



OPEN Impact of transfer learning methods and dataset characteristics on generalization in birdsong classification

Burooj Ghani^{1✉}, Vincent J. Kalkman¹, Bob Planqué², Willem-Pier Vellinga², Lisa Gill^{3,5} & Dan Stowell^{1,4}

Animal sounds can be recognised automatically by machine learning, and this has an important role to play in biodiversity monitoring. Yet despite increasingly impressive capabilities, bioacoustic species classifiers still exhibit imbalanced performance across species and habitats, especially in complex soundscapes. In this study, we explore the effectiveness of transfer learning in large-scale bird sound classification across various conditions, including single- and multi-label scenarios, and across different model architectures such as CNNs and Transformers. Our experiments demonstrate that both finetuning and knowledge distillation yield strong performance, with cross-distillation proving particularly effective in improving in-domain performance on Xeno-canto data. However, when generalizing to soundscapes, shallow finetuning exhibits superior performance compared to knowledge distillation, highlighting its robustness and constrained nature. Our study further investigates how to use multi-species labels, in cases where these are present but incomplete. We advocate for more comprehensive labeling practices within the animal sound community, including annotating background species and providing temporal details, to enhance the training of robust bird sound classifiers. These findings provide insights into the optimal reuse of pretrained models for advancing automatic bioacoustic recognition.

Keywords Deep learning, Feature embeddings, Bioacoustics, Classification, Transfer learning, Knowledge distillation, Passive acoustic monitoring

Biodiversity is declining globally at an alarming rate, making it crucial for research to provide reliable estimates of the health of global ecosystems¹. While several modalities of data streams are in place to cater to this formidable challenge of monitoring biodiversity, sound-based monitoring methods coupled with machine learning provide a robust, fast, and scalable solution^{2,3}. Population monitoring depends on estimating species occupancy in specific locations to assess changes over time. Animal vocalizations serve as a good (though incomplete) non-invasive proxy for occupancy, as their presence in sound clips indicate species presence^{3,4}. The advent of deep learning in bioacoustics has sparked the development of advanced machine learning methods, enabling automated analysis of large acoustic datasets^{5,6}. Using these methods, researchers can automatically detect and categorize animal vocalizations, significantly saving time and effort while facilitating the study of lesser-known species⁷. For example, many lesser-known or cryptic species inhabit dense forests, remote areas, or aquatic environments, making direct observation or tracking challenging. Automated bioacoustic methods can identify their vocalizations within large soundscape datasets, even when these species are sparsely represented.

Birdsong classification has gained significant attention due to its applications in biodiversity monitoring and conservation, enabled by the increasing accuracy of birdsong species classifiers. Kahl et al.⁸ introduced BirdNET, a neural network model specifically designed for bird sound classification. Trained on a large dataset of bird sounds, BirdNET v2.2 can recognize over 3000 bird species. Stowell et al.⁹ reported the application of modern machine learning techniques to the task of acoustic bird detection, demonstrating that very high retrieval rates

¹Naturalis Biodiversity Center, Leiden, The Netherlands. ²Xeno-Canto Foundation, The Hague, The Netherlands.

³Landesbund für Vogel- und Naturschutz, Hilpoltstein, Germany. ⁴Department of Cognitive Science and Artificial Intelligence, Tilburg University, Tilburg, The Netherlands. ⁵BIOTOPIA Lab, Naturkundemuseum Bayern, Munich, Germany. ✉email: burooj.ghani@naturalis.nl

can be achieved. Their work highlighted the potential of deep learning in remote monitoring, overcoming the challenges posed by varying environments and species diversity.

However, the development of these tools typically depends on two key resources. The first is the availability of well-annotated training data, since the strong performance of deep learning typically requires much larger data volumes than prior methods. There is ample training data available for some taxonomic groups: for bird species, these are overwhelmingly sourced from community collections like Xeno-canto¹⁰ or the Macaulay Library¹¹. These collections, however, suffer from some limitations. Firstly, sufficient training data is often lacking for rare and endangered species, which are often the prime target of conservation efforts⁹. Secondly, the labelling (annotation) of collections may also be incomplete or approximate, lacking precise temporal details and omission of species audible in the background. Despite these limitations, such community audio collections are the key ingredient in all high-performing general-purpose birdsong classifiers.

The second key resource, increasing in importance, is the availability of machine learning models trained by others. *Transfer learning*—re-using the information encoded in a previously-trained machine learning algorithm—has been widely adopted in the field of audio classification due to its effectiveness in addressing the limitations of small and imbalanced datasets¹². Transfer learning is also increasingly crucial as the idea of *foundation models* rapidly evolves, offering large-scale, flexible pre-training adaptable to numerous downstream tasks. Transfer learning enables these foundation models, while scale amplifies their effectiveness¹³. Increasing the scale—whether in terms of the size of the model, the volume of training data, or the computational resources used—enhances the performance of foundation models. In bioacoustics, these models can learn broad, generalizable representations, useful for tasks like species detection, classification, and individual or call type identification, even extending to unknown taxa¹⁴.

The dominant approach to transfer learning is *finetuning*, which means additional iterations of training applied to the pre-trained neural network using a new dataset. Notably, models pre-trained on large audio datasets have shown significant improvements in performance when fine-tuned for specific tasks.

Researchers have successfully applied transfer learning to various audio tasks, such as speech recognition¹⁵, music genre classification¹⁶, and environmental sound event classification¹⁷. Gemmeke et al.¹⁸ introduced AudioSet, a large-scale dataset containing over 2 million human-annotated 10s sound clips drawn from YouTube videos. This dataset has served as a benchmark for training robust audio classification models. For instance, Hershey et al.¹⁹ explored various convolutional neural network (CNN) architectures pre-trained on AudioSet and demonstrated their efficacy in various audio recognition tasks. Gong et al.²⁰ proposed the PSLA model, a CNN-based approach trained on AudioSet. PSLA's architecture includes specialized layers for capturing both local and global features of audio signals, resulting in superior performance on various sound event classification benchmarks. Koutini et al.²¹ presented the PaSST model, which leverages an optimized transformer-based architecture for audio classification. Pre-trained on AudioSet, PaSST effectively captures long-term dependencies in audio signals, making it suitable for complex audio classification tasks.

Transfer learning has also been utilized in bioacoustics. Çoban, et al. employ the VGGish model, pre-trained on AudioSet, to classify ecoacoustic recordings from the North Slope of Alaska into eight different sound classes, including birdsongs and waterfowl sounds²². Ghani, et al. showed in¹² that embeddings derived from global bird sound classifiers can be effectively used for few-shot transfer learning across various taxonomic groups, including marine mammals, bats, and frogs. Lauha, et al. demonstrated that finetuning a bird sound recognition model on local data improved the model performance in the specific acoustic environment²³. In²⁴ Dufourq, Emmanuel, et al. demonstrated that transfer learning, specifically by adapting CNN architectures pre-trained on ImageNet, can effectively simplify the creation of bioacoustic classifiers for passive acoustic monitoring (PAM). The study highlighted that this approach reduces the complexity of CNN design and hyper-parameter tuning, making it more accessible and efficient even with very small datasets.

As an alternative to finetuning, the information could be transferred from the pretrained model to an entirely new model. *Knowledge distillation* is a technique where a new “student” model is trained to replicate the behavior of an existing “teacher” model^{25,26}. This approach has been employed to create small and efficient models that retain high performance while being more suitable for deployment in resource-constrained environments²⁷. This method has been widely adopted in various domains, including image and audio classification, to compress large models without significant loss in accuracy. Knowledge distillation has shown promising results in audio classification tasks. For instance, Choi et al.²⁸ applied knowledge distillation to develop a compact audio event classification model that achieved competitive performance with reduced computational requirements.

To train a sound recognition algorithm, the key resource is a large annotated dataset. However other datasets and models trained for related tasks are available and these could serve as a starting point allowing for building better models in a more efficient way. For a given machine learning goal, the question is thus how best to (re-) use these resources, since multiple strategies are available even within the domain of transfer learning. In this paper we explore this by empirically comparing strategies of finetuning against knowledge distillation. Our goal is to develop a highly accurate classifier for European bird species. While training from multiple diverse data sources is recommended for strong generalisation, training from a single common dataset is important for research study, benchmarking, and for pretraining. Xeno-canto is a uniquely valuable birdsong collection and is used as the core data to train all successful general-purpose birdsong classifiers. Its role in the development of computational bioacoustics has been analogous to the role of ImageNet in computer vision. In the present study we investigate how to train a model on Xeno-canto alone, the best to generalise to soundscape data. The outcome is a high-performing European bird species classifier, trained using fully open Xeno-canto data.

Method

The data processing, methodology, and evaluation workflow for this study are outlined in Fig. 1.

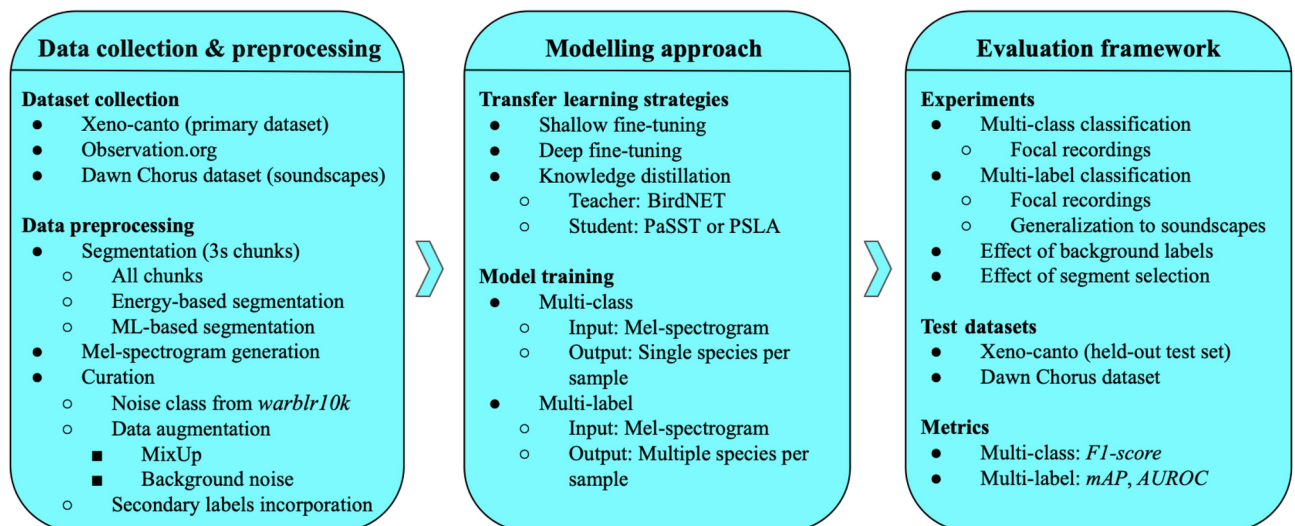


Fig. 1. Overview of the data processing, methodology, and evaluation pipeline.

Transfer learning

Transfer learning is a machine learning technique where knowledge gained from one domain is applied to another domain. “Knowledge” here is a shorthand term to express that a trained system has stored latent information about how to extract general or specific information from its input data, based on the training data’s characteristic patterns and distributions. The approach leverages a model developed for one task as a starting point for another, related task, often compensating for limited or inadequate data in the destination domain¹⁷. The primary idea is to utilize the patterns learned from a large dataset in the source domain and fine-tune them to improve performance on a different but related problem in the target domain. In the context of deep learning, transfer learning most commonly involves using a pre-trained network (trained on a large dataset such as ImageNet or AudioSet) and finetuning it on a smaller, task-specific dataset. Note however that there can be many ways to achieve transfer learning. For example, Ntalampiras (2018) extracted features from a (non-deep) classifier trained on music, as input to a bird classifier.²⁹

Birdsong classification is a challenging task due to the limited availability of labeled bird sound recordings and the complexity and variability of bird vocalizations. Transfer learning is particularly useful in this context for several reasons. First, data scarcity is a common issue, as birdsong datasets are often small and imbalanced. Transfer learning allows us to leverage pre-trained models that have been trained on large datasets, mitigating this problem. Second, these pre-trained models have learned to extract rich features that capture important patterns and characteristics in audio signals, which can significantly enhance the accuracy of birdsong classification. Additionally, starting with a pre-trained model reduces the time and computational resources needed to train a model from scratch, as the initial layers already contain useful representations. In this study, we explore different transfer learning strategies, including deep finetuning, shallow finetuning, and knowledge distillation, to capitalize on the advantages of using pre-trained models for birdsong classification. These strategies not only address data scarcity and improve feature richness but also streamline the training process, making transfer learning particularly beneficial for this task.

We next describe the transfer learning strategies considered in the present work.

Deep finetuning

Deep finetuning involves adjusting all layers of a pre-trained model on a specific dataset, in this case, the birdsong dataset. This comprehensive adjustment enables the model to better capture and adapt to the unique characteristics of bird sounds, adapting its “early” and “late” processing stages and thus potentially improving its performance on this specialized task. However, this approach has two potential drawbacks. Retraining the entire network may be computationally costly, especially as larger and larger networks are introduced. The method also risks overfitting to the new dataset, since there are many free parameters: if all layers are free to change, the pretrained knowledge may be fully erased by the later training, a phenomenon referred to as “catastrophic forgetting”^{30,31}. To mitigate this, deep finetuning needs careful implementation and/or a relatively large dataset.

The optimisation for deep finetuning is the same as for ordinary training of a deep neural network, in which we define a loss function L dependent on the input (spectrogram) X and the ground-truth label y :

$$L(X, y) = L_g(\delta(f(X)), y)$$

Here L_g is a standard loss function such as cross-entropy, and f representing the neural network. The final nonlinearity of the neural network is shown explicitly as the function δ , which we will later refer to, separately from the unnormalised outputs (“logits”) of f . The parameters of f (not shown) are chosen to minimise L .

In the context of deep learning, finetuning pre-trained models has been shown to yield significant improvements in various tasks. For instance, Howard and Ruder³² demonstrated that finetuning the entire model, as opposed to only the final layers, can result in substantial gains in performance on text classification tasks. Similarly, in the field of computer vision, deep finetuning of convolutional neural networks (CNNs) has been proven effective in adapting models to new domains, as evidenced by Yosinski et al.³³, who highlighted the benefits of finetuning across different tasks and datasets.

Moreover, recent advancements in finetuning techniques, such as those described by Raffel et al.³⁴ in their work on the T5 model, underscore the importance of appropriately scaling the model and dataset size to achieve optimal results. Applying these principles to the task of birdsong recognition, deep finetuning has the potential to harness the power of pre-trained models to enhance performance, provided that sufficient data and computational resources are available to support the process.

Shallow finetuning

In shallow finetuning, only the final layers of a pre-trained model are adjusted to fit on the target dataset, while the early layers remain unchanged. The unchanged part of the pre-trained model can be considered as a feature extractor, providing a rich set of learned features that should be useful for the new task. This method is less computationally intensive compared to deep finetuning, since only a small/shallow part of the network needs to be trained. It can be particularly effective when working with smaller datasets, as it is strongly constrained to leverage the general features previously learned by the lower layers of the model.

To express shallow finetuning mathematically, we separate the network f into two parts:

$$L(X, y) = L_g(\delta(f_2(f_1(X))), y)$$

where f_1 is constant, and the parameters of f_2 are optimised.

Shallow finetuning has been demonstrated to be effective in various contexts. White et al.³⁵ showed that adapting trained CNNs to new marine acoustic environments by only finetuning the top layers allowed the models to effectively generalize to different conditions without extensive retraining. This strategy not only saves computational resources but also reduces the risk of overfitting, making it suitable for tasks with limited data.

Additionally, Ghani et al.¹² highlighted the benefits of using global birdsong embeddings to enhance transfer learning for bioacoustic classification. By leveraging pre-trained embeddings and finetuning only the top layers, their approach achieved superior performance with minimal computational overhead. Similarly, Williams et al.³⁶ emphasized the advantages of using pre-trained models on diverse sound datasets, including tropical reef and bird sounds, to improve transfer learning in marine bioacoustics. Their work supports the efficacy of shallow finetuning in adapting models to new acoustic environments.

These studies underscore that shallow finetuning can effectively utilize the robust, general features captured by the lower layers of pre-trained models, making it a practical and efficient approach for tasks like birdsong recognition, especially when computational resources and dataset sizes are limited.

Knowledge distillation

Knowledge distillation is a technique where a pre-trained model is not itself changed; instead, it acts as a “teacher” and guides the training of a student model (Fig. 2, lower panel). In doing so the knowledge is transferred to the student model, to obtain a competitive performance versus the teacher model^{25,26}. This is often applied by training the student on “soft targets” from the teacher model: the fuzzy predictions rather than binarised 1-or-0 decisions. We interpret this as training the student to approximate the continuously-varying prediction function that the teacher outputs.

While model compression for efficiency and speed is a significant motivation behind this process, i.e. training smaller models from large²⁷, improved generalization and transfer learning are also key benefits. Learning from the nuanced information in the teacher’s soft targets helps the student model to learn the rich similarity information between classes that the teacher model has established^{26,37}. Interestingly, recent research has shown that cross-model distillation between different architectures, such as convolutional neural networks (CNNs) and transformer models, can yield superior performance, leveraging the strengths of both types of models to create a more robust and effective student model³⁸. This can be interpreted as one model learning to approximate information encoded not only in the data, but also in the biases implicit in the architecture of a very different network.

The original knowledge distillation framework was proposed by²⁶. Although diverse and more complex attention distillation techniques have since been studied³⁹, the authors in³⁸ reported that the original method is more effective, provided that certain details such as “consistent teacher” training are used⁴⁰. Therefore, we employ this approach. During the training of the student model, we use the same input with “consistent teaching”, meaning that the teacher and student are exposed to identical mel spectrograms, even when data augmentation is used to modify the training data. To train the student model the following loss is minimized, composed of two parts³⁸:

$$L(X, y) = \lambda L_g(\delta(f_s(X)), y) + (1 - \lambda) L_{kd}(\delta(f_s(X)), \delta(f_t(X)/\tau))$$

The coefficient λ represents the weighting whose value, following³⁸, we fix at 0.5. The terms L_g and L_{kd} correspond to the ground truth and distillation losses, respectively. In other words, the knowledge distillation can be seen as a weighted sum of ground truth loss L_g and distillation loss L_{kd} . The symbol δ denotes the activation function, while f_s and f_t are the student and teacher models, respectively. The parameter τ , referred to as the “temperature,” is set to $\tau = 1.0$ in alignment with³⁸, which showed that using this value yields optimal

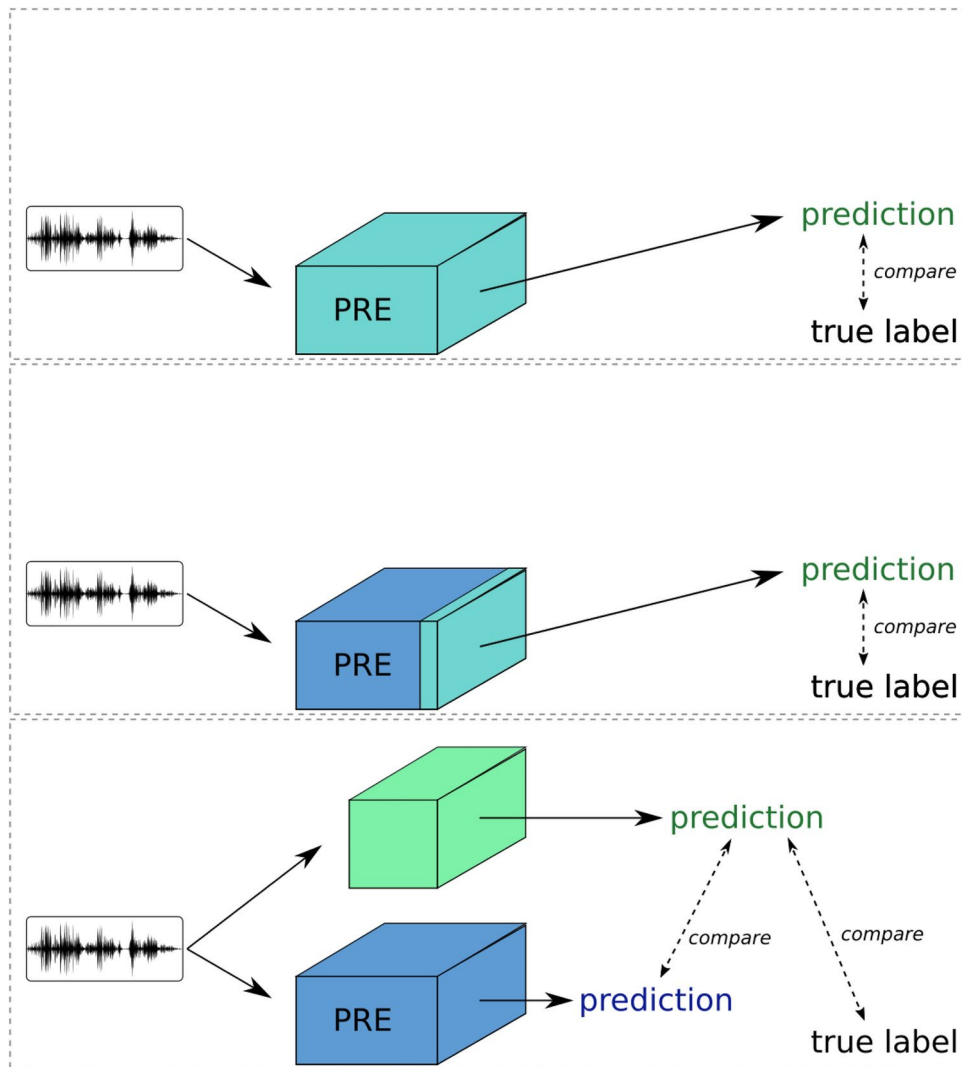


Fig. 2. Transfer learning strategies. Light-coloured blocks are neural networks being trained; dark-coloured blocks are ‘frozen’ and unchanging during transfer learning. Deep finetuning (top) retrain all layers. Shallow finetuning (middle) uses most of the pretrained model as a fixed feature extractor, retraining the final layer(s) on the new dataset. Knowledge distillation (lower) differs from both of these, training a new model to produce the same outputs as the teacher model, and also to match ground-truth labels when they are available.

results for knowledge distillation when a CNN serves as the teacher and a transformer acts as the student. It is important to note that the teacher model remains frozen during the training of the student model.

We employ the cross-entropy (CE) loss as L_g , combined with the softmax activation function, for multi-class classification (single-label AvesEcho v0). For the multi-label classification (AvesEcho v1), we utilize the binary cross-entropy (BCE) loss as L_g with the sigmoid activation function. Additionally, we use Kullback-Leibler divergence as the knowledge distillation loss L_{kd} for multi-class classification, and binary cross-entropy (BCE) loss for multi-label classification.

Audio data

To build a model allowing for the monitoring of European breeding birds we selected species in the 2nd European Breeding Bird Atlas (EBBA2)⁴¹ identified as breeding in Europe. This also includes all birds normally wintering or migrating in our target geography (running up to the Ural in the east). This list includes all birds found breeding in Europe in the period 2013–2017 including established non-native species.

We then downloaded corresponding audio data from Xeno-canto⁴², the citizen-science open data source for animal sound recordings. We used the Xeno-canto API to select sound recordings preferentially of high quality, and recorded in Europe. Xeno-canto (XC) is the largest open access database of nature sounds and available recordings cover all European breeding birds, although for some very rare species the number of recordings is limited. Xeno-canto has the benefit that the overall quality of recordings is good and the identification is reliable due to peer-review by other workers. A further advantage is that the data is largely open access so that future models can make use of the same source of training data.

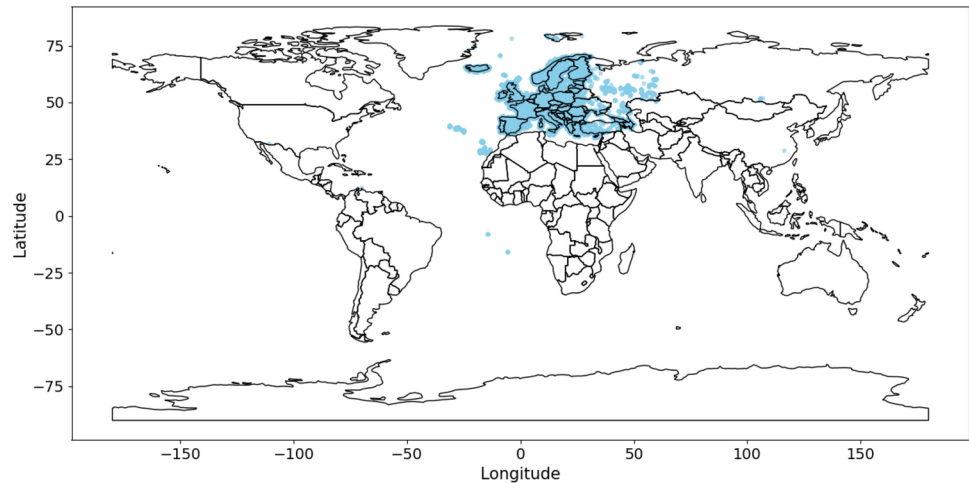


Fig. 3. Geographic distribution of our data sourced from Xeno-canto.

Dataset source	Multi-label	Classes	Sound Files	Total Duration (hours)
Xeno-canto*	Yes	585	127,161	1802
Xeno-canto*	No	438	120,807	1726
Observation.org	No	406	65,075	385
Dawn Chorus	Yes	50	236	4

Table 1. Summary of sound files and duration for the datasets used in this work. Datasets marked with an asterisk (*) have all classes derived from XC, except for the noise class. The noise class, representing non-bird sounds, is derived from the *warblrb10k* dataset, specifically from the subset labeled as “0” indicating “no birds present”.

In the following we report on two studies of neural networks, based on slightly differing datasets. First, *multi-class* classification, based on the assumption that only one foreground bird is active. To achieve a minimum requirement of at least 10 recordings per species in order to train the recognition algorithm, we curated a list of 438 bird species for the multi-class model (AvesEcho v0). Second, to train a *multi-label* model in which multiple species can be simultaneously active, we relaxed the selection criteria to include more species from the EBBA2 list, and also to include recordings from outside Europe for species where we did not have enough recordings. Additionally, we added recordings from outside Europe for all species with less than 50 recordings to make the total for a species up to 50. This curation resulted in a list of 585 species for the multi-label model (AvesEcho v1). Fig. 3 shows geographical distribution of the data sourced from Xeno-canto for the multi-label model.

For both of these datasets, a noise class is included, which is derived from the *warblrb10k* dataset⁴³, specifically from the subset labeled as “0” indicating “no birds present,” and so consisting just noise.

Additionally, we also curated a dataset of 406 species from Observation.org⁴⁴ to test our single-species models on recordings from a source other than Xeno-canto.

We also present results using a Dawn Chorus dataset developed by the Dawn Chorus project⁴⁵, a Citizen Science and Art initiative by the Naturkundemuseum Bayern/BIOTOPIA LAB and the LBV (The Bavarian Association for the Protection of Birds and Nature). The recordings were mostly short smartphone-based soundscapes collected during the dawn chorus - a time when many birds are most vocally active - by a worldwide community of enthusiasts. The dataset used in this study is a collection of 1-minute bird vocalisation recordings annotated by volunteers, including both professional and lay ornithologists. Only annotations from volunteers who scored 70% or higher on a species recognition questionnaire were included to ensure a high level of accuracy. The dataset comprises 236 annotated recordings from Germany, selected from a large worldwide pool of recordings⁴⁶ made available through the Dawn Chorus project. Annotations are conducted in 3-second chunks, following the BirdNET scheme, and are based on species identification through audio cues. The recordings in the dataset are primarily in uncompressed WAV format, although other formats like MP3 and FLAC are also present, particularly from earlier collection efforts. This subset represents a diverse range of bird species from Germany. Each recording in the dataset is linked to additional metadata, such as location and environmental conditions, though these were not utilized in the current analysis.

Table 1 provides a summary for different datasets used in this work and Fig. 4 shows a histogram illustrating the polyphony present in the Dawn Chorus dataset.

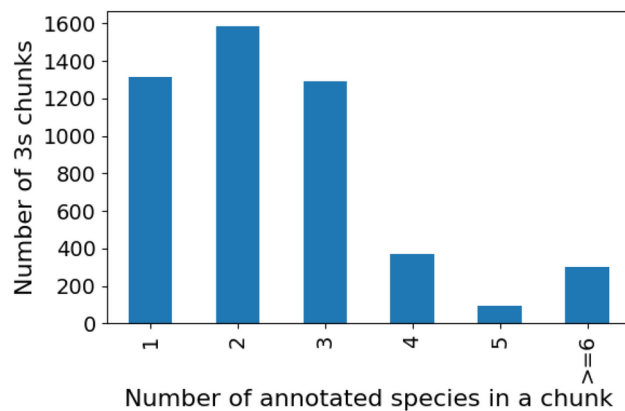


Fig. 4. Histogram depicting the number of labeled species within 3-second chunks from the Dawn Chorus dataset, illustrating the degree of polyphony present.

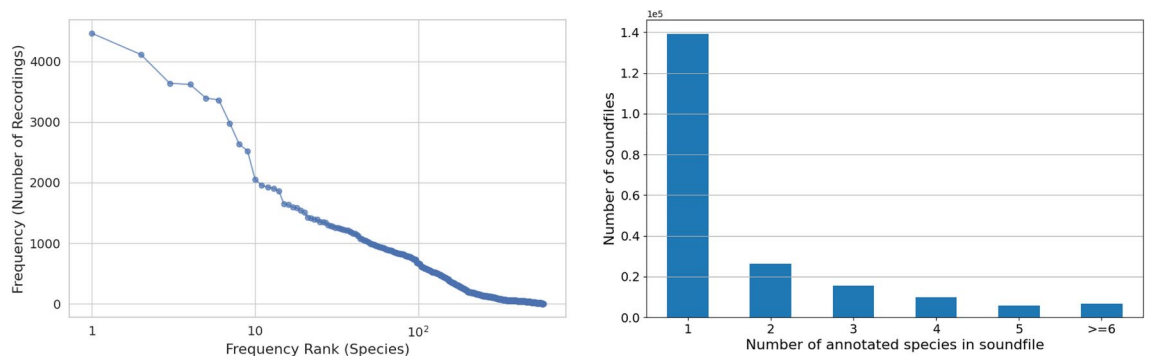


Fig. 5. (Left) Distribution of the number of recordings per species in the Xeno-Canto dataset for multi-label case. The plot shows that the distribution is unbalanced. (Right) Distribution of the total number of species annotated per file in our Xeno-canto data. The large peak at 1 indicates that many sound files are annotated as having no other species audible in the background.

Challenges with Xeno-canto data for model training

While Xeno-canto offers significant advantages, such as being the largest open-access database of nature sounds, covering all European breeding birds, and providing high-quality, peer-reviewed, and reliable identifications, there are several challenges when using this data for training machine learning models.

Firstly, the data is weakly labelled: while the species of the vocalizing bird is known, the exact timing of these vocalizations within the recordings is not, complicating segment-based model training. Secondly, many background species are missing in the annotations, potentially impairing the detection of background vocalizations. This arises from Xeno-canto's history, with its original design allowing only for one species per file. For machine learning this is a type of *label noise*: since many occurrences are in effect mislabelled as absences, it can bias a learning algorithm against making correct positive predictions. Thirdly, Xeno-Canto recordings are primarily focal recordings, designed to capture a target bird's sound as clearly as possible, making these vocalizations louder and more prominent than those typically found in soundscapes. This discrepancy poses difficulties in recognizing fainter bird sounds. Finally, the dataset is limited for certain species, such as *Stercorarius pomarinus* and *Puffinus mauretanicus*, which have very few recordings available. Fig. 5 illustrates the distribution of the number of recordings per species within the dataset curated to train the model. This distribution highlights the variability in the amount of available data across different bird species. This scarcity necessitates either the removal of these classes or the collection of additional data to ensure accurate classification.

Despite these limitations, Xeno-canto powers most general-purpose birdsong classifiers, and this yields better bioacoustic recognition than models trained on more generic audio data such as AudioSet¹². In the present study we aimed to evaluate how to make best use of the characteristics of Xeno-canto data, including the possible use of background labels.

Pretrained models

BirdNET⁸ utilizes EfficientNet, a well-established CNN architecture. It has been trained on an extensive dataset that includes recordings from Xeno-canto (XC), the Macaulay Library, and labeled soundscape data from various global locations, enabling it to recognize many thousands of bird species. Beyond avian identification, BirdNET is also capable of identifying human speech, dogs, and numerous frog species. This broad training set

allows BirdNET to be highly versatile, supporting a range of downstream applications. To achieve this versatility, BirdNET strikes a balance between accuracy and computational efficiency, ensuring it can be deployed on a variety of devices with limited processing power. In this work, we report on BirdNET version 2.2, which offers improved performance and broader species coverage. The BirdNET codebase, along with comprehensive documentation, is publicly available on GitHub⁴⁷, facilitating further research and development in bioacoustic monitoring and related fields.

PaSST²¹ leverages the strengths of transformer architectures for audio classification. The model, similar to the AST⁴⁸, is inspired by the Vision Transformer (ViT)⁴⁹. It extracts patches from the spectrogram, adds trainable positional encodings, and applies transformer layers to the flattened sequence of patches. Transformer layers utilize a global attention mechanism, which causes the computation and memory requirements to increase quadratically as the input sequence lengthens³⁷. PaSST attempts to mitigate this by introducing an innovative training technique known as “Patchout,” which selectively drops patches of the spectrogram during training, enhancing efficiency and robustness. In this way, the model significantly reduces compute complexity and memory requirements for training these transformers for audio processing tasks. The model is competitive with SOTA and obtains a mAP of 47.1 using their best single model on the AudioSet-2M. The implementation of PaSST, along with its training framework and performance benchmarks is available for public access on GitHub⁵⁰.

PSLA²⁰ is an EfficientNet model trained on AudioSet. We include this model to have a comparison with the BirdNET model which employs the same CNN-based architecture. By pretraining on large-scale audio datasets, PSLA captures a wide range of audio features, which are further refined through strategic sampling and labeling processes. To generate predictions, it uses an attention layer on the final embeddings, distinguishing it slightly from BirdNET, which uses a single dense layer for the same purpose. The model is again competitive with SOTA and obtains a mAP of 44.3 using their best single model on the AudioSet-2M. In this work, PSLA is incorporated into our pipeline as simply a model pre-trained on AudioSet, and we apply finetuning after our own pre-processing steps, rather than using the pre-processing pipeline followed in the PSLA paper. The PSLA framework, along with its implementation details and performance metrics is documented and accessible for public access on GitHub⁵¹.

Data preprocessing

For our experiments, we used two different configurations for generating mel-spectrograms, tailored to the requirements of the PaSST and PSLA models. For the PaSST configuration, we used 128 mel bands, a sampling rate of 16,000 Hz, a window length of 400 samples, a hop size of 160 samples, and an FFT size of 512. With a 3-second input signal, these settings result in a mel-spectrogram with 298 time frames. For the PSLA configuration, we used 128 mel bands, a sampling rate of 32,000 Hz, a window length of 800 samples, a hop size of 320 samples, and an FFT size of 1024. Similarly, a 3-second input signal with these settings also results in a mel-spectrogram with 298 time frames. Both configurations ensure that the mel-spectrograms have the same number of time frames, facilitating a consistent comparison between the models. Unlike PaSST and PSLA, BirdNET includes a pre-processing stem (a spectrogram layer). For BirdNET, the data is fed as an audio time series with a fixed input size of 3 seconds, using a sampling rate of 48,000 Hz.

Additionally, we also experimented with various data augmentation techniques to improve model robustness and performance. After evaluating several options, we found that two augmentations provided the best results: MixUp⁵² and background noise addition. Both of these augmentations are applied directly to the raw waveforms before spectrogram conversion.

MixUp is a data augmentation technique where two samples are combined to create a new synthetic sample. We apply MixUp with a probability of 0.6, meaning that 60% of the time, two random samples are mixed together to create a new training example. A value *lam* is sampled from a Beta distribution to linearly combine two input samples and their labels, ensuring smooth interpolation between them rather than a hard selection, with *lam* controlling the ratio of the mix. Since the process operates in floating-point format, there is no risk of clipping during training. Additionally, no explicit SNR is controlled for mixup, as the relative contributions of the two signals are determined by *lam*.

For background noise addition, we augment the input sample by mixing them with noise samples with a probability of 0.5. In this way, a random half of training samples are mixed with noise. Background noise is scaled dynamically to a random SNR sampled from [3 dB, 30 dB]. This ensures variability while maintaining the original signal's intelligibility. The augmentation is applied dynamically (“on-the-fly”) with a probability, and clipping is inherently avoided in floating-point processing. The noise samples are derived from the *warblrb10k* dataset, which contains a variety of background sounds, ensuring that the model becomes more resilient to environmental noise and can better generalize to real-world conditions.

The probability values for both MixUp and background noise were manually tuned in pilot testing. By incorporating these augmentations, we aim to enhance the models' ability to generalize, ultimately leading to more robust performance across different acoustic environments.

Experimental setup

When training a classifier, one must choose whether to train for single-label prediction, with the assumption that only one label is active, or for multi-label prediction. Single-label is appropriate for focal recordings in which there is one target of interest, which is the case for a large fraction of Xeno-canto recordings (Fig. 5), but multi-label is appropriate for soundscapes more generally, as multiple species can be simultaneously active. We investigated both configurations, using single-label as a first study and focusing on multi-label for our more detailed study.

For evaluation, we process the recordings in 3-second chunks and apply max pooling across all chunks within a recording. Instead of evaluating detections on each chunk individually, we aggregate the detections across the entire recording. This approach is meaningful given that our test set from Xeno-Canto is weakly labeled. By max pooling, we can confidently say that if a species is detected in any of the chunks, it is present in the recording as a whole. Secondly, by using max pooling, we are not merely aggregating detections from all chunks, rather, we are specifically reporting the highest detection score for each species across all chunks of the recording. This approach gives us a more robust and confident measure of the species' presence, as it focuses on the strongest evidence within the entire recording.

For single-label classification, this means we output one label with the highest score across the chunks, while for multi-label classification, we capture multiple labels, each with their highest scores across all chunks.

Single-label classification

To establish a clear starting point and maintain simplicity, we employed single-label classification (multi-class), assuming only one species is active at any given time. This assumption aligns with the majority of items in our training data. We explored various transfer learning paradigms, transferring from BirdNET, which is trained exclusively on bird sounds, as well as models trained on AudioSet, a dataset of general audio.

Each model (both CNN and transformer-based) takes a Mel-spectrogram as an input. Among the transfer learning techniques, we experimented with shallow finetuning (where only the last layer is trained while the pre-trained backbone remains frozen), deep finetuning (where all parts of the model are trained), and teacher-student model distillation (where training is transferred from a pre-trained model, the teacher, to another model, the student).

We have used the F1-score to report performance for single-label classification, as it provides a straightforward and balanced measure of precision and recall. A model's prediction for a given audio clip is selected as simply the one species that is most strongly predicted.

Multi-label classification

In many soundscapes, multiple bird species vocalize simultaneously. Consequently, unlike the previous single-label model, we trained our classifiers in a multi-label fashion. The training data typically include either a single species or one "foreground" species along with some labeled background species, prompting us to explore the best ways to utilize such annotations. Additionally, we expanded the number of classes for this model to 585. Pretraining is commonly used to enhance classifier performance. For birdsong audio, the question arises whether pretraining is necessary, and if so, whether it should be derived from birdsong or more general audio. Ghani et al. (2023)¹² demonstrated that pretrained features from BirdNET are both powerful and flexible. Therefore, we compared BirdNET pretraining (birdsong-based) with AudioSet pretraining (based on general audio from human soundscapes). These considerations and combinations were explored by training multiple variants of the machine learning algorithm, using the same standard training set and protocol.

Evaluating multi-label classification is slightly more involved than single-label. Evaluation can be based on thresholded predictions (i.e. the predicted species are all species for which the model output exceeds a fixed threshold value), or on the continuous-valued model outputs themselves. Thresholded metrics (such as accuracy, precision, and recall) lead to results strongly dependent on the chosen decision threshold. Selecting an optimal threshold is complex and varies based on the application, such as prioritizing recall for detecting rare species or precision for monitoring diverse species. To avoid the challenges of threshold selection, for our main evaluation we favour threshold-free metrics: the area under the ROC curve (AUROC) and average precision (AP; mAP when averaging the AP for all classes). These provide a more general evaluation by integrating performance across all possible thresholds, making them ideal for robust model comparison¹⁴. Besides, AUROC, also has a helpful probabilistic interpretation: it represents the likelihood that a randomly selected positive example is ranked above a negative one⁵³. For these reasons, we have chosen to use AUROC and mAP for our main evaluation, though we also use thresholded metrics in some targeted investigations of differences between models.

To assess the impact of data quantity on training a robust multi-label model, we analyzed the model's performance across different data regimes. We categorized the number of training sound samples per species into three arbitrary groups: low (<50 samples), medium (50-200 samples), and high (>200 samples).

Effect of secondary labels

We further designed some conditions to explore the impact of training with secondary ("background") labels from the Xeno-canto database. Given that many background species are missing in the annotations, this creates label noise that could bias the learning algorithm. We aimed to determine if incorporating these secondary labels can enhance the detection of background vocalizations, despite their potential for noise. Our experiments compared models trained with primary labels only against those trained with both primary and secondary labels to assess the overall effect on model performance.

Additional to using secondary labels in the *training* phase, we also investigated how our a trained model performs on the *test* set, contrasting the test set items where the target species is labelled as primary against those where it is labelled as secondary. For this analysis, we selected recordings from the test set where only two species were labeled: one as the primary (foreground; main focus of the recording) and the other as the secondary (background species). For each species in our selected subset, we examined the model's recall in two scenarios:

- Background Instances: Where the species occurs as a single background species in the recording in the test set as per the annotations.
- Foreground Instances: Where the species is the main focus of the recording, with just one other species labeled in the background.

This selection allowed us to constrain the experiment, ensuring that we could directly assess the impact of using secondary labels on the model’s performance for both foreground and background instances of the same species. Any differences in performance would then indicate differences in acoustic characteristics, for the same bird labelled in differing situation.

Weak labelling: effect of segment selection

Xeno-canto is a weakly labeled database, where the species of the vocalizing bird is known, but the exact timing of these vocalizations within the recordings is not. This complicates segment-based model training. In this section, we examine the impact of using a detector to discard segments that are likely devoid of bird sounds. By filtering out these no-bird segments, we aim to assess whether this preprocessing step improves the training performance of our models. Our experiments compared model performance with and without the use of the detector to determine its effectiveness in enhancing detection performance.

To extract 3-second segments from our Xeno-Canto dataset using a detector, we employed two methods: energy-based and ML-based.

The energy-based method is adapted from a previously established approach⁵⁴. This method uses a energy-based heuristic to identify and select segments likely to contain bird vocalizations while discarding those with non-bird sounds. For further details on the extraction process, please refer to Appendix B in⁵⁴.

In the ML-based method we employed a pre-trained deep learning model PANNs⁵⁵ which is trained using AudioSet. We selected the chunks for which the score for the *bird* class crossed 0.3.

Results

Single-label classification

We trained four different single-label classifiers using the set of transfer learning techniques, then evaluated them using held-out Xeno-canto data (Table 2).

The best-performing model in our experiments was the PaSST, where knowledge was distilled from BirdNET, achieving an F1-score of 0.704. This suggests that leveraging the domain-specific knowledge embedded in BirdNET through distillation can effectively enhance the performance of a transformer-based architecture like PaSST. The second model involved deep finetuning of all layers of the PaSST model pretrained on AudioSet, which yielded an F1-score of 0.679. Although this model did not outperform the knowledge-distilled PaSST, it still demonstrated robust performance. The third model, which involved training a dense layer on top of the BirdNET backbone, resulted in an F1-score of 0.676. Notably, this score is nearly identical to that of the PaSST model pretrained on AudioSet and deeply finetuned. This suggests that BirdNET, with its architecture and pretraining specifically on bird vocalizations, remains highly competitive even with minimal additional training (shallow finetuning). The fourth model, a PSLA architecture pretrained on AudioSet and deeply finetuned, achieved an F1-score of 0.632. This was the lowest score among the models tested, indicating a significant drop in performance compared to the other models.

The performance of our models was also evaluated on a separate dataset extracted from *Observation.org*, which, similar to Xeno-Canto, consists of focal recordings. This allowed us to test the models in a multi-class scenario. As shown in Table 2, the highest F1-score was achieved using the shallow finetuned model where a dense layer was trained on top of the BirdNET backbone, scoring 0.52. The knowledge-distilled transformer model (PaSST with BirdNET pretraining) performed second best with an F1-score of 0.49, followed by the deeply finetuned PaSST model pretrained on AudioSet with a score of 0.45. The PSLA model, also deeply finetuned on AudioSet, had the lowest performance, with an F1-score of 0.42. These results further emphasize the effectiveness of shallow finetuning on bird-specific pretraining, particularly for multi-class bird species classification tasks.

Multi-label classification

We trained multi-label classifiers using the set of transfer learning techniques, again evaluating them on a held-out Xeno-canto test set (Table in Fig. 6).

Similarly as the single-label case, the PaSST model with knowledge distillation from BirdNET achieved the best performance, with a mAP of 0.71 and an AUROC of 0.95, requiring 10 epochs and just over 23 hours of total training time. The BirdNET model with shallow finetuning, while having a slightly lower mAP (0.68), matched the AUROC (0.95) and completed training significantly faster, in less than 1 hour. Other models, such as the PSLA and PaSST models pretrained on AudioSet, showed varying performance, with the PSLA model using shallow finetuning on AudioSet scoring the lowest with a mAP of 0.03. Overall, the results highlight the

Model	Pretraining source	Transfer method	F1-score	
			Xeno-canto	Observation.org
PaSST	BirdNET	Knowledge distillation	0.704	0.49
PaSST	AudioSet	Deep finetuning	0.679	0.45
BirdNET	BirdNET	Shallow finetuning	0.676	0.52
PSLA	AudioSet	Deep finetuning	0.632	0.42

Table 2. Performance of models on held-out test dataset from Xeno-canto and the dataset derived from Observation.org for single-label classification. The table compares F1-score across different models and transfer learning methods. Entries are bold-faced if the model scored the highest for the respective dataset.

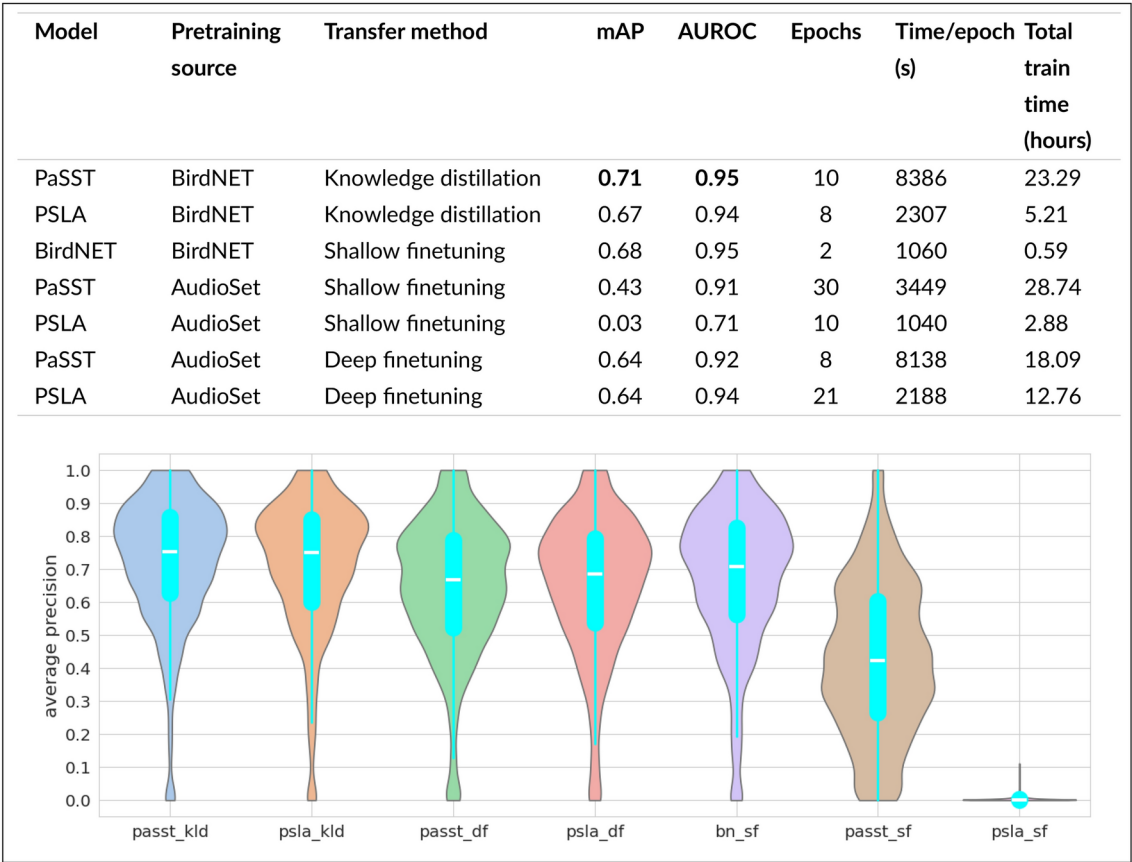


Fig. 6. Performance of different models evaluated on the held-out test dataset from Xeno-canto. **(a)** Table showing various metrics including mAP, AUROC, epochs, time per epoch, and total training time. **(b)** Species-wise distribution of scores using the approaches listed in **(a)**. Entries are bold-faced if the model scored the highest.

Model	Pretraining source	Transfer method	mAP	AUROC
BirdNET	BirdNET	Shallow finetuning	0.311	0.836
PaSST	BirdNET	Knowledge distillation	0.293	0.792
PSLA	BirdNET	Knowledge distillation	0.273	0.793
PaSST	AudioSet	Deep finetuning	0.291	0.794
PSLA	AudioSet	Deep finetuning	0.256	0.768

Table 3. Comparison of models on the Dawn Chorus dataset. Entries are bold-faced if the model scored the highest.

effectiveness of knowledge distillation from BirdNET, particularly with the PaSST model, in achieving high multi-label classification performance. The inability of the shallow fine-tuned PSLA to train might suggest that the embeddings’ space in PSLA is heavily influenced by the attention pooling layer, perhaps making it intrinsically tied to the specific dataset used during training.

Inspecting the species-wise distribution of scores using different approaches, we found that the distribution of performance per-species was generally consistent between models, and coherent with the overall performance scores (Fig. 6).

We then evaluated our multi-label models on the dawn chorus dataset, to evaluate how well they generalise to new conditions and polyphonic multi-species soundscapes (Table 3). The BirdNET model with shallow finetuning achieved the highest mAP (0.311) and AUROC (0.836), outperforming the other models. The PaSST model, which had knowledge distilled from BirdNET, performed slightly lower with a mAP of 0.293 and an AUROC of 0.792. Similarly, the PSLA model using knowledge distillation from BirdNET showed comparable results with a mAP of 0.273 and AUROC of 0.793. Models pre-trained on AudioSet, specifically PaSST and PSLA with deep finetuning, scored lower, with the PSLA model obtaining the lowest mAP (0.256) and AUROC (0.768).

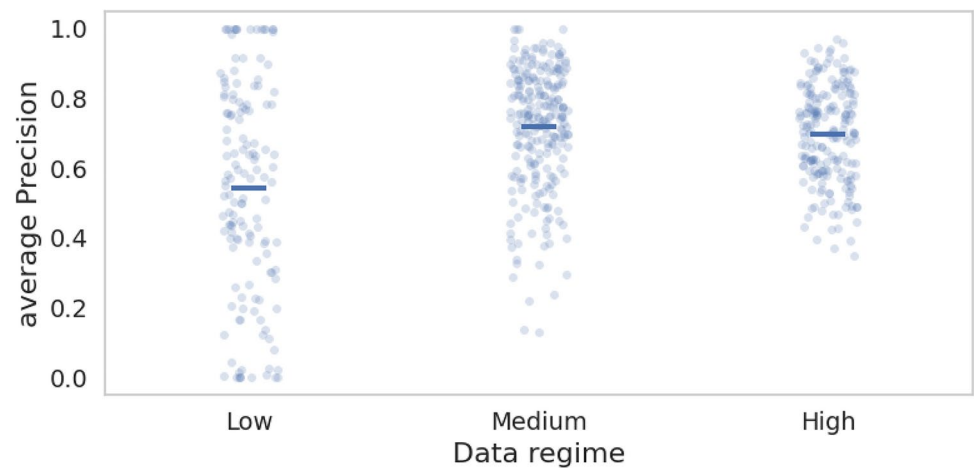


Fig. 7. Distribution of species-specific average precision scores using the shallow fine-tuned BirdNET model, grouped by data availability for each species (low, medium, and high data regimes based on the amount of available data in the XC-derived dataset).

Dataset	Train using secondary labels	Precision	Recall	mAP	AUROC
Xeno-Canto	Yes	0.44	0.70	0.68	0.96
	No	0.59	0.61	0.68	0.95
Dawn chorus	Yes	0.74	0.45	0.31	0.84
	No	0.88	0.14	0.31	0.81

Table 4. Comparison of training with and without using secondary labels on the XC-derived test set and the dawn chorus datasets. The reported performance is using the shallow finetuned BirdNET model. To compute the precision and recall, a common threshold of 0.2 was set.

In order to understand the importance of data quantity in training a robust multi-label model, we performed an analysis of the model’s performance across different data regimes (Fig. 7). The results reveal a clear dependency on the amount of training data available for each species. Species with a high number of recordings consistently achieved higher average precision scores, indicating strong model performance in this regime. In contrast, species in the low data regime, with as few as 10 recordings, showed significantly lower average precision and greater variability, highlighting the model’s difficulty in accurately predicting species with limited data. The medium data regime displayed intermediate results, with moderate average precision and some variability, further underscoring the importance of data quantity in achieving reliable model performance. These findings emphasize the critical role of sufficient training data in enhancing model performance for multi-label bird species classification.

Effect of secondary labels

In this experiment, we explored the impact of incorporating secondary labels while training on the performance of our bird species recognition model. Our aim was to determine whether these additional labels, despite the label noise, could enhance the model’s ability to detect background vocalizations. The results of this experiment are summarized in Table 4. The overall mAP remained essentially unchanged (0.68 vs. 0.68 for XC, 0.31 vs. 0.31 for dawn chorus), indicating that the introduction of secondary labels did not significantly affect the model’s ability to rank predictions accurately across the datasets.

However, we observed a notable decrease in precision for both the XC-derived test set (0.44 with secondary labels vs. 0.59 without) and the dawn chorus soundscapes dataset (0.74 with secondary labels vs. 0.88 without) when using secondary labels. This suggests that the model became more prone to predicting the presence of species even in cases where they were not present, likely due to the noise and incompleteness of these labels.

Conversely, recall increased in both test scenarios (0.70 with secondary labels vs. 0.61 without for XC, and 0.45 with secondary labels vs. 0.14 without for dawn chorus). This indicates that the model became more sensitive to detecting species, successfully identifying more true positives, including those that may have been underrepresented or missing from the primary labels.

The AUROC score, on the other hand, increased when incorporating secondary labels (0.96 with secondary labels vs. 0.95 without for XC, and 0.84 with secondary labels vs. 0.81 without for dawn chorus). This suggests an overall improvement in the model’s ability to distinguish between species presence and absence across various decision thresholds, indicating enhanced discrimination capability despite the label noise.

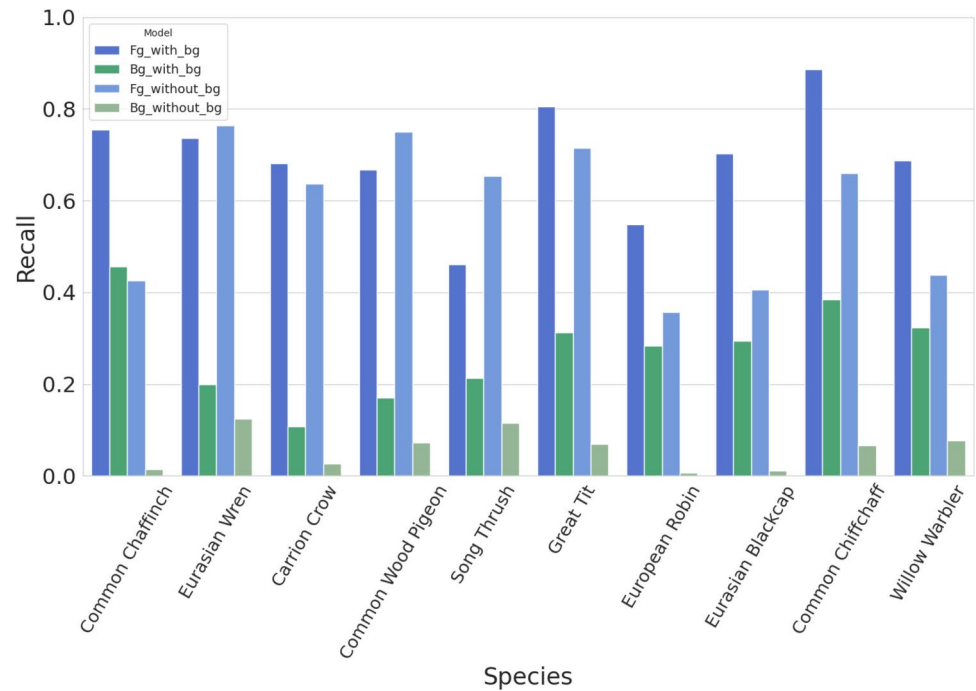


Fig. 8. Recall of foreground and background predictions for common bird species in our XC-derived test set, where exactly two species are labeled as present: one in the foreground and one in background. For each species, recall is measured when instances are labeled as either foreground or background. Darker shades of blue and green represent predictions from a model trained using both foreground and background labels, while lighter shades of blue and green represent predictions from a model trained using only foreground labels. In both cases, the model is a shallow fine-tuned BirdNET.

Dataset	Detection	mAP	AUROC
Xeno-Canto	Peak detector	0.70	0.95
	PANNs detector	0.67	0.95
	No	0.68	0.95
Dawn chorus	Peak detector	0.32	0.82
	PANNs detector	0.30	0.82
	No	0.31	0.81

Table 5. Comparison of training with and without using detection algorithms on the XC-derived test set and the dawn chorus datasets. The reported performance is using the shallow finetuned BirdNET model.

To investigate this balance in more detail, we selected 10 common species for analysis, then analysed model predictions for all Xeno-canto sound recordings that contain one of these 10 species as primary or secondary, restricting our attention to the sound recordings containing exactly two species in total. This scenario enables us to probe how predictive performance varies whether a sound is primary or secondary (implying that it may be foreground or background). For all 10 species, the model’s accuracy in identifying instances where the species was labeled as secondary improved when the model was trained using secondary labels (Fig. 8). This suggests that the inclusion of secondary labels during training helped the model better recognize species even when they are not the primary focus of a recording. The results for foreground instances were more mixed. For some species, the accuracy increased when secondary labels were included in the training phase, indicating that the model became better at recognizing these species even when they were in focus. However, for other species, the accuracy of foreground instances decreased. This mixed result suggests that, while the model’s overall sensitivity improved, the specific characteristics of certain species when they are the focus of a recording might have been less distinguishable when trained with noisy or incomplete secondary labels.

Weak labelling: effect of segment selection

We compared the performance of models trained with and without the use of the detector, evaluating their performance on both the XC-derived test set and the dawn chorus datasets. The results, as shown in Table 5 indicate that the effect of segment selection using a detector is at best marginal. For the XC-derived test set, training with the energy-based detector yielded a slight improvement in mAP (0.70 vs. 0.68), while the AUROC

remained unchanged at 0.95. Similarly, for the dawn chorus dataset, using the energy-based detector led to a small increase in both mAP (0.32 vs. 0.31) and AUROC (0.82 vs. 0.81). Surprisingly, we did not notice any noticeable gain when employing the PANNs detector for selecting chunks.

These results suggest that while a detector does contribute to a slight improvement in model performance, the overall effect is small to negligible. This could indicate that the models are robust enough to handle the weak labeling inherent in the Xeno-Canto dataset. Thus, while segment selection might offer a modest boost in performance, its impact may be limited depending on the nature of the dataset and the model used.

Discussion

Transfer learning

We have shown that transfer learning leads to a consistent strong performance in birdsong classification, across a variety of conditions: both single- and multi-label, both CNN and Transformer models. This is the case for the common finetuning method, but also for knowledge distillation. In particular, we investigated cross-model distillation, which is knowledge distillation from one type of deep learning architecture to a different one. We find that cross-model distillation can yield classifiers that give better performance than those trained with finetuning, when evaluated in-domain on Xeno-canto data (Table 2, Fig. 6).

However, contrary to indications from related literature on cross-distillation³⁸, we find that the generalisation to novel soundscape datasets, while strong, does not outperform shallow finetuning (Tables 2, 3). It thus appears that shallow finetuning has better generalisation capabilities than knowledge distillation. This is likely caused by the fact that shallow finetuning is the more constrained form of training, having fewer degrees of freedom.

Despite this concern, knowledge distillation brings the significant advantage that it liberates us from the design decisions of previous neural network architectures, enabling automatic bioacoustic recognition to keep pace with developments in the general field of machine learning. For example, future researchers will need to consider not only CNNs and Transformers, but also alternatives such as state-space models (Mamba) or graph neural networks^{56,57}.

As we have demonstrated, it is feasible to train a bird classifier made from a high-performing modern Transformer architecture (PaSST) while reusing the knowledge contained in a completely different CNN-type architecture (BirdNET). Knowledge distillation takes more computational effort than shallow finetuning—as expected, since the latter trains only a small part of a network—yet this is not the dominant factor for computational load in our studies. Rather, the choice of neural network architecture has the strongest influence on the total training time required. Thus, future work should investigate how to make use of knowledge distillation for flexible knowledge transfer, while also constraining the algorithm to enhance generalisation.

Our analysis further reveals that while bird-specific pretraining generally leads to superior performance, models pretrained on broader audio datasets like AudioSet also achieve good results. For example, the BirdNET model with shallow finetuning consistently outperforms others, particularly in challenging environments such as the dawn chorus, which is known for its high polyphony. Despite this, models like PaSST pretrained on AudioSet demonstrate that general audio pretraining can still produce competitive results, especially in scenarios where specialized bird sound datasets may not be readily available. This suggests that transfer learning is a flexible method, in which even general audio models can adapt effectively to specific bioacoustic tasks.

While transfer learning is advantageous for developing models in low data regimes, our results reveal that the quantity of data still plays a role in model performance. As shown in Figure 7, species with ample training data achieve strong average precision scores, whereas those with limited data show greater variability—some perform well, but others poorly. This suggests that while transfer learning helps mitigate data scarcity, other factors, such as vocalization complexity, recording quality, background noise, and species-specific variability, might also impact performance. These findings highlight the need to consider data quality and species-specific characteristics alongside transfer learning to achieve reliable model performance, particularly in low-data scenarios.

Use of single-species and polyphonic training data

Xeno-canto is the most important data source for training birdsong classifiers, yet it contains a mixture of single-species and multi-species recordings, the latter often incompletely labelled. Thus we investigated how best to use this data for recognition in polyphonic birdsong soundscapes. Our results suggest that while the introduction of secondary labels increases the model's sensitivity (as seen in higher recall), it also introduces a higher risk of false positives (reflected in lower precision). The stability of mAP implies that, although the model's confidence may have decreased, its overall ranking of species likelihoods across the dataset did not suffer. This trade-off between precision and recall must be carefully considered when designing bird species recognition models, especially in scenarios where the completeness and accuracy of labels are in question.

Model predictions on Xeno-canto data are more confident and more accurate for species that have been labelled as primary (foreground, or focal) species, versus secondary species (Fig. 8). This tells us that there is some difference in the acoustic data between these two cases. So what is this difference? Although the (meta) data do not answer this question explicitly, community practice in Xeno-canto is to label as primary the most salient species heard, or the species that the recordist was targeting. Secondary species might be overlapping in time with the primary, which can lead to acoustic masking for a listener, or might be separate in time; the data doubtlessly contain examples of each case.

The more general differences are likely to be the interrelated factors of loudness and distance. Loudness can be probed via the relative amplitude of recordings, though we remind the reader that the data in question are not provided with calibrated amplitudes. Distance affects loudness as well as other characteristics such as reverberation and the atmospheric absorption of high frequencies. The task of species classification is, formally, not dependent on loudness: sounds should be labelled as present irrespective of their (relative) amplitudes, and

machine learning algorithms should then learn to perform the task invariant to amplitudes. However, even with deep learning systems the strength or confidence in detecting a signal is related to the signal-to-noise ratio and the distinctiveness of the signal—though, as for human listeners, this relationship is nonlinear. Further, quieter sounds have a lower probability of having been positively labelled in the training data. Thus in many audio recognition tasks it is typical that the louder a sound, the more confidently an automatic classifier can pick up on it.

To produce discrete predictions, classifier outputs are thresholded in practice. Calibrating thresholds to local data conditions has been identified to be crucial, if automatic recognition is to produce reliable biodiversity observations^{58,59}. In this work we do not attempt to address this calibration issue, but the issue is nonetheless evident in the fact that our results on generalisation to novel soundscapes have relatively low mAP (evaluated on thresholded detections) even when they have strong AUC scores. It remains an open question how far we can automate or avoid such calibration.

Our study finds positive results when using the labels of secondary (“background”) species in Xeno-canto sound recordings, despite the fact that such labels may be incomplete. Therefore it would be valuable for the Xeno-canto community, and those who wish to make use of its data, to invest effort into labelling the sound recordings more completely. Our observations demonstrate the value of labelling all species audible in the background; it would also be useful to annotate temporal details (start time and end time of vocalisations), which would help to overcome the issue of weak labelling in the training of classifiers.

Towards the future of biodiversity monitoring frameworks

We have highlighted the potential of transfer learning to address challenges in bioacoustic monitoring. A key direction for future work lies in the development of foundation models or models capable of learning robust embedding spaces. These models provide high-quality, generalizable features, which can be leveraged for few-shot learning. This approach is particularly promising for addressing the underrepresentation of rare or lesser-studied species in bioacoustic datasets, as it allows for effective classification with minimal labeled data.

An important next step is the integration of human-in-the-loop annotation systems into bioacoustic pipelines. Such systems, which connect to experts with limited time, enable iterative improvements to model performance through active learning. Such agile modeling paradigms can use a combination of a few labeled examples and hard cases - instances where models struggle - to iteratively refine classifiers. This approach not only enhances model accuracy but also makes efficient use of scarce expert annotations.

These frameworks have the potential to significantly improve weakly labeled datasets like Xeno-canto, where annotations are currently provided only at the recording level. By incorporating active learning and expert-driven feedback, these systems could facilitate the addition of detailed temporal annotations, allowing for the identification and localization of specific vocalizations. In this way such advancements could lay the groundwork for scalable biodiversity monitoring frameworks. These frameworks that combine transfer learning, foundation models, and human-in-the-loop active learning can allow for adaptable pipelines capable of handling diverse data and addressing complex ecological questions. By leveraging these technologies, we could move toward large-scale, high-resolution biodiversity monitoring systems that are accessible to both researchers and citizen scientists, opening new avenues for bioacoustic research and conservation efforts.

Conclusion

This study demonstrates the effectiveness of transfer learning in advancing birdsong classification, highlighting their potential for addressing challenges in bioacoustic monitoring. While shallow finetuning excels in generalization to novel soundscapes, knowledge distillation offers flexibility for leveraging diverse architectures. Our findings emphasize the importance of data quality, species-specific characteristics, and the integration of background species' labels to improve model performance. We report that incorporating background species' labels from polyphonic recordings increases sensitivity, though at the cost of reduced precision, highlighting the trade-offs in leveraging weakly labeled datasets. Looking ahead, combining transfer learning with foundation models and human-in-the-loop systems presents a promising path toward scalable, adaptable biodiversity monitoring frameworks. These advancements hold significant potential for empowering conservation research and addressing ecological challenges at scale.

Data availability

The Xeno-canto datasets were downloaded using the XC API <https://xeno-canto.org/explore/api>. The Observation.org data was downloaded using GBIF generating the following DOI - <https://doi.org/10.15468/dl.3cnds7>. The Dawn Chorus dataset was graciously provided by Lisa Gill, as detailed in the dataset section, and might be available upon reasonable request and with permission of the providers.

Received: 23 September 2024; Accepted: 2 May 2025

Published online: 09 May 2025

References

1. Raven, P. & Wackernagel, M. Maintaining biodiversity will define our long-term success. *Plant Diversity*. **42**(4), 211–220 (2020).
2. Rutz, C. et al. Using machine learning to decode animal communication. *Science*. **381**(6654), 152–155 (2023).
3. Rasmussen, J. H., Stowell, D. & Briefer, E. F. Sound evidence for biodiversity monitoring. *Science*. **385**(6705), 138–140 (2024).
4. Wood, C. M. et al. Detecting small changes in populations at landscape scales: a bioacoustic site-occupancy framework. *Ecological Indicators*. **98**, 492–507 (2019).
5. Stowell, D. Computational bioacoustics with deep learning: a review and roadmap. *PeerJ*. **10**, e13152 (2022).
6. Bick, A. et al. National-scale acoustic monitoring of avian biodiversity and phenology. *bioRxiv* 2024;p. 2024–05.

7. Brunk, K. M. et al. Quail on fire: changing fire regimes may benefit mountain quail in fire-adapted forests. *Fire Ecology*. **19**(1), 1–13 (2023).
8. Kahl, S., Wood, C. M., Eibl, M. & Klinck, H. BirdNET: A deep learning solution for avian diversity monitoring. *Ecological Informatics*. **61**, 101236 (2021).
9. Stowell, D., Wood, M. D., Pamula, H., Stylianou, Y. & Glotin, H. Automatic acoustic detection of birds through deep learning: the first bird audio detection challenge. *Methods in Ecology and Evolution*. **10**(3), 368–380 (2019).
10. Xeno-canto Foundation and Naturalis Biodiversity Center, xeno-canto; <https://xeno-canto.org>.
11. The Cornell Lab of Ornithology, Macaulay Library; <https://www.macaulaylibrary.org>.
12. Ghani, B., Denton, T., Kahl, S. & Klinck, H. Global birdsong embeddings enable superior transfer learning for bioacoustic classification. *Scientific Reports*. **13**(1), 22876 (2023).
13. Bommasani, R. et al. On the opportunities and risks of foundation models. arXiv preprint [arXiv:2108.07258](https://arxiv.org/abs/2108.07258) 2021.
14. van Merriënboer, B., Hamer, J., Dumoulin, V., Triantafyllou, E. & Denton, T. Birds, bats and beyond: Evaluating generalization in bioacoustics models. *Frontiers in Bird Science*. **3**, 1369756 (2024).
15. Wang, D., & Zheng, T. F. Transfer learning for speech and language processing. In: 2015 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA) IEEE; 2015. p. 1225–1237.
16. Ghosal, D. & Kolekar, M. H. Music Genre Recognition Using Deep Neural Networks and Transfer Learning. In: Interspeech; 2018. p. 2087–2091.
17. Tsalera, E., Papadakis, A. & Samarakou, M. Comparison of pre-trained CNNs for audio classification using transfer learning. *Journal of Sensor and Actuator Networks*. **10**(4), 72 (2021).
18. Gemmeke, J. F. et al. Audio set: An ontology and human-labeled dataset for audio events. In: 2017 IEEE international conference on acoustics, speech and signal processing (ICASSP) IEEE; 2017. p. 776–780.
19. Hershey, S. et al. CNN architectures for large-scale audio classification. In: 2017 IEEE international conference on acoustics, speech and signal processing (icassp) IEEE; 2017. p. 131–135.
20. Gong, Y., Chung, Y. A. & Glass, J. Psl: Improving audio tagging with pretraining, sampling, labeling, and aggregation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **29**, 3292–3306 (2021).
21. Koutini, K., Schlüter, J., Eghbal-Zadeh, H. & Widmer, G. Efficient training of audio transformers with patchout. arXiv preprint [arXiv:2110.05069](https://arxiv.org/abs/2110.05069) 2021;.
22. Çoban EB, Pir D, So R, Mandel MI. Transfer learning from youtube soundtracks to tag arctic ecoacoustic recordings. In: ICASSP 2020–2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) IEEE; 2020. p. 726–730.
23. Lauha, P. et al. Domain-specific neural networks improve automated bird sound recognition already with small amount of local data. *Methods in Ecology and Evolution*. **13**(12), 2799–2810 (2022).
24. Dufourq, E., Batist, C., Foquet, R. & Durbach, I. Passive acoustic monitoring of animal populations with transfer learning. *Ecological Informatics*. **70**, 101688 (2022).
25. Gou, J., Yu, B., Maybank, S. J. & Tao, D. Knowledge distillation: A survey. *International Journal of Computer Vision*. **129**(6), 1789–1819 (2021).
26. Hinton, G., Vinyals, O. & Dean, J. Distilling the knowledge in a neural network. arXiv preprint [arXiv:1503.02531](https://arxiv.org/abs/1503.02531) 2015;.
27. Priebe, D., Ghani, B. & Stowell, D. Efficient speech detection in environmental audio using acoustic recognition and knowledge distillation. *Sensors*. **24**(7), 2046 (2024).
28. Choi, K., Kersner M., Morton, J. & Chang, B. Temporal knowledge distillation for on-device audio classification. In: ICASSP 2022–2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) IEEE; 2022. p. 486–490.
29. Ntalampiras, S. Bird species identification via transfer learning from music genres. *Ecological informatics*. **44**, 76–81 (2018).
30. McCloskey, M. & Cohen, N. J. Catastrophic interference in connectionist networks: The sequential learning problem. In: Psychology of learning and motivation, vol. 24 Elsevier; 1989. p. 109–165.
31. Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., & Bengio, Y. An empirical investigation of catastrophic forgetting in gradient-based neural networks. arXiv preprint [arXiv:1312.6211](https://arxiv.org/abs/1312.6211) 2013;.
32. Howard, J. & Ruder, S. Universal language model fine-tuning for text classification. arXiv preprint [arXiv:1801.06146](https://arxiv.org/abs/1801.06146) 2018.
33. Yosinski, J., Clune, J., Bengio, Y., & Lipson H. How transferable are features in deep neural networks? Advances in neural information processing systems 2014;27.
34. Raffel, C. et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*. **21**(140), 1–67 (2020).
35. White, E. L., Klinck, H., Bull, J. M., White, P. R. & Risch, D. One size fits all? Adaptation of trained CNNs to new marine acoustic environments. *Ecological Informatics*. **78**, 102363 (2023).
36. Williams B, van Merriënboer B, Dumoulin V, Hamer J, Triantafyllou E, Fleishman AB, et al. Leveraging tropical reef, bird and unrelated sounds for superior transfer learning in marine bioacoustics. arXiv preprint [arXiv:2404.16436](https://arxiv.org/abs/2404.16436) 2024.
37. Schmid, F., Koutini, K. & Widmer, G. Efficient large-scale audio tagging via transformer-to-cnn knowledge distillation. In: ICASSP 2023–2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) IEEE; 2023. p. 1–5.
38. Gong, Y., Khurana, S., Rouditchenko, A. & Glass, J. Cmkd: Cnn/transformer-based cross-model knowledge distillation for audio classification. arXiv preprint [arXiv:2203.06760](https://arxiv.org/abs/2203.06760) 2022.
39. Touvron, H. et al. Training data-efficient image transformers & distillation through attention. In: International conference on machine learning PMLR; 2021. p. 10347–10357.
40. Beyer, L. et al. Knowledge distillation: A good teacher is patient and consistent. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition; 2022. p. 10925–10934.
41. Keller, V. et al. *European breeding bird atlas 2: distribution, abundance and change* (European Bird Census Council & Lynx Edicions, Barcelona, 2020).
42. Vellinga, W. P. & Planqué, R. The Xeno-canto Collection and its Relation to Sound Recognition and Classification. In: CLEF (Working Notes); 2015.
43. Warblr the bird recognition app, warblrb10k dataset; <https://dcase.community/challenge2018/task-bird-audio-detection#audio-datasets>.
44. Observation International, Observation.org; <https://observation.org>.
45. Naturkundemuseum Bayern, BIOTOPIA LAB, and Landesbund für Vogelschutz, Dawn Chorus project; 2024. <https://dawn-chorus.org/>.
46. Naturkundemuseum Bayern, BIOTOPIA LAB, and Landesbund für Vogelschutz, Dawn Chorus project - Explore the dawn chorus recordings. <https://explore.dawn-chorus.org>.
47. Kahl, S. BirdNET-Analyzer; 2023. <https://github.com/kahst/BirdNET-Analyzer>.
48. Gong, Y., Chung, Y. A. & Glass, J. Ast: Audio spectrogram transformer. arXiv preprint [arXiv:2104.01778](https://arxiv.org/abs/2104.01778) 2021;.
49. Dosovitskiy, A., et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929) 2020;.
50. Koutini, K. PaSST; 2021. <https://github.com/kkoutini/PaSST>.
51. Gong, Y. PSLA; 2022. <https://github.com/YuanGongND/psla>.
52. Zhang, H., Cisse, M., Dauphin, Y. N. & Lopez-Paz, D. mixup: Beyond empirical risk minimization. arXiv preprint [arXiv:1710.09412](https://arxiv.org/abs/1710.09412) 2017;.

53. Hanley, J. A. & McNeil, B. J. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*. **143**(1), 29–36 (1982).
54. Boudiaf, M., Denton, T., Van Merriënboer, B., Dumoulin, V. & Triantafillou, E. In search for a generalizable method for source free domain adaptation. In: International Conference on Machine Learning PMLR; 2023. p. 2914–2931.
55. Kong, Q. et al. Panns: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*. **28**, 2880–2894 (2020).
56. Bhati, S. et al. DASS: Distilled Audio State Space Models Are Stronger and More Duration-Scalable Learners. arXiv preprint [arXiv:2407.04082](https://arxiv.org/abs/2407.04082) 2024.
57. Castro-Ospina, A. E., Solarte-Sanchez, M. A., Vega-Escobar, L. S., Isaza, C. & Martínez-Vargas, J. D. Graph-Based Audio Classification Using Pre-Trained Models and Graph Neural Networks. *Sensors*. **24**(7), 2106 (2024).
58. Sethi, S. S et al. Large-scale avian vocalization detection delivers reliable global biodiversity insights. Proceedings of the National Academy of Sciences 2024 Aug;121(33). <https://doi.org/10.1073/pnas.2315933121>.
59. Wood, C. M & Kahl, S. Guidelines for appropriate use of BirdNET scores and other detector outputs. *Journal of Ornithology* 2024;p. 1–6.

Acknowledgements

The Dawn Chorus data were collected through the Dawn Chorus Project <https://dawn-chorus.org/> by Naturk undemuseum Bayern/BIOTOPIA LAB and LBV. We extend our gratitude to Lisa Gill for enabling us to access the Dawn Chorus annotation data subset. We thank the many volunteers who contributed sound recordings and annotations to Xeno-canto, to Observation.org, to Warblr, and to the Dawn Chorus project.

Author contributions

Burooj Ghani: Conceptualization, Methodology, Software, Validation, Investigation, Data curation, Writing - original draft, Visualization. Vincent J. Kalkman: Conceptualization, Investigation, Data curation, Supervision, Writing - review & editing, Project administration. Bob Planqué: Data curation, Writing - review & editing. Willem-Pier Vellinga: Data curation, Writing - review & editing. Lisa Gill: Data curation, Writing - review & editing. Dan Stowell: Conceptualization, Methodology, Investigation, Writing - original draft, Supervision, Project administration.

Funding

Burooj Ghani is supported by the EU-funded projects Mambo and Guarden. Specifically, Mambo is funded under grant agreement No. 101060639, and Guarden is supported by grant agreement No. 101060693.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to B.G.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025