC-200-Gen correspond to this strategy in Table 1.

We also evaluate our methods in two adaptation scenarios, one with a single source (iPhone 14) and target, and another with multiple sources and a single target (iPhone 14), with probabilities set at 0.5 and 1.0 for applying Microphone Conversion. MC-100-Adapt and MC-200-Adapt in Table 1 and 2, respectively, correspond to the single-target and multi-target strategy. Finally, a ‘Real’ performance metric is introduced as an ideal case, where training and inference device matches.

### 4.2. Recent Approaches

We examine top-performing methods from recent DCASE Challenges that mitigate device variability. Freq-MixStyle[13] and Residual Normalization[14] utilize frequency-wise statistics for better generalization. An extended version, Relaxed Instance Frequency-wise Normalization (RFN)[15], is applied after each bottleneck block, resulting in 5 RFN blocks.



**Fig. 3:** 2D t-SNE[7] visualization of intermediate embeddings of the baseline ResNet50 on the development data. (a) and (b) illustrate the full development set and whistle sound samples drawn from it, respectively. (c) and (d) are sound samples converted to iPhone 14 using (a) and (b).

### 3.3. Network Architecture

We adapt the CycleGAN implementation from the original author’s code<sup>2</sup>. Our generator includes two up- and down-sampling layers and nine residual blocks with instance normalization, as in [5]. The tanh layer is omitted. For the discriminator, we employ a 16x16 PatchGAN[8] with instance normalization for better convergence and to minimize blurriness in the output.

### 3.4. Visualization of Generated Samples

Figure 1 demonstrates that the Microphone Conversion network accurately replicates the spectro-temporal characteristics of each target device. We visualize difference spectra between devices using Welch’s method, and Figure 2 confirms the network’s ability to capture each device’s unique frequency response. We average spectrogram samples along with temporal dimension to derive spectra of different devices. Feature-level analysis, shown in Figure 3, validates that our method produces outputs closely aligned with the target device, establishing its efficacy for data augmentation.

<sup>2</sup><https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix>

**Table 1:** Results for generalization/adaptation capability of previous methods and our methods on the validation set. Source device(S) is iPhone 14, and target devices(T1~T6) are Galaxy S22, iPad 7, Galaxy Tab A8, Apple Watch SE, Macbook Pro('20), and LG Gram('20), respectively. The last column shows an average and 95% confidence interval of the performance.

Method	F1 Score							
	S	T1	T2	T3	T4	T5	T6	Overall (- S)
Baseline	0.982	0.409	0.709	0.248	0.471	0.687	0.491	0.503 ± 0.167
Gaussian Noise	0.983	0.708	0.918	0.576	0.565	0.780	0.683	0.705 ± 0.127
Reverberation	0.980	0.895	0.852	0.539	0.832	0.736	0.360	0.702 ± 0.202
Pitch Shift	0.981	0.471	0.744	0.221	0.658	0.648	0.442	0.531 ± 0.183
SpecAugment	<b>0.985</b>	0.372	0.762	0.214	0.363	0.634	0.324	0.445 ± 0.199
MixUp	0.983	0.336	0.677	0.213	0.449	0.656	0.387	0.453 ± 0.175
FilterAugment	0.981	0.964	0.891	0.586	0.874	0.794	0.642	0.792 ± 0.143
Freq-MixStyle	0.974	0.839	0.879	0.795	0.902	<b>0.885</b>	0.832	0.855 ± 0.038
RFN	0.980	0.919	0.909	0.742	0.907	0.829	0.614	0.820 ± 0.116
MC-100-Gen	0.981	0.958	<b>0.912</b>	0.894	0.899	0.831	0.852	0.891 ± 0.043
MC-200-Gen	0.982	<b>0.969</b>	0.909	<b>0.903</b>	<b>0.912</b>	0.859	<b>0.887</b>	<b>0.907 ± 0.035</b>
MC-100-Adapt (p=0.5)	-	0.956	0.905	0.904	0.880	0.902	0.890	0.906 ± 0.025
MC-100-Adapt (p=1.0)	-	0.891	0.906	0.855	0.834	0.811	0.808	0.851 ± 0.039
MC-200-Adapt (p=0.5)	-	0.965	<b>0.922</b>	<b>0.906</b>	0.910	<b>0.908</b>	<b>0.907</b>	<b>0.920 ± 0.022</b>
MC-200-Adapt (p=1.0)	-	0.935	0.917	0.864	0.873	0.880	0.848	0.886 ± 0.032
Real	0.983	0.982	0.972	0.985	0.979	0.983	0.986	0.981 ± 0.005

**Table 2:** Results for the second adaptation scenario. Training Set 1 and 2 are (GalaxyS22, LG Gram, GalaxyTabA8) and (iPad7, AppleWatchSE, MacbookPro), respectively.

Method	F1 Score	
	Set 1	Set 2
Baseline	0.917	0.975
MC-100-Adapt (p=0.5)	0.976	<b>0.978</b>
MC-100-Adapt (p=1.0)	0.940	0.971
MC-200-Adapt (p=0.5)	0.977	<b>0.978</b>
MC-200-Adapt (p=1.0)	<b>0.979</b>	<b>0.978</b>
Real	0.983	0.983

FilterAugment[16] adjusts spectrograms via randomly generated filters, and we apply its optimal linear type[17]. We also test standard data augmentations like Gaussian noise, room impulse response, pitch shift, SpecAugment[18], and MixUp[19], all with a 0.5 probability, except for RFN.

### 4.3. Results

Our evaluation, summarized in Tables 1 and 2, reveals SEC system performance vulnerabilities when faced with heterogeneous recording devices. For instance, there is up to a 73.7% performance drop compared to the ‘Real’ metric in the case of T3. However, models perform better on devices like

T2 and T5 that share similar features with the training data, as indicated by their closeness in the t-SNE feature space, showing a modest performance drop of 26.3% and 29.5%.

Traditional data augmentations provide limited improvement. Among recent approaches, Freq-MixStyle is most effective, achieving an average F1 score of 85.5% with a 3.8% confidence interval. MC-200-Gen (Ours), excels with an average F1 score of 90.7% and minimal variability (3.5%).

In adaptation scenarios, SEC models using an adaptation strategy outperform MC-200-Adapt (p=0.5) on 4 out of 6 targets, gaining 1.3% in overall performance. They also nearly match the ‘Real’ performance metric when trained on diverse devices. Longer training for our Microphone Conversion networks boosts both generalization and adaptability, hinting at the potential for further optimization.

## 5. CONCLUSIONS AND LIMITATIONS

To tackle device variability in SEC systems, we developed a specialized sound event dataset recorded across multiple real-world devices in an anechoic chamber. Our introduced Microphone Conversion method significantly improves SEC performance, outperforming recent approaches in both generalization and adaptation tasks. However, the CycleGAN component of our solution assumes a one-to-one domain mapping, requiring separate models for each domain pair. Future work could explore integrating impulse response with CycleGAN for more versatile domain mapping.

## 6. REFERENCES

- [1] Irene Martín-Morató, Francesco Paissan, Alberto Ancilotto, Toni Heittola, Annamaria Mesaros, Elisabetta Farella, Alessio Brutti, and Tuomas Virtanen, “Low-complexity acoustic scene classification in dcase 2022 challenge,” 2022.
- [2] Nicolas Turpault, Romain Serizel, Ankit Parag Shah, and Justin Salamon, “Sound event detection in domestic environments with weakly labeled data and soundscape synthesis,” in *Workshop on Detection and Classification of Acoustic Scenes and Events*, New York City, United States, October 2019.
- [3] Akhil Mathur, Anton Isopoussu, Fahim Kawsar, Nadia Berthouze, and Nicholas D Lane, “Mic2mic: using cycle-consistent generative adversarial networks to overcome microphone variability in speech systems,” in *Proceedings of the 18th international conference on information processing in sensor networks*, 2019, pp. 169–180.
- [4] Saurabh Kataria, Jesús Villalba, Piotr Żelasko, Laureano Moro-Velázquez, and Najim Dehak, “Deep feature cyclegrams: Speaker identity preserving non-parallel microphone-telephone domain adaptation for speaker verification,” *arXiv preprint arXiv:2104.01433*, 2021.
- [5] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [6] Satoshi Nakamura, Kazuo Hiyane, Futoshi Asano, Takanobu Nishiura, and Takeshi Yamada, “Acoustical sound database in real environments for sound scene understanding and hands-free speech recognition,” in *LREC*, 2000.
- [7] Laurens van der Maaten and Geoffrey Hinton, “Visualizing data using t-sne,” *Journal of Machine Learning Research*, vol. 9, no. 86, pp. 2579–2605, 2008.
- [8] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros, “Image-to-image translation with conditional adversarial networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1125–1134.
- [9] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley, “Least squares generative adversarial networks,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2794–2802.
- [10] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb, “Learning from simulated and unsupervised images through adversarial training,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2107–2116.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep Residual Learning for Image Recognition,” in *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition*. June 2016, CVPR ’16, pp. 770–778, IEEE.
- [12] Ilya Loshchilov and Frank Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations*, 2019.
- [13] Florian Schmid, Shahed Masoudian, Khaled Koutini, and Gerhard Widmer, “CP-JKU submission to dcase22: Distilling knowledge for low-complexity convolutional neural networks from a patchout audio transformer,” Tech. Rep., DCASE2022 Challenge, June 2022.
- [14] Byeonggeun Kim, Seunghan Yang, Jangho Kim, and Simyung Chang, “QTI submission to DCASE 2021: Residual normalization for device-imbalanced acoustic scene classification with efficient design,” Tech. Rep., DCASE2021 Challenge, June 2021.
- [15] Byeonggeun Kim, Seunghan Yang, Jangho Kim, Hyunsin Park, Juntae Lee, and Simyung Chang, “Domain Generalization with Relaxed Instance Frequency-wise Normalization for Multi-device Acoustic Scene Classification,” in *Proc. Interspeech 2022*, 2022, pp. 2393–2397.
- [16] Hyeonuk Nam, Byeong-Yun Ko, Gyeong-Tae Lee, Seong-Hu Kim, Won-Ho Jung, Sang-Min Choi, and Yong-Hwa Park, “Heavily augmented sound event detection utilizing weak predictions,” Tech. Rep., DCASE2021 Challenge, June 2021.
- [17] Hyeonuk Nam, Seong-Hu Kim, and Yong-Hwa Park, “Filteraugment: An acoustic environmental data augmentation method,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 4308–4312.
- [18] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le, “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech 2019*, 2019, pp. 2613–2617.
- [19] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz, “mixup: Beyond empirical risk minimization,” in *International Conference on Learning Representations*, 2018.