

# SSL-NET: A SYNERGISTIC SPECTRAL AND LEARNING-BASED NETWORK FOR EFFICIENT BIRD SOUND CLASSIFICATION

Yiyuan Yang    Kaichen Zhou    Niki Trigoni    Andrew Markham

Department of Computer Science, University of Oxford, United Kingdom

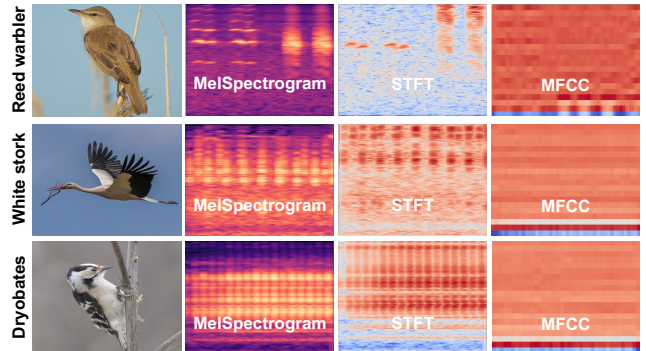
## ABSTRACT

Efficient and accurate bird sound classification is of importance for ecology, habitat protection and scientific research, as it plays a central role in monitoring the distribution and abundance of species. However, prevailing methods typically demand extensively labeled audio datasets and have highly customized frameworks, imposing substantial computational and annotation loads. In this study, we present an efficient and general framework called SSL-Net, which combines spectral and learned features to identify different bird sounds. Encouraging empirical results gleaned from a standard field-collected bird audio dataset validate the efficacy of our method in extracting features efficiently and achieving heightened performance in bird sound classification, even when working with limited sample sizes. Furthermore, we present three feature fusion strategies, aiding engineers and researchers in their selection through quantitative analysis.

**Index Terms**— Audio pattern recognition, Spectral information, Pretrained model, Feature fusion

## 1. INTRODUCTION

There is a significant and increasing need for automated pipelines for bird sound classification, as these techniques hold great potential for habitat/species protection and advancing interdisciplinary scientific research. Recently, numerous large-scale bird audio datasets created through crowdsourcing have played a key role in increasing research efforts [1, 2]. These datasets facilitate the development and deployment of data-driven models aimed at solving the challenge of bird audio classification [3]. Whilst many bespoke models have been proposed for bird sound classification, we note that training these from scratch is particularly label inefficient [4]. Therefore, we see that pretrained models are increasingly used to accelerate research by acting as well-trained feature extractors. These general models are tailored to different domains through fine-tuning and parameter-efficient classification heads [5]. For example, pretrained models from computer vision (e.g. ResNet [6]) show reasonable out-of-the-box classification performance by treating simple spectrographic features as images. Conversely, in the general acoustic analysis field, some recent works like LEAF [7] and BEATs [8]



**Fig. 1.** Visualization examples of bird sound spectral features, including MelSpectrogram, STFT and MFCC.

have made substantial progress in learning features directly from raw waveforms without any pre-transformation.

Although the above datasets and algorithms have contributed to bird sound classification tasks, it is not immediately clear which approach to follow to build an efficient (both in terms of training sample size and model parameter count) and accurate model, with low effort. Bird sound classification is a specialized sub-field of acoustic classification which has significant challenges - a species may have a wide repertoire of vocalizations, yet may sound quite similar to another species. This requires highly discriminative feature extraction, yet we require our approach to be label-efficient, due to the high cost of annotating extensive datasets [9]. This contradiction can be resolved to an extent by using pretrained models, but we first need to transform the raw acoustic data to a more compact representation. To perform this transformation, we note that there is an inherent tension between spectral feature extraction and learnable feature extraction, with each having different relative merits. On the one hand, using spectral features such as a MelSpectrogram (MEL), Short-Time Fourier Transform (STFT) or Mel-Frequency Cepstral Coefficient (MFCC) is well grounded in acoustic signal processing. Yet, even choosing which one of these transformations and its parameters involves significant effort and tuning. On the other hand, learnable features such as BEATs or LEAF operate directly on the raw waveform and are well-trained on large corpora of general acoustic data. However, its direct use for bird sound classification remains to be investigated. Over-

all, while spectral feature-based methods currently achieve better classification performance, pretrained models are more sample-efficient and parameter-efficient and can yield diverse representation outcomes.

In this study, we propose an efficient **Synergistic Spectral and Learning-based Network** called **SSL-Net** to handle the challenges bird sound classification. Instead of attempting to extensively tune models and architectures for maximum performance, we show how to synergize multiple approaches to yield efficient and accurate classification. By using spectral representations to compensate for the distribution bias of learned features, our method achieves leading results. Importantly, it does so with low practitioner effort, low model complexity and high sample efficiency. Specifically, our main contributions are:

- We present an efficient bird sound classifier based on synergistic spectral and learned information, which uses spectral representations to mitigate distribution bias in the pretrained model features and enhances performance and robustness with minimal labels.
- We put forth different fusion strategies to mitigate the distribution bias within spectral and learned features, thereby yielding enhanced feature representations for the classification task.
- Through comprehensive evaluations and ablation experiments with real-world bird audio data, our framework aids engineers in designing efficient and accurate classifiers with low effort.

## 2. RELATED WORK

AI-based bird sound classifiers have been extensively researched. Many of them leverage acoustic characteristics (e.g., MEL, STFT, and MFCC shown in Fig. 1) as feature maps, subsequently integrating them as images and using machine learning-based models for classification. For example, Xie and Zhu [10] proposed to use handcrafted features combined with deep learning to classify 14 bird sounds and achieved good performance. Besides, Tuncer *et al.* [11] presented a multileveled and handcrafted features-based machine learning model and good results in a dataset of 18 categories. Moreover, Nanni *et al.* [12] and Jasim *et al.* [13] designed different enhanced convolutional neural networks with spectral features to improve the model classification ability as well as robustness respectively. There are also some related and similar researches [14, 15, 16, 17].

In addition, the audio-pretrained models developed over the past two years (e.g., Panns [18], LEAF [7] and BEATs [8]) are different from the above-mentioned methods. They rely on data-driven mining of representations that deviate from the traditional spectral acoustic features and have a wide range of applications, including acoustic feature extraction, audio classification, emotion analysis, keyword spotting, etc. Therefore, it is possible to efficiently fine-tune the generic audio-pretrained model in the bird sound classification tasks.

## 3. METHODOLOGY

In this section, we introduce the proposed SSL-Net as illustrated in Fig. 2. It comprises two branches: the learned branch and the spectral branch, integrated through a feature fusion module. The former leverages an audio-pretrained model followed by an encoder that generates a semantic feature map. The latter utilizes acoustic features in conjunction with an encoder to derive a spectral feature map. Within the feature fusion module, we explore three distinct fusion strategies and apply them to merge the feature maps obtained from both branches for the final classification task.

### 3.1. Learned branch

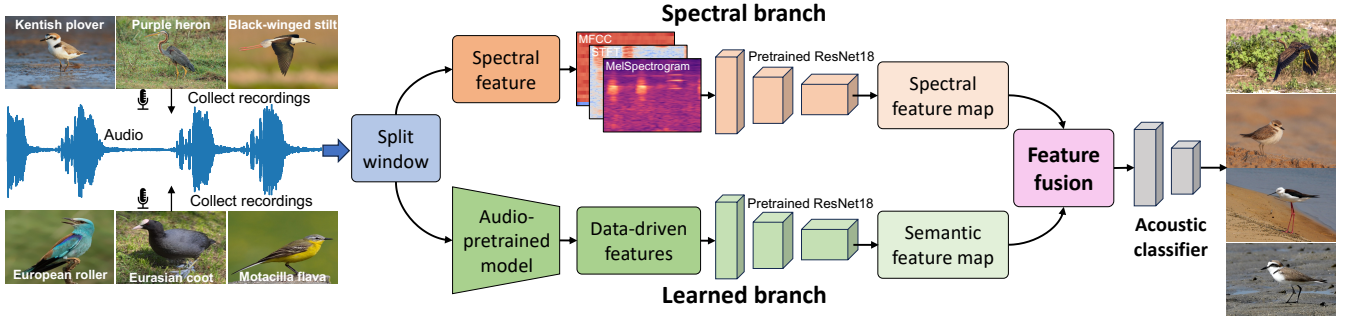
The learned branch focuses on leveraging audio-pretrained models for general acoustic feature extraction. It is composed of two components: an audio-pretrained model and a pretrained model-based encoder. This approach serves two purposes: 1) enhancing the efficiency of feature extraction while 2) concurrently diminishing the required training sample size. We use BEATs [8] as the backbone of the acoustic pretrained model. It extracts data-driven representations  $\mathbf{a}_{\text{sem}}$  from split original audio data  $\mathbf{x}$ , as  $\mathbf{a}_{\text{sem}} = \text{BEATs}(\mathbf{x})$ . After resizing representation  $\mathbf{a}_{\text{sem}}$  to the same size, the pretrained ResNet18 [6] backbone is used as the encoder for further feature mining and to obtain the semantic feature map  $\mathbf{f}_{\text{sem}}$ .

Since the full module of this branch uses pretrained models, it is very efficient without adding more weights. However, the semantic feature map is biased for bird sounds because it is pretrained using other general species' sounds dataset.

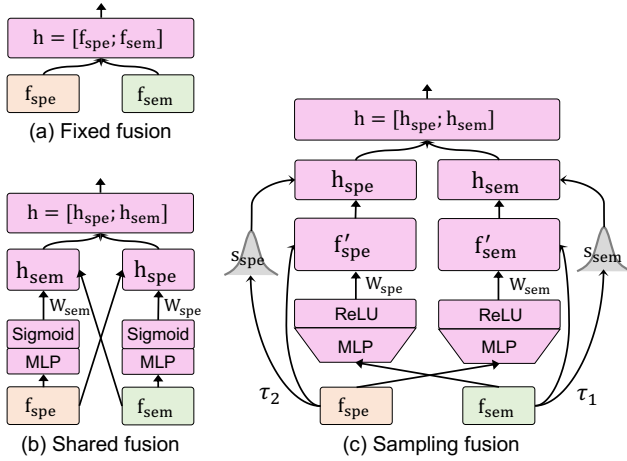
### 3.2. Spectral branch

As the learned branch primarily encapsulates information from the general sound domain, it inherently lacks insights into the target domain (i.e., bird sounds). To mitigate it, we introduce the spectral branch which serves to compensate for the overall distribution by incorporating information pertinent to the target domain. It includes two modules: a spectral features extractor and a pretrained encoder. The former extractor processes the segmented audio signal, extracting key acoustic attributes based on parameters suitable for bird audio, namely the MEL, STFT, and MFCC. Subsequently, these three features are resized to uniform dimensions and concatenated as individual channels. This composite representation  $\mathbf{a}_{\text{spe}}$  is then input into a pretrained ResNet18 [6] for further extraction. The output of the last convolutional layer is used as the spectral feature map  $\mathbf{f}_{\text{spe}}$ .

Similar to the first branch, the spectral branch exclusively employs pretrained models and incorporates acoustic spectral features that require no iterative training. Consequently, this entire branch operates with remarkable efficiency. Furthermore, it can mitigate the biased distribution in the preceding branch within the target domain. More details about the branches will be shown in the ablation experiment in Sec. 4.



**Fig. 2.** The architecture of the proposed SSL-Net for efficient bird sound classification. It is made up of a **learned branch** constructed using the audio-pretrained model and a **spectral branch** built on various acoustic spectral representations. For the **feature fusion**, we present three feasible and selectable strategies, the specifics are illustrated in Fig. 3. Moreover, a lightweight **acoustic classifier** is used for the final sound classification.



**Fig. 3.** Flow charts of the proposed fixed, shared, and sampling feature fusion strategies.

### 3.3. Feature fusion module

Building on top of the extracted semantic feature  $\mathbf{f}_{\text{sem}}$  and spectral feature  $\mathbf{f}_{\text{spe}}$  from the two branches, we design three feature fusion strategies with the same input and output, as shown in Fig. 3, to facilitate the transfer from the general acoustic domain to the target domain and merge the complementary features effectively.

1) **Fixed fusion** is a straightforward approach, as shown in Fig. 3(a). Two representations  $\mathbf{f}_{\text{sem}}$  and  $\mathbf{f}_{\text{spe}}$  are directly concatenated into a new fusion feature  $\mathbf{h}$ . This most efficient fusion can be modeled as  $\mathbf{h} = \text{Concat}[\mathbf{f}_{\text{spe}}; \mathbf{f}_{\text{sem}}]$ .

2) **Shared fusion** is similar to the attention mechanism [19], as shown in Fig. 3(b). This strategy enables the fusion process to be simultaneously trained alongside other modules by conditioning and reweighting both features concurrently. First, we perform the following operations on the semantic feature  $\mathbf{f}_{\text{sem}}$  and spectral feature  $\mathbf{f}_{\text{spe}}$ .

$$\mathbf{W}_{\text{sem}} = \text{Sigmoid}(\text{Linear}(\mathbf{f}_{\text{spe}})), \quad (1)$$

$$\mathbf{W}_{\text{spe}} = \text{Sigmoid}(\text{Linear}(\mathbf{f}_{\text{sem}})), \quad (2)$$

where  $\text{Linear}(\cdot)$  denotes a multilayer perceptron (MLP).  $\mathbf{W}_{\text{sem}}$  and  $\mathbf{W}_{\text{spe}}$  are the new fusion weights. Sigmoid is used to ensure the weights are in the range of 0 to 1. The next step is to perform element-wise multiplication of the weights with the features and get representations  $\mathbf{h}_{\text{sem}}$  and  $\mathbf{h}_{\text{spe}}$ . The two representations are concatenated into the final feature  $\mathbf{h}$ .

$$\mathbf{h} = \text{Concat}[\mathbf{f}_{\text{spe}} \odot \mathbf{W}_{\text{spe}}; \mathbf{f}_{\text{sem}} \odot \mathbf{W}_{\text{sem}}]. \quad (3)$$

3) **Sampling fusion** refers to learning a stochastic function that samples in two features, as shown in Fig. 3(c). Compared with the shared fusion, it could achieve more robust feature fusion. Firstly, we implement the operations on the semantic feature  $\mathbf{f}_{\text{sem}}$  and spectral feature  $\mathbf{f}_{\text{spe}}$ .

$$\mathbf{f}'_{\text{sem}} = \mathbf{f}_{\text{sem}} \odot \text{ReLU}(\text{Linear}(\mathbf{f}_{\text{spe}})), \quad (4)$$

$$\mathbf{f}'_{\text{spe}} = \mathbf{f}_{\text{spe}} \odot \text{ReLU}(\text{Linear}(\mathbf{f}_{\text{sem}})), \quad (5)$$

where the  $\text{Linear}(\cdot)$  will boost the dimensions to double for the following distribution operations.

Next, the original features,  $\mathbf{f}_{\text{sem}}$  and  $\mathbf{f}_{\text{spe}}$ , are individually sampled using the Gumbel-Softmax (GS) distribution [20]. It is selected due to its differentiability during the sampling procedure. This property enables backpropagation, thus facilitating end-to-end model training.

$$\mathbf{s}_{\text{sem}} \sim p(\mathbf{s}_{\text{sem}} | \mathbf{f}_{\text{sem}}, \tau_1) = \text{GS}(\mathbf{f}_{\text{sem}}), \quad (6)$$

$$\mathbf{s}_{\text{spe}} \sim p(\mathbf{s}_{\text{spe}} | \mathbf{f}_{\text{spe}}, \tau_2) = \text{GS}(\mathbf{f}_{\text{spe}}), \quad (7)$$

where  $\tau_1$  and  $\tau_2$  represent the temperature parameter in GS distribution. Finally, we multiply the GS distribution with new representations correspondingly and get the final fusion feature representations  $\mathbf{h}_{\text{sem}}$  and  $\mathbf{h}_{\text{spe}}$ , which is:

$$\mathbf{h} = \text{Concat}[\mathbf{f}'_{\text{spe}} \odot \mathbf{s}_{\text{spe}}; \mathbf{f}'_{\text{sem}} \odot \mathbf{s}_{\text{sem}}]. \quad (8)$$

### 3.4. Final acoustic classifier

We apply a lightweight MLP to classify the fused features and get the final classification score. The iterative training process is performed using the cross-entropy loss function.

**Table 1.** The effectiveness of different single branch methods with WMWB dataset [21]. The results at \* are different due to the different classifier with the same number of parameters.

Branch	Model	Feature	Metric				Param (#)
			Acc $\uparrow$ (%)	Prc $\uparrow$ (%)	Rec $\uparrow$ (%)	F1 $\uparrow$	
Baseline [21]	ResNet50	MEL	77.70	74.68	72.33	73.60	542,292
	ResNet18	MEL*	76.10	72.23	71.98	72.42	
Learned	ResNet18	LEAF	83.60	<b>86.16</b>	77.14	78.50	68,500
		BEATs	83.23	80.56	78.61	79.02	
Spectral	ResNet18	MEL*	83.68	81.52	74.43	76.78	
		STFT	83.45	83.25	73.98	74.99	
		MFCC	83.31	81.17	72.36	74.26	
		ALL	<b>84.02</b>	83.11	<b>82.70</b>	<b>81.63</b>	

**Table 2.** The effectiveness of different fusion strategies with different branches. We use accuracy, precision, recall and F1-score as evaluation metrics.

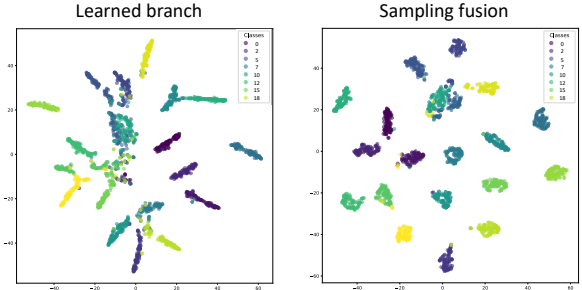
Fusion	Branch	Metric				Param (#)
		Acc $\uparrow$ (%)	Prc $\uparrow$ (%)	Rec $\uparrow$ (%)	F1 $\uparrow$	
Fixed	LEAF+BEATs	83.48	85.34	79.06	79.88	134,036
	Spectral+LEAF	83.33	87.78	84.24	83.71	
	Spectral+BEATs	83.65	86.92	86.28	86.12	
Shared	LEAF+BEATs	84.62	85.59	80.89	81.90	2,260,964
	Spectral+LEAF	<b>87.29</b>	90.38	85.81	86.53	
	Spectral+BEATs	86.54	89.60	88.29	88.68	
Sampling	LEAF+BEATs	83.37	84.79	80.03	81.59	4,262,708
	Spectral+LEAF	86.06	<b>91.98</b>	86.14	87.48	
	Spectral+BEATs	85.70	89.69	<b>88.77</b>	<b>88.79</b>	

#### 4. EXPERIMENTS


In this section, we validate the proposed SSL-Net using the Western Mediterranean Wetland Birds (WMWB) dataset [21] collected from Xeno-canto [2], which consists of 20 classes over 200 minutes. We also perform full ablation experiments and trade-offs to evaluate each module in the pipeline.

Table 1 presents the results from the single branch methods, where 200 labeled samples per species are used. The learned branch methods and spectral branch methods outperform the baseline (ResNet50) with fewer learnable parameters. Besides, the learned branch achieves better performance than a single spectral feature, which demonstrates the adaptability and transferability of the generic audio-pretrained model to the fine-grained task of bird sound classification. In addition, the spectral branch utilizing three features outperforms the methods solely from the learned branch, as spectral features are better suited for the target domain (bird sound).

Table 2 shows the effectiveness of different fusion strategies with different branches. The majority of BEATs+Spectral-based methods achieved optimal results, which shows the effectiveness of the BEATs as a generic audio-pretrained model. However, the results of BEATs+LEAF (fusing two learned branches) are comparatively subpar, due to both representations being data-driven general acoustic representations, which introduces distributional bias in the target domain. As for fusion strategies, the sampling fusion excels in classification performance, but the shared fusion, while slightly sacrificing classification performance, reduces trainable parameters and model complexity, occupying the efficiency-classification performance trade-off peak. Lastly, the simplest but quite efficient fixed fusion, using the fewest



**Fig. 4.** Visualisation of the single branch features and fused features before final classification using t-SNE.



**Fig. 5.** The impact of the number of labeled samples on performance with different methods.

parameters, also achieved good results. In summary, three fusion strategies effectively synergize distinct representations from different domains, resulting in state-of-the-art outcomes.

To further study the advantages of fusion, we visualize the pre-fusion (learned branch for example) and post-fusion features (sampling fusion for example), as shown in Fig 4. Clearly, the fusion module yields more discriminative features, allowing for more efficient classification with fewer learnable parameters while ensuring robust classification.

Finally, we explore the impact of the sample quantity, as depicted in Fig 5. The performance of all methods improved as the number of labeled samples increased. Among them, the fusion-based approaches show the most pronounced variation, ultimately achieving optimal performance with a larger quantity of samples. This suggests that our proposed methods can help researchers and engineers select appropriate models in the case of limited samples with low effort.

#### 5. CONCLUSION

In the paper, we propose a synergistic spectral and learning-based network (SSL-Net) to handle bird sound classification. Our approach with three diverse fusion strategies efficiently bridges the gap between generic audio and bird audio by leveraging spectral representations to address distribution bias in the pretrained model features. The results from the audio dataset of 20 bird species, collected in the real field, affirm both the effectiveness and efficiency of our method within the limitations of the labeled sample size.

## 6. REFERENCES

- [1] Petteri Lehtikainen, Meeri Rannisto, Ulisses Camargo, Aki Aintila, Patrik Lauha, Esko Piirainen, Panu Somervuo, and Otso Ovaskainen, “A successful crowdsourcing approach for bird sound classification,” *Citizen science*, vol. 8, no. 1, 2023.
- [2] Xeno-Canto Foundation, “Xeno-canto: Sharing bird sounds from around the world,” <https://xeno-canto.org/>.
- [3] Dan Stowell, Emmanouil Benetos, and Lisa F Gill, “On-bird sound recordings: automatic acoustic recognition of activities and contexts,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1193–1206, 2017.
- [4] Emily Vera Hockman, “Acoustic monitoring of wildlife in inaccessible areas and automatic detection of bird songs from continuous recordings,” 2018.
- [5] Xiao Wang, Guangyao Chen, Guangwu Qian, Pengcheng Gao, Xiao-Yong Wei, Yaowei Wang, Yonghong Tian, and Wen Gao, “Large-scale multi-modal pre-trained models: A comprehensive survey,” *Machine Intelligence Research*, pp. 1–36, 2023.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [7] Neil Zeghidour, Olivier Teboul, Félix de Chaumont Quiry, and Marco Tagliasacchi, “Leaf: A learnable frontend for audio classification,” *ICLR*, 2021.
- [8] Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, and Furu Wei, “Beats: Audio pre-training with acoustic tokenizers,” *arXiv preprint arXiv:2212.09058*, 2022.
- [9] Hanlin Wang, Yingfan Xu, Yan Yu, Yucheng Lin, and Jianghong Ran, “An efficient model for a vast number of bird species identification based on acoustic features,” *Animals*, vol. 12, no. 18, pp. 2434, 2022.
- [10] Jie Xie and Mingying Zhu, “Handcrafted features and late fusion with deep learning for bird sound classification,” *Ecological Informatics*, vol. 52, pp. 74–81, 2019.
- [11] Turker Tuncer, Erhan Akbal, and Sengul Dogan, “Multi-level ternary pattern and iterative relief based bird sound classification,” *Applied Acoustics*, vol. 176, pp. 107866, 2021.
- [12] Loris Nanni, Gianluca Maguolo, Sheryl Brahnem, and Michelangelo Paci, “An ensemble of convolutional neural networks for audio classification,” *Applied Sciences*, vol. 11, no. 13, pp. 5796, 2021.
- [13] Hasan Abdullah Jasim, Saadaldeen R Ahmed, Abdullahi Abdu Ibrahim, and Adil Deniz Duru, “Classify bird species audio by augment convolutional neural network,” in *2022 International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*. IEEE, 2022, pp. 1–6.
- [14] Aymen Saad, Javed Ahmed, and Ahmed Elaraby, “Classification of bird sound using high-and low-complexity convolutional neural networks,” *Traitement du Signal*, vol. 39, no. 1, pp. 187, 2022.
- [15] A Noumida and Rajeev Rajan, “Multi-label bird species classification from audio recordings using attention framework,” *Applied Acoustics*, vol. 197, pp. 108901, 2022.
- [16] Rajeev Rajan, Jisna Johnson, and Noumida Abdul Kareem, “Bird call classification using dnn-based acoustic modelling,” *Circuits, Systems, and Signal Processing*, pp. 1–12, 2022.
- [17] Jie Xie and Mingying Zhu, “Sliding-window based scale-frequency map for bird sound classification using 2d-and 3d-cnn,” *Expert Systems with Applications*, vol. 207, pp. 118054, 2022.
- [18] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D Plumbley, “Panns: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [19] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [20] Eric Jang, Shixiang Gu, and Ben Poole, “Categorical reparameterization with gumbel-softmax,” *arXiv preprint arXiv:1611.01144*, 2016.
- [21] Juan Gómez-Gómez, Ester Vidaña-Vila, and Xavier Sevillano, “Western mediterranean wetlands bird species classification: evaluating small-footprint deep learning approaches on a new annotated dataset,” *arXiv preprint arXiv:2207.05393*, 2022.