

# Mind the Domain Gap: a Systematic Analysis on Bioacoustic Sound Event Detection

Jinhua Liang<sup>†</sup> Ines Nolasco<sup>†</sup> Burooj Ghani<sup>◇</sup> Huy Phan<sup>‡</sup> Emmanouil Benetos<sup>†</sup> Dan Stowell<sup>\*◇</sup>

<sup>†</sup>*Centre for Digital Music, Queen Mary University of London, London, UK*

<sup>◇</sup>*Naturalis Biodiversity Center, Leiden, Netherlands*

<sup>\*</sup>*Department of Cognitive Science and Artificial Intelligence, Tilburg University, Tilburg, Netherlands*

<sup>‡</sup>*Amazon AGI, Cambridge, UK*

**Abstract**—Detecting the presence of animal vocalisations in nature is essential to study animal populations and their behaviors. A recent development in the field is the introduction of the task known as few-shot bioacoustic sound event detection, which aims to train a versatile animal sound detector using only a small set of audio samples. Previous efforts in this area have utilized different architectures and data augmentation techniques to enhance model performance. However, these approaches have not fully bridged the domain gap between source and target distributions, limiting their applicability in real-world scenarios. In this work, we introduce a new dataset designed to augment the diversity and breadth of classes available for few-shot bioacoustic event detection, building on the foundations of our previous datasets. To establish a robust baseline system tailored for the DCASE 2024 Task 5 challenge, we delve into an array of acoustic features and adopt negative hard sampling as our primary domain adaptation strategy. This approach, chosen in alignment with the challenge’s guidelines that necessitate the independent treatment of each audio file, sidesteps the use of transductive learning to ensure compliance while aiming to enhance the system’s adaptability to domain shifts. Our experiments show that the proposed baseline system achieves a better performance compared with the vanilla prototypical network. The findings also confirm the effectiveness of each domain adaptation method by ablating different components within the networks. This highlights the potential to improve few-shot bioacoustic sound event detection by further reducing the impact of domain shift.

**Index Terms**—Bioacoustic sound event detection, few-shot learning, domain adaptation

## I. INTRODUCTION

The automatic detection and analysis of bioacoustic sound events is a valuable tool for assessing animal populations and their behaviours, facilitating conservation efforts, and providing insights into ecological dynamics [1]–[3]. By leveraging advanced machine learning and computational bioacoustics techniques, it is possible to decode the complex acoustic signals of diverse species [4]. This approach can allow us to overcome some of the challenges to a considerable extent, such as background noise, simultaneous calls from various sources, and the major issue of limited annotated data in the face of an expanding volume of recordings [5]. Despite the advancements, the domain-specific challenges of bioacoustic sound event detection – such as distinguishing between various vocalisation

types, dialect types, species versus individuals identification, models capable of decoding both soundscapes and focal recordings – underscore the need for general-purpose systems that are capable of adapting to these specialized tasks [6]–[8]. Addressing these fragmentation and specific challenges in bioacoustic sound event detection, few-shot learning (FSL) emerges as a promising framework capable of transcending the varied subdomains within computational bioacoustics [9], [10]. FSL in sound event detection leverages the concept that a system can be trained on a small, representative set of data samples to detect and classify novel, unseen sound events. By training a single model to adapt to multiple datasets with minimal examples, FSL proposes a robust and versatile solution [11], [12]. This innovation not only addresses the lack of large-scale annotated datasets but also the practicality of applying deep learning techniques across the fragmented landscape of bioacoustic research, offering a pathway to more scalable and efficient detection systems [9].

To diversify methods and accelerate progress in FSL within the bioacoustics domain, an open-science task was organized within the Detection and Classification of Acoustic Scenes and Events (DCASE) framework, focusing on bioacoustic sound event detection (SED). This challenge, open to all , invited participants to develop few shot systems that given only 5 examples of an animal sound (per recording) the system is able to detect the same sound in the rest of the recording. The task featured public datasets, evaluation metrics, documentation, and baseline systems [13]–[15]. The few-shot bioacoustic sound event detection task highlights some unique features: 1) diverse target sounds: sounds of different animals vary greatly due to the distinct mechanism of their vocalisation; 2) sparse distribution of animal vocalisations: the presence of animal sounds is much rarer than the background slices in the recording; and 3) dynamic acoustic environment: the context of target sounds changes drastically due to the heterogeneous nature of various microphones, different sound sources, and environments. We assume these unique features pose a challenge of domain shift, hindering the generalization ability of few-shot bioacoustic detectors.

In this work we explore the problem of domain shift in the field of bioacoustics and present a new baseline system for

<sup>‡</sup>The work does not relate to Huy Phan’s work at Amazon.

the few-shot bioacoustic event detection task in the DCASE 2024 challenge<sup>1</sup>. This is achieved by adapting the system proposed in [16]. Diverse acoustic features are investigated and two domain adaptation techniques, namely, negative hard sampling [17] and transductive learning [18] are systematically ablated to study their impact on the performance on the few-shot bioacoustic event detection task. Additionally, to mitigate the domain shift, we introduce a new dataset for DCASE Task 5 for 2024 by extending the dataset from the previous years challenges. We conduct an additional ablation study and compare the performance on different versions of the dataset to highlight the challenge of domain shift and analyse the mechanism for a good performance in this task. In adherence to the DCASE 2024 Task 5 challenge rules, which require treating each audio file independently to prevent potential class overlap in the evaluation set, our baseline system does not incorporate transductive learning - a technique we evaluated in our 2022 analysis. Instead, our domain adaptation efforts for the 2024 challenge are centered around negative hard sampling, ensuring compliance with the challenge’s guidelines and focusing on methodologies conducive to advancing few-shot bioacoustic event detection research within the stipulated framework.

## II. DATASET

In this work we make use of the development dataset from the DCASE 2022’s few-shot bioacoustic event detection task [14]. As described in [9], each subset of data represents different bioacoustic sources and contains long duration recordings of different species, recorded in various locations and with different recording setups and devices. Due to this diversity, the data is able to represent the nature of the bioacoustics domain which is commonly described more as a set of computational related tasks instead of single one. The training set recordings are paired with multi-class annotations while in the validation set each audio recording is only annotated for a single class of interest. This means that in the validation set, other salient/foreground events exist however these are not the target of the detection task. Table I summarizes the content of each set of data, and it serves to further illustrate the heterogeneous nature of the data. We note that the level of the annotations can also vary across the different subsets. In some, the target classes are different species (see table I for WMW, BV and PB), however in other sets the target classes are different call types of the same species (see table I for ME).

Focusing on the domain shift problem, this dataset provides a realistic example of the expected degree to which data distributions can shift and vary across the different sets of data. In fact the training set has minimum overlap between classes and characteristics of the data in the validation set which constitutes a challenging scenario for system development, particularly since the expectation of similar performance on training and validation sets cannot be guaranteed here. To further highlight the occurrence of domain shifts in bioacoustic

datasets, we also consider the recently updated version of the validation set<sup>2</sup>, see Table I, Validation set 2024. The validation set has been extended to include more flight calls recordings (PB data) and recordings of two new species: Red Deer (RD) and Pilot Whales (PW). It is also of note that with the addition of Pilot whale recordings in the validation set, we are introducing an important domain shift since the acoustic underwater soundscape is completely different from the other acoustic environments.

## III. PROPOSED SYSTEM

To automatically detect the presence of an animal in a recording, the bioacoustic sound event detection system should learn the concept of the vocalisation of a specific species from a limited number of references. This poses a unique challenge in the task of few-shot bioacoustic sound event detection: the target domain could be misaligned with the source domain from which the model learns the concept of animal vocalisation. In Sect. III-A we formulate the problem of few-shot bioacoustic sound event detection and introduce the architecture of the proposed baseline in Sect. III-B. To mitigate the domain gap between a source and a target domain, we elaborate two designs in Sect. III-C.

### A. Problem formulation

Suppose a dataset  $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{D}|}$  where  $\mathbf{x}_i$  denotes the  $i$ -th audio example,  $\mathbf{y}_i \in \mathcal{C}$  denotes the set of discrete labels of the  $i$ -th example, and  $\mathcal{C}$  denotes the label set of  $|\mathcal{C}|$  classes,  $\mathcal{C} = \{1, \dots, |\mathcal{C}|\}$ .

For two non-overlapping subsets  $\mathcal{D}_s, \mathcal{D}_q \subset \mathcal{D}$ , a classical few-shot learning setting can be formulated as a “ $N$ -way  $K$ -shot” problem where a classification task is constituted by: (i) a label subset  $\mathcal{C}_s$  of  $N$  classes sampled from  $\mathcal{C}$ , (ii)  $K$  examples (known as *support* data) sampled from  $\mathcal{D}_s$  for each class in  $\mathcal{C}_s$ , and (iii)  $Q$  examples (known as *query* data) sampled from  $\mathcal{D}_q$  for each class in  $\mathcal{C}_s$ . With a model family chosen as  $p_\theta(\mathbf{x})$ , with unknown parameters  $\theta \in \Theta$ , the problem boils down to maximizing the average likelihood of all the samples from  $\mathcal{D}_q$  under the model parameter:

$$\theta^* = \underset{\theta \in \Theta}{\operatorname{argmin}} \frac{1}{|\mathcal{D}_q|} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_q} \mathcal{L}(\mathbf{y}|\mathbf{x}, \theta) \quad (1)$$

### B. Model architecture

We adopt prototypical networks [19] to directly learn the concept of animal vocalisation in the latent space. For a class  $n \in \mathcal{C}_s$ , let  $\mathcal{S}_n$  be the subset of support examples containing the sound events of this class,  $|\mathcal{S}_n| = K$ . Then the set of sound events is represented as  $\mathcal{S}_n^+ \subset \mathcal{S}_n$  where the set of negative segmentation is denoted as  $\mathcal{S}_n^- = \mathcal{S}_n / \mathcal{S}_n^+$ .

The positive prototype  $\mathbf{a}_n^+$  of class  $n$  is then derived as the mean of embedding vectors of the support examples in  $\mathcal{S}_n^+$ . Formally,

$$\mathbf{a}_n^+ = \frac{1}{K} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{S}_n^+} f_\theta(\mathbf{x}), \quad (2)$$

<sup>1</sup><https://dcase.community/challenge2024/>

<sup>2</sup><https://doi.org/10.5281/zenodo.10829604>

TABLE I  
SUMMARY OF DATASET CHARACTERISTICS.

	Name and species	Mic type	# Audio files	Total duration	# Labels	# Events	Mean Event duration (s)
Training set	BV: BirdVox-DCASE-10h	fixed	5	10 hours	11	9026	0.15
	HT: Hyenas	various	5	5 hours	5	611	1.42
	MT: Meerkats	animal mounted	2	70 mins	4	1294	0.14
	JD: Jackdaws	mobile	1	10 mins	1	357	0.12
	WMW: Western Mediterranean Birds	various	161	5 hours	26	2941	1.54
Validation set 2022	HB: Humberg mosquitoes	handheld	10	2.38 hours	1	712	11.6
	PB: Polish Baltic Sea bird flight calls	fixed	6	3 hours	2	292	0.11
	ME: Meerkats	animal mounted	2	20 mins	2	73	0.19
Validation set 2024	Validation set 2022 +						
	PB24: Polish Baltic Sea bird flight calls	fixed	4	120 mins	2	350	0.08
	RD: Red Deer	fixed	6	18 hours	1	1372	1.52
	PW: Pilot Whales	animal mounted	15	24 hours	1	705	2.21

where  $f$  denotes the embedding mapping realized by the model whose parameters are denoted collectively as  $\theta$ . Likewise, the negative prototype  $\mathbf{a}_n^-$  of class  $n$  is calculated by

$$\mathbf{a}_n^- = \frac{1}{K} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{S}_n^-} f_\theta(\mathbf{x}), \quad (3)$$

Given a query example  $\mathbf{x}_q$ , the model performs  $N$  binary classifications independently by producing a probability distribution based on a softmax over distances between  $\mathbf{x}_q$  and the positive and negative prototypes in the latent space. More specifically, the probability that  $\mathbf{x}_q$  is classified as class  $n \in \mathcal{C}_s$  is calculated as

$$p_\theta(\hat{\mathbf{y}}_{q,n} | \mathbf{x}_q) = \frac{e^{-d(f_\theta(\mathbf{x}_q), \mathbf{a}_n^+)}}{\sum_{i \in \mathcal{C}_s} \left[ e^{-d(f_\theta(\mathbf{x}_q), \mathbf{a}_i^+)} + e^{-d(f_\theta(\mathbf{x}_q), \mathbf{a}_i^-)} \right]}, \quad (4)$$

where  $\hat{\mathbf{y}}_{q,n}$  is the predicted label for  $\mathbf{x}_q$  w.r.t. class  $n$ ,  $d$  is a distance function, such as  $\ell_2$  or cosine distance. The network is trained to minimize the negative log-probability of the true class over the  $N \times Q$  query examples:

$$\mathcal{L}(\theta) = \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}} \sum_{n \in \mathcal{C}_s} -\log p_\theta(\hat{\mathbf{y}}_n = \mathbf{y}_n | \mathbf{x}), \quad (5)$$

where  $\mathcal{Q}$  is the set of query examples,  $|\mathcal{Q}| = N \times Q$ .

While prototypical networks perform well in few-shot learning settings [20], their performance still suffers when the domain of the development dataset is drastically different from the one of the test dataset. To mitigate the domain misalignment issue, we adopt two strategies in the following part.

### C. Domain adaptation

**Negative hard sampling.** As shown in eqn. (4), both positive and negative prototypes are important to predict the probability distribution over each class. Bioacoustic sound events are sparsely distributed throughout a recording. While these sound events are hardly available during training, there are many negative audio segments available. This drives us to investigate how to create “meaningful” negative prototypes by using hard sampling techniques [17] during inference stage. Instead of using all negative audio segments, we randomly sample a proportion of segmentations for negative prototype measurement. In other words, we inject the randomness to the

cluster of negative segments such that the trained model can generalise to evaluation set well. In addition, we enhanced the trained prototypical network with different sets of negative segments to further improve the performance of the proposed model on the evaluation set.

**Transductive learning.** Acoustic environment of test recordings may vary from the development domain due to many factors, such as the difference of recording contexts, recording devices, and pre-processing methods. To alleviate the domain mismatch caused by different acoustic environments, we apply transductive inference [18] by learning not only from the labelled events of the training set but the first five labelled examples for each audio recording in the validation set.

## IV. EXPERIMENT

### A. Experiment setup

For a fair comparison, we adopted the same convolutional neural network as the backbone for feature extraction. Following [16], this network consists of three  $3 \times 3$  convolutional layers, with dimension of 64, 128, 64. Each of the convolutional layer is followed by batch normalisation, leak ReLU activation, and a  $2 \times 2$  max-pooling layer. For each convolutional block, we added a skip connection followed by a max-pooling layer to avoid catastrophic forgetting.

We trained and evaluated the models on DCASE 2022 Task5 and DCASE 2022 Task5 datasets, respectively. For all experiments, we resampled audio recordings to 22.5 kHz sampling rate and applied short-time Fourier transformation (STFT) with 1024 of window length and 256 of hop length. We set the number of frequency bins as 128 to extract log-Mel spectrograms (Log Mel) from audio examples. To investigate the impact of acoustic features on this task, we also experimented diverse acoustic features, including the Mel spectrogram (Mel), Mel-frequency cepstral coefficients (MFCC), delta features of MFCC (delta MFCC), per-channel energy normalization (PCEN), and their combinations as input features. Specifically, the dimension of MFCC features is set to 32 in the experiments.

We trained the models using RTX A5000 GPUs. Following [16], we set the initiate learning rate to 0.001 with exponential decay of 0.65 for every 10 epochs. We adopted early stopping by monitoring the accuracy score and took the

TABLE II  
BENCHMARKING THE BASELINE SYSTEM (%) ON DCASE 2022 TASK 5 AND DCASE 2024 TASK 5 VALIDATION SETS. THE GRAY BAR HIGHLIGHTS THE PERFORMANCE SCORE OF OUR PROPOSED BASELINE FOR DCASE 2024 TASK 5.

Negative hard sampling	Transductive learning	DCASE 2022 Task 5			DCASE 2024 Task 5		
		Precision	Recall	$F_1$ -score	Precision	Recall	$F_1$ -score
		55.92±2.6	40.61±5.6	46.78±2.9	44.94±3.39	45.89±4.80	45.23±0.48
	✓	55.17±3.6	43.59±0.4	48.66±1.6			
✓		48.33±2.3	56.64±2.0	52.09±0.7	<b>56.18±0.61</b>	<b>48.64±0.23</b>	<b>52.14±0.20</b>
✓	✓	<b>66.43±3.6</b>	<b>61.28±1.5</b>	<b>63.67±1.0</b>			

TABLE III  
PERFORMANCE COMPARISON (%) OF SYSTEMS WITH DIVERSE ACOUSTIC FEATURES ON THE DCASE 2022 TASK 5 VALIDATION SET.

Feature	Precision	Recall	$F_1$ -score
Mel	38.56±1.2	49.01±1.5	43.14±0.9
Log mel	66.43±3.6	61.28±1.5	<b>63.67±1.0</b>
Log mel + MFCC	58.89±1.1	<b>65.70±1.0</b>	62.10±0.5
Log mel + delta MFCC	61.22±1.4	62.75±0.6	61.96±0.7
PCEN	<b>68.00±2.2</b>	53.70±1.7	59.97±0.7
PCEN + MFCC	63.88±1.7	57.65±0.7	60.59±0.6
PCEN + delta MFCC	59.47±5.1	52.22±1.8	55.47±1.5

checkpoint of the highest accuracy score as our best model. During evaluation, we fixed the set of negative audio segments across models with the same number of negative segments. For each experiment, we calculated the average performance scores with 95% confidence score over five independent trials to avoid impact of randomness. Our code and implementation can be found at this url<sup>3</sup>.

### B. Experimental results

Table II details the baseline system’s performance on the DCASE 2022 Task 5 and DCASE 2024 Task 5 validation sets, alongside an ablation study for domain adaptation techniques applied to the DCASE 2022 Task 5 dataset. The baseline system achieved a precision of 48.33%, a recall of 56.64%, and an  $F_1$ -score of 52.09% for DCASE 2022 Task 5, and a precision of 56.18%, a recall of 48.64%, and an  $F_1$ -score of 52.14% for DCASE 2024 Task 5, outperforming the vanilla prototypical network (the one without negative sampling) by a large margin. We explored integrating transductive learning with the baseline system, which showed improved performance with an increase in precision by 18.1%, recall by 4.64%, and  $F_1$ -score by 11.58% absolute, indicating the impact of domain shift in few-shot bioacoustic sound event detection.

Given the DCASE task rules, which stipulate treating each audio file independently so that each item in the evaluation set is handled as an independent few-shot scenario, we adhered to these guidelines in our approach to the DCASE 2024 Task 5 challenge. Consequently, the version of transductive learning we used for our 2022 analysis was not applied for the 2024 dataset. For the DCASE 2024 Task 5, we focused our domain adaptation efforts on negative hard sampling, excluding transductive learning from our methodology. This strategic decision reflects our commitment to adhering to the DCASE task rules while still advancing research in domain adaptation

<sup>3</sup>[https://github.com/c4dm/dcase-few-shot-bioacoustic/tree/main/baselines/dcase2024\\_task5](https://github.com/c4dm/dcase-few-shot-bioacoustic/tree/main/baselines/dcase2024_task5)




Fig. 1. Ablation study on the number of negative audio segments and number of enhanced models.

for few-shot bioacoustic sound event detection; various forms of domain adaptation remain as future research possibilities.

Table III compares diverse acoustic features on DCASE 2022 Task 5 validation set. The system with log Mel features achieves the best  $F_1$ -score of compared to those of other acoustic features. Among these models, the combination of log Mel and MFCC yielded the best recall score of 65.70% while PCEN had the best precision score of 68.00%. Provided that we expect a balanced performance of bioacoustic sound event detectors, we applied log Mel as the input acoustic features for the baseline. In addition, the precision scores of log Mel and PCEN are lower than their counterpart combined with MFCC, respectively. This probably indicates that equipping input features with MFCC is beneficial to models’ performance in terms of the precision score.

Fig. 1 illustrates the ablation study on the number of negative audio segments and enhanced models. Enhanced models hereby are referred to as the model of the same checkpoint with different sets of negative segments. As the number of sets of negative audio segments increases, the system’ performance varies less. Likewise, as the number of negative segments increases, the performance score of systems becomes more stable. In addition, the system of three models with 150 negative segments achieved the best  $F_1$ -score performance score. This guides us to create the baseline system of the DCASE 2024 Task 5 dataset.

Fig. 2 compares the performance of the baseline on DCASE 2024 Task 5 validation set by species. It can be observed that the baseline yielded the best performance on the species of humbug mosquitoes while performing the worst on that of pilot




Fig. 2. Species-wise performance of the baseline on DCASE 2024 Task 5 validation set.

whales. The latter observation is in line with our expectation: the model did not learn any underwater sounds from the training set and thus cannot be adapted to the vocalisation of marine animals well.

## V. CONCLUSION

In this work, we benchmarked the task of few-shot bioacoustic sound event detection using the recently updated DCASE 2024 Task 5 dataset. We pinpointed domain shift as a significant challenge within this task. To substantiate our hypothesis, we improved on the previous dataset, aimed at enhancing the diversity and enlarging the size of the preceding version. We experimented with a variety of acoustic features as inputs for the baseline model to assess the effects of feature engineering on bioacoustic computations. Additionally, we improved the prototypical network by integrating negative hard sampling and transductive learning, aiming to better adapt the network to the target domain. In compliance with the guidelines of DCASE Task 5, we adopted the enhanced prototypical network with negative hard sampling as the new baseline for the DCASE 2024 Task 5 competition. Our goal is to stimulate future research on domain adaptation in the context of few-shot bioacoustic sound event detection.

## VI. ACKNOWLEDGEMENTS

We would like to thank Haohe Liu for his valuable discussion on this work, and Hanna Pamula, Helen Whitehead and Frants Jensen for contributing the new data for the 2024 validation set. The research utilised Queen Mary’s Apocrita HPC facility, supported by QMUL Research-IT, <http://doi.org/10.5281/zenodo.438045>. J. Liang is supported by the Engineering and Physical Sciences Research Council [grant number EP/T518086/1]. I. Nolasco is supported by the Engineering and Physical Sciences Research Council [grant number EP/R513106/1]. B. Ghani is supported by EU-funded projects MAMBO and GUARDEN. E. Benetos is supported by a RAEng/Leverhulme Trust Research Fellowship [grant number LTRF2223-19-106]. For the purpose of open access, the authors

have applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising.

## REFERENCES

- [1] P. E. Caiger, M. J. Dean, A. I. DeAngelis, L. T. Hatch, A. N. Rice, J. A. Stanley, C. Tholke, D. R. Zemeckis, and S. M. Van Parijs, “A decade of monitoring atlantic cod *gadus morhua* spawning aggregations in massachusetts bay using passive acoustics,” *Marine Ecology Progress Series*, vol. 635, pp. 89–103, 2020.
- [2] S. Gillings and C. Scott, “Nocturnal flight calling behaviour of thrushes in relation to artificial light at night,” *Ibis*, vol. 163, no. 4, pp. 1379–1393, 2021.
- [3] D. Tuia, B. Kellenberger, S. Beery, B. R. Costelloe, S. Zuffi, B. Risse, A. Mathis, M. W. Mathis, F. Van Langevelde, T. Burghardt *et al.*, “Perspectives in machine learning for wildlife conservation,” *Nature communications*, vol. 13, no. 1, pp. 1–15, 2022.
- [4] T. Ganchev, *Computational bioacoustics: biodiversity monitoring and assessment*. Walter de Gruyter GmbH & Co KG, 2017, vol. 4.
- [5] D. Stowell, “Computational bioacoustics with deep learning: a review and roadmap,” *PeerJ*, vol. 10, p. e13152, 2022.
- [6] B. Ghani, T. Denton, S. Kahl, and H. Klinck, “Global birdsong embeddings enable superior transfer learning for bioacoustic classification,” *Scientific Reports*, vol. 13, no. 1, p. 22876, 2023.
- [7] J. Hamer, E. Triantafyllou, B. van Merriënboer, S. Kahl, H. Klinck, T. Denton, and V. Dumoulin, “Birb: A generalization benchmark for information retrieval in bioacoustics,” *arXiv preprint arXiv:2312.07439*, 2023.
- [8] M. Boudiaf, T. Denton, B. Van Merriënboer, V. Dumoulin, and E. Triantafyllou, “In search for a generalizable method for source free domain adaptation,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 2914–2931.
- [9] I. Nolasco, S. Singh, V. Morfi, V. LOSTANLEN, A. Strandburg-Peshkin, E. Vidana-Vila, L. Gill, H. Pamula, H. Whitehead, I. Kiskin *et al.*, “Learning to detect an animal sound from five examples,” *Ecological informatics*, vol. 77, p. 102258, 2023.
- [10] R. Li, J. Liang, and H. Phan, “Few-shot bioacoustic event detection: Enhanced classifiers for prototypical networks,” in *Proceedings of the 7th Detection and Classification of Acoustic Scenes and Events 2022 Workshop (DCASE2022)*, Nancy, France, November 2022.
- [11] J. Liang, H. Phan, and E. Benetos, “Leveraging label hierarchies for few-shot everyday sound recognition,” in *Proceedings of the 7th Detection and Classification of Acoustic Scenes and Events 2022 Workshop (DCASE2022)*, Nancy, France, November 2022.
- [12] J. Liang, X. Liu, H. Liu, H. Phan, E. Benetos, M. D. Plumbley, and W. Wang, “Adapting Language-Audio Models as Few-Shot Audio Learners,” in *Proc. INTERSPEECH 2023*, 2023, pp. 276–280.
- [13] V. Morfi, I. Nolasco, V. LOSTANLEN, S. Singh, A. Strandburg-Peshkin, D. Benvent, and D. Stowell, “Few-shot bioacoustic event detection: A new task at the dcase 2021 challenge.”
- [14] I. Nolasco, S. Singh, E. Vidana-Vila, E. Grout, J. Morford, M. Emmerson, F. Jensen, H. Whitehead, A. Strandburg-Peshkin *et al.*, “Few-shot bioacoustic event detection at the dcase 2022 challenge,” *arXiv preprint arXiv:2207.07911*, 2022.
- [15] I. Nolasco, B. Ghani, S. Singh, E. Vidana-Vila, H. Whitehead, E. Grout, M. Emmerson, F. Jensen, I. Kiskin, J. Morford *et al.*, “Few-shot bioacoustic event detection at the dcase 2023 challenge,” *arXiv preprint arXiv:2306.09223*, 2023.
- [16] H. Liu, X. Liu, X. Mei, Q. Kong, W. Wang, and M. D. Plumbley, “Surrey system for dcase 2022 task 5 : Few-shot bioacoustic event detection with segment-level metric learning,” DCASE2022 Challenge Technical Report, Tech. Rep., 2022.
- [17] J. Robinson, C.-Y. Chuang, S. Sra, and S. Jegelka, “Contrastive learning with hard negative samples,” *arXiv preprint arXiv:2010.04592*, 2020.
- [18] A. Arnold, R. Nallapati, and W. W. Cohen, “A comparative study of methods for transductive transfer learning,” in *Seventh IEEE international conference on data mining workshops (ICDMW 2007)*. IEEE, 2007, pp. 77–82.
- [19] J. Snell, K. Swersky, and R. S. Zemel, “Prototypical networks for few-shot learning,” 2017.
- [20] J. Liang, H. Phan, and E. Benetos, “Learning from taxonomy: Multi-label few-shot classification for everyday sound recognition,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 1–5.