

NEAL: An open-source tool for audio annotation

Anthony Gibbons ^{Corresp., 1}, Ian Donohue ², Courtney Gorman ², Emma King ², Andrew Parnell ¹

¹ Hamilton Institute, Department of Mathematics and Statistics, Maynooth University, Kildare, Ireland

² Zoology, School of Natural Sciences, Trinity College Dublin, Dublin, Ireland

Corresponding Author: Anthony Gibbons
Email address: anthony.gibbons.2022@mumail.ie

Passive acoustic monitoring is used widely in ecology, biodiversity, and conservation studies. Data sets collected via acoustic monitoring are often extremely large and built to be processed automatically using Artificial Intelligence and Machine learning models, which aim to replicate the work of domain experts. These models, being supervised learning algorithms, need to be trained on high quality annotations produced by experts. Since the experts are often resource-limited, a cost-effective process for annotating audio is needed to get maximal use out of the data. We present an open-source interactive audio data annotation tool, NEAL (Nature+Energy Audio Labeller). Built using R and the associated Shiny framework, the tool provides a reactive environment where users can quickly annotate audio files and adjust settings that automatically change the corresponding elements of the user interface. The app has been designed with the goal of having both expert birders and citizen scientists contribute to acoustic annotation projects. The popularity and flexibility of R programming in bioacoustics means that the Shiny app can be modified for other bird labelling data sets, or even to generic audio labelling tasks. We demonstrate the app by labelling data collected from wind farm sites across Ireland.

1 NEAL: An open-source tool for audio 2 annotation

3 Anthony Gibbons¹, Anthony Gibbons ¹, Ian Donohue ², Courtney E.
4 Gorman ², Emma King², and Andrew Parnell ¹

5 ¹Hamilton Institute, Department of Mathematics and Statistics, Maynooth University,
6 Kildare, Ireland

7 ²Zoology, School of Natural Sciences, Trinity College Dublin, Dublin, Ireland

8 Corresponding author:

9 Anthony Gibbons¹

10 Email address: anthony.gibbons.2022@mumail.ie

11 ABSTRACT

12 Passive acoustic monitoring is used widely in ecology, biodiversity, and conservation studies. Data sets
13 collected via acoustic monitoring are often extremely large and built to be processed automatically using
14 Artificial Intelligence and Machine learning models, which aim to replicate the work of domain experts.
15 These models, being supervised learning algorithms, need to be trained on high quality annotations
16 produced by experts. Since the experts are often resource-limited, a cost-effective process for annotating
17 audio is needed to get maximal use out of the data. We present an open-source interactive audio data
18 annotation tool, *NEAL* (Nature+Energy Audio Labeller). Built using R and the associated Shiny framework,
19 the tool provides a reactive environment where users can quickly annotate audio files and adjust settings
20 that automatically change the corresponding elements of the user interface. The app has been designed
21 with the goal of having both expert birders and citizen scientists contribute to acoustic annotation projects.
22 The popularity and flexibility of R programming in bioacoustics means that the Shiny app can be modified
23 for other bird labelling data sets, or even to generic audio labelling tasks. We demonstrate the app by
24 labelling data collected from wind farm sites across Ireland.

25 INTRODUCTION

26 Passive acoustic recording is now a staple of ecological monitoring Ross and Allen (2014); Hagens et al.
27 (2018); Sugai et al. (2018); Rogers et al. (2013). Remote sensors can collect vast quantities of high quality
28 audio data at a cost affordable to both academic researchers and citizen scientists. However, the volume
29 of data collected can quickly surpass feasible manual labelling ability, and automatic methods are actively
30 being developed Morgan and Braasch (2021); Brunoldi et al. (2016); Baumgartner et al. (2019).

31 Application of deep learning, in particular convolutional neural networks (CNNs) Lecun et al. (1998),
32 to image processing for audio classification is growing in both popularity and effectiveness Hershey et al.
33 (2017). High performing models have been produced for bat classification Mac Aodha et al. (2018),
34 insects Yin et al. (2021), aquatic mammals Thomas et al. (2019); and bird calls, ranging from relatively
35 few Stowell et al. (2019) to tens Salamon et al. (2017) to even thousands Kahl et al. (2021) of classes.
36 While custom CNNs can be trained from scratch, a lack of large *labelled* bioacoustic datasets Baker and
37 Vincent (2019) can impede model performance.

38 In recent years, audio classification models have benefited from Transfer Learning Ntalampiras
39 (2018); Zhong et al. (2020). This is where a new model is created with the assistance of a neural network
40 pre-trained using a relatively large dataset for a similar task (such as another audio classification Hershey
41 et al. (2017) or even image classification Simonyan and Zisserman (2014)). The new model adjusts the
42 predictions of the larger model by performing further training with the small amount of labelled data
43 available for the task of interest. In both the custom CNN and transfer learning settings, some training
44 data are still required to tune to the specific application area, and routine test data should be annotated to
45 monitor performance over time. For novel tasks such as medical imaging classification, manufacturing
46 defect detection, agricultural yield prediction and marine image classification, domain specialists are

47 often needed to label the initial training data and evaluate the model output over time Sarma et al. (2021);
48 Ng (2021); Srivastava et al. (2022); Langenkämper et al. (2019).

49 Here, we focus on bird species detection, which requires experts to manually label files so that they
50 can be input into a supervised learning algorithm. We present an audio annotation tool that aims to reduce
51 the bottlenecks associated with audio annotation, improving the efficiency of the expert's time, which is
52 often at a premium, and providing finer granularity of labelled data (time-frequency limits, species, call
53 type, additional notes) so multiple classification tasks can be carried out on the same data simultaneously.

54 In much of the existing audio labelling software, supplemental information important to decision
55 making, such as the exact time of the recording, general location, geographic coordinates and shortest
56 distance to the coastline are often not readily available to the user, reducing the effectiveness of the
57 application as a decision tool. Giving annotators this temporal and spatial information can help contextu-
58 alise hard-to-classify sound clips. In relatively complicated soundscapes, such as wind farms or urban
59 environments, users often lack the flexibility to temporarily filter out noise and focus on the sound of
60 interest, further increasing labelling time. We include the ability to display various metadata and two
61 methods of filtering audio in our app.

62 We present *NEAL* (Nature+Energy Audio Labeller), an interactive Shiny app designed for audio data
63 annotation. It allows users to visually and audibly interpret audio files and label any sounds observed
64 and offers time and frequency granularity for precise labeling. The app incorporates metadata to inform
65 labellers, as well as labeller confidence for each annotation to provide quality annotations.

66 Some of the key strengths of *NEAL* are that:

- 67 • it can be used locally and on deployed servers;
- 68 • it has automatic frequency filtering of selected areas of the spectrogram to remove unwanted sounds
69 during analysis;
- 70 • it displays clear metadata for each audio file, giving context on the geography, habitat and time of
71 year recordings were taken;
- 72 • labellers can specify label confidence behind each annotation, as opposed to each individual
73 classification (e.g. species of a bird vocalisation) being assumed to have 100% confidence associated
74 with it;
- 75 • users can search through existing labels and navigate to those of interest;
- 76 • user annotations can be downloaded in bulk.

77 A comparison between *NEAL* and popular existing tools for bioacoustic annotation projects (Audacity
78 Team (2021); Marsland et al. (2019); Fukuzawa et al. (2020)) is shown in Table 1. While the Shiny app
79 was built to have a user-friendly layout for manual annotation of bird vocalisations with the option of
80 local or server-side use, it doesn't yet have vast functionality in terms of bulk classification or training
81 custom species classification models in-app.

82 We present the *NEAL* App (Figure 1), together with open-source code and sample data to allow for
83 further modification or deployment. Whilst our focus during development was on bird call detection, the
84 app can be easily changed to enable labelling of other audio tasks. This facilitates the adoption of the
85 Shiny app in projects where complicated soundscapes and data quality may differ greatly among sites and
86 equipment.

87 We provide the overall layout of a labelling project on the Shiny application, as well as a brief overview
88 of the procedures involved and some of the computational workarounds to avoid computation waiting
89 time. We then demonstrate a workflow of classifying bird species on wind farm sites across Ireland, with
90 step-by-step directions of how the app is utilized. We expand on the input and output formats of the data
91 to allow custom projects to be established easily. Source code for the *NEAL* App, as well as a link to a
92 working demo on an RStudio Server, is available at <https://github.com/gibbonal/neal>.

93 METHODS

94 User Interface (UI)

95 The app was built in R R Core Team (2022) using the Shiny Chang et al. (2021) framework. Shiny is
96 itself a package in R. No knowledge of HTML, CSS, or JavaScript is necessary to build a simple

Table 1. Comparison of NEAL with other free labelling software

Feature	NEAL	Audacity	AviaNZ	Koe
Date released	2022	2000	2019	2019
Platform(s) compatible	  	  	  	  
Built with	R, JavaScript	C, C++, Python	Python, C	Python, JavaScript
Label format	bounding boxes	bounding boxes ¹	bounding boxes	time segments
Label confidence slider	✓	✗	✗ ²	✗
Band-pass filter for selected spectrogram area	✓	✗ ³	✓	✗
Changeable class list	✓	✗	✓	✓
Dynamic class +/-	✓	✗	✓ ⁴	✗
Site metadata display	✓	✓	✓	✓
In-app label editing	✓	✓	✓	✓
Operates locally	✓	✓	✓	✓
Deployment to server	✓	✗	✗	✓ ⁵
Changeable spectrogram colour palettes	✓	✓	✓	✓
Bulk export annotations	✓	✗ ⁶	✗ ⁷	✓
Filter labels by multiple fields	✓	✗	✗	✓
Multiple concurrent (collaborating) users	✗ ⁸	✗	✗	✓
Visualise multiple spectrograms (comparison)	✗	✓	✓ ⁹	✓
Bulk classification	✗	✗ ¹⁰	✓	✓
Analysing sequence structure	✗	✗	✗	✓
Train a species recogniser in-app	✗	✗	✓	✗

¹ not by default

² colour codes

³ Not by default but there may be plugins available

⁴ new species added will appear in the list

⁵ also operates on its own server

⁶ individual *label tracks* for each file can be downloaded but must be named appropriately

⁷ per-species annotations can be exported from batch processed files

⁸ multiple users can work on a single server but this can be slow and updates are not immediate

⁹ Quick review mode after batch processing

¹⁰ Can view multiple files at once but these must be manually labelled

⁹⁷ application in Shiny, but small amounts were used here to enhance certain features. One of the many
⁹⁸ benefits of the Shiny framework is that its end-users do not need any knowledge of R programming to
⁹⁹ interact with the data and provide annotations. Shiny has already been used in developing ecology-related
¹⁰⁰ apps and decision-support tools, such as in species-habitat modeling Wszola et al. (2017), conservation
¹⁰¹ management Pascal et al. (2020) and forest structure assessment Silva et al. (2022).

¹⁰² NEAL makes use of several open-source R packages, such as Attali (2021); Bailey (2022); Pedersen
¹⁰³ et al. (2021); Chang (2021); Chang and Borges Ribeiro (2021); Granjon (2021); Perrier et al. (2022);
¹⁰⁴ Aden-Buie (2022); Littlefield and Fay (2021); Silva (2021); Ligges et al. (2018); Sueur et al. (2008);
¹⁰⁵ Wickham (2016); Garnier et al. (2021); Chang et al. (2020); Schloerke (2020); Wickham et al. (2021);
¹⁰⁶ Wickham (2019); Firke (2021); Xie et al. (2021). The most notable are:

¹⁰⁷ • *shinyjs*, which allows for custom JavaScript (JS) plugins, giving extended functionality using only
¹⁰⁸ a small amount of JS code. This includes toggling pause/play of the embedded audio element using
¹⁰⁹ the spacebar, disabling UI elements and reset buttons.

¹¹⁰ • *shinyBS* contains extra user-interface (UI) objects such as collapsible panels, giving the app a more
¹¹¹ compact layout. In particular, the metadata panel, label edit and label summary tables are contained
¹¹² within these collapsible panels and can be opened as needed.

¹¹³ • *shinyFiles* for file and directory navigation. This is especially useful for deployments to server where




Figure 1. Main components of the App User Interface. The main interactive element of the app is the **spectrogram**, providing a visual representation of the sound. This plot can be drawn on with boxes to filter audio or make annotations. The button on the left underneath the spectrogram is for **saving** the current selection on the plot, which will be labelled with the class from the grid of categories below. Various **audio playback** widgets are grouped together to the right of the save button. **File navigation** allows the user to proceed to the next file in the workflow folder, or select from a drop-down menu of available files, the second of which displays the number of annotations present in each file. The collapsed **side panel** on the left hand side has extra configuration settings expected to be changed at most once or twice per session. The **class list** is dynamic: the core (green) categories are selected from a drop-down list in the Configuration tab, while the miscellaneous (orange) categories are static in the current implementation. Custom categories can be added or removed using the category widgets below. Users can provide more information than just the primary category (e.g. bird species) of sound identified; these can include **call type**, free text **additional notes** and **label confidence**.

114 the folder structure may not be as familiar to users. The folder selected using `shinyDirButton`
 115 and `shinyDirChoose` is then searched for audio files.

116 • *shinydashboard* and *shinydashboardPlus* for the dashboard layout, sidebar and header tabs. Moving
 117 the less-used settings and widgets to the sidebar and grouping them into collapsible menus reduces
 118 clutter in the Shiny app's main body.

119 • *tuneR* for reading and writing audio files. It handles the audio files as `Wave` objects during
 120 preprocessing - including dB gain, segmentation and normalization - before passing them to
 121 `seewave` functions.

122 • *seewave* provides audio waveform manipulation functions, such as spectrogram computation. A
 123 noise-reduced or frequency-filtered spectrogram can be reconstructed as a `Wave` object using the
 124 `istft` function.

125 • Graphics are implemented using *ggplot2*. The spectrogram is rendered using `geom_raster` and
 126 bounding box annotations are drawn with `geom_rect` and `geom_label`. Additional colour
 127 palettes are provided via *viridis*.

128 A full list of packages employed is available in the source code.

129 Display

130 Upon opening the app, users are presented with the audio data from the chosen file in the form of
 131 a spectrogram, which highlights the sound intensity of many frequency levels over time (Figure 1).
 132 Spectrograms are one of the most visually perceptible forms for audio data Lin et al. (2012). The audio
 133 file can be played with an embedded audio player underneath the visual.

134 **General labelling workflow**

135 The standard use of the app is as follows:

- 136 ① Once the user has opened the app and clicked **start labelling**, the first audio file in their workflow
137 is loaded and the corresponding spectrogram is generated.
- 138 ② The **spectrogram parameters** such as colour palette and contrast may need to be adjusted until the
139 user is comfortable with the visual distinction of the sounds present.
- 140 ③ The user plays the audio and, when they come across a sound of interest, they **draw a bounding**
141 **box** around the vocalisation. This updates the audio player with a **temporary filtered audio file**,
142 keeping only the times and frequencies within the box drawn.
- 143 ④ Once they have identified the class of sound from the class list, they draw a tight box around the
144 vocalisation and click **save selection**. This will draw a permanent (unless deleted) bounding box
145 labelled with the given class. If they are unhappy with any of the bounding boxes, they can be
146 deleted using the **delete** button in the **label buttons** section.
- 147 ⑤ If a new class is being added or deleted, or the base species list is changed, this will affect the **class**
148 **list** - the grid of classes to choose from. This is only a minor computation and the user should not
149 see any delay unless there are several bounding boxes present in the plot.
- 150 ⑥ The user **continues annotating** the file by repeating steps ② to ⑤ until no more unlabelled
151 sounds of interest remain.
- 152 ⑦ The user clicks on the arrow to **proceed to the next file** and returns to step ②.

153 The Shiny framework uses reactive programming, meaning that inputs (in the user interface) changed
154 by the user automatically affect those outputs to which they are connected. This gives a smooth experience
155 for users, where the app does not have to be refreshed manually whenever settings are adjusted. In the case
156 of the app, rendering the spectrogram is the most computationally intensive process. Avoiding the plot
157 refreshing every time an unrelated input is changed is key to a seamless user experience. The back-end of
158 the app includes some modular code to split up dependencies (user inputs), which affect outputs of the
159 user interface. We elaborate on how this code works below.

160 Figure 2 shows the full layout of the workflow including some of the back-end components. These
161 components are invisible to the user but are required to reduce redundancy in computation. ① to ⑤
162 correspond to those points in the workflow above.

163 **Back-end workflow computation**

164 ***Spectrogram plot rendering***

165 While the computation of spectrograms is relatively fast, primarily due to efficiencies gained by imple-
166 menting the Fast Fourier Transform (FFT) algorithm Brigham and Morrow (1967), rendering the result in
167 R using ggplot is slow and creates a noticeable bottleneck for the annotation procedure. As an example,
168 a 15 second audio clip with a sample rate of 24 kHz will produce a spectrogram with dimensions of
169 128×5622 when applying the `seewave::spectro` function with an FFT window of 256 points and a
170 window overlap of 75% between two contiguous windows. The `geom_raster` function would then
171 have to plot $\approx 720,000$ equally-sized tiles, which is four times that with the linear interpolation needed
172 for adequate resolution, and then colour by the amplitude (in dB) and apply the chosen colour scale.

173 This computation must be done at least once for each audio file opened, and may be re-run several
174 times when adjusting FFT parameters, colour palette or zooming in on selected areas of the plot. The
175 common ggplot procedure of adding plot components (such as bounding boxes or frequency guides) to the
176 one plot is not ideal here due to the hindrance caused by frequent interactions with, and thus changes to,
177 the spectrogram plot. If even a small parameter was changed (such as adding a class to the class list UI),
178 the entire plot would have to be re-rendered. Avoiding these barriers to labeller efficiency is paramount to
179 the viability of the Shiny app as an audio annotation tool.

180 To reduce much of this unnecessary computation, we split the computation work of the spectrogram
181 generation into overlapping panels, shown in Figure 3. These are often independently generated or




Figure 2. Workflow diagram. Demonstration of the front and back-end workflow of the NEAL App. Any analyses or adjustments carried out by the user in the User Interface only affect the relevant processes in the back end. This keeps the user focused only on the elements being updated while reducing downtime compared to the event that the entire page is refreshed every time. Each of the steps ① to ⑤ in the front end—the actions and changes performed by the user via the UI—have corresponding processes in the background which are in decreasing order of computational demand. The advantage is that these back-end processes have little to no overlap, meaning they can be run separately when the front-end changes are unrelated. If the change is minor (far right, in white) we do not want this to result in an avoidable refresh involving large computations (shown to the left of the figure in orange).

182 refreshed by different widgets in the UI. The panels Figure 3c and Figure 3d have a transparent background
 183 and identical axes and padding to the main spectrogram Figure 3b. This alignment ensures that all plot
 184 layers match exactly.




Figure 3. Spectrogram plot. The full spectrogram plot (a) is a combination of overlapping panels. From back to front: (b) highlights the main spectrogram, which is the slowest to render; (c) shows the labels panel displaying bounding box annotations; and (d) emphasizes the time/frequency guides which gauge the time and frequency ranges the sound clip occupies, as well as tracing the shape of the bounding box should a label be saved.

185 If the user wants to get a closer look at a particular sound in the spectrogram, a section can be zoomed
 186 in on by selecting the area of interest and double clicking on the selected area. The plot should then

187 re-render to show only the selected ranges in time and frequency. The parameters in FFT settings can be
 188 tuned further to investigate more complex sounds at this level of magnification. In particular, increasing the
 189 FFT window size can put more emphasis on frequency resolution and less on time resolution. Decreasing
 190 this parameter has the opposite effect. To zoom out, the user double clicks any point on the plot, or clicks
 191 the Reset Zoom button.

192 **Filtered audio production**

193 To display the audio to end-users, it is converted to a spectrogram, a visual representation of sound. This
 194 is achieved using the Short-Time Fourier Transform (STFT) algorithm Allen (1977), which applies the
 195 FFT algorithm on successive windows of the audio waveform. The output of this transform is a 2D array
 196 of complex numbers, before the modulus is taken and the results (scaled to dB) are printed to the screen.

197 When selecting the audio area of interest, it is often helpful to remove or significantly reduce obvious
 198 noise which can hinder clear identification. By default, the app keeps the selected area of the (complex)
 199 spectrogram the same, sets all values outside this area to zero and reconstructs a filtered audio clip
 200 using the Inverse Short-time Fourier Transform (ISTFT). When the user selects some time/frequency
 201 range, only those time/frequency segments will be played. Alternatively, if the audio frequency range
 202 slider is adjusted, the entire duration of the audio file is included, but only those frequencies within the
 203 filtered range. A similar reconstruction is performed through the spectrogram noise reduction techniques
 204 carried out by noise reduction of the complex spectrogram data. Figure 4 illustrates a comparison of the
 205 spectrogram and audio player appearances.




Figure 4. Audio Player and spectrogram. The three versions of the audio UI element the user will see (in Chrome), with an example corresponding spectrogram. The spectrogram has been zoomed in to 0-5kHz (y-axis) and 3-6 seconds (x-axis) to more clearly illustrate the filtering features: (a) shows the default audio player and spectrogram shown to user; (b) displays the audio player and spectrogram with some frequency filtering applied, i.e., the audio within the greyed out box is reduced or removed and the remaining audio is reconstructed; (c) conveys the player and spectrogram when row-wise noise reduction is applied. (a) shows the raw audio file from the workflow folder, while (b) and (c) show the temporary filtered audio files.

206 The refined audio can often sound unnatural or distorted when the phase of the sounds outside the
 207 selection are collapsed to zero, so the app has an alternative option to reduce the magnitude of the audio
 208 outside the selected region by a factor of 100. The user deselects *Zero Audio outside Selected* in the
 209 **Spectrogram Settings**, and when the plot is clicked outside the selected box, the audio is reset.

210 **Bounding box display**

211 The annotations created by the NEAL app are bounding boxes: rectangular regions defined using *time* and
 212 *frequency* coordinates by the lower left and upper right corners of the object. The purpose of providing a
 213 bounding box as the main label is to maximize the detail captured in a sound event annotation project.
 214 Figure 5 describes the varying definitions of sound classification projects, based on the diagram in Stowell
 215 (2022).

216 Classifying the sound clip as a whole (Figure 5a) gives no indication of the duration or regularity (how
 217 often it occurs) of a sound event, or whether multiple species are present in a single audio file. Annotation
 218 setups where labels transcribe the time dimension with either no frequency resolution (Figure 5b) or with
 219 a small number of set frequency bins is called sound event detection. This does not have the flexibility to




Figure 5. Three common types of sound event detection and classification. (a) shows the most common setup where an entire audio clip is labelled either using binary (e.g. bird/not bird) or multi-class classification. It does not give any explicit temporal resolution; (b) displays sound event detection (SED). Labels typically have a fixed width (time duration), but the resulting automatic detection models classify successive time windows (small fractions of the file duration) for each class, with consecutive windows belonging to the same class being joined together to form a larger sound event. In the above example, the blackbird classification is longer than the other two detections. This illustration does not include any frequency resolution (the annotations span the entire y-axis) but a similar setup would have a small number of frequency bins where a species is expected to occupy a given range; (c) is object detection setup using bounding boxes. This clearly shows highest resolution detail among the annotation examples shown, with the possibility of overlapping sound events in both time and frequency domains. In this toy example, the two wren annotations have differing frequency ranges which gives richer frequency information. Adapted from Stowell (2022).

220 account for overlapping vocalisations in the frequency dimension, or a large number of classes where the
 221 frequency bins are difficult to distinguish.

222 Time- and frequency-specific bounding boxes (Figure 5c) therefore give millisecond level timestamps
 223 (both start and end) for sound events, as well as maximum, minimum and median frequencies of the
 224 labelled sound clip. If not used directly for object detection Zsebők et al. (2019), frequency metrics can
 225 serve as auxiliary input to train machine learning models for species detection Lostanlen et al. (2019).

226 **Class list update**

227 The class list is a group of buttons representing the available classes to label the identified sounds in the
 228 audio files. A predefined list of classes is important for consistency within and among labellers, reducing
 229 the time taken to repetitively type class names, as well as minimizing the need to correct typos during
 230 post hoc analysis.

231 The list is displayed in CSS `grid` (default) or `flex` containers. The flexbox layout makes the buttons
 232 as wide as the text for each class and automatically wraps row-wise. The grid layout has a custom number
 233 of columns, which can be adjusted using the *Number of Columns* parameter in **Other Settings**. The grid
 234 layout is neater while the flexbox layout is more compact. The differences can be seen in Figure 6.

235 The different groups of classes are colour coordinated to distinguish the main classes, miscellaneous
 236 categories and manually added classes. Core classes in the selected site species list have green borders,
 237 while classes manually added to the list are in blue. Miscellaneous sounds such as human, insect or
 238 weather noise are outlined in orange. If a label is present in the file that is not in the collective class list,
 239 its bounding box is coloured grey in the plot.

240 When a new class is added to the list, the list is searched and if it is not already present, the class is
 241 appended. If the class is found, an error message is thrown saying the class is already present in the list.

242 The British Trust for Ornithology has a list of 2-letter species codes for over 250 bird species. The
 243 app has the ability for users to view classes in the displayed list with their corresponding code, if
 244 they have a matching species in the BTO list. `bto_codes.csv` acts as a lookup table, which was
 245 sourced from [https://www.bto.org/sites/default/files/u16/downloads/forms_](https://www.bto.org/sites/default/files/u16/downloads/forms_instructions/bto_bird_species_codes.pdf)
 246 `instructions/bto_bird_species_codes.pdf`. It has two columns, **bto_code** and **species_name**:
 247 the first column has the two letter code associated with the species name in the second column. An
 248 example of the conversion is shown on Figure 6.




Figure 6. Class list layouts. Shows the possible class list layouts for a small example list of classes, as well as the miscellaneous categories and one manually added class. (a) is the default layout of the classes in a grid, with the number of columns decreased to 3 for display purposes. (b) is the same list in flexbox layout. Note the reduced amount of empty space from shorter class names. (c) shows the same flexbox layout with class names changed to their corresponding BTO codes, where applicable. The space is compacted further which becomes more apparent with longer class lists.

249 Other workflow features

250 At the top left of the screen, there is a hidden/collapsed sidebar that contains multiple adjustable drop-down
 251 menus of settings. These include **Configuration:** general setup inputs for the annotation project, such as
 252 the directory containing the audio files, uploading more files to the user's folder, uploading annotations
 253 from previous projects and the column of `species_list.csv` with the relevant classes for labelling,
 254 which is expanded upon in Extension to other audio labelling projects. **Sound Settings:** adjustable inputs
 255 for processing incoming audio files, including dB gain (or amplification factor) and the length of audio
 256 clips to display in the spectrogram (the default is 15 seconds). **Spectrogram Settings:** adjust the visual
 257 aspects of the spectrogram plot, e.g. the colour palette, plot height and whether to include time and
 258 frequency guidelines for boxes drawn. **FFT Settings:** parameters for the Short-Time Fourier Transform
 259 (STFT) calculation via the `spectro` function. **Other Settings:** miscellaneous widgets for adjusting the
 260 class list, label summary table and label edit table.

261 Metadata is a helpful addition to annotation workflows, providing context for users who may not have
 262 been involved in the project design or data collection phases. The app attempts to link audio files to a
 263 predefined list of recorders that were deployed in order to give information to end-users on aspects of
 264 the study site such as habitat type, location and distance to coastline, then displays them in a collapsible
 265 panel to be consulted as desired. For example, the location of the recordings may give important context
 266 for detecting species, in particular for migratory birds. Addressing the information needs of labellers by
 267 providing all available metadata for context can improve the accuracy of the classifications Mortimer and
 268 Greene (2017).

269 With a key objective of the app being labeller efficiency, using keyboard shortcuts instead of repetitive
 270 mouse drags and clicks enables increased precision and productivity. This was not included in the main
 271 procedure but is a useful addition to efficient labeller workflow. Available hotkeys include saving and
 272 deleting labels, navigating to the previous or next file, pausing and playing the current audio file and
 273 adding or deleting classes in the class list.

274 The label edit table is an interactive revision tool that describes all labels for the current file. This can
 275 be enabled in the Other Settings tab in the sidebar, where it then appears at the bottom of the app. It will
 276 only display when the current file has at least one label. Some of the fields from saved annotations can be
 277 edited, such as the time/frequency limits of the bounding boxes, class identified and label confidence.

278 The summary table is another revision tool providing an abridged description of all files in the
 279 workflow folder. It can also be enabled through the Other Settings tab. The user can filter these files by
 280 name, number of labels present per file and even split the label counts per file by the classes present. Users
 281 can quickly navigate between a large number of files to find other examples of previously labelled species.

282 When the labelling project is complete and the user wishes to access their labels (especially on a

server), they can be exported using the Download button. This is located in the **Label buttons** section towards the bottom of the main body of the app. These labels are exported to the Downloads folder of the user's local machine. This download feature works both on a server or local deployment.

CASE STUDY

All examples and results in this paper come from data recorded for the Nature+Energy project MaREI, the SFI Research Centre for Energy, Climate and Marine (2022), part of which involves developing an acoustic monitoring system for on-shore wind farm sites in Ireland. Over 1,700 hours of audio was collected between April and July 2022. Wildlife Acoustics Song Meter Mini Bat (<https://www.wildlifeacoustics.com/products/song-meter-mini-bat>) units were mounted at approximately 1.5 metres above the ground, placed between a wind turbine and some linear feature such as hedgerows or gaps in tree cover. For the first stage of the project, we selected five recording sites (one recorder per location) representing the variety of habitat types present.

Nature+Energy is a collaboration between academic researchers and industry partners that aims to measure and enhance biodiversity at onshore wind farms throughout Ireland. The project focuses on exemplar wind farm sites, located in typical habitats where wind farms are situated in Ireland.

Wind energy is undergoing substantial growth in Ireland, providing over 30 percent of its energy needs in 2020 and expected to reach 80 percent by 2030 Department of the Environment, Climate and Communications; Department of the Taoiseach (2019). With the increase in wind turbines comes the possibility of biodiversity degradation, with bird and bat populations being particularly vulnerable to collision mortality Richardson et al. (2021); Choi et al. (2020). The need for monitoring systems is crucial to inform habitat management planning, potentially reducing habitat decline, fatalities and issues specific to each site. Through the Nature + Energy project, a number of acoustic sensors were set up as a non-invasive monitoring method at select study sites across the country. A description of each recorder, and the sites where each was deployed, is given in Table 2.

Table 2. Wind Farm Study Site info

Site	Land Use Type	Area (Ha)	No. Turbines	Hours recorded	
				Acoustic	Ultrasonic
Carnsore Point	Agricultural (coastal, pasture)	80	14	547	19.6
Cloosh Valley	Commercial forestry on peat substrate	1378	36	438	11.9
Rahora	Agricultural (inland, arable)	24	5	170	14.2
Richfield	Agricultural (coastal, arable)	112	18	850	52.4
Teevurcher	Agricultural (inland, pasture)	60	5	156	1.7

Passive acoustic monitoring provides rich insight into species abundance, and temporal and spatial granularity Warren et al. (2021). Due to the nature of where the recorders are located, anthropogenic noise from the wind turbines themselves has an adverse effect on data quality, especially for labelling species with calls occupying lower frequencies. There is a need for frequency granularity to conceal other species or noise sources occurring simultaneously. Metadata, as mentioned in Other workflow features, is a helpful addition to give labellers insight on aspects of the study site such as habitat type, time of day/year, and location.

Project annotation procedure

A version of the app was deployed to an Rstudio Connect server integrating user login via Auth0 (auth0.com). Users (bird experts) were assigned a sample of audio files from multiple wind farm sites that were believed to contain sounds of interest. A summary of the actions involved in labelling a selected segment of audio with a bounding box (Steps 3 to 6 in Detailed labeller instructions) is outlined in Figure 7 below.

Results of labelling

A small sample of the results obtained from labelling are shown in Figure 8. These examples illustrate the diversity of bird vocalisations found on wind farms, even within the same species. The structure, duration and frequency ranges of bird vocalisations are shown to vary, the latter of which could not be extracted




Figure 7. Spectrogram labelling example.

324 via one dimensional labels. Noise from the wind turbines on the study sites is especially noticeable in the
 325 examples on the left.

326 A sample of the count data for non noise-related classes labelled are shown in Table 3. There are
 327 several zeroes in the table, indicating the diversity of the selected study sites, which cover different
 328 geographic regions and habitat types. Further details on the bird species identified by NEAL are during
 329 the case study are included in Bird conservation status in the appendix. It displays bird conservation data
 330 and information on their distribution across Ireland, and was obtained from Gilbert et al. (2021).

331 Some common species such as the Blackbird, Meadow Pipit and Robin are present across almost
 332 all sites, while the Yellowhammer is only present at one site due to it being a rare species that feeds on
 333 cereal crops. The annotations that were discerned to be birds but could not specifically be identified
 334 (representing about 8% of non-noise related labels collected, not included in the table) can be returned to
 335 the same labeller for deeper analysis or given to another labeller for a second opinion. Both could make
 336 use of the output of a trained species recognition machine learning model to inform the decision.

337 Counts of annotations from the first set of data gathered from deployed recorders are shown in Figure 9.
 338 It displays the activity periods of selected species from the recordings. The recorders were deployed
 339 from April-July, which in Ireland is a period when day length is at its greatest; on June 21st, the summer
 340 solstice and longest day of the year, Ireland receives approximately 17 hours of daylight. This figure
 341 provides an example of the type of data that can be extracted from acoustic recordings, and it could
 342 potentially be used to aid management decisions, especially for particularly vulnerable species such as
 343 raptors. Some species such as the Yellowhammer tend to be identified throughout the day, whereas others
 344 such as the Robin and Chaffinch occur in shorter bursts of activity interspersed with lulls. If multiple
 345 recorders were placed on the sites, it could also inform how bird species are using the site throughout
 346 the day - for example, to determine whether they are just passing through the study area or are foraging
 347 multiple locations throughout the day.

348 EXTENSION TO OTHER AUDIO LABELLING PROJECTS

349 Our app is built so that it can be configured for other audio labelling tasks beyond the bird species
 350 classification demonstrated in our case study. Here we provide some guidance on setting up the input files
 351 if they are being expanded to similar areas or changed to cover other domains.

352 To proceed, a new user merely needs to download (or clone) the GitHub repository. Upon running, the
 353 app automatically will create a new user if one is not found in the `www` directory. The user can then add au-
 354 dio files into the `XXX` directory. The expected audio filename format is `RECORDERNAME_YYYYMMDD_HHMMSS.wav`.
 355 If the files are split up into smaller segments, the additional convention `_start_MM_SS` should be




Figure 8. Example bounding box annotations. Example annotations of sample audio collected from wind farm sites across Ireland. (a), (c) and (e) in the first column show different vocalisations of the chaffinch. (a) can be readily identified both visually and through listening to the audio, whereas (c) and (e) are increasingly difficult to distinguish from noise. Adjusting the contrast or applying noise reduction in the spectrogram settings may help with identifying examples similar to these. The structure of the vocalisations in (a) and (c) are quite dissimilar despite coming from the same species, which can be optionally captured using the **call_type** input. (b) shows a distinct robin vocalisations which was easily identified given little background noise. (d) includes three separate vocalisations of a linnet, since the sounds are separated by a natural pause. (f) contains three vocalisations of a blackbird. This sits partially in the range of wind turbine noise, and would be difficult to identify if the noise was more prominent or the bird was located further from the recorder.




Figure 9. Distribution of selected labelled species by timestamp in recording.

356 appended before the filename extension to indicate the number of minutes/seconds into the recording that
 357 they start. This allows the date and time to be parsed from the above filename format and shown in the
 358 metadata panel.

Table 3. Example raw species counts from labelled data

Class	Carnsore Point	Cloosh Valley	Rahora	Richfield	Teevurcher	Conservation Status
Blackbird	37	0	107	73	96	Green
Chaffinch	0	141	0	32	0	Green
Dunnock	26	1	0	12	0	Green
Goldfinch	12	0	0	0	0	Green
Hooded Crow	0	0	13	1	16	Green
House Sparrow	0	0	20	0	0	Amber
Linnet	49	0	15	42	2	Amber
Meadow Pipit	17	1	33	15	51	Red
Pied Wagtail	7	0	1	4	2	Green
Robin	0	62	4	49	7	Green
Rook	0	0	10	83	0	Green
Sedge Warbler	11	0	0	1	0	Green
Starling	0	0	37	0	0	Amber
Stonechat	30	0	0	0	0	Green
Wren	26	0	12	143	0	Green
Yellowhammer	0	0	60	0	0	Red
Blue Tit	0	0	0	7	0	Green
Buzzard	0	0	2	0	1	Green
Coal Tit	0	6	0	1	0	Green
Curlew	1	0	0	0	0	Red
Goldcrest	0	3	0	0	0	Green
Grasshopper Warbler	3	0	0	0	0	Amber
Great Tit	0	0	0	5	0	Green
Jackdaw	0	0	0	0	1	Green
Magpie	0	0	1	0	0	Green
Mallard	3	0	0	1	0	Green
Oystercatcher	1	0	0	0	0	Amber
Pheasant	0	0	1	0	0	Green
Skylark	1	0	0	0	0	Amber
Song Thrush	3	0	0	0	0	Green
Swallow	1	0	0	0	9	Amber
Teal	0	1	0	0	0	Amber
Woodpigeon	0	0	4	1	0	Green

The table details the most and least common species identified by NEAL. The top species were those with a total count of more than 10, while less frequent species are shown below. Some bird species, such as the Robin (*Erithacus rubecula*), regularly exploit a wide variety of habitats and food resources, and thus are common and widely distributed. Others however, have more specific habitat requirements. Oystercatchers (*Haematopus ostralegus*), for example, are limited to coastal habitats where they feed on large invertebrates such as mussels. NEAL identified several birds of conservation concern in Ireland (Status **Red**), including the Yellowhammer, Curlew, and Meadow Pipit.

359 Users can upload annotations from previous labelling projects using the upload button in the Configu-
 360 ration tab in the sidebar panel. These labels could be from manual work or generated by an ML model or
 361 some other automatic method. The CSV label data should at least contain the following:

- 362 • **date_time:** Date and time label was made in yyyy-mm-dd HH:MM:SS format
- 363 • **file_name:** Name of the corresponding audio file in the Data Folder
- 364 • **start_time:** Start time (in seconds) of the bounding box
- 365 • **end_time:** End time (in seconds) of the bounding box
- 366 • **start_freq:** Start frequency (in kHz) of the bounding box

- 367 • **end_freq**: End frequency (in kHz) of the bounding box
- 368 • **class_label**: Class of sound (e.g. bird species) assigned to this bounding box

369 Labels will be stored in `labels/labels_<username>.csv` if such a folder by this name exists
370 in the above directory. Otherwise the labels are saved in `labels/labels_tmp.csv`. The information
371 contained in each label includes: the time the annotation was made; for which audio file; the start and end
372 time/frequency bounds of each bounding box; the class name; label confidence; labeller and optional call
373 type and additional notes.

374 The app uses multiple read-only input data tables to accompany the folder of audio files. These
375 include the recorder (study site) location data, species lists and BTO codes. If any of these are edited
376 or replaced for a new annotation project the files should be located in the parent folder of the GitHub
377 repository, and saved in `.csv` format using `UTF-8-BOM` encoding. This ensures that column names and
378 text strings are read correctly.

379 NEAL allows for multiple species lists to be included, all of which are stored in `species_list.csv`.
380 The file has one column for each site being studied, with the first entry of each column being the site
381 name. Columns can have different numbers of rows (species), i.e. all those one expects to find at the
382 given site. The species names can be stored in the column in any order as they are sorted alphabetically at
383 run time. Columns can be appended to it by uploading via a widget in the Configuration tab. Each of the
384 columns in the uploaded data not already present in the species data are appended.

385 The `location_list.csv` file has a list of recorders that were deployed along with accompanying
386 metadata about the sites studied. If the current audio file in the above format matches a recorder in the list,
387 the following columns of the file are passed to the metadata tab:

- 388 • **recorder_name**: prefix for audio file names recorded by this device
- 389 • **lat**: latitude of the recorder for the study period, in decimal degrees
- 390 • **long**: longitude of the recorder for the study period, in decimal degrees
- 391 • **location_name**: name of the study site
- 392 • **location_county**: county of study site
- 393 • **habitat_type**: primary habitat type of the study site
- 394 • **dist_to_coastline**: approximate distance to the nearest coastline in kilometres

395 If any of this information is unavailable, the column name in the metadata panel will still appear but the
396 body of text will be blank. Extra columns can be added to the file, where they will be printed verbatim to
397 the metadata panel.

398 CONCLUSION

399 NEAL is an open-source application for visually examining and annotating audio data. The tool was
400 designed with the primary goal of improving labeller efficiency and consistency. Vocalisations are tagged
401 with comprehensive annotations providing time and frequency detail, accurate to several decimal places.
402 Call type and an open text field for notes attempt to capture multiple modes of information, with the
403 possibility of performing multiple labelling tasks simultaneously. A labeller providing classifications of
404 different bird species, as well as the call type and identifying sources of noise in each sound clip, has
405 generated three potentially separate sets of labels without expending more effort than generating only one.
406 These can each serve as targets for a machine learning algorithm to predict.

407 The Shiny package for R allows for an interactive front-end, while its reactive ability reduces
408 unnecessary computational expense. Its no-code interface allows domain experts to interact with the data
409 without any knowledge of R programming. While the app was designed mainly for classifying bird audio,
410 it can be expanded to projects with data focusing on bats, frogs, small mammals and insects - all are
411 popular in bioacoustics research.

412 There is room for extending the functionality of the app. Expanded contextual information such as
413 weather data, proximity to Special Areas of Conservation and habitat type could prove particularly useful.

414 New interactive tables and visualisations (even the static summary visualisations presented in this paper)
415 could be used to investigate outliers and labeller inconsistency. The app being open-source means others
416 can contribute or request features, driving innovation for future releases. R is one of the most popular
417 programming languages in bioacoustics and new features to the app can be easily added on through the
418 Shiny app's modular design.

419 ACKNOWLEDGEMENTS

420 We thank our labellers Harry Hussey, David Kelly, Seán Ronayne and Mark Shorten, who annotated the
421 majority of the audio files, as well as providing beta-testing of the app in its early stages. We also thank
422 the on-site personnel and surveyors for carrying out recorder maintenance, i.e. data and battery transfers,
423 as the wind farms are widely spread across Ireland. Finally, we would like to acknowledge the significant
424 work done by Dr. Aoibheann Gaughran on the Nature + Energy project, and for her encouragement and
425 comments on the Shiny app from its inception. This paper is dedicated to her memory.

426 FUNDING INFORMATION

427 This publication was produced by the Nature+Energy Project, funded by Science Foundation Ireland
428 (12/RC/2302_P2) and MaREI, the SFI Research Centre for Energy, Climate and Marine Research and
429 Innovation, with additional funding from Microsoft and the SFI CONNECT Centre for Future Networks
430 and Communications (13/RC/2077_P2). In addition, Andrew Parnell's work was supported by: a Science
431 Foundation Ireland Career Development Award (17/CDA/4695); SFI Centre for Research Training in
432 Foundations of Data Science 18/CRT/6049, and SFI Research Centre awards I-Form 16/RC/3872 and
433 Insight 12/RC/2289_P2. For the purpose of Open Access, the author has applied a CC BY public copyright
434 licence to any Author Accepted Manuscript version arising from this submission.

435 REFERENCES

- 436 Aden-Buie, G. (2022). *shinyThings: Reusable Shiny Modules and Other Shiny Things*. R package version
437 0.4.0.
- 438 Allen, J. (1977). Short term spectral analysis, synthesis, and modification by discrete fourier transform.
439 *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 25(3):235–238.
- 440 Attali, D. (2021). *shinyjs: Easily Improve the User Experience of Your Shiny Apps in Seconds*. R package
441 version 2.1.0.
- 442 Audacity Team (Copyright ©1999-2021). Audacity®.
- 443 Bailey, E. (2022). *shinyBS: Twitter Bootstrap Components for Shiny*. R package version 0.61.1.
- 444 Baker, E. and Vincent, S. (2019). A deafening silence: a lack of data and reproducibility in published
445 bioacoustics research? *Biodiversity Data Journal*, 7:e36783.
- 446 Baumgartner, M. F., Bonnell, J., Van Parijs, S. M., Corkeron, P. J., Hotchkin, C., Ball, K., Pelletier, L.-P.,
447 Partan, J., Peters, D., Kemp, J., Pietro, J., Newhall, K., Stokes, A., Cole, T. V. N., Quintana, E., and
448 Kraus, S. D. (2019). Persistent near real-time passive acoustic monitoring for baleen whales from a
449 moored buoy: System description and evaluation. *Methods in Ecology and Evolution*, 10(9):1476–1489.
- 450 Brigham, E. O. and Morrow, R. E. (1967). The fast fourier transform. *IEEE Spectrum*, 4(12):63–70.
- 451 Brunoldi, M., Bozzini, G., Casale, A., Corvisiero, P., Grosso, D., Magnoli, N., Alessi, J., Bianchi, C. N.,
452 Mandich, A., Morri, C., Povero, P., Wurtz, M., Melchiorre, C., Viano, G., Cappanera, V., Fanciulli, G.,
453 Bei, M., Stasi, N., and Taiuti, M. (2016). A permanent automated real-time passive acoustic monitoring
454 system for bottlenose dolphin conservation in the mediterranean sea. *PLOS ONE*, 11(1):1–35.
- 455 Chang, W. (2021). *shinythemes: Themes for Shiny*. R package version 1.2.0.
- 456 Chang, W. and Borges Ribeiro, B. (2021). *shinydashboard: Create Dashboards with 'Shiny'*. R package
457 version 0.7.2.
- 458 Chang, W., Cheng, J., Allaire, J., Sievert, C., Schloerke, B., Xie, Y., Allen, J., McPherson, J., Dipert, A.,
459 and Borges, B. (2021). *shiny: Web Application Framework for R*. R package version 1.7.1.
- 460 Chang, W., Luraschi, J., and Mastny, T. (2020). *profvis: Interactive Visualizations for Profiling R Code*.
461 R package version 0.3.7.
- 462 Choi, D. Y., Wittig, T. W., and Kluever, B. M. (2020). An evaluation of bird and bat mortality at wind
463 turbines in the northeastern united states. *PLOS ONE*, 15(8):1–22.

- 464 Department of the Environment, Climate and Communications; Department of the Taoiseach (2019).
465 Climate action plan 2019.
- 466 Firke, S. (2021). *janitor: Simple Tools for Examining and Cleaning Dirty Data*. R package version 2.1.0.
- 467 Fukuzawa, Y., Webb, W. H., Pawley, M. D., Roper, M. M., Marsland, S., Brunton, D. H., and Gilman, A.
468 (2020). Koe: Web-based software to classify acoustic units and analyse sequence structure in animal
469 vocalizations. *Methods in Ecology and Evolution*, 11(3):431–441.
- 470 Garnier, Simon, Ross, Noam, Rudis, Robert, Camargo, Pedro, A., Sciaini, Marco, Scherer, and Cédric
471 (2021). *viridis - Colorblind-Friendly Color Maps for R*. R package version 0.6.2.
- 472 Gilbert, G., Stanbury, A., and Lewis, L. (2021). Birds of conservation concern in ireland 4:2020-2026.
473 *Irish Birds Number 43 2021*, 43:1–22.
- 474 Granjon, D. (2021). *shinydashboardPlus: Add More 'AdminLTE2' Components to 'shinydashboard'*. R
475 package version 2.0.3.
- 476 Hagens, S. V., Rendall, A. R., and Whisson, D. A. (2018). Passive acoustic surveys for predicting species'
477 distributions: Optimising detection probability. *PLOS ONE*, 13(7):1–16.
- 478 Hershey, S., Chaudhuri, S., Ellis, D. P. W., Gemmeke, J. F., Jansen, A., Moore, R. C., Plakal, M., Platt,
479 D., Saurous, R. A., Seybold, B., Slaney, M., Weiss, R. J., and Wilson, K. (2017). Cnn architectures
480 for large-scale audio classification. In *2017 IEEE International Conference on Acoustics, Speech and*
481 *Signal Processing (ICASSP)*, pages 131–135.
- 482 Kahl, S., Wood, C. M., Eibl, M., and Klinck, H. (2021). Birdnet: A deep learning solution for avian
483 diversity monitoring. *Ecological Informatics*, 61:101236.
- 484 Langenkämper, D., Simon-Lledó, E., Hosking, B., Jones, D. O. B., and Nattkemper, T. W. (2019). On the
485 impact of citizen science-derived data quality on deep learning based classification in marine images.
486 *PLOS ONE*, 14(6):1–16.
- 487 Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document
488 recognition. *Proceedings of the IEEE*, 86(11):2278–2324.
- 489 Ligges, U., Krey, S., Mersmann, O., and Schnackenberg, S. (2018). *tuneR: Analysis of Music and Speech*.
- 490 Lin, K.-H., Zhuang, X., Goudeseune, C., King, S., Hasegawa-Johnson, M., and Huang, T. S. (2012).
491 Improving faster-than-real-time human acoustic event detection by saliency-maximized audio visual-
492 ization. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*,
493 pages 2277–2280.
- 494 Littlefield, T. and Fay, C. (2021). *keys: Keyboard Shortcuts for 'shiny'*. R package version 0.1.1.
- 495 Lostanlen, V., Salamon, J., Farnsworth, A., Kelling, S., and Bello, J. P. (2019). Robust sound event
496 detection in bioacoustic sensor networks. *PLOS ONE*, 14(10):1–31.
- 497 Mac Aodha, O., Gibb, R., Barlow, K. E., Browning, E., Firman, M., Freeman, R., Harder, B., Kinsey, L.,
498 Mead, G. R., Newson, S. E., Pandourski, I., Parsons, S., Russ, J., Szodoray-Paradi, A., Szodoray-Paradi,
499 F., Tilova, E., Girolami, M., Brostow, G., and Jones, K. E. (2018). Bat detective—deep learning tools
500 for bat acoustic signal detection. *PLOS Computational Biology*, 14(3):1–19.
- 501 MaREI, the SFI Research Centre for Energy, Climate and Marine (2022). Nature + energy.
- 502 Marsland, S., Priyadarshani, N., Juodakis, J., and Castro, I. (2019). Avianz: A future-proofed program for
503 annotation and recognition of animal sounds in long-time field recordings. *Methods in Ecology and*
504 *Evolution*, 10:1189–1195.
- 505 Morgan, M. and Braasch, J. (2021). Long-term deep learning-facilitated environmental acoustic monitor-
506 ing in the capital region of new york state. *Ecological Informatics*, 61:101242.
- 507 Mortimer, J. A. and Greene, T. C. (2017). Investigating bird call identification uncertainty using data from
508 processed audio recordings. *New Zealand journal of ecology*, 41(1):126–133.
- 509 Ng, A. (2021). MlOps: From model-centric to data-centric ai.
- 510 Ntalampiras, S. (2018). Bird species identification via transfer learning from music genres. *Ecological*
511 *Informatics*, 44:76–81.
- 512 Pascal, L., Memarzadeh, M., Boettiger, C., Lloyd, H., and Chadès, I. (2020). A shiny r app to solve the
513 problem of when to stop managing or surveying species under imperfect detection. *Methods in Ecology*
514 *and Evolution*, 11(12):1707–1715.
- 515 Pedersen, T. L., Nijs, V., Schaffner, T., and Nantz, E. (2021). *shinyFiles: A Server-Side File System Viewer*
516 *for Shiny*. R package version 0.9.1.
- 517 Perrier, V., Meyer, F., and Granjon, D. (2022). *shinyWidgets: Custom Inputs Widgets for Shiny*. R package
518 version 0.6.3.

- 519 R Core Team (2022). *R: A Language and Environment for Statistical Computing*. R Foundation for
520 Statistical Computing, Vienna, Austria.
- 521 Richardson, S. M., Lintott, P. R., Hosken, D. J., Economou, T., and Mathews, F. (2021). Peaks in bat
522 activity at turbines and the implications for mitigating the impact of wind energy developments on bats.
523 *Scientific Reports*, 11(1):3636.
- 524 Rogers, T. L., Ciaglia, M. B., Klinck, H., and Southwell, C. (2013). Density can be misleading for
525 low-density species: Benefits of passive acoustic monitoring. *PLOS ONE*, 8(1):1–11.
- 526 Ross, J. C. and Allen, P. E. (2014). Random forest for improved analysis efficiency in passive acoustic
527 monitoring. *Ecological Informatics*, 21:34–39. Ecological Acoustics.
- 528 Salamon, J., Bello, J. P., Farnsworth, A., and Kelling, S. (2017). Fusing shallow and deep learning for
529 bioacoustic bird species classification. In *2017 IEEE International Conference on Acoustics, Speech
530 and Signal Processing (ICASSP)*, pages 141–145.
- 531 Sarma, K. V., Raman, A. G., Dhinagar, N. J., Priester, A. M., Harmon, S., Sanford, T., Mehralivand, S.,
532 Turkbey, B., Marks, L. S., Raman, S. S., Speier, W., and Arnold, C. W. (2021). Harnessing clinical
533 annotations to improve deep learning performance in prostate segmentation. *PLOS ONE*, 16(6):1–15.
- 534 Schloerke, B. (2020). *reactlog: Reactivity Visualizer for 'shiny'*. R package version 1.1.0.
- 535 Silva, C. A., Hudak, A. T., Vierling, L. A., Valbuena, R., Cardil, A., Mohan, M., de Almeida, D. R. A.,
536 Broadbent, E. N., Almeyda Zambrano, A. M., Wilkinson, B., Sharma, A., Drake, J. B., Medley, P. B.,
537 Vogel, J. G., Prata, G. A., Atkins, J. W., Hamamura, C., Johnson, D. J., and Klauberg, C. (2022).
538 treetop: A shiny-based application and r package for extracting forest information from lidar data for
539 ecologists and conservationists. *Methods in Ecology and Evolution*, 13(6):1164–1176.
- 540 Silva, P. (2021). *imola: CSS Layouts (Grid and Flexbox) Implementation for R/Shiny*. R package version
541 0.3.2.
- 542 Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image
543 recognition.
- 544 Srivastava, A. K., Safaei, N., Khaki, S., Lopez, G., Zeng, W., Ewert, F., Gaiser, T., and Rahimi, J.
545 (2022). Winter wheat yield prediction using convolutional neural networks from environmental and
546 phenological data. *Scientific Reports*, 12(1):3215.
- 547 Stowell, D. (2022). Computational bioacoustics with deep learning: a review and roadmap. *PeerJ*,
548 10:e13152.
- 549 Stowell, D., Wood, M. D., Pamuła, H., Stylianou, Y., and Glotin, H. (2019). Automatic acoustic detection
550 of birds through deep learning: The first bird audio detection challenge. *Methods in Ecology and
551 Evolution*, 10(3):368–380.
- 552 Sueur, J., Aubin, T., and Simonis, C. (2008). Seewave: a free modular tool for sound analysis and
553 synthesis. *Bioacoustics*, 18:213–226.
- 554 Sugai, L. S. M., Silva, T. S. F., Ribeiro, José Wagner, J., and Llusia, D. (2018). Terrestrial Passive
555 Acoustic Monitoring: Review and Perspectives. *BioScience*, 69(1):15–25.
- 556 Thomas, M., Martin, B., Kowarski, K., Gaudet, B., and Matwin, S. (2019). Marine mammal species
557 classification using convolutional neural networks and a novel acoustic representation.
- 558 Warren, V. E., Širović, A., McPherson, C., Goetz, K. T., Radford, C. A., and Constantine, R. (2021).
559 Passive acoustic monitoring reveals spatio-temporal distributions of antarctic and pygmy blue whales
560 around central new zealand. *Frontiers in Marine Science*, 7.
- 561 Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York.
- 562 Wickham, H. (2019). *stringr: Simple, Consistent Wrappers for Common String Operations*. R package
563 version 1.4.0.
- 564 Wickham, H., François, R., Henry, L., and Müller, K. (2021). *dplyr: A Grammar of Data Manipulation*.
565 R package version 1.0.7.
- 566 Wszola, L. S., Simonsen, V. L., Stuber, E. F., Gillespie, C. R., Messinger, L. N., Decker, K. L., Lusk,
567 J. J., Jorgensen, C. F., Bishop, A. A., and Fontaine, J. J. (2017). Translating statistical species-habitat
568 models to interactive decision support tools. *PLoS ONE*, 12.
- 569 Xie, Y., Cheng, J., and Tan, X. (2021). *DT: A Wrapper of the JavaScript Library 'DataTables'*. R package
570 version 0.20.
- 571 Yin, M. S., Haddawy, P., Nirandmongkol, B., Kongthaworn, T., Chaisumritchoke, C., Supratak, A., Sa-
572 ngamuang, C., and Sriwichai, P. (2021). A lightweight deep learning approach to mosquito classification
573 from wingbeat sounds. In *Proceedings of the Conference on Information Technology for Social Good*,

- 574 GoodIT '21, page 37–42, New York, NY, USA. Association for Computing Machinery.
- 575 Zhong, M., LeBien, J., Campos-Cerqueira, M., Dodhia, R., Lavista Ferres, J., Velez, J. P., and Aide, T. M.
576 (2020). Multispecies bioacoustic classification using transfer learning of deep convolutional neural
577 networks with pseudo-labeling. *Applied Acoustics*, 166:107375.
- 578 Zsebők, S., Nagy-Egri, M. F., Barnaföldi, G. G., Laczi, M., Nagy, G., Éva Vaskuti, and Garamszegi,
579 L. Z. (2019). Automatic bird song and syllable segmentation with an open-source deep-learning object
580 detection method – a case study in the collared flycatcher. *Ornis Hungarica*, 27(2):59–66.

581 A DETAILED LABELLER INSTRUCTIONS

582 A more detailed version of the general workflow defined above in General labelling workflow carried out
583 by labellers is below. Items without numbers are optional steps.

- 584 ① Go to the deployed app at the given URL. Login via the Auth0 page.
- 585 ② Click the user icon in the top right, then start labelling. The first audio file and corresponding
586 spectrogram should load
- 587 ③ Navigate to the next file with little or no annotations. This can be inspected with the dropdown
588 menu, displaying the number of annotations in brackets.
- 589 ④ Tune parameters in the sidebar to the desired configuration. In particular, the spectrogram parameters
590 such as colour palette and contrast may be adjusted until the user is comfortable with the visual
591 distinction of the sounds present.
- 592 ⑤ If the user can already visually detect sound events (e.g. bird vocalisations or common forms of
593 noise), these can be annotated as described in steps ③-⑥.
- 594 ③ Play the audio until the user comes across a sound of interest. Once identified, place the mouse
595 at the top left corner of the vocalisation you identified and drag to the bottom right of the object,
596 drawing a moderately tight box around it.
- 597 ④ The audio player now updates to have a **filtered audio file**, reconstructed using only the times
598 and frequencies within the box drawn. This should assist the user in identifying the sound by
599 removing much of the noise present from other frequencies, and cropping the audio to the small
600 clip of interest. To return to the raw audio file, click anywhere on the plot.
- 601 ⑤ Once they have identified sound, if it is in the class list. If it is present, click the class button;
602 otherwise add it using the text box below and add it to the list.
- 603 ⑥ If any additional information can be extracted from the audio, such as the call type of the species or
604 miscellaneous notes, or the labeller's confidence of a particular annotation is less than certain, they
605 can be included using the text fields on the bottom right of the app's main page.
- 606 ⑥ When the desired class is selected from the class list, redraw a tight box around the vocalisation and
607 click **Save Selection** directly below the plot. This will save the annotation and draw a bounding
608 box with the chosen label onto the plot.
- 609 ⑦ If the user wishes to change any of the bounding boxes, they can be deleted using the **delete** button
610 in the **label buttons** section and and redrawn using the above steps. Alternatively they can be edited
611 using the label edit table which can be enabled in the **Other Settings** tab. This is described more in
612 the Extension to other audio labelling projects section.
- 613 ⑦ **Continue annotating** the file by repeating steps ③ to ⑥ until no more unlabelled sounds of
614 interest remain.
- 615 ⑧ **Proceed to the next file** using the file navigation menu or adjacent navigation buttons and go back
616 to step ③.
- 617 ⑨ Once the user finishes their session, click the user icon again, followed by **End Labelling** to finish.

618 **B BIRD CONSERVATION STATUS****Table 4. Bird species identified with conservation status and distribution across Ireland**

Common Name	Scientific Name	Bird Family	Conservation Status	Ireland Distribution
Blackbird	<i>Turdus merula</i>	Thrushes	Green	Widespread and common resident
Blue Tit	<i>Cyanistes caeruleus</i>	Tits	Green	Widespread and common resident
Buzzard	<i>Buteo buteo</i>	Raptors	Green	Widespread and common resident
Chaffinch	<i>Fringilla coelebs</i>	Finches	Green	Widespread and common resident
Coal Tit	<i>Parus ater</i>	Tits	Green	Widespread and common resident
Curlew	<i>Numenius arquata</i>	Waders	Red	Declining resident population & winter visitor to wetlands
Dunnock	<i>Prunella modularis</i>	Dunnocks	Green	Widespread and common resident
Goldcrest	<i>Regulus regulus</i>	Kinglets	Green	Common resident (coniferous forests)
Goldfinch	<i>Carduelis carduelis</i>	Finches	Green	Widespread and common resident
Grasshopper Warbler	<i>Locustella naevia</i>	Warblers	Amber	Widespread summer visitor
Great Tit	<i>Parus major</i>	Tits	Green	Widespread and common resident
Hooded Crow	<i>Corvus cornix</i>	Crows	Green	Widespread and common resident
House Sparrow	<i>Passer domesticus</i>	Sparrows	Amber	Widespread and common resident
Jackdaw	<i>Corvus monedula</i>	Crows	Green	Widespread and common resident
Linnet	<i>Carduelis cannabina</i>	Finches	Amber	Widespread and common resident
Magpie	<i>Pica pica</i>	Crows	Green	Widespread and common resident
Mallard	<i>Anas platyrhynchos</i>	Ducks	Green	Common resident (wetlands)
Meadow Pipit	<i>Anthus pratensis</i>	Pipits	Red	Common resident (rough pastures and uplands)
Oystercatcher	<i>Haematopus ostralegus</i>	Waders	Amber	Resident & winter visitor (coast and inland lakes)
Pheasant	<i>Phasianus colchicus</i>	Game Birds	Green	Introduced- Widespread and common resident
Pied Wagtail	<i>Motacilla alba yarrellii</i>	Wagtails	Green	Widespread and common resident
Robin	<i>Erithacus rubecula</i>	Chats	Green	Widespread and common resident
Rook	<i>Corvus frugilegus</i>	Crows	Green	Widespread (absent from expansive uplands)
Sedge Warbler	<i>Acrocephalus schoenobaenus</i>	Warblers	Green	Widespread and common summer visitor
Skylark	<i>Alauda arvensis</i>	Skylarks	Amber	Common resident (uplands and areas of farmland)
Song Thrush	<i>Turdus philomelos</i>	Thrushes	Green	Widespread and common resident
Starling	<i>Sturnus vulgaris</i>	Starling	Amber	Widespread and common resident
Stonechat	<i>Saxicola rubicola</i>	Chats	Green	Widespread resident (scrubland, mainly near the coast)
Swallow	<i>Hirundo rustica</i>	Swallows & Martins	Amber	Common summer visitor
Teal	<i>Anas crecca</i>	Ducks	Amber	Small numbers throughout Ireland
Woodpigeon	<i>Columba palumbus</i>	Pigeons & Doves	Green	Widespread and common resident
Wren	<i>Troglodytes troglodytes</i>	Wrens	Green	Widespread and common resident
Yellowhammer	<i>Emberiza citrinella</i>	Buntings	Red	Declining resident mainly in the east and south. Strongly tied to cereal cultivation.