



Domain Adaptation Method and Modality Gap Impact in Audio-Text Models for Prototypical Sound Classification

Emiliano Acevedo¹, Martín Rocamora^{1,2}, Magdalena Fuentes³

¹Facultad de Ingeniería, Universidad de la República, Uruguay

²Music Technology Group, Universitat Pompeu Fabra, Spain

³MARL-IDM, New York University, USA

eacevedo@fing.edu.uy, martin.rocamora@upf.edu, mfuentes@nyu.edu

Abstract

Audio-text models are widely used in zero-shot environmental sound classification as they alleviate the need for annotated data. However, we show that their performance severely drops in the presence of background sound sources. Our analysis reveals that this degradation is primarily driven by SNR levels of background soundscapes, and independent of background type. To address this, we propose a novel method that quantifies and integrates the contribution of background sources into the classification process, improving performance without requiring model retraining. Our domain adaptation technique enhances accuracy across various backgrounds and SNR conditions. Moreover, we analyze the modality gap between audio and text embeddings, showing that narrowing this gap improves classification performance. The method generalizes effectively across state-of-the-art prototypical approaches, showcasing its scalability and robustness for diverse environments.

Index Terms: audio-text models, modality gap, domain adaptation, zero-shot sound classification

1. Introduction

Audio-text models (ATMs), i.e. embedding models with semantic knowledge of language and sound, have become popular in zero-shot (ZS) sound classification as they relieve the need for collecting and annotating audio recordings [1, 2]. Aligning audio and text modalities facilitates prototype approaches, where *anchors*, specific embedding points of one modality, act as class prototypes for the other within a shared latent space. A common method uses text *anchors* to define audio class prototypes and then classify unlabeled audio based on its proximity to the nearest prototype. Such models have potential in applications such as noise pollution analysis [3], surveillance systems [4], hearing aid technology [5], and smart city initiatives [6]. Among audio-text models, LION-CLAP [7] stands out in multiple tasks [8] as one of the best-performing due to training on one of the largest available audio-text datasets, and careful selection of hyperparameters for both audio and text encoders.

Despite the real-world nature of sound classification (i.e. sounds are rarely encountered in isolation but within a background soundscape), most ATMs have been used for ZS classification with isolated sound sources. In this work, we show that the presence of background sound sources significantly degrades ATMs performance. Recently, domain adaptation (DA) techniques for ATMs have been proposed [9, 10], focusing only on optimizing the text encoder’s input or feature space by learning domain-specific tokens. Furthermore, as noted in [11], ATMs struggle to fully understand natural language, particularly contextual concepts such as sequential or concurrent sound events. Our analysis further reveals that treating the background

as a high-level text concept (e.g., “the sound of a park”) is misleading and often results in misclassifications driven by background sound sources. These findings underscore the need for more precise characterizations of background soundscapes to achieve effective DA in these scenarios. Motivated by these limitations, we propose a method that quantifies the contribution of background sound sources in an audio recording to improve ZS classification. By treating the background as a combination of individual sounds, our method consistently improves performance across various types and amplitude levels of background soundscapes without requiring model retraining.

Besides the challenges of real-world sound classification, audio-text models suffer from a common issue of multimodal models, the modality gap (i.e. a semantic misalignment of modalities in the embedding space). This is a known consequence of models trained with contrastive loss [12, 13]. Previous solutions, such as projection-based and nearest-neighbor decoding [14], have addressed this issue. Building on these, we show that the recent method proposed by [15], which leverages unlabeled audio, is similarly effective for narrowing the gap, resulting in improved performance.

This work makes three key contributions: (i) a quantification of background sound sources’ influence in zero-shot classification, (ii) a novel method that incorporates background profiles to improve performance in domain shifts, and (iii) an analysis of the modality gap’s impact on sound classification. Our code is made available as open source for further research.¹

2. Impact of Background Sounds

To assess the impact of background sound on ZS sound classification using ATMs, we synthetically generated soundscapes using the *Scraper* library [16]. The foreground samples were drawn from UrbanSound8k [17], while the background soundscape samples were taken from TAU Urban Acoustic Scenes 2019 [18]. We generated new soundscapes by pairing each UrbanSound8k foreground sound with one random background sample from three indoor (Shopping Mall, Metro Station, Airport) and three outdoor (Park, Public Square, Street Traffic) environments, varying the signal-to-noise ratio (SNR) across three levels, namely 6, 8, and 10 dB. The selected values ensure that foreground sounds remain salient sounds, while preserving the background’s texture-like qualities. Values below 6 dB make the foreground indistinguishable from the background, whereas values above 10 dB minimize the impact of the background, consistent with the insights from previous work [19].

Figure 1 shows the classification performance of different sources from UrbanSound8k in the presence of background, compared to a reference performance of the sounds in isolation.

¹github.com/eacevedo1/AudioText-ContextDA.git

In every case, we observe a notable decline in performance once the background soundscape is introduced. Performance degradation is tied to SNR levels—lower SNR (higher background presence) results in poorer performance. This decrease is consistent across both indoor and outdoor environments, with SNR being the primary factor affecting the model.




Figure 1: Impact of background soundscape on zero-shot sound classification accuracy across different acoustic environments. The plot compares model performance under 6 dB, 8 dB, and 10 dB SNR conditions against the reference performance (i.e. isolated sources).

In addition to evaluating downstream task performance, we measured the cosine distance between text *anchors* and generated soundscapes, which averaged 0.76—slightly higher than the 0.73 observed with isolated foreground sounds. To investigate whether explicitly providing background information could improve performance, we modified the prompt to include the background type (e.g., ‘This is a sound of LABEL in the BACKGROUND TYPE’). This adjustment improved the alignment between text *anchors* and audio embeddings, reducing the cosine distance to 0.67. However, this improvement did not result in better classification performance across all background types. Moreover, we conducted ZS classification on the background soundscapes alone, achieving an accuracy of only 33.0%, which is significantly lower than 63.3% reported in the acoustic scene classification baseline [18]. This suggests that the model struggles to differentiate between various background types. These results imply that the concept of “background” is not well represented as a high-level concept within the model.

However, we observed that some background information was encoded within the model. For instance, when queried with *Airport* soundscapes, the model returned classes such as *bell*, *pedestrians* and *phone ringing*. Similarly, the prompt ‘Public square sound in the background’ yielded results like *people talking*, *pedestrians*, and *footsteps*. These findings suggest that the model encodes some relationships between background soundscapes and individual classes. In addition, we observed that the model’s classification errors were systematically linked to the background sound sources. For example, when adding *Street Traffic* soundscapes, the model top confusions are with *sirens*, *engine idling*, and *car horns*, or when adding *Shopping Mall* the model main confusions are *street music*, *air conditioner* and *children playing*. Adding a background soundscape transforms the au-

dio into polyphonic, which poses significant challenges for prototypical approaches, as they struggle to represent overlapping sources as a single point in the latent space. This limitation arises from the uniform aggregation of both temporal and timbral information, which fails to capture the inherent complexity of polyphonic acoustic scenes, as previously noted in [20]. Each individual background sound source influences the position of the query audio embedding with respect to different text *anchors* (i.e. the classes), diminishing confidence in a single foreground sound. This leads to unintended confusion between classes in the task and those present in the background soundscape that are acoustically similar or overlapping.

To study the effect of polyphony, we generated a dataset of 1000 samples without background, varying the number of classes per audio (C/A) by combining UrbanSound8k foreground sounds. We performed multi-label sound classification on this dataset, and results are shown in Table 1. We observe that with a fixed classification threshold, the average number of predicted classes per audio (Pred. C/A) decreases as the complexity of the audio scene increases, resulting in lower accuracy. We include both ZS and supervised approaches (i.e., using the class centroid based on annotations) to demonstrate that this effect is not specific to the audio-text model but is inherent to the prototype-based methods.

C/A	Zero-Shot		Supervised	
	Accuracy	Pred. C/A	Accuracy	Pred. C/A
1	59.5	0.9	71.4	0.8
2	50.8	0.8	43.7	0.6
3	48.5	0.6	39.3	0.4

Table 1: Performance comparison between zero-shot (text *anchors*) and Supervised (class centroid) modes on multi-label audio classification with varying numbers of classes per audio (C/A).

In summary, our experiments show that ATMs struggle to represent background as a coherent concept, making DA difficult. Adding background soundscapes turns audio into polyphonic, causing confusion and reducing classification confidence, which ultimately degrades model performance. Hence, to effectively address DA, the background should be treated as a combination of sound sources rather than a standalone concept, reducing its impact on the final classification.

3. Method Proposed

Sound classification task in this context implies accurately identifying a foreground sound embedded within a background soundscape (e.g., detecting a dog barking in a metro station soundscape). The audio to be classified consists of a mixture of the foreground sound and the background soundscape, which we refer to as test audio soundscape as shown in Figure 2. To address DA, we first define the profile concept as a vector of cosine similarity scores between a test audio embedding and the class prototypes. We begin the method by calculating the profile for the test audio soundscape.

Rather than just using this profile for classification, we refine it by subtracting the “contribution” of the background soundscape. To achieve this, we compute the background profile separately, using one of two approaches: (i) *Text* approach, where textual prompts describing the background are compared to class prototypes to get the profile, and (ii) *background Audio* approach, where an actual background

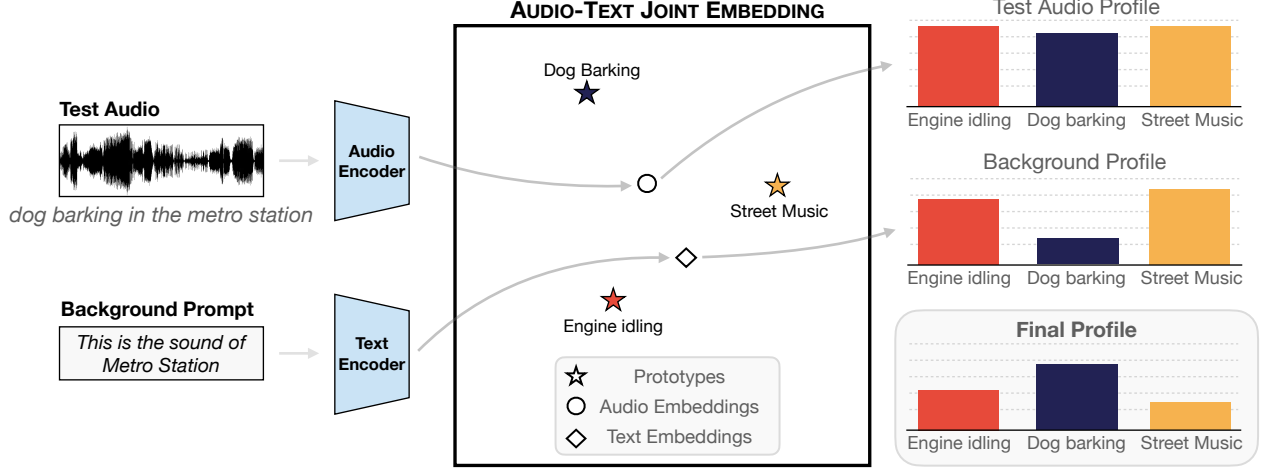


Figure 2: *Proposed method for domain adaptation in prototypical sound classification. The method quantifies the background’s contribution to the test audio soundscape and refines class estimation by subtracting its influence. First, the test audio profile is calculated by measuring the similarity between the test audio embedding and the class prototypes. Next, a background profile is generated, either from text descriptions or audio recordings, and subtracted from the test audio profile to minimize the impact of background classes.*

audio recording is used to get the profile. Both approaches can use single or average multiple prompts or recordings for a more robust background profile. For instance, the best results for the *Text* approach were obtained by averaging the profiles of three prompts: ‘This is a sound of BACKGROUND TYPE’, ‘BACKGROUND TYPE sounds in the background’, and ‘This is a sound of BACKGROUND TYPE in the background’.

After computing both the test audio and background profiles, we combine them to obtain a refined classification profile. This is done by subtracting the background profile from the soundscape profile, reducing the influence of background classes. The relationship is defined as:

$$P_f = P_s - P_b \times \tau \quad (1)$$

where P_s and P_b represent the test audio profile and the background profile, respectively, and the parameter τ controls the degree of background influence on the final classification.

4. On the Audio-Text Modality Gap

So far, we have mentioned how prototypical approaches in ATMs struggle to represent polyphonic acoustic scenes and to effectively leverage text. Apart from this, ATMs also face the challenge of the *modality gap*, wherein embeddings from different modalities are separated into distinct regions within the latent space. Despite its strong performance, LION-CLAP is not immune to the modality gap. On the UrbanSound8k dataset, we observed that the audio and text embeddings are linearly separable, with intra-modality distances consistently smaller than inter-modality distances. In particular, for any given class, the corresponding text *anchors* are typically farther from the audio samples of that class than from audio centroids of other classes, highlighting the persistent nature of this issue.

This modality misalignment suggests that relying solely on text *anchors* might limit the performance of zero-shot sound classification. Motivated by this, we explored recent work [15], which introduced a multimodal prototypical approach that we called *Text-Guided Audio Prototype* (TGAP). This method gen-

erates class prototypes by retrieving the N closest unlabeled audio embeddings to a text prompt and calculating their centroid, which becomes the prototype for classification. The method is fully unsupervised, without the need for re-training. The TGAP method consistently improves sound classification performance over simple ZS classification across various datasets. To understand the source of this improvement, we examined the cosine distances from (i) text-based prototypes, (ii) TGAP-generated centroids and (iii) real centroids to corresponding class audio samples, as illustrated in Figure 3.




Figure 3: *Comparison of cosine distances between audio samples and three types of class prototypes across multiple sound classes: text-based anchors (red), TGAP estimated centroids (blue), and real audio centroids (yellow) in UrbanSound8k.*

The experiment shows that text *anchors* prototypes (red) have the highest cosine distances, indicating poor alignment with audio samples. TGAP prototypes (blue) significantly reduce this distance, demonstrating improved alignment with the audio embeddings. Real audio centroids (yellow), derived using labels, represent the case of no modality gap. Then, TGAP significantly reduces the modality gap compared to text-based

Background	Zero-Shot			TGAP		
	Baseline	Text	Audio	Baseline	Text	Audio
Park	61.0	62.2	64.3	65.7	67.7	72.2
Airport	57.1	58.8	60.3	65.1	66.3	70.2
Shopping Mall	55.8	57.5	60.2	63.2	65.2	68.7
Street Traffic	61.1	63.5	66.8	68.8	70.6	72.6
Public Square	55.6	60.5	62.1	64.1	65.8	71.3
Metro Station	56.2	58.3	63.0	60.5	63.0	68.1
Zero-Shot Reference: 73.1			TGAP Reference: 76.9			

Table 2: Performance comparison of Zero-Shot and TGAP approach across various context backgrounds using SNR of 6 dB. Using no domain adaptation (Baseline), domain adaptation methods using text (Text) and audio (Audio) to get the background profile.

prototypes. These findings highlight the critical role of modality alignment in zero-shot sound classification and the effectiveness of TGAP in narrowing the modality gap.

5. Results and discussion

The results of our proposed method for ZS sound classification are shown in Table 2. We compare three settings: no domain adaptation (Baseline), text-based adaptation (Text), and audio-based adaptation (Audio). Both adaptation approaches consistently improve performance across all background environments, with audio-based adaptation achieving the best results. The proposed method is flexible and effective with both text *anchors* and TGAP prototypes. Applying the method to TGAP confirms its superior performance across different background types. Notably, the reduction in the modality gap provided by TGAP, also consistently boosts the performance compared to ZS. Audio-based adaptation also performs better with the TGAP method. We attribute this to the fact that actual background recordings offer a more precise and accurate representation than text descriptions. As mentioned earlier, text struggles to fully capture the background concept, making audio a more reliable source for adaptation.

All the results above were computed with a fixed SNR of 6dB. Figure 4 explores performance across varying SNRs (6, 8 and 10 dB) for both ZS and TGAP methods averaging over all background types. Our proposed adaptation consistently improves performance at all SNR levels, with audio-based adaptation consistently yielding the highest gains. Regarding the parameter τ , we optimized it using a grid search (0 to 1 in 0.1 steps), identifying 0.2 for text-based and 0.7 for audio-based adaptation as generally optimal. These values achieved a balance in background subtraction, as lower τ underestimated background effects, while higher τ overcompensated, reducing performance. The optimal value showed slight variations depending on the specific context and SNR's.

While the synthetic dataset allowed precise control over factors such as sound sources, polyphony, and SNR's, we also validated the generalization of our approach using the TUT datasets (DCASE SED Challenge 2016 [21] and 2017 [22]), which introduced real-world complexity with labeled foreground and background audio. The audio samples were segmented into 10-second chunks, with labels assigned based on annotations, and chunks without any class presence were categorized as background for the audio-based adaptation. Due to the dataset's structure, TGAP could not be applied, as isolated class instances within each chunk were unavailable. Applying the zero-shot method, the results obtained were 32.4% with the baseline, 32.3% with text-based adaptation, and 42.4%

with audio-based adaptation, showing that text-based adaptation does not hurt performance, while audio-based adaptation significantly improves it.

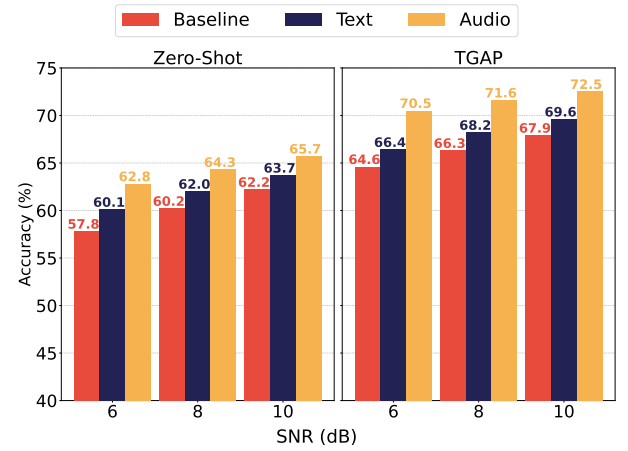


Figure 4: Average performance across all backgrounds under varying SNRs. The Zero-Shot method is shown on the left, and TGAP on the right, with Baseline, Text, and Audio evaluation setups in each.

In brief, our proposed DA method proves effective across different background environments and SNR levels, demonstrating the method's robustness and versatility for broader sound classification tasks. We consistently improve Zero-Shot and TGAP performance. While both text and audio approaches enhance accuracy, audio-based adaptation delivers the most significant gains.

6. Conclusions and future work

We show how prototypical audio-text models face challenges in coherently representing background as a concept, which leads to confusion and performance degradation when classifying sounds in real-world settings. To address these challenges, we propose a domain adaptation method that significantly improves classification accuracy in various background environments and SNR levels. Notably, audio-based adaptations outperform text-based approaches, delivering the most substantial gains. Our experiments also underscore the crucial role of modality alignment in zero-shot sound classification, showing that models capable of minimizing the modality gap experience a marked performance boost. An automatic selection of this parameter τ is left for future research.

7. References

- [1] A. Guzhov, F. Raue, J. Hees, and A. Dengel, "Audioclip: Extending clip to image, text and audio," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 976–980.
- [2] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, "Clap learning audio concepts from natural language supervision," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [3] G. Rey Gozalo and J. M. Barrigón Morillas, "Analysis of sampling methodologies for noise pollution assessment and the impact on the population," *International journal of environmental research and public health*, vol. 13, no. 5, p. 490, 2016.
- [4] P. Foggia, N. Petkov, A. Saggese, N. Strisciuglio, and M. Vento, "Audio surveillance of roads: A system for detecting anomalous sounds," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 1, pp. 279–288, 2016.
- [5] J. Z. Huang, H. Chhabria, and D. Jain, "“Not There Yet”: Feasibility and Challenges of Mobile Sound Recognition to Support Deaf and Hard-of-Hearing People," in *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility*, ser. ASSETS '23. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: <https://doi.org/10.1145/3597638.3608431>
- [6] J. Renaud, R. Karam, M. Salomon, and R. Couturier, "Deep learning and gradient boosting for urban environmental noise monitoring in smart cities," *Expert Systems with Applications*, vol. 218, p. 119568, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0957417423000696>
- [7] Y. Wu*, K. Chen*, T. Zhang*, Y. Hui*, T. Berg-Kirkpatrick, and S. Dubnov, "Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2023.
- [8] R. O. Araz, D. Bogdanov, P. Alonso-Jiménez, and F. Font, "Evaluation of deep audio representations for semantic sound similarity," *IN PRESS*, 2024.
- [9] S. Deshmukh, R. Singh, and B. Raj, "Domain adaptation for contrastive audio-language models," *arXiv preprint arXiv:2402.09585*, 2024.
- [10] A. Hanif, M. T. Agro, M. A. Qazi, and H. Aldarmaki, "Palm: Few-shot prompt learning for audio language models," *arXiv preprint arXiv:2409.19806*, 2024.
- [11] H.-H. Wu, O. Nieto, J. P. Bello, and J. Salamon, "Audio-text models do not yet leverage natural language," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [12] A. Fahim, A. Murphy, and A. Fyshe, "Its not a modality gap: Characterizing and addressing the contrastive gap," *arXiv preprint arXiv:2405.18570*, 2024.
- [13] V. W. Liang, Y. Zhang, Y. Kwon, S. Yeung, and J. Y. Zou, "Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 17 612–17 625. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/702f4db7543a7432431df588d57bc7c9-Paper-Conference.pdf
- [14] T. Kouzelis and V. Katsouros, "Weakly-supervised automated audio captioning via text only training," *arXiv preprint arXiv:2309.12242*, 2023.
- [15] S. S. Kushwaha and M. Fuentes, "A multimodal prototypical approach for unsupervised sound classification," *arXiv preprint arXiv:2306.12300*, 2023.
- [16] J. Salamon, D. MacConnell, M. Cartwright, P. Li, and J. P. Bello, "Scaper: A library for soundscape synthesis and augmentation," in *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2017, pp. 344–348.
- [17] J. Salamon, C. Jacoby, and J. P. Bello, "A dataset and taxonomy for urban sound research," in *Proceedings of the 22nd ACM international conference on Multimedia*, 2014, pp. 1041–1044.
- [18] A. Mesaros, T. Heittola, and T. Virtanen, "A multi-device dataset for urban acoustic scene classification," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop (DCASE2018)*, November 2018, pp. 9–13. [Online]. Available: <https://arxiv.org/abs/1807.09840>
- [19] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [20] D. Dogan, H. Xie, T. Heittola, and T. Virtanen, "Multi-label zero-shot audio classification with temporal attention," *arXiv preprint arXiv:2409.00408*, 2024.
- [21] A. Mesaros, T. Heittola, and T. Virtanen, "TUT database for acoustic scene classification and sound event detection," in *24th European Signal Processing Conference 2016 (EUSIPCO 2016)*, Budapest, Hungary, 2016.
- [22] A. Mesaros, T. Heittola, A. Diment, B. Elizalde, A. Shah, E. Vincent, B. Raj, and T. Virtanen, "DCASE 2017 challenge setup: Tasks, datasets and baseline system," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2017 Workshop (DCASE2017)*, November 2017, pp. 85–92.