

Computational bioacoustics with deep learning: a review and roadmap

Dan Stowell^{1,2}

¹ Department of Cognitive Science and Artificial Intelligence, Tilburg University, Tilburg, The Netherlands

² Naturalis Biodiversity Center, Leiden, The Netherlands

ABSTRACT

Animal vocalisations and natural soundscapes are fascinating objects of study, and contain valuable evidence about animal behaviours, populations and ecosystems. They are studied in bioacoustics and ecoacoustics, with signal processing and analysis an important component. Computational bioacoustics has accelerated in recent decades due to the growth of affordable digital sound recording devices, and to huge progress in informatics such as big data, signal processing and machine learning. Methods are inherited from the wider field of deep learning, including speech and image processing. However, the tasks, demands and data characteristics are often different from those addressed in speech or music analysis. There remain unsolved problems, and tasks for which evidence is surely present in many acoustic signals, but not yet realised. In this paper I perform a review of the state of the art in deep learning for computational bioacoustics, aiming to clarify key concepts and identify and analyse knowledge gaps. Based on this, I offer a subjective but principled roadmap for computational bioacoustics with deep learning: topics that the community should aim to address, in order to make the most of future developments in AI and informatics, and to use audio data in answering zoological and ecological questions.

Subjects Animal Behavior, Bioinformatics, Computational Biology, Ecology, Data Mining and Machine Learning

Keywords Acoustics, Deep learning, Bioacoustics, Animal vocal behaviour, Sound, Machine learning, Passive acoustic monitoring

INTRODUCTION

Bioacoustics—the study of animal sound—offers a fascinating window into animal behaviour, and also a valuable evidence source for monitoring biodiversity (*Marler & Slabbekoorn, 2004; Laiolo, 2010; Marques et al., 2012; Brown & Riede, 2017*). Bioacoustics has long benefited from computational analysis methods including signal processing, data mining and machine learning (*Towsey et al., 2012; Ganchev, 2017*). Within machine learning, *deep learning* (DL) has recently revolutionised many computational disciplines: early innovations, motivated by the general aims of artificial intelligence (AI) and developed for image or text processing, have cascaded through to many other fields (*LeCun, Bengio & Hinton, 2015; Goodfellow, Bengio & Courville, 2016*). This includes audio domains such as automatic speech recognition and music informatics (*Abeßer, 2020; Manilow, Seetharman & Salamon, 2020*).

Submitted 22 December 2021

Accepted 1 March 2022

Published 21 March 2022

Corresponding author

Dan Stowell,

d.stowell@tilburguniversity.edu

Academic editor

Yilun Shang

Additional Information and
Declarations can be found on
page 31

DOI 10.7717/peerj.13152

© Copyright

2022 Stowell

Distributed under

Creative Commons CC-BY 4.0

OPEN ACCESS

Computational bioacoustics is now also benefiting from the power of DL to solve and automate problems that were previously considered intractable. This is both enabled and demanded by the twenty-first century data deluge: digital recording devices, data storage and sharing have become dramatically more widely available, and affordable for large-scale bioacoustic monitoring, including continuous audio capture ([Ranft, 2004](#); [Roe et al., 2021](#); [Webster & Budney, 2017](#); [Roch et al., 2017](#)). The resulting deluge of audio data means that a common bottleneck is the lack of person-time for trained analysts, heightening the importance of methods that can automate large parts of the workflow, such as machine learning.

The revolution in this field is real, but it is recent: reviews and textbooks as recently as 2017 did not give much emphasis to DL as a tool, even when focussing on machine learning for bioacoustics ([Ganchev, 2017](#); [Stowell, 2018](#)). [Mercado & Sturdy \(2017\)](#) reviewed the ways in which artificial neural networks (hereafter, *neural networks* or NNs) had been used by bioacoustics researchers; however, that review concerns the pre-deep-learning era of neural networks, which has some foundational aspects in common but many important differences, both conceptual and practical.

Many in bioacoustics are now grappling with deep learning, and there is much interesting work which uses and adapts DL for the specific requirements of bioacoustic analysis. Yet the field is immature and there are few reference works. This review aims to provide an overview of the emerging field of deep learning for computational bioacoustics, reviewing the state of the art and clarifying key concepts. A particular goal of this review is to identify knowledge gaps and under-explored topics that could be addressed by research in the next few years. Hence, after stating the survey methodology, I summarise the current state of the art, outlining standard good practice and the tasks themes addressed. I then offer a ‘roadmap’ of work in deep learning for computational bioacoustics, based on the survey and thematic analysis, and drawing in topics from the wider field of deep learning as well as broad topics in bioacoustics.

SURVEY METHODOLOGY

Deep learning is a recent and rapidly-moving field. Although deep learning has been applied to audio tasks (such as speech, music) for more than ten years, its application in wildlife bioacoustics is recent and immature. Key innovations in acoustic deep learning include [Hershey et al. \(2017\)](#) which represents the maturing of audio recognition based on convolutional neural networks (CNNs)—it introduced a dataset (*AudioSet*) and a NN architecture (*VGGish*) both now widely-used; and convolutional-recurrent neural network (CRNN) methods ([Çakır et al., 2017](#)). The organisers of the *BirdCLEF* data challenge announced “the arrival of deep learning” in 2016 ([Goëau et al., 2016](#)). Hence I chose to constrain keyword-based literature searches, using both Google Scholar and Web of Science, to papers published no earlier than 2016. The query used was:

```
(bioacoust* OR ecoacoust* OR vocali* OR “animal calls” OR
“passive acoustic monitoring” OR “soundscape”) AND (“deep learning” OR
“convolutional neural network” OR “recurrent neural network”) AND (animal
OR bird* OR cetacean* OR insect* OR mammal*)
```

With Google Scholar, this search yielded 989 entries. Many were excluded due to being off-topic, or being duplicates, reviews, abstract-only, not available in English, or unavailable. Various preprints (non-peer reviewed) were encountered (in arXiv, biorXiv and more): I did not exclude all of these, but priority was given to peer-reviewed published work. With Web of Science, the same search query yielded 56 entries. After merging and deduplication, this yielded 162 articles. This sub-field is a rapidly-growing one: the set of selected articles grows from 5 in 2016 through to 64 in 2021. The bibliography file published along with this paper lists all these articles, plus other articles added for context while writing the review.

STATE OF THE ART AND RECENT DEVELOPMENTS

I start with a standard recipe abstracted from the literature, and the taxonomic coverage of the literature, before reviewing those topics in bioacoustic DL that have received substantial attention and are approaching maturity. To avoid repetition, some of the more tentative or unresolved topics will be deferred to the later ‘roadmap’ section, even when discussed in existing literature.

The standard recipe for bioacoustic deep learning

Deep learning is flexible and can be applied to many different tasks, from classification/regression through to signal enhancement and even synthesis of new data. The ‘workhorse’ of DL is however classification, by which we mean assigning data items one or more ‘labels’ from a fixed list (*e.g.*, a list of species, individuals or call types). This is the topic of many DL breakthroughs, and many other tasks have been addressed in part by using the power of DL classification—even image generation ([Goodfellow et al., 2014](#)).

Classification is indeed the main use of DL seen in computational bioacoustics.

A typical ‘recipe’ for bioacoustic classification using deep learning, applicable to very many of the surveyed articles from recent years, is as follows. Some of the terminology may be unfamiliar, and I will expand upon it in later sections of the review:

- Use one of the well-known CNN architectures (ResNet, VGGish, Inception, MobileNet), perhaps pretrained from AudioSet. (These are conveniently available within the popular DL Python frameworks PyTorch, Keras, TensorFlow).
- The input will be spectrogram data, typically divided into audio clips of fixed size such as 1 s or 10 s, which is done so that a ‘batch’ of spectrograms fits easily into GPU memory. The spectrograms may be standard (linear-frequency), or mel spectrograms, or log-frequency spectrograms. The “pixels” in the spectrogram are magnitudes: typically these are log-transformed before use, but might not be, or alternatively transformed by per-channel energy normalisation (PCEN). There is no strong consensus on the ‘best’ spectrogram format—it is likely a simple empirical choice based on the frequency bands of interest in your chosen task and their dynamic ranges.
- The list of labels to be predicted could concern species, individuals, call types, or something else. It may be a binary (yes/no) classification task, which could be used for detecting the presence (occupancy) of some sound. In many cases a list of species is

used: modern DL can scale to many hundreds of species. The system may be configured to predict a more detailed output such as a transcription of multiple sound events; I return to this later.

- Use data augmentation to artificially make a small bioacoustic training dataset more diverse (noise mixing, time shifting, mixup).
- Although a standard CNN is common, CRNNs are also relatively popular, adding a recurrent layer (LSTM or GRU) after the convolutional layers, which can be achieved by creating a new network from scratch or by adding a recurrent layer to an off-the-shelf network architecture.
- Train your network using standard good practice in deep learning (for example: Adam optimiser, dropout, early stopping, and hyperparameter tuning) (*Goodfellow, Bengio & Courville, 2016*).
- Following good practice, there should be separate data(sub)sets for training, validation (used for monitoring the progress of training and for selecting hyperparameters), and final testing/evaluation. It is especially beneficial if the testing set represents not just unseen data items but novel conditions, to better estimate the true generalisability of the system (*Stowell et al., 2019b*). However, it is still common for the training/validation/testing data to be sampled from the same pool of source data.
- Performance is measured using standard metrics such as accuracy, precision, recall, F-score, and/or area under the curve (AUC or AUROC). Since bioacoustic datasets are usually “unbalanced”, having many more items of one category than another, it is common to account for this—for example by using *macro-averaging*, calculating performance for each class and then taking the average of those to give equal weight to each class (*Mesaros, Heittola & Virtanen, 2016*).

This standard recipe will work well for many bioacoustic classification tasks, including noisy outdoor sound scenes. (Heavy rain and wind remains a problem across all analysis methods, including DL). It can be implemented using a handful of well-known Python libraries: PyTorch/TensorFlow/Keras, librosa or another library for sound file processing, plus a data augmentation tool such as SpecAugment, audiomentations or kapre. The creation and data augmentation of spectrograms is specific to audio domains, but the CNN architecture and training is standard across DL for images, audio, video and more, which has the benefit of being able to inherit good practice from this wider field.

Data augmentation helps with small and also with unbalanced datasets, common in bioacoustics. The commonly-used augmentation methods (time-shifting, sound mixing, noise mixing) are “no regret” in that it is extremely unlikely these modifications will damage the semantic content of the audio. Other modifications such as time warping, frequency shifting, frequency warping, modify sounds in ways which could alter subtle cues such as those that might distinguish individual animals or call types from one another. Hence the appropriate choice of augmentation methods is audio-specific and even animal-sound-specific.

The standard recipe does though have its limits. The use of the mel frequency scale, AudioSet pretraining, and magnitude-based spectrograms (neglecting some details of phase or temporal fine structure) all bias the process towards aspects of audio that are easily perceptible to humans, and thus may overlook some details that are important for fine high-resolution discriminations or for matching animal perception (*Morfi, Lachlan & Stowell, 2021a*). All the common CNN architectures have small-sized convolutional filters, biasing them towards objects that are compact in the spectrogram, potentially an issue for broad-band sound events.

There are various works that pay close attention to the parameters of spectrogram generation, or argue for alternative representations such as wavelets. This engineering work can lead to improved performance in each chosen task, especially for difficult cases such as echolocation clicks. However, as has been seen in previous eras of audio analysis, these are unlikely to overturn standard use of spectrograms since the improvement rarely generalises across many tasks. Networks using raw waveforms as input may overcome many of these concerns in future, though they require larger training datasets; pretrained raw-waveform networks may be a useful tool to look forward to in the near term.

Taxonomic coverage

Species/taxa whose vocalisations have been analysed through DL include:

- Birds—the most commonly studied group, covered by at least 65 of the selected papers. Specific examples will be cited below, and some overviews can be found in the outcomes of data challenges and workshops (*Stowell et al., 2019a, Joly et al., 2019*).
- Cetaceans and other marine mammals—another very large subfield, covered by 30 papers in the selected set. Again, data challenges and workshops are devoted to these taxa (*Frazao, Padovese & Kirsebom, 2020*).
- Bats (*Mac Aodha et al., 2018; Chen et al., 2020c, Fujimori et al., 2021; Kobayashi et al., 2021; Zhang et al., 2020; Zualkernan et al., 2020, 2021*).
- Terrestrial mammals (excluding bats): including primates (*Bain et al., 2021; Dufourq et al., 2021; Romero-Mujalli et al., 2021; Oikarinen et al., 2019; Tzirakis et al., 2020*), elephants (*Bjorck et al., 2019*), sheep (*Wang et al., 2021*), cows (*Jung et al., 2021*), koalas (*Himawan et al., 2018*).
- A particular subset of work focusses on mouse and rat ultrasonic vocalisations (USVs). These have been of interest particularly in laboratory mice studies, hence a vigorous subset of literature based on rodent USVs primarily recorded in laboratory conditions (*Coffey, Marx & Neumaier, 2019; Fonseca et al., 2021; Ivanenko et al., 2020; Premoli et al., 2021; Smith & Kristensen, 2017; Steinfath et al., 2021*).
- Anurans (*Colonna et al., 2016; Dias, Ponti & Minghim, 2021; Hassan, Ramli & Jaafar, 2017; Islam & Valles, 2020; LeBien et al., 2020; Xie et al., 2021b, 2020, 2021c*).

- Insects (*Hibino, Suzuki & Nishino, 2021; Khalighifar et al., 2021; Kiskin et al., 2021; Kiskin et al., 2020; Rigakis et al., 2021; Sinka et al., 2021; Steinfath et al., 2021*).
- Fish (*Guyot et al., 2021; Ibrahim et al., 2018; Waddell, Rasmussen & Širović, 2021*).

Many works cover more than one taxon, since DL enables multi-species recognition across a large number of categories and benefits from large and diverse data. Some works sidestep taxon considerations by focusing on the overall ecosystem or soundscape level (“ecoacoustic” approaches) (*Sethi et al., 2020; Heath et al., 2021; Fairbrass et al., 2019*).

The balance of emphasis across taxa has multiple drivers. Many of the above taxa are important for biodiversity and conservation monitoring (including birds, bats, insects), or for comparative linguistics and behaviour studies (songbirds, cetaceans, primates, rodents). For some taxa, acoustic communication is a rich and complex part of their behaviour, and their vocalisations have a high complexity which is amenable to signal analysis (*Marler & Slabbekoorn, 2004*). On the other hand, progress is undeniably driven in part by practical considerations, such as the relative ease of recording terrestrial and diurnal species. Aside from standard open science practices such as data sharing, progress in bird sound classification has been stimulated by large standardised datasets and automatic recognition challenges, notably the *BirdCLEF* challenge conducted annually since 2014 (*Goëau et al., 2014; Joly et al., 2019*). This dataset- and challenge-based progress follows a pattern of work seen in many applications of machine learning. Nonetheless, the allocation of research effort does not necessarily match up with the variety or importance of taxa—a topic I will return to.

Having summarised a standard recipe and the taxonomic coverage of the literature, I next review the themes that have received detailed attention in the literature on DL for computational bioacoustics.

Neural network architectures

The “architecture” of a neural network (NN) is the general layout of the nodes and their interconnections, often arranged in sequential layers of processing (*Goodfellow, Bengio & Courville, 2016*). Early work applying NNs to animal sound made use of the basic “multi-layer perceptron” (MLP) architecture (*Koops, Van Balen & Wiering, 2015; Houégnigan et al., 2017; Hassan, Ramli & Jaafar, 2017; Mercado & Sturdy, 2017*), with manually-designed summary features (such as syllable duration, peak frequency) as input. However, the MLP is superseded and dramatically outperformed by CNN and (to a lesser extent) recurrent neural network (RNN) architectures, both of which can take advantage of the sequential/grid structure in raw or lightly-preprocessed data, meaning that the input to the CNN/RNN can be time series or time-frequency spectrogram data (*Goodfellow, Bengio & Courville, 2016*). This change—removing the step in which the acoustic data is reduced to a small number of summary features in a manually-designed feature extraction process—keeps the input in a much higher dimensional format, allowing for much richer information to be presented. Neural networks are highly nonlinear and can make use of subtle variation in this “raw” data. CNNs and RNNs apply assumptions about the sequential/grid structure of the data, allowing efficient training. For example,

CNN classifiers are by design invariant to time-shift of the input data. This embodies a reasonable assumption (most sound stimuli do not change category when moved slightly later/earlier in time), and results in a CNN having many fewer free parameters than the equivalent MLP, thus being easier to train.

One of the earliest works applying a CNN to bioacoustic data was to classify among 10 anuran species (Colonna *et al.*, 2016). In the same year, 3 of the 6 teams in the 2016 BirdCLEF challenge submitted CNN systems taking spectrograms as input, including the highest-scoring team (Goëau *et al.*, 2016). Reusing a high-performing CNN architecture from elsewhere is very popular now, but was possible even in 2016: one of the submitted systems re-used a 2012 CNN designed for images, called AlexNet. Soon after, Salamon *et al.* (2017a) and Knight *et al.* (2017) also found that a CNN outperformed the previous “shallow” paradigm of bioacoustic machine learning.

CNNs are now dominant: at least 83 of the surveyed articles made use of CNNs (sometimes in combination with other modules). Many articles empirically compare the performance of selected NN architectures for their tasks, and configuration options such as the number of CNN layers (Wang *et al.*, 2021; Li *et al.*, 2021; Zualkernan *et al.*, 2020). Oikarinen *et al.* (2019) studied an interesting dual task of simultaneously inferring call type and caller ID from devices carried by pairs of marmoset monkeys, evaluating different types of output layer for this dual-task scenario.

While many of the surveyed articles used a self-designed CNN architecture, there is a strong move towards using, or at least evaluating, off-the-shelf CNN architectures (Lasseck, 2018; Zhong *et al.*, 2020a; Guyot *et al.*, 2021; Dias, Ponti & Minghim, 2021; Li *et al.*, 2021; Kiskin *et al.*, 2021; Bravo Sanchez *et al.*, 2021; Gupta *et al.*, 2021). These are typically CNNs that have been influential in DL more widely, and are now available conveniently in DL frameworks (Table 1). They can even be downloaded already pretrained on standard datasets, to be discussed further below. The choice of CNN architecture is rarely a decision that can be made from first principles, aside from general advice that the size/complexity of the CNN should generally scale with that of the task being attempted (Kaplan *et al.*, 2020). Some of the popular recent architectures (notably ResNet and DenseNet) incorporate architectural modifications to make it feasible to train very deep networks; others (MobileNet, EfficientNet, Xception) are designed for efficiency, reducing the number of computations needed to achieve a given level of accuracy (Canziani, Paszke & Culurciello, 2016).

The convolutional layers in a CNN layer typically correspond to non-linear filters with small “receptive fields” in the axes of the input data, enabling them to make use of local dependencies within spectrogram data. However, it is widely understood that sound scenes and vocalisations can be driven by dependencies over both short and very long timescales. This consideration about time series in general was the inspiration for the design of recurrent neural networks (RNNs), with the LSTM and GRU being popular embodiments (Hochreiter & Schmidhuber, 1997): these networks have the capacity to pass information forwards (and/or backwards) arbitrarily far in time while making inferences. Hence, RNNs have often been explored to process sound, including animal sound (Xian *et al.*, 2016; Wang *et al.*, 2021; Madhusudhana *et al.*, 2021; Islam & Valles, 2020;

Table 1 Off-the-shelf CNN architectures used in the surveyed literature, and the number of articles using them. This is indicative only, since not all articles clearly state whether an off-the-shelf model is used, some articles use modified/derived versions, and some use multiple architectures.

CNN architecture	Num articles
ResNet (<i>He et al., 2016</i>)	23
VGG (<i>Simonyan & Zisserman, 2014</i>) or VGGish (<i>Hershey et al., 2017</i>)	17
DenseNet (<i>Huang, Liu & Weinberger, 2016</i>)	7
AlexNet (<i>Krizhevsky, Sutskever & Hinton, 2012</i>)	5
Inception (<i>Szegedy et al., 2014</i>)	4
LeNet (<i>LeCun et al., 1998</i>)	3
MobileNet (<i>Sandler et al., 2018</i>)	2
EfficientNet (<i>Tan & Le, 2019</i>)	2
Xception (<i>Chollet, 2017</i>)	2
U-net (<i>Ronneberger, Fischer & Brox, 2015</i>)	2
bulbul (<i>Grill & Schlüter, 2017</i>)	2
Self-designed CNN	19

Garcia et al., 2020; Ibrahim et al., 2018). An RNN alone is not often found to give strong performance. However, in around 2017 it was observed that adding an RNN layer after the convolutional layers of a CNN could give strong performance in multiple audio tasks, with an interpretation that the RNN layer(s) perform temporal integration of the information that has been preprocessed by the early layers (*Cakir et al., 2017*). This “CRNN” approach has since been applied variously in bioacoustics, often with good results (*Himawan et al., 2018; Morfi & Stowell, 2018; Gupta et al., 2021; Xie et al., 2020; Tzirakis et al., 2020; Li et al., 2019*). However, CRNNs can be more computationally intensive to train than CNNs, and the added benefit is not universally clear.

In 2016 an influential audio synthesis method entitled *WaveNet* showed that it was possible to model long temporal sequences using CNN layers with a special ‘dilated’ structure, enabling many hundreds of time steps to be used as context for prediction (*van den Oord et al., 2016*). This inspired a wave of work replacing recurrent layers with 1-D temporal convolutions, sometimes called temporal CNN (TCN or TCNN) (*Bai, Kolter & Koltun, 2018*). Note that whether applied to spectrograms or waveform data, these are 1-D (time only) convolutions, not the 2-D (time-frequency) convolutions more commonly used. TCNs can be faster to train than RNNs, with similar or superior results. TCNs have been used variously in bioacoustics since 2021, and this is likely to continue (*Steinfath et al., 2021; Fujimori et al., 2021; Roch et al., 2021; Xie et al., 2021b; Gupta et al., 2021; Gillings & Scott, 2021; Bhatia, 2021*). *Gupta et al. (2021)* compare CRNN against CNN+TCN, and also standard CNN architectures (ResNet, VGG), with CRNN the strongest method in their evaluation.

Innovations in NN architectures continue to be explored. *Vesperini et al. (2018)* apply capsule networks for bird detection. *Gupta et al. (2021)* apply Legendre memory units, a novel type of recurrent unit, in birdsong species classification. When we later review

“object detection”, we will encounter some custom architectures for that task. In wider DL, especially text processing, it is popular to use a NN architectural modification referred to as “attention” ([Chorowski et al., 2015](#)). The structure of temporal sequences is highly variable, yet CNN and RNN architectures implicitly assume that the pattern of previous timesteps that are important predictors is fixed. Attention networks go beyond this by combining inputs in a weighted combination whose weights are determined on-the-fly. (Note that this is not in any strong sense a model of auditory attention as considered in cognitive science.) This approach was applied to spectrograms by [Ren et al. \(2018\)](#), and used for bird vocalisations by [Morfi, Lachlan & Stowell \(2021a\)](#). A recent trend in DL has been to use attention (as opposed to convolution or recurrence) as the fundamental building block of an NN architecture, known as “transformer” layers ([Vaswani et al., 2017](#)). Transformers are not yet widely explored in bioacoustic tasks, but given their strong performance in other domains we can expect their use to increase. The small number of recent studies shows encouraging results ([Elliott et al., 2021](#); [Wolters et al., 2021](#)).

Many studies compare NN architectures empirically, usually from a manually-chosen set of options, perhaps with evaluation over many hyperparameter settings such as the number of layers. There are too many options to search them all exhaustively, and too little guidance on how to choose a network *a priori*. [Brown, Montgomery & Garg \(2021\)](#) propose one way to escape this problem: a system to automatically construct the workflow for a given task, including NN architecture selection.

Acoustic features: spectrograms, waveforms, and more

In the vast majority of studies surveyed, the magnitude spectrogram is used as input data. This is a representation in which the raw audio time series has been lightly processed to a 2D grid, whose values indicate the energy present at a particular time and frequency. Prior to the use of DL, the spectrogram would commonly be used as the source for subsequent feature extraction such as peak frequencies, sound event durations, and more. Using the spectrogram itself allows a DL system potentially to make use of diverse information in the spectrogram; it also means the input is a similar format to a digital image, thus taking advantage of many of the innovations and optimisations taking place in image DL.

Standard options in creating a spectrogram include the window length for the short-time Fourier transforms used (and thus the tradeoff of time- vs frequency-resolution), and the shape of the window function ([Jones & Baraniuk, 1995](#)). Mild benefits can be obtained by careful selection of these parameters, and have been argued for in DL ([Heuer et al., 2019](#); [Knight et al., 2020](#)). A choice more often debated is whether to use a standard spectrogram with its linear frequency axis, or to use a (pseudo-)logarithmically-spaced frequency axis such as the mel spectrogram ([Xie et al., 2019](#); [Zuolkernan et al., 2020](#)) or constant-Q transform (CQT) ([Himawan et al., 2018](#)). The mel spectrogram uses the mel scale, originally intended as an approximation of human auditory selectivity, and thus may seem an odd choice for non-human data. Its use likely owes a lot to

convenience, but also to the fact that pitch shifts of harmonic signals correspond to linear shifts on a logarithmic scale—potentially a good match for CNNs which are designed to detect linearly-shifted features reliably. [Zualkernan et al. \(2020\)](#) even found a mel spectrogram representation useful for bat signals, with of course a modification of the frequency range. The literature presents no consensus, with evaluations variously favouring the mel ([Xie et al., 2019](#); [Zualkernan et al., 2020](#)), logarithmic ([Himawan et al., 2018](#); [Smith & Kristensen, 2017](#)), or linear scale ([Bergler et al., 2019b](#)). There is likely no representation that will be consistently best across all tasks and taxa. Some studies take advantage of multiple spectrogram representations of the same waveform, by “stacking” a set of spectrograms into a multi-channel input (processed in the same fashion as the colour channels in an RGB image) ([Thomas et al., 2019](#); [Xie et al., 2021c](#)). The channels are extremely redundant with one another; this stacking allows the NN flexibly to use information aggregated across these closely-related representations, and thus gain a small informational advantage.

ML practitioners must concern themselves with how their data are normalised and preprocessed before input to a NN. Standard practice is to transform input data to have zero mean and unit variance, and for spectrograms perhaps to apply light noise-reduction such as by median filtering. In practice, spectral magnitudes can have dramatically varying dynamic ranges, noise levels and event densities. [Lostanlen et al. \(2019a, 2019b\)](#) give theoretical and empirical arguments for the use of *per-channel energy normalisation* (PCEN), a simple adaptive normalisation algorithm. Indeed PCEN has been deployed by other recent works, and found to permit improved performance of deep bioacoustic event detectors ([Allen et al., 2021](#); [Morfi et al., 2021b](#)).

As an aside, previous eras of acoustic analysis have made widespread use of mel-frequency cepstral coefficients (MFCCs), a way of compressing spectral information into a small number of standardised measurements. MFCCs have occasionally been used in bioacoustic DL ([Colonna et al., 2016](#); [Kojima et al., 2018](#); [Jung et al., 2021](#)). However, they are likely to be a poor match to CNN architectures since sounds are not usually considered shift-invariant along the MFCC coefficient axis. Deep learning evaluations typically find that MFCCs are outperformed by less-preprocessed representations such as the (closely-related) mel spectrogram ([Zualkernan et al., 2020](#); [Elliott et al., 2021](#)).

Other types of time-frequency representation are explored by some authors as input to DL, such as wavelets ([Smith & Kristensen, 2017](#); [Kiskin et al., 2020](#)) or traces from a sinusoidal pitch tracking algorithm ([Jancovic & Köküer, 2019](#)). These can be motivated by considerations of the target signal, such as chirplets as a match to the characteristics of whale sound ([Glotin, Ricard & Balestrierio, 2017](#)).

However, the main alternative to spectrogram representations is in fact to use the raw waveform as input. This is now facilitated by NN architectures such as WaveNet and TCN mentioned above. DL based on raw waveforms is often found to require larger datasets for training than that based on spectrograms; one of the main attractions is to remove yet another of the manual preprocessing steps (the spectrogram transformation),

allowing the DL system to extract information in the fashion needed. A range of recent studies use TCN architectures (also called 1-dimensional CNNs) applied to raw waveform input (*Ibrahim et al., 2018; Li et al., 2019; Fujimori et al., 2021; Roch et al., 2021; Xie et al., 2021b*). *Ibrahim et al. (2018)* compares an RNN against a TCN, both applied to waveforms for fish classification; *Li et al. (2019)* applies a TCN with a final recurrent layer to bird sound waveforms. *Steinfath et al. (2021)* offer either spectrogram or waveform input for their CNN segmentation method. *Bhatia (2021)* investigates bird sound synthesis using multiple methods including WaveNet. Transformer architectures can also be applied directly to waveform data (*Elliott et al., 2021*).

Some recent work has proposed trainable representations that are intermediate between raw waveform and spectrogram methods (*Balestriero et al., 2018; Ravanelli & Bengio, 2018; Zeghidour et al., 2021*). These essentially act as parametric filterbanks, whose filter parameters are optimised along with the other NN layer parameters. *Balestriero et al. (2018)* introduce their own trainable filterbank and achieve promising results for bird audio detection. *Bravo Sanchez et al. (2021)* applies a representation called SincNet, achieving competitive results on birdsong classification with a benefit of short training time. *Zeghidour et al. (2021)* apply SincNet but also introduce an alternative learnable audio front-end “LEAF”, which incorporates a trainable PCEN layer as well as a trainable filterbank, finding strong performance on a bird audio detection task.


To summarise this discussion: in many cases a spectrogram representation is appropriate for bioacoustic DL, often with (pseudo-)logarithmic frequency axis such as mel spectrogram or CQT spectrogram. PCEN appears often to be useful for spectrogram preprocessing. Methods using raw waveforms and adaptive front-ends are likely to gain increased prominence, especially if incorporated into some standard off-the-shelf NN architectures that are found to work well across bioacoustic tasks.

Classification, detection, clustering


The most common tasks considered in the literature, by far, are classification and detection. These tasks are fundamental building blocks of many workflows; they are also the tasks that are most comprehensively addressed by the current state of the art in deep learning.

The terms *classification* and *detection* are used in various ways, sometimes interchangeably. In this review I interpret ‘classification’ as in much of ML, the prediction of one or more categorical labels such as species or call type. Classification is very commonly investigated in bioacoustic DL. It is most widely used for species classification—typically within a taxon family, such as in the BirdCLEF challenge (*Joly et al., 2021*) (see above for other taxon examples). Other tasks studied are to classify among individual animals (*Oikarinen et al., 2019; Ntalampiras & Potamitis, 2021*), call types (*Bergler et al., 2019a; Waddell, Rasmussen & Širović, 2021*), sex and strain (within-species) (*Ivanenko et al., 2020*), or behavioural states (*Wang et al., 2021; Jung et al., 2021*). Some work broadens the focus beyond animal sound to classify more holistic soundscape categories such as biophony, geophony, anthropophony (*Fairbrass et al., 2019; Mishachandar & Vairamuthu, 2021*).

(a) Binary classification



(b) SED (multi-species)



(c) Object detection




Figure 1 Three common approaches to implementation of sound detection. Adapted from *Stowell et al. (2016b)*. Full-size DOI: 10.7717/peerj-13152/fig-1

There are three different ways to define a ‘detection’ task that are common in the surveyed literature (Fig. 1):

1. The first is **detection as binary classification**: for a given audio clip, return a binary yes/no decision about whether the signal of interest is detected within (*Stowell et al., 2019a*). This output would be described by a statistical ecologist as “occupancy” information (presence/absence). It is simple to implement since binary classification is a fundamental task in DL, and does not require data to be labelled in high-resolution detail. Perhaps for these reasons it is widely used in the surveyed literature (e.g. *Mac Aodha et al., 2018; Prince et al., 2019; Kiskin et al., 2021; Bergler et al., 2019b; Himawan et al., 2018; Lostanlen et al., 2019b*).
2. The second is detection as transcription, returning slightly more detail: the start and end times of sound events (*Morfi et al., 2019; Morfi et al., 2021b*). In the DCASE series of challenges and workshops (Detection and Classification of Acoustic Scenes and Events), the task of transcribing sound events, potentially for multiple classes in parallel, is termed **sound event detection (SED)**, and in the present review I will use that terminology. It has typically been approached by training DL to label each small time step (e.g. a segment of 10 ms or 1 s) as positive or negative, and sequences of positives are afterwards merged into predicted event regions (*Kong, Xu & Plumbley, 2017; Madhusudhana et al., 2021; Marchal, Fabianek & Aubry, 2021*).
3. The third is the form common in **image object detection**, which consists of estimating multiple bounding boxes indicating object locations within an image. Transferred to spectrogram data, each bounding box would represent time and frequency bounds for

an “object” (a sound event). This has not often been used in bioacoustics but may be increasing in interest ([Venkatesh, Moffat & Miranda, 2021](#); [Shrestha et al., 2021](#); [Romero-Mujalli et al., 2021](#); [Zsebök et al., 2019](#); [Coffey, Marx & Neumaier, 2019](#)).

For all three of these task settings, CNN-based networks are found to have strong performance, outperforming other ML techniques ([Marchal, Fabianek & Aubry, 2021](#); [Knight et al., 2017](#); [Prince et al., 2019](#)). Since the data format in each of the three task settings is different, the final (output) layers of a network take a slightly different form, as do the loss function used to optimise them ([Mesaros et al., 2019](#)). (Other settings are possible, for example pixel-wise segmentation of arbitrary spectral shapes ([Narasimhan, Fern & Raich, 2017](#))).

The detection problem is important in the marine context, where surveys over very large scales of time and space are common, yielding very large data-processing demands, and the sounds to be detected may be very sparsely-occurring ([Frazao, Padovese & Kirsebom, 2020](#)). For this reason there are numerous works applying DL to detection of cetacean sounds underwater ([Jiang et al., 2019](#); [Bergler et al., 2019b](#); [Best et al., 2020](#); [Shiu et al., 2020](#); [Zhong et al., 2020a](#); [Ibrahim et al., 2021](#); [Vickers et al., 2021b](#); [Zhong et al., 2021](#); [Roch et al., 2021](#); [Vickers et al., 2021a](#); [Allen et al., 2021](#); [Madhusudhana et al., 2021](#)).

In bioacoustics it is common to follow a two-step “detect then classify” workflow ([Waddell, Rasmussen & Širović, 2021](#); [LeBien et al., 2020](#); [Schröter et al., 2019](#); [Jiang et al., 2019](#); [Koumura & Okanoya, 2016](#); [Zhong et al., 2021](#); [Padovese et al., 2021](#); [Frazao, Padovese & Kirsebom, 2020](#); [Garcia et al., 2020](#); [Marchal, Fabianek & Aubry, 2021](#); [Coffey, Marx & Neumaier, 2019](#)). A notable benefit of the two-step approach is that for sparsely-occurring sounds, the detection stage can be tuned to reject the large number of ‘negative’ sound clips, with advantages for data storage/transmission, but also perhaps easing the process of training and applying the classifier, to make finer discriminations at the second step. Combined detection and classification is also feasible, and the SED and image object detection methods imported from neighbouring disciplines often include detection and classification within one NN architecture ([Kong, Xu & Plumbley, 2017](#); [Narasimhan, Fern & Raich, 2017](#); [Shrestha et al., 2021](#); [Venkatesh, Moffat & Miranda, 2021](#)).

When no labels are available even to train a classifier, unsupervised learning methods can be applied such as clustering algorithms. The use of DL directly to drive clustering is not heavily studied. A typical approach could be to use an unsupervised algorithm such as an autoencoder (an algorithm trained to compress and then decode data); and then to apply a standard clustering algorithm to the autoencoder-transformed representation of the data, on the assumption that this representation will be well-behaved in terms of clustering similar items together ([Coffey, Marx & Neumaier, 2019](#); [Ozanich et al., 2021](#)).

Signal processing using deep learning

Applications of DL have also been studied in computational bioacoustics, which do not come under the standard descriptions of classification, detection, or clustering. A theme

common to the following less-studied tasks is that they relate variously to signal processing, manipulation or generation.

Denoising and source separation are preprocessing steps that have been used to improve the quality of a sound signal before analysis, useful in difficult signal-to-noise ratio (SNR) conditions (Xie, Colonna & Zhang, 2021a). For automatic analysis, it is worth noting that such preprocessing steps are not always necessary or desirable, since they may remove information from the signal, and DL recognition may often work well despite noise. Denoising and source separation typically use lightweight signal processing algorithms, especially when used as a front-end for automatic recognition (Xie, Colonna & Zhang, 2021a; Lin & Tsao, 2020). However, in many audio fields there is a move towards using CNN-based DL for signal enhancement and source separation (Manilow, Seetharman & Salamon, 2020). Commonly, this works on the spectrogram (rather than the raw audio). Instead of learning a function that maps the spectrogram onto a classification decision, denoising works by mapping the spectrogram onto a spectrogram as output, where the pixel magnitudes are altered for signal enhancement. DL methods for this are based on denoising autoencoders and/or more recently the *u-net*, which is a specialised CNN architecture for mapping back to the same domain (Jansson et al., 2017). In bioacoustics, some work has reported good performance of DL denoising as a preprocessing step for automatic recognition, both underwater (Vickers et al., 2021b; Yang et al., 2021) and for bird sound (Sinha & Rajan, 2018).

Privacy in bioacoustic analysis is not a mainstream issue. However, Europe's General Data Protection Regulation (GDPR) drives some attention to this matter, which is well-motivated as acoustic monitoring devices are deployed in larger numbers and with increased sophistication (Le Cornu, Mitchell & Cooper, 2021). One strategy is to detect speech in bioacoustic recordings, in order to delete the respective recording clips, investigated for bee hive sound (Janetzky et al., 2021). Another is to approach the task as denoising or source separation, with speech the "noise" to remove. Cohen-Hadria et al. (2019) take this latter approach for urban sound monitoring, and to recreate the complete anonymised acoustic scene they go one step further by blurring the speech signal content and mixing it back into the soundscape. This is perhaps more than needed for most monitoring, but may be useful if the presence of speech is salient for downstream analysis, such as investigating human-animal interactions.

Data compression is another concern of relevance to deployed monitoring projects. If sound is to be compressed and transmitted back for some centralised analysis, there is a question about whether audio compression codecs will impact DL analysis. Heath et al. (2021) investigate this and concur with previous non-DL work that compression such as MP3 can have surprisingly small effect on analysis; they also obtain good performance using a CNN AudioSet embedding as a compressed 'fingerprint'. Bjorck et al. (2019) use DL more directly to optimise a codec, producing a compressed representation of elephant sounds that (unlike the fingerprint) can be decoded to the audio clip.

Synthesis of animal sounds receives occasional attention, and could be useful among other things for playback experimental stimuli. Bhatia (2021) studies birdsong synthesis

using modern DL methods, including a WaveNet and generative adversarial network (GAN) method.

Small data: data augmentation, pre-training, embeddings

The DL revolution has been powered in part by the availability of large labelled datasets. However, a widespread and persistent issue in bioacoustic projects is the *lack* of large labelled datasets: the species/calls may be rare or hard to capture, meaning not many audio examples are held; or the sound events may require a subject expert to annotate them with the correct labels (for training), and this expert time is often in short supply. Such constraints are felt for fine categorical distinctions such as those between conspecific individuals or call types, and also for large-scale monitoring in which the data volume far exceeds the person hours available. There are various strategies for effective work in such situations, including data mining and ecoacoustic methods; here I focus on techniques concerned with making DL feasible.

Data augmentation is a technique which artificially increases the size of a dataset (usually the training set) by taking the data samples and applying small irrelevant modifications to create additional data samples. For audio, this can include shifting the audio in time, adding low-amplitude noise, mixing audio files together (sometimes called ‘mixup’), or more complicated operations such as small warpings of the time or frequency axis in a spectrogram (Lasseck, 2018). The important consideration is that the modifications should not change the meaning (the label) of the data item. Importantly, in some animal vocalisations this may exclude frequency shifts. Data augmentation was in use even in 2016 at the “arrival” of DL (Goëau *et al.*, 2016), and is now widespread, used in many of the papers surveyed. Various authors study the specific combinations of data augmentation, both for terrestrial and underwater sound (Lasseck, 2018; Li *et al.*, 2021; Padovese *et al.*, 2021). Data augmentation, using the basic set of augmentations mentioned above, should be a standard part of training most bioacoustic DL systems. Software packages are available to implement audio data augmentation directly (for example *SpecAugment*, *kapre* or *audiomentions* for Python). Beyond standard practice, data augmentation can even be used to estimate the impact of confounding factors in datasets (Stowell *et al.*, 2019b).

A second widespread technique is **pretraining**: instead of training for some task by starting from a random initialisation of the NN, one starts from a NN that has previously been trained for some other, preferably related, task. The principle of “transfer learning” embodied here is that the two tasks will have some common aspects—such as the patterns of time-frequency correlations in a spectrogram, which at their most basic may have similar tendencies across many datasets—and that a NN can benefit from inheriting some of this knowledge gained from other tasks. This becomes particularly useful when large well-annotated datasets can be used for pretraining. Early work used pretraining from *image* datasets such as ImageNet, which gave substantial performance improvements even though images are quite different from spectrograms (Lasseck, 2018). Although ImageNet pretraining is still occasionally used (Disabato *et al.*, 2021; Fonseca *et al.*, 2021), many authors now pretrain using Google’s AudioSet (a diverse dataset of audio from YouTube

videos (Hershey et al., 2017; Çoban et al., 2020; Kahl et al., 2021). A similar but more recent dataset is VGG-Sound (Chen et al., 2020a), used by Bain et al. (2021). Practically, off-the-shelf networks with these well-known datasets are widely available in standard toolkits. Although publicly-available bioacoustics-specific datasets (such as that from BirdCLEF) are now large, they are rarely explored as a source of pretraining—perhaps because they are not as diverse as AudioSet/VGG-Sound, or perhaps as a matter of convenience. Ntalampiras (2018) explored transfer learning from a music genre dataset. Contrary to the experiences of others, Morgan & Braasch (2021) report that pretraining was not of benefit in their task, perhaps because the dataset was large enough in itself (150 h annotated). Another alternative is to pretrain from simulated sound data, such as synthetic underwater clicks or chirps (Glotin, Ricard & Balestrieri, 2017; Yang et al., 2021).

Closely related to pretraining is the popular and important concept of **embeddings**, and (related) **metric learning**. The common use of this can be stated simply: instead of using standard acoustic features as input, and training a NN directly to predict our labels of interest, we train a NN to convert the acoustic features into some partially-digested vector coordinates, such that this new representation is useful for classification or other tasks. The “embedding” is the space of these coordinates.

The simplest way to create an embedding is to take a pretrained network and remove the “head”, the final classification layers. The output from the “body” is a representation intermediate between the acoustic input and the highly-reduced semantic output from the head, and thus can often be a useful high-dimensional feature representation. This has been explored in bioacoustics and ecoacoustics using AudioSet embeddings, and found useful for diverse tasks (Sethi et al., 2021; Sethi et al., 2020; Çoban et al., 2020; Heath et al., 2021).

An alternative approach is to train an autoencoder directly to encode and decode items in a dataset, and then use the autoencoder’s learnt representation (from its encoder) as an embedding (Ozanich et al., 2021; Rowe et al., 2021). This approach can be applied even to unlabeled data, though it may not be clear how to ensure this encodes semantic information. It can be used as unsupervised analysis to be followed by clustering (Ozanich et al., 2021).

A third strategy for DL embedding is the use of so-called Siamese networks and triplet networks. These are not really a separate class of network architectures—typically a standard CNN is the core of the NN. The important change is the loss function: unlike most other tasks, training is not based on whether the network can correctly label a single item, but on the vector coordinates produced for a pair/triplet of items, and their distances from one another. In Siamese networks, training proceeds pairwise, with some pairs intended to be close together (e.g. same class) or far apart (e.g. different class). In triplet networks, training uses triplets with one selected as the ‘anchor’, one positive instance to be brought close, and one negative instance to be kept far away. In all cases, each of the items is projected through the NN independently, before the comparison is made. The product of such a procedure is this NN trained directly to produce an embedding in which location, or at least distance, carries semantic information. These and other

embeddings can be used for downstream tasks by applying simple classification/clustering/regression algorithms to the learnt representation. A claimed benefit of Siamese/triplet networks is that they can train relatively well with small or unbalanced datasets, and this has been reported to be the case in terrestrial and underwater projects (*Thakur et al., 2019; Nanni et al., 2020; Clementino & Colonna, 2020; Acconciaioco & Ntalampiras, 2021; Zhong et al., 2021*).

Other strategies to counter data scarcity have been investigated for bioacoustics:

- multi-task learning—another form of transfer learning, this involves training on multiple tasks simultaneously (*Morfi & Stowell, 2018; Zeghidour et al., 2021; Cramer et al., 2020*);
- semi-supervised learning, which supplements labelled data with unlabelled data (*Zhong et al., 2020b; Bergler et al., 2019a*);
- weakly-supervised learning, which allows for labelling that is imprecise or lacks detail (e.g. lacks start and end time of sound events) (*Kong, Xu & Plumbley, 2017; Knight et al., 2017; Morfi & Stowell, 2018; LeBien et al., 2020*);
- self-supervised learning, which uses some aspect of the data itself as a substitute for supervised labelling. For example, *Baevski et al. (2020)* and *Saeed, Grangier & Zeghidour (2021)* present different self-supervised learning methods to pretrain a system, for use when large amounts of audio are available but no labels. In both of these works, the pretraining process optimises a NN to determine, for a given audio recording, which of a set of short audio segments genuinely comes from that recording. This *contrastive learning* task acts as a substitute for a truly semantic task, and performs well for speech and other audio. Since this is at heart an unsupervised learning approach, with no “guidance” on which aspects of the data are of interest, it remains to be seen how well it performs in bioacoustics where the key information may be only a small part of the overall energy of the signal;
- few-shot learning, in which a system is trained across multiple similar tasks, in such a way that for a new unseen task (e.g. a new type of call to be detected) the system can perform well even with only one or very few examples of the new task (*Morfi et al., 2021b; Acconciaioco & Ntalampiras, 2021*). A popular method for few-shot learning is to create embeddings using prototypical networks, which involve a customised loss function that aims to create an embedding having good “prototypes” (cluster centroids). *Pons, Serrà & Serra (2019)* determined this to outperform transfer learning for small-data scenarios, and it is the baseline considered in a recent few-shot learning bioacoustic challenge (*Morfi et al., 2021b*).

In general, these approaches are less commonly studied, and many authors in bioacoustics use off-the-shelf pretrained embeddings. However, many of the above techniques are useful to enable training despite dataset limitations; hence, they can themselves be used in creating embeddings, and could be part of future work on creating high-quality embeddings.

Generalisation and domain shift

Concern about whether the a DL system's quality of performance will generalise to new data is a widespread concern, especially when small datasets are involved. A more specific concern is whether performance will generalise to new conditions *in which attributes of the input data have changed*: for example changes in the background soundscape, the sub-population of a species, the occurrence frequency of certain events, or the type of microphone used. All of these can change the overall distribution of basic acoustic attributes, so-called *domain shift*, which can have undesirable impacts on the outputs of data-driven inference ([Morgan & Braasch, 2021](#)).

It is increasingly common to evaluate DL systems, not only on a test set which is kept separate from the training data, but also on test set(s) which differ in some respects from the training data, such as location, SNR, or season ([Shiu et al., 2020](#); [Vickers et al., 2021b](#); [Çoban et al., 2020](#); [Allen et al., 2021](#); [Khalighifar et al., 2021](#)). This helps to avoid the risk of overestimating generalisation performance in practice.

Specific DL methods can be used explicitly to account for domain shift. *Domain adaptation* methods may automatically adapt the NN parameters ([Adavanne et al., 2017](#); [Best et al., 2020](#)). Explicitly including contextual correlates as input to the NN is an alternative strategy for automatic adaptation ([Lostanlen et al., 2019b](#), [Roch et al., 2021](#)). Where a small amount of human input about the new domain is possible, fine-tuning (limited retraining) or active learning (interactive feedback on predictions) have been explored ([Çoban et al., 2020](#); [Allen et al., 2021](#); [Ryazanov et al., 2021](#)). [Stowell et al. \(2019b\)](#) designed a public “bird audio detection” challenge specifically to stimulate the development cross-condition (cross-dataset) generalisable methods. In that challenge, however, the leading submissions did not employ explicit domain adaptation, instead relying on the implicit generality of transfer learning (pretraining) from general-purpose datasets, as well as data augmentation to simulate diverse conditions during training.

Open-set and novelty

One problem with the standard recipe (and in fact many ML methods) is that by default, recognition is limited to a pre-specified and fixed set of labels. When recording in the wild, it is surely possible to encounter species or individuals not accounted for in the training set, which should be identified. This is common *e.g.* for individual ID ([Ptacek et al., 2016](#)). Detecting new sound types beyond the known set of target classes is referred to as *open set* recognition, perhaps related to the more general topic of novelty detection which aims to detect any novel occurrence in data. [Cramer et al. \(2020\)](#) argue that hierarchical classification is useful for this, in that a sound may be strongly classified to a higher-level taxon even when the lower-level class is novel. [Ntalampiras & Potamitis \(2021\)](#) apply novelty detection based on a CNN autoencoder (an algorithm trained to compress and then recreate data). Since the method is trained to reconstruct the training examples with low error, the authors use the assumption that novel sounds will be reconstructed with high error, and thus use this as a trigger for detecting novelty.

More broadly, the aforementioned topic of embeddings offers a useful route to handling open-set classification. A good embedding should provide a somewhat semantic

representation even of new data, such that even novel classes will cluster well in the space (standard clustering algorithms such as k-nearest neighbours can be applied). This is advocated by [Thakur et al. \(2019\)](#), using triplet learning, and later [Acconciaco & Ntalampiras \(2021\)](#) using Siamese learning. Novelty and open-set issues are likely to be an ongoing concern, in practice if not in theory, though the increasing popularity of general-purpose embeddings indeed offers part of the solution.

Context and auxiliary information

Deep learning implementations almost universally operate on segments of audio or spectrogram (e.g. 1 or 10 s per datum) rather than a continuous data stream. This is true even for RNNs which in theory can have unbounded time horizons. Yet it is clear from basic considerations that animal vocalisations, and their accurate recognition, may depend strongly on contextual factors originating outside a short window of temporal attention, whether this be prior soundscape activity or correlates such as date/time, location or weather.

[Lostanlen et al. \(2019b\)](#) add a “context-adaptive neural network” layer to their CNN, whose weights are dynamically adapted at prediction time by an auxiliary network taking long-term summary statistics of spectrotemporal features as input. Similarly, [Roch et al. \(2021\)](#) input acoustic context to their CNN based on estimates of the local signal-to-noise ratio (SNR). [Madhusudhana et al. \(2021\)](#) apply a CNN (DenseNet) to acoustic data, and then postprocess the predictions of that system using an RNN, to incorporate longer-term temporal context into the final output. Note that this CNN and RNN is not an integrated CRNN but two separate stages, with the consequence that the RNN can be applied over differing (longer) timescales than the CNN.

Animal taxonomy is another form of contextual information which may help to inform or constrain inferences. Although taxonomy is rarely the strongest determinant of vocal repertoire, it may offer a partial guide. Hierarchical classification is used in many fields, including bioacoustics; [Cramer et al. \(2020\)](#) propose a method that explicitly encodes taxonomic relationships between classes into the training of a CNN, evaluated using bird calls and song. [Nolasco & Stowell \(2022\)](#) propose a different method, and evaluate across a broader hierarchy, covering multiple taxa at the top level and individual animal identity at the lowest level.

Perception

The work so far discussed uses DL as a practical tool. Deep learning methods are loosely inspired by ideas from natural perception and cognition ([LeCun, Bengio & Hinton, 2015](#)), but there is no strong assumption that bioacoustic DL implements the same processes as natural hearing. Further, since current DL models are hard to interpret, it would be hard to validate whether or not that assumption held.

Even so, a small stream of research aims to use deep learning to model animal acoustic perception. DL can model highly non-linear phenomena, so perhaps could replicate many of the subtleties of natural hearing, which simpler signal processing models do not. Such models could then be studied or used as a proxy for animal judgment.

Morfi, Lachlan & Stowell (2021a) use triplet loss to train a CNN to produce the same decisions as birds in a two-alternative forced-choice experiment. *Simon et al. (2021)* train a CNN from sets of bat echolocation call reflections, to classify flowers as bat-pollinated or otherwise—a simplified version of an object recognition task that a nectarivorous/frugivorous bat presumably solves. *Francl & McDermott (2020)* study sound localisation, finding that a DL trained to localise sounds in a (virtual) reverberant environment exhibits some phenomena known from human acoustic perception.

On-device deep learning

Multiple studies focus on how to run bioacoustic DL on a small hardware device, for affordable/flexible monitoring in the field. Many projects do not need DL running in real-time on device: they can record audio to storage or transmit it to a base station, for later processing (*Roe et al., 2021; Heath et al., 2021*). However, implementing DL on-device allows for live readouts and rapid responses (*Mac Aodha et al., 2018*), potential savings in power or data transmission costs, and enables some patterns of deployment that might not otherwise be possible. One benefit of wide interest might be to perform a first step of detection/filtering and discard many hours of uninformative audio, to extend deployment durations before storage is full and reduce transmission bandwidths: this is traditionally performed with simple energy detection, but could be enhanced by lightweight ML algorithms, perhaps similar to “keyword spotting” in domestic devices (*Zhang et al., 2017*).

The Raspberry Pi is a popular small Linux device, and although low-power it can have much of the functionality of a desktop computer, such as running Python or R scripts; thus the Raspberry Pi has been used for acoustic monitoring and other deployments (*Jolles, 2021*). Similar devices are NVIDIA Jetson Nano and Google Coral (the latter carries a TPU unit on board intended for DL processing). *Zualkernan et al. (2021)* evaluate these three for running a bat detection algorithm on-device.

Even more constrained devices offer lower power consumption (important for remote deployment powered by battery or solar power), lower ecological footprint, and smaller form factor; often based on the ARM Cortex-M family of processors. The AudioMoth is a popular example (*Hill et al., 2018*). It is too limited to run many DL algorithms; however *Prince et al. (2019)* were able to implement a CNN (depthwise-separable to reduce the complexity), applied to mel frequency features, and report that it outperformed a hidden Markov model detector on-device, although “at the cost of both size and speed”: it was not efficient enough to run in real-time on AudioMoth.

Programming frameworks help to make such low-level implementations possible: *Disabato et al. (2021)* use ARM CMSIS (Common Microcontroller Software Interface Standard—a software interface) to implement a bird detector, and *Zualkernan et al. (2021)* use TensorFlow Lite to implement a bat species classifier. As in the more general case, off-the-shelf NN architectures can be useful, including MobileNet and SqueezeNet which are designed to be small/efficient (*Vidaña-Vila et al., 2020*). However, all three of the bioacoustic studies just mentioned, while inspired by these, implemented their own CNN designs and feature modifications to shrink the footprint even further.

Small-footprint device implementations offer the prospect of DL with reduced demands for power, bandwidth and storage. However, [Lostanlen et al. \(2021b\)](#) argue that energy efficiency is not enough, and that fundamental resource requirements such as the rare minerals required for batteries are a constraint on wider use of computational bioacoustic monitoring. They propose batteryless acoustic sensing, using novel devices capable of intermittent computing whenever power becomes available. It remains to be seen whether this intriguing proposal can be brought together with the analytical power of DL.

Workflows and other practicalities

As DL comes into increased use in practice, questions shift from the proof-of-concept to the integration into broader workflows (e.g. of biodiversity monitoring), and other practicalities. Many of the issues discussed above arise from such considerations. Various authors offer recommendations and advice for ecologists using DL ([Knight et al., 2017](#); [Rumelt, Basto & Roncal, 2021](#); [Maegawa et al., 2021](#)). Others investigate integration of a CNN detector/classifier into an overall workflow including data acquisition, selection and labeling ([LeBien et al., 2020](#); [Morgan & Braasch, 2021](#); [Ruff et al., 2021](#)). [Brown, Montgomery & Garg \(2021\)](#) go further and investigate the automation of designing the overall workflow, arguing that “[t]here is merit to searching for workflows rather than blindly using workflows from literature. In almost all cases, workflows selected by [their proposed] search algorithms (even random search, given enough iterations) outperformed those based on existing literature.”

One aspect of workflow is the user interface (UI) through which an algorithm is configured and applied, and its outputs explored. Many DL researchers provide their algorithms as Python scripts or suchlike, a format which is accessible by some but not by all potential users. Various authors provide graphical user interfaces (GUIs) for the algorithms they publish, and to varying extents study efficient graphical interaction ([Jiang et al., 2019](#); [Coffey, Marx & Neumaier, 2019](#); [Steinfath et al., 2021](#); [Ruff et al., 2021](#)).

A ROADMAP FOR BIOACOUSTIC DEEP LEARNING

I next turn to the selection of topics that are unresolved and/or worthy of further development: recommended areas of focus in the medium-term for research in deep learning applied within computational bioacoustics. The gaps are identified and confirmed through the literature survey, although there will always be a degree of subjectivity in the thematic synthesis.

Let us begin with some principles. Firstly, AI does not replace expertise, even though this may be implied by the standard recipe and general approach (i.e. using supervised learning to reproduce expert labels). Instead, through DL we train sophisticated but imperfect agents, with differing sets of knowledge. For example, a bird classifier derived from an AudioSet embedding may have one type of expertise, while a raw waveform system trained from scratch has a different expertise. As the use of these systems becomes even more standardised, they take on the role of expert peers, with whom we consult and debate. The move to active learning, which deserves more attention, cements this role by allowing DL agents to learn from criticism of their decisions. Hence, DL does not

displace the role of experts, nor even of crowdsourcing; future work in the field will integrate the benefits of all three (*Kitzes & Schricker, 2019*). Secondly, open science is a vital component of progress. We have seen that the open publication of datasets, NN architectures, pretrained weights, and other source code has been crucial in the development of bioacoustic DL. There is a move toward open sharing of data, but in bioacoustics this is incomplete (*Baker & Vincent, 2019*). Sharing audio and metadata, and the standardisation of metadata, will help us to move far beyond the limitations of single datasets.

Maturing topics? Architectures and features

Let us also briefly revise core topics within bioacoustic DL that are frequently discussed, but can be considered to be maturing, and thus not of high urgency.

The vast majority of the surveyed work uses spectrograms or mel spectrograms as the input data representation. Although some authors raise the question of whether a species-customised spectrogram should be more appropriate than the human-derived mel spectrogram, for many tasks such alterations are unlikely to make a strong difference: as long as the spectrogram represents sufficient detail, and a DL algorithm can reasonably be trained, the empirical performance is likely to be similar across many different spectrogram representations. Preprocessing such as noise reduction and PCEN is often found to be useful and will continue to be applied. Methods based on raw waveforms, or adaptive front-ends such as SincNet or LEAF, are certainly of interest, and further exploration of these in bioacoustics is anticipated. They may be particularly useful for tasks requiring fine-grained distinctions.

Commonly-used acoustic “features” in future are likely to include off-the-shelf deep embeddings, even more commonly than now. Whether the input to those features is waveform or spectrogram will be irrelevant to users. AudioSet and VGG-Sound are the most commonly-used datasets for such pretraining; note however that these cannot cover all bioacoustic needs—for example ultrasound—and so it seems likely that bioacoustics-specific embeddings will be useful in at least some niches.

CNNs have become dominant in many applications of DL, and this applies to bioacoustic DL too. The more recent use of one-dimensional temporal convolutions (TCNs) is likely to continue, because of their simplicity and relative efficiency. However, looking forward it is not in fact clear whether CNNs will retain their singular dominance. In NLP and other domains, NN architectures based on “attention” (transformers/ perceivers, discussed above) have displaced CNN as a basic architecture. CNNs fit well with waveform and spectrogram data, and thus are likely to continue to contribute to NN architectures for sound, perhaps combined with transformer layers (cf. *Baevski et al. (2020)*). For example, *Wolters et al. (2021)* propose to address sound event detection by using a CNN together with a perceiver network: their results imply that a perceiver is a good way to process variable-length spectrogram data into per-event summary representations.

A similar lesson applies to RNNs, except that RNNs have a more varied history of popularity. Recent CRNNs make good use of recurrent layers; but TCNs seem to threaten

to displace them. I suggest that although RNNs embody a very general idea about sequential data, they are a special case of more general *computation with memory*. Transformers and other attention NNs show a different approach to allowing a computation to refer back to previous time steps flexibly. (See also the Legendre memory unit explored by [Gupta et al. \(2021\)](#).) All are special cases, and future work in DL may move more towards the general goal of differentiable neural computing ([Graves, Wayne & Danihelka, 2014](#)). The fluctuating popularity of recurrence and attention depends on their convenience and reusability as efficient modules in a DL architecture, and the future of DL-with-memory is likely to undergo many changes. Computational bioacoustics will continue to use these and integrate short- and long-term memory with other contextual data.

Learning without large datasets

Bioacoustics in general will benefit from the increasing open availability of data. However, this does not dissipate the oft-studied issue of small data: project-specific recognition tasks will continue to arise, including high-resolution discrimination tasks, and tasks for which transfer learning is inappropriate (*e.g.* due to the risk of bias) ([Morfi, Lachlan & Stowell, 2021a](#)). Many approaches to dealing with small datasets have been surveyed in the preceding text; important for future work is for these approaches to be integrated together, and for their advantages and disadvantages to be clarified. As shown in data challenges such as BirdCLEF and DCASE, pre-training, embeddings, multi-task learning and data augmentation all offer low-risk methods for improved generalisation.

Few-shot learning is a recent topic of interest; it is not clear whether it will continue long-term to be a separate “task” or will integrate with wider approaches, but it reflects a common need in bioacoustic practice. Active learning (AL) is also a paradigm of recent interest, and of high importance. It moves beyond the basic non-interactive model of most machine learning, in which a fixed training set is the only information available to a classifier. In AL, there is a human-machine interaction of multiple iterations, in which (some) predictions from a system are shown to a user for feedback, and the user’s feedback about correct and mistaken identifications is fed into the next round of optimisation of the algorithm. I identify AL as high importance because it offers a principled way to make the most efficient use of a person’s time in labelling or otherwise interacting with a system ([Qian et al., 2017](#)). It can be a highly effective way to deal with large datasets, including domain shift and other issues. It has been used in bioacoustic DL ([Steinfath et al., 2021](#); [Allen et al., 2021](#)) but is under-explored, in part because its interactive nature makes it slightly more complex to design an AL evaluation. It may be that future work will use something akin to few-shot learning as the first step in an AL process.

A very different approach to reduce the dependence on large data is to create entirely simulated datasets that can be used for training. This is referred to as *sim2real* in DL, and its usefulness depends on whether it is feasible to create good simulations of the phenomena to be analysed. It goes beyond data augmentation in generating new data points rather than modifying existing ones. It may thus be able to generate higher diversity of training data, at a cost of lower realism. One notable advantage of *sim2real* is that

any confounds or biases in the training data can be directly controlled. Simulated datasets have been explored in training DL detectors of marine sounds, perhaps because this class of signals can be modelled using chirp/impulse/sinusoidal synthesis ([Glotin, Ricard & Balestrieri, 2017](#); [Yang et al., 2021](#); [Li et al., 2020](#)). Simulation is also especially relevant for spatial sound scenes, since the spatial details of natural sound scenes are hard to annotate ([Gao et al., 2020](#); [Simon et al., 2021](#)). Simulation, often involving composing soundscapes from a library of sound clips, has been found useful in urban and domestic sound analysis ([Salamon et al., 2017b](#); [Turpault et al., 2021](#)). Such results imply that wider use in bioacoustic DL may be productive, even when simulation of the sound types in question is not perfect.

Equal representation

Deep learning systems are well-known to be powerful but with two important weaknesses: (1) in most cases they must be treated as ‘black boxes’ whose detailed behaviour in response to new data is an empirical question; (2) they can carry a high risk of making biased decisions, which usually occurs because they faithfully reproduce biases in the training data ([Koenecke et al., 2020](#)). Our concern here is to create DL systems that can be a *reliable* guide to animal vocalisations, especially if used as a basis for conservation interventions. Hence we should ensure that the tools we create lead to an equal representation in terms of their sensitivity, error rates, *etc.* ([Hardt et al., 2016](#)).

[Baker & Vincent \(2019\)](#) point out that research output in bioacoustics is strongly biased: its taxonomic balance is unrepresentative of the audible animal kingdom, whether considered in terms of species diversity, biomass, or conservation importance. The same is true in the sub-field of DL applied to bioacoustics, both for datasets and research papers (see taxa listed above). Baker advocates for further attention to insect sound, and insects are recognised more broadly as under-studied ([Montgomery et al., 2020](#)); [Linke et al. \(2018\)](#) make a related case for freshwater species.

Equal representation (taxonomic, geographic, *etc.*) can be inspected in a dataset, and we should make further efforts to join forces and create more diverse open datasets, covering urban and remote locations, rich and poor countries. [Baker & Vincent \(2019\)](#) argue that the general field of bioacoustic research suffers from a lack of data deposition, with only 21% of studied papers publishing acoustic recordings for others to use. Addressing this gap in open science practice may in fact be our most accessible route to better coverage in acoustic data.

Equal representation should also be evaluated in feature representations such as widely-used embeddings. The representational capacities of an embedding derive from the dataset, the NN architecture and the training regime, and any of these factors could introduce biases that represent some acoustic behaviours better than others.

Beyond equal representation, it may of course remain important to develop targeted methods, such as those targeted at rare species ([Znidarsic et al., 2020](#); [Wood et al., 2021](#)). Since rare occurrences are intrinsically difficult to create large datasets for, this is worthy of further study. This review lists many methods that may help when rare species are of interest, but the best use of them is not yet resolved. To give examples beyond the

bioacoustic literature: [Beery et al. \(2020\)](#) explore the use of synthetic examples for rarely-observed categories (in camera trap images); and [Baumann et al. \(2020\)](#) consider frameworks for evaluating rare sound event detection.

Interfaces and visualisation

Many bioacoustic DL projects end with their outputs as custom Python scripts: this is good practice in computer science/DL, for reproducibility, but not immediately accessible to a broad community of zoologists/conservationists. User interfaces (UIs) are a non-trivial component in bridging this gap. Since the potential users of DL may wish to use it *via* R, Python, desktop apps, smartphone apps, or websites, there remains no clear consensus on what kinds of UI will be most appropriate for bioacoustic DL, besides the general wish to integrate with existing audio editing/annotation tools. It seems likely that in future many algorithms will be available as installable packages or web APIs, and accessed variously through R/Python/desktop/*etc.* as preferred. Some existing work creates and even evaluates interfaces (discussed above, Workflow section), but more work on this is merited, including (a) research on efficient human-computer interaction for bioacoustics, and (b) visualisation tools making use of large-scale DL processing (cf. [Kholghi et al., 2018](#); [Znidarsic et al., 2020](#); [Phillips, Towsey & Roe, 2018](#)).

One domain in which user interaction is particularly important is active learning (AL), since it involves an iterative human-computer interaction. The machine learning components in AL can be developed without UI work, but interaction with sound data has idiosyncratic characteristics (temporal regions, spectrograms, simultaneously-occurring sounds) which suggest that productive bioacoustic AL will involve UI designs that specifically enhance this interaction.

Beyond human-computer interaction is animal-computer interaction, for example using robotic animal agents in behavioural studies ([Simon et al., 2019](#); [Slonina et al., 2021](#)). These studies offer the prospect of new insights about animal behaviour, and they might use DL in future to provide sophisticated vocal interaction.

The most common formulation of DL tasks, *via* fixed sets of training data and evaluation data, become less relevant when considering active learning and other interactive situations. There will need to be further consideration of the format, for example of data-driven challenges, and potentially DL techniques such as reinforcement learning (not reviewed here since not used in the current literature) ([Teşileanu, Ölveczky & Balasubramanian, 2017](#)).

Under-explored machine learning tasks

The following tasks are known in the literature, but according to the present survey are not yet mature, and also worthy of further work because of their importance or generality.

Individual ID

Automatically recognising/discriminating between individual animals has been addressed by many studies in bioacoustics, whether for understanding animal communication or for monitoring/censusing animals ([Ptacek et al., 2016](#); [Vignal, Mathevon & Mottin, 2008](#); [Searby, Jouventin & Aubin, 2004](#); [Linhart et al., 2019](#); [Adi, Johnson & Osiejuk, 2010](#); [Fox,](#)

2008; Beecher, 1989). Acoustic recognition of individuals can be a non-invasive replacement for survey techniques involving physical capture of individuals; it thus holds potential for improved monitoring with lower disturbance of wild populations. Thus far DL has only rarely been applied to individual ID in acoustic surveying, though this will surely change (Ntalampiras & Potamitis, 2021; Nolasco & Stowell, 2022). For some species, inter-individual distinctions may be as simple as a difference in fundamental frequency, and thus addressable using a simple algorithm; however, in general the attraction for DL here comes from its capacity to distinguish highly nonlinear patterns and its potential to generalise well to out-of-sample data. Within-species acoustic differences between individuals are typically fine-scale differences, requiring finer distinctions than species distinctions. This makes the task harder than species classification.

I suggest that these characteristics make the task of general-purpose automatic discrimination of individual animals, a useful focus for DL development. A DL system that can address this task usefully is one that can make use of diverse fine acoustic distinctions. Its inferences will be of use in ethology as well as in biodiversity monitoring. Cross-species approaches and multi-task learning can help to bridge bioacoustic considerations across the various taxon groups commonly studied (Nolasco & Stowell, 2022). A complete approach to individual recognition would also handle the open-set issue well, since novel individuals may often be encountered in the wild. There are not many bioacoustic datasets labelled with individual ID, and increased open data sharing can help with this.

Sound event detection and object detection

For many reasons it can be useful to create a detailed “transcript” of the sound events within a recording, going beyond basic classification. As with individual ID, this more detailed analysis can feed into both ethological and biodiversity analyses; its development goes hand-in-hand with higher-resolution bioacoustic DL.

As described in the earlier discussion of detection, there are at least two main approaches to this in existing literature. One version of SED (Fig. 1B) follows the same model as automatic music transcription or speaker diarisation in other domains, and uses similar DL architectures to solve the problem (Mesaros et al., 2021; Morfi & Stowell, 2018; Morfi et al., 2021b). An alternative approach inherits directly from object detection based on bounding boxes in images (Fig. 1C). This fits well when data annotations are given as time-frequency bounding boxes drawn on spectrograms. Solutions typically adapt well-known image object detector architectures such as YOLO and Faster R-CNN, which are quite different from the architectures used in other tasks (Venkatesh, Moffat & Miranda, 2021; Shrestha et al., 2021; Zsebök et al., 2019; Coffey, Marx & Neumaier, 2019). These two approaches each have their advantages. For example, frequency ranges in sound events can sometimes be useful information, but can sometimes be ill-defined/unnecessary, and not present in many datasets. Future work in bioacoustic DL should take the best of each paradigm, perhaps with a unified approach that can be applied whether or not frequency bounds are included in the data about a sound event.

Spatial acoustics

On a fine scale, the spatial arrangement of sound sources in a scene can be highly informative, for example in attributing calls to individuals and/or counting individuals correctly. It can also be important for behavioural and evolutionary ecological analysis (Jain & Balakrishnan, 2011). Spatial location can be handled using multi-microphone arrays, including stereo or ambisonic microphones. It is often analysed in terms of the direction-of-arrival (DoA) and/or range (distance) relative to the sensor. Taken together, the DoA and the range imply the Cartesian location; but either of them can be useful on its own.

The standard approach to spatial analysis uses signal processing algorithms, even when the data are later to be classified using machine learning (Kojima *et al.*, 2018). However, this may change. For example, Houégnigan *et al.* (2017) train an MLP and Van Komen *et al.* (2020) a CNN, to estimate the range (distance) of underwater synthetic sound events. Yip *et al.* (2019) perform a similar task using terrestrial recordings: using calibrated microphone recordings of two bird species, they obtain useful estimates of distance by deep learning regression from the sound level measurements. In other domains of acoustics *e.g.* speech and urban sound, there is already a strong move to supplant signal processing with DL for spatial tasks (Hammer *et al.*, 2021; Adavanne *et al.*, 2018; Shimada *et al.*, 2021). Important to note is that these works usually deal with *indoor* sound. Indoor and outdoor environments have very different acoustic propagation effects, meaning that the generalisation to outdoor sound may not be trivial (Traer & McDermott, 2016).

Spatial inference can also be combined with SED (*e.g.* in the DCASE challenge “sound event localisation and detection” task or SELD), combining the two-step process (*e.g.* Kojima *et al.*, 2018) into a single joint estimation task (Shimada *et al.*, 2021).

It is clear that many bioacoustic datasets and research questions will continue to be addressed in a spatially-agnostic fashion. Although some spatial attributes such as distance can be estimated from single-channel recordings (as above), multi-channel audio is usually required for robust spatial inference. Spatial acoustic considerations are quite different in terrestrial and marine sound, and more commonly considered in the latter. However, the development of DL tasks such as distance estimation and SELD (in parallel to SED) could benefit bioacoustics generally, with local spatial information used more widely in analysis.

The discussion thus far does not address the broader geographic-scale distribution of populations, which statistical ecologists may estimate from observations. Although machine observations will increasingly feed into such work, the use of DL in statistical ecology is outside the scope of this review (but cf. Kitzes & Schrickler, 2019).

Useful integration of outputs

As DL becomes increasingly used in practice, there will inevitably be further work on integrating it into practical workflows, discussed earlier. However, there are some gaps to be bridged, worthy of specific attention.

An important issue is the calibration of the outputs of automatic inference. [Kitzes & Schricker \(2019\)](#) state the problem:

“We wish to specifically highlight one subtler challenge, however, which we believe is substantially hindering progress: the need for better approaches for dealing with uncertainty in these indirect observations. [...] First, machine learning classifiers must be specifically designed to return probabilistic, not binary, estimates of species occurrence in an image or recording. Second, statistical models must be designed to take this probabilistic classifier output as input data, instead of the more usual binary presence–absence data. The standard statistical models that are widely used in ecology and conservation, including generalized linear mixed models, generalized additive models and generalized estimating equations, are not designed for this type of input.” ([Kitzes & Schricker, 2019](#))

In fact, although many ML algorithms do output strict binary decisions, DL classifiers/detectors do not: they output numerical values between zero and one, which we can interpret as probabilities, or convert into binary decisions by thresholding. However, the authors’ first point does not disappear since the outputs from DL systems are not always well-calibrated probabilities: they may under- or over-confident, depending on subtleties of how they have been trained (such as regularisation) ([Niculescu-Mizil & Caruana, 2005](#)). This does not present an issue when evaluating DL by standard metrics, but becomes clear when combining many automatic detections to form abundance estimates. DL outputs, interpreted as probabilities, may be under- or over-confident, or biased in favour of some categories and against others. Measuring (mis)calibration is the first step, and postprocessing the outputs can help ([Niculescu-Mizil & Caruana, 2005](#)). Evaluating systematic biases is also important: DL can be expected to exhibit higher sensitivity towards sounds well-represented in its training data, and this has been seen in practice ([Lostanlen et al., 2018](#)). Birdsong species classifiers are strongest for single-species recordings, and even with current DL they show reduced performance in denser soundscape recordings—an important concern given that much birdsong is heard in dense dawn choruses ([Joly et al., 2019](#)). Evaluating and improving upon these biases is vital.

The spatial reliability of detection is one particular facet of this. For manual surveys, there is well-developed statistical methodology to measure how detection probability relates to the distance to the observer, and how this might vary with species and habitat type ([Johnston et al., 2014](#)). The same must be applied to automatic detectors. We have an advantage of reproducibility: we can assume that distance curves and calibration curves for a given DL algorithm, analysing audio from a given device model, will be largely consistent. Thus such measurements applied to a widely-used DL algorithm and recording device would be widely useful. Some work does evaluate the performance of bioacoustic DL systems and how they degrade over distance ([Maegawa et al., 2021](#); [Lostanlen et al., 2021a](#)). This can be developed further, in both simulated and real acoustic environments.

Under the model of detection as binary classification, our observations are “occupancy” (presence/absence) measurements. These can be used to estimate population distributions,

but are less informative than observed abundances of animals. Under the more detailed models of detection, we can recover individual calls/song bouts and then count them, though of course these do not directly reflect the number of animals unless we can use calling-rate information collected separately (Stevenson *et al.*, 2015). Routes toward bridging this gap using DL include applying “language models” of vocal sequences and interactions; the use of spatial information to segregate calls per individual; and direct inference of animal abundance, skipping the intermediate step of call detections. Counting and density estimation using DL has been explored for image data (e.g. Arteta, Lempitsky & Zisserman, 2016), a kind of highly nonlinear regression task. Early work explores this for audio, using a CNN to predict the numbers of targeted bird/anuran species in a sound clip (Dias, Ponti & Minghim, 2021). Sethi *et al.* (2021) suggest that regression directly from deep acoustic embedding features to species relative density can work well, especially for common species with strong temporal occurrence patterns.

As DL tools become integrated into various workflows, the issue of standardised data exchange becomes salient. Standards bodies such as Biodiversity Information Standards (“TDWG”) provide guidance on formats for biodiversity data exchange, including those for machine observations (<https://www.tdwg.org/>). These standards are useful, but may require further development: for example, the probabilistic rather than binary output referenced by Kitzes & Schriker (2019) needs to be usefully represented. The attribution of observations to a specific algorithm (trained using specific datasets...) requires a refinement of the more conventional attribution metadata schemes used for named persons. Such attribution can perhaps already be represented by standards such as the W3C Provenance Ontology, though such usage is not widespread (<https://www.w3.org/TR/prov-o/>). An explicit machine-readable representation of such provenance relationships would greatly simplify tasks such as merging of multiple machine-observation datasets, or investigating the causal effects of algorithmic and training-data choices on deployed systems.

Taken together, these integration-related technical topics are important for closing the loop between bioacoustic monitoring, data repositories, policy and interventions. They are thus salient for bringing bioacoustic DL into full service to help address the biodiversity crisis. The international body for biodiversity is the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services (IPBES). The IPBES recently raised the alarm about the unprecedented and accelerating species extinction rates; among other things it called for enhanced environmental monitoring and evaluation (IPBES, 2019). For computational bioacoustics to address such global challenges fully, we must focus openly and collaboratively on integration work.

Behaviour and multi-agent interactions

Animal behaviour research (ethology) can certainly benefit from automatic detection of vocalisations, for intra- and inter-species vocal interactions and other behaviour. This will increasingly make use of SED/SELD/object-detection to transcribe sound scenes. Prior ethology works have used correlational analysis, Markov models and network analysis, though it is difficult to construct general-purpose data-driven models of vocal sequencing

(*Kershenbaum et al., 2014; Stowell, Gill & Clayton, 2016a*). Deep learning offers the flexibility to model multi-agent (*i.e.* multi-animal) sound scenes and interactions, with recent work including neural point process models that may offer new tools (*Xiao et al., 2019; Chen, Amos & Nickel, 2020b*).

Ethology is not the only reason to consider (vocal) behaviour in DL. The modelling just mentioned is analogous to the so-called “language model” that is typically used in automatic speech recognition (ASR) technology: when applied to new sound recordings, it acts as a prior on the temporal structure of sound events, which helps to disambiguate among potential transcriptions (*O’Shaughnessy, 2003*). This structural prior is missing in most approaches to acoustic detection/classification, which often implies that each sound event is assumed to occur with conditional independence from others. Note that language modelling in ASR considers only one voice. A grand challenge in bioacoustic DL could be to construct DL “language models” that incorporate flexible, open-set, agent-based models of vocal sequences and interactions; and to integrate these with SED/SELD/object-detection methods for sound scene transcription. Note that SED/SELD/object-detection paradigms will also need to be improved: for example the standard approach to SED is not only closed-set, but does not transcribe overlapping sound events within the same category (*Stowell & Clayton, 2015*). Analogies with natural sound scene parsing may help to design useful approaches (*Chait, 2020*).

Low impact

When advocating for computational work that might be large-scale or widely deployed, we have a duty to consider the wider impacts of deploying such technology: carbon footprint, and resource usage (*e.g.* rare earth minerals and e-waste considerations of electronics). *Lostanlen et al. (2021b)* offer a very good summary of these considerations in bioacoustic monitoring hardware, as well as a novel proposition to develop batteryless bioacoustic devices.

For DL, impacts are incurred while training a NN, and while applying it in practice: their relative significance depends on whether training or inference will be run many times (*Henderson et al., 2020*). Happily, the power (and thus carbon emission) impacts of training DL can be reduced through some of the techniques that are already in favour for cross-task generalisation: using pretrained networks rather than starting training from random intialisation, and using pretrained embeddings as fixed feature transformations. Data augmentation during training can *increase* power usage by artificially increasing the training set size, but this can be offset if the required number of training epochs is reduced. The increasing trend of NN architectures designed to have a low number of parameters or calculations (ResNet, MobileNet, EfficientNet) also helps to reduce the power intensity.

The question of running DL algorithms on-device brings further interesting resource tradeoffs. Small devices may be highly efficient, and might necessarily run small-footprint NNs. (Note that a given algorithm may have differing footprint when run as a Python script on a general-purpose device, *vs* a low-level implementation for a custom device). Running DL on-device also offers the ability to reduce storage and/or communication

overheads, by discarding irrelevant data at an early stage. Alternatively, in many cases it may be more efficient to use fixed recording schedules and analyse data later in batches (Dekkers *et al.*, 2022). The question becomes still more complex when considering networking options such as GSM/LoRa or star- vs mesh-networking.

Our domain has only just begun to spell out these factors coherently. Smart bioacoustic monitoring has potential to provide rapid-response ecosystem monitoring, in support of nature-based solutions in climate change and biodiversity. This motivates further development in low-impact bioacoustic DL paradigms.

CONCLUSIONS

In computational bioacoustics, as in other fields, DL has enabled a leap in the performance of automatic systems. Bioacoustics will continue to benefit from wider developments in DL, including methods adapted from image recognition, speech and general audio. It will continue to be driven by the availability of data, developments in audio hardware and processing power, and the demands of national and international biodiversity accounting. However, we cannot respond to these drivers merely by adopting techniques from neighbouring fields. The roadmap presented here identifies topics meriting study within bioacoustics, arising from the specific characteristics of the data and questions we face.

ADDITIONAL INFORMATION AND DECLARATIONS

Funding

The author received no funding for this work.

Competing Interests

Dan Stowell is an Academic Editor for PeerJ.

Author Contributions

- Dan Stowell conceived and designed the experiments, performed the experiments, analyzed the data, prepared figures and/or tables, authored or reviewed drafts of the paper, and approved the final draft.

Data Availability

The following information was supplied regarding data availability:

This is a review paper and there is no raw data.

Supplemental Information

Supplemental information for this article can be found online at <http://dx.doi.org/10.7717/peerj.13152#supplemental-information>.

REFERENCES

- Abeßer J. 2020. A review of deep learning based methods for acoustic scene classification. *Applied Sciences* 10(6):2020 DOI 10.3390/app10062020.

- Acconciaco M, Ntalampiras S. 2021.** One-shot learning for acoustic identification of bird species in non-stationary environments. In: *2020 25th International Conference on Pattern Recognition (ICPR)*. Piscataway: IEEE, 755–762.
- Adavanne S, Drossos K, Çakr E, Virtanen T. 2017.** Stacked convolutional and recurrent neural networks for bird audio detection. In: *Proceedings of EUSIPCO 2017. Special Session on Bird Audio Signal Processing*, 1729–1733.
- Adavanne S, Politis A, Virtanen T, Adavanne S, Politis A, Virtanen T. 2018.** Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network. In: *2018 26th European Signal Processing Conference (EUSIPCO)*. Piscataway: IEEE, 1462–1466.
- Adi K, Johnson MT, Osiejuk TS. 2010.** Acoustic censusing using automatic vocalization classification and identity recognition. *Journal of the Acoustical Society of America* 127(2):874–883 DOI 10.1121/1.3273887.
- Allen AN, Harvey M, Harrell L, Jansen A, Merckens KP, Wall CC, Cattiau J, Oleson EM. 2021.** A convolutional neural network for automated detection of humpback whale song in a diverse, long-term passive acoustic dataset. *Frontiers in Marine Science* 8:165 DOI 10.3389/fmars.2021.607321.
- Arteta C, Lempitsky V, Zisserman A. 2016.** Counting in the wild. In: *European Conference on Computer Vision*, Berlin: Springer, 483–498.
- Baevski A, Zhou Y, Mohamed A, Auli M. 2020.** wav2vec 2.0: a framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems (NeurIPS)* 33:12449–12460.
- Bai S, Kolter JZ, Koltun V. 2018.** An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *ArXiv preprint. Available at arXiv:1803.01271*.
- Bain M, Nagrani A, Schofield D, Berdugo S, Bessa J, Owen J, Hockings KJ, Matsuzawa T, Hayashi M, Biro D, Carvalho S, Zisserman A. 2021.** Automated audiovisual behavior recognition in wild primates. *Science Advances* 7(46):eabi4883 DOI 10.1126/sciadv.abi4883.
- Baker E, Vincent S. 2019.** A deafening silence: a lack of data and reproducibility in published bioacoustics research? *Biodiversity Data Journal* 7:e36783 DOI 10.3897/BDJ.7.e36783.
- Balestriero R, Cosentino R, Glotin H, Baraniuk R. 2018.** Spline filters for end-to-end deep learning. In: *Proceedings of the 35th International Conference on Machine Learning (ICML), volume 80 of Proceedings of Machine Learning Research*.
- Baumann J, Lohrenz T, Roy A, Fingscheidt T. 2020.** Beyond the dcase 2017 challenge on rare sound event detection: a proposal for a more realistic training and test framework. In: *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE.
- Beecher MD. 1989.** Signalling systems for individual recognition: an information theory approach. *Animal Behaviour* 38(2):248–261 DOI 10.1016/S0003-3472(89)80087-9.
- Beery S, Liu Y, Morris D, Piavis J, Kapoor A, Meister M, Perona P. 2020.** Synthetic examples improve generalization for rare classes. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, Vol. abs/1904.05916. Piscataway: IEEE, 863–873.
- Bergler C, Schmitt M, Cheng RX, Schröter H, Maier A, Barth V, Weber M, Nöth E. 2019a.** Deep representation learning for orca call type classification. In: Ekštejn K, ed. *Text, Speech, and Dialogue. TSD 2019. Lecture Notes in Computer Science*. Vol. 11697. Cham: Springer, 274–286 DOI 10.1007/978-3-030-27947-9_23.
- Bergler C, Schröter H, Cheng RX, Barth V, Weber M, Nöth E, Hofer H, Maier A. 2019b.** Orca-spot: an automatic killer whale sound detection toolkit using deep learning. *Scientific Reports* 9(1):1–17 DOI 10.1038/s41598-019-47335-w.

- Best P, Ferrari M, Poupard M, Paris S, Marxer R, Symonds H, Spong P, Glotin H. 2020.** Deep learning and domain transfer for orca vocalization detection. In: *International Joint Conference on Neural Networks*.
- Bhatia R. 2021.** Bird song synthesis using neural vocoders. Master's thesis, Itä-Suomen yliopisto.
- Bjorck J, Rappazzo BH, Chen D, Bernstein R, Wrege PH, Gomes CP. 2019.** Automatic detection and compression for passive acoustic monitoring of the African forest elephant. *Proceedings of the AAAI Conference on Artificial Intelligence* **33**:476–484 DOI [10.1609/aaai.v33i01.3301476](https://doi.org/10.1609/aaai.v33i01.3301476).
- Bravo Sanchez FJ, Hossain MR, English NB, Moore ST. 2021.** Bioacoustic classification of avian calls from raw sound waveforms with an open-source deep learning architecture. *Scientific Reports* **11**(1):1–12 DOI [10.1038/s41598-021-95076-6](https://doi.org/10.1038/s41598-021-95076-6).
- Brown A, Montgomery J, Garg S. 2021.** Automatic construction of accurate bioacoustics workflows under time constraints using a surrogate model. *Applied Soft Computing* **107944**(12):107944 DOI [10.1016/j.asoc.2021.107944](https://doi.org/10.1016/j.asoc.2021.107944).
- Brown CH, Riede T. 2017.** *Comparative bioacoustics: an overview*. Oak Park: Bentham Science Publishers.
- Cakir E, Adavanne S, Parascandolo G, Drossos K, Virtanen T. 2017.** Convolutional recurrent neural networks for bird audio detection. In: *Proceedings of EUSIPCO 2017*. Piscataway: IEEE, 1744–1748.
- Çakır E, Parascandolo G, Heittola T, Huttunen H, Virtanen T. 2017.** Convolutional recurrent neural networks for polyphonic sound event detection. *ArXiv preprint. Available at arXiv:1702.06286*.
- Canziani A, Paszke A, Culurciello E. 2016.** An analysis of deep neural network models for practical applications. *ArXiv preprint. Available at arXiv:1605.07678*.
- Chait M. 2020.** How the brain discovers structure in sound sequences. *Acoustical Science and Technology* **41**(1):48–53 DOI [10.1250/ast.41.48](https://doi.org/10.1250/ast.41.48).
- Chen H, Xie W, Vedaldi A, Zisserman A. 2020a.** VGGSound: a large-scale audio-visual dataset. In: *Proceedings of the 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 721–725.
- Chen RTQ, Amos B, Nickel M. 2020b.** Neural spatio-temporal point processes. *ArXiv preprint. Available at arXiv:2011.04583*.
- Chen X, Zhao J, Chen YH, Zhou W, Hughes AC. 2020c.** Automatic standardized processing and identification of tropical bat calls using deep learning approaches. *Biological Conservation* **241**(2):108269 DOI [10.1016/j.biocon.2019.108269](https://doi.org/10.1016/j.biocon.2019.108269).
- Chollet F. 2017.** Xception: deep learning with depthwise separable convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Piscataway: IEEE, 1251–1258.
- Chorowski JK, Bahdanau D, Serdyuk D, Cho K, Bengio Y. 2015.** Attention-based models for speech recognition. In: *Advances in Neural Information Processing systems*. 577–585.
- Clementino T, Colonna J. 2020.** Using triplet loss for bird species recognition on BirdCLEF 2020. In: *CLEF (Working Notes)*.
- Çoban EB, Pir D, So R, Mandel MI. 2020.** Transfer learning from Youtube soundtracks to tag arctic ecoacoustic recordings. In: *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 726–730.
- Coffey KR, Marx RG, Neumaier JF. 2019.** DeepSqueak: a deep learning-based system for detection and analysis of ultrasonic vocalizations. *Neuropsychopharmacology* **44**(5):859–868 DOI [10.1038/s41386-018-0303-6](https://doi.org/10.1038/s41386-018-0303-6).

- Cohen-Hadria A, Cartwright M, McFee B, Bello JP. 2019.** Voice anonymization in urban sound recordings. In: *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP)*. Piscataway: IEEE, 1–6.
- Colonna J, Peet T, Ferreira CA, Jorge AM, Gomes EF, Gama J. 2016.** Automatic classification of anuran sounds using convolutional neural networks. In: *Proceedings of the Ninth International c* Conference on Computer Science & Software Engineering*. 73–78.
- Cramer J, Lostanlen V, Farnsworth A, Salamon J, Bello JP. 2020.** Chirping up the right tree: incorporating biological taxonomies into deep bioacoustic classifiers. In: *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 901–905.
- Dekkers G, Rosas F, van Waterschoot T, Vanrumste B, Karsmakers P. 2022.** Dynamic sensor activation and decision-level fusion in wireless acoustic sensor networks for classification of domestic activities. *Information Fusion* 77(5):196–210 DOI 10.1016/j.inffus.2021.07.022.
- Dias FF, Ponti MA, Minghim R. 2021.** A classification and quantification approach to generate features in soundscape ecology using neural networks. *Neural Computing and Applications* 34:1923–1937 DOI 10.1007/s00521-021-06501-w.
- Disabato S, Canonaco G, Flikkema PG, Roveri M, Alippi C. 2021.** Birdsong detection at the edge with deep learning. In: *2021 IEEE International Conference on Smart Computing (SMARTCOMP)*. Piscataway: IEEE, 9–16.
- Dufourq E, Durbach I, Hansford JP, Hoepfner A, Ma H, Bryant JV, Stender CS, Li W, Liu Z, Chen Q, Zhou Z, Turvey ST, Pettorelli N, Astaras C. 2021.** Automated detection of Hainan Gibbon calls for passive acoustic monitoring. *Remote Sensing in Ecology and Conservation* 7(3):475–487 DOI 10.1002/rse2.201.
- Elliott D, Otero CE, Wyatt S, Martino E. 2021.** Tiny transformers for environmental sound classification at the edge. *ArXiv preprint*. Available at [arXiv:2103.12157](https://arxiv.org/abs/2103.12157).
- Fairbrass AJ, Firman M, Williams C, Brostow GJ, Titheridge H, Jones KE. 2019.** CityNet—deep learning tools for urban ecoacoustic assessment. *Methods in Ecology and Evolution* 10(2):186–197 DOI 10.1111/2041-210X.13114.
- Fonseca AH, Santana GM, Ortiz GMB, Bampi S, Dietrich MO. 2021.** Analysis of ultrasonic vocalizations from mice using computer vision and machine learning. *Elife* 10:e59161 DOI 10.7554/eLife.59161.
- Fox EJS. 2008.** Call-independent identification in birds. PhD thesis, University of Western Australia, Perth, Australia.
- Francl A, McDermott JH. 2020.** Deep neural network models of sound localization reveal how perception is adapted to real-world environments. *bioRxiv* DOI 10.1101/2020.07.21.214486.
- Frazao F, Padovese B, Kirsebom OS. 2020.** Workshop report: detection and classification in marine bioacoustics with deep learning. *ArXiv preprint*. Available at [arXiv:2002.08249](https://arxiv.org/abs/2002.08249).
- Fujimori K, Raytchev B, Kaneda K, Yamada Y, Teshima Y, Fujioka E, Hiryu S, Tamaki T. 2021.** Localization of flying bats from multichannel audio signals by estimating location map with convolutional neural networks. *Journal of Robotics and Mechatronics* 33(3):515–525 DOI 10.20965/jrm.2021.p0515.
- Ganchev T. 2017.** *Computational bioacoustics: biodiversity monitoring and assessment*. Vol. 4. Berlin: Walter de Gruyter.
- Gao R, Chen C, Al-Halah Z, Schissler C, Grauman K. 2020.** Visualechoes: spatial image representation learning through echolocation. In: *European Conference on Computer Vision*, Berlin: Springer, 658–676.

- Garcia HA, Couture T, Galor A, Topple JM, Huang W, Tiwari D, Ratilal P. 2020.** Comparing performances of five distinct automatic classifiers for fin whale vocalizations in beamformed spectrograms of coherent hydrophone array. *Remote Sensing* **12**(2):326 DOI [10.3390/rs12020326](https://doi.org/10.3390/rs12020326).
- Gillings S, Scott C. 2021.** Nocturnal flight calling behaviour of thrushes in relation to artificial light at night. *Ibis* **163**(4):1379–1393 DOI [10.1111/ibi.12955](https://doi.org/10.1111/ibi.12955).
- Glotin H, Ricard J, Balestriero R. 2017.** Fast chirplet transform injects priors in deep learning of animal calls and speech. In: *Proceedings of the 2017 ICLR Workshop*.
- Goëau H, Glotin H, Vellinga W-P, Planqué R, Joly A. 2016.** LifeCLEF bird identification task 2016: the arrival of deep learning. In: *Working Notes of CLEF 2016-Conference and Labs of the Evaluation Forum, Évora, Portugal*. 440–449.
- Goëau H, Glotin H, Vellinga W-P, Rauber A. 2014.** LifeCLEF bird identification task 2014. In: *CLEF Working Notes 2014*.
- Goodfellow I, Bengio Y, Courville A. 2016.** *Deep learning*. Cambridge: MIT Press.
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y. 2014.** Generative adversarial nets. In: *Advances in Neural Information Processing Systems*. 2672–2680.
- Graves A, Wayne G, Danihelka I. 2014.** Neural Turing machines. *ArXiv preprint*. Available at [arXiv:1410.5401](https://arxiv.org/abs/1410.5401).
- Grill T, Schlüter J. 2017.** Two convolutional neural networks for bird detection in audio signals. In: *Proceedings of EUSIPCO 2017, Special Session on Bird Audio Signal Processing*, 1764–1768.
- Gupta G, Kshirsagar M, Zhong M, Gholami S, Ferres JL. 2021.** Comparing recurrent convolutional neural networks for large scale bird species classification. *Scientific Reports* **11**(1):1–12 DOI [10.1038/s41598-021-96446-w](https://doi.org/10.1038/s41598-021-96446-w).
- Guyot P, Alix F, Guerin T, Lambeaux E, Rotureau A. 2021.** Fish migration monitoring from audio detection with CNNs. In: *Audio Mostly 2021*. 244–247.
- Hammer H, Chazan SE, Goldberger J, Gannot S. 2021.** Dynamically localizing multiple speakers based on the time-frequency domain. *EURASIP Journal on Audio, Speech, and Music Processing* **2021**(1):276 DOI [10.1186/s13636-021-00203-w](https://doi.org/10.1186/s13636-021-00203-w).
- Hardt M, Price E, Price E, Srebro N. 2016.** Equality of opportunity in supervised learning. In: Lee D, Sugiyama M, Luxburg U, Guyon I, Garnett R, eds. *Advances in Neural Information Processing Systems*. Vol. 29. Red Hook: Curran Associates, Inc.
- Hassan N, Ramli DA, Jaafar H. 2017.** Deep neural network approach to frog species recognition. In: *2017 IEEE 13th International Colloquium on Signal Processing & Its Applications (CSPA)*. Piscataway: IEEE, 173–178.
- He K, Zhang X, Ren S, Sun J. 2016.** Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Piscataway: IEEE, 770–778.
- Heath BE, Sethi SS, Orme CDL, Ewers RM, Picinali L. 2021.** How index selection, compression, and recording schedule impact the description of ecological soundscapes. *Ecology and Evolution* **11**(19):13206–13217 DOI [10.1002/ece3.8042](https://doi.org/10.1002/ece3.8042).
- Henderson P, Hu J, Romoff J, Brunskill E, Jurafsky D, Pineau J. 2020.** Towards the systematic reporting of the energy and carbon footprints of machine learning. *ArXiv preprint*. Available at [arXiv:2002.05651](https://arxiv.org/abs/2002.05651).
- Hershey S, Chaudhuri S, Ellis DP, Gemmeke JF, Jansen A, Moore RC, Plakal M, Platt D, Saurous RA, Seybold B, Slaney M, Weiss RJ, Wilson K. 2017.** CNN architectures for

- large-scale audio classification. In: *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 131–135.
- Heuer S, Tafo P, Holzmann H, Dahlke S. 2019.** New aspects in birdsong recognition utilizing the Gabor transform. In: *Proceedings of ICA 2019*, Aachen: Universitätsbibliothek der RWTH Aachen, 2917–2924.
- Hibino S, Suzuki C, Nishino T. 2021.** Classification of singing insect sounds with convolutional neural network. *Acoustical Science and Technology* **42(6)**:354–356 DOI [10.1250/ast.42.354](https://doi.org/10.1250/ast.42.354).
- Hill AP, Prince P, Covarrubias EP, Doncaster CP, Snaddon JL, Rogers A. 2018.** AudioMoth: evaluation of a smart open acoustic device for monitoring biodiversity and the environment. *Methods in Ecology and Evolution* **9(5)**:1199–1211 DOI [10.1111/2041-210X.12955](https://doi.org/10.1111/2041-210X.12955).
- Himawan I, Towsey M, Law B, Roe P. 2018.** Deep learning techniques for koala activity detection. In: *Proceedings of INTERSPEECH 2018*, Baixas: International Speech Communication Association, 2107–2111.
- Hochreiter S, Schmidhuber J. 1997.** Long short-term memory. *Neural Computation* **9(8)**:1735–1780 DOI [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735).
- Houégnigan L, Safari P, Nadeu C, van der Schaar M, André M. 2017.** A novel approach to real-time range estimation of underwater acoustic sources using supervised machine learning. In: *OCEANS 2017-Aberdeen*. Piscataway: IEEE, 1–5.
- Huang G, Liu Z, Weinberger KQ. 2016.** Densely connected convolutional networks. *ArXiv preprint*. Available at [arXiv:1608.06993](https://arxiv.org/abs/1608.06993).
- Ibrahim AK, Zhuang H, Chérubin LM, Erdol N, O’Corry-Crowe G, Ali AM. 2021.** A multimodel deep learning algorithm to detect north Atlantic right whale up-calls. *The Journal of the Acoustical Society of America* **150(2)**:1264–1272 DOI [10.1121/10.0005898](https://doi.org/10.1121/10.0005898).
- Ibrahim AK, Zhuang H, Chérubin LM, Schärer-Umpierre MT, Erdol N. 2018.** Automatic classification of grouper species by their sounds using deep neural networks. *The Journal of the Acoustical Society of America* **144(3)**:EL196–EL202 DOI [10.1121/1.5054911](https://doi.org/10.1121/1.5054911).
- IPBES. 2019.** Global assessment report on biodiversity and ecosystem services. Technical report, Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services (IPBES). Available at <https://www.ipbes.net/news/Media-Release-Global-Assessment>.
- Islam S, Valles D. 2020.** Houston toad and other chorusing amphibian species call detection using deep learning architectures. In: *2020 10th Annual Computing and Communication Workshop and Conference (CCWC)*. Piscataway: IEEE, 511–516.
- Ivanenko A, Watkins P, van Gerven MA, Hammerschmidt K, Englitz B. 2020.** Classifying sex and strain from mouse ultrasonic vocalizations using deep learning. *PLoS Computational Biology* **16(6)**:e1007918 DOI [10.1371/journal.pcbi.1007918](https://doi.org/10.1371/journal.pcbi.1007918).
- Jain M, Balakrishnan R. 2011.** Microhabitat selection in an assemblage of crickets (orthoptera: Ensifera) of a tropical evergreen forest in southern India. *Insect Conservation and Diversity* **4(2)**:152–158 DOI [10.1111/j.1752-4598.2010.00118.x](https://doi.org/10.1111/j.1752-4598.2010.00118.x).
- Jancovic P, Köküer M. 2019.** Bird species recognition using unsupervised modeling of individual vocalization elements. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **27(5)**:932–947 DOI [10.1109/TASLP.2019.2904790](https://doi.org/10.1109/TASLP.2019.2904790).
- Janetzky P, Davidson P, Steininger M, Krause A, Hotho A. 2021.** Detecting presence of speech in acoustic data obtained from beehives. In: *Proceedings of the 2021 DCASE Workshop*.
- Jansson A, Humphrey E, Montecchio N, Bittner R, Kumar A, Weyde T. 2017.** Singing voice separation with deep u-net convolutional networks. In: *Proceedings of ISMIR 2017*.

- Jiang JJ, Bu LR, Duan FJ, Wang XQ, Liu W, Sun ZB, Li CY. 2019. Whistle detection and classification for whales based on convolutional neural networks. *Applied Acoustics* 150(1–4):169–178 DOI 10.1016/j.apacoust.2019.02.007.
- Johnston A, Newson SE, Risely K, Musgrove AJ, Massimino D, Baillie SR, Pearce-Higgins JW. 2014. Species traits explain variation in detectability of UK birds. *Bird Study* 61(3):340–350 DOI 10.1080/00063657.2014.941787.
- Jolles JW. 2021. Broad-scale applications of the Raspberry Pi: a review and guide for biologists. *Methods in Ecology and Evolution* 12(9):1562–1579 DOI 10.1111/2041-210X.13652.
- Joly A, Goëau H, Glotin H, Spampinato C, Bonnet P, Vellinga W-P, Lombardo J-C, Planqué R, Palazzo S, Müller H. 2019. Biodiversity information retrieval through large scale content-based identification: a long-term evaluation. In: *Information Retrieval Evaluation in a Changing World*, Berlin: Springer, 389–413.
- Joly A, Goëau H, Kahl S, Picck L, Lorieul T, Cole E, Deneu B, Servajean M, Durso A, Bolon I, Glotin H, Planqué R, de Castañeda RR, Vellinga WP, Klinck H, Denton T, Eggel I, Bonnet P, Müller H. 2021. Overview of lifeclef 2021: an evaluation of machine-learning based species identification and species distribution prediction. In: *International Conference of the Cross-Language Evaluation Forum for European Languages*, Berlin: Springer, 371–393.
- Jones DL, Baraniuk RG. 1995. An adaptive optimal-kernel time-frequency representation. *IEEE Transactions on Signal Processing* 43(10):2361–2371 DOI 10.1109/78.469854.
- Jung D-H, Kim NY, Moon SH, Jhin C, Kim H-J, Yang J-S, Kim HS, Lee TS, Lee JY, Park SH. 2021. Deep learning-based cattle vocal classification model and real-time livestock monitoring system with noise filtering. *Animals* 11(2):357 DOI 10.3390/ani11020357.
- Kahl S, Wood CM, Eibl M, Klinck H. 2021. BirdNET: a deep learning solution for avian diversity monitoring. *Ecological Informatics* 61(7):101236 DOI 10.1016/j.ecoinf.2021.101236.
- Kaplan J, McCandlish S, Henighan T, Brown TB, Chess B, Child R, Gray S, Radford A, Wu J, Amodei D. 2020. Scaling laws for neural language models. *ArXiv preprint. Available at arXiv:2001.08361*.
- Kershenbaum A, Blumstein DT, Roch MA, Akçay Ç, Backus G, Bee MA, Bohn K, Cao Y, Carter G, Cäsar C, Coen M, DeRuiter SL, Doyle L, Edelman S, Ferrer-i-Cancho R, Freeberg TM, Garland EC, Gustison M, Harley HE, Huetz C, Hughes M, Hyland Bruno J, Ilany A, Jin DZ, Johnson M, Ju C, Karnowski J, Lohr B, Manser MB, McCowan B, Mercado E III, Narins PM, Piel A, Rice M, Salmi R, Sasahara K, Sayigh L, Shiu Y, Taylor C, Vallejo EE, Waller S, Zamora-Gutierrez V. 2014. Acoustic sequences in non-human animals: a tutorial review and prospectus. *Biological Reviews* 91(1):13–52 DOI 10.1111/brv.12160.
- Khalighifar A, Jiménez-García D, Campbell LP, Ahadji-Dabla KM, Aboagye-Antwi F, Ibarra-Juárez LA, Peterson AT. 2021. Application of deep learning to community-science-based mosquito monitoring and detection of novel species. *Journal of Medical Entomology* 59(1):355–362 DOI 10.1093/jme/tjab161.
- Kholghi M, Phillips Y, Towsey M, Sitbon L, Roe P. 2018. Active learning for classifying long-duration audio recordings of the environment. *Methods in Ecology and Evolution* 9(9):1948–1958 DOI 10.1111/2041-210X.13042.
- Kiskin I, Sinka M, Cobb AD, Rafique W, Wang L, Zilli D, Gutteridge B, Dam R, Marinos T, Li Y, Msaky D, Kaindoa E, Killeen G, Herreros-Moya E, Willis KJ, Roberts SJ. 2021. HumBugDB: a large-scale acoustic mosquito dataset. *ArXiv preprint. Available at arXiv:2110.07607*.

- Kiskin I, Zilli D, Li Y, Sinka M, Willis K, Roberts S. 2020.** Bioacoustic detection with wavelet-conditioned convolutional neural networks. *Neural Computing and Applications* 32(4):915–927 DOI [10.1007/s00521-018-3626-7](https://doi.org/10.1007/s00521-018-3626-7).
- Kitzes J, Schricker L. 2019.** The necessity, promise and challenge of automated biodiversity surveys. *Environmental Conservation* 46(4):247–250 DOI [10.1017/S0376892919000146](https://doi.org/10.1017/S0376892919000146).
- Knight E, Hannah K, Foley G, Scott C, Brigham R, Bayne E. 2017.** Recommendations for acoustic recognizer performance assessment with application to five common automated signal recognition programs. *Avian Conservation and Ecology* 12(2):14 DOI [10.5751/ACE-01114-120214](https://doi.org/10.5751/ACE-01114-120214).
- Knight EC, Poo Hernandez S, Bayne EM, Bulitko V, Tucker BV. 2020.** Pre-processing spectrogram parameters improve the accuracy of bioacoustic classification using convolutional neural networks. *Bioacoustics-the International Journal of Animal Sound and Its Recording* 29(3):337–355 DOI [10.1080/09524622.2019.1606734](https://doi.org/10.1080/09524622.2019.1606734).
- Kobayashi K, Masuda K, Haga C, Matsui T, Fukui D, Machimura T. 2021.** Development of a species identification system of Japanese bats from echolocation calls using convolutional neural networks. *Ecological Informatics* 62:101253 DOI [10.1016/j.ecoinf.2021.101253](https://doi.org/10.1016/j.ecoinf.2021.101253).
- Koenecke A, Nam A, Lake E, Nudell J, Quartey M, Mengesha Z, Toups C, Rickford JR, Jurafsky D, Goel S. 2020.** Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences of the United States of America* 117(14):7684–7689 DOI [10.1073/pnas.1915768117](https://doi.org/10.1073/pnas.1915768117).
- Kojima R, Sugiyama O, Hoshihara K, Suzuki R, Nakadai K. 2018.** Hark-bird-box: a portable real-time bird song scene analysis system. In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Piscataway: IEEE, 2497–2502.
- Kong Q, Xu Y, Plumbley M. 2017.** Joint detection and classification convolutional neural network on weakly labelled bird audio detection. In: *Proceedings of EUSIPCO 2017, Special Session on Bird Audio Signal Processing*, 1749–1753.
- Koops HV, Van Balen J, Wiering F. 2015.** Automatic segmentation and deep learning of bird sounds. In: *International Conference of the Cross-Language Evaluation Forum for European Languages*, Berlin: Springer, 261–267.
- Koumura T, Okanoya K. 2016.** Automatic recognition of element classes and boundaries in the birdsong with variable sequences. *PLOS ONE* 11(7):e0159188 DOI [10.1371/journal.pone.0159188](https://doi.org/10.1371/journal.pone.0159188).
- Krizhevsky A, Sutskever I, Hinton GE. 2012.** Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25:1097–1105.
- Laiolo P. 2010.** The emerging significance of bioacoustics in animal species conservation. *Biological Conservation* 143(7):1635–1645 DOI [10.1016/j.biocon.2010.03.025](https://doi.org/10.1016/j.biocon.2010.03.025).
- Lasseck M. 2018.** Audio-based bird species identification with deep convolutional neural networks. In: *Working Notes of CLEF, 2018*.
- Le Cornu T, Mitchell C, Cooper N. 2021.** Audio for machine learning: The law and your reputation. Technical report, Audio Analytic Ltd. Available at <https://www.audioanalytic.com/be-careful-about-how-you-get-that-data-set-ai-industry-warned-by-us-federal-trade-commission/>.
- LeBien J, Zhong M, Campos-Cerqueira M, Velez JP, Dodhia R, Ferres JL, Aide TM. 2020.** A pipeline for identification of bird and frog species in tropical soundscape recordings using a convolutional neural network. *Ecological Informatics* 59(5):101113 DOI [10.1016/j.ecoinf.2020.101113](https://doi.org/10.1016/j.ecoinf.2020.101113).

- LeCun Y, Bengio Y, Hinton G. 2015.** Deep learning. *Nature* **521(7553)**:436–444
DOI [10.1038/nature14539](https://doi.org/10.1038/nature14539).
- LeCun Y, Bottou L, Bengio Y, Haffner P. 1998.** Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86(11)**:2278–2324 DOI [10.1109/5.726791](https://doi.org/10.1109/5.726791).
- Li L, Qiao G, Liu S, Qing X, Zhang H, Mazhar S, Niu F. 2021.** Automated classification of *Tursiops aduncus* whistles based on a depth-wise separable convolutional neural network and data augmentation. *The Journal of the Acoustical Society of America* **150(5)**:3861–3873
DOI [10.1121/10.0007291](https://doi.org/10.1121/10.0007291).
- Li P, Liu X, Palmer K, Fleishman E, Gillespie D, Nosal E-M, Shiu Y, Klinck H, Cholewiak D, Helble T, Roch MA. 2020.** Learning deep models from synthetic data for extracting dolphin whistle contours. In: *2020 International Joint Conference on Neural Networks (IJCNN)*. Piscataway: IEEE, 1–10.
- Li S, Li X, Xing Z, Zhang Z, Wang Y, Li R, Guo R, Xie J. 2019.** Intelligent audio bird repeller for transmission line tower based on bird species variation. In: *IOP Conference Series: Materials Science and Engineering*, Vol. 592. Bristo: IOP Publishing, 12142.
- Lin T-H, Tsao Y. 2020.** Source separation in ecoacoustics: a roadmap towards versatile soundscape information retrieval. *Remote Sensing in Ecology and Conservation* **6(3)**:236–247
DOI [10.1002/rse2.141](https://doi.org/10.1002/rse2.141).
- Linhart P, Osiejuk T, Budka M, Šálek M, Špinka M, Policht R, Syrová M, Blumstein DT. 2019.** Measuring individual identity information in animal signals: overview and performance of available identity metrics. *Methods in Ecology and Evolution* **10(9)**:1558–1570
DOI [10.1111/2041-210X.13238](https://doi.org/10.1111/2041-210X.13238).
- Linke S, Gifford T, Desjonquères C, Tonolla D, Aubin T, Barclay L, Karaconstantis C, Kennard MJ, Rybak F, Sueur J. 2018.** Freshwater ecoacoustics as a tool for continuous ecosystem monitoring. *Frontiers in Ecology and the Environment* **16(4)**:231–238
DOI [10.1002/fee.1779](https://doi.org/10.1002/fee.1779).
- Lostanlen V, Arnaud P, du Gardin M, Godet L, Lagrange M. 2021a.** An interactive human—animal—robot approach to distance sampling in bioacoustics. In: *Proceedings of the 3rd International Workshop on Vocal Interactivity in-and-between Humans, Animals and Robots (VIHAR 2021)*.
- Lostanlen V, Bernabeu A, Béchenec JL, Briday M, Faucou S, Lagrange M. 2021b.** Energy efficiency is not enough: towards a batteryless internet of sounds. In: *International Workshop on the Internet of Sounds*. New York: Association for Computing Machinery.
- Lostanlen V, Palmer K, Knight E, Clark C, Klinck H, Farnsworth A, Wong T, Cramer J, Bello JP. 2019a.** Long-distance detection of bioacoustic events with per-channel energy normalization. In: *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2019 Workshop (DCASE2019)*. 144–148.
- Lostanlen V, Salamon J, Farnsworth A, Kelling S, Bello JP. 2018.** Birdvox-full-night: a dataset and benchmark for avian flight call detection. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 266–270.
- Lostanlen V, Salamon J, Farnsworth A, Kelling S, Bello JP. 2019b.** Robust sound event detection in bioacoustic sensor networks. *PLOS ONE* **14(10)**:1–31 DOI [10.1371/journal.pone.0214168](https://doi.org/10.1371/journal.pone.0214168).
- Mac Aodha O, Gibb R, Barlow KE, Browning E, Firman M, Freeman R, Harder B, Kinsey L, Mead GR, Newson SE, Pandourski I, Parsons S, Russ J, Szodoray-Paradi A, Szodoray-Paradi F, Tilova E, Girolami M, Brostow G, Jones KE, Fenton B. 2018.** Bat detective—deep learning tools for bat acoustic signal detection. *PLOS Computational Biology* **14(3)**:e1005995
DOI [10.1371/journal.pcbi.1005995](https://doi.org/10.1371/journal.pcbi.1005995).

- Madhusudhana S, Shiu Y, Klinck H, Fleishman E, Liu X, Nosal E-M, Helble T, Cholewiak D, Gillespie D, Širović A, Roch MA. 2021.** Improve automatic detection of animal call sequences with temporal context. *Journal of the Royal Society Interface* **18(180)**:20210297 DOI [10.1098/rsif.2021.0297](https://doi.org/10.1098/rsif.2021.0297).
- Maegawa Y, Ushigome Y, Suzuki M, Taguchi K, Kobayashi K, Haga C, Matsui T. 2021.** A new survey method using convolutional neural networks for automatic classification of bird calls. *Ecological Informatics* **61(4)**:101164 DOI [10.1016/j.ecoinf.2020.101164](https://doi.org/10.1016/j.ecoinf.2020.101164).
- Manilow E, Seetharman P, Salamon J. 2020.** Open source tools & data for music source separation. Available at <https://source-separation.github.io/tutorial>.
- Marchal J, Fabianek F, Aubry Y. 2021.** Software performance for the automated identification of bird vocalisations: the case of two closely related species. *Bioacoustics-the International Journal of Animal Sound and Its Recording* **68(3)**:1–17 DOI [10.1080/09524622.2021.1945952](https://doi.org/10.1080/09524622.2021.1945952).
- Marler PR, Slabbekoorn H. 2004.** *Nature's music: the science of birdsong*. Cambridge: Academic Press.
- Marques TA, Thomas L, Martin SW, Mellinger DK, Ward JA, Moretti DJ, Harris D, Tyack PL. 2012.** Estimating animal population density using passive acoustics. *Biological Reviews* **88(2)**:287–309 DOI [10.1111/brv.12001](https://doi.org/10.1111/brv.12001).
- Mercado E III, Sturdy CB. 2017.** Chapter 10: classifying animal sounds with neural networks. In: Brown CH, Riede T, eds. *Comparative Bioacoustics: An Overview*. Oak Park: Bentham Science Publishers, 415–461.
- Mesaros A, Diment A, Elizalde B, Heittola T, Vincent E, Raj B, Virtanen T. 2019.** Sound event detection in the DCASE 2017 challenge. *IEEE/ACM Transactions on Audio, Speech and Language Processing* **27(6)**:992–1006 DOI [10.1109/TASLP.2019.2907016](https://doi.org/10.1109/TASLP.2019.2907016).
- Mesaros A, Heittola T, Virtanen T. 2016.** Metrics for polyphonic sound event detection. *Applied Sciences* **6(6)**:162 DOI [10.3390/app6060162](https://doi.org/10.3390/app6060162).
- Mesaros A, Heittola T, Virtanen T, Plumbley MD. 2021.** Sound event detection: a tutorial. *IEEE Signal Processing Magazine* **38(5)**:67–83 DOI [10.1109/MSP.2021.3090678](https://doi.org/10.1109/MSP.2021.3090678).
- Mishachandar B, Vairamuthu S. 2021.** Diverse ocean noise classification using deep learning. *Applied Acoustics* **181**:108141 DOI [10.1016/j.apacoust.2021.108141](https://doi.org/10.1016/j.apacoust.2021.108141).
- Montgomery GA, Dunn RR, Fox R, Jongejans E, Leather SR, Saunders ME, Shortall CR, Tingley MW, Wagner DL. 2020.** Is the insect apocalypse upon us? How to find out. *Biological Conservation* **241(3)**:108327 DOI [10.1016/j.biocon.2019.108327](https://doi.org/10.1016/j.biocon.2019.108327).
- Morfi V, Bas Y, Pamula H, Glotin H, Stowell D. 2019.** Nips4bplus: a richly annotated birdsong audio dataset. *PeerJ Computer Science* **5(6)**:e223 DOI [10.7717/peerj-cs.223](https://doi.org/10.7717/peerj-cs.223).
- Morfi V, Lachlan RF, Stowell D. 2021a.** Deep perceptual embeddings for unlabelled animal sound events. *The Journal of the Acoustical Society of America* **150(1)**:2–11 DOI [10.1121/10.0005475](https://doi.org/10.1121/10.0005475).
- Morfi V, Nolasco I, Lostanlen V, Singh S, Strandburg-Peshkin A, Gill L, Pamula H, Benvent D, Stowell D. 2021b.** Few-shot bioacoustic event detection: a new task at the DCASE, 2021 challenge. In: *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2021 Workshop (DCASE2021)*, Barcelona, Spain, 145–149.
- Morfi V, Stowell D. 2018.** Deep learning for audio event detection and tagging on low-resource datasets. *Applied Sciences* **8(8)**:1397 DOI [10.3390/app8081397](https://doi.org/10.3390/app8081397).
- Morgan M, Braasch J. 2021.** Long-term deep learning-facilitated environmental acoustic monitoring in the capital region of New York State. *Ecological Informatics* **61**:101242 DOI [10.1016/j.ecoinf.2021.101242](https://doi.org/10.1016/j.ecoinf.2021.101242).

- Nanni L, Rigo A, Lumini A, Brahnam S. 2020.** Spectrogram classification using dissimilarity space. *Applied Sciences* **10**(12):4176 DOI [10.3390/app10124176](https://doi.org/10.3390/app10124176).
- Narasimhan R, Fern XZ, Raich R. 2017.** Simultaneous segmentation and classification of bird song using CNN. In: *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 146–150.
- Niculescu-Mizil A, Caruana R. 2005.** Predicting good probabilities with supervised learning. In: *Proceedings of the 22nd International Conference on Machine Learning*, New York: ACM, 625–632.
- Nolasco I, Stowell D. 2022.** Rank-based loss for learning hierarchical representations. *ArXiv preprint*. Available at [arXiv:2110.05941](https://arxiv.org/abs/2110.05941).
- Ntalampiras S. 2018.** Bird species identification via transfer learning from music genres. *Ecological Informatics* **44**(12):76–81 DOI [10.1016/j.ecoinf.2018.01.006](https://doi.org/10.1016/j.ecoinf.2018.01.006).
- Ntalampiras S, Potamitis I. 2021.** Acoustic detection of unknown bird species and individuals. *CAAI Transactions on Intelligence Technology* **6**(3):291–300 DOI [10.1049/cit2.12007](https://doi.org/10.1049/cit2.12007).
- Oikarinen T, Srinivasan K, Meisner O, Hyman JB, Parmar S, Fanucci-Kiss A, Desimone R, Landman R, Feng G. 2019.** Deep convolutional network for animal sound classification and source attribution using dual audio recordings. *The Journal of the Acoustical Society of America* **145**(2):654–662 DOI [10.1121/1.5087827](https://doi.org/10.1121/1.5087827).
- O’Shaughnessy D. 2003.** Interacting with computers by voice: automatic speech recognition and synthesis. *Proceedings of the IEEE* **91**(9):1272–1305 DOI [10.1109/JPROC.2003.817117](https://doi.org/10.1109/JPROC.2003.817117).
- Ozanich E, Thode A, Gerstoft P, Freeman LA, Freeman S. 2021.** Deep embedded clustering of coral reef bioacoustics. *The Journal of the Acoustical Society of America* **149**(4):2587–2601 DOI [10.1121/10.0004221](https://doi.org/10.1121/10.0004221).
- Padovese B, Frazao F, Kirsebom OS, Matwin S. 2021.** Data augmentation for the classification of North Atlantic right whales upcalls. *The Journal of the Acoustical Society of America* **149**(4):2520–2530 DOI [10.1121/10.0004258](https://doi.org/10.1121/10.0004258).
- Phillips YF, Towsey M, Roe P. 2018.** Revealing the ecological content of long-duration audio-recordings of the environment through clustering and visualisation. *PLOS ONE* **13**(3):1–27 DOI [10.1371/journal.pone.0193345](https://doi.org/10.1371/journal.pone.0193345).
- Pons J, Serrà J, Serra X. 2019.** Training neural audio classifiers with few data. Available at <https://arxiv.org/abs/1810.10274>.
- Premoli M, Baggi D, Bianchetti M, Gnutti A, Bondaschi M, Mastinu A, Migliorati P, Signoroni A, Leonardi R, Memo M, Bonini SA. 2021.** Automatic classification of mice vocalizations using machine learning techniques and convolutional neural networks. *PLOS ONE* **16**(11):e0244636 DOI [10.1371/journal.pone.0244636](https://doi.org/10.1371/journal.pone.0244636).
- Prince P, Hill A, Piña Covarrubias E, Doncaster P, Snaddon JL, Rogers A. 2019.** Deploying acoustic detection algorithms on low-cost, open-source acoustic sensors for environmental monitoring. *Sensors* **19**(3):553 DOI [10.3390/s19030553](https://doi.org/10.3390/s19030553).
- Ptacek L, Machlica L, Linhart P, Jaska P, Muller L. 2016.** Automatic recognition of bird individuals on an open set using as-is recordings. *Bioacoustics-the International Journal of Animal Sound and Its Recording* **25**(1):55–73 DOI [10.1080/09524622.2015.1089524](https://doi.org/10.1080/09524622.2015.1089524).
- Qian K, Zhang Z, Baird A, Schuller B. 2017.** Active learning for bird sound classification via a kernel-based extreme learning machine. *The Journal of the Acoustical Society of America* **142**(4):1796–1804 DOI [10.1121/1.5004570](https://doi.org/10.1121/1.5004570).
- Ranft R. 2004.** Natural sound archives: past, present and future. *Anais da Academia Brasileira de Ciências* **76**(2):456–460 DOI [10.1590/S0001-37652004000200041](https://doi.org/10.1590/S0001-37652004000200041).

- Ravanelli M, Bengio Y. 2018. Speech and speaker recognition from raw waveform with SincNet. *ArXiv preprint. Available at [arXiv:1812.05920](https://arxiv.org/abs/1812.05920).*
- Ren Z, Kong Q, Qian K, Plumbley MD, Schuller B. 2018. Attention-based convolutional neural networks for acoustic scene classification. In: *Proceedings of DCASE 2018*.
- Rigakis I, Potamitis I, Tatlas N-A, Potirakis SM, Ntalampiras S. 2021. TreeVibes: modern tools for global monitoring of trees for borers. *Smart Cities* 4(1):271–285 DOI 10.3390/smartcities4010017.
- Roch MA, Lindeneau S, Aurora GS, Frasier KE, Hildebrand JA, Glotin H, Baumann-Pickering S. 2021. Using context to train time-domain echolocation click detectors. *The Journal of the Acoustical Society of America* 149(5):3301–3310 DOI 10.1121/10.0004992.
- Roch MA, Miller P, Helble TA, Baumann-Pickering S, Širović A. 2017. Organizing metadata from passive acoustic localizations of marine animals. *The Journal of the Acoustical Society of America* 141(5):3605 DOI 10.1121/1.4987714.
- Roe P, Eichinski P, Fuller RA, McDonald PG, Schwarzkopf L, Towsey M, Truskinger A, Tucker D, Watson DM. 2021. The Australian Acoustic Observatory. *Methods in Ecology and Evolution* 12(10):1802–1808 DOI 10.1111/2041-210X.13660.
- Romero-Mujalli D, Bergmann T, Zimmermann A, Scheumann M. 2021. Utilizing DeepSqueak for automatic detection and classification of mammalian vocalizations: a case study on primate vocalizations. *Scientific Reports* 11(1):1–15 DOI 10.1038/s41598-021-03941-1.
- Ronneberger O, Fischer P, Brox T. 2015. U-net: convolutional networks for biomedical image segmentation. Available at <http://arxiv.org/abs/1505.04597>.
- Rowe B, Eichinski P, Zhang J, Roe P. 2021. Acoustic auto-encoders for biodiversity assessment. *Ecological Informatics* 62(1):101237 DOI 10.1016/j.ecoinf.2021.101237.
- Ruff ZJ, Lesmeister DB, Appel CL, Sullivan CM. 2021. Workflow and convolutional neural network for automated identification of animal sounds. *Ecological Indicators* 124(2):107419 DOI 10.1016/j.ecolind.2021.107419.
- Rumelt RB, Basto A, Roncal CM. 2021. Automated audio recording as a means of surveying tinamous (Tinamidae) in the Peruvian Amazon. *Ecology and Evolution* 11(19):13518–13531 DOI 10.1002/ece3.8078.
- Ryazanov I, Nylund AT, Basu D, Hassellöv I-M, Schliep A. 2021. Deep learning for deep waters: an expert-in-the-loop machine learning framework for marine sciences. *Journal of Marine Science and Engineering* 9(2):169 DOI 10.3390/jmse9020169.
- Saeed A, Grangier D, Zeghidour N. 2021. Contrastive learning of general-purpose audio representations. In: *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE.
- Salamon J, Bello JP, Farnsworth A, Kelling S. 2017a. Fusing shallow and deep learning for bioacoustic bird species classification. In: *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 141–145.
- Salamon J, MacConnell D, Cartwright M, Li P, Bello JP. 2017b. Scaper: a library for soundscape synthesis and augmentation. In: *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. Piscataway: IEEE, 344–348.
- Sandler M, Howard A, Zhu M, Zhmoginov A, Chen L-C. 2018. Mobilenetv2: inverted residuals and linear bottlenecks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Piscataway: IEEE, 4510–4520.
- Schröter H, Nöth E, Maier A, Cheng R, Barth V, Bergler C. 2019. Segmentation, classification, and visualization of orca calls using deep learning. In: *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 8231–8235.

- Searby A, Jouventin P, Aubin T. 2004. Acoustic recognition in macaroni penguins: an original signature system. *Animal Behaviour* 67(4):615–625 DOI 10.1016/j.anbehav.2003.03.012.
- Sethi SS, Ewers RM, Jones NS, Sleutel J, Shabrani A, Zulkifli N, Picinali L. 2021. Soundscapes predict species occurrence in tropical forests. *Oikos* 2022(3):1 DOI 10.1111/oik.08525.
- Sethi SS, Jones NS, Fulcher BD, Picinali L, Clink DJ, Klinck H, Orme CDL, Wrege PH, Ewers RM. 2020. Characterizing soundscapes across diverse ecosystems using a universal acoustic feature set. *Proceedings of The National Academy of Sciences of The United States of America* 117(29):17049–17055 DOI 10.1073/pnas.2004702117.
- Shimada K, Koyama Y, Takahashi N, Takahashi S, Mitsufuji Y. 2021. Accdoa: activity-coupled cartesian direction of arrival representation for sound event localization and detection. In: 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). Piscataway: IEEE, 915–919.
- Shiu Y, Palmer KJ, Roch MA, Fleishman E, Liu X, Nosal E-M, Helble T, Cholewiak D, Gillespie D, Klinck H. 2020. Deep neural networks for automated detection of marine mammal species. *Scientific Reports* 10(1):607 DOI 10.1038/s41598-020-57549-y.
- Shrestha R, Glackin C, Wall J, Cannings N. 2021. Bird audio diarization with faster R-CNN. In: *International Conference on Artificial Neural Networks*, Berlin: Springer, 415–426.
- Simon R, Bakunowski K, Reyes-Vasques A, Tschapka M, Knörnschild M, Steckel J, Stowell D. 2021. Acoustic traits of bat-pollinated flowers compared to flowers of other pollination syndromes and their echo-based classification using convolutional neural networks. *PLoS Computational Biology* 17(12):e1009706 DOI 10.1371/journal.pcbi.1009706.
- Simon R, Varkevisser J, Mendoza E, Hochradel K, Scharff C, Riebel K, Halfwerk W. 2019. Development and application of a robotic zebra finch (RoboFinch) to study multimodal cues in vocal communication. *PeerJ Preprints* 7:e28004v2 DOI 10.7287/peerj.preprints.28004v3.
- Simonyan K, Zisserman A. 2014. Very deep convolutional networks for large-scale image recognition. *ArXiv preprint. Available at arXiv:1409.1556*.
- Sinha R, Rajan P. 2018. A deep autoencoder approach to bird call enhancement. In: 2018 IEEE 13th International Conference on Industrial and Information Systems (ICIIS). Piscataway: IEEE, 22–26.
- Sinka ME, Zilli D, Li Y, Kiskin I, Kirkham D, Rafique W, Wang L, Chan H, Gutteridge B, Herreros-Moya E, Portwood H, Roberts S, Willis KJ. 2021. HumBug-an acoustic mosquito monitoring tool for use on budget smartphones. *Methods in Ecology and Evolution* 12(10):1848–1859 DOI 10.1111/2041-210X.13663.
- Slonina Z, Bonzini AA, Brown J, Wang S, Farkhatdinov I, Althoefer K, Jamone L, Versace E. 2021. Using robochick to identify the behavioral features promoting social interactions. In: 2021 IEEE International Conference on Development and Learning (ICDL). Piscataway: IEEE, 1–6.
- Smith AA, Kristensen D. 2017. Deep learning to extract laboratory mouse ultrasonic vocalizations from scalograms. In: 2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). Piscataway: IEEE, 1972–1979.
- Steinfath E, Palacios-Muñoz A, Rottschäfer JR, Yuezak D, Clemens J. 2021. Fast and accurate annotation of acoustic signals with deep neural networks. *eLife* 10:e68837 DOI 10.7554/eLife.68837.
- Stevenson BC, Borchers DL, Altwegg R, Swift RJ, Gillespie DM, Measey GJ. 2015. A general framework for animal density estimation from acoustic detections across a fixed microphone array. *Methods in Ecology and Evolution* 6(1):38–48 DOI 10.1111/2041-210X.12291.
- Stowell D. 2018. Chapter 11: computational bioacoustic scene analysis. In: *Computational Analysis of Sound Scenes and Events*, Berlin: Springer, 303–333.

- Stowell D, Clayton D. 2015.** Acoustic event detection for multiple overlapping similar sources. In: *2015 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. Piscataway: IEEE.
- Stowell D, Gill LF, Clayton D. 2016a.** Detailed temporal structure of communication networks in groups of songbirds. *Journal of the Royal Society Interface* **13(119)**:20160296 DOI [10.1098/rsif.2016.0296](https://doi.org/10.1098/rsif.2016.0296).
- Stowell D, Petrusková T, Šálek M, Linhart P. 2019a.** Automatic acoustic identification of individuals in multiple species: improving identification across recording conditions. *Journal of The Royal Society Interface* **16(153)**:20180940 DOI [10.1098/rsif.2018.0940](https://doi.org/10.1098/rsif.2018.0940).
- Stowell D, Stylianou Y, Wood M, Pamuła H, Glotin H. 2019b.** Automatic acoustic detection of birds through deep learning: the first bird audio detection challenge. *Methods in Ecology and Evolution* **16(153)**:368–380 DOI [10.1111/2041-210X.13103](https://doi.org/10.1111/2041-210X.13103).
- Stowell D, Wood M, Stylianou Y, Glotin H. 2016b.** Bird detection in audio: a survey and a challenge. In: *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*. Piscataway: IEEE.
- Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A. 2014.** Going deeper with convolutions. *ArXiv preprint*. Available at [arXiv:1409.4842](https://arxiv.org/abs/1409.4842).
- Tan M, Le Q. 2019.** Efficientnet: rethinking model scaling for convolutional neural networks. In: *International Conference on Machine Learning*. PMLR. 6105–6114.
- Teşileanu T, Ölveczky B, Balasubramanian V. 2017.** Rules and mechanisms for efficient two-stage learning in neural circuits. *eLife* **6**:e20944 DOI [10.7554/eLife.20944](https://doi.org/10.7554/eLife.20944).
- Thakur A, Thapar D, Rajan P, Nigam A. 2019.** Deep metric learning for bioacoustic classification: overcoming training data scarcity using dynamic triplet loss. *The Journal of the Acoustical Society of America* **146(1)**:534–547 DOI [10.1121/1.5118245](https://doi.org/10.1121/1.5118245).
- Thomas M, Martin B, Kowarski K, Gaudet B, Matwin S. 2019.** Marine mammal species classification using convolutional neural networks and a novel acoustic representation. In: *Proceedings of the ECMLPKDD 2019*.
- Towsey M, Planitz B, Nantes A, Wimmer J, Roe P. 2012.** A toolbox for animal call recognition. *Bioacoustics* **21(2)**:107–125 DOI [10.1080/09524622.2011.648753](https://doi.org/10.1080/09524622.2011.648753).
- Traer J, McDermott JH. 2016.** Statistics of natural reverberation enable perceptual separation of sound and space. *Proceedings of The National Academy of Sciences of The United States of America* **113(48)**:E7856–E7865 DOI [10.1073/pnas.1612524113](https://doi.org/10.1073/pnas.1612524113).
- Turpault N, Serizel R, Wisdom S, Erdogan H, Hershey JR, Fonseca E, Seetharaman P, Salamon J. 2021.** Sound event detection and separation: a benchmark on DESED synthetic soundscapes. In: *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Piscataway: IEEE, 840–844.
- Tzirakis P, Shiarella A, Ewers R, Schuller BW. 2020.** Computer audition for continuous rainforest occupancy monitoring: the case of Bornean Gibbons' call detection. In: *Proceedings of InterSpeech*. 1211–1215.
- van den Oord A, Dieleman S, Zen H, Simonyan K, Vinyals O, Graves A, Kalchbrenner N, Senior A, Kavukcuoglu K. 2016.** Wavenet: a generative model for raw audio. Technical report, Google DeepMind.
- Van Komen DF, Neilsen TB, Howarth K, Knobles DP, Dahl PH. 2020.** Seabed and range estimation of impulsive time series using a convolutional neural network. *The Journal of the Acoustical Society of America* **147(5)**:EL403–EL408 DOI [10.1121/10.0001216](https://doi.org/10.1121/10.0001216).

- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I. 2017. Attention is all you need. In: *Advances in Neural Information Processing Systems*. 5998–6008.
- Venkatesh S, Moffat D, Miranda ER. 2021. You only hear once: a YOLO-like algorithm for audio segmentation and sound event detection. *ArXiv preprint*. Available at [arXiv:2109.00962](https://arxiv.org/abs/2109.00962).
- Vesperini F, Gabrielli L, Principi E, Squartini S. 2018. A capsule neural networks based approach for bird audio detection. In: *Proceedings of the International Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE) 2018*.
- Vickers W, Milner B, Gorpincenko A, Lee R. 2021a. Methods to improve the robustness of right whale detection using CNNs in changing conditions. In: *Proceedings of EUSIPCO 2020*, European Association for Signal Processing, 106–110.
- Vickers W, Milner B, Risch D, Lee R. 2021b. Robust north Atlantic right whale detection using deep learning models for denoising. *The Journal of the Acoustical Society of America* **149**(6):3797–3812 DOI [10.1121/10.0005128](https://doi.org/10.1121/10.0005128).
- Vidaña-Vila E, Navarro J, Borda-Fortuny C, Stowell D, Alsina-Pagès RM. 2020. Low-cost distributed acoustic sensor network for real-time urban sound monitoring. *Electronics* **9**(12):2119 DOI [10.3390/electronics9122119](https://doi.org/10.3390/electronics9122119).
- Vignal C, Mathevon N, Mottin S. 2008. Mate recognition by female zebra finch: analysis of individuality in male call and first investigations on female decoding process. *Behavioural Processes* **77**(2):191–198 DOI [10.1016/j.beproc.2007.09.003](https://doi.org/10.1016/j.beproc.2007.09.003).
- Waddell EE, Rasmussen JH, Širović A. 2021. Applying artificial intelligence methods to detect and classify fish calls from the northern gulf of Mexico. *Journal of Marine Science and Engineering* **9**(10):1128 DOI [10.3390/jmse9101128](https://doi.org/10.3390/jmse9101128).
- Wang K, Wu P, Cui H, Xuan C, Su H. 2021. Identification and classification for sheep foraging behavior based on acoustic signal and deep learning. *Computers and Electronics in Agriculture* **187**(195):106275 DOI [10.1016/j.compag.2021.106275](https://doi.org/10.1016/j.compag.2021.106275).
- Webster MS, Budney GF. 2017. Chapter 11: sound archives and media specimens in the 21st century. In: Brown CH, Riede T, eds. *Comparative Bioacoustics: An Overview*. Oak Park: Bentham Science Publishers.
- Wolters P, Daw C, Hutchinson B, Phillips L. 2021. Proposal-based few-shot sound event detection for speech and environmental sounds with perceivers. Available at <http://arxiv.org/abs/2107.13616>.
- Wood CM, Kahl S, Chaon P, Peery MZ, Klinck H. 2021. Survey coverage, recording duration and community composition affect observed species richness in passive acoustic surveys. *Methods in Ecology and Evolution* **12**(5):885–896 DOI [10.1111/2041-210X.13571](https://doi.org/10.1111/2041-210X.13571).
- Xian Y, Pu Y, Gan Z, Lu L, Thompson A. 2016. Modified DCTNet for audio signals classification. *Acoustical Society of America Journal* **140**(4):3405 DOI [10.1121/1.4970932](https://doi.org/10.1121/1.4970932).
- Xiao S, Yan J, Farajtabar M, Song L, Yang X, Zha H. 2019. Learning time series associated event sequences with recurrent point process networks. *IEEE Transactions on Neural Networks and Learning Systems* **30**(10):3124–3136 DOI [10.1109/TNNLS.2018.2889776](https://doi.org/10.1109/TNNLS.2018.2889776).
- Xie J, Colonna JG, Zhang J. 2021a. Bioacoustic signal denoising: a review. *Artificial Intelligence Review* **54**(5):3575–3597 DOI [10.1007/s10462-020-09932-4](https://doi.org/10.1007/s10462-020-09932-4).
- Xie J, Hu K, Guo Y, Zhu Q, Yu J. 2021b. On loss functions and CNNs for improved bioacoustic signal classification. *Ecological Informatics* **64**(5):101331 DOI [10.1016/j.ecoinf.2021.101331](https://doi.org/10.1016/j.ecoinf.2021.101331).
- Xie J, Hu K, Zhu M, Guo Y. 2020. Bioacoustic signal classification in continuous recordings: syllable-segmentation vs sliding-window. *Expert Systems with Applications* **152**(6):113390 DOI [10.1016/j.eswa.2020.113390](https://doi.org/10.1016/j.eswa.2020.113390).

- Xie J, Hu K, Zhu M, Yu J, Zhu Q. 2019. Investigation of different CNN-based models for improved bird sound classification. *IEEE Access* 7:175353–175361 DOI 10.1109/ACCESS.2019.2957572.
- Xie J, Zhu M, Hu K, Zhang J, Hines H, Guo Y. 2021c. Frog calling activity detection using lightweight CNN with multi-view spectrogram: a case study on kroombit tinker frog. *Machine Learning with Applications* 7:100202 DOI 10.1016/j.mlwa.2021.100202.
- Yang W, Chang W, Song Z, Zhang Y, Wang X. 2021. Transfer learning for denoising the echolocation clicks of finless porpoise (*neophocaena phocaenoides sunameri*) using deep convolutional autoencoders. *The Journal of the Acoustical Society of America* 150(2):1243–1250 DOI 10.1121/10.0005887.
- Yip DA, Knight EC, Haave-Audet E, Wilson SJ, Charchuk C, Scott CD, Sólymos P, Bayne EM, Rowcliffe M, Quick N. 2019. Sound level measurements from audio recordings provide objective distance estimates for distance sampling wildlife populations. *Remote Sensing in Ecology and Conservation* 6(3):301–315 DOI 10.1002/rse2.118.
- Zeghidour N, Teboul O, Qutry FdC, Tagliasacchi M. 2021. LEAF: a learnable frontend for audio classification. In: *ICLR 2021*.
- Zhang K, Liu T, Song S, Zhao X, Sun S, Metzner W, Feng J, Liu Y. 2020. Separating overlapping bat calls with a bi-directional long short-term memory network. *Integrative Zoology* DOI 10.1111/1749-4877.12549.
- Zhang Z, Liu D, Han J, Schuller B. 2017. Learning audio sequence representations for acoustic event classification. *ArXiv preprint. Available at arXiv:1707.08729*.
- Zhong M, Castellote M, Dodhia R, Lavista Ferres J, Keogh M, Brewer A. 2020a. Beluga whale acoustic signal classification using deep learning neural network models. *The Journal of the Acoustical Society of America* 147(3):1834–1841 DOI 10.1121/10.0000921.
- Zhong M, LeBien J, Campos-Cerqueira M, Dodhia R, Ferres JL, Velev JP, Aide TM. 2020b. Multispecies bioacoustic classification using transfer learning of deep convolutional neural networks with pseudo-labeling. *Applied Acoustics* 166(25):107375 DOI 10.1016/j.apacoust.2020.107375.
- Zhong M, Torterotot M, Branch TA, Stafford KM, Royer J-Y, Dodhia R, Lavista Ferres J. 2021. Detecting, classifying, and counting blue whale calls with Siamese neural networks. *The Journal of the Acoustical Society of America* 149(5):3086–3094 DOI 10.1121/10.0004828.
- Znidarsic E, Towsey M, Roy W, Darling SE, Truskinger A, Roe P, Watson DM. 2020. Using visualization and machine learning methods to monitor low detectability species—the least bittern as a case study. *Ecological Informatics* 55:101014 DOI 10.1016/j.ecoinf.2019.101014.
- Zsebök S, Nagy-Egri MF, Barnaföldi GG, Laczi M, Nagy G, Vaskuti É, Garamszegi LZ. 2019. Automatic bird song and syllable segmentation with an open-source deep-learning object detection method—a case study in the collared flycatcher. *Ornis Hungarica* 27(2):59–66 DOI 10.2478/orhu-2019-0015.
- Zualkernan I, Judas J, Mahbub T, Bhagwagar A, Chand P. 2020. A tiny CNN architecture for identifying bat species from echolocation calls. In: *2020 IEEE/ITU International Conference on Artificial Intelligence for Good (AI4G)*. Piscataway: IEEE, 81–86.
- Zualkernan I, Judas J, Mahbub T, Bhagwagar A, Chand P. 2021. An AIoT system for bat species classification. In: *2020 IEEE International Conference on Internet of Things and Intelligence System (IoT&IS)*. Piscataway: IEEE, 155–160.