



HAL
open science

Self-Supervised Learning for Few-Shot Bird Sound Classification

Ilyass Moummad, Nicolas Farrugia, Romain Serizel

► **To cite this version:**

Ilyass Moummad, Nicolas Farrugia, Romain Serizel. Self-Supervised Learning for Few-Shot Bird Sound Classification. ICASSPW 2024 SASB, Apr 2024, Seoul, South Korea. hal-04925735

HAL Id: hal-04925735

<https://imt-atlantique.hal.science/hal-04925735v1>

Submitted on 2 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SELF-SUPERVISED LEARNING FOR FEW-SHOT BIRD SOUND CLASSIFICATION

Ilyass Moummad, Nicolas Farrugia

IMT Atlantique, Lab-STICC,
UMR CNRS 6285, Brest, France,

Romain Serizel

University of Lorraine, CNRS, Inria,
Loria, 54000, Nancy, France,

ABSTRACT

Self-supervised learning (SSL) in audio holds significant potential across various domains, particularly in situations where abundant, unlabeled data is readily available at no cost. This is pertinent in bioacoustics, where biologists routinely collect extensive sound datasets from the natural environment. In this study, we demonstrate that SSL is capable of acquiring meaningful representations of bird sounds from audio recordings without the need for annotations. Our experiments showcase that these learned representations exhibit the capacity to generalize to new bird species in few-shot learning (FSL) scenarios. Additionally, we show that selecting windows with high bird activation for self-supervised learning, using a pretrained audio neural network, significantly enhances the quality of the learned representations.¹

Index Terms— Self-supervised learning, data augmentation, few-shot learning, bird sounds.

1. INTRODUCTION

Bioacoustics, the study of sounds produced by living organisms, provides valuable insights into species diversity, behavior, and ecosystem health. Bird activity, in particular, can reflect changes in habitats and climate [1]. Deep learning, capable of automatically learning patterns from large datasets, emerges as a powerful tool for the automatic identification of bird species based on their unique acoustic signatures [2].

In the pursuit of leveraging technological advancements, BirdCLEF emerges as an annual competition dedicated to identifying endangered bird species through bioacoustics [3]. The data for BirdCLEF is sourced from the collection of bird sounds *Xeno-Canto* (<https://xeno-canto.org/>). BirdCLEF aims to enhance techniques in bioacoustics to tackle the challenges associated with the identification of bird species by providing labeled datasets of bird sounds. In this context, BirdNet [1] emerges as a model specifically trained on thousands of hours of *Xeno-Canto* bird data using supervised learning (SL). However, the requirement for abundant labeled training data in SL hinders progress, particularly for species

with limited documentation [4]. Nolasco et al. [4] address this problem by reformulating bioacoustic sound event detection as a collection of small tasks, each having few labeled data per animal sound, aligning with the FSL framework.

With the goal of advancing FSL for bioacoustic, Ghani et al. [5] show that embeddings derived from large-scale audio models, particularly BirdNet and Google Perch, trained on bird sounds using SL, exhibit strong performance in bioacoustic few-shot classification tasks. Considering the broader scope of FSL for audio classification, MetaAudio [6] proposes a benchmark of few-shot classification, covering a variety of domains including bird sounds. The authors split the training set of BirdCLEF 2020 [3], into three distinct sets of classes, representing new training, validation and test sets. The goal is to transfer knowledge from a set of classes, addressing the challenge of n -way k -shot sound classification.

This study explores SSL of bird sound representations on the MetaAudio split of BirdCLEF. Challenges arise in defining a pretext task to learn representations of bird sounds. We investigate three methods: SimCLR [7], Barlow Twins [8], and FroSSL [9]. Focusing on pretext-invariant representation learning, we exclude reconstruction methods due to their requirements for large datasets, and their demonstrated poor performance in low-shot classification [10].

Data augmentation plays a crucial role in the SSL of invariant representations [7]. The specific invariance that a model needs to learn for efficient representations in bioacoustics is not fully understood [2]. In this context, we operate under the assumption that domain-agnostic data augmentation, which does not assume any specific strong invariance on bird sounds, should prove effective for this task.

Beyond this aspect, in practical scenarios with lengthy audio files, particularly in bioacoustics, selecting training segments becomes crucial for the encoder’s quality, as choosing segments with only background noise or silence will be detrimental to the performance. Denton et al. [11] apply a wavelet peak detector to find frames with high energy. A recent foundation model in computer vision, DINOv2, selects training images using a pretrained model on a large uncurated dataset to retrieve similar images in a curated dataset [12]. Similarly, we use a pretrained audio neural network (PANN) on AudioSet [13], a dataset for general-purpose audio tagging [14], to select the segment with highest bird class activation.

¹Our code is released at:
<https://github.com/ilyassmoummad/ssl4birdsounds>
This work is co-funded by AI@IMT program and OSO-AI company

2. METHOD

In this section, we describe one SL and three SSL methods for pre-training representations of bird sounds. Subsequently, we describe the few-shot classification task for evaluating the learned representations.

2.1. Representation Learning

We describe three SSL methods: SimCLR [7] (sample contrastive), Barlow Twins [8] (dimension contrastive), and FroSSL [9] (both sample and dimension contrastive), along with SupCon [15] as a supervised method for reference. Let X be a batch of N examples, X_1 and X_2 represent the two resulting batches obtained by applying two transformations to each example in X . Let Z_1 and Z_2 be the two batches of embeddings extracted from X_1 and X_2 , respectively.

2.1.1. SimCLR

Let z_i and z_j be the ℓ^2 -normalized embeddings of two views x_i and x_j coming from X_1 and X_2 , respectively. The loss function for SimCLR [7] computed for two similar views (i, j) originating from the same original example as:

$$\mathcal{L}_{(i,j)}^{SimCLR} = -\log \frac{\exp(z_i \cdot z_j / \tau)}{\sum_{k \neq i} \exp(z_i \cdot z_k / \tau)}, \quad (1)$$

where $\tau \in \mathbb{R}^{+*}$ is a scalar temperature parameter and $z_k \in \text{concat}([Z_1, Z_2])$. The final loss SimCLR for the augmented batch is written as:

$$\mathcal{L}^{SimCLR} = \sum_{i,j} \frac{1}{2} (\mathcal{L}_{(i,j)}^{SimCLR} + \mathcal{L}_{(j,i)}^{SimCLR}), \quad (2)$$

SimCLR pulls in the embedding space similar views while pushing apart dissimilar ones.

2.1.2. Barlow Twins

Let $C \in \mathbb{R}^{N \times D}$ (with D the feature space dimension) be the cross-correlation matrix computed between Z_1 and Z_2 (assumed to be mean-centered along the batch dimension) computed as: $C = \frac{1}{N} Z_1^t \cdot Z_2$ where t is the transpose operator and \cdot is the matrix multiplication. The loss function for Barlow Twins (BT) [8] for the batch of embeddings Z_1 and Z_2 with cross-correlation matrix C is computed as:

$$\mathcal{L}^{BT} = \sum_i (1 - C_{ii})^2 + \lambda \sum_i \sum_{j \neq i} C_{ij}^2, \quad (3)$$

BT causes the embedding vectors of two transformed versions of an example to be similar, while minimizing the redundancy between the components of these vectors.

2.1.3. FroSSL

Z_1 and Z_2 are mean-centered along the batch dimension and normalized by their respective Frobenius norms. The Frobenius norm of a matrix $Z \in \mathbb{R}^{N \times D}$ is defined as: $\|Z\|_F^2 = \sum_{i=1}^N \sum_{j=1}^D Z_{ij}^2 = \sum_k^{\min(N,D)} \dots$. The loss function for FroSSL is defined as:

$$\mathcal{L}^{Fro} = MSE(Z_1, Z_2) + \lambda R(Z_1, Z_2), \quad (4)$$

where:

$$MSE(Z_1, Z_2) = \frac{1}{N} \sum_{i=1}^N \|Z_1^i - Z_2^i\|_2^2, \quad (5)$$

and:

$$R(Z_1, Z_2) = \log(\|Z_1^t \cdot Z_1\|_F^2) + \log(\|Z_2^t \cdot Z_2\|_F^2), \quad (6)$$

FroSSL minimizes the Euclidean distance between the embedding vectors of two transformed versions of an example, and the Frobenius norms of their respective covariances. Given the duality of Frobenius norm: $\|Z_i^t \cdot Z_i\|_F^2 = \|Z_i \cdot Z_i^t\|_F^2$, FroSSL can be considered both sample and dimension contrastive. We refer the reader to the original work for more theoretical details [9]. In the original formulation, the authors did not introduce the hyperparameter tradeoff λ between the invariance term MSE and the variance term R . Their intuition and experiments suggested that the optimal tradeoff occurred when the two terms were equally weighted. However, we found in our exploration that the coefficient of the variance term has to be small for FroSSL to converge.

2.1.4. SupCon

SupCon [15] is the supervised extension of SimCLR where the two views are two examples sharing the same label. Using the same notations as Eq 1, the loss of SupCon is defined as:

$$\mathcal{L}^{SupCon} = \sum_i \frac{-1}{|P(i)|} \sum_{j \in P(i)} \log \frac{\exp(z_i \cdot z_j / \tau)}{\sum_{k \neq i} \exp(z_i \cdot z_k / \tau)}, \quad (7)$$

where $P(i) = \{j : y_j = y_i\}$ is the set of indices of all views sharing the same label with i . $|P(i)|$ is its cardinality.

2.2. Few-Shot Learning

We consider FSL tasks sampled from the test split. A n -way k -shot few-shot task consists of a support set of k labeled examples from n classes to learn on, and a query set of test examples. Following our previous work [16], we adopt a simple inference method known as the nearest prototype classifier. This approach involves computing Euclidean distances to make predictions. Let \bar{z}_i be the prototype for each class label i from the support set, for each query from the query set of feature vector q , we predict its label i_q as:

$$i_q = \arg \min_i \|q - \bar{z}_i\|_2, \quad (8)$$

3. EXPERIMENTS

In this section, we describe the dataset, the strategy for selecting windows, data augmentations, and implementation details.

3.1. Data

In our study, we conduct experiments on BirdCLEF2020 (Pruned) dataset from MetaAudio [6]. This dataset is derived from the training set of BirdCLEF2020 by removing samples longer than 3 minutes and classes with less than 50 samples. It comprises 44,543/6,222/12,600 examples across 501/72/142 classes for the train, validation, and test sets, respectively.

Table 1. Data distribution in BirdCLEF 2020.

Dataset	N°Classes	N°Samples	Duration
BirdCLEF 2020	960	72,305	3s-30m
BirdCLEF 2020 (Pruned)	715	63,365	3s-3m

3.1.1. Data Selection

In our approach to handling audio signals of varying durations during training, we pad short signals to a fixed duration using a circular padding. For long audio signals, we employ either temporal proximity or PANN selection (PS). Temporal proximity randomly selects two nearby segments, encouraging the model to exploit temporal consistency [17]. PS involves dividing the audio signal into chunks of a fixed duration, passing them through a PANN on AudioSet, and selecting the chunk with the highest activation for the bird class. It’s worth mentioning that PANN undergoes supervised training to classify over 500 diverse classes, among which the bird class constitutes approximately 1% of the dataset.

3.1.2. Data Augmentation

We explore a variety of data augmentations drawn from the literature on SSL for audio [17, 18]. Through ablation experiments, not shown in this work, we observe that certain techniques, such as frequency shift, additive white Gaussian noise, compression, and random resized crop, are detrimental to learning representations for bird sounds.

Loris et al. [19] explore three data augmentation families for animal sounds within a supervised classification framework, comparing spectrogram augmentations to image and signal augmentations. While their experiments demonstrate the superiority of spectrogram augmentations, the absence of an ablation study assessing the individual impact of each technique raises uncertainties about their suitability for SSL. It is important to note that label information in supervised settings

may mitigate negative effects of inappropriate data augmentation, a condition not applicable to self-supervised invariant learning (Table 4). Therefore, we present a simple domain-agnostic data augmentation pipeline without specific hypothesis on spectrotemporal properties of bird sounds. It incorporates time shift [18] of spectrograms along the time axis to encourage the model to learn features invariant to temporal shifts. It also includes spectrogram mixing [18, 17], a process of random interpolation of spectrograms. This generates noisy views to encourage the model to learn features robust to various forms of noise. Finally, the pipeline includes the random masking of consecutive frequency bands and time steps using SpecAugment [20, 18], guiding the model to focus on capturing higher-level and more abstract representations.

3.2. Implementation Details

Following MetaAudio [6], we set the duration of training examples to 5 seconds. Signals are sampled at 16 kHz. We compute mel spectrogram features from FFT of size 1024, a hop length of 320, and a number of mels of 128. In order to generate two views, we use one of the two strategies: Temporal proximity, where we select two windows with a minimum of 60% overlap; or PS, where we choose a single window using CNN14, a PANN on Audioset. We then apply data augmentation to the resulting windows. Time shift is circular such that elements shifted beyond the last timestep are reintroduced at the first timestep. Spectrogram mixing linearly interpolates samples with random elements in the batch using a coefficient sampled uniformly between 0.6 and 1. SpecAugment masks three consecutive blocks of 30 frequency bands and 10 timesteps. The hyperparameters for data augmentation were tuned on the validation split.

For all our experiments, we train MobileNetV3-Large [21] from scratch using AdamW optimizer with a batch size of 256, a learning rate of 1e-3, and a weight decay of 1e-6 for 100 epochs. For BT and FroSSL, we employ a regularization coefficient of 1e-2. In the case of SimCLR and its supervised counterpart SupCon, a temperature of 1 is utilized. For inference, when PS is employed, the chosen segment with PANN serves as the prototype. Otherwise, we divide long test audio files into 5s chunks, average their feature vectors to get one prototype (for both query and test audio files). The features are then centered using the mean of the support set prototypes and are ℓ^2 -normalized before computing the Euclidean distance to make the predictions. We conduct 10,000 5-way 1-shot test tasks for each run, and all our experiments represent the average results obtained from three runs.

4. RESULTS AND DISCUSSION

We conduct three experiments. First, we train our models using temporal proximity and compare them with MetaAudio [6], as well as our few-shot inference method using the

pretrained feature extractor of CNN14, the same model used for PS. Second, we use PS exclusively for selecting training segments, and we also apply it for both selecting training and testing segments. Thirdly, we perform an ablation study for the augmentation techniques employed in our study.

Table 2 shows that SupCon outperforms the meta-learning approaches. In addition, all three SSL methods outperform CNN14 inference, with BT achieving the highest accuracy of 48.90%. This result underscores the feasibility of learning representations of bird sounds without labeled data. The lower performance of SimCLR and FroSSL in comparison to BT could be explained by the following: When training on long audio files with short events, sampling silence or background noise is possible. If two different samples share this trait in a training batch, the sample-contrastive losses separate them in the embedding space, potentially affecting the quality of the learned representation.

Table 2. Results of 5-way 1-shot tasks from BirdClef2020

Method	Accuracy
Supervised	
FO-MAML [6]	56.26±0.45
FO-Meta-Curvature [6]	61.34±0.46
ProtoNets [6]	56.11±0.46
SimpleShot CL2N [6]	57.66±0.43
Meta-Baseline [6]	57.28±0.41
SupCon (ours)	64.55±0.42
Self-Supervised (ours)	
SimCLR	43.28±0.43
FroSSL	38.24±0.39
BT	48.90±0.42
CNN14 (PANN)	32.97±0.35

Average accuracies along with 95% CIs using 10,000 test tasks

Table 3 illustrates the impact of PS as window selection strategy when applied during training, as well as during both training and test. PS enhances the results of all the methods we experimented with, underscoring the importance of an effective window selection. During inference, PS significantly improves the performance of the nearest prototype classifier by using the windows with the highest PANN activation of the bird class as prototypes, in contrast to averaging all the windows. This approach may be less conservative than averaging, and the expected performance improvement was not evident considering the results obtained with PANN alone (Table 2). The results highlight the complementarity of PS and SSL.

Table 4 presents the effects of ablating each data augmentation. The results underscore the crucial role played by the combination of both time shifting and spectrogram mixing. This underlines the intuitive idea that a model should be robust to background noise and temporal variation. Together, they contribute significantly to the model’s capacity to learn strong representations of bird sounds. It is important to note

Table 3. Effect of PANN selection of windows

Method	Accuracy
Supervised	
SupCon	64.55±0.42
+PS on Train	67.68±0.41
+PS on Train & Test	73.07±0.44
Self-Supervised	
SimCLR	43.28±0.43
+PS on Train	55.98±0.43
+PS on Train & Test	64.19±0.47
FroSSL	38.24±0.39
+PS on Train	46.23±0.40
+PS on Train & Test	56.79±0.49
BT	48.90±0.42
+PS on Train	53.82±0.42
+PS on Train & Test	63.61±0.47

Table 4. Effect of Data Augmentation

	SupCon	SimCLR	FroSSL	BT	RI
	73.07	64.19	56.79	63.61	24.53
-SA	71.20	57.85	50.14	57.73	
-SM	71.27	26.97	27.93	26.60	
-TS	72.25	58.15	27.40	27.48	

SA, SM, TS, RI stand for SpecAugment, Spectrogram Mixing, Time Shift, Random Initialization, respectively.

that when data augmentation is not enough to learn strong invariances, SimCLR performs the best among the three SSL methods. This may be due to the negative pushing force of its loss, which separates augmented examples coming from different original samples in the embedding space.

5. CONCLUSIONS AND PERSPECTIVES

In this work, we have presented a benchmark of three SSL approaches for few-shot bird sound classification. We have demonstrated that relying mainly on SSL can be competitive with SL. These results using systems relying on the efficient neural network MobileNetV3, provide a promising approach for bioacoustic applications. By employing only time shifting, spectrogram mixing and masking strategies without assuming specific invariances related to bird sounds, we established that meaningful representations can be learned, showcasing promising performance in few-shot learning scenarios. Additionally, leveraging a PANN as relevant segment selection for training and decision at test time significantly improves the performance of classification. In future work, we seek to investigate methods for assessing the quality of representations learned through SSL without relying on performance assessment on a validation set, which can substantially differ for the evaluation scenario as in our case.

6. REFERENCES

- [1] Stefan Kahl, Connor M Wood, Maximilian Eibl, and Holger Klinck, “BirdNET: A deep learning solution for avian diversity monitoring,” *Ecological Informatics*, 2021.
- [2] Dan Stowell, “Computational bioacoustics with deep learning: a review and roadmap,” *PeerJ*, 2022.
- [3] Stefan Kahl, Mary Clapp, W Alexander Hopping, Hervé Goëau, Hervé Glotin, Robert Planqué, Willem-Pier Vellinga, and Alexis Joly, “Overview of birdclef 2020: Bird sound recognition in complex acoustic environments,” in *Conference and Labs of the Evaluation Forum*, 2020.
- [4] Ines Nolasco, Shubhr Singh, Veronica Morfi, Vincent LOSTANLEN, Ariana Strandburg-Peshkin, Ester Vidaña-Vila, Lisa Gill, Hanna Pamuła, Helen Whitehead, Ivan Kiskin, et al., “Learning to detect an animal sound from five examples,” *Ecological Informatics*, 2023.
- [5] Burooj Ghani, Tom Denton, Stefan Kahl, and Holger Klinck, “Global birdsong embeddings enable superior transfer learning for bioacoustic classification,” *arXiv:2307.06292*, 2023.
- [6] Calum Heggan, Sam Budgett, Timothy Hospedales, and Mehrdad Yaghoobi, “Metaaudio: A few-shot audio classification benchmark,” in *International Conference on Artificial Neural Networks*, 2022.
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*, 2020.
- [8] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny, “Barlow twins: Self-supervised learning via redundancy reduction,” in *International Conference on Machine Learning*, 2021.
- [9] Oscar SKEAN, Aayush Dhakal, Nathan Jacobs, and Luis Gonzalo Sanchez Giraldo, “FroSSL: Frobenius Norm Minimization for Self-Supervised Learning,” *arXiv:2310.02903*, 2023.
- [10] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas, “Masked siamese networks for label-efficient learning,” in *European Conference on Computer Vision*, 2022.
- [11] Tom Denton, Scott Wisdom, and John R Hershey, “Improving bird classification with unsupervised sound separation,” in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022.
- [12] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al., “Dinov2: Learning robust visual features without supervision,” *arXiv:2304.07193*, 2023.
- [13] 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*, 2020.
- [14] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2017.
- [15] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan, “Supervised contrastive learning,” *Advances in NeurIPS*, 2020.
- [16] Ilyass Moummad, Romain Serizel, and Nicolas Farrugia, “Regularized Contrastive Pre-training for Few-shot Bioacoustic Sound Detection,” in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing*, 2024.
- [17] Aren Jansen, Manoj Plakal, Ratheet Pandya, Daniel PW Ellis, Shawn Hershey, Jiayang Liu, R Channing Moore, and Rif A Saurous, “Unsupervised learning of semantic audio representations,” in *IEEE international Conference on Acoustics, Speech and Signal Processing*, 2018.
- [18] Eduardo Fonseca, Diego Ortego, Kevin McGuinness, Noel E O’Connor, and Xavier Serra, “Unsupervised contrastive learning of sound event representations,” in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing*, 2021.
- [19] Loris Nanni, Gianluca Maguolo, and Michelangelo Paci, “Data augmentation approaches for improving animal audio classification,” *Ecological Informatics*, 2020.
- [20] Daniel S Park, Yu Zhang, Chung-Cheng Chiu, Youzheng Chen, Bo Li, William Chan, Quoc V Le, and Yonghui Wu, “SpecAugment on large scale datasets,” in *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020.
- [21] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al., “Searching for mobilenetv3,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019.