

---

# **opensoundscape**

***Release 0.9.0***

**Apr 28, 2023**



---

## Contents

---

<b>1</b>	<b>Mac and Linux</b>	<b>3</b>
1.1	Installation via Anaconda . . . . .	3
1.2	Installation via <code>venv</code> . . . . .	3
<b>2</b>	<b>Windows</b>	<b>5</b>
2.1	Get Ubuntu shell . . . . .	5
2.2	Download Anaconda . . . . .	6
2.3	Install OpenSoundscape in virtual environment . . . . .	6
<b>3</b>	<b>Contributors</b>	<b>9</b>
3.1	Poetry installation . . . . .	9
3.2	Contribution workflow . . . . .	10
<b>4</b>	<b>Jupyter</b>	<b>11</b>
4.1	Use virtual environment . . . . .	11
4.2	Create independent kernel . . . . .	11
<b>5</b>	<b>Audio and spectrograms</b>	<b>13</b>
5.1	Quick start . . . . .	14
5.2	Audio loading . . . . .	14
5.3	Audio methods . . . . .	16
5.4	Spectrogram creation . . . . .	23
5.5	Spectrogram methods . . . . .	25
<b>6</b>	<b>Manipulating audio annotations</b>	<b>31</b>
6.1	Download example files . . . . .	31
6.2	View a subset of annotations . . . . .	34
6.3	Saving annotations to Raven-compatible file . . . . .	34
6.4	1. Split Audio object, then split annotations to match . . . . .	35
6.5	Split annotations using splitting parameters . . . . .	36
6.6	Match up audio files and Raven annotations . . . . .	37
6.7	Create label dataframes . . . . .	39
6.8	Load Raven annotations . . . . .	39
6.9	Create label dataframes . . . . .	39
6.10	Sanity check: look at spectrograms of clips labeled 0 and 1 . . . . .	40
<b>7</b>	<b>Prediction with pre-trained CNNs</b>	<b>43</b>

7.1	Load required packages . . . . .	43
7.2	generate predictions with the model . . . . .	44
7.3	Overlapping prediction clips . . . . .	44
7.4	Inspect samples generated during prediction . . . . .	45
7.5	Options for prediction . . . . .	46
7.6	Using models from older OpenSoundscape versions . . . . .	47
<b>8</b>	<b>Beginner friendly training and prediction with CNNs</b>	<b>53</b>
8.1	Prepare audio data . . . . .	54
8.2	Create and train a model . . . . .	55
8.3	Visualize gradient activation maps . . . . .	61
8.4	Prediction . . . . .	63
8.5	Multi-class models . . . . .	65
8.6	Save and load models . . . . .	67
8.7	Predict using saved (or pre-trained) model . . . . .	68
8.8	Continue training from saved model . . . . .	69
8.9	Next steps . . . . .	69
<b>9</b>	<b>Preprocessing audio samples with OpenSoundscape</b>	<b>71</b>
9.1	Modifying the preprocessor of the CNN class . . . . .	72
9.2	Download labeled audio files . . . . .	73
9.3	Load dataframe of files and labels . . . . .	73
9.4	Initialize preprocessor . . . . .	74
9.5	Generate a sample from a Dataset . . . . .	75
9.6	Subset samples from a Dataset . . . . .	77
9.7	About Pipelines . . . . .	79
9.8	About actions . . . . .	80
9.9	View default parameters for an Action . . . . .	80
9.10	Modify Action parameters . . . . .	81
9.11	Bypass actions . . . . .	81
9.12	Example: return Spectrogram instead of Tensor . . . . .	83
9.13	analyse the output at steps of interest . . . . .	85
9.14	Modify the sample rate . . . . .	87
9.15	Modify spectrogram window length and overlap . . . . .	88
9.16	Bandpass spectrograms . . . . .	89
9.17	Change the output shape . . . . .	90
9.18	Turn all augmentation on or off . . . . .	91
9.19	Modify augmentation parameters . . . . .	92
9.20	remove an action by its name . . . . .	94
9.21	add an action at a specific position . . . . .	94
9.22	Overlay augmentation . . . . .	95
<b>10</b>	<b>Advanced CNN training</b>	<b>107</b>
10.1	Prepare audio data . . . . .	108
10.2	Creating a model . . . . .	109
10.3	Model training parameters . . . . .	109
10.4	Selecting CNN architectures . . . . .	111
10.5	Multi-target training with ResampleLoss . . . . .	113
10.6	Training and predicting with custom preprocessors . . . . .	113
<b>11</b>	<b>RIBBIT Pulse Rate model demonstration</b>	<b>115</b>
11.1	Import packages . . . . .	115
11.2	Download example audio . . . . .	116
11.3	Select model parameters . . . . .	117
11.4	Search for pulsing vocalizations with <code>ribbit()</code> . . . . .	118



11.5	Analyzing a set of files . . . . .	120
11.6	Run RIBBIT on multiple species simultaneously . . . . .	121
11.7	Detail view of RIBBIT method . . . . .	124
11.8	Time to experiment for yourself . . . . .	126
<b>12</b>	<b>Spectrogram</b>	<b>129</b>
<b>13</b>	<b>CNN</b>	<b>135</b>
<b>14</b>	<b>Annotations</b>	<b>143</b>
<b>15</b>	<b>ARU</b>	<b>149</b>
<b>16</b>	<b>Machine Learning utils</b>	<b>151</b>
<b>17</b>	<b>CNN Architectures</b>	<b>153</b>
<b>18</b>	<b>Logging with WandB (Weights and Biases)</b>	<b>159</b>
<b>19</b>	<b>Data Selection</b>	<b>161</b>
<b>20</b>	<b>Datasets</b>	<b>163</b>
<b>21</b>	<b>CAM (Class Activation Maps)</b>	<b>165</b>
<b>22</b>	<b>Loss</b>	<b>167</b>
<b>23</b>	<b>Safe Dataset</b>	<b>169</b>
<b>24</b>	<b>Sample</b>	<b>171</b>
<b>25</b>	<b>Sampling</b>	<b>173</b>
<b>26</b>	<b>Metrics</b>	<b>175</b>
<b>27</b>	<b>Image Augmentation</b>	<b>177</b>
<b>28</b>	<b>Actions</b>	<b>179</b>
<b>29</b>	<b>Preprocessors</b>	<b>185</b>
<b>30</b>	<b>preprocessors.utils</b>	<b>187</b>
<b>31</b>	<b>Tensor Augment</b>	<b>189</b>
<b>32</b>	<b>RIBBIT</b>	<b>191</b>
<b>33</b>	<b>Signal Processing</b>	<b>195</b>
<b>34</b>	<b>Taxa</b>	<b>201</b>
<b>35</b>	<b>Localization</b>	<b>203</b>
<b>36</b>	<b>utils</b>	<b>211</b>
<b>37</b>	<b>Index</b>	<b>215</b>
<b>38</b>	<b>Modules</b>	<b>217</b>

<b>Python Module Index</b>	<b>219</b>
<b>Index</b>	<b>221</b>

OpenSoundscape is free and open source software for the analysis of bioacoustic recordings ([GitHub](#)). Its main goals are to allow users to train their own custom species classification models using a variety of frameworks (including convolutional neural networks) and to use trained models to predict whether species are present in field recordings. OpSo can be installed and run on a single computer or in a cluster or cloud environment.

OpenSoundscape is developed and maintained by the [Kitzes Lab](#) at the University of Pittsburgh.

The Installation section below provides guidance on installing OpSo. The Tutorials pages below are written as Jupyter Notebooks that can also be downloaded from the [project repository](#) on GitHub.



# CHAPTER 1

---

## Mac and Linux

---

OpenSoundscape can be installed on Mac and Linux machines with Python 3.8, or 3.9 using the pip command `pip install opensoundscape==0.9.0`. We recommend installing OpenSoundscape in a virtual environment to prevent dependency conflicts.

Below are instructions for installation with two package managers:

- `conda`: Python and package management through Anaconda, a package manager popular among scientific programmers
- `venv`: Python's included virtual environment manager, `venv`

Feel free to use another virtual environment manager (e.g. `virtualenvwrapper`) if desired.

### 1.1 Installation via Anaconda

- Install Anaconda if you don't already have it.
  - Download the installer [here](#), or
  - follow the [installation instructions](#) for your operating system.
- Create a Python (3.8, or 3.9) conda environment for opensoundscape: `conda create --name opensoundscape pip python=3.9`
- Activate the environment: `conda activate opensoundscape`
- Install opensoundscape using pip: `pip install opensoundscape==0.9.0`
- Deactivate the environment when you're done using it: `conda deactivate`

### 1.2 Installation via venv

Download Python 3.8 or 3.9 from [this website](#).

Run the following commands in your bash terminal:

- Check that you have installed Python 3.8 or 3.9.\_: `python3 --version`
- Change directories to where you wish to store the environment: `cd [path for environments folder]`
  - Tip: You can use this folder to store virtual environments for other projects as well, so put it somewhere that makes sense for you, e.g. in your home directory.
- Make a directory for virtual environments and cd into it: `mkdir .venv && cd .venv`
- Create an environment called opensoundscape in the directory: `python3 -m venv opensoundscape`
- Activate/use the environment: `source opensoundscape/bin/activate`
- Install OpenSoundscape in the environment: `pip install opensoundscape==0.9.0`
- Once you are done with OpenSoundscape, deactivate the environment: `deactivate`
- To use the environment again, you will have to refer to absolute path of the virtual environments folder. For instance, if I were on a Mac and created `.venv` inside a directory `/Users/MyFiles/Code` I would activate the virtual environment using: `source /Users/MyFiles/Code/.venv/opensoundscape/bin/activate`

For some of our functions, you will need a version of `ffmpeg` `>= 0.4.1`. On Mac machines, `ffmpeg` can be installed via `brew`.

We recommend that Windows users install and use OpenSoundscape using Windows Subsystem for Linux, because some of the machine learning and audio processing packages required by OpenSoundscape do not install easily on Windows computers. Below we describe the typical installation method. This gives you access to a Linux operating system (we recommend Ubuntu 20.04) in which to use Python and install and use OpenSoundscape. Using Ubuntu 20.04 is as simple as opening a program on your computer.

### 2.1 Get Ubuntu shell

If you don't already use Windows Subsystem for Linux (WSL), activate it using the following:

- Search for the “Powershell” program on your computer
- Right click on “Powershell,” then click “Run as administrator” and in the pop-up, allow it to run as administrator
- Install WSL2 (more information: <https://docs.microsoft.com/en-us/windows/wsl/install-win10>):

```
wsl --install
```

- Restart your computer

Once you have WSL, follow these steps to get an Ubuntu shell on your computer:

- Open Windows Store, search for “Ubuntu” and click “Ubuntu 20.04 LTS”
- Click “Get”, wait for the program to download, then click “Launch”
- An Ubuntu shell will open. Wait for Ubuntu to install.
- Set username and password to something you will remember
- Run `sudo apt update` and type in the password you just set

## 2.2 Download Anaconda

We recommend installing OpenSoundscape in a package manager. We find that the easiest package manager for new users is “Anaconda,” a program which includes Python and tools for managing Python packages. Below are instructions for downloading Anaconda in the Ubuntu environment.

- Open [this page](#) and scroll down to the “Anaconda Installers” section. Under the Linux section, right click on the link “64-Bit (x86) Installer” and click “Copy link”
- Download the installer:
  - Open the Ubuntu terminal
  - Type in `wget` then paste the link you copied, e.g.: (the filename of your file may differ)

```
wget https://repo.anaconda.com/archive/Anaconda3-2020.07-Linux-x86_64.sh
```

- Execute the downloaded installer, e.g.: (the filename of your file may differ)
- ```
bash Anaconda3-2020.07-Linux-x86_64.sh
```
- Press ENTER, read the installation requirements, press Q, then type “yes” and press enter to install
  - Wait for it to install
  - If your download hangs, press CTRL+C, `rm -rf ~/anaconda3` and try again
  - Type “yes” to initialize conda
    - If you skipped this step, initialize your conda installation: `run source ~/anaconda3/bin/activate` and then after that command has run, `conda init`.
  - Remove the downloaded file after installation, e.g. `rm Anaconda3-2020.07-Linux-x86_64.sh`
  - Close and reopen terminal window to have access to the initialized Anaconda distribution

You can now manage packages with `conda`.

## 2.3 Install OpenSoundscape in virtual environment

- Create a Python (3.8, or 3.9) conda environment for opensoundscape: `conda create --name opensoundscape pip python=3.9`
- Activate the environment: `conda activate opensoundscape`
- Install opensoundscape using pip: `pip install opensoundscape==0.9.0`

If you see an error that says “No matching distribution found...”, your best bet is to use these commands to download then install the package:

```
cd
git clone https://github.com/kitzeslab/opensoundscape.git
cd opensoundscape/
pip install .
```

If you run into this error and you are on a Windows 10 machine:



```
(opensoundscape_environment) username@computername:~$ pip install opensoundscape==0.9.
↪0
WARNING: Retrying (Retry(total=4, connect=None, read=None, redirect=None,
↪status=None)) after connection broken by 'NewConnectionError('<pip._vendor.urllib3.
↪connection.HTTPSConnection object at 0x7f7603c5da90>: Failed to establish a new
↪connection: [Errno -2] Name or service not known')': /simple/opensoundscape/
```

You may be able to solve it by going to System Settings, searching for “Proxy Settings,” and beneath “Automatic proxy setup,” turning “Automatically detect settings” OFF. Restart your terminal for changes to take effect. Then activate the environment and install OpenSoundscape using pip.



Contributors and advanced users can use this workflow to install OpenSoundscape using Poetry. Poetry installation allows direct use of the most recent version of the code. This workflow allows advanced users to use the newest features in OpenSoundscape, and allows developers/contributors to build and test their contributions.

### 3.1 Poetry installation

- Install `poetry`
- Create a new virtual environment for the OpenSoundscape installation. If you are using Anaconda, you can create a new environment with `conda create -n opso-dev python=3.9` where `opso-dev` is the name of the new virtual environment. Use `conda activate opso-dev` to enter the environment to work on OpenSoundscape and `conda deactivate opso-dev` to return to your base Python installation. If you are not using Anaconda, other packages such as `virtualenv` should work as well. Ensure that the Python version is compatible with the current version of OpenSoundscape.
- **Internal Contributors:** Clone this github repository to your machine: `git clone https://github.com/kitzeslab/opensoundscape.git`
- **External Contributors:** Fork this github repository and clone the fork to your machine
- Ensure you are in the top-level directory of the clone
- Switch to the development branch of OpenSoundscape: `git checkout develop`
- Install OpenSoundscape using `poetry install`. This will install OpenSoundscape and its dependencies into the `opso-dev` virtual environment. By default it will install OpenSoundscape in develop mode, so that updated code in the repository can be imported without reinstallation.
  - If you are on a Mac and `poetry install` fails to install `numba`, contact one of the developers for help troubleshooting your issues.
- Install the `ffmpeg` dependency. On a Mac, `ffmpeg` can be installed using Homebrew.
- Run the test suite to ensure that everything installed properly. From the top-level directory, run the command `pytest`.

## 3.2 Contribution workflow

### 3.2.1 Contributing to code

Make contributions by editing the code in your repo. Create branches for features by starting with the `develop` branch and then running `git checkout -b feature_branch_name`. Once work is complete, push the new branch to remote using `git push -u origin feature_branch_name`. To merge a feature branch into the development branch, use the GitHub web interface to create a merge or a pull request. Before opening a PR, do the following to ensure the code is consistent with the rest of the package:

- Run the test suite using `pytest`
- Format the code with `black` style (from the top level of the repo): `black .`

### 3.2.2 Contributing to documentation

Build the documentation using `sphinx-build docs docs/_build`

To use OpenSoundscape in JupyterLab or in a Jupyter Notebook, you may either start Jupyter from within your OpenSoundscape virtual environment and use the “Python 3” kernel in your notebooks, or create a separate “OpenSoundscape” kernel using the instructions below

The following steps assume you have already used your operating system-specific installation instructions to create a virtual environment containing OpenSoundscape and its dependencies.

### 4.1 Use virtual environment

- Activate your virtual environment
- Start JupyterLab or Jupyter Notebook from inside the conda environment, e.g.: `jupyter lab`
- Copy and paste the JupyterLab link into your web browser

With this method, the default “Python 3” kernel will be able to import `opensoundscape` modules.

### 4.2 Create independent kernel

Use the following steps to create a kernel that appears in any notebook you open, not just notebooks opened from your virtual environment.

- Activate your virtual environment to have access to the `ipykernel` package
- Create `ipython` kernel with the following command, replacing `ENV_NAME` with the name of your OpenSoundscape virtual environment.

```
python -m ipykernel install --user --name=ENV_NAME --display-name=OpenSoundscape
```

- Now when you make a new notebook on JupyterLab, or change kernels on an existing notebook, you can choose to use the “OpenSoundscape” Python kernel

Contributors: if you include Jupyter's `autoreload`, any changes you make to the source code installed via poetry will be reflected whenever you run the `%autoreload` line magic in a cell:

```
%load_ext autoreload
%autoreload
```

---

## Audio and spectrograms

---

This tutorial demonstrates how to use OpenSoundscape to open and modify audio files and spectrograms.

Audio files can be loaded into OpenSoundscape and modified using its `Audio` class. The class gives access to modifications such as trimming short clips from longer recordings, splitting a long clip into multiple segments, bandpassing recordings, and extending the length of recordings by looping them. Spectrograms can be created from `Audio` objects using the `Spectrogram` class. This class also allows useful features like measuring the amplitude signal of a recording, trimming a spectrogram in time and frequency, and converting the spectrogram to a saveable image.

To download the tutorial as a Jupyter Notebook, click the “Edit on GitHub” button at the top right of the tutorial. Using it requires that you install OpenSoundscape and follow the instructions for using it in Jupyter.

As an example, we will download a file from the Kitzes Lab box location using the code below, and use it throughout the tutorial. To use your own file for the following examples, change the string assigned to `audio_filename` to any audio file on your computer.

```
[1]: import subprocess
subprocess.run(['curl',
               'https://drive.google.com/uc?export=download&id=1hBUyhgPf-
↳ vxet8BITSMZ0XsGiU-a7SG',
               '-L', '-o', 'lmin_audio.wav'])
```

| % Total | % Received | % Xferd | Average Speed | Time    | Time    | Time     | Current |
|---------|------------|---------|---------------|---------|---------|----------|---------|
|         |            |         | Dload Upload  | Total   | Spent   | Left     | Speed   |
| 0       | 0          | 0       | 0             | 0       | 0:00:01 | --:--:-- | 0       |
| 100     | 3750k      | 100     | 3750k         | 0       | 0       | 1603k    | 0       |
|         |            |         |               | 0:00:02 | 0:00:02 | --:--:-- | 3124k   |

```
[1]: CompletedProcess(args=['curl', 'https://drive.google.com/uc?export=download&
↳ id=1hBUyhgPf-vxet8BITSMZ0XsGiU-a7SG', '-L', '-o', 'lmin_audio.wav'], returncode=0)
```

```
[2]: audio_filename = './lmin_audio.wav'
```

## 5.1 Quick start

Import the `Audio` and `Spectrogram` classes from `OpenSoundscape`. (For more information about Python imports, review [this article](#).)

```
[3]: # import Audio and Spectrogram classes from OpenSoundscape
from opensoundscape import Audio, Spectrogram
```

These classes provide a variety of tools to load and manipulate audio and spectrograms. The code below demonstrates a basic pipeline:

- load an audio file
- generate a spectrogram with default parameters
- create a 224px X 224px-sized image of the spectrogram
- save the image to a file

```
[4]: from pathlib import Path

# Settings
image_shape = (224, 224) # (height, width) not (width, height)
image_save_path = Path('./saved_spectrogram.png')

# Load audio file as Audio object
audio = Audio.from_file(audio_filename)

# Create Spectrogram object from Audio object
spectrogram = Spectrogram.from_audio(audio)

# Convert Spectrogram object to Python Imaging Library (PIL) Image
image = spectrogram.to_image(shape=image_shape, invert=True)

# Save image to file
image.save(image_save_path)
```

The above function calls could even be condensed to a single line:

```
[5]: Spectrogram.from_audio(Audio.from_file(audio_filename)).to_image(shape=image_shape,
    ↪invert=True).save(image_save_path)
```

Clean up by deleting the spectrogram saved above.

```
[6]: image_save_path.unlink()
```

## 5.2 Audio loading

The `Audio` class in `OpenSoundscape` allows loading and manipulation of audio files.

### 5.2.1 Load .wav(s)

Load the example audio from file:



```
[34]: audio_object = Audio.from_file(audio_filename)
      audio_object
```

```
[34]: <Audio(samples=(1920000,), sample_rate=32000)>
```

The interactive playback widget is displayed when an Audio object is returned from a notebook cell. We can also create the widget with `.show_widget()`:

```
[38]: # create playback widget with normalized playback level
      audio_object.show_widget(normalize=True)
```

```
#do more things
print('continue...')
```

```
<IPython.lib.display.Audio object>
```

```
continue...
```

### 5.2.2 Load a segment of a file

We can directly load a section of a `.wav` file very quickly (even if the audio file is large) using the `offset` and `duration` parameters.

For example, let's load 1 second of audio from 2.0-3.0 seconds after the start of the file:

```
[39]: audio_segment = Audio.from_file(audio_filename, offset=2.0, duration=1.0)
      audio_segment.duration
```

```
[39]: 1.0
```

### 5.2.3 Audio properties

The properties of an Audio object include its `samples` (the actual audio data) and the `sample_rate` (the number of audio samples taken per second, required to understand the samples). After an audio file has been loaded, these properties can be accessed using the `samples` and `sample_rate` attributes, respectively.

```
[40]: print(f"How many samples does this audio object have? {len(audio_object.samples)}")
      print(f"What is the sampling rate? {audio_object.sample_rate}")
```

```
How many samples does this audio object have? 1920000
```

```
What is the sampling rate? 32000
```

### 5.2.4 Resample audio during load

By default, an audio object is loaded with the same sample rate as the source recording.

The `sample_rate` parameter of `Audio.from_file` allows you to re-sample the file during the creation of the object. This is useful when working with multiple files to ensure that all files have a consistent sampling rate.

Let's load the same audio file as above, but specify a sampling rate of 22050 Hz.

```
[41]: audio_object_resample = Audio.from_file(audio_filename, sample_rate=22050)
      audio_object_resample.sample_rate
```

```
[41]: 22050
```

## 5.2.5 Load audio from a specific real-world time from AudioMoth recordings

OpenSoundscape parses metadata of files recorded on AudioMoth recorders, and can use the metadata to extract pieces of audio corresponding to specific real-world times. (Note that AudioMoth internal clocks can drift an estimated 10-60 seconds per month).

```
[42]: from datetime import datetime; import pytz

start_time = pytz.timezone('UTC').localize(datetime(2020,4,4,10,25))
audio_length = 5 #seconds
path = '/path/to/audiomoth_recording.WAV' #an AudioMoth recording

#this line is commented because it will fail unless you specify a valid path in the_
↪line above
# Audio.from_file(path, start_timestamp=start_time,duration=audio_length)
```

For other options when loading audio objects, see the `Audio.from_file()` documentation.

## 5.3 Audio methods

The `Audio` class gives access to a variety of tools to change audio files, load them with special properties, or get information about them. Various examples are shown below.

For a description of the entire `Audio` object API, see the [API documentation](#).

### 5.3.1 NOTE: Out-of-place operations

Functions that modify `Audio` (and `Spectrogram`) objects are “out of place”, meaning that they return a new, modified instance of `Audio` instead of modifying the original instance. This means that running a line

```
audio_object.resample(22050) # WRONG!
```

will **not** change the sample rate of `audio_object`! If your goal was to overwrite `audio_object` with the new, resampled audio, you would instead write

```
audio_object = audio_object.resample(22050)
```

### 5.3.2 Save audio to file

Opensoundscape currently supports saving `Audio` objects to `.wav` formats **only**. It does **not** currently support saving metadata (tags) along with wav files - only the samples and sample rate will be preserved in the file.

```
[43]: audio_object.save('./my_audio.wav')
```

clean up: delete saved file

```
[44]: from pathlib import Path
Path('./my_audio.wav').unlink()
```

### 5.3.3 Measurements of Audio objects

Print the duration in seconds, rms (root mean square) value, and dBFS measurements:

```
[45]: print(f"duration: {audio_object.duration}")
      print(f"rms: {audio_object.rms:0.3f}")
      print(f"dBFS: {audio_object.dBFS:0.1f}")
```

```
duration: 60.0
rms: 0.013
dBFS: -34.5
```

### 5.3.4 Trim

The `.trim()` method extracts audio from a specified time period in seconds (relative to the start of the audio object).

```
[46]: trimmed = audio_object.trim(0,5)
      trimmed.duration
```

```
[46]: 5.0
```

### 5.3.5 Split Audio into clips

The `.split()` method divides audio into even-lengthed clips, optionally with overlap between adjacent clips (default is no overlap). See the function's documentation for options on how to handle the last clip.

The function returns a list containing Audio objects for each clip and a DataFrame giving the start and end times of each clip with respect to the original file.

#### split

```
[47]: #split into 5-second clips with no overlap between adjacent clips
      clips, clip_df = audio_object.split(clip_duration=5, clip_overlap=0, final_clip=None)

      #check the duration of the Audio object in the first returned element
      print(f"duration of first clip: {clips[0].duration}")

      print(f"head of clip_df")
      clip_df.head(3)
```

```
duration of first clip: 5.0
head of clip_df
```

```
[47]:   start_time  end_time
0         0.0        5.0
1         5.0       10.0
2        10.0       15.0
```

#### split with overlap

if we want overlap between consecutive clips

Note that a negative “overlap” value would leave *gaps* between consecutive clips.

```
[48]: _, clip_df = audio_object.split(clip_duration=5, clip_overlap=2.5, final_clip=None)
print(f"head of clip_df")
clip_df.head()
```

head of clip\_df

```
[48]:      start_time  end_time
0          0.0      5.0
1          2.5      7.5
2          5.0     10.0
3          7.5     12.5
4         10.0     15.0
```

## split and save

The `Audio.split_and_save()` method splits audio into clips and immediately saves them to files in a specified location. You provide it with a naming prefix, and it will add on a suffix indicating the start and end times of the clip (eg `_5.0-10.0s.wav`). It returns just a `DataFrame` with the paths and start/end times for each clip (it does not return `Audio` objects).

The splitting options are the same as `.split()`: `clip_duration`, `clip_overlap`, and `final_clip`

```
[49]: #split into 5-second clips with no overlap between adjacent clips
Path('./temp_audio').mkdir(exist_ok=True)
clip_df = audio_object.split_and_save(
    destination='./temp_audio',
    prefix='audio_clip_',
    clip_duration=5,
    clip_overlap=0,
    final_clip=None
)

print(f"head of clip_df")
clip_df.head()
```

head of clip\_df

```
[49]:      file      start_time  end_time
0  ./temp_audio/audio_clip__0.0s_5.0s.wav      0.0      5.0
1  ./temp_audio/audio_clip__5.0s_10.0s.wav      5.0     10.0
2  ./temp_audio/audio_clip__10.0s_15.0s.wav     10.0     15.0
3  ./temp_audio/audio_clip__15.0s_20.0s.wav     15.0     20.0
4  ./temp_audio/audio_clip__20.0s_25.0s.wav     20.0     25.0
```

The folder `temp_audio` should now contain 12 5-second clips created from the 60-second audio file.

clean up: delete temp folder of saved audio clips

```
[50]: from shutil import rmtree
rmtree('./temp_audio')
```

## split\_and\_save dry run

we can use the `dry_run=True` option to produce only the `clip_df` but not actually process the audio. this is useful as a quick test to see if the function is behaving as expected, before doing any (potentially slow) splitting on huge audio files.

Just for fun, we'll use an overlap of -5 in this example (5 second gap between each consecutive clip)

This function returns a DataFrame of clips, but does not actually process the audio files or write any new files.

```
[51]: clip_df = audio_object.split_and_save(
        destination='./temp_audio',
        prefix='audio_clip_',
        clip_duration=5,
        clip_overlap=-5,
        final_clip=None,
        dry_run=True,
    )
clip_df
```

|                                          | start_time | end_time |
|------------------------------------------|------------|----------|
| file                                     |            |          |
| ./temp_audio/audio_clip__0.0s_5.0s.wav   | 0.0        | 5.0      |
| ./temp_audio/audio_clip__10.0s_15.0s.wav | 10.0       | 15.0     |
| ./temp_audio/audio_clip__20.0s_25.0s.wav | 20.0       | 25.0     |
| ./temp_audio/audio_clip__30.0s_35.0s.wav | 30.0       | 35.0     |
| ./temp_audio/audio_clip__40.0s_45.0s.wav | 40.0       | 45.0     |
| ./temp_audio/audio_clip__50.0s_55.0s.wav | 50.0       | 55.0     |

### 5.3.6 Extend and loop

The `.extend()` method extends an audio file to a desired length by adding silence to the end.

The `.loop()` method extends an audio file to a desired length (or number of repetitions) by looping the audio.


`extend()` example: create an Audio object twice as long as the original, extending with silence (0 valued samples)

```
[55]: import matplotlib.pyplot as plt
%config InlineBackend.figure_format = 'retina'

# create an audio object twice as long, extending the end with silence (zero-values)
extended = trimmed.extend(trimmed.duration * 2)

print(f"duration of original clip: {trimmed.duration}")
print(f"duration of extended clip: {extended.duration}")
print(f"samples of extended clip:")
plt.plot(extended.samples)
plt.show()
```

```
duration of original clip: 5.0
duration of extended clip: 10.0
samples of extended clip:
```




Looping example: create an audio object 1.5x as long, extending the end by looping

```
[56]: looped = trimmed.loop(trimmed.duration * 1.5)
      print(looped.duration)
      plt.plot(looped.samples)
```

7.5

```
[56]: [<matplotlib.lines.Line2D at 0x2af395430>]
```




create an audio object that loops the original object 5 times and plot the samples

```
[57]: looped = trimmed.loop(n=5)
      print(looped.duration)
      plt.plot(looped.samples)
```

25.0

```
[57]: [<matplotlib.lines.Line2D at 0x297fc4df0>]
```



### 5.3.7 Resample

The `.resample()` method resamples the audio object to a new sampling rate (can be lower or higher than the original sampling rate)

```
[58]: resampled = trimmed.resample(sample_rate=48000)
      resampled.sample_rate
```

```
[58]: 48000
```

### 5.3.8 Generate a frequency spectrum


The `.spectrum()` method provides an easy way to compute a Fourier Transform on an audio object to measure its frequency composition.

```
[59]: # calculate the fft
      fft_spectrum, frequencies = trimmed.spectrum()

      #plot settings
      from matplotlib import pyplot as plt
      plt.rcParams['figure.figsize']=[15,5] #for big visuals
      %config InlineBackend.figure_format = 'retina'

      # plot
      plt.plot(frequencies,fft_spectrum)
      plt.ylabel('Fast Fourier Transform (V**2/Hz)')
      plt.xlabel('Frequency (Hz)')
```

```
[59]: Text(0.5, 0, 'Frequency (Hz)')
```



### 5.3.9 Bandpass

Bandpass the audio file to limit its frequency range to 1000 Hz to 5000 Hz. The bandpass operation uses a Butterworth filter with a user-provided order.


```
[60]: # apply a bandpass filter
bandpassed = trimmed.bandpass(low_f = 1000, high_f = 5000, order=9)

# calculate the bandpassed audio's spectrum
fft_spectrum, frequencies = bandpassed.spectrum()

# plot
print('spectrum after bandpassing the audio:')
plt.plot(frequencies,fft_spectrum)
plt.ylabel('Fast Fourier Transform (V**2/Hz)')
plt.xlabel('Frequency (Hz)')
```

spectrum after bandpassing the audio:

```
[60]: Text(0.5, 0, 'Frequency (Hz)')
```





## 5.4 Spectrogram creation

### 5.4.1 Load spectrogram

A Spectrogram object can be created from an audio object using the `from_audio()` method.

```
[61]: audio_object = Audio.from_file(audio_filename)
      spectrogram_object = Spectrogram.from_audio(audio_object)
```

### 5.4.2 Spectrogram properties

To check the time and frequency axes of a spectrogram, you can look at its `times` and `frequencies` attributes. The `times` attribute is the list of the spectrogram windows' centers' times in seconds relative to the beginning of the audio. The `frequencies` attribute is the list of frequencies represented by each row of the spectrogram. These are not the actual values of the spectrogram — just the values of the axes.


```
[62]: spec = Spectrogram.from_audio(Audio.from_file(audio_filename))
      print(f'the first few times: {spec.times[0:5]}')
      print(f'the first few frequencies: {spec.frequencies[0:5]}')

the first few times: [0.008 0.016 0.024 0.032 0.04 ]
the first few frequencies: [ 0.   62.5 125. 187.5 250. ]
```

### 5.4.3 Plot spectrogram

A Spectrogram object can be visualized using its `plot()` method.

```
[63]: audio_object = Audio.from_file(audio_filename)
      spectrogram_object = Spectrogram.from_audio(audio_object)
      spectrogram_object.plot()
```



### 5.4.4 Spectrogram parameters

Spectrograms are created using “windows”. A window is a subset of consecutive samples of the original audio that is analyzed to create one pixel in the horizontal direction (one “column”) on the resulting spectrogram. The appearance of a spectrogram depends on two parameters that control the size and spacing of these windows:

### Samples per window, `window_samples`

This parameter is the length (in audio samples) of each spectrogram window. Choosing the value for `window_samples` represents a trade-off between frequency resolution and time resolution:

- Larger value for `window_samples` → higher frequency resolution (more rows in a single spectrogram column)
- Smaller value for `window_samples` → higher time resolution (more columns in the spectrogram per second)

### Overlap of consecutive windows, `overlap_samples`

`overlap_samples`: this is the number of audio samples that will be re-used (overlap) between two consecutive Spectrogram windows. It must be less than `window_samples` and greater than or equal to zero. Zero means no overlap between windows, while a value of `window_samples/2` would give 50% overlap between consecutive windows. Using higher overlap percentages can sometimes yield better time resolution in a spectrogram, but will take more computational time to generate.

### Spectrogram parameter tradeoffs

When there is zero overlap between windows, the number of columns per second is equal to the size in Hz of each spectrogram row. Consider the relationship between time resolution (columns in the spectrogram per second) and frequency resolution (rows in a given frequency range) in the following example:

- Let `sample_rate=48000`, `window_samples=480`, and `overlap_samples=0`
- Each window (“spectrogram column”) represents  $480/48000 = 1/100 = 0.01$  seconds of audio
- There will be  $1/(\text{length of window in seconds}) = 1/0.01 = 100$  columns in the spectrogram per second.
- Each pixel will span 100 Hz in the frequency dimension, i.e., the lowest pixel spans 0-100 Hz, the next lowest 100-200 Hz, then 200-300 Hz, etc.

If `window_samples=4800`, then the spectrogram would have better time resolution (each window represents only  $4800/48000 = 0.001$ s of audio) but worse frequency resolution (each row of the spectrogram would represent 1000 Hz in the frequency range).

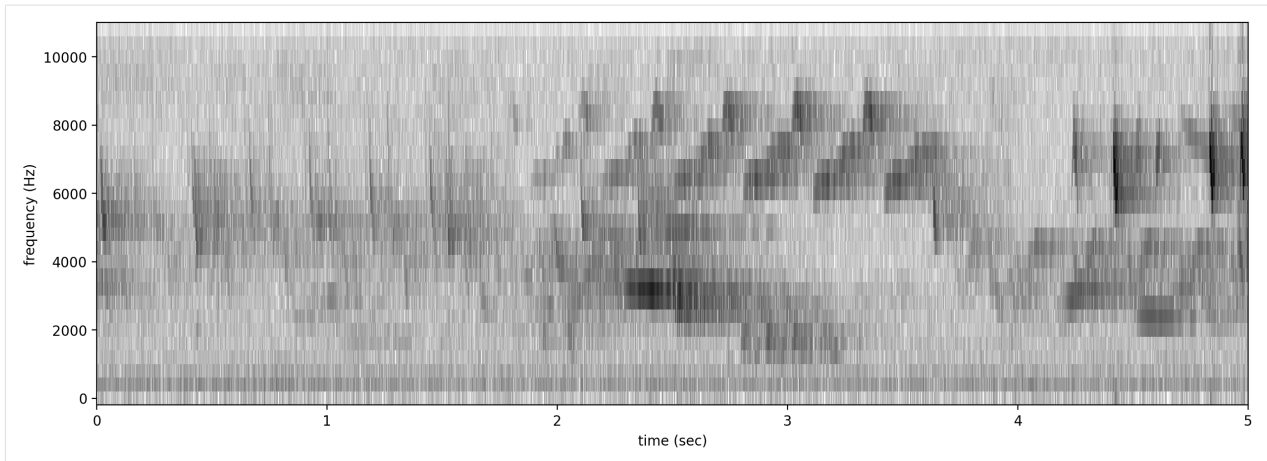
As an example, let’s create two spectrograms, one with high time resolution and another with high frequency resolution.

```
[64]: # Load audio
audio = Audio.from_file(audio_filename, sample_rate=22000).trim(0,5)
```

#### Create a spectrogram with high time resolution

Using `window_samples=55` and `overlap_samples=0` gives  $55/22000 = 0.0025$  seconds of audio per window, or  $1/0.0025 = 400$  windows per second. Each spectrogram pixel spans 400 Hz.

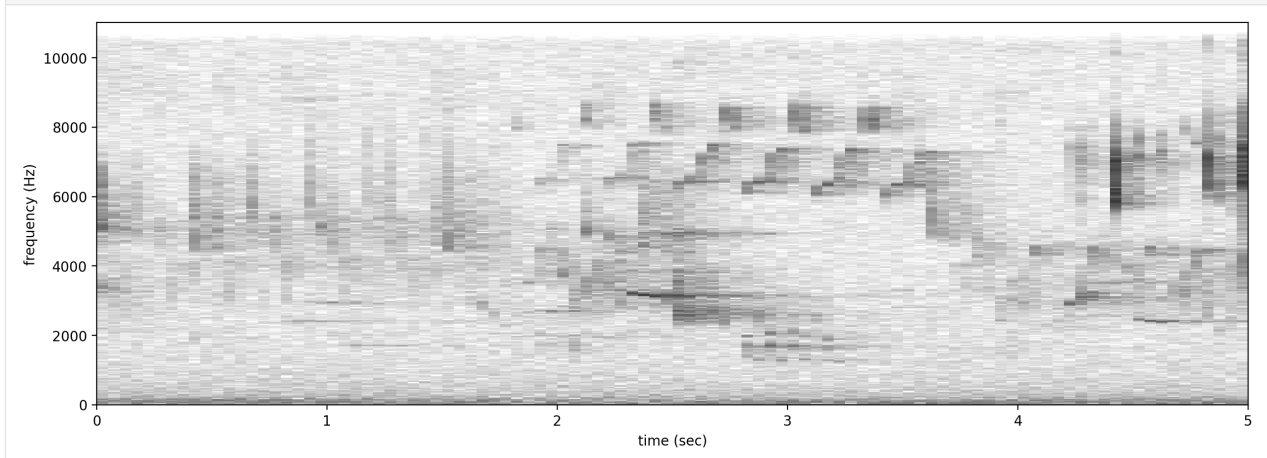
```
[65]: spec = Spectrogram.from_audio(audio, window_samples=55, overlap_samples=0)
spec.plot()
```



### Create a spectrogram with high frequency resolution

Using `window_samples=1100` and `overlap_samples=0` gives  $1100/22000 = 0.05$  seconds of audio per window, or  $1/0.05 = 20$  windows per second. Each spectrogram pixel spans 20 Hz.

```
[66]: spec = Spectrogram.from_audio(audio, window_samples=1100, overlap_samples=0)
spec.plot()
```



For other options when loading spectrogram objects from audio objects, see the `from_audio()` documentation.


## 5.5 Spectrogram methods

The tools and features of the spectrogram class are demonstrated here, including plotting; how spectrograms can be generated from modified audio; saving a spectrogram as an image; customizing a spectrogram; trimming and bandpassing a spectrogram; and calculating the amplitude signal from a spectrogram.

### 5.5.1 Plot

A `Spectrogram` object can be plotted using its `plot()` method.

```
[67]: audio_object = Audio.from_file(audio_filename)
      spectrogram_object = Spectrogram.from_audio(audio_object)
      spectrogram_object.plot()
```



## 5.5.2 Load modified audio

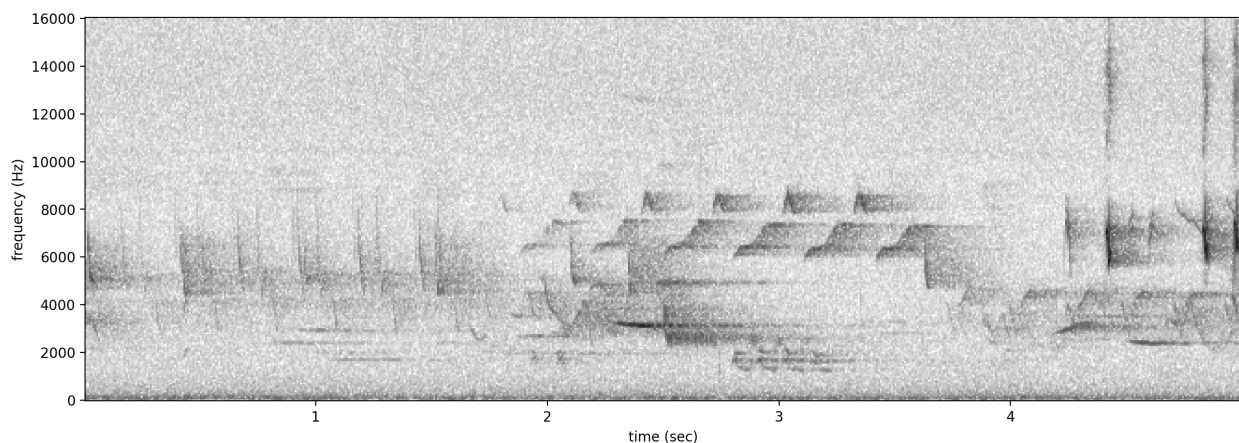
Sometimes, you may wish to trim or modify an audio object before creating a spectrogram. In this case, you should first modify the Audio object, then call `Spectrogram.from_audio()`.

For example, the code below demonstrates creating a spectrogram from a 5 second long trim of the audio object. Compare this plot to the plot above.

```
[68]: # Trim the original audio
      trimmed = audio_object.trim(0, 5)

      # Create a spectrogram from the trimmed audio
      spec = Spectrogram.from_audio(trimmed)

      # Plot the spectrogram
      spec.plot()
```



### 5.5.3 Save spectrogram to file

To save the created spectrogram, first convert it to an image. It will no longer be an OpenSoundscape Spectrogram object, but instead a Python Image Library (PIL) Image object.

```
[69]: print("Type of `spectrogram_audio` (before conversion):", type(spectrogram_object))
spectrogram_image = spectrogram_object.to_image()
print("Type of `spectrogram_image` (after conversion):", type(spectrogram_image))

Type of `spectrogram_audio` (before conversion): <class 'opensoundscape.spectrogram.
↪Spectrogram'>
Type of `spectrogram_image` (after conversion): <class 'PIL.Image.Image'>
```

Save the PIL Image using its `save()` method, supplying the filename at which you want to save the image.

```
[70]: image_path = Path('./saved_spectrogram.png')
spectrogram_image.save(image_path)
```

To save the spectrogram at a desired size, specify the image shape when converting the Spectrogram to a PIL Image.

```
[71]: image_shape = (512, 512)
large_image_path = Path('./saved_spectrogram_large.png')
spectrogram_image = spectrogram_object.to_image(shape=image_shape)
spectrogram_image.save(large_image_path)
```


Delete the files created above.

```
[72]: image_path.unlink()
large_image_path.unlink()
```

### 5.5.4 Trim

Spectrograms can be trimmed in time using `trim()`. Trim the above spectrogram to zoom in on one vocalization.

```
[73]: spec_trimmed = spec.trim(1.7, 3.9)
spec_trimmed.plot()
```





### 5.5.5 Bandpass


Spectrograms can be trimmed in frequency using `bandpass()`. This simply subsets the Spectrogram array rather than performing an audio-domain filter.

For instance, the vocalization zoomed in on above is the song of a Black-and-white Warbler (*Mniotilta varia*), one of the highest-frequency bird songs in our area. Set its approximate frequency range.

```
[74]: baww_low_freq = 5500
      baww_high_freq = 9500
```

Bandpass the above time-trimmed spectrogram in frequency as well to limit the spectrogram view to the vocalization of interest.

```
[75]: spec_bandpassed = spec_trimmed.bandpass(baww_low_freq, baww_high_freq)
      spec_bandpassed.plot()
```




### 5.5.6 Sum the columns of a spectrogram

The `.amplitude()` method sums the columns of the spectrogram to create a one-dimensional amplitude versus time vector.

Note: the amplitude of the Spectrogram (and FFT) has units of power ( $V^{**2}$ ) over frequency (Hz) on a logarithmic scale


```
[76]: # calculate amplitude signal
      high_freq_amplitude = spec_trimmed.amplitude()

      # plot
      from matplotlib import pyplot as plt
      plt.plot(spec_trimmed.times, high_freq_amplitude)
      plt.xlabel('time (sec)')
      plt.ylabel('amplitude')
      plt.show()
```



It is also possible to get the amplitude signal from a restricted range of frequencies, for instance, to look at the amplitude in the frequency range of a species of interest. For example, get the amplitude signal from the 8000 Hz to 8500 Hz range of the audio (displayed below):

```
[77]: spec_bandpassed = spec_trimmed.bandpass(8000, 8500)
spec_bandpassed.plot()
```




Get and plot the amplitude signal of only 8-8.5 kHz.

```
[78]: # Get amplitude signal
high_freq_amplitude = spec_trimmed.amplitude(freq_range=[8000,8500])

# Get amplitude signal
high_freq_amplitude = spec_trimmed.amplitude(freq_range=[8000,8500])

# Plot signal
plt.plot(spec_trimmed.times, high_freq_amplitude)
plt.xlabel('time (sec)')
plt.ylabel('amplitude')
plt.show()
```



Amplitude signals like these can be used to identify periodic calls, like those by many species of frogs. A pulsing-call identification pipeline called [RIBBIT](#) is implemented in OpenSoundscape.

Amplitude signals may not be the most reliable method of identification for species like birds. In this case, it is possible to create a machine learning algorithm to identify calls based on their appearance on spectrograms.

The developers of OpenSoundscape have trained machine learning models for over 500 common North American bird species; for examples of how to download demonstration models, see the [Prediction with pre-trained CNNs](#) tutorial.

### 5.5.7 clean up

```
[79]: #delete the file we downloaded for the tutorial
Path('lmin_audio.wav').unlink()
```



---

## Manipulating audio annotations

---

This notebook demonstrates how to use the `annotations` module of `OpenSoundscape` to:

- load annotations from Raven files
- create a set of one-hot labels corresponding to fixed-length audio clips
- split a set of labeled audio files into clips and create labels dataframe for all clips

The audio recordings used in this notebook were recorded by Andrew Spencer and are available under a Creative Commons License (CC BY-NC-ND 2.5) from [xeno-canto.org](http://xeno-canto.org). Annotations were performed in Raven Pro software by our team.

```
[1]: from opensoundscape.audio import Audio
      from opensoundscape.spectrogram import Spectrogram
      from opensoundscape.annotations import BoxedAnnotations

      import numpy as np
      import pandas as pd
      from glob import glob

      from matplotlib import pyplot as plt
      plt.rcParams['figure.figsize']=[15,5] #for big visuals
      %config InlineBackend.figure_format = 'retina'
```

### 6.1 Download example files

Run the code below to download a set of example audio and raven annotations files for this tutorial.

```
[2]: import subprocess
      subprocess.run(['curl','https://drive.google.com/uc?export=download&
      ↪id=1ZTlV9KzWU0lWsjZpnlrgA92LUuX0lu8a','-L','-o','gwwa_audio_and_raven_annotations.
      ↪tar.gz']) # Download the data
      subprocess.run(["tar","-xzf", "gwwa_audio_and_raven_annotations.tar.gz"]) # Unzip the
      ↪downloaded tar.gz file
```

(continues on next page)

(continued from previous page)

```
subprocess.run(["rm", "gwwa_audio_and_raven_annotations.tar.gz"]) # Remove the file_
↳ after its contents are unzipped
```

| % Total | % Received | % Xferd | Average Speed |        | Time     | Time    | Time     | Current |
|---------|------------|---------|---------------|--------|----------|---------|----------|---------|
|         |            |         | Dload         | Upload | Total    | Spent   | Left     | Speed   |
| 0       | 0          | 0       | 0             | 0      | --:--:-- | 0:00:02 | --:--:-- | 0       |
| 100     | 5432k      | 100     | 5432k         | 0      | 0        | 1764k   | 0        | 0:00:03 |

```
[2]: CompletedProcess(args=['rm', 'gwwa_audio_and_raven_annotations.tar.gz'], returncode=0)
```

### 6.1.1 Load a single Raven annotation table from a txt file

We can use the BoxedAnnotation class's `from_raven_file` method to load a Raven .txt file into OpenSoundscape. This table contains the frequency and time limits of rectangular “boxes” representing each annotation that was created in Raven.

Note that we need to specify the name of the column containing annotations, since it can be named anything in Raven. The column will be renamed to “annotation”.


This table looks a lot like what you would see in the Raven interface.

```
[3]: # specify an audio file and corresponding raven annotation file
audio_file = './gwwa_audio_and_raven_annotations/GWWA_XC/13738.wav'
annotation_file = './gwwa_audio_and_raven_annotations/GWWA_XC_AnnoTables/13738.Table.
↳ 1.selections.txt'
```

Let's look at a spectrogram of the audio file to see what we're working with.

```
[4]: Spectrogram.from_audio(Audio.from_file(audio_file)).plot()

/Users/SML161/miniconda3/envs/opso_dev/lib/python3.9/site-packages/matplotlib_inline/
↳ config.py:68: DeprecationWarning: InlineBackend._figure_format_changed is
↳ deprecated in traitlets 4.1: use @observe and @unobserve instead.
def _figure_format_changed(self, name, old, new):
```



Now, let's load the annotations from the Raven annotation file.

```
[14]: #create an object from Raven file
annotations = BoxedAnnotations.from_raven_files([annotation_file], audio_files=[audio_
↳ file])
```

(continues on next page)

(continued from previous page)

```
#inspect the object's .df attribute, which contains the table of annotations
annotations.df.head()
```

```
[14]:
```

|   | file                                              | annotation | start_time | \ |
|---|---------------------------------------------------|------------|------------|---|
| 0 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 0.459636   |   |
| 1 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 6.705283   |   |
| 2 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | ?          | 13.464641  |   |
| 3 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 20.128208  |   |
| 4 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 26.047590  |   |

|   | end_time  | low_f  | high_f  | Channel | Notes | index | View          | Selection |
|---|-----------|--------|---------|---------|-------|-------|---------------|-----------|
| 0 | 2.298182  | 4029.8 | 17006.4 | 1       | NaN   | 0     | Spectrogram 1 | 1         |
| 1 | 8.246417  | 4156.6 | 17031.7 | 1       | NaN   | 1     | Spectrogram 1 | 2         |
| 2 | 15.005775 | 3903.1 | 17082.4 | 1       | NaN   | 2     | Spectrogram 1 | 3         |
| 3 | 21.601748 | 4055.2 | 16930.3 | 1       | NaN   | 3     | Spectrogram 1 | 4         |
| 4 | 27.521131 | 4207.2 | 17057.1 | 1       | NaN   | 4     | Spectrogram 1 | 5         |

Note: if you **do not have an annotation column**, e.g., if you are annotating the sounds of a single species, pass the argument `annotation_column_idx=None`. The resulting dataframe will have an empty annotation column.

We could instead choose to only load the necessary columns (`start_time`, `end_time`, `low_f`, `high_f`, and `annotation`) using the `keep_extra_columns=None`.

In this example, we use `keep_extra_columns=[ 'Notes' ]` to keep only the Notes column.

```
[24]: annotations = BoxedAnnotations.from_raven_files([annotation_file], keep_extra_columns=[
↳ 'Notes'], audio_files=[audio_file])
annotations.df.head()
```

```
[24]:
```

|   | file                                              | annotation | start_time | \ |
|---|---------------------------------------------------|------------|------------|---|
| 0 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 0.459636   |   |
| 1 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 6.705283   |   |
| 2 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | ?          | 13.464641  |   |
| 3 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 20.128208  |   |
| 4 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 26.047590  |   |

|   | end_time  | low_f  | high_f  | index | Notes |
|---|-----------|--------|---------|-------|-------|
| 0 | 2.298182  | 4029.8 | 17006.4 | 0     | NaN   |
| 1 | 8.246417  | 4156.6 | 17031.7 | 1     | NaN   |
| 2 | 15.005775 | 3903.1 | 17082.4 | 2     | NaN   |
| 3 | 21.601748 | 4055.2 | 16930.3 | 3     | NaN   |
| 4 | 27.521131 | 4207.2 | 17057.1 | 4     | NaN   |

## 6.1.2 Convert or correct annotations

We can provide a DataFrame (e.g., from a .csv file) or a dictionary to convert original values to new values.

Let's load up a little .csv file that specifies a set of conversions we'd like to make. The .csv file should have two columns, but it doesn't matter what they are called. If you create a table in Microsoft Excel, you can export it to a .csv file to use it as your conversion table.

```
[25]: conversion_table = pd.read_csv('./gwwa_audio_and_raven_annotations/conversion_table.
↳ csv')
conversion_table
```

```
[25]:
```

|   | original  | new  |
|---|-----------|------|
| 0 | gwwa_song | gwwa |

Alternatively, we could simply write a Python dictionary for the conversion table. For instance:

```
[26]: conversion_table = {
      "GWWA_song": "GWWA",
      "?": np.nan
    }
```

Now, we can apply the conversions in the table to our annotations.

This will create a new `BoxedAnnotations` object rather than modifying the original object (an “out of place operation”).

```
[27]: annotations_corrected = annotations.convert_labels(conversion_table)
      annotations_corrected.df
```

```
[27]:
```

|   | file                                              | annotation | start_time | \ |
|---|---------------------------------------------------|------------|------------|---|
| 0 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA       | 0.459636   |   |
| 1 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA       | 6.705283   |   |
| 2 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | NaN        | 13.464641  |   |
| 3 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA       | 20.128208  |   |
| 4 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA       | 26.047590  |   |
| 5 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA       | 33.121946  |   |
| 6 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA       | 42.967925  |   |
| 7 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA       | 52.417508  |   |

|   | end_time  | low_f  | high_f  | index | Notes |
|---|-----------|--------|---------|-------|-------|
| 0 | 2.298182  | 4029.8 | 17006.4 | 0     | NaN   |
| 1 | 8.246417  | 4156.6 | 17031.7 | 1     | NaN   |
| 2 | 15.005775 | 3903.1 | 17082.4 | 2     | NaN   |
| 3 | 21.601748 | 4055.2 | 16930.3 | 3     | NaN   |
| 4 | 27.521131 | 4207.2 | 17057.1 | 4     | NaN   |
| 5 | 34.663079 | 4207.2 | 17082.4 | 5     | NaN   |
| 6 | 44.427946 | 4181.9 | 17057.1 | 6     | NaN   |
| 7 | 53.891048 | 4232.6 | 16930.3 | 7     | NaN   |

## 6.2 View a subset of annotations

We can specify a list of classes to view annotations of.

For example, we can subset to only annotations marked as “?” - perhaps we’re interested in looking at these annotations in Raven again to determine what class they really were.

```
[28]: classes_to_keep = ['?']
      annotations_only_unsure = annotations.subset(classes_to_keep)
      annotations_only_unsure.df
```

```
[28]:
```

|   | file                                              | annotation | start_time | \ |
|---|---------------------------------------------------|------------|------------|---|
| 2 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | ?          | 13.464641  |   |

|   | end_time  | low_f  | high_f  | index | Notes |
|---|-----------|--------|---------|-------|-------|
| 2 | 15.005775 | 3903.1 | 17082.4 | 2     | NaN   |

## 6.3 Saving annotations to Raven-compatible file

We can always save our `BoxedAnnotations` object to a Raven-compatible .txt file, which can be opened in Raven along with an audio file just like the .txt files Raven creates itself. You must specify a path for the save file that ends with

.txt.

```
[29]: annotations_only_unsure.to_raven_files('./gwwa_audio_and_raven_annotations/') #13738_
      ↪ unsure.txt')
```

### 6.3.1 Split an audio clip and its annotations

Often, we want to train or validate models on short audio segments (e.g., 5 seconds) rather than on long files (e.g., 2 hours).

We can easily create tables of “one hot” encoded labels for a series of audio segments within each annotated file using `BoxedAnnotations.one_hot_clip_labels()`

#### \*What is one-hot encoding?\*

The functions below demonstrate the creation of one-hot encoded labels.

This machine learning term, “one-hot encoding,” refers to a way to format a table of labels in which:

- \* Each row represents a single sample, like a single 5-second long clip
- \* Each column represents a single possible class (e.g. one of multiple species)
- \* A “0” in a row and column means that in that sample, the class is not present
- \* A “1” is “hot,” meaning that in that sample, the class *IS* present.

For example, let’s say we had a 15-second audio clip that we were splitting into three 5s clips. Let’s say we are training a classifier to identify coyotes and dogs, and we labeled the clip and found:

- \* a coyote howled from 2.5 to 4 seconds into the clip (so, only the first clip contains it)
- \* a dog barked from 4 seconds to 10 seconds into the clip (so, both the first and second clips contain it)
- \* and there was silence for the last 5 seconds of the clip (so, the third clip has neither coyotes nor dogs in it).

The one-hot encoded labels file for this example would look like:

```
[30]: pd.DataFrame({
      "start_time": [0, 5, 10],
      "end_time": [5, 10, 15],
      "COYOTE": [1, 0, 0],
      "DOG": [1, 1, 0]
    })
```

```
[30]:
```

|   | start_time | end_time | COYOTE | DOG |
|---|------------|----------|--------|-----|
| 0 | 0          | 5        | 1      | 1   |
| 1 | 5          | 10       | 0      | 1   |
| 2 | 10         | 15       | 0      | 0   |

## 6.4 1. Split Audio object, then split annotations to match

First, split an Audio object with `Audio.split()`, which returns two things:

1. A list of audio clip objects
2. A dataframe of start/end times

```
[34]: # Load the Audio and Annotations
      audio = Audio.from_file(audio_file)
      annotations = BoxedAnnotations.from_raven_files([annotation_file], audio_files=[audio_
      ↪ file])
```

Different overlap and duration settings produce different results:

```
[ ]: # Split the audio into 5 second clips with no overlap (we use _ because we don't_
      ↪really need to save the audio clip objects for this demo)
_, clip_df_short_with_overlaps = audio.split(
    clip_duration=3.0, # How long each clip should be
    clip_overlap=1.0 # By how many seconds each subsequent clip should overlap
)
clip_df_short_with_overlaps.head()
```

Next, extract annotations for each clip using `BoxedAnnotations.one_hot_labels_like()`.

```
[ ]: # Split the annotations using the returned clip_df
labels_df = annotations.one_hot_labels_like(
    clip_df,
    min_label_overlap=0.25, # Minimum label overlap
    class_subset=['GWWA_song']
)

#the returned dataframe of one-hot labels (0/1 for each class and each clip) has rows_
↪corresponding to each audio clip
labels_df.head()
```

## 6.5 Split annotations using splitting parameters

This function requires that we specify the minimum overlap of the label (in seconds) with the clip for the clip to be labeled positive. It also requires that we either (1) specify the list of classes for one-hot labels or (2) specify `class_subset=None`, which will make a column for every unique label in the annotations. In this example, that would include a “?” class

```
[37]: labels_df = annotations.one_hot_clip_labels(
    full_duration=60, # The duration of the entire audio file
    clip_duration=5,
    clip_overlap=0,
    class_subset=['GWWA_song'],
    min_label_overlap=0.25,
)
labels_df.head()
```

```
[37]:
```

|                                                    |            |          | GWWA_song |
|----------------------------------------------------|------------|----------|-----------|
| file                                               | start_time | end_time |           |
| ./gwwa_audio_and_raven_annotations/GWWA_XC/1373... | 0          | 5        | 1.0       |
|                                                    | 5          | 10       | 1.0       |
|                                                    | 10         | 15       | 0.0       |
|                                                    | 15         | 20       | 0.0       |
|                                                    | 20         | 25       | 1.0       |

### 6.5.1 A data munging example: pairing Raven files and audio to create a labeled dataset

In practice, we have tons of audio files with their corresponding Raven files. We need to:

- Pair up all the audio files with their Raven annotation files
- Create a dataframe of labels corresponding to short segments of each audio file

Let's walk through the steps required to do this. But be warned, pairing Raven files and audio files might require more finagling than shown here.

## 6.6 Match up audio files and Raven annotations

The first step in the process is associating audio files with their corresponding Raven files. Perhaps not every audio file is annotated, and perhaps some audio files have been annotated multiple times. This code walks through some steps of sorting through these data to pair files.

**Caveat: you'll need to be careful** using the code below, depending on how your audio and Raven files are named and organized.

In this example, we'll assume that each audio file has the same name as its Raven annotation file (ignoring the extensions like ".Table.1.selections.txt"), which is the default naming convention when using Raven. We'll also start by assuming that the audio filenames *are unique (!)* - that is, no two audio files have the same name.

First, find all the Raven files and all the audio files.

```
[54]: # Specify folder containing Raven annotations
raven_files_dir = "./gwwa_audio_and_raven_annotations/GWWA_XC_AnnoTables/"

# Find all .txt files
# We'll naively assume all files with the suffix ".txt" are Raven files!
# A better assumption could be to search for files with the suffix ".selections.txt"
raven_files = glob(f"{raven_files_dir}/*.txt")
print(f"found {len(raven_files)} annotation files")

# Specify folder containing audio files
audio_files_dir = "./gwwa_audio_and_raven_annotations/GWWA_XC/"

# Find all audio files (we'll assume they are .wav, .WAV, or .mp3)
audio_files = glob(f"{audio_files_dir}/*.wav")+glob(f"{audio_files_dir}/*.WAV")+glob(f
↪ "{audio_files_dir}/*.mp3")
print(f"found {len(audio_files)} audio files")

found 3 annotation files
found 3 audio files
```

Next, starting by assuming that audio files have unique names, use the audio filenames to pair up the annotation files. Then, double-check that our assumption is correct.

```
[55]: # Pair up the Raven and audio files based on the audio file name
from pathlib import Path
audio_df = pd.DataFrame({'audio_file': audio_files})
audio_df.index = [Path(f).stem for f in audio_files]

# Check that there aren't duplicate audio file names
print('\n audio files with duplicate names:')
audio_df[audio_df.index.duplicated(keep=False)]

audio files with duplicate names:
```

```
[55]: Empty DataFrame
Columns: [audio_file]
Index: []
```

Seeing that no audio files have duplicate names, check to make sure the same is true for Raven files.

```
[56]: raven_df = pd.DataFrame({'raven_file':raven_files})
raven_df.index = [Path(f).stem.split('.')[0] for f in raven_files]

#check that there aren't duplicate audio file names
print('\n raven files with duplicate names:')
raven_df[raven_df.index.duplicated(keep=False)]
```

raven files with duplicate names:

```
[56]:                                     raven_file
13738  ./gwwa_audio_and_raven_annotations/GWWA_XC_Ann...
13738  ./gwwa_audio_and_raven_annotations/GWWA_XC_Ann...
```

Since we found some duplicate Raven files, resolve this issue by deleting the extra Raven file, which in this case was named “selections2”.

```
[57]: #remove the second selection table for file 13738.wav
raven_df=raven_df[raven_df.raven_file.apply(lambda x: "selections2" not in x)]
```

Once we’ve resolved any issues with duplicate names, we can match up Raven and audio files.

```
[58]: paired_df = audio_df.join(raven_df,how='outer')
```

Check if any audio files don’t have Raven annotation files:

```
[59]: print(f"audio files without raven file: {len(paired_df[paired_df.raven_file.
↪apply(lambda x:x!=x)])}")
paired_df[paired_df.raven_file.apply(lambda x:x!=x)]

audio files without raven file: 2
```

```
[59]:                                     audio_file raven_file
135601  ./gwwa_audio_and_raven_annotations/GWWA_XC/135...      NaN
13742   ./gwwa_audio_and_raven_annotations/GWWA_XC/137...      NaN
```

Check if any Raven files don’t have audio files:

```
[60]: #look at unmatched raven files
print(f"raven files without audio file: {len(paired_df[paired_df.audio_file.
↪apply(lambda x:x!=x)])}")

paired_df[paired_df.audio_file.apply(lambda x:x!=x)]

raven files without audio file: 1
```

```
[60]:                                     audio_file                                     raven_file
16989      NaN  ./gwwa_audio_and_raven_annotations/GWWA_XC_Ann...
```

In this example, let’s discard any unpaired Raven or audio files.

```
[61]: paired_df = paired_df.dropna()
```

```
[62]: paired_df
```

```
[62]:                                     audio_file \
13738  ./gwwa_audio_and_raven_annotations/GWWA_XC/137...

                                     raven_file
13738  ./gwwa_audio_and_raven_annotations/GWWA_XC_Ann...
```



## 6.7 Create label dataframes

Now we have a set of paired up Raven and audio files.

Let's create label dataframes representing 3-second segments of each audio file

```
[63]: # Choose settings for audio splitting
clip_duration = 3
clip_overlap = 0
final_clip = None
```

Next, set up the settings for annotation splitting:

- Whether to use a subset of classes
- How many seconds a label should overlap a clip, at minimum, in order for that clip to be labeled

```
[64]: # Choose settings for annotation splitting
class_subset = None #Equivalent to a list of all classes: ['GWWA_song', '?']
min_label_overlap = 0.1
```

## 6.8 Load Raven annotations

```
[71]: boxed_annotations = BoxedAnnotations.from_raven_files(paired_df.raven_file,paired_df.
↳audio_file)
boxed_annotations.df.head(3)
```

```
[71]:
```

|   | file                                              | annotation | start_time | \ |
|---|---------------------------------------------------|------------|------------|---|
| 0 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 0.459636   |   |
| 1 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | GWWA_song  | 6.705283   |   |
| 2 | ./gwwa_audio_and_raven_annotations/GWWA_XC/137... | ?          | 13.464641  |   |

|   | end_time  | low_f  | high_f  | Channel | Notes | index | View          | Selection |
|---|-----------|--------|---------|---------|-------|-------|---------------|-----------|
| 0 | 2.298182  | 4029.8 | 17006.4 | 1       | NaN   | 0     | Spectrogram 1 | 1         |
| 1 | 8.246417  | 4156.6 | 17031.7 | 1       | NaN   | 1     | Spectrogram 1 | 2         |
| 2 | 15.005775 | 3903.1 | 17082.4 | 1       | NaN   | 2     | Spectrogram 1 | 3         |

## 6.9 Create label dataframes

```
[75]: label_df = boxed_annotations.one_hot_clip_labels(
clip_duration=clip_duration,
clip_overlap=clip_overlap,
min_label_overlap=min_label_overlap,
class_subset=class_subset,
final_clip=final_clip,
)
```

```
[76]: label_df.head(2)
```

```
[76]:
```

| file                                               | start_time | end_time | GWWA_song | \ |
|----------------------------------------------------|------------|----------|-----------|---|
| ./gwwa_audio_and_raven_annotations/GWWA_XC/1373... | 0.0        | 3.0      | 1.0       |   |
|                                                    | 3.0        | 6.0      | 0.0       |   |

(continues on next page)

(continued from previous page)

| file                                               | start_time | end_time | ?   |
|----------------------------------------------------|------------|----------|-----|
| ./gwwa_audio_and_raven_annotations/GWWA_XC/1373... | 0.0        | 3.0      | 0.0 |
|                                                    | 3.0        | 6.0      | 0.0 |

## 6.10 Sanity check: look at spectrograms of clips labeled 0 and 1


```
[78]: # ignore the "?" annotations for this visualization
label_df = label_df[label_df["?"]==0]
```



Note: replace the “GWWA\_song” here with a class name from your own dataset.

```
[87]: # plot spectrograms for 3 random positive clips
positives = label_df[label_df['GWWA_song']==1].sample(3,random_state=0)
print("spectrograms of 3 random positive clips (label=1)")
for clip, t0, t1 in positives.index.values:
    Spectrogram.from_audio(Audio.from_file(clip,offset=t0,duration=t1-t0)).plot()


# plot spectrograms for 5 random negative clips
negatives = label_df[label_df['GWWA_song']==0].sample(3,random_state=0)
print("spectrogram of 3 random negative clips (label=0)")
for clip, t0, t1 in negatives.index.values:
    Spectrogram.from_audio(Audio.from_file(clip,offset=t0,duration=t1-t0)).plot()
```


spectrograms of 3 random positive clips (label=1)





spectrogram of 3 random negative clips (label=0)





Clean up: remove the sounds that we downloaded for this tutorial as well as the audio files we created.

```
[ ]: import shutil
shutil.rmtree('./gwa_audio_and_raven_annotations')
```

---

## Prediction with pre-trained CNNs

---

This notebook contains all the code you need to use a pre-trained OpenSoundscape convolutional neural network model (CNN) to make predictions on your own data. Before attempting this tutorial, install OpenSoundscape by following the instructions on the OpenSoundscape website, [opensoundscape.org](https://opensoundscape.org). More detailed tutorials about data preprocessing, training CNNs, and customizing prediction methods can also be found on this site.

### 7.1 Load required packages

The `cnn` module provides a function `load_model` to load saved opensoundscape models

```
[1]: from opensoundscape.ml.cnn import load_model
import opensoundscape
```

load some additional packages and perform some setup for the Jupyter notebook.

```
[2]: # Other utilities and packages
import torch
from pathlib import Path
import numpy as np
import pandas as pd
from glob import glob
import subprocess
```

```
[3]: #set up plotting
from matplotlib import pyplot as plt
plt.rcParams['figure.figsize']=[15,5] #for large visuals
%config InlineBackend.figure_format = 'retina'
```

For this example, let's create an untrained model and save it. This 2-class model is not actually good at recognizing any particular species, but it's useful for illustrating how prediction works.

```
[4]: from opensoundscape import CNN
CNN('resnet18', ['classA', 'classB'], 5.0).save('./temp.model')
```

### 7.1.1 Load a saved model

load the model object using the `load_model` function imported above

(if the model was created with an older version of opensoundscape, see instructions below)

```
[5]: model = load_model('./temp.model')
```

### 7.1.2 Choose audio files for prediction

Create a list of audio files to predict on. They can be of any length. Consider using `glob` to find many files at once.

For this example, let's download a 1-minute audio clip from the Kitzes Lab box to use as an example.

```
[6]: subprocess.run(['curl',
                    'https://drive.google.com/uc?export=download&id=1hBUyhgPf-
                    ↪vxeet8BITSMZ0XsGiU-a7SG',
                    '-L', '-o', '1min_audio.wav'])
```

| % Total | % Received | % Xferd | Average Speed | Time    | Time     | Time     | Current |
|---------|------------|---------|---------------|---------|----------|----------|---------|
|         |            |         | Dload Upload  | Total   | Spent    | Left     | Speed   |
| 0       | 0          | 0       | 0             | 0       | --:--:-- | --:--:-- | 0       |
| 100     | 3750k      | 100     | 3750k         | 0       | 0        | 2302k    | 0       |
|         |            |         |               | 0:00:01 | 0:00:01  | --:--:-- | 6148k   |

```
[6]: CompletedProcess(args=['curl', 'https://drive.google.com/uc?export=download&
    ↪id=1hBUyhgPf-vxeet8BITSMZ0XsGiU-a7SG', '-L', '-o', '1min_audio.wav'], returncode=0)
```

use `glob` to create a list of all files matching a pattern in a folder:

```
[7]: from glob import glob
audio_files = glob('./*.wav') #match all .wav files in the current directory
audio_files

[7]: ['./1min_audio.wav']
```

## 7.2 generate predictions with the model

The model returns a dataframe with a MultiIndex of file, start\_time, and end\_time. There is one column for each class.

```
[8]: scores = model.predict(audio_files)
scores.head()
```

```
[8]:
```

|                  |            |          | classA    | classB    |
|------------------|------------|----------|-----------|-----------|
| file             | start_time | end_time |           |           |
| ./1min_audio.wav | 0.0        | 5.0      | -0.007259 | 0.189202  |
|                  | 5.0        | 10.0     | -0.573235 | 0.185848  |
|                  | 10.0       | 15.0     | -0.514158 | 0.300901  |
|                  | 15.0       | 20.0     | -0.659756 | 0.067114  |
|                  | 20.0       | 25.0     | -0.718755 | -0.011302 |

### 7.3 Overlapping prediction clips

```
[9]: scores = model.predict(audio_files, overlap_fraction=0.5)
scores.head()
```

```
[9]:
```

|                  |            |          | classA    | classB   |
|------------------|------------|----------|-----------|----------|
| file             | start_time | end_time |           |          |
| ./lmin_audio.wav | 0.0        | 5.0      | -0.007259 | 0.189202 |
|                  | 2.5        | 7.5      | -0.270948 | 0.161639 |
|                  | 5.0        | 10.0     | -0.573235 | 0.185848 |
|                  | 7.5        | 12.5     | -0.428415 | 0.321352 |
|                  | 10.0       | 15.0     | -0.514158 | 0.300901 |

## 7.4 Inspect samples generated during prediction

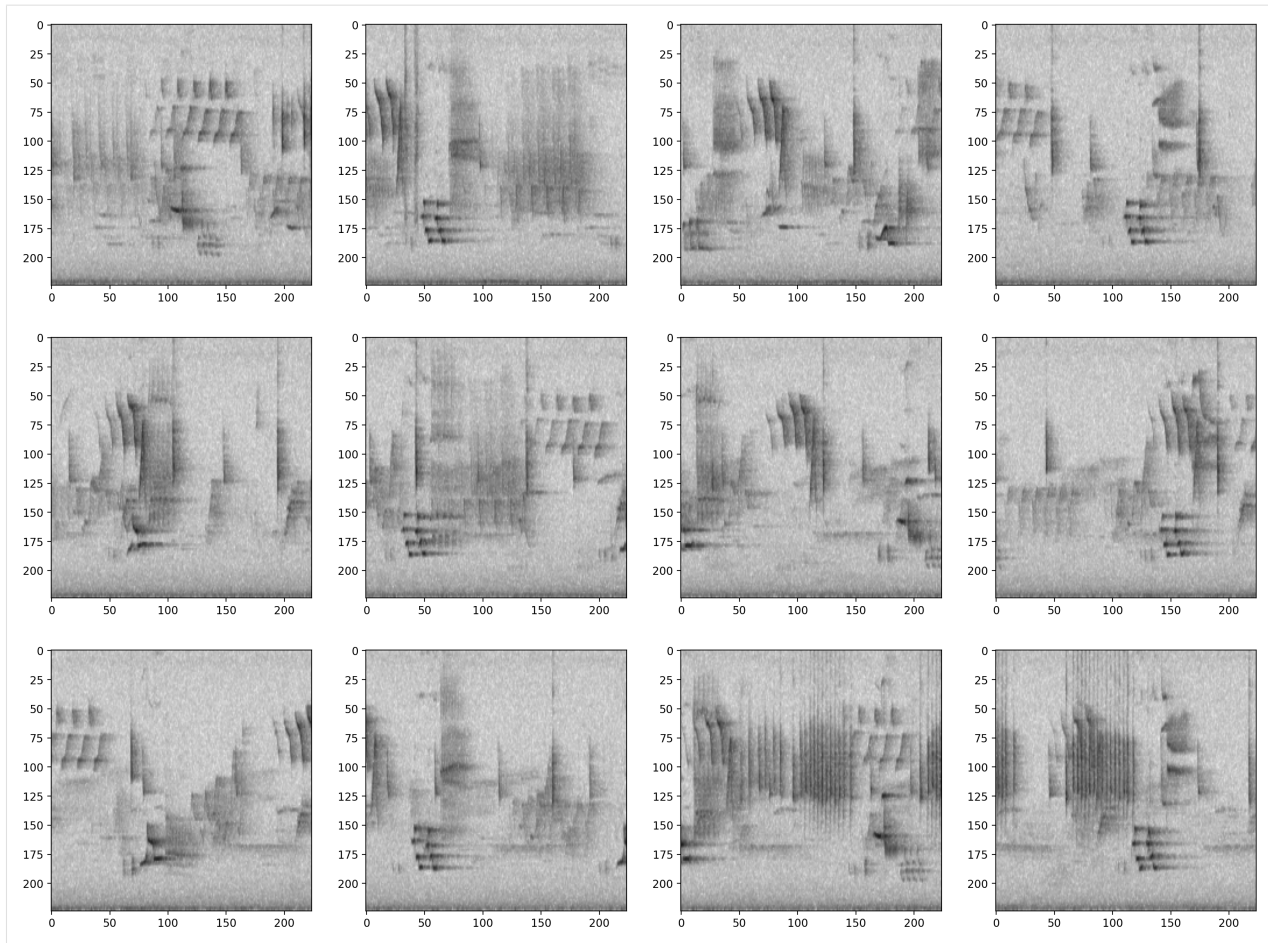
```
[10]: from opensoundscape.preprocess.utils import show_tensor_grid
from opensoundscape.ml.datasets import AudioSplittingDataset

#generate a dataset with the samples we wish to generate and the model's preprocessor
inspection_dataset = AudioSplittingDataset(audio_files, model.preprocessor)
inspection_dataset.bypass_augmentations = True

samples = [sample.data for sample in inspection_dataset]
_ = show_tensor_grid(samples, 4)

/Users/SML161/miniconda3/envs/opso_dev/lib/python3.9/site-packages/matplotlib_inline/
↳ config.py:68: DeprecationWarning: InlineBackend._figure_format_changed is
↳ deprecated in traitlets 4.1: use @observe and @unobserve instead.
def _figure_format_changed(self, name, old, new):
```





## 7.5 Options for prediction

The code above returns the raw predictions of the model without any post-processing (such as a softmax layer or a sigmoid layer).

For details on how to post-process prediction scores and to generate binary 0/1 predictions of class presence, see the “Basic training and prediction with CNNs” tutorial notebook. But, as a quick example here, let’s add a softmax layer to make the prediction scores for both classes sum to 1.

We can also convert our continuous scores into True/False (or 1/0) predictions for the presence of each class in each sample. Think about whether each clip should be labeled with only one class (use `metrics.predict_single_target_labels`) or whether each clip could contain zero, one, or multiple classes (use `metrics.predict_multi_target_labels`)

```
[11]: scores = model.predict(
    audio_files,
    activation_layer='softmax',
)
```

As before, the `scores` are continuous variables, but now have been softmaxed:

```
[12]: scores.head()
```



```
[12]:
```

|  | file             | start_time | end_time | classA   | classB   |
|--|------------------|------------|----------|----------|----------|
|  | ./lmin_audio.wav | 0.0        | 5.0      | 0.451042 | 0.548958 |
|  |                  | 5.0        | 10.0     | 0.318845 | 0.681155 |
|  |                  | 10.0       | 15.0     | 0.306814 | 0.693187 |
|  |                  | 15.0       | 20.0     | 0.325882 | 0.674118 |
|  |                  | 20.0       | 25.0     | 0.330162 | 0.669838 |

Now let's use the `predict_single_target_labels(scores)` function to label the highest scoring class 1 for each sample, and other classes 0.

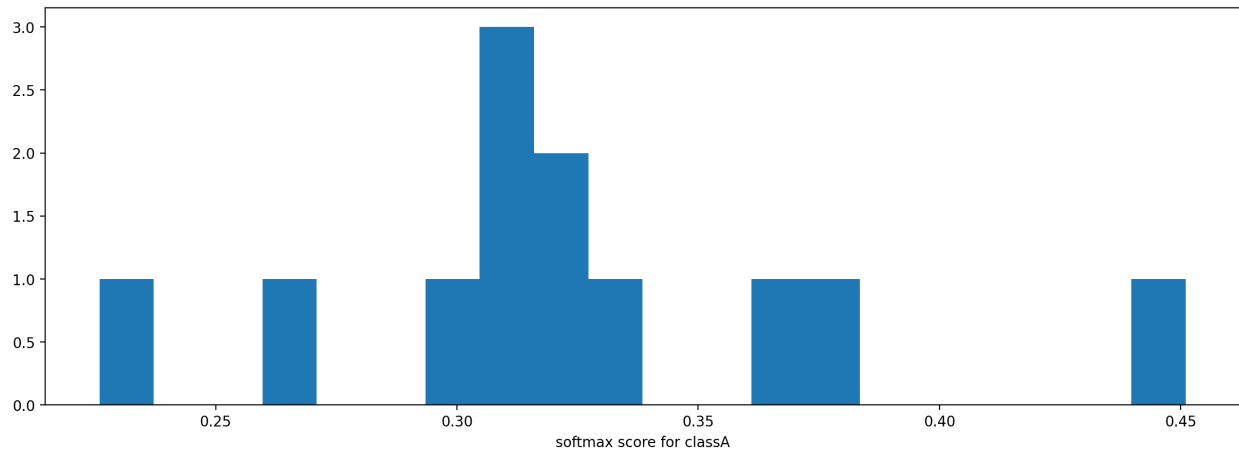
```
[13]: from opensoundscape.metrics import predict_single_target_labels
predicted_labels = predict_single_target_labels(scores)
predicted_labels.head()
```

```
[13]:
```

|  | file             | start_time | end_time | classA | classB |
|--|------------------|------------|----------|--------|--------|
|  | ./lmin_audio.wav | 0.0        | 5.0      | 0      | 1      |
|  |                  | 5.0        | 10.0     | 0      | 1      |
|  |                  | 10.0       | 15.0     | 0      | 1      |
|  |                  | 15.0       | 20.0     | 0      | 1      |
|  |                  | 20.0       | 25.0     | 0      | 1      |

It is sometimes helpful to look at a histogram of the scores:

```
[14]: _ = plt.hist(scores['classA'], bins=20)
_ = plt.xlabel('softmax score for classA')
```



## 7.6 Using models from older OpenSoundscape versions

### 7.6.1 Models from OpenSoundscape 0.4.x and 0.5.x

Models trained and saved with OpenSoundscape versions 0.4.x and 0.5.x need to be loaded in a different way, and require that you know the architecture of the saved model.

For example, one set of our publicly available [binary models for 500 species](#) was created with an older version of OpenSoundscape. These models require a little bit of manipulation to load into OpenSoundscape 0.5.x and onward.

First, let's download one of these models (it's stored in a .tar format) and save it to the same directory as this notebook in a file called `opso_04_model_acanthis-flammea.tar`

```
[15]: subprocess.run(['curl',
                    'https://drive.google.com/uc?export=download&id=lp0RYBtdQ_
↪0A7L6PVUx6OkZPIOqNZZJa3',
                    '-L', '-o', 'opso_04_model_acanthis-flammea.tar'])
```

| % Total | % Received | % Xferd | Average Speed | Time  | Time     | Time    | Current |
|---------|------------|---------|---------------|-------|----------|---------|---------|
|         |            |         | Dload Upload  | Total | Spent    | Left    | Speed   |
| 0       | 0          | 0       | 0             | 0     | --:--:-- | 0:00:02 | 0       |
| 100     | 42.9M      | 100     | 42.9M         | 0     | 0        | 0:00:05 | 16.4M   |

```
[15]: CompletedProcess(args=['curl', 'https://drive.google.com/uc?export=download&
↪id=lp0RYBtdQ_0A7L6PVUx6OkZPIOqNZZJa3', '-L', '-o', 'opso_04_model_acanthis-flammea.
↪tar'], returncode=0)
```

From the model notes page, we know that this is a single-target model with a resnet18 architecture trained on 5 second files. Let's load the model with `load_outdated_model`. We also need to make sure we use the same preprocessing settings as the original model. In this case, the original model used the same preprocessing settings as the default CNN.preprocessor.

```
[16]: from opensoundscape.ml.cnn import load_outdated_model
```

```
[17]: model = load_outdated_model('./opso_04_model_acanthis-flammea.tar', 'resnet18', 5.0)

#invert values to match the convention of OpenSoundscape 0.7.x (lowest values = quiet,
↪ highest = loud)
model.preprocessor.pipeline.to_tensor.set(invert=True)
```

mismatched keys:  
<All keys matched successfully>

/Users/SML161/opensoundscape/opensoundscape/ml/cnn.py:1792: UserWarning: After\_  
↪loading a model, you still need to ensure that your preprocessing (model.  
↪preprocessor) matches the settings used to createthe original model.  
warnings.warn(

Again, you may need to modify `model.preprocessor` to match the settings used to train the model.

The model is now fully compatible with OpenSoundscape, and can be used as above. For example:

```
[18]: scores = model.predict(audio_files)
scores.head()
```

```
[18]:
```

| file             | start_time | end_time | acanthis-flammea-absent \ |
|------------------|------------|----------|---------------------------|
| ./lmin_audio.wav | 0.0        | 5.0      | 6.254239                  |
|                  | 5.0        | 10.0     | 4.935342                  |
|                  | 10.0       | 15.0     | 6.227312                  |
|                  | 15.0       | 20.0     | 5.256021                  |
|                  | 20.0       | 25.0     | 4.836051                  |

| file             | start_time | end_time | acanthis-flammea-present |
|------------------|------------|----------|--------------------------|
| ./lmin_audio.wav | 0.0        | 5.0      | -5.859637                |
|                  | 5.0        | 10.0     | -4.917323                |
|                  | 10.0       | 15.0     | -5.752949                |
|                  | 15.0       | 20.0     | -5.732774                |
|                  | 20.0       | 25.0     | -5.272484                |

if we save the model using `model.save(path)`, we can re-load the full model object later using `load_model()` rather than repeating the procedure above.

## 7.6.2 Loading models from OpenSoundscape 0.6.0

If you saved a model with OpenSoundscape 0.6.0 and want to use it in 0.7.0 or above, you will need to re-load the model using the original OpenSoundscape version that it was created with and save the model's weights explicitly. Here's an example of code you could run *in an environment with opensoundscape version 0.6.0* to export a model for use in later OpenSoundscape versions:

```
# This code will only work in an environment with OpenSoundscape version 0.6.0
# Use it if you need to save a model created in OpenSoundscape v0.6.0 for use in
↳later opso versions

import torch
from opensoundscape.torch.models.cnn import Resnet18Binary #choose the class used to
↳create the model

model = Resnet18Binary(classes=['negative','positive']) #provide the list of classes
↳for the model
model.load('/path/to/saved.model')

dict_to_save = {
    'network_state_dict':model.network.state_dict(),
    'classes': model.classes,
}
torch.save(dict_to_save, '/path/to/model_dict.pt')
```

Then, you will be able to create a new model object in OpenSoundscape  $\geq 0.7.0$  and load the weights from the state dict as demonstrated above. Make sure to specify the correct architecture and sample duration when you create the CNN object.

```
#run this code in an envrionment with a newer OpenSoundscape version  $\geq 0.9.0$ 

import torch
from opensoundscape.ml.cnn import CNN

model_dict = torch.load('/path/to/model_dict.pt')
classes = model_dict["classes"]

#remove the 'feature' prefix on weights and replace the 'classifier' prefix with 'fc'
model_dict['network_state_dict'] = {
    k.replace('feature.','').replace('classifier.','fc.'):v
    for k, v in model_dict['network_state_dict'].items()
}

architecture = 'resnet18' #match this with the original model!

sample_duration = 5.0 #match this with the original model!

model = CNN('resnet18',classes,sample_duration)
model.network.load_state_dict(model_dict['network_state_dict'])

#invert values to match the convention of OpenSoundscape 0.7.x
model.preprocessor.pipeline.to_tensor.set(invert=True)

#save the model object so that we can simply reload it with load_model() in the
↳future:
model.save('/path/to/saved_full_object.model')
```

(continues on next page)

(continued from previous page)

```
# Next time, we can just load the full model object directly:
from opensoundscape.ml.cnn import load_model
model = load_model('/path/to/saved_full_object.model')
```

### 7.6.3 Loading models from OpenSoundscape 0.6.1 and 0.6.2

If you saved a model with OpenSoundscape 0.6.1 or 0.6.2 and want to use it in 0.7.0 or above, you will need to re-load the model using the original OpenSoundscape version that it was created with and save the model's weights explicitly. Here's an example of code you could run *in an environment with opensoundscape version 0.6.1 or 0.6.2* to export a model for use in later OpenSoundscape versions:

```
# This code will only work in an environment with OpenSoundscape version 0.6.1 or 0.6.
↪2
# Use it if you need to save a model created in OpenSoundscape v0.6.1 or 0.6.2 for_
↪use in later opso
versions

import torch
from opensoundscape.ml.cnn import load_model #opensoundscape.torch.models.cnn in v0.7.
↪0 or v0.8.0
model = load_model('/path/to/saved.model')

dict_to_save = {
    'network_state_dict':model.network.state_dict(),
    'classes': model.classes,
    '
}
torch.save(dict_to_save, '/path/to/model_dict.pt')
```

Then, you will be able to create a new model object in OpenSoundscape 0.9.0 and load the weights from the state dict as demonstrated above. Make sure to specify the correct architecture and sample duration when you create the CNN object.

```
#run this code in an envrionment with a newer OpenSoundscape version >=0.9.0

import torch
from opensoundscape.ml.cnn import CNN

model_dict = torch.load('/path/to/model_dict.pt')
classes = model_dict["classes"]

architecture = 'resnet18' #match this with the original model!

sample_duration = 5.0 #match this with the original model!

model = CNN('resnet18',classes,sample_duration)
model.network.load_state_dict(model_dict['network_state_dict'])

#invert values to match the convention of OpenSoundscape 0.7.x
model.preprocessor.pipeline.to_tensor.set(invert=True)

#save the model object so that we can simply reload it with load_model() in the_
↪future:
model.save('/path/to/saved_full_object.model')
```

(continues on next page)

(continued from previous page)

```
# Next time, we can just load the full model object directly:
from opensoundscape.ml.cnn import load_model
model = load_model('/path/to/saved_full_object.model')
```

OpenSoundscape model objects include helper functions `.save_weights()` and `.load_weights()` which allow you to save and load platform/class independent dictionaries for increased flexibility. The weights saved and loaded by these functions are simply a dictionary of keys and numeric values, so they don't depend on the existence of particular classes in the code base. We recommend saving both the full model object (`.save()`) and the raw weights (`.save_weights()`) for models you plan to use in the future.

## 7.6.4 Clean up: delete model objects

```
[19]: from pathlib import Path
      for p in Path('.').glob('*.model'):
          p.unlink()
      for p in Path('.').glob('*.tar'):
          p.unlink()
      Path('lmin_audio.wav').unlink()
```



---

## Beginner friendly training and prediction with CNNs

---

Convolutional Neural Networks (CNNs) are a popular tool for developing automated machine learning classifiers on images or image-like samples. By converting audio into a two-dimensional frequency vs. time representation such as a spectrogram, we can generate image-like samples that can be used to train CNNs. This tutorial demonstrates the basic use of OpenSoundscape's preprocessors and cnn modules for training CNNs and making predictions using CNNs.

Under the hood, OpenSoundscape uses Pytorch for machine learning tasks. By using the class `opensoundscape.torch.models.cnn.CNN`, you can train and predict with PyTorch's powerful CNN architectures in just a few lines of code.

First, let's import some utilities.

```
[1]: # the cnn module provides classes for training/predicting with various types of CNNs
    from opensoundscape import CNN

    #other utilities and packages
    import torch
    import pandas as pd
    from pathlib import Path
    import numpy as np
    import pandas as pd
    import random
    import subprocess

    #set up plotting
    from matplotlib import pyplot as plt
    plt.rcParams['figure.figsize']=[15,5] #for large visuals
    %config InlineBackend.figure_format = 'retina'
```

Set manual seeds for pytorch and python. These ensure the training results are reproducible. You probably don't want to do this when you actually train your model, but it's useful for debugging.

```
[2]: torch.manual_seed(0)
    random.seed(0)
    np.random.seed(0)
```

## 8.1 Prepare audio data

### 8.1.1 Download labeled audio files

Training a machine learning model requires some pre-labeled data. These data, in the form of audio recordings or spectrograms, are labeled with whether or not they contain the sound of the species of interest. These data can be obtained from online databases such as Xeno-Canto.org, or by labeling one's own ARU data using a program like Cornell's Raven sound analysis software.

The Kitzes Lab has created a small labeled dataset of short clips of American Woodcock vocalizations. You have two options for obtaining the folder of data, called `woodcock_labeled_data`:

1. Run the following cell to download this small dataset. These commands require you to have `tar` installed on your computer, as they will download and unzip a compressed file in `.tar.gz` format.
2. Download a `.zip` version of the files by clicking [here](#). You will have to unzip this folder and place the unzipped folder in the same folder that this notebook is in.

**Note:** Once you have the data, you do not need to run this cell again.

```
[3]: subprocess.run(['curl','https://drive.google.com/uc?export=download&id=1Ly2M--
↳dKzpx331cfUFdVuiP96QKGJz_P','-L','-o','woodcock_labeled_data.tar.gz']) # Download
↳the data
subprocess.run(["tar","-xzf", "woodcock_labeled_data.tar.gz"]) # Unzip the downloaded
↳tar.gz file
subprocess.run(["rm", "woodcock_labeled_data.tar.gz"]) # Remove the file after its
↳contents are unzipped
```

| % Total | % Received | % Xferd | Average Speed | Time  | Time  | Time | Current |
|---------|------------|---------|---------------|-------|-------|------|---------|
|         |            |         | Dload Upload  | Total | Spent | Left | Speed   |
| 0       | 0          | 0       | 0             | 0     | 0     | 0    | 0       |
| 100     | 9499k      | 100     | 9499k         | 0     | 0     | 669k | 0       |

```
[3]: CompletedProcess(args=['rm', 'woodcock_labeled_data.tar.gz'], returncode=0)
```

### 8.1.2 Generate one-hot encoded labels

The folder contains 2s long audio clips taken from an autonomous recording unit. It also contains a file `woodcock_labels.csv` which contains the names of each file and its corresponding label information, created using a program called `Specky`.

```
[4]: #load Specky output: a table of labeled audio files
specky_table = pd.read_csv(Path("woodcock_labeled_data/woodcock_labels.csv"))
specky_table.head()
```

```
[4]:
```

|   | filename                             | woodcock | sound_type |
|---|--------------------------------------|----------|------------|
| 0 | d4c40b6066b489518f8da83af1ee4984.wav | present  | song       |
| 1 | e84a4b60a4f2d049d73162ee99a7ead8.wav | absent   | na         |
| 2 | 79678c979ebb880d5ed6d56f26ba69ff.wav | present  | song       |
| 3 | 49890077267b569e142440fa39b3041c.wav | present  | song       |
| 4 | 0c453a87185d8c7ce05c5c5ac5d525dc.wav | present  | song       |

This table must provide an accurate path to the files of interest. For this self-contained tutorial, we can use relative paths (starting with a dot and referring to files in the same folder), but you may want to use absolute paths for your training.



```
[5]: #update the paths to the audio files
specky_table.filename = ['./woodcock_labeled_data/'+f for f in specky_table.filename]
specky_table.head()
```

```
[5]:
```

|   | filename                                          | woodcock | sound_type |
|---|---------------------------------------------------|----------|------------|
| 0 | ./woodcock_labeled_data/d4c40b6066b489518f8da8... | present  | song       |
| 1 | ./woodcock_labeled_data/e84a4b60a4f2d049d73162... | absent   | na         |
| 2 | ./woodcock_labeled_data/79678c979ebb880d5ed6d5... | present  | song       |
| 3 | ./woodcock_labeled_data/49890077267b569e142440... | present  | song       |
| 4 | ./woodcock_labeled_data/0c453a87185d8c7ce05c5c... | present  | song       |

We then modify these labels, replacing present with 1 and absent with zero. Ones and zeros are the way that presences and absences are represented in a machine learning model.

```
[6]: # create a new dataframe with the filenames from the previous table as the index
labels = pd.DataFrame(index=specky_table['filename'])

#convert 'present' to 1 and 'absent' to 0
labels['woodcock']=[1 if l=='present' else 0 for l in specky_table['woodcock']]

#look at the first rows
labels.head(3)
```

```
[6]:
```

|                                                    | woodcock |
|----------------------------------------------------|----------|
| filename                                           |          |
| ./woodcock_labeled_data/d4c40b6066b489518f8da83... | 1        |
| ./woodcock_labeled_data/e84a4b60a4f2d049d73162e... | 0        |
| ./woodcock_labeled_data/79678c979ebb880d5ed6d56... | 1        |

### 8.1.3 Split into training and validation sets

We use a utility from `sklearn` to randomly divide the labeled samples into two sets. The first set, `train_df`, will be used to train the CNN, while the second set, `valid_df`, will be used to test how well the model can predict the classes of samples that it was not trained with.

During the training process, the CNN will go through all of the samples once every “epoch” for several (sometimes hundreds of) epochs. Each epoch usually consists of a “learning” step and a “validation” step. In the learning step, the CNN iterates through all of the training samples while the computer program is modifying the weights of the convolutional neural network. In the validation step, the program performs prediction on all of the validation samples and prints out metrics to assess how well the classifier generalizes to unseen data.

**Note:** using the `random_state` argument with a fixed number means that the “random” split will be exactly the same each time we run it. This is useful for reproducible results, but to get a different split each time you would not use the `random_state` argument.

```
[7]: from sklearn.model_selection import train_test_split
train_df, validation_df = train_test_split(labels, test_size=0.2, random_state=1)
```

## 8.2 Create and train a model

Now, we create a convolutional neural network model object, train it on the `train_dataset` with validation from `validation_dataset`

## 8.2.1 Set up a one-class CNN

The purpose of this model is to predict the presence or absence of a single species, so it has one class “woodcock”. Its also possible to train models to recognize multiple species - we call these “multi-class models” and each category of sounds it learns to recognize is a “class”.

The model object should be initialized with a list of class names that matches the class names in the training dataset. Here we’ll use the resnet18 architecture, a popular and powerful architecture that makes a good starting point. For more details on other CNN architectures, see the “Advanced CNN Training” tutorial.

```
[8]: # Create model object
classes = train_df.columns #in this case, there's just one class: ["woodcock"]
model = CNN('resnet18', classes=classes, sample_duration=2.0)
```

**CAVEAT:** the default audio preprocessing in this class bandpasses spectrograms to 0-11025 Hz. If your audio has a sample rate of less than 22050 Hz, the preprocessing will raise an error because the spectrogram will not contain the expected frequencies. In this case you could change the parameters of the bandpass action, or simply disable the bandpass action:

```
model.preprocessor.pipeline.bandpass.bypass=True
```

## 8.2.2 Inspect training images

Before creating a machine learning algorithm, we strongly recommend making sure the images coming out of the preprocessor look like you expect them to. Here we generate images for a few samples.

```
[9]: #helper functions to visualize processed samples
from opensoundscape.preprocess.utils import show_tensor_grid, show_tensor
from opensoundscape.ml.datasets import AudioFileDataset
```

Now, let’s check what the samples generated by our model look like

```
[10]: #pick some random samples from the training set


sample_of_4 = train_df.sample(n=4)

#generate a dataset with the samples we wish to generate and the model's preprocessor
inspection_dataset = AudioFileDataset(sample_of_4, model.preprocessor)

#generate the samples using the dataset
samples = [sample.data for sample in inspection_dataset]
labels = [list(sample.labels[sample.labels>0].index) for sample in inspection_dataset]

#display the samples
_ = show_tensor_grid(samples, 4, labels=labels)

/Users/SML161/miniconda3/envs/opso_dev/lib/python3.9/site-packages/matplotlib_inline/
↳ config.py:68: DeprecationWarning: InlineBackend._figure_format_changed is
↳ deprecated in traitlets 4.1: use @observe and @unobserve instead.
    def _figure_format_changed(self, name, old, new):
Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or
↳ [0..255] for integers).
Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or
↳ [0..255] for integers).
```




The dataset allows you to turn all augmentation off or on as desired. Inspect the unaugmented images as well:

```
[11]: #turn augmentation off for the dataset
inspection_dataset.bypass_augmentations = True

#generate the samples without augmentation
samples = [sample.data for sample in inspection_dataset]
labels = [list(sample.labels[sample.labels>0].index) for sample in inspection_dataset]

#display the samples
_ = show_tensor_grid(samples, 4, labels=labels)
```



### 8.2.3 Initialize Weights & Biases logging session

Weights and Biases (WandB) is a powerful and beautiful tool for logging and visualizing the machine learning model training process.

OpenSoundscape supports WandB integration, making it easy to visualize preprocessed samples and progress during training.

To use WandB, first create an account at <https://wandb.ai/>. The first time you use it on any computer, you'll need to run `wandb.login()` either in the command line or in Python, and enter the API key from your [settings](#) page. The “Entity” or team option allows runs and projects to be shared across members in a group, making it easy to collaborate and see progress and results of other team members’ runs.

```
[12]: try:
import wandb
wandb.login()
wandb_session = wandb.init(
    entity='kitzeslab',
    project='opensoundscape_tutorials',
    name='cnn training tutorial'
)
```

(continues on next page)

(continued from previous page)

```
except:
    wandb_session = None
```

Failed to detect the name of this notebook, you can set it manually with the WANDB\_↵NOTEBOOK\_NAME environment variable to enable code saving.

wandb: Currently logged in as: samlapp (kitzeslab). Use `wandb login --relogin` to↵↵force relogin

VBox(children=(Label(value='Waiting for wandb.init()...\r'), FloatProgress(value=0.↵01670685833333335, max=1.0)...

<IPython.core.display.HTML object>

<IPython.core.display.HTML object>

<IPython.core.display.HTML object>

<IPython.core.display.HTML object>

## 8.2.4 Train the model

Depending on the speed of your computer, training the CNN may take a few minutes.

We'll only train for 5 epochs on this small dataset as a demonstration, but you'll probably need to train for tens (or hundreds) of epochs on hundreds (or thousands) of training files to create a useful model.

Batch size refers to the number of samples that are simultaneously processed by the model. In practice, using larger batch sizes (64+) improves stability and generalizability of training, particularly for architectures (such as ResNet) that contain a 'batch norm' layer. Here we use a small batch size to keep the computational requirements for this tutorial low.

```
[13]: model.train(
    train_df=train_df,
    validation_df=validation_df,
    save_path='./binary_train/', #where to save the trained model
    epochs=10,
    batch_size=8,
    save_interval=5, #save model every 5 epochs (the best model is always saved in↵
    ↵addition)
    num_workers=0, #specify 4 if you have 4 CPU processes, eg; 0 means only the root↵
    ↵process
    wandb_session=wandb_session
)

# let wandb know that we finished training successfully
wandb.unwatch(model.network)
wandb.finish()
```

wandb: logging graph, to disable use `wandb.watch(log\_graph=False)`

```
Training Epoch 0
Epoch: 0 [batch 0/3, 0.00%]
    DistLoss: 0.694
Metrics:
Metrics:
    MAP: 0.637
```

(continues on next page)

(continued from previous page)

```
Validation.
Metrics:
    MAP: 1.000

Training Epoch 1
Epoch: 1 [batch 0/3, 0.00%]
    DistLoss: 0.619
Metrics:
Metrics:
    MAP: 0.882

Validation.
Metrics:
    MAP: 1.000

Training Epoch 2
Epoch: 2 [batch 0/3, 0.00%]
    DistLoss: 0.749
Metrics:
Metrics:
    MAP: 0.862

Validation.
Metrics:
    MAP: 1.000

Training Epoch 3
Epoch: 3 [batch 0/3, 0.00%]
    DistLoss: 0.405
Metrics:
Metrics:
    MAP: 0.929

Validation.
Metrics:
    MAP: 1.000

Training Epoch 4
Epoch: 4 [batch 0/3, 0.00%]
    DistLoss: 0.279
Metrics:
Metrics:
    MAP: 0.945

Validation.
Metrics:
    MAP: 0.967

Training Epoch 5
Epoch: 5 [batch 0/3, 0.00%]
    DistLoss: 0.619
Metrics:
Metrics:
    MAP: 0.948

Validation.
Metrics:
```

(continues on next page)

(continued from previous page)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre> MAP: 0.967  Training Epoch 6 Epoch: 6 [batch 0/3, 0.00%]     DistLoss: 0.672 Metrics: Metrics:     MAP: 0.955  Validation. Metrics:     MAP: 0.967  Training Epoch 7 Epoch: 7 [batch 0/3, 0.00%]     DistLoss: 0.160 Metrics: Metrics:     MAP: 0.957  Validation. Metrics:     MAP: 0.927  Training Epoch 8 Epoch: 8 [batch 0/3, 0.00%]     DistLoss: 0.427 Metrics: Metrics:     MAP: 0.956  Validation. Metrics:     MAP: 1.000  Training Epoch 9 Epoch: 9 [batch 0/3, 0.00%]     DistLoss: 0.178 Metrics: Metrics:     MAP: 1.000  Validation. Metrics:     MAP: 1.000  Best Model Appears at Epoch 0 with Validation score 1.000. </pre> |
| <IPython.core.display.HTML object>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <IPython.core.display.HTML object>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <IPython.core.display.HTML object>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <IPython.core.display.HTML object>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

### 8.2.5 Inspect metrics and samples using wandb

to see the wandb webpage generated by the original run of this notebook, visit [https://wandb.ai/kitzeslab/opensoundscape\\_tutorials/runs/rv67reaq?workspace=user-samlapp](https://wandb.ai/kitzeslab/opensoundscape_tutorials/runs/rv67reaq?workspace=user-samlapp)

```
[14]: wandb_session
[14]: <wandb.sdk.wandb_run.Run at 0x2aeeed070>
```

## 8.3 Visualize gradient activation maps

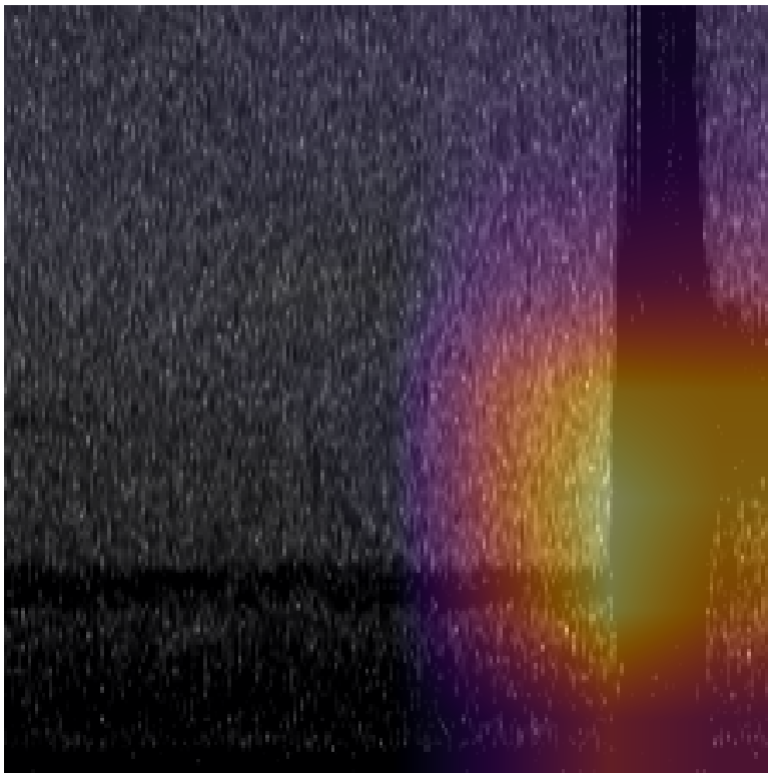
using gradient activation maps such as GradCAM, we can check what parts of the sample the model is paying attention to when it correctly (or incorrectly) labels a sample with a class.

Here, we can see that the American Woodcock vocalization in the sample is activating the network

```
[15]: samples = model.generate_cams(samples=train_df.head(1))
      samples[0].cam.plot()

Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or
↪ [0..255] for integers).
```

activation for class None



```
[15]: (<Figure size 1500x500 with 1 Axes>,
      <AxesSubplot:title={'center':'activation for class None'}>)
```


You can also use `generate_cams(...backprop=True)` and `cam.plot(...mode='backprop')` or `mode='backprop_and_activation'` to see pixel-level activations of gradients. This will work best with a model trained more thoroughly than our toy model in this example.

### 8.3.1 Plot the loss history

We can plot the loss from each epoch to check that our loss is declining. Loss should decline as the model learns, but may have ups and downs along the way.

```
[16]: plt.scatter(model.loss_hist.keys(), model.loss_hist.values())
plt.xlabel('epoch')
plt.ylabel('loss')
```

```
[16]: Text(0, 0.5, 'loss')
```



### 8.3.2 Printing and Logging outputs

We can log the outputs of the training process to a file, and/or print them. We can independently modify how much content is logged/printed with the model's attributes `model.verbose` and `model.logging_level`. Content increases from level 0 (nothing) to 1 (standard), 2, 3, etc. For instance, let's train for an epoch with lots of logged content but no printed output:

```
[17]: model.logging_level = 3 #request lots of logged content
model.log_file = './binary_train/training_log.txt' #specify a file to log output to
Path(model.log_file).parent.mkdir(parents=True, exist_ok=True) #make the folder ./
↪binary_train

model.verbose = 0 #don't print anything to the screen during training
model.train(
    train_df=train_df,
    validation_df=validation_df,
    save_path='./binary_train/', #where to save the trained model
    epochs=1,
    batch_size=8,
    save_interval=5, #save model every 5 epochs (the best model is always saved in_
↪addition)
    num_workers=0, #specify 4 if you have 4 CPU processes, eg; 0 means only the root_
↪process
)
```



## 8.4 Prediction

We haven't actually trained a useful model in 5 epochs, but we can use the trained model to demonstrate how prediction works and show several of the settings useful for prediction.

We will run prediction on two one-minute clips of field data recorded by an AudioMoth acoustic recorded. The two files are located in `woodcock_labeled_data/field_data`

### 8.4.1 Predict on the field data

To run prediction, also known as “inference”, with a CNN, we simply call model's `predict` method and pass it a list of file paths (or a dataframe with file paths in the index).

The `predict` function will internally split audio files into the appropriate length clips for prediction and generate prediction scores for each clip.

- By default, there is no overlap between these clips, but we can specify a fraction of overlap with consecutive clips with the `overlap_fraction` argument (eg, 0.5 for 50% overlap).
- Additionally, if we want to predict on audio files that are already trimmed to the same duration as the training files, we can specify `split_files_into_clips=False`.

Calling `.predict()` will return a dataframe with numeric (continuous score) predictions from the model for each sample and class (by default these are raw outputs from the model).

Let's predict on the two field recordings:

```
[18]: from glob import glob
      field_recordings = glob('./woodcock_labeled_data/field_data/*')
      field_recordings

[18]: ['./woodcock_labeled_data/field_data/60s_field_data_sample_1.wav',
      './woodcock_labeled_data/field_data/60s_field_data_sample_2.wav']

[19]: prediction_scores_df = model.predict(field_recordings)
```

The `predict` function generated a dataframe with rows for each 2-second segment of each 1-minute audio clip. Let's look at the first few rows:

```
[20]: prediction_scores_df.head()

[20]:
```

| file                                               | start_time | end_time | woodcock |
|----------------------------------------------------|------------|----------|----------|
| ./woodcock_labeled_data/field_data/60s_field_da... | 0.0        | 2.0      | 1.785296 |
|                                                    | 2.0        | 4.0      | 1.707386 |
|                                                    | 4.0        | 6.0      | 1.753464 |
|                                                    | 6.0        | 8.0      | 2.018827 |
|                                                    | 8.0        | 10.0     | 1.863675 |

Prediction can also be monitored with WandB by passing `wandb_session` to the `wandb_session` argument. This is especially useful for prediction runs that may take a long time.

### 8.4.2 Generate boolean predicted class labels (0/1) from the continuous scores

**Note:** Presence/absence predictions always have some error rates, sometimes large ones. It is not generally advisable to use these binary predictions as scientific observations without a thorough understanding of the model's false-positive and false-negative rates.

There are two different ways we might want to predict class labels which reflect the nature of the classes themselves:

**single target** means that out of a set of classes, one and only one should be chosen for each sample. For instance, if our classes were days of the week, any single should be labeled with one and only one day of the week. In opensoundscape, use the function `generate_single_target_labels()` to convert scores to predicted single target labels. For each sample, the class with the highest score will receive a label of 1 and all other classes will receive a label of 0.

**multi-target** means that a sample can have 0, 1, or more than 1 labels. For instance, if our classes were the types of flowers in a photo, any given photo might have none of the classes, one class, or multiple different classes at once. In opensoundscape, use the function `generate_multi_target_labels()` to convert scores to predicted multi target labels. For each sample and each class, the class will be labeled 1 if its score is higher than a user-specified threshold and 0 otherwise. You can choose to use a single threshold for all classes, or specify a different threshold for each class.

```
[21]: from opensoundscape.metrics import predict_single_target_labels

score_df = model.predict(field_recordings)

pred_df = predict_single_target_labels(score_df)
pred_df.head()
```

| file                                               | start_time | end_time | woodcock |
|----------------------------------------------------|------------|----------|----------|
| ./woodcock_labeled_data/field_data/60s_field_da... | 0.0        | 2.0      | 1        |
|                                                    | 2.0        | 4.0      | 1        |
|                                                    | 4.0        | 6.0      | 1        |
|                                                    | 6.0        | 8.0      | 1        |
|                                                    | 8.0        | 10.0     | 1        |

The `predict_multi_target_labels` function allows you to select a threshold. If a score exceeds that threshold, the binary prediction is 1; otherwise, it is 0. You can also specify a list of thresholds, with one for each class

```
[22]: from opensoundscape.metrics import predict_multi_target_labels

multi_target_pred_df = predict_multi_target_labels(score_df, threshold=0.99)
multi_target_pred_df.head()
```

| file                                               | start_time | end_time | woodcock |
|----------------------------------------------------|------------|----------|----------|
| ./woodcock_labeled_data/field_data/60s_field_da... | 0.0        | 2.0      | 1        |
|                                                    | 2.0        | 4.0      | 1        |
|                                                    | 4.0        | 6.0      | 1        |
|                                                    | 6.0        | 8.0      | 1        |
|                                                    | 8.0        | 10.0     | 1        |

Note that it is possible both the negative and positive classes are predicted to be present. This is because `multi_target` labeling assumes that the classes are not mutually exclusive. For a presence/absence model like the one above, `single_target` labeling is more appropriate.

### 8.4.3 Change the activation layer

We can modify the final activation layer to change the scores returned by the `predict()` function. Note that this does not impact the results of the binary predictions (described above), which are always calculated using a sigmoid transformation (for multi-target models) or softmax function (for single-target models).

Options include:

- `None`: default. Just the raw outputs of the network, which are in `(-inf, inf)`

- 'softmax': scores across all classes will sum to 1 for each sample
- 'softmax\_and\_logit': softmax the scores across all classes so they sum to 1, then apply the “logit” transformation to these scores, taking them from [0,1] back to (-inf,inf)
- 'sigmoid': transforms each score individually to [0, 1] without requiring that all scores sum to 1

In this case, since we are just looking at the output of one class, we can use the ‘sigmoid’ activation layer to put scores on the interval [0,1]

Let’s generate binary 0/1 predictions on the validation set. Since these samples are the same length as the training files, we’ll specify `split_files_into_clips=False` (we just want one prediction per file, we don’t want to divide each file into shorter clips).

```
[23]: valid_scores = model.predict(
        validation_df,
        activation_layer='sigmoid',
        split_files_into_clips=False
    )
```

Compare the softmax scores to the true labels for this dataset, side-by-side:

```
[24]: valid_scores.columns = ['pred_woodcock']
validation_df.join(valid_scores).sample(5)
```

```
[24]:
```

| filename                                           | woodcock | pred_woodcock |
|----------------------------------------------------|----------|---------------|
| ./woodcock_labeled_data/4afa902e823095e03ba23eb... | 1        | 0.989541      |
| ./woodcock_labeled_data/92647ab903049a9ee4125ab... | 1        | 0.969963      |
| ./woodcock_labeled_data/75b2f63e032dbd6d1979004... | 1        | 0.957937      |
| ./woodcock_labeled_data/882de25226ed989b31274ee... | 1        | 0.999741      |
| ./woodcock_labeled_data/ad14ac7ffa729060712b442... | 0        | 0.360559      |

We can directly compare our model’s confidence that woodcock is present with the original labels

## 8.4.4 Parallelizing prediction

Two parameters can be used to increase prediction efficiency, depending on the computational resources available:

- `num_workers`: Pytorch’s method of parallelizing across cores (CPUs) - choose 0 to predict on the root process, or >1 if you want to use more than 1 CPU process.
- `batch_size`: number of samples to predict on simultaneously. You can try increasing this by factors of two until you get a memory error, which means your batch size is too large for your system.

```
[25]: score_df = model.predict(
        validation_df,
        batch_size=8,
        num_workers=0,
    )
```

## 8.5 Multi-class models

A multi-class model can have any number of classes, and can be either

- multi-target: any number of classes can be positive for one sample
- single-target: exactly one class is positive for each sample

Models that are multi-target benefit from a modified loss function, and we have implemented a special class that is specifically designed for multi-target problems called `ResampleLoss`. We can use it as follows:

```
[26]: from opensoundscape.ml.cnn import use_resample_loss
model = CNN('resnet18', classes, 2.0, single_target=False)
use_resample_loss(model)
print("model.single_target:", model.single_target)

model.single_target: False
```

## 8.5.1 Train

Training looks the same as in one-class models.

```
[27]: model.train(
    train_df,
    validation_df,
    save_path='./multilabel_train/',
    epochs=1,
    batch_size=64,
    save_interval=100,
    num_workers=0,
)
```

```
Training Epoch 0
Epoch: 0 [batch 0/1, 0.00%]
    DistLoss: nan
Metrics:
Metrics:
    MAP: nan
```

Validation.

```
/Users/SML161/opensoundscape/opensoundscape/ml/cnn.py:700: UserWarning: Recieved_
↪empty list of predictions (or all nan)
  warnings.warn("Recieved empty list of predictions (or all nan)")
```

Best Model Appears at Epoch 0 with Validation score 0.000.

Note: since we used the same data as above, we just trained a 1 class model with “resample loss”. You should not actually use resample loss for single class models!

## 8.5.2 Predict

Prediction looks the same as demonstrated above, but make sure to think carefully:

- What `activation_layer` do you want?
- If creating boolean (0/1 or True/False) predictions for each sample and class, is my model `single-target` (use `metrics.predict_single_target_labels`) or `multi-target` (use `metrics.predict_multi_target_labels`)?

For more detail on these choices, see the sections about activation layers and boolean predictions above.

## 8.6 Save and load models

Models can be easily saved to a file and loaded at a later time. If the model was saved with OpenSoundscape version  $\geq 0.6.1$ , the entire model object will be saved - including the class, cnn architecture, loss function, and training/validation datasets. Models saved with earlier versions of OpenSoundscape do not contain all of this information and may require that you know their class and architecture (see below).

### 8.6.1 Save and load a model

OpenSoundscape saves models automatically during training:

- The model saves a copy of itself `self.save_path` to `epoch-X.model` automatically during training every `save_interval` epochs
- The model keeps the file `best.model` updated with the weights that achieve the best score on the validation dataset. By default the model is evaluated using the mean average precision (MAP) score, but you can overwrite `model.eval()` if you want to use a different metric for the best model.

You can also save the model manually at any time with `model.save(path)`

```
[28]: model1 = CNN('resnet18', classes, 2.0, single_target=False)
# Save every 2 epochs
model1.train(
    train_df,
    validation_df,
    epochs=3,
    batch_size=8,
    save_path='./binary_train/',
    save_interval=2,
    num_workers=0
)
model1.save('./binary_train/my_favorite.model')
```

```
Training Epoch 0
Epoch: 0 [batch 0/3, 0.00%]
    DistLoss: 0.835
Metrics:
Metrics:
    MAP: 0.751

Validation.
Metrics:
    MAP: 1.000

Training Epoch 1
Epoch: 1 [batch 0/3, 0.00%]
    DistLoss: 0.411
Metrics:
Metrics:
    MAP: 0.819

Validation.
Metrics:
    MAP: 1.000

Training Epoch 2
```

(continues on next page)

(continued from previous page)

```
Epoch: 2 [batch 0/3, 0.00%]
      DistLoss: 0.450
Metrics:
Metrics:
      MAP: 0.978

Validation.
Metrics:
      MAP: 1.000

Best Model Appears at Epoch 0 with Validation score 1.000.
```

## Load

Re-load a saved model with the `load_model` function:

```
[29]: from opensoundscape.ml.cnn import load_model
      model = load_model('./binary_train/best.model')
```

## Note on saving models and version compatability

Loading a model in a different version of OpenSoundscape than the version that saved the model may not work. To use a model across different versions of OpenSoundscape, you should save the model.network’s state dict using `model.save_weights(path)` as described in the “predicting with pre-trained models” tutorial. You can load weights from a saved state dict with `model.load_weights(path)`. We recommend saving both the full model object (`.save()`) and the raw weights (`.save_weights()`) for models you plan to use in the future.

Models saved with OpenSoundscape 0.4.x and 0.5.x can be loaded with `load_outdated_model` - but be sure to update the model.preprocessor after loading to match the settings used during training. See the tutorial “predicting with pre-trained models” for more details on loading models from earlier OpenSoundscape versions.

## 8.7 Predict using saved (or pre-trained) model

Using a saved or downloaded model to run predictions on audio files is as simple as

1. Loading a previously saved model
2. Generating a list of files for prediction
3. Running `model.predict()` on the preprocessor

```
[30]: # load the saved model
      model = load_model('./binary_train/best.model')

      #predict on a dataset
      scores = model.predict(field_recordings, activation_layer='sigmoid')
```

NOTE: See the tutorial “predicting with pre-trained models” for loading and using models from earlier OpenSoundscape versions

## 8.8 Continue training from saved model

Similar to predicting using a saved model, we can also continue to train a model after loading it from a saved file.

Note that `.load()` loads the entire model object, which includes optimizer parameters and learning rate parameters from the saved model, in addition to the network weights.

```
[31]: # Create architecture
model = load_model('./binary_train/best.model')

# Continue training from the checkpoint where the model was saved
model.train(train_df, validation_df, save_path='.', epochs=0)
```

```
Best Model Appears at Epoch 0 with Validation score 0.000.
```

## 8.9 Next steps

You now have seen the basic usage of training CNNs with OpenSoundscape and generating predictions.

Additional tutorials you might be interested in are: \* [Custom preprocessing](#): how to change spectrogram parameters, modify augmentation routines, etc. \* [Custom training](#): how to modify and customize model training \* [Predict with pre-trained CNNs](#): details on how to predict with pre-trained CNNs. Much of this information was covered in the tutorial above, but this tutorial also includes information about using models made with previous versions of OpenSoundscape

Finally, clean up and remove files created during this tutorial:

```
[32]: import shutil
dirs = ['./multilabel_train', './binary_train', './woodcock_labeled_data']
for d in dirs:
    try:
        shutil.rmtree(d)
    except:
        pass
```

```
[ ]:
```





---

## Preprocessing audio samples with OpenSoundscape

---

`Preprocessors` in `OpenSoundscape` perform all of the preprocessing steps from loading a file from disk, up to providing a sample to the machine learning algorithm for training or prediction. They are designed to be flexible and customizable. These classes are used internally by classes such as `opensoundscape.torch.models.cnn.CNN` when (a) training a machine learning model in `OpenSoundscape`, or (b) making predictions with a machine learning model in `OpenSoundscape`.

`Datasets` are PyTorch's way of handling a list of inputs to preprocess. In `OpenSoundscape`, there are two built-in classes (`AudioFileDataset` and `AudioSplittingDataset`) which use a `Preprocessor` to generate samples from a list of file paths.

While the `CNN` class in `OpenSoundscape` contains a default `Preprocessor`, you may want to modify or create your own `Preprocessor` depending on the specific way you wish to generate samples. `Preprocessors` are designed to be flexible and modular, so that each step of the preprocessing pipeline can be modified or removed. This notebook demonstrates:

- preparation of audio data to be used by a `preprocessor`
- how “Actions” are strung together in a `Preprocessor` to define how samples are generated
- modifying the parameters of actions
- turning Actions on and off
- modifying the order and contents of a `Preprocessor`
- use of the `SpectrogramPreprocessor` class, including examples of:
  - modifying audio and spectrogram parameters
  - changing the output image shape
  - changing the output type
  - turning augmentation on and off
  - modifying augmentation parameters
  - using the “overlay” augmentation
- writing custom preprocessors and actions

it also uses the Dataset classes to demonstrate - how to load one sample per file path - how to load long audio files as a series of shorter clips

## 9.1 Modifying the preprocessor of the CNN class

When training a CNN model in OpenSoundscape, you will create an object of the CNN class. There are two ways to modify the preprocessing:

- 1) modify the model.preprocessor directly The model contains a preprocessor object that you can modify, for instance:

```
model.preprocessor.pipeline.bandpass.bypass = True
```

- 2) overwrite the preprocessor with a new one:

```
my_preprocessor = SpectrogramPreprocessor(...) #this tutorial will help you with_
↳how to make this object
model.preprocessor = my_preprocessor
```

Note that if you want to create a preprocessor with overlay augmentation, it's easiest to use option 2 and initialize the preprocessor with an overlay\_df.

Note on augmentations: - While training, the CNN class will use all actions in the preprocessor's pipeline. - When running validation or prediction, by default, the CNN will bypass actions with action.is\_augmentation==True.

First, import some packages.

```
[1]: # Preprocessor classes are used to load, transform, and augment audio samples for use_
↳in a machine learning model
from opensoundscape.preprocess.preprocessors import SpectrogramPreprocessor
from opensoundscape.ml.datasets import AudioFileDataset, AudioSplittingDataset

# helper function for displaying a sample as an image
from opensoundscape.preprocess.utils import show_tensor, show_tensor_grid

#other utilities and packages
import torch
import pandas as pd
from pathlib import Path
import numpy as np
import random
import subprocess
import IPython.display as ipd
```

Set up plotting

```
[2]: #set up plotting
from matplotlib import pyplot as plt
plt.rcParams['figure.figsize']=[15,5] #for large visuals
%config InlineBackend.figure_format = 'retina'
```

Set manual seeds for pytorch and python. These ensure the training results are reproducible. You probably don't want to do this when you actually train your model, but it's useful for debugging.

```
[3]: torch.manual_seed(0)
      np.random.seed(0)
      random.seed(0)
```

### 9.1.1 Preparing audio data

## 9.2 Download labeled audio files

The Kitzes Lab has created a small labeled dataset of short clips of American Woodcock vocalizations. You have two options for obtaining the folder of data, called `woodcock_labeled_data`:

1. Run the following cell to download this small dataset. These commands require you to have `tar` installed on your computer, as they will download and unzip a compressed file in `.tar.gz` format.
2. Download a `.zip` version of the files by clicking [here](#). You will have to unzip this folder and place the unzipped folder in the same folder that this notebook is in.

**Note:** Once you have the data, you do not need to run this cell again.

```
[4]: subprocess.run(['curl', 'https://drive.google.com/uc?export=download&id=1ly2M--
      ↳dKzpx331cfUFdVuiP96QKGJz_P', '-L', '-o', 'woodcock_labeled_data.tar.gz']) # Download
      ↳the data
      subprocess.run(["tar", "-xzf", "woodcock_labeled_data.tar.gz"]) # Unzip the downloaded
      ↳tar.gz file
      subprocess.run(["rm", "woodcock_labeled_data.tar.gz"]) # Remove the file after its
      ↳contents are unzipped
```

| % Total | % Received | % Xferd | Average Speed | Time  | Time     | Time    | Current                          |
|---------|------------|---------|---------------|-------|----------|---------|----------------------------------|
|         |            |         | Dload Upload  | Total | Spent    | Left    | Speed                            |
| 0       | 0          | 0       | 0             | 0     | --:--:-- | 0:00:03 | --:--:-- 0                       |
| 100     | 9499k      | 100     | 9499k         | 0     | 0        | 1856k   | 0 0:00:05 0:00:05 --:--:-- 8699k |

```
[4]: CompletedProcess(args=['rm', 'woodcock_labeled_data.tar.gz'], returncode=0)
```

## 9.3 Load dataframe of files and labels

We need a dataframe with file paths in the index, so we manipulate the included `one_hot_labels.csv` slightly:

```
[5]: # load one-hot labels dataframe
      labels = pd.read_csv('./woodcock_labeled_data/one_hot_labels.csv').set_index('file')

      # prepend the folder location to the file paths
      labels.index = pd.Series(labels.index).apply(lambda f: './woodcock_labeled_data/'+f)

      #inspect
      labels.head()
```

```
[5]:
```

|                                                    | present | absent |
|----------------------------------------------------|---------|--------|
| file                                               |         |        |
| ./woodcock_labeled_data/d4c40b6066b489518f8da83... | 1       | 0      |
| ./woodcock_labeled_data/e84a4b60a4f2d049d73162e... | 0       | 1      |
| ./woodcock_labeled_data/79678c979ebb880d5ed6d56... | 1       | 0      |
| ./woodcock_labeled_data/49890077267b569e142440f... | 1       | 0      |
| ./woodcock_labeled_data/0c453a87185d8c7ce05c5c5... | 1       | 0      |

### 9.3.1 Intro to Preprocessors

Preprocessors prepare samples for use by machine learning algorithms by performing a sequential procedure on each sample, like a recipe. The procedure is defined by a **Pipeline** which contains a sequential set of steps called **Actions**. There are 3 important characteristics of Preprocessors and Actions:

- [1] A Preprocessor has a `pipeline` which defines a list of `Actions` to perform on each sample
- [2] `Actions` contain parameters that modify their behavior in the attribute `.params`. You can modify parameter values directly or use the action's `.set()` method to change parameter values.
- [3] Preprocessing can be performed with or without augmentation. The Preprocessor's `.bypass_augmentations` boolean variable will determine whether `Actions` in the pipeline with attribute `.is_augmentation==True` are performed or bypassed
- [3] `SpecPreprocessor` (the default Preprocessor class) loads audio in two distinct modes: (a) loading one sample per file, and (b) splitting files into clips, and creating a sample from each clip. You can see examples of each mode below. By default, OpenSoundscape's CNN class loads one sample per file during training and splits files into clips during prediction.

In this notebook, you will see how to edit, add, remove, and bypass `Actions` in the `pipeline` to modify the Preprocessing procedure.

The CNN class in OpenSoundscape has an internal Preprocessor object which it use to generate samples during training, validation, and prediction. We can modify or overwrite the `cnn` model's preprocessor object if we want to change how it generates samples.

The starting point for most preprocessors will be the `SpecPreprocessor` class, which loads audio files, creates spectrograms from the audio, performs various augmentations, and returns a `pytorch Tensor`.

## 9.4 Initialize preprocessor

We need to tell the preprocessor the duration (in seconds) of each sample it should create.

```
[6]: pre = SpectrogramPreprocessor(sample_duration=2.0)
```

### 9.4.1 Initialize a Dataset

A Dataset pairs a set of samples (possibly including labels) with a Preprocessor

The Dataset draws samples from it's `.df` attribute which must be a very specific dataframe:

- the index of the dataframe provides paths to audio samples
- the columns are the class names
- the values are 0 (absent/False) or 1 (present/True) for each sample and each class.

For example, we've set up the labels dataframe with files as the index and classes as the columns, so we can use it to make an instance of `SpecPreprocessor`:

```
[7]: dataset = AudioFileDataset(labels, pre)
```

## 9.5 Generate a sample from a Dataset

We can ask a dataset for a specific sample using its numeric index, like accessing an element of a list. Each sample is a dictionary with two keys: ‘X’, the Tensor of the sample, and ‘y’, the Tensor of labels of the sample. The shape of ‘X’ is [channels, height, width] and the shape of ‘y’ is [number of classes].

```
[8]: dataset[0] #loads and preprocesses the sample at row 0 of dataset.df
[8]: AudioSample(source=woodcock_labeled_data/d4c40b6066b489518f8da83af1ee4984.wav, start_
      ↳time=0.0,end_time=2.0, labels=present      1
      absent      0
      Name: ./woodcock_labeled_data/d4c40b6066b489518f8da83af1ee4984.wav, dtype: int64)
```

Using a helper function, we can easily visualize a set of samples on a grid. We *highly* recommend inspecting your preprocessed samples in this way before training or predicting with a machine learning model. By inspecting the samples, you can confirm that your labeled data is reasonable and that the preprocessing is representing your samples in a reasonable way.

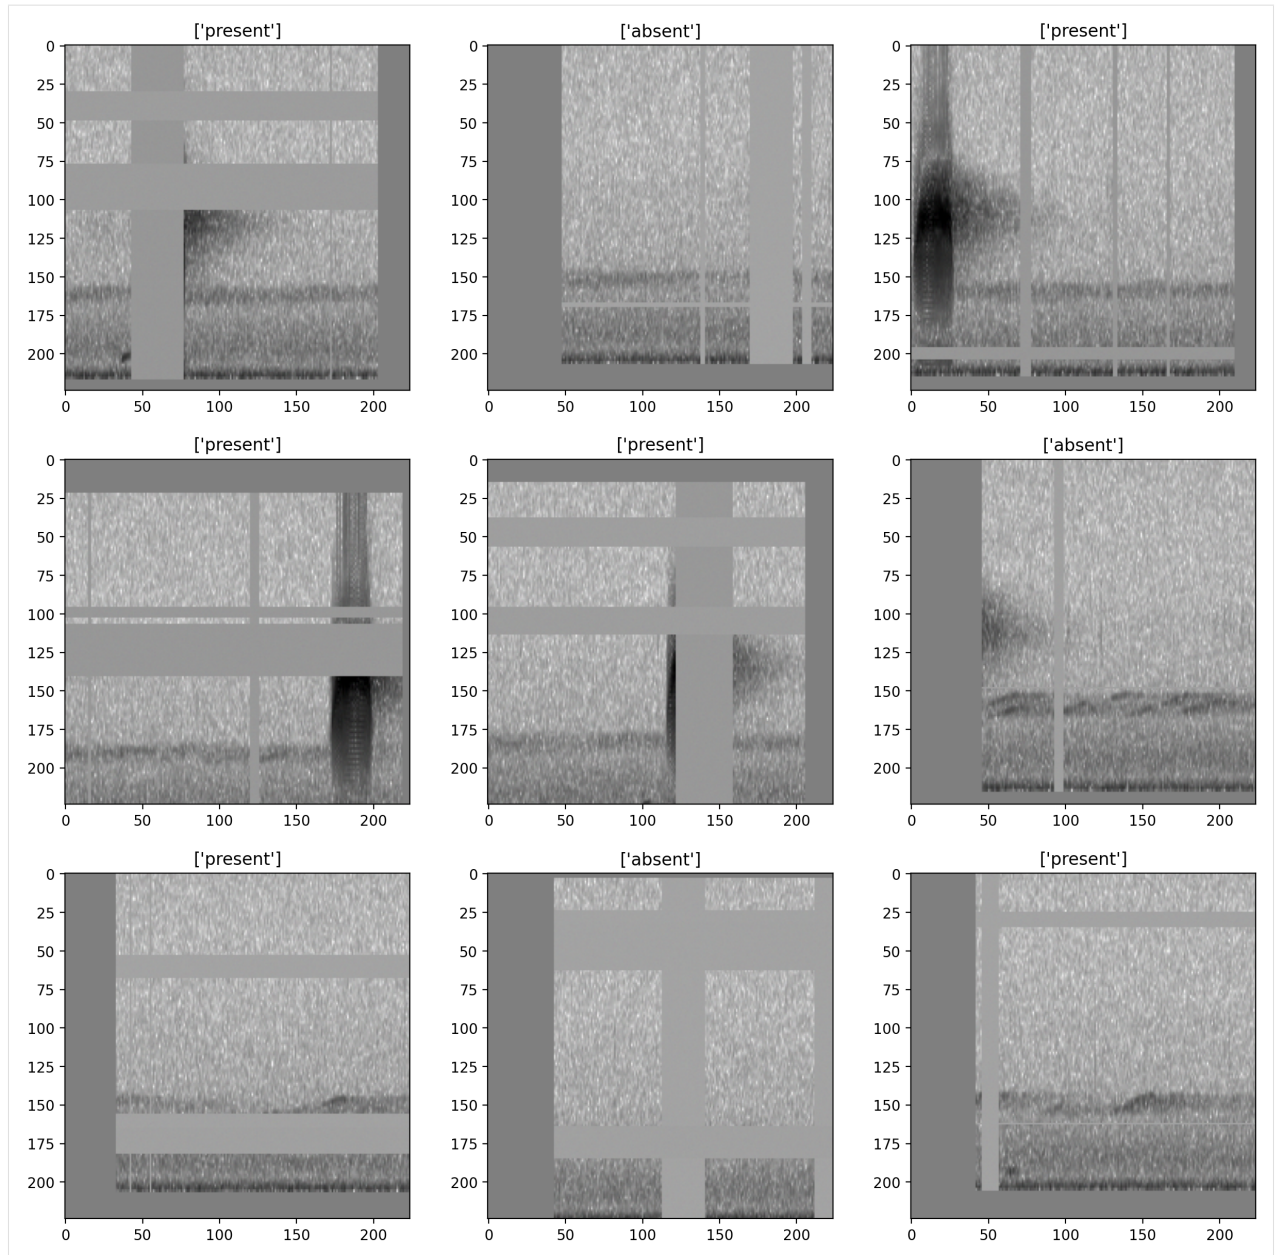
```
[9]: from opensoundscape.preprocess.utils import show_tensor_grid

pre = SpectrogramPreprocessor(sample_duration=2.0)
dataset = AudioFileDataset(labels,pre)

tensors = [dataset[i].data for i in range(9)]
sample_labels = [list(dataset[i].labels[dataset[i].labels>0].index) for i in range(9)]

_ = show_tensor_grid(tensors,3,labels=sample_labels)

/Users/SML161/miniconda3/envs/opso_dev/lib/python3.9/site-packages/matplotlib_inline/
↳config.py:68: DeprecationWarning: InlineBackend._figure_format_changed is
↳deprecated in traitlets 4.1: use @observe and @unobserve instead.
    def _figure_format_changed(self, name, old, new):
Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or
↳[0..255] for integers).
Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or
↳[0..255] for integers).
Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or
↳[0..255] for integers).
```




let's repeat the exercise of inspecting preprocessed samples, this time without augmentation

```
[10]: dataset.bypass_augmentations = True

tensors = [dataset[i].data for i in range(9)]
sample_labels = [list(dataset[i].labels[dataset[i].labels>0].index) for i in range(9)]

_ = show_tensor_grid(tensors, 3, labels=sample_labels)
```





during machine learning tasks with Pytorch, a `DataLoader` is often used on top of a `Dataset` to “batch” samples - that is, to prepare multiple samples at once. A batch returned by a `DataLoader` will have an extra leading dimension for both ‘X’ and ‘y’; for instance, a `batch_size` of 16 would produce ‘X’ with shape [16, 3, 224, 224] for 3-channel 224x224 tensors and ‘y’ with shape [16, 5] if the labels contain 5 classes (columns). OpenSoundscape uses `DataLoaders` internally to create batches of samples during CNN training and prediction.

## 9.6 Subset samples from a Dataset

Preprocessors allow you to select a subset of samples using `sample()` and `head()` methods (like Pandas `DataFrames`).

(note that these methods subset files from the index, they do not subset individual clips from files)

```
[11]: len(dataset)
```

```
[11]: 29
```

Select the first 10 samples (non-random)

```
[12]: len(dataset.head(10))
```

```
[12]: 10
```

Randomly select an absolute number of samples

```
[13]: len(dataset.sample(n=10))
```

```
[13]: 10
```

Randomly select a fraction of samples

```
[14]: len(dataset.sample(frac=0.5))
```

```
[14]: 14
```

## 9.6.1 Loading many fixed-duration samples from longer audio files

When preprocessing should result in many fixed-length samples per input file, instead of one sample per file, we use `AudioSplittingDataset` instead of `AudioFileDataset`. This dataset can be customized with parameters for:

- fractional overlap between consecutive samples
- how to handle remaining audio at the end of a file (if it is shorter than the desired sample duration)

The `CNN.predict()` function uses `AudioSplittingDataset` internally, so that the user can specify long audio file paths and get back predictions on fixed-length clips. (If one sample per file is desired, you can pass the argument `split_files_into_clips=False` to `CNN.predict`)

Here's an example of how to use `AudioSplittingDataset` to create several samples from a long audio file:

(Note that you never have to manually create `AudioSplittingDataset` or `AudioFileDataset` objects to train and predict with the CNN class, they are created internally.)


```
[15]: prediction_df = pd.DataFrame(index=['./woodcock_labeled_data/field_data/60s_field_
↳data_sample_1.wav'])
```

```
[17]: pre = SpectrogramPreprocessor(sample_duration=2.0)
      splitting_dataset = AudioSplittingDataset(prediction_df, pre, overlap_fraction=0.5)
      splitting_dataset.bypass_augmentations = True

      #get the first 9 samples and plot them
      tensors = [splitting_dataset[i].data for i in range(9)]

      _ = show_tensor_grid(tensors, 3)
```





### 9.6.2 Pipelines and actions

Each Preprocessor class has a `pipeline` which is an ordered set of operations that are performed on each sample, in the form of a `pandas.Series` object. Each element of the series is an object of class `Action` (or one of its subclasses) and represents a transformation on the sample.

## 9.7 About Pipelines

The preprocessor's Pipeline is the ordered list of Actions that the preprocessor performs on each sample.

- The Pipeline is stored in the `preprocessor.pipeline` attribute.

- You can modify the contents or order of Preprocessor Actions by overwriting the preprocessor's `.pipeline` attribute. When you modify this attribute, you must provide `pd.Series` with elements `name:Action`, where each Action is an instance of a class that sub-classes `opensoundscape.preprocess.BaseAction`.

Let's inspect the current pipeline of our preprocessor.

```
[18]: # inspect the current pipeline (ordered sequence of Actions to take)
preprocessor = SpectrogramPreprocessor(sample_duration=2)
preprocessor.pipeline

[18]: load_audio          Action calling <bound method Audio.from_file o...
random_trim_audio      Augmentation Action calling <function trim_aud...
trim_audio             Action calling <function trim_audio at 0x17eb8...
to_spec                Action calling <bound method Spectrogram.from...
bandpass               Action calling <function Spectrogram.bandpass ...
to_tensor              Action
time_mask              Augmentation Action calling <function time_mas...
frequency_mask         Augmentation Action calling <function frequenc...
add_noise              Augmentation Action calling <function tensor_a...
rescale                Action calling <function scale_tensor at 0x17e...
random_affine          Augmentation Action calling <function torch_ra...
dtype: object
```

## 9.8 About actions

Each element of the preprocessor's pipeline (a `pd.Series`) contains a name (string) and an action (Action)

- Each Action takes a sample (and its labels), performs some transformation to them, and returns the sample (and its labels).
- You can generate an Action based on a function like this : `Action(fn=my_function, other parameters...)`. The function you pass (`my_function` in this case) must expect the sample as the first argument. It can then take additional parameters. For instance, if we define the function:

```
def multiply(x,n):
    return x*n
```

then we can create an action to multiply by 3 with `action=Action(fn=multiply,n=3)`

- Any customizable parameters for performing the Action are stored in a dictionary, `.params`. These parameters can be modified directly (e.g. `Action.params.param1=value1`) or using the Action's `.set()` method (e.g. `action.set(param=value, param2=value2, ...)`)
- You can bypass an action in a pipeline by changing `Action.bypass` to `True`
- You can declare whether an Action is an augmentation (should not be performed if `bypass_augmentation=True`) using its `.is_augmentation` boolean attribute

### 9.8.1 Modifying Actions

## 9.9 View default parameters for an Action

the `.params` attribute of an Action is a pandas Series containing parameters that can be modified

```
[19]: #since the pipeline is a series, we can access elements like pipeline.to_spec as well,
      ↪as pipeline['to_spec']
      preprocessor.pipeline.to_spec.params

[19]: window_type          hann
      window_samples      None
      window_length_sec   None
      overlap_samples      None
      overlap_fraction    None
      fft_size            None
      decibel_limits      (-100, -20)
      dB_scale            True
      scaling              spectrum
      dtype: object
```

## 9.10 Modify Action parameters

we can modify parameters with the Action's `.set()` method:

```
[20]: preprocessor.pipeline.to_spec.set(dB_scale=False)
```

or by accessing the parameter directly (params is a pandas Series)

```
[21]: preprocessor.pipeline.to_spec.params.window_samples = 512
      preprocessor.pipeline.to_spec.params['overlap_fraction'] = 0.75

      preprocessor.pipeline.to_spec.params

[21]: window_type          hann
      window_samples      512
      window_length_sec   None
      overlap_samples      None
      overlap_fraction    0.75
      fft_size            None
      decibel_limits      (-100, -20)
      dB_scale            False
      scaling              spectrum
      dtype: object
```

## 9.11 Bypass actions

Actions can be bypassed by changing the attribute `.bypass=True`. A bypassed action is never performed regardless of the `.perform_augmentations` attribute.

```
[22]: preprocessor = SpectrogramPreprocessor(sample_duration=2.0)

      #turn off augmentations other than noise
      preprocessor.pipeline.add_noise.bypass=True
      preprocessor.pipeline.time_mask.bypass=True
      preprocessor.pipeline.frequency_mask.bypass=True

      #printing the pipeline will show which actions are bypassed
      preprocessor.pipeline
```

```
[22]: load_audio      Action calling <bound method Audio.from_file o...
      random_trim_audio  Augmentation Action calling <function trim_aud...
      trim_audio        Action calling <function trim_audio at 0x17eb8...
      to_spec           Action calling <bound method Spectrogram.from_...
      bandpass          Action calling <function Spectrogram.bandpass ...
      to_tensor          Action
      time_mask          ## Bypassed ## Augmentation Action calling <fu...
      frequency_mask     ## Bypassed ## Augmentation Action calling <fu...
      add_noise          ## Bypassed ## Augmentation Action calling <fu...
      rescale            Action calling <function scale_tensor at 0x17e...
      random_affine       Augmentation Action calling <function torch_ra...
      dtype: object
```


create a Dataset with this preprocessor and our label dataframe

```
[24]: dataset = AudioFileDataset(labels,preprocessor)


print('random affine off')
preprocessor.pipeline.random_affine.bypass = True
show_tensor(dataset[0].data,invert=True,transform_from_zero_centered=True)
plt.show()

print('random affine on')
preprocessor.pipeline.random_affine.bypass = False
show_tensor(dataset[0].data,invert=True,transform_from_zero_centered=True)

random affine off
```



random affine on



To view whether an individual Action in a pipeline is on or off, inspect its `bypass` attribute:

```
[25]: # The AudioLoader Action that is still on
preprocessor.pipeline.load_audio.bypass
```

```
[25]: False
```

```
[26]: # The frequency_mask Action that we turned off
preprocessor.pipeline.frequency_mask.bypass
```

```
[26]: True
```

### 9.11.1 Modifying the pipeline

Sometimes, you may want to change the order or composition of the Preprocessor's pipeline. You can simply overwrite the `.pipeline` attribute, as long as it is a pandas Series of names:Actions

## 9.12 Example: return Spectrogram instead of Tensor

Here's an example where we replace the pipeline with one that just loads audio and converts it to a Spectrogram, returning a Spectrogram instead of a Tensor:

```
[30]: #initialize a preprocessor
preprocessor = SpectrogramPreprocessor(2.0)
print('original pipeline:')
[print(p) for p in pre.pipeline]


#overwrite the pipeline with a slice of the original pipeline
print('\nnew pipeline:')
preprocessor.pipeline = preprocessor.pipeline[0:4]
[print(p) for p in preprocessor.pipeline]

print('\nWe now have a preprocessor that returns Spectrograms instead of Tensors:')
dataset = AudioFileDataset(labels,preprocessor)
print(f"Type of returned sample: {type(dataset[0].data)}")
dataset[0].data.plot()
```

original pipeline:  
Action calling <bound method Audio.from\_file of <class 'opensoundscape.audio.Audio'>>  
Augmentation Action calling <function trim\_audio at 0x17eb853a0>  
Action calling <function trim\_audio at 0x17eb853a0>  
Action calling <bound method Spectrogram.from\_audio of <class 'opensoundscape.  
↳spectrogram.Spectrogram'>>  
Action calling <function Spectrogram.bandpass at 0x17da20820>  
Action  
Augmentation Action calling <function time\_mask at 0x17eb94a60>  
Augmentation Action calling <function frequency\_mask at 0x17eb94af0>  
Augmentation Action calling <function tensor\_add\_noise at 0x17eb94b80>  
Action calling <function scale\_tensor at 0x17eb949d0>  
Augmentation Action calling <function torch\_random\_affine at 0x17eb948b0>

new pipeline:  
Action calling <bound method Audio.from\_file of <class 'opensoundscape.audio.Audio'>>  
Augmentation Action calling <function trim\_audio at 0x17eb853a0>  
Action calling <function trim\_audio at 0x17eb853a0>  
Action calling <bound method Spectrogram.from\_audio of <class 'opensoundscape.  
↳spectrogram.Spectrogram'>>

We now have a preprocessor that returns Spectrograms instead of Tensors:  
Type of returned sample: <class 'opensoundscape.spectrogram.Spectrogram'>



### 9.12.1 Analyzing/debugging the pipeline

In order to debug the Preprocessor's pipeline you can utilize the `trace` argument to save and review the output of action step in the pipeline as part of the sample information returned by the preprocessor.

```
[31]: # initialize a preprocessor
preprocessor = SpectrogramPreprocessor(2.0)
pre.pipeline
```

```
[31]: load_audio          Action calling <bound method Audio.from_file o...
random_trim_audio      Augmentation Action calling <function trim_aud...
trim_audio             Action calling <function trim_audio at 0x17eb8...
to_spec                Action calling <bound method Spectrogram.from_...
bandpass               Action calling <function Spectrogram.bandpass ...
to_tensor              Action
time_mask              Augmentation Action calling <function time_mas...
frequency_mask         Augmentation Action calling <function frequenc...
add_noise              Augmentation Action calling <function tensor_a...
rescale                Action calling <function scale_tensor at 0x17e...
random_affine          Augmentation Action calling <function torch_ra...
dtype: object
```

```
[35]: # pass a sample through the preprocessor's pipeline
sample = preprocessor.forward(labels.index.values[0], trace=True)
sample.trace
```

```
/Users/SML161/opensoundscape/opensoundscape/preprocess/preprocessors.py:127:
↳ DeprecationWarning: The default dtype for empty Series will be 'object' instead of
↳ 'float64' in a future version. Specify a dtype explicitly to silence this warning.
sample.trace = pd.Series(index=self.pipeline.index)
```

```
[35]: load_audio          <Audio(samples=(88200,), sample_rate=44100)>
random_trim_audio      <Audio(samples=(88200,), sample_rate=44100)>
trim_audio             <Audio(samples=(88200,), sample_rate=44100)>
to_spec                <Spectrogram(spectrogram=(257, 343), frequenci...
bandpass               <Spectrogram(spectrogram=(129, 343), frequenci...
to_tensor              [[tensor(0.3032), tensor(0.3257), tensor(0.30...
time_mask              [[tensor(0.3032), tensor(0.3257), tensor(0.30...
frequency_mask         [[tensor(0.3032), tensor(0.3257), tensor(0.30...
add_noise              [[tensor(0.3031), tensor(0.3306), tensor(0.29...
rescale                [[tensor(-0.3937), tensor(-0.3389), tensor(-0...
random_affine          [[tensor(0.), tensor(0.), tensor(0.), tensor(...
dtype: object
```


## 9.13 analyse the output at steps of interest

```
[38]: # Initial audio
sample.trace['load_audio']
```


```
[38]: <Audio(samples=(88200,), sample_rate=44100)>
```

```
[39]: # Initial spectrogram
sample.trace["to_spec"].plot()
```





```
[40]: # After applyin frequency mask
show_tensor(sample.trace["frequency_mask"], invert=True, transform_from_zero_
         ↳ centered=True)
```



### 9.13.1 adding the preprocessor to a CNN

You can always overwrite the preprocessor of a CNN model object with a new one:



```
my_preprocessor = SpectrogramPreprocessor(...)
...
model.preprocessor = my_preprocessor
```

**WARNING:** Be careful! If your new preprocessor has a different sample duration (eg 3 seconds instead of 2) or shape (eg [100,100,3] instead of [224,224,1]), these new values will also take effect when using the CNN.

### 9.13.2 Customizing preprocessing to achieve better machine learning outcomes

The right choice of preprocessing depends heavily on the characteristics of the sounds you wish to study. The best way to tune preprocessing parameters is to visually inspect samples created by your preprocessing procedure and tweak parameters to achieve visual clarity of the sounds of interest in your samples. We find these heuristics to be a good starting point:

- The duration of a sample should be approximately 2-5x the duration of the target sound. For instance, a very short nocturnal flight call lasting 0.1 seconds might be best visualized with a 0.3 second sample\_duration. Meanwhile, a 10-second bout of ruffed grouse drumming might deserve a 20 second sample\_duration.
- The frequency range of a sample should be wider than the target sound, but not by more than 1 order of magnitude. For instance, sounds that are low-pitched will be more clearly visualized when bandpassing a spectrogram to the low frequencies. If you use a 0-10,000 Hz spectrogram for a 500 Hz target sound, your target sound will only occupy a small fraction of your sample.
- Spectrogram parameters should be matched to the temporal or spectral features of the target sound. Modify the Spectrogram’s window\_samples to achieve high enough time resolution (lower value of window\_samples) or frequency resolution (higher value of window\_samples) to see features of your target sound clearly on the resulting sample. For example, a rapid trill with a pulse repetition rate of 50 Hz will only be distinctive on a spectrogram if the Spectrogram windows are less than  $1 / (50 * 2) = 0.01$  seconds in duration. On the other hand, visualizing a distinctive harmonic “ladder” structure of a nasal sound might require long spectrogram windows which will increase frequency resolution.

Augmentations are Actions that are only performed during training, not during prediction. These actions manipulate the sample in some randomized way, so that each time the same sample is provided to the model as training data, the actual values of the sample are different. This prevents over-training of a model on a training set and effectively increases the size of a training dataset. In general, you can expect that a basic set of augmentations (such as those included by default in the SpecPreprocessor and CNN classes) will be necessary to train a useful machine learning model. In particular, “overlay” augmentations which blend together multiple samples often increase the generalizability (transferability) of a model. You might choose to use audio from your target system (for instance, field recordings at your study site) to make the training data look more similar to the data that the model will be applied to.

Below are various examples of how to modify parameters of the Actions to achieve different preprocessing outcomes.

## 9.14 Modify the sample rate

Resample all loaded audio to a specified rate during the load\_audio action

```
[41]: pre = SpectrogramPreprocessor(sample_duration=2)
pre.pipeline.load_audio.set(sample_rate=24000)
```

## 9.15 Modify spectrogram window length and overlap

(see `Spectrogram.from_audio()` for detailed documentation)


```
[43]: dataset = AudioFileDataset(labels, SpectrogramPreprocessor(sample_duration=2))
dataset.bypass_augmentations=True

print('default parameters:')
show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
plt.show()


print('high time resolution, low frequency resolution:')
dataset.preprocessor.pipeline.to_spec.set(window_samples=64)

show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)

default parameters:
```



```
high time resolution, low frequency resolution:
```



## 9.16 Bandpass spectrograms

Trim spectrograms to a specified frequency range:


```
[44]: dataset = AudioFileDataset(labels, SpectrogramPreprocessor(2.0))

print('default parameters:')
show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)

print('bandpassed to 2-4 kHz:')
dataset.preprocessor.pipeline.bandpass.set(min_f=2000, max_f=4000)
show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
```

Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or [0..255] for integers).

```
default parameters:
bandpassed to 2-4 kHz:
```



## 9.17 Change the output shape

Change the shape of the output sample - note that the shape argument expects (height, width, channels)

```
[48]: dataset.preprocessor.out_shape
```


```
[48]: (224, 224, 3)
```

```
[50]: dataset = AudioFileDataset(labels, SpectrogramPreprocessor(2.0))
```

```
dataset.preprocessor.out_shape = [100, 500, 3]
```

```
show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
```

Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or [0..255] for integers).




## 9.18 Turn all augmentation on or off

augmentation is controlled by the `preprocessor.bypass_augmentation` boolean (aka True/False) variable. By default, augmentations are performed. A CNN will internally manipulate this attribute to perform augmentations during training but not during validation or prediction.


```
[52]: dataset = AudioFileDataset(labels, SpectrogramPreprocessor(2.0))

dataset.bypass_augmentations = True
show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
```



```
[53]: dataset.bypass_augmentations = False
show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
```

Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or [0..255] for integers).



## 9.19 Modify augmentation parameters


SpectrogramPreprocessor includes several augmentations with customizable parameters. Here we provide a couple of illustrative examples - see any action's documentation for details on how to use its parameters.

```
[55]: #initialize a preprocessor
preprocessor = SpectrogramPreprocessor(2.0)

#turn off augmentations other than overlay
preprocessor.pipeline.random_affine.bypass=True
preprocessor.pipeline.time_mask.bypass=True
preprocessor.pipeline.add_noise.bypass=True

# allow up to 20 horizontal masks, each spanning up to 0.1x the height of the image.
preprocessor.pipeline.frequency_mask.set(max_width = 0.03, max_masks=20)

#preprocess the same sample 4 times
dataset = AudioFileDataset(labels,preprocessor)
tensors = [dataset[0].data for i in range(4)]
fig = show_tensor_grid(tensors,2)
plt.show()
```




turn off frequency mask and turn on gaussian noise

```
[56]: dataset.preprocessor.pipeline.add_noise.bypass = False
dataset.preprocessor.pipeline.frequency_mask.bypass = True

# increase the intensity of gaussian noise added to the image
dataset.preprocessor.pipeline.add_noise.set(std=0.2)
show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
```

Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or [0..255] for integers).





## 9.20 remove an action by its name

```
[57]: preprocessor.remove_action('add_noise')
preprocessor.pipeline
```

```
[57]: load_audio          Action calling <bound method Audio.from_file o...
random_trim_audio       Augmentation Action calling <function trim_aud...
trim_audio              Action calling <function trim_audio at 0x17eb8...
to_spec                 Action calling <bound method Spectrogram.from_...
bandpass                Action calling <function Spectrogram.bandpass ...
to_tensor                Action
time_mask               ## Bypassed ## Augmentation Action calling <fu...
frequency_mask          ## Bypassed ## Augmentation Action calling <fu...
rescale                 Action calling <function scale_tensor at 0x17e...
random_affine           ## Bypassed ## Augmentation Action calling <fu...
dtype: object
```

## 9.21 add an action at a specific position

specify the action in the pipeline you want to insert before or after



```
[58]: from opensoundscape.preprocess.actions import Action, tensor_add_noise

preprocessor.insert_action(
    action_index='add_noise_NEW', #give it a name
    action=Action(tensor_add_noise, std=0.01), #the action object
    after_key='to_tensor', #where to put it (can also use before_key=...)
)
```

```
[59]: preprocessor.pipeline
```

```
[59]: load_audio          Action calling <bound method Audio.from_file o...
random_trim_audio      Augmentation Action calling <function trim_aud...
trim_audio             Action calling <function trim_audio at 0x17eb8...
to_spec                Action calling <bound method Spectrogram.from_...
bandpass               Action calling <function Spectrogram.bandpass ...
to_tensor              Action
add_noise_NEW          Action calling <function tensor_add_noise at 0...
time_mask              ## Bypassed ## Augmentation Action calling <fu...
frequency_mask         ## Bypassed ## Augmentation Action calling <fu...
rescale                Action calling <function scale_tensor at 0x17e...
random_affine          ## Bypassed ## Augmentation Action calling <fu...
dtype: object
```

it will complain if you use a non-unique index

```
[60]: from opensoundscape.preprocess.actions import Action, tensor_add_noise

try:
    preprocessor.insert_action(
        action_index='add_noise_NEW', #using the same name as a current action will_
        ↪lead to an AssertionError
        action=Action(tensor_add_noise, std=0.01), #the action object
        after_key='to_tensor', #where to put it (can also use before_key=...)
    )
except AssertionError:
    print("raised AssertionError, as expected")

raised AssertionError, as expected
```

## 9.22 Overlay augmentation

Overlay is a powerful Action that allows additional samples to be overlayed or blended with the original sample.

The additional samples are chosen from the `overlay_df` that is provided to the preprocessor when it is initialized. The index of the `overlay_df` must be paths to audio files. The dataframe can be simply an index containing audio files with no other columns, or it can have the same columns as the sample dataframe for the preprocessor.

Samples for overlays are chosen based on their class labels, according to the parameter `overlay_class`:

- None - Randomly select any file from `overlay_df`
- "different" - Select a random file from `overlay_df` containing none of the classes this file contains
- specific class name - always choose files from this class

By default, the overlay Action does **not** change the labels of the sample it modifies. However, if you wish to add the labels from overlayed samples to the original sample's labels, you can set `update_labels=True` (see example below).


```
[61]: #initialize a preprocessor and provide a dataframe with samples to use as overlays
preprocessor = SpectrogramPreprocessor(2.0, overlay_df=labels)

#remove augmentations other than overlay
for name in ['random_affine', 'time_mask', 'frequency_mask', 'add_noise']:
    preprocessor.remove_action(name)
```

Let's change `overlay_weight` to


To demonstrate this, let's show what happens if we overlay samples from the "negative" class, resulting in the final sample having a higher or lower signal-to-noise ratio. By default, the overlay Action chooses a random file from the overlay dataframe. Instead, choose a sample from the class called "present" using the `overlay_class` parameter.

```
[63]: preprocessor.pipeline.overlay.set(overlay_class='present')
tensors = []
overlay_weights = [0.01, 0.4, 0.6, 0.8]
for w in overlay_weights:
    preprocessor.pipeline.overlay.set(overlay_weight=w)
    dataset = AudioFileDataset(labels, preprocessor)
    np.random.seed(0) #get the same overlay every time
    tensors.append(dataset[2].data)
_ = show_tensor_grid(tensors, 2, labels=overlay_weights)
```




As demonstrated above, you can choose a specific class to choose samples from. Here, instead, we choose samples from the “absent” class.

```
[64]: dataset.preprocessor.pipeline.overlay.set(
        overlay_class='absent',
        overlay_weight=0.4
    )
show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
```



By default, or by specifying `overlay_class=None`, the overlay sample is chosen randomly from the `overlay_df` with no restrictions.

```
[65]: dataset.preprocessor.pipeline.overlay.set(overlay_class=None)
      show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
```




The 'different' option for `overlay_class` chooses a sample to overlay that has non-overlapping labels with the original sample.

In the case of this example, this has the same effect as drawing samples from the "negative" class as demonstrated above. In multi-class examples, this would draw from any of the samples not labeled with the class(es) of the original sample.

We'll again use `overlay_weight=0.8` to exaggerate the importance of the overlaid sample (80%) compared to the original sample (20%).

```
[66]: dataset.preprocessor.pipeline.overlay.set(update_labels=False, overlay_class='different
      ↪', overlay_weight=0.8)
      show_tensor(dataset[0].data, invert=True, transform_from_zero_centered=True)
```



By default, the overlay Action does **not** change the labels of the sample it modifies.

For instance, if the overlaid sample has labels [1,0] and the original sample has labels [0,1], the default behavior will return a sample with labels [0,1] not [1,1].

If you wish to add the labels from overlaid samples to the original sample's labels, you can set `update_labels=True`.

```
[76]: print('default: labels do not update')
dataset.preprocessor.pipeline.overlay.set(update_labels=False, overlay_class='different
→')
print(f"\t resulting labels: \n {dataset[1].labels}")

print('Using update_labels=True')
dataset.preprocessor.pipeline.overlay.set(update_labels=True, overlay_class='different
→')
print(f"\t resulting labels: \n {dataset[1].labels}")
```

```
default: labels do not update
    resulting labels:
present      0
absent       1
Name: ./woodcock_labeled_data/e84a4b60a4f2d049d73162ee99a7ead8.wav, dtype: int64
Using update_labels=True
    resulting labels:
present      1
absent       1
```

(continues on next page)

(continued from previous page)

```
Name: ./woodcock_labeled_data/e84a4b60a4f2d049d73162ee99a7ead8.wav, dtype: int64
```

This example is a single-target problem: the two classes represent “woodcock absent” and “woodcock present.” Because the labels are mutually exclusive, labels [1,1] do not make sense. So, for this single-target problem, we would **not** want to use `update_labels=True`, and it would probably make most sense to only overlay absent recordings, e.g., `overlay_class='absent'`.

### 9.22.1 Creating a new Preprocessor class

If you have a specific augmentation routine you want to perform, you may want to create your own Preprocessor class rather than modifying an existing one.

Your subclass might add a different set of Actions, define a different pipeline, or even override the `__getitem__` method of `BasePreprocessor`.

Here’s an example of a customized preprocessor that subclasses `AudioToSpectrogramPreprocessor` and creates a pipeline that depends on the `magic_parameter` input.


```
[77]: from opensoundscape.preprocess.actions import Action, tensor_add_noise
class MyPreprocessor(SpectrogramPreprocessor):
    """Child of AudioToSpectrogramPreprocessor with weird augmentation routine"""

    def __init__(
        self,
        magic_parameter,
        sample_duration,
        out_shape=[224, 224, 1],
    ):

        super(MyPreprocessor, self).__init__(
            sample_duration=sample_duration,
            out_shape=out_shape,
        )


        for i in range(magic_parameter):
            action = Action(tensor_add_noise, std=0.1*magic_parameter)
            self.insert_action(f'noise_{i}', action)

[79]: dataset = AudioFileDataset(labels, MyPreprocessor(sample_duration=2.0, magic_
    ↪parameter=1))
show_tensor(dataset[0].data, invert=False)
```



```
[80]: dataset = AudioFileDataset(labels, MyPreprocessor(sample_duration=2.0, magic_
    ↪parameter=4))
show_tensor(dataset[0].data, invert=False)
```





### 9.22.2 Defining new Actions

You can usually define a new action simply by passing a method to `Action()`. However, you can also write a subclass of `Action` for more advanced use cases.

```
[91]: from opensoundscape.preprocess.actions import Action

class AudioGate(BaseAction):
    """Replace audio samples below a threshold with 0, but only if label[0]==1

    Audio in, Audio out

    Args:
        threshold: sample values below this will become 0
    """
    def __init__(self, threshold):
        super(AudioGate, self).__init__()
        self.params['threshold'] = threshold

    def go(self, sample, **kwargs):
        threshold = 0.2
        if sample.labels[0] == 1:
            samples = np.array([0 if np.abs(s) < threshold else s for s in audio.
↪ samples])
```

(continues on next page)

(continued from previous page)

```
sample.data = Audio(samples, audio.sample_rate)
```

Test it out:

```
[96]: from opensoundscape.audio import Audio
      from opensoundscape.sample import AudioSample

gate_action = AudioGate(threshold=0.2)


print('histogram of samples')
audio = Audio.from_file('./woodcock_labeled_data/01c5d0c90bd4652f308fd9c73feb1bf5.wav
↪')
sample = AudioSample(source=audio, labels={0:1,1:0})

_ = plt.hist(audio.samples, bins=100)
plt.semilogy()
plt.show()


print('histogram of samples after audio gate')
gate_action.go(sample)
_ = plt.hist(sample.data.samples, bins=100)
plt.semilogy()
plt.show()

print('histogram of samples after audio gate, when labels[0]==0')
print('histogram of samples')
audio = Audio.from_file('./woodcock_labeled_data/01c5d0c90bd4652f308fd9c73feb1bf5.wav
↪')
sample = AudioSample(source=audio, labels={0:0,1:1})
gate_action.go(sample)
_ = plt.hist(sample.data.samples, bins=100)
plt.semilogy()
```

histogram of samples




histogram of samples after audio gate



histogram of samples after audio gate, when labels[0]==0  
 histogram of samples

[96]: []



### 9.22.3 Add custom Action to a preprocessor

For instance, if you want to use your custom Action while training a `cnn`, you can add it to the `cnn.preprocessor`'s pipeline.

In this example, we put the custom `AudioGate` action before the `to_spec` action.

```
[97]: gate_action = AudioGate(threshold=0.2)
preprocessor.insert_action(
    action_index='custom_audio_gate', #give it a name
    action=gate_action,
    before_key='to_spec', #where to put it (can also use before_key=...)
)
```

```
[98]: preprocessor.pipeline
```

```
[98]: load_audio          Action calling <bound method Audio.from_file o...
random_trim_audio       Augmentation Action calling <function trim_aud...
trim_audio              Action calling <function trim_audio at 0x17eb8...
custom_audio_gate       Action
```

(continues on next page)

(continued from previous page)

```
to_spec          Action calling <bound method Spectrogram.from_...
bandpass         Action calling <function Spectrogram.bandpass ...
to_tensor                Action
overlay            Augmentation Action calling <function overlay ...
rescale           Action calling <function scale_tensor at 0x17e...
dtype: object
```

Clean up files created during this tutorial:

```
[99]: import shutil
      shutil.rmtree('./woodcock_labeled_data')
```

---

## Advanced CNN training

---

This notebook demonstrates how to use classes from `opensoundscape.torch.models.cnn` and architectures created using `opensoundscape.torch.architectures.cnn_architectures` to

- choose between single-target and multi-target model behavior
- modify learning rates, learning rate decay schedule, and regularization
- choose from various CNN architectures
- train a multi-target model with a special loss function
- use strategic sampling for imbalanced training data
- customize preprocessing: train on spectrograms with a bandpassed frequency range

Rather than demonstrating their effects on training (model training is slow!), most examples in this notebook either don't train the model or "train" it for 0 epochs for the purpose of demonstration.

For introductory demos (basic training, prediction, saving/loading models), see the “Beginner-friendly training and prediction with CNNs” tutorial ([cnn.ipynb](#)).

```
[24]: from opensoundscape.preprocess import preprocessors
      from opensoundscape.ml import cnn, cnn_architectures

      import torch
      import pandas as pd
      from pathlib import Path
      import numpy as np
      import random
      import subprocess

      from matplotlib import pyplot as plt
      plt.rcParams['figure.figsize']=[15,5] #for big visuals
      %config InlineBackend.figure_format = 'retina'
```

## 10.1 Prepare audio data

### 10.1.1 Download labeled audio files

The Kitzes Lab has created a small labeled dataset of short clips of American Woodcock vocalizations. You have two options for obtaining the folder of data, called `woodcock_labeled_data`:

1. Run the following cell to download this small dataset. These commands require you to have `tar` installed on your computer, as they will download and unzip a compressed file in `.tar.gz` format.
2. Download a `.zip` version of the files by clicking [here](#). You will have to unzip this folder and place the unzipped folder in the same folder that this notebook is in.

If you already have these files, you can skip or comment out this cell

```
[25]: subprocess.run(['curl', 'https://drive.google.com/uc?export=download&id=1Ly2M--
↳dKzpx331cfUFdVuiP96QKGJz_P', '-L', '-o', 'woodcock_labeled_data.tar.gz']) # Download
↳the data
subprocess.run(["tar", "-xzf", "woodcock_labeled_data.tar.gz"]) # Unzip the downloaded
↳tar.gz file
subprocess.run(["rm", "woodcock_labeled_data.tar.gz"]) # Remove the file after its
↳contents are unzipped
```

| % Total | % Received | % Xferd | Average Speed | Time  | Time     | Time    | Current                          |
|---------|------------|---------|---------------|-------|----------|---------|----------------------------------|
|         |            |         | Dload Upload  | Total | Spent    | Left    | Speed                            |
| 0       | 0          | 0       | 0             | 0     | --:--:-- | 0:00:03 | --:--:-- 0                       |
| 100     | 9499k      | 100     | 9499k         | 0     | 0        | 2223k   | 0 0:00:04 0:00:04 --:--:-- 18.7M |

```
[25]: CompletedProcess(args=['rm', 'woodcock_labeled_data.tar.gz'], returncode=0)
```

### 10.1.2 Load dataframe of files and one-hot labels

We need a dataframe with file paths in the index, so we manipulate the included `one_hot_labels.csv` slightly

See the “Basic training and prediction with CNNs” tutorial for more details.

```
[26]: # load one-hot labels dataframe
labels = pd.read_csv('./woodcock_labeled_data/one_hot_labels.csv').set_index('file')[[
↳'present']]
# prepend the folder location to the file paths
labels.index = pd.Series(labels.index).apply(lambda f: './woodcock_labeled_data/'+f)
#create class list
classes = labels.columns
#inspect
labels.head()
```

```
[26]:
```

|                                                    | present |
|----------------------------------------------------|---------|
| file                                               |         |
| ./woodcock_labeled_data/d4c40b6066b489518f8da83... | 1       |
| ./woodcock_labeled_data/e84a4b60a4f2d049d73162e... | 0       |
| ./woodcock_labeled_data/79678c979ebb880d5ed6d56... | 1       |
| ./woodcock_labeled_data/49890077267b569e142440f... | 1       |
| ./woodcock_labeled_data/0c453a87185d8c7ce05c5c5... | 1       |

### 10.1.3 Split into train and validation sets

Randomly split the data into training data and validation data.

```
[27]: from sklearn.model_selection import train_test_split
train_df, valid_df = train_test_split(labels, test_size=0.2, random_state=0)
print(f"created train_df (len {len(train_df)}) and valid_df (len {len(valid_df)})")

created train_df (len 23) and valid_df (len 6)
```

## 10.2 Creating a model

We initialize a model object by specifying the architecture, a list of classes, and the duration of individual samples in seconds

```
[28]: arch = cnn_architectures.resnet50(num_classes=len(classes))
model = cnn.CNN(arch, classes, sample_duration=2.0)
```

Alternatively, we can specify the name of an architecture as a string (see Cnn Architectures below for details or use `cnn_architectures.list_architectures()` for options)

```
[29]: model = cnn.CNN('resnet18', classes, 2.0)
```

### 10.2.1 Single-target versus multi-target

One important decision is whether your model is single-target (exactly one label per sample) or multi-target (any number of labels per sample, including 0). Single-target models have a softmax activation layer which forces the sum of all class scores to be 1.0. By default, models are created as multi-target, but you can set `single_target=True` either when creating the object or afterwards.

```
[30]: #change the model to be single_target
model.single_target = True

#or specify single_target when you create the object
model = cnn.CNN(arch, classes, 2.0)
```

## 10.3 Model training parameters

We can modify various parameters about model training, including:

- The learning rate
- The learning rate schedule
- Weight decay for regularization

Let's take a peek at the current parameters, stored in a dictionary.

```
[31]: model.optimizer_params
[31]: {'lr': 0.01, 'momentum': 0.9, 'weight_decay': 0.0005}
```

### 10.3.1 Learning rates

The learning rate determines how much the model's weights change every time it calculates the loss function.

Faster learning rates improve the speed of training and help the model leave local minima as it learns to classify, but if the learning rate is too fast, the model may not successfully fit the data or its fitting might be unstable.

Often after training a model for a while at a relatively high learning rate (think 0.01), we might want to “fine tune” the model by training for a few epochs with a lower learning rate. Let’s set a low learning rate for fine tuning:

```
[32]: model.optimizer_params['lr']=0.001
```

### 10.3.2 Separate learning rates for feature and classifier blocks

For ResNet architectures, we can modify the learning rates for the feature extraction and classification blocks of the network separately. For example, we can specify a relatively fast learning rate for classifier and slower one for features, if we think the features from a pre-trained model are close to optimal but we have a different set of classes than the pre-trained model. We first use a helper function to separate the feature and classifier parameters, then specify parameters for each:

```
[33]: from opensoundscape.ml.cnn import separate_resnet_feat_clf
```

```
[34]: r18_model = cnn.CNN('resnet18', classes, 2.0)
      print(r18_model.optimizer_params)

      separate_resnet_feat_clf(r18_model) #in place operation!

      #now we can specify separate parameters for the 'feature' and 'classifier' portions_
      ↪ of the network
      r18_model.optimizer_params['feature']['lr'] = 0.001
      r18_model.optimizer_params['classifier']['lr'] = 0.01

      r18_model.optimizer_params

      {'lr': 0.01, 'momentum': 0.9, 'weight_decay': 0.0005}

[34]: {'feature': {'lr': 0.001, 'momentum': 0.9, 'weight_decay': 0.0005},
      'classifier': {'lr': 0.01, 'momentum': 0.9, 'weight_decay': 0.0005}}
```

### 10.3.3 Learning rate schedule

It’s often helpful to decrease the learning rate over the course of training. By reducing the amount that the model’s weights are updated as time goes on, this causes the learning to gradually switch from coarsely searching across possible weights to fine-tuning the weights.

By default, the learning rates are multiplied by 0.7 (the learning rate “cooling factor”) once every 10 epochs (the learning rate “update interval”).

Let’s modify that for a very fast training schedule, where we want to multiply the learning rates by 0.1 every epoch.

```
[35]: model.lr_cooling_factor = 0.1
      model.lr_update_interval = 1
```

### 10.3.4 Regularization weight decay

Pytorch optimizers perform [L2 regularization](#), giving the optimizer an incentive for the model to have small weights rather than large weights. The goal of this regularization is to reduce overfitting to the training data by reducing the complexity of the model.



Depending on how much emphasis you want to place on the L2 regularization, you can change the weight decay parameter. By default, it is 0.0005. The higher the value for the “weight decay” parameter, the more the model training algorithm prioritizes smaller weights.

```
[36]: model.optimizer_params['weight_decay']=0.001
```

## 10.4 Selecting CNN architectures

The `opensoundscape.torch.architectures.cnn_architectures` [https://github.com/kitzeslab/opensoundscape/blob/master/opensoundscape/torch/architectures/cnn\\_architectures.py](https://github.com/kitzeslab/opensoundscape/blob/master/opensoundscape/torch/architectures/cnn_architectures.py) module provides functions to create several common CNN architectures. These architectures are built in to pytorch, but the OpenSoundscape module helps us out by reshaping the final layer to match the number of classes we have.

You could also create a custom architecture by subclassing an existing pytorch model or writing one from scratch (the minimum requirement is that it subclasses `torch.nn.Module` - it should at least have `.forward()` and `.backward()` methods).

In general, we can create any pytorch model architecture and pass it to the `architecture` argument when creating a model in opensoundscape. We can choose whether to use pre-trained (ImageNet) weights or start from scratch (`weights=None` for random weights). For instance, lets create an alexnet architecture with random weights:

```
[37]: my_arch = cnn_architectures.alexnet(num_classes=len(classes), weights=None)
```

For convenience, we can also initialize a model object by providing the name of an architecture as a string, rather than the architecture object. For a list of valid architecture names, use `cnn_architectures.list_architectures()`. Note that these will use default architecture parameters, including using pre-trained ImageNet weights. If you don’t want to use pre-trained weights, follow the method above of creating the architecture and passing it to the initialization of CNN.

```
[38]: print(cnn_architectures.list_architectures())

['resnet18', 'resnet34', 'resnet50', 'resnet101', 'resnet152', 'alexnet', 'vgg11_bn',
↳ 'squeezenet1_0', 'densenet121', 'inception_v3', 'efficientnet_b0', 'efficientnet_b4
↳ ', 'efficientnet_widese_b0', 'efficientnet_widese_b4']
```

```
[39]: model = cnn.CNN(architecture='resnet18', classes=classes, sample_duration=2.0)
```

### 10.4.1 Pretrained weights

In OpenSoundscape, by default, model architectures are initialized with weights pretrained on the [ImageNet](#) image database. It takes some time for pytorch to download these weights from an online repository the first time an instance of a particular architecture is created with pretrained weights - pytorch will do this automatically and only once.

Using pretrained weights often speeds up training significantly, as the representation learned from ImageNet is a good start at beginning to interpret spectrograms, even though they are not true “pictures.”

If you prefer not to use pre-trained weights, or if you don’t have an internet connection, you can specify `weights` argument to `None`, when creating an architecture:

```
[40]: arch = cnn_architectures.alexnet(num_classes=10, weights=None)
```

## 10.4.2 Freezing the feature extractor

Convolutional Neural Networks can be thought of as having two parts: a **feature extractor** which learns how to represent/“see” the input data, and a **classifier** which takes those representations and transforms them into predictions about the class identity of each sample.

You can freeze the feature extractor if you only want to train the final classification layer of the network but not modify any other weights. This could be useful for applying pre-trained classifiers to new data, i.e. “transfer learning”. To do so, set the `freeze_feature_extractor` argument to `True` when you create an architecture.

```
[41]: # See "InceptionV3 architecture" section below for more information
arch = cnn_architectures.resnet50(num_classes=10, freeze_feature_extractor=True,
    ↪weights=None)
```

## 10.4.3 InceptionV3 class

The Inception architecture requires slightly different training and preprocessing from the ResNet architectures and the other architectures implemented in OpenSoundscape (see below), because:

- 1) the input image shape must be 299x299, and
- 2) Inception’s forward pass gives output + auxiliary output instead of a single output

The InceptionV3 class in `cnn` handles the necessary modifications in training and prediction for you, so use that instead of CNN:

```
[42]: from opensoundscape.ml.cnn import InceptionV3

#generate an Inception model
model = InceptionV3(classes=classes, weights=None, sample_duration=2)

#train and validate for 1 epoch
#note that Inception will complain if batch_size=1
model.train(train_df, valid_df, epochs=1, batch_size=4)

#predict
scores = model.predict(valid_df)

/Users/SML161/miniconda3/envs/opso_dev/lib/python3.9/site-packages/torchvision/models/
↪inception.py:43: FutureWarning: The default weight initialization of inception_v3_
↪will be changed in future releases of torchvision. If you wish to keep the old_
↪behavior (which leads to long initialization times due to scipy/scipy#11299),_
↪please set init_weights=True.
    warnings.warn(

Training Epoch 0
Epoch: 0 [batch 0/6, 0.00%]
    DistLoss: 0.863
Metrics:
Metrics:
    MAP: 0.914

Validation.
Metrics:
    MAP: 1.000

Best Model Appears at Epoch 0 with Validation score 1.000.
```

### 10.4.4 Changing the architecture of an existing model (not recommended)

The architecture is stored in the model object's `.network` attribute. We can access parameters of the network or even replace it entirely.

Note that replacing the architecture will completely remove anything the model has “learned” since the learned weights are a part of the architecture.

```
[43]: #initialize the AlexNet architecture
new_arch = cnn_architectures.densenet121(num_classes=2, weights=None)

# replace the alexnet architecture with the densenet architecture
model.network = new_arch
```

## 10.5 Multi-target training with ResampleLoss

Training multi-target models (a.k.a. multi-label: there can be any number of positive labels on each sample) is challenging and can benefit from using a modified loss function. OpenSoundscape provides a loss function designed for training multi-target models. We recommend using this loss function when training multi-target models. You can add it to a class with an in-place helper function:

```
[44]: from opensoundscape.ml.cnn import use_resample_loss
```

```
[45]: model = cnn.CNN('resnet18', classes, 2.0)
use_resample_loss(model)
print(model.loss_cls)

#use as normal...
#model.train(...)
#model.predict(...)
```

<class 'opensoundscape.ml.loss.ResampleLoss'>

## 10.6 Training and predicting with custom preprocessors

The preprocessing tutorial gives in-depth descriptions of how to customize your preprocessing pipeline.

Here, we'll just give a quick example of tweaking the preprocessing pipeline: providing the CNN with a bandpassed spectrogram object instead of the full frequency range.

It's good practice to create the validation from the training dataset (after any modifications are made), so that they perform the same preprocessing. You may or may not want to use augmentation on the validation dataset.

### 10.6.1 Example: Training on bandpassed spectrograms

```
[47]: model = cnn.CNN('resnet18', classes, 2.0)

# change the min and max frequencies for the spectrogram bandpass action
model.preprocessor.pipeline.bandpass.set(min_f=3000, max_f=5000)

# inspect a few preprocessed samples (see basic CNN training and prediction tutorial_
↪ for details)
```

(continues on next page)


(continued from previous page)

```
from opensoundscape.preprocess.utils import show_tensor_grid
from opensoundscape.ml.datasets import AudioFileDataset
sample_of_4 = train_df.sample(n=4)
inspection_dataset = AudioFileDataset(sample_of_4, model.preprocessor)
samples = [sample.data for sample in inspection_dataset]
labels = [list(sample.labels[sample.labels>0].index) for sample in inspection_dataset]
_ = show_tensor_grid(samples,4,labels=labels)

# now we can train and validate on the bandpassed spectrograms
model.train(train_df, valid_df, epochs=0)
```

```
/Users/SML161/miniconda3/envs/opso_dev/lib/python3.9/site-packages/matplotlib_inline/
↳ config.py:68: DeprecationWarning: InlineBackend._figure_format_changed is
↳ deprecated in traitlets 4.1: use @observe and @unobserve instead.
def _figure_format_changed(self, name, old, new):
Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or
↳ [0..255] for integers).
Clipping input data to the valid range for imshow with RGB data ([0..1] for floats or
↳ [0..255] for integers).
```

Best Model Appears at Epoch 0 with Validation score 0.000.



If we predict using this model for prediction, it will use the same preprocessor settings, bandpassing the prediction samples in the same way as the training samples.

## 10.6.2 clean up

remove files

```
[48]: import shutil
shutil.rmtree('./woodcock_labeled_data')

for p in Path('.').glob('*.model'):
    p.unlink()
```

---

## RIBBIT Pulse Rate model demonstration

---

RIBBIT (Repeat-Interval Based Bioacoustic Identification Tool) is a tool for detecting vocalizations that have a repeating structure.

This tool is useful for detecting vocalizations of frogs, toads, and other animals that produce vocalizations with a periodic structure. In this notebook, we demonstrate how to select model parameters for the Great Plains Toad, then run the model on data to detect vocalizations.

This work is described in:

- 2021 paper, “Automated detection of frog calls and choruses by pulse repetition rate”
- 2020 poster, “Automatic Detection of Pulsed Vocalizations”

RIBBIT is also available as an R package.

This notebook demonstrates how to use the RIBBIT tool implemented in opensoundscape as `opensoundscape.ribbit.ribbit()`

For help installing OpenSoundscape, see the [documentation](#)

### 11.1 Import packages

```
[1]: # suppress warnings
import warnings
warnings.simplefilter('ignore')

#import packages
import numpy as np
from glob import glob
import pandas as pd
from matplotlib import pyplot as plt
import subprocess

#local imports from opensoundscape
```

(continues on next page)

(continued from previous page)

```
from opensoundscape.audio import Audio
from opensoundscape.spectrogram import Spectrogram
from opensoundscape.ribbit import ribbit
```

```
[2]: # display settings
plt.rcParams['figure.figsize']=[15,8]
pd.set_option('display.precision', 2)
%config InlineBackend.figure_format = 'retina'
```

## 11.2 Download example audio

First, let's download some example audio to work with.

You can run the cell below, **OR** visit this link to download the data (whichever you find easier):

<https://pitt.box.com/shared/static/0xclmulc4gy0obewtzbzyfnszczwgr9we.zip>

If you download using the link above, first un-zip the folder (double-click on mac or right-click -> extract all on Windows). Then, move the `great_plains_toad_dataset` folder to the same location on your computer as this notebook. Then you can skip this cell:

```
[3]: #download files from box.com to the current directory
subprocess.run(['curl', 'https://drive.google.com/uc?export=download&
↳id=1Zjeyp6Zjdk0zugK9jhFBIuWHE__a0SKu', '-L', '-o', 'great_plains_toad_dataset.tar.gz
↳']) # Download the data
subprocess.run(["tar", "-xzf", "great_plains_toad_dataset.tar.gz"]) # Unzip the_
↳downloaded tar.gz file
subprocess.run(["rm", "great_plains_toad_dataset.tar.gz"]) # Remove the file after_
↳its contents are unzipped
```

| % Total | % Received | % Xferd | Average Speed | Time  | Time     | Time    | Current                          |
|---------|------------|---------|---------------|-------|----------|---------|----------------------------------|
|         |            |         | Dload Upload  | Total | Spent    | Left    | Speed                            |
| 0       | 0          | 0       | 0             | 0     | --:--:-- | 0:00:04 | --:--:-- 0                       |
| 100     | 11.6M      | 100     | 11.6M         | 0     | 0        | 1857k   | 0 0:00:06 0:00:06 --:--:-- 8476k |

```
[3]: CompletedProcess(args=['rm', 'great_plains_toad_dataset.tar.gz'], returncode=0)
```

now, you should have a folder in the same location as this notebook called `great_plains_toad_dataset`

if you had trouble accessing the data, you can try using your own audio files - just put them in a folder called `great_plains_toad_dataset` in the same location as this notebook, and this notebook will load whatever is in that folder

### 11.2.1 Load an audio file and create a spectrogram

```
[4]: audio_path = np.sort(glob('./great_plains_toad_dataset/*'))[0]

#load the audio file into an OpenSoundscape Audio object
audio = Audio.from_file(audio_path)

#trim the audio to the time from 0-3 seconds for a closer look
audio = audio.trim(0,3)


#create a Spectrogram object
spectrogram = Spectrogram.from_audio(audio)
```

## 11.2.2 Show the Great Plains Toad spectrogram as an image

A spectrogram is a visual representation of audio with frequency on the vertical axis, time on the horizontal axis, and intensity represented by the color of the pixels

```
[5]: spectrogram.plot()
```

```
/Users/SML161/miniconda3/envs/opso_dev/lib/python3.9/site-packages/matplotlib_inline/
↳ config.py:68: DeprecationWarning: InlineBackend._figure_format_changed is_
↳ deprecated in traitlets 4.1: use @observe and @unobserve instead.
def _figure_format_changed(self, name, old, new):
```



## 11.3 Select model parameters

RIBBIT requires the user to select a set of parameters that describe the target vocalization. Here is some detailed advice on how to use these parameters.

**Signal Band:** The signal band is the frequency range where RIBBIT looks for the target species. Based on the spectrogram above, we can see that the Great Plains Toad vocalization has the strongest energy around 2000-2500 Hz, so we will specify `signal_band = [2000, 2500]`. It is best to pick a narrow signal band if possible, so that the model focuses on a specific part of the spectrogram and has less potential to include erroneous sounds.

**Noise Bands:** Optionally, users can specify other frequency ranges called noise bands. Sounds in the `noise_bands` are *subtracted* from the `signal_band`. Noise bands help the model filter out erroneous sounds from the recordings, which could include confusion species, background noise, and popping/clicking of the microphone due to rain, wind, or digital errors. It's usually good to include one noise band for very low frequencies – this specifically eliminates popping and clicking from being registered as a vocalization. It's also good to specify noise bands that target confusion species. Another approach is to specify two narrow `noise_bands` that are directly above and below the `signal_band`.

**Pulse Rate Range:** This parameters specifies the minimum and maximum pulse rate (the number of pulses per second, also known as pulse repetition rate) RIBBIT should look for to find the focal species. Looking at the spectrogram above, we can see that the pulse rate of this Great Plains Toad vocalization is about 15 pulses per second. By looking

at other vocalizations in different environmental conditions, we notice that the pulse rate can be as slow as 10 pulses per second or as fast as 20. So, we choose `pulse_rate_range = [10, 20]` meaning that RIBBIT should look for pulses no slower than 10 pulses per second and no faster than 20 pulses per second.

**Clip Duration:** This parameter tells the algorithm how many seconds of audio to analyze at one time. Generally, you should choose a `clip_duration` that is ~2x longer than the target species vocalization, or a little bit longer. For very slowly pulsing vocalizations, choose a longer window so that at least 5 pulses can occur in one window (0.5 pulses per second -> 10 second window). Typical values for `clip_duration` are 0.3 to 10 seconds. Here, because the The Great Plains Toad has a vocalization that continues on for many seconds (or minutes!), we chose a 2-second window which will include plenty of pulses.

- we can also set `clip_overlap` if we want overlapping clips. For instance, a `clip_duration` of 2 with `clip_overlap` of 1 results in 50% overlap of each consecutive clip. This can help avoid sounds being split up across two clips, and therefore not being detected.
- `final_clip` determines what should be done when there is less than `clip_duration` audio remaining at the end of an audio file. We'll just use `final_clip=None` to discard any remaining audio that doesn't make a complete clip.

**Plot:** We can choose to show the power spectrum of pulse repetition rate for each window by setting `plot=True`. The default is not to show these plots (`plot=False`).

```
[6]: # minimum and maximum rate of pulsing (pulses per second) to search for
pulse_rate_range = [8,15]

# look for a vocalization in the range of 1000-2000 Hz
signal_band = [1800,2400]

# subtract the amplitude signal from these frequency ranges
noise_bands = [ [0,1000], [3000,3200] ]

#divides the signal into segments this many seconds long, analyzes each independently
clip_duration = 2 #seconds
clip_overlap = 0 #seconds

#if True, it will show the power spectrum plot for each audio segment
show_plots = True
```

## 11.4 Search for pulsing vocalizations with `ribbit()`

This function takes the parameters we chose above as arguments, performs the analysis, and returns two arrays: - **scores:** the pulse rate score for each window - **times:** the start time in seconds of each window

The scores output by the function may be very low or very high. They do not represent a “confidence” or “probability” from 0 to 1. Instead, the relative values of scores on a set of files should be considered: when RIBBIT detects the target species, the scores will be significantly higher than when the species is not detected.

The file `gpt0.wav` has a Great Plains Toad vocalizing only at the beginning. Let's analyze the file with RIBBIT and look at the scores versus time.

```
[7]: #get the audio file path
audio_path = np.sort(glob('./great_plains_toad_dataset/*'))[0]

#make the spectrogram
spec = Spectrogram.from_audio(audio.from_file(audio_path))
```

(continues on next page)



(continued from previous page)

```


#run RIBBIT
score_df = ribbit(
    spec,
    pulse_rate_range=pulse_rate_range,
    signal_band=signal_band,
    clip_duration=clip_duration,
    noise_bands=noise_bands,
    plot=False
)

#show the spectrogram
print('spectrogram of 10 second file with Great Plains Toad at the beginning')
spec.plot()

# plot the score vs time of each window
plt.scatter(score_df['start_time'],score_df['score'])
plt.xlabel('window start time (sec)')
plt.ylabel('RIBBIT score')
plt.title('RIBBIT scores for 10 second file with Great Plains Toad at the beginning')

```


spectrogram of 10 second file with Great Plains Toad at the beginning



```

[7]: Text(0.5, 1.0, 'RIBBIT scores for 10 second file with Great Plains Toad at the_
↪beginning')

```



as we hoped, RIBBIT outputs a high score during the vocalization (the window from 0-2 seconds) and a low score when the frog is not vocalizing

## 11.5 Analyzing a set of files

```
[8]: # set up a dataframe for storing files' scores and labels
df = pd.DataFrame(index = glob('./great_plains_toad_dataset/*'), columns=['score',
    ↳ 'label'])

# label is 1 if the file contains a Great Plains Toad vocalization, and 0 if it does_
    ↳ not
df['label'] = [1 if 'gpt' in f else 0 for f in df.index]

# calculate RIBBIT scores
for path in df.index:

    #make the spectrogram
    spec = Spectrogram.from_audio(audio.from_file(path))

    #run RIBBIT
    score_df = ribbit(
        spec,
        pulse_rate_range=[8,20],
        signal_band=[1900,2400],
        clip_duration=clip_duration,
        noise_bands=[[0,1500],[2500,3500]],
        plot=False)

    # use the maximum RIBBIT score from any window as the score for this file
    # multiply the score by 10,000 to make it easier to read
```

(continues on next page)

(continued from previous page)

```
df.at[path, 'score'] = max(score_df['score']) * 10000

print("Files sorted by score, from highest to lowest:")
df.sort_values(by='score', ascending=False)
```

Files sorted by score, from highest to lowest:

```
[8]:
```

|                                           | score  | label |
|-------------------------------------------|--------|-------|
| ./great_plains_toad_dataset/gpt0.wav      | 107.65 | 1     |
| ./great_plains_toad_dataset/gpt3.wav      | 29.31  | 1     |
| ./great_plains_toad_dataset/gpt2.wav      | 16.69  | 1     |
| ./great_plains_toad_dataset/gpt1.wav      | 10.13  | 1     |
| ./great_plains_toad_dataset/negative9.wav | 3.04   | 0     |
| ./great_plains_toad_dataset/negative8.wav | 0.89   | 0     |
| ./great_plains_toad_dataset/negative4.wav | 0.76   | 0     |
| ./great_plains_toad_dataset/negative2.wav | 0.65   | 0     |
| ./great_plains_toad_dataset/negative1.wav | 0.3    | 0     |
| ./great_plains_toad_dataset/negative3.wav | 0.3    | 0     |
| ./great_plains_toad_dataset/gpt4.wav      | 0.12   | 1     |
| ./great_plains_toad_dataset/negative6.wav | 0.06   | 0     |
| ./great_plains_toad_dataset/pops2.wav     | 0.0    | 0     |
| ./great_plains_toad_dataset/pops1.wav     | 0.0    | 0     |
| ./great_plains_toad_dataset/negative5.wav | 0.0    | 0     |
| ./great_plains_toad_dataset/silent.wav    | 0.0    | 0     |
| ./great_plains_toad_dataset/negative7.wav | 0.0    | 0     |
| ./great_plains_toad_dataset/water.wav     | 0.0    | 0     |

So, how good is RIBBIT at finding the Great Plains Toad?

We can see that the scores for all of the files with Great Plains Toad (gpt) score above 10 except `gpt4.wav` (which contains only a very quiet and distant vocalization). All files that do not contain the Great Plains Toad score less than 3.5. So, RIBBIT is doing a good job separating Great Plains Toads vocalizations from other sounds!

Notably, noisy files like `pops1.wav` score low even though they have lots of periodic energy - our `noise_bands` successfully rejected these files. Without using `noise_bands`, files like these would receive very high scores. Also, some birds in “negatives” files that have periodic calls around the same pulse rate as the Great Plains Toad received low scores. This is also a result of choosing a tight `signal_band` and strategic `noise_bands`. You can try adjusting or eliminating these bands to see their effect on the audio.

(HINT: eliminating the `noise_bands` will result in high scores for the “pops” files)

## 11.6 Run RIBBIT on multiple species simultaneously

If you want to search for multiple species, its best to combine the analysis into one function - that way you only have to load each audio file (and make it's spectrogram) one time, instead of once for each species. (If you have thousands of audio files, this might be a big time saver.)

This code gives a quick exmaple of how you could use a pre-made dataframe (could load it in from a spreadsheet, for instance) of parameters for a set of species to run RIBBIT on all of them.

Note that this example assumes you are using the same spectrogram settings for each species - this might not be the case in practice, if some species require high time-resolution spectrograms and others require high frequency-resolution spectrograms.

```
[9]: #we'll create a dataframe here, but you could also load it from a spreadsheet
species_df = pd.DataFrame(columns=['pulse_rate_range', 'signal_band', 'clip_duration',
↪ 'noise_bands'])
```

(continues on next page)

(continued from previous page)

```
species_df.loc['great_plains_toad']={
    'pulse_rate_range':[8,20],
    'signal_band':[1900,2400],
    'clip_duration':2.0,
    'noise_bands':[[0,1500],[2500,3500]]
}
```

```
species_df.loc['bird_series']={
    'pulse_rate_range':[8,11],
    'signal_band':[5000,6500],
    'clip_duration':2.0,
    'noise_bands':[[0,4000]]
}
```

```
species_df
```

```
[9]:
```

|                   | pulse_rate_range | signal_band  | clip_duration \ |
|-------------------|------------------|--------------|-----------------|
| great_plains_toad | [8, 20]          | [1900, 2400] | 2.0             |
| bird_series       | [8, 11]          | [5000, 6500] | 2.0             |

|                   | noise_bands               |
|-------------------|---------------------------|
| great_plains_toad | [[0, 1500], [2500, 3500]] |
| bird_series       | [[0, 4000]]               |

now let's analyze each audio file for each species.

We'll save the results in a table that has a column for each species.

```
[10]: # set up a dataframe for storing files' scores and labels
df = pd.DataFrame(index = glob('./great_plains_toad_dataset/*'), columns=species_df.
    ↪index.values)

# calculate RIBBIT scores
for path in df.index:

    for species, species_params in species_df.iterrows():
        #use RIBBIT for each species in species_df
        #make the spectrogram
        spec = Spectrogram.from_audio(audio.from_file(path))

        #run RIBBIT
        score_df = ribbit(
            spec,
            pulse_rate_range=species_params['pulse_rate_range'],
            signal_band=species_params['signal_band'],
            clip_duration=species_params['clip_duration'],
            noise_bands=species_params['noise_bands'],
            plot=False)

        # use the maximum RIBBIT score from any window as the score for this file
        # multiply the score by 10,000 to make it easier to read
        df.at[path,species] = max(score_df['score']) * 10000

print("Files with scores for each species, sorted by 'bird_series' score:")
df.sort_values(by='bird_series',ascending=False)
```


Files with scores for each species, sorted by 'bird\_series' score:

```
[10]:
```

|                                           | great_plains_toad | bird_series |
|-------------------------------------------|-------------------|-------------|
| ./great_plains_toad_dataset/negative5.wav | 0.0               | 93.81       |
| ./great_plains_toad_dataset/negative1.wav | 0.3               | 72.63       |
| ./great_plains_toad_dataset/negative3.wav | 0.3               | 5.05        |
| ./great_plains_toad_dataset/negative7.wav | 0.0               | 2.87        |
| ./great_plains_toad_dataset/negative9.wav | 3.04              | 0.09        |
| ./great_plains_toad_dataset/negative2.wav | 0.65              | 0.02        |
| ./great_plains_toad_dataset/negative6.wav | 0.06              | 0.01        |
| ./great_plains_toad_dataset/pops2.wav     | 0.0               | 0.01        |
| ./great_plains_toad_dataset/negative8.wav | 0.89              | 0.0         |
| ./great_plains_toad_dataset/negative4.wav | 0.76              | 0.0         |
| ./great_plains_toad_dataset/water.wav     | 0.0               | 0.0         |
| ./great_plains_toad_dataset/pops1.wav     | 0.0               | 0.0         |
| ./great_plains_toad_dataset/silent.wav    | 0.0               | 0.0         |
| ./great_plains_toad_dataset/gpt4.wav      | 0.12              | 0.0         |
| ./great_plains_toad_dataset/gpt0.wav      | 107.65            | 0.0         |
| ./great_plains_toad_dataset/gpt1.wav      | 10.13             | 0.0         |
| ./great_plains_toad_dataset/gpt3.wav      | 29.31             | 0.0         |
| ./great_plains_toad_dataset/gpt2.wav      | 16.69             | 0.0         |

looking at the highest scoring file for 'bird\_series', it has the trilled bird sound at 5-6.5 kHz

```
[11]: Spectrogram.from_audio(audio.from_file('./great_plains_toad_dataset/negative5.wav')).
      ↪ plot()
```



### 11.6.1 Warning

when loading a dataframe from a file, lists of numbers like [8,20] might be read in as strings (“[8,20]”) rather than a list of numbers. Here’s a handy little piece of code that will load the values in the desired format

```
[12]: #let's say we have the species df saved as a csv file
species_df.index.name='species'
species_df.to_csv('species_df.csv')

#define the conversion parameters for each column
import ast
generic = lambda x: ast.literal_eval(x)
conv = {
    'pulse_rate_range':generic,
    'signal_band':generic,
    'noise_bands':generic
}
#tell pandas to use them when loading the csv
species_df=pd.read_csv('./species_df.csv',converters=conv).set_index('species')

#now the species_df has numeric values instead of strings
species_df
```

```
[12]:
```

|                   | pulse_rate_range | signal_band  | clip_duration | \ |
|-------------------|------------------|--------------|---------------|---|
| species           |                  |              |               |   |
| great_plains_toad | [8, 20]          | [1900, 2400] | 2.0           |   |
| bird_series       | [8, 11]          | [5000, 6500] | 2.0           |   |

|                   | noise_bands               |
|-------------------|---------------------------|
| species           |                           |
| great_plains_toad | [[0, 1500], [2500, 3500]] |
| bird_series       | [[0, 4000]]               |

## 11.7 Detail view of RIBBIT method


Now, let's look at one 10 second file and tell ribbit to plot the power spectral density for each window (`plot=True`). This way, we can see if peaks are emerging at the expected pulse rates. Since our `window_length` is 2 seconds, each of these plots represents 2 seconds of audio. The vertical lines on the power spectral density represent the lower and upper `pulse_rate_range` limits.

In the file `gpt0.wav`, the Great Plains Toad vocalizes for a couple seconds at the beginning, then stops. We expect to see a peak in the power spectral density at 15 pulses/sec in the first 2 second window, and maybe a bit in the second, but not later in the audio.


```
[13]: #create a spectrogram from the file, like above:
# 1. get audio file path
audio_path = np.sort(glob('./great_plains_toad_dataset/*'))[0]
# 2. make audio object and trim (this time 0-10 seconds)
audio = Audio.from_file(audio_path).trim(0,10)
# 3. make spectrogram
spectrogram = Spectrogram.from_audio(audio)

clip_df = ribbit(
    spectrogram,
    pulse_rate_range=pulse_rate_range,
    signal_band=signal_band,
    clip_duration=clip_duration,
    noise_bands=noise_bands,
    plot=show_plots)
```


window: 0.0 to 2.0 sec



window: 2.0 to 4.0 sec



window: 4.0 to 6.0 sec



## 11.8 Time to experiment for yourself

Now that you know the basics of how to use RIBBIT, you can try using it on your own data. We recommend spending some time looking at different recordings of your focal species before choosing parameters. Experiment with the noise



bands and window length, and get in touch if you have questions!

Sam's email: sam . lapp [at] pitt.edu

this cell will delete the folder `great_plains_toad_dataset`. Only run it if you wish delete that folder and the example audio inside it.

```
[14]: from pathlib import Path
import shutil
shutil.rmtree('./great_plains_toad_dataset/')
Path('./species_df.csv').unlink()
```

Audio



## Spectrogram

spectrogram.py: Utilities for dealing with spectrograms

```
class opensoundscape.spectrogram.MelSpectrogram(spectrogram, frequencies, times, decibel_limits, window_samples=None, overlap_samples=None, window_type=None, audio_sample_rate=None, scaling=None)
```

Immutable mel-spectrogram container

A mel spectrogram is a spectrogram with pseudo-logarithmically spaced frequency bins (see literature) rather than linearly spaced bins.

See Spectrogram class and Librosa's melspectrogram for detailed documentation.

NOTE: Here we rely on scipy's spectrogram function (via Spectrogram) rather than on librosa's \_spectrogram or melspectrogram, because the amplitude of librosa's spectrograms do not match expectations. We only use the mel frequency bank from Librosa.

```
classmethod from_audio(audio, window_type='hann', window_samples=None, window_length_sec=None, overlap_samples=None, overlap_fraction=None, fft_size=None, decibel_limits=(-100, -20), dB_scale=True, scaling='spectrum', n_mels=64, norm='slaney', htk=False)
```

Create a MelSpectrogram object from an Audio object

First creates a spectrogram and a mel-frequency filter bank, then computes the dot product of the filter bank with the spectrogram.

A Mel spectrogram is a spectrogram with a quasi-logarithmic frequency axis that has often been used in language processing and other domains.

The kwargs for the mel frequency bank are documented at: - <https://librosa.org/doc/latest/generated/librosa.feature.melspectrogram.html#librosa.feature.melspectrogram> - <https://librosa.org/doc/latest/generated/librosa.filters.mel.html#librosa.filters.mel>

### Parameters

- **audio** – Audio object
- **window\_type="hann"** – see `scipy.signal.spectrogram` docs for description
- **window\_samples** – number of audio samples per spectrogram window (pixel) - Defaults to 512 if `window_samples` and `window_length_sec` are None - Note: cannot specify both `window_samples` and `window_length_sec`
- **window\_length\_sec** – length of a single window in seconds - Note: cannot specify both `window_samples` and `window_length_sec` - Warning: specifying this parameter often results in less efficient spectrogram computation because `window_samples` will not be a power of 2.
- **overlap\_samples** – number of samples shared by consecutive windows - Note: must not specify both `overlap_samples` and `overlap_fraction`
- **overlap\_fraction** – fractional temporal overlap between consecutive windows - Defaults to 0.5 if `overlap_samples` and `overlap_fraction` are None - Note: cannot specify both `overlap_samples` and `overlap_fraction`
- **fft\_size** – see `scipy.signal.spectrogram`'s *nfft* parameter
- **decibel\_limits** – limit the dB values to (min,max) (lower values set to min, higher values set to max)
- **dB\_scale** – If True, rescales values to decibels,  $x=10*\log_{10}(x)$  - if `dB_scale` is False, `decibel_limits` is ignored
- **scaling="spectrum"** – (“spectrum” or “density”) see `scipy.signal.spectrogram` docs
- **n\_mels** – Number of mel bands to generate [default: 128] Note: `n_mels` should be chosen for compatibility with the Spectrogram parameter *window\_samples*. Choosing a value  $> \sim \text{window\_samples}/10$  will result in zero-valued rows while small values blend rows from the original spectrogram.
- **norm='slanley'** – mel filter bank normalization, see Librosa docs
- **htk** – use HTK mel-filter bank instead of Slaney, see Librosa docs [default: False]

**Returns** `opensoundscape.spectrogram.MelSpectrogram` object

**plot** (*inline=True, fname=None, show\_colorbar=False*)

Plot the mel spectrogram with `matplotlib.pyplot`

We can't use `pcolormesh` because it will smash pixels to achieve a linear y-axis, rather than preserving the mel scale.

#### Parameters

- **inline=True** –
- **fname=None** – specify a string path to save the plot to (ending in .png/.pdf)
- **show\_colorbar** – include image legend colorbar from `pyplot`

**class** `opensoundscape.spectrogram.Spectrogram` (*spectrogram, frequencies, times, decibel\_limits, window\_samples=None, overlap\_samples=None, window\_type=None, audio\_sample\_rate=None, scaling=None*)

Immutable spectrogram container

Can be initialized directly from spectrogram, frequency, and time values or created from an Audio object using the `.from_audio()` method.

**frequencies**

(list) discrete frequency bins generated by fft times: (list) time from

**beginning of file to the center of each window spectrogram**

a 2d array containing

**10\*log10**

minimum and maximum decibel values in

**Type** fft

**.spectrogram window\_samples**

number of samples per window when spec was created [default: none]

**overlap\_samples**

number of samples overlapped in consecutive windows when spec was created [default: none]

**window\_type**

window fn used to make spectrogram, eg 'hann' [default: none]

**audio\_sample\_rate**

sample rate of audio from which spec was created [default: none]

**scaling**

Selects between computing the power spectral density ('density') where Sxx has units

**of V\*\*2/Hz and computing the power spectrum**

**Type** 'spectrum'

**is measured in V and fs is measured in Hz. [default**

spectrum]

**amplitude** (*freq\_range=None*)

create an amplitude vs time signal from spectrogram

by summing pixels in the vertical dimension

**Args** *freq\_range=None*: sum Spectrogram only in this range of [low, high] frequencies in Hz (if None, all frequencies are summed)

**Returns** a time-series array of the vertical sum of spectrogram value

**bandpass** (*min\_f, max\_f, out\_of\_bounds\_ok=True*)

extract a frequency band from a spectrogram

cropps the 2-d array of the spectrograms to the desired frequency range

**Parameters**

- **min\_f** – low frequency in Hz for bandpass
- **max\_f** – high frequency in Hz for bandpass
- **out\_of\_bounds\_ok** – (bool) if False, raises ValueError if min\_f or max\_f are not within the range of the original spectrogram's frequencies [default: True]

**Returns** bandpassed spectrogram object

**duration**

calculate the ammount of time represented in the spectrogram

Note: time may be shorter than the duration of the audio from which the spectrogram was created, because the windows may align in a way such that some samples from the end of the original audio were discarded

```
classmethod from_audio (audio, window_type='hann', window_samples=None, window_length_sec=None, overlap_samples=None, overlap_fraction=None, fft_size=None, decibel_limits=(-100, -20), dB_scale=True, scaling='spectrum')
```

create a Spectrogram object from an Audio object

#### Parameters

- **audio** – Audio object
- **window\_type="hann"** – see `scipy.signal.spectrogram` docs
- **window\_samples** – number of audio samples per spectrogram window (pixel) - Defaults to 512 if `window_samples` and `window_length_sec` are None - Note: cannot specify both `window_samples` and `window_length_sec`
- **window\_length\_sec** – length of a single window in seconds - Note: cannot specify both `window_samples` and `window_length_sec` - Warning: specifying this parameter often results in less efficient spectrogram computation because `window_samples` will not be a power of 2.
- **overlap\_samples** – number of samples shared by consecutive windows - Note: must not specify both `overlap_samples` and `overlap_fraction`
- **overlap\_fraction** – fractional temporal overlap between consecutive windows - Defaults to 0.5 if `overlap_samples` and `overlap_fraction` are None - Note: cannot specify both `overlap_samples` and `overlap_fraction`
- **fft\_size** – see `scipy.signal.spectrogram`'s *nfft* parameter
- **decibel\_limits** – limit the dB values to (min,max) (lower values set to min, higher values set to max)
- **dB\_scale** – If True, rescales values to decibels,  $x=10*\log_{10}(x)$  - if `dB_scale` is False, `decibel_limits` is ignored
- **scaling="spectrum"** – (“spectrum” or “density”) see `scipy.signal.spectrogram` docs

**Returns** `opensoundscape.spectrogram.Spectrogram` object

```
limit_db_range (min_db=-100, max_db=-20)
```

Limit the decibel values of the spectrogram to range from `min_db` to `max_db`

values less than `min_db` are set to `min_db` values greater than `max_db` are set to `max_db`

similar to Audacity's gain and range parameters

#### Parameters

- **min\_db** – values lower than this are set to this
- **max\_db** – values higher than this are set to this

**Returns** Spectrogram object with db range applied

```
linear_scale (feature_range=(0, 1))
```

Linearly rescale spectrogram values to a range of values using `in_range` as `decibel_limits`

**Parameters** **feature\_range** – tuple of (low,high) values for output

**Returns** Spectrogram object with values rescaled to `feature_range`

```
min_max_scale (feature_range=(0, 1))
```

Linearly rescale spectrogram values to a range of values using `in_range` as minimum and maximum

**Parameters** **feature\_range** – tuple of (low,high) values for output

**Returns** Spectrogram object with values rescaled to feature\_range

**net\_amplitude** (*signal\_band*, *reject\_bands=None*)

create amplitude signal in *signal\_band* and subtract amplitude from *reject\_bands*

rescale the signal and reject bands by dividing by their bandwidths in Hz (amplitude of each *reject\_band* is divided by the total bandwidth of all *reject\_bands*. amplitude of *signal\_band* is divided by bandwidth of *signal\_band*.)

#### Parameters

- **signal\_band** – [low,high] frequency range in Hz (positive contribution)
- **band** (*reject*) – list of [low,high] frequency ranges in Hz (negative contribution)

return: time-series array of net amplitude

**plot** (*inline=True*, *fname=None*, *show\_colorbar=False*)

Plot the spectrogram with matplotlib.pyplot

#### Parameters

- **inline=True** –
- **fname=None** – specify a string path to save the plot to (ending in .png/.pdf)
- **show\_colorbar** – include image legend colorbar from pyplot

**to\_image** (*shape=None*, *channels=1*, *colormap=None*, *invert=False*, *return\_type='pil'*)

Create an image from spectrogram (array, tensor, or PIL.Image)

Linearly rescales values in the spectrogram from self.decibel\_limits to [0,255] (PIL.Image) or [0,1] (array/tensor)

Default of self.decibel\_limits on load is [-100, -20], so, e.g., -20 db is loudest -> black, -100 db is quietest -> white

#### Parameters

- **shape** – tuple of output dimensions as (height, width) - if None, retains original shape of self.spectrogram
- **channels** – eg 3 for rgb, 1 for greyscale - must be 3 to use colormap
- **colormap** – if None, greyscale spectrogram is generated Can be any matplotlib colormap name such as 'jet'
- **return\_type** – type of returned object - 'pil': PIL.Image - 'np': numpy.ndarray - 'torch': torch.tensor

#### Returns

- **PIL.Image** with **c** channels and shape **w,h** given by *shape* and values in [0,255]
- np.ndarray with shape [c,h,w] and values in [0,1]
- or torch.tensor with shape [c,h,w] and values in [0,1]

**Return type** Image/array with type depending on *return\_type*

**trim** (*start\_time*, *end\_time*)

extract a time segment from a spectrogram

#### Parameters

- **start\_time** – in seconds
- **end\_time** – in seconds

**Returns** spectrogram object from extracted time segment

**window\_length**

type: calculate length of a single fft window, in seconds

**window\_start\_times**

get start times of each window, rather than midpoint times

**window\_step**

calculate time difference (sec) between consecutive windows' centers



classes for pytorch machine learning models in opensoundscape

For tutorials, see notebooks on opensoundscape.org

```
class opensoundscape.ml.cnn.CNN(architecture, classes, sample_duration, single_target=False, preprocessor_class=<class 'opensoundscape.preprocess.preprocessors.SpectrogramPreprocessor'>, sample_shape=(224, 224, 3))
```

Generic CNN Model with .train(), .predict(), and .save()

flexible architecture, optimizer, loss function, parameters

for tutorials and examples see opensoundscape.org

#### Parameters

- **architecture** – *EITHER* a pytorch model object (subclass of torch.nn.Module), for example one generated with the *cnn\_architectures* module *OR* a string matching one of the architectures listed by *cnn\_architectures.list\_architectures()*, eg 'resnet18'. - If a string is provided, uses default parameters  
(including pretrained weights, *weights*="DEFAULT") Note: if num channels != 3, copies weights from original channels by averaging (<3 channels) or recycling (>3 channels)
- **classes** – list of class names. Must match with training dataset classes if training.
- **single\_target** –
  - True: model expects exactly one positive class per sample
  - False: samples can have any number of positive classes
 [default: False]
- **preprocessor\_class** – class of Preprocessor object
- **sample\_shape** – tuple of height, width, channels for created sample [default: (224,224,3)] #TODO: consider changing to (ch,h,w) to match torchww

**eval** (*targets, scores, logging\_offset=0*)

compute single-target or multi-target metrics from targets and scores

By default, the overall model score is “map” (mean average precision) for multi-target models (self.single\_target=False) and “f1” (average of f1 score across classes) for single-target models).

Override this function to use a different set of metrics. It should always return (1) a single score (float) used as an overall metric of model quality and (2) a dictionary of computed metrics

#### Parameters

- **targets** – 0/1 for each sample and each class
- **scores** – continuous values in 0/1 for each sample and class
- **logging\_offset** – modify verbosity - for example, -1 will reduce the amount of printing/logging by 1 level

**classmethod from\_torch\_dict** (*path*)

load a model saved using CNN.save\_torch\_dict()

**Parameters** *path* – path to file saved using CNN.save\_torch\_dict()

**Returns** new CNN instance

Note: if you used .save() instead of .save\_torch\_dict(), load the model using cnn.load\_model(). Note that the model object will not load properly across different versions of OpenSoundscape. To save and load models across different versions of OpenSoundscape, use .save\_torch\_dict(), but note that preprocessing and other customized settings will not be retained.

**generate\_cams** (*samples, method='gradcam', classes=None, target\_layers=None, guided\_backprop=False, split\_files\_into\_clips=True, bypass\_augmentations=True, batch\_size=1, num\_workers=0*)

Generate a activation and/or backprop heatmaps for each sample

#### Parameters

- **samples** – (same as CNN.predict()) the files to generate predictions for. Can be: - a dataframe with index containing audio paths, OR - a dataframe with multi-index (file, start\_time, end\_time), OR - a list (or np.ndarray) of audio file paths
- **method** – method to use for activation map. Can be str (choose from below) or a class of pytorch\_grad\_cam (any subclass of BaseCAM), or None if None, activation maps will not be created [default: 'gradcam']

**str can be any of the following:** "gradcam": pytorch\_grad\_cam.GradCAM, "hirescam": pytorch\_grad\_cam.HiResCAM, "scorecam": pytorch\_grad\_cam.ScoreCAM, "gradcam++": pytorch\_grad\_cam.GradCAMPlusPlus, "ablationcam": pytorch\_grad\_cam.AblationCAM, "xgradcam": pytorch\_grad\_cam.XGradCAM, "eigencam": pytorch\_grad\_cam.EigenCAM, "eigengradcam": pytorch\_grad\_cam.EigenGradCAM, "layercam": pytorch\_grad\_cam.LayerCAM, "fullgrad": pytorch\_grad\_cam.FullGrad, "gradcamelementwise": pytorch\_grad\_cam.GradCAMElementWise,

- **classes** (*list*) – list of classes, will create maps for each class [default: None] if None, creates an activation map for the highest scoring class on a sample-by-sample basis
- **target\_layers** (*list*) – list of target layers for GradCAM - if None [default] attempts to use architecture's default target\_layer Note: only architectures created with opensoundscape 0.9.0+ will have a default target layer. See pytorch\_grad\_cam docs for suggestions. Note: if multiple layers are provided, the activations are merged across

layers (rather than returning separate activations per layer)

- **guided\_backprop** – bool [default: False] if True, performs guided backpropagation for each class in classes. AudioSamples will have attribute .gbp\_maps, a pd.Series indexed by class name
- **split\_files\_into\_clips** (*bool*) – see CNN.predict()
- **bypass\_augmentations** (*bool*) – whether to bypass augmentations in preprocessing see CNN.predict
- **batch\_size** – number of samples to simultaneously process, see CNN.predict()
- **num\_workers** – parallel CPU threads for preprocessing, see CNN.predict()

**Returns** a list of cam class activation objects. see the cam class for more details

See pytorch\_grad\_cam documentation for references to the source of each method.

**load\_weights** (*path*, *strict=True*)

load network weights state dict from a file

For instance, load weights saved with .save\_weights() in-place operation

#### Parameters

- **path** – file path with saved weights
- **strict** – (bool) see torch.load()

**predict** (*samples*, *batch\_size=1*, *num\_workers=0*, *activation\_layer=None*, *split\_files\_into\_clips=True*, *overlap\_fraction=0*, *final\_clip=None*, *bypass\_augmentations=True*, *invalid\_samples\_log=None*, *raise\_errors=False*, *wandb\_session=None*, *return\_invalid\_samples=False*)

Generate predictions on a dataset

Choose to return any combination of scores, labels, and single-target or multi-target binary predictions. Also choose activation layer for scores (softmax, sigmoid, softmax then logit, or None). Binary predictions are performed post-activation layer

Note: the order of returned dataframes is (scores, preds, labels)

#### Parameters

- **samples** – the files to generate predictions for. Can be: - a dataframe with index containing audio paths, OR - a dataframe with multi-index (file, start\_time, end\_time), OR - a list (or np.ndarray) of audio file paths
- **batch\_size** – Number of files to load simultaneously [default: 1]
- **num\_workers** – parallelization (ie cpus or cores), use 0 for current process [default: 0]
- **activation\_layer** – Optionally apply an activation layer such as sigmoid or softmax to the raw outputs of the model. options: - None: no activation, return raw scores (ie logit, [-inf:inf]) - 'softmax': scores all classes sum to 1 - 'sigmoid': all scores in [0,1] but don't sum to 1 - 'softmax\_and\_logit': applies softmax first then logit [default: None]
- **split\_files\_into\_clips** – If True, internally splits and predicts on clips from longer audio files Otherwise, assumes each row of *samples* corresponds to one complete sample
- **overlap\_fraction** – fraction of overlap between consecutive clips when predicting on clips of longer audio files. For instance, 0.5 gives 50% overlap between consecutive clips.
- **final\_clip** – see *opensoundscape.utils.generate\_clip\_times\_df*

- **bypass\_augmentations** – If False, Actions with `is_augmentation==True` are performed. Default True.
- **invalid\_samples\_log** – if not None, samples that failed to preprocess will be listed in this text file.
- **raise\_errors** – if True, raise errors when preprocessing fails if False, just log the errors to `unsafe_samples_log`
- **wandb\_session** – a wandb session to log to - pass the value returned by `wandb.init()` to progress log to a Weights and Biases run - if None, does not log to wandb
- **return\_invalid\_samples** – bool, if True, returns second argument, a set containing file paths of samples that caused errors during preprocessing [default: False]

**Returns** df of post-activation\_layer scores - if `return_invalid_samples` is True, returns (df,invalid\_samples) where `invalid_samples` is a set of file paths that failed to preprocess

**Effects:** (1) wandb logging If `wandb_session` is provided, logs progress and samples to Weights and Biases. A random set of samples is preprocessed and logged to a table. Progress over all batches is logged. After prediction, top scoring samples are logged. Use `self.wandb_logging` dictionary to change the number of samples logged or which classes have top-scoring samples logged.

(2) unsafe sample logging If `unsafe_samples_log` is not None, saves a list of all file paths that failed to preprocess in `unsafe_samples_log` as a text file

**Note: if loading an audio file raises a PreprocessingError, the scores** for that sample will be `np.nan`

**save** (*path*, *save\_train\_loader=False*, *save\_hooks=False*, *as\_torch\_dict=False*)

save model with weights using `torch.save()`

load from saved file with `torch.load(path)` or `cnn.load_model(path)`

Note: saving and loading model objects across OpenSoundscape versions will not work properly. Instead, use `.save_torch_dict` and `.load_torch_dict` (but note that customizations to preprocessing, training params, etc will not be retained using those functions).

For maximum flexibility in further use, save the model with both `.save()` and `.save_torch_dict()`

#### Parameters

- **path** – file path for saved model object
- **save\_train\_loader** – retrain `.train_loader` in saved object [default: False]
- **save\_hooks** – retain forward and backward hooks on modules [default: False] Note: True can cause issues when using `wandb.watch()`

**save\_torch\_dict** (*path*)

save model to file for use in other opso versions

**WARNING: this does not save any preprocessing or augmentation** settings or parameters, or other attributes such as the training parameters or loss function. It only saves architecture, weights, classes, sample shape, sample duration, and `single_target`.

To save the entire pickled model object (recover all parameters and settings), use `model.save()` instead. Note that models saved with `model.save()` will not work across different versions of OpenSoundscape.

To recreate the model after saving with this function, use `CNN.from_torch_dict(path)`

**Parameters** **path** – file path for saved model object

**Effects:** saves a file using `torch.save()` containing model weights and other information

### **save\_weights** (*path*)

save just the weights of the network

This allows the saved weights to be used more flexibly than `model.save()` which will pickle the entire object. The weights are saved in a pickled dictionary using `torch.save(self.network.state_dict())`

**Parameters** *path* – location to save weights file

**train** (*train\_df*, *validation\_df=None*, *epochs=1*, *batch\_size=1*, *num\_workers=0*,  
*save\_path='.'*, *save\_interval=1*, *log\_interval=10*, *validation\_interval=1*, *invalid\_samples\_log='./invalid\_training\_samples.log'*, *raise\_errors=False*,  
*wandb\_session=None*)  
 train the model on samples from *train\_dataset*

If customized loss functions, networks, optimizers, or schedulers are desired, modify the respective attributes before calling `.train()`.

### **Parameters**

- **train\_df** – a dataframe of files and labels for training the model - either has index *file* or multi-index (*file,start\_time,end\_time*)
- **validation\_df** – a dataframe of files and labels for evaluating the model [default: None means no validation is performed]
- **epochs** – number of epochs to train for (1 epoch constitutes 1 view of each training sample)
- **batch\_size** – number of training files simultaneously passed through forward pass, loss function, and backpropagation
- **num\_workers** – number of parallel CPU tasks for preprocessing Note: use 0 for single (root) process (not 1)
- **save\_path** – location to save intermediate and best model objects [default=".", ie current location of script]
- **save\_interval** – interval in epochs to save model object with weights [default:1]  
 Note: the best model is always saved to `best.model` in addition to other saved epochs.
- **log\_interval** – interval in batches to print training loss/metrics
- **validation\_interval** – interval in epochs to test the model on the validation set  
 Note that model will only update it's best score and save `best.model` file on epochs that it performs validation.
- **invalid\_samples\_log** – file path: log all samples that failed in preprocessing (file written when training completes) - if None, does not write a file
- **raise\_errors** – if True, raise errors when preprocessing fails if False, just log the errors to `unsafe_samples_log`
- **wandb\_session** – a wandb session to log to - pass the value returned by `wandb.init()` to progress log to a Weights and Biases run - if None, does not log to wandb  
 For example:
 

```
` import wandb
wandb.login(key=api_key) #find your api_key at https://wandb.ai/settings
session = wandb.init(entity='mygroup',project='project1',name='first_run')
... model.train(...,wandb_session=session)
session.finish()
`
```

**Effects:** If `wandb_session` is provided, logs progress and samples to Weights and Biases. A random set of training and validation samples are preprocessed and logged to a table. Training progress, loss, and metrics are also logged. Use `self.wandb_logging` dictionary to change the number of samples logged.

```
class opensoundscape.ml.cnn.InceptionV3(classes, sample_duration, single_target=False,
   preprocessor_class=<class 'opensound-
   scape.preprocess.preprocessors.SpectrogramPreprocessor'>,
   freeze_feature_extractor=False,
   weights='DEFAULT', sample_shape=(299, 3))
```

Child of CNN class for InceptionV3 architecture

**classmethod from\_torch\_dict()**

load a model saved using CNN.save\_torch\_dict()

**Parameters** *path* – path to file saved using CNN.save\_torch\_dict()

**Returns** new CNN instance

Note: if you used .save() instead of .save\_torch\_dict(), load the model using cnn.load\_model(). Note that the model object will not load properly across different versions of OpenSoundscape. To save and load models across different versions of OpenSoundscape, use .save\_torch\_dict(), but note that preprocessing and other customized settings will not be retained.

```
opensoundscape.ml.cnn.load_model(path, device=None)
```

load a saved model object

Note: saving and loading model objects across OpenSoundscape versions will not work properly. Instead, use .save\_torch\_dict and .load\_torch\_dict (but note that customizations to preprocessing, training params, etc will not be retained using those functions).

For maximum flexibility in further use, save the model with both .save() and .save\_torch\_dict()

**Parameters**

- **path** – file path of saved model
- **device** – which device to load into, eg 'cuda:1'
- **[default – None]** will choose first gpu if available, otherwise cpu

**Returns** a model object with loaded weights

```
opensoundscape.ml.cnn.load_outdated_model(path, architecture, sample_duration,
   model_class=<class 'opensound-
   scape.ml.cnn.CNN'>, device=None)
```

load a CNN saved with a version of OpenSoundscape <0.6.0

This function enables you to load models saved with opso 0.4.x and 0.5.x. If your model was saved with .save() in a previous version of OpenSoundscape >=0.6.0, you must re-load the model using the original package version and save it's network's state dict, i.e., *torch.save(model.network.state\_dict(),path)*, then load the state dict to a new model object with *model.load\_weights()*. See the *Predict with pre-trained CNN* tutorial for details.

For models created with the same version of OpenSoundscape as the one you are using, simply use *opensoundscape.torch.models.cnn.load\_model()*.

Note: for future use of the loaded model, you can simply call *model.save(path)* after creating it, then reload it with *model = load\_model(path)*. The saved model will be fully compatible with opensoundscape >=0.7.0.

Examples: ““ #load a binary resnet18 model from opso 0.4.x, 0.5.x, or 0.6.0 from opensoundscape import CNN  
model = load\_outdated\_model('old\_model.tar',architecture='resnet18') ““

#load a resnet50 model of class CNN created with opso 0.5.0 from opensoundscape import CNN model\_050 =  
load\_outdated\_model('opso050\_pytorch\_model\_r50.model',architecture='resnet50') ““

**Parameters**

- **path** – path to model file, ie .model or .tar file

- **architecture** – see CNN docs (pass None if the class `__init__` does not take architecture as an argument)
- **sample\_duration** – length of samples in seconds
- **model\_class** – class to construct. Normally CNN.
- **device** – optionally specify a device to map tensors onto,
- **'cpu', 'cuda** (*eg*) – 0', 'cuda:1'[default: None] - if None, will choose cuda:0 if cuda is available, otherwise chooses cpu

**Returns** a cnn model object with the weights loaded from the saved model

`opensoundscape.ml.cnn.separate_resnet_feat_clf(model)`

Separate feature/classifier training params for a ResNet model

**Parameters** **model** – an opso model object with a pytorch resnet architecture

**Returns** model with modified `.optimizer_params` and `._init_optimizer()` method

**Effects:** creates a new `self.opt_net` object that replaces the old one resets `self.current_epoch` to 0

`opensoundscape.ml.cnn.use_resample_loss(model)`

Modify a model to use ResampleLoss for multi-target training

ResampleLoss may perform better than BCE Loss for multitarget problems in some scenarios.





functions and classes for manipulating annotations of audio

includes BoxedAnnotations class and utilities to combine or “diff” annotations, etc.

**class** `opensoundscape.annotations.BoxedAnnotations` (*df*)

container for “boxed” (frequency-time) annotations of audio (for instance, annotations created in Raven software)

includes functionality to load annotations from Pandas DataFrame or Raven Selection tables (.txt files), output one-hot labels for specific clip lengths or clip start/end times, apply corrections/conversions to annotations, and more.

Contains some analogous functions to Audio and Spectrogram, such as `trim()` [limit time range] and `bandpass()` [limit frequency range]

**bandpass** (*low\_f, high\_f, edge\_mode='trim'*)

Bandpass a set of annotations, analogous to `Spectrogram.bandpass()`

Reduces the range of annotation boxes overlapping with the bandpass limits, and removes annotation boxes entirely if they lie completely outside of the bandpass limits.

Out-of-place operation: does not modify itself, returns new object

### Parameters

- **low\_f** – low frequency (Hz) bound
- **high\_f** – high frequency (Hz) bound
- **edge\_mode** – what to do when boxes overlap with edges of trim region - ‘trim’: trim boxes to bounds - ‘keep’: allow boxes to extend beyond bounds - ‘remove’: completely remove boxes that extend beyond bounds

**Returns** a copy of the BoxedAnnotations object on the bandpassed region

**convert\_labels** (*conversion\_table*)

modify annotations according to a conversion table

Changes the values of ‘annotation’ column of dataframe. Any labels that do not have specified conversions are left unchanged.

Returns a new BoxedAnnotations object, does not modify itself (out-of-place operation). So use could look like: `my_annotations = my_annotations.convert_labels(table)`

**Parameters** `conversion_table` – current values -> new values. can be either -  
`pd.DataFrame` with 2 columns [current value, new values] or - dictionary {current values:  
new values}

**Returns** new BoxedAnnotations object with converted annotation labels

**classmethod** `from_raven_files` (`raven_paths`, `audio_files=None`, `annotation_column_idx=7`,  
`keep_extra_columns=True`)

load annotations from Raven .txt files

**Parameters**

- **raven\_paths** – list of raven .txt file paths (as str or `pathlib.Path`)
- **audio\_files** – (list) optionally specify audio files corresponding to each raven file (length should match `raven_paths`)
- **annotation\_column\_idx** – (int) position of column containing annotations - [default: 7] will be correct if the first user-created column in Raven contains annotations - pass `None` to load the raven file without explicitly assigning a column as the annotation column. The resulting object’s `.df` will have an `annotation` column with nan values!
- **keep\_extra\_columns** – keep or discard extra Raven file columns (always keeps `start_time`, `end_time`, `low_f`, `high_f`, `annotation` `audio_file`). [default: `True`] - `True`: keep all - `False`: keep none - or iterable of specific columns to keep

**Returns** BoxedAnnotations object containing annotations from the Raven files (the `.df` attribute is a dataframe containing each annotation)

**global** `one_hot_labels` (`classes`)

get a list of one-hot labels for entire set of annotations :param classes: iterable of class names to give 0/1 labels

**Returns** list of 0/1 labels for each class

**one\_hot\_clip\_labels** (`clip_duration`, `clip_overlap`, `min_label_overlap`, `min_label_fraction=1`,  
`full_duration=None`, `class_subset=None`, `final_clip=None`)

Generate one-hot labels for clips of fixed duration

wraps `utils.make_clip_df()` with `self.one_hot_labels_like()` - Clips are created in the same way as `Audio.split()` - Labels are applied based on overlap, using `self.one_hot_labels_like()`

**Parameters**

- **clip\_duration** (`float`) – The duration in seconds of the clips
- **clip\_overlap** (`float`) – The overlap of the clips in seconds [default: 0]
- **min\_label\_overlap** – minimum duration (seconds) of annotation within the time interval for it to count as a label. Note that any annotation of length less than this value will be discarded. We recommend a value of 0.25 for typical bird songs, or shorter values for very short-duration events such as chip calls or nocturnal flight calls.
- **min\_label\_fraction** – [default: `None`] if  $\geq$  this fraction of an annotation overlaps with the time window, it counts as a label regardless of its duration. Note that *if either* of the two criteria (overlap and fraction) is met, the label is 1. if `None` (default), this criterion is not used (i.e., only `min_label_overlap` is used). A value of 0.5 for this parameter would

ensure that all annotations result in at least one clip being labeled 1 (if there are no gaps between clips).

- **full\_duration** – The amount of time (seconds) to split into clips for each file float or *None*; if *None*, attempts to get each file’s duration using *librosa.get\_duration(path=file)*
- **class\_subset** – list of classes for one-hot labels. If *None*, classes will be all unique values of *self.df[‘annotation’]*
- **final\_clip** (*str*) – Behavior if final\_clip is less than clip\_duration seconds long. By default, discards remaining time if less than clip\_duration seconds long [default: *None*]. Options: - *None*: Discard the remainder (do not make a clip) - “extend”: Extend the final clip beyond full\_duration to reach

clip\_duration length

- “remainder”: Use only remainder of full\_duration (final clip will be shorter than clip\_duration)
- “full”: Increase overlap with previous clip to yield a clip with clip\_duration length

**Returns** dataframe with index [‘file’, ‘start\_time’, ‘end\_time’] and columns=classes

**one\_hot\_labels\_like** (*clip\_df*, *min\_label\_overlap*, *min\_label\_fraction=None*,  
*class\_subset=None*, *warn\_no\_annotations=False*)  
create a dataframe of one-hot clip labels based on given starts/ends

Uses start and end clip times from *clip\_df* to define a set of clips for each file. Then extracts annotations overlapping with each clip.

Required overlap to consider an annotation to overlap with a clip is defined by user: an annotation must satisfy the minimum time overlap OR minimum % overlap to be included (doesn’t require both conditions to be met, only one)

*clip\_df* can be created using *opensoundscape.utils.make\_clip\_df*

See also: *.one\_hot\_clip\_labels()*, which creates even-lengthed clips automatically and can often be used instead of this function.

#### Parameters

- **clip\_df** – dataframe with (file, start\_time, end\_time) MultiIndex specifying the temporal bounds of each clip (*clip\_df* can be created using *opensoundscape.helpers.make\_clip\_df*)
- **min\_label\_overlap** – minimum duration (seconds) of annotation within the time interval for it to count as a label. Note that any annotation of length less than this value will be discarded. We recommend a value of 0.25 for typical bird songs, or shorter values for very short-duration events such as chip calls or nocturnal flight calls.
- **min\_label\_fraction** – [default: *None*] if  $\geq$  this fraction of an annotation overlaps with the time window, it counts as a label regardless of its duration. Note that *if either* of the two criteria (overlap and fraction) is met, the label is 1. if *None* (default), this criterion is not used (i.e., only *min\_label\_overlap* is used). A value of 0.5 for this parameter would ensure that all annotations result in at least one clip being labeled 1 (if there are no gaps between clips).
- **class\_subset** – list of classes for one-hot labels. If *None*, classes will be all unique values of *self.df[‘annotation’]*

- **warn\_no\_annotations** – bool [default:False] if True, raises warnings for any files in clip\_df with no corresponding annotations in self.df

**Returns** DataFrame of one-hot labels w/ multi-index of (file, start\_time, end\_time), a column for each class, and values of 0=absent or 1=present

**subset** (*classes*)

subset annotations to those from a list of classes

out-of-place operation (returns new filtered BoxedAnnotations object)

**Parameters**

- **classes** – list of classes to retain (all others are discarded)
- **the list can include nan or None if you want to keep them (-)** –

**Returns** new BoxedAnnotations object containing only annotations in *classes*

**to\_raven\_files** (*save\_dir*)

save annotations to a Raven-compatible tab-separated text files

Creates one file per unique audio file in ‘file’ column of self.df

**Parameters** **save\_dir** – directory for saved files - can be str or pathlib.Path

**Outcomes:** creates files containing the annotations for each audio file in a format compatible with Raven Pro/Lite. File is tab-separated and contains columns matching the Raven defaults.

Note: Raven Lite does not support additional columns beyond a single annotation column. Additional columns will not be shown in the Raven Lite interface.

**trim** (*start\_time, end\_time, edge\_mode='trim'*)

Trim the annotations of each file in time

Trims annotations from outside of the time bounds. Note that the annotation start and end times of different files may not represent the same real-world times. This function only uses the numeric values of annotation start and end times in the annotations, which should be relative to the beginning of the corresponding audio file.

For zero-length annotations (start\_time = end\_time), start and end times are inclusive on the left and exclusive on the right, ie [lower,upper). For instance start\_time=0, end\_time=1 includes zero-length annotations at 0 but excludes zero-length annotations at 1.

Out-of-place operation: does not modify itself, returns new object

**Parameters**

- **start\_time** – time (seconds) since beginning for left bound
- **end\_time** – time (seconds) since beginning for right bound
- **edge\_mode** – what to do when boxes overlap with edges of trim region - ‘trim’: trim boxes to bounds - ‘keep’: allow boxes to extend beyond bounds - ‘remove’: completely remove boxes that extend beyond bounds

**Returns** a copy of the BoxedAnnotations object on the trimmed region. - note that, like Audio.trim(), there is a new reference point for 0.0 seconds (located at start\_time in the original object). For example, calling .trim(5,10) will result in an annotation previously starting at 6s to start at 1s in the new object.

**unique\_labels()**

get list of all unique labels

ignores null/Falsy labels by performing `.df.dropna()`

`opensoundscape.annotations.categorical_to_one_hot(labels, class_subset=None)`

transform multi-target categorical labels (list of lists) to one-hot array

#### Parameters

- **labels** – list of lists of categorical labels, eg `[[‘white’,‘red’],[‘green’,‘white’]]` or `[[0,1,2],[3]]`
- **classes=None** – list of classes for one-hot labels. if None, taken to be the unique set of values in *labels*

**Returns** 2d array with 0 for absent and 1 for present *class\_subset*: list of classes corresponding to columns in the array

**Return type** `one_hot`

`opensoundscape.annotations.diff(base_annotations, comparison_annotations)`

look at differences between two BoxedAnnotations objects Not Implemented.

Compare different labels of the same boxes (Assumes that a second annotator used the same boxes as the first, but applied new labels to the boxes)

`opensoundscape.annotations.one_hot_labels_on_time_interval(df, class_subset, start_time, end_time, min_label_overlap, min_label_fraction=None)`

generate a dictionary of one-hot labels for given time-interval

Each class is labeled 1 if any annotation overlaps sufficiently with the time interval. Otherwise the class is labeled 0.

#### Parameters

- **df** – DataFrame with columns ‘start\_time’, ‘end\_time’ and ‘annotation’
- **classes** – list of classes for one-hot labels. If None, classes will be all unique values of `self.df[‘annotation’]`
- **start\_time** – beginning of time interval (seconds)
- **end\_time** – end of time interval (seconds)
- **min\_label\_overlap** – minimum duration (seconds) of annotation within the time interval for it to count as a label. Note that any annotation of length less than this value will be discarded. We recommend a value of 0.25 for typical bird songs, or shorter values for very short-duration events such as chip calls or nocturnal flight calls.
- **min\_label\_fraction** – [default: None] if  $\geq$  this fraction of an annotation overlaps with the time window, it counts as a label regardless of its duration. Note that *if either* of the two criteria (overlap and fraction) is met, the label is 1. if None (default), the criterion is not used (only `min_label_overlap` is used). A value of 0.5 would ensure that all annotations result in at least one clip being labeled 1 (if no gaps between clips).

**Returns** label 0/1 } for all classes

**Return type** dictionary of {class

`opensoundscape.annotations.one_hot_to_categorical(one_hot, classes)`

transform one\_hot labels to multi-target categorical (list of lists)

**Parameters**

- **one\_hot** – 2d array with 0 for absent and 1 for present
- **classes** – list of classes corresponding to columns in the array

**Returns**

**list of lists of categorical labels for each sample, eg** `[[‘white’,‘red’],[‘green’,‘white’]]` or `[[0,1,2],[3]]`

**Return type** labels

Utilities specifically for audio files recorded by AudioMoths and other ARUs

`opensoundscape.aru.audiomoth_start_time` (*file*, *filename\_timezone*='UTC', *to\_utc*=False)  
 parse audiomoth file name into a time stamp

AudioMoths create their file name based on the time that recording starts. This function parses the name into a timestamp. Older AudioMoth firmwares used a hexadecimal unix time format, while newer firmwares use a human-readable naming convention. This function handles both conventions.

### Parameters

- **file** – (str) path or file name from AudioMoth recording
- **filename\_timezone** – (str) name of a pytz time zone (for options see `pytz.all_timezones`). This is the time zone that the AudioMoth uses to record its name, not the time zone local to the recording site. Usually, this is 'UTC' because the AudioMoth records file names in UTC.
- **to\_utc** – if True, converts timestamps to UTC localized time stamp. Otherwise, will return timestamp localized to *timezone* argument [default: False]

**Returns** localized datetime object - if *to\_utc*=True, datetime is always “localized” to UTC

`opensoundscape.aru.parse_audiomoth_metadata` (*metadata*)  
 parse a dictionary of AudioMoth .wav file metadata

-parses the comment field -adds keys for *gain\_setting*, *battery\_state*, *recording\_start\_time* -if available (firmware >=1.4.0), adds temperature

Notes on comment field: - Starting with Firmware 1.4.0, the audiomoth logs Temperature to the metadata (wav header) eg “and temperature was 11.2C.”

- At some point the firmware shifted from writing “gain setting 2” to “medium gain setting” and later to just “medium gain”. Should handle both modes.
- In later versions of the firmware, “state” was ommitted from “battery state” in the comment field. Handles either phrasing.

**Tested for AudioMoth firmware versions:** 1.5.0 through 1.8.1

**Parameters** `metadata` – dictionary with audiomoth metadata

**Returns** metadata dictionary with added keys and values

`opensoundscape.aru.parse_audiomoth_metadata_from_path` (*file\_path*)

Wrapper function to parse audiomoth metadata from file path



Utilities for .ml

**class** `opensoundscape.ml.utils.BaseModule` (*\*args, \*\*kwargs*)  
Base class for a pytorch model pipeline class.

All child classes should define load, save, etc

`opensoundscape.ml.utils.apply_activation_layer` (*x, activation\_layer=None*)  
applies an activation layer to a set of scores

### Parameters

- **x** – input values
- **activation\_layer** –
  - None [default]: return original values
  - 'softmax': apply softmax activation
  - 'sigmoid': apply sigmoid activation
  - 'softmax\_and\_logit': apply softmax then logit transform

**Returns** values with activation layer applied

`opensoundscape.ml.utils.cas_dataloader` (*dataset, batch\_size, num\_workers*)  
Return a dataloader that uses the class aware sampler

Class aware sampler tries to balance the examples per class in each batch. It selects just a few classes to be present in each batch, then samples those classes for even representation in the batch.

### Parameters

- **dataset** – a pytorch dataset type object
- **batch\_size** – see DataLoader
- **num\_workers** – see DataLoader

`opensoundscape.ml.utils.get_batch(array, batch_size, batch_number)`

get a single slice of a larger array

using the batch size and batch index, from zero

**Parameters**

- **array** – iterable to split into batches
- **batch\_size** – num elements per batch
- **batch\_number** – index of batch

**Returns** one batch (subset of array)

Note: the final elements are returned as the last batch even if there are fewer than `batch_size`

**Example**

if `array=[1,2,3,4,5,6,7]` then:

- `get_batch(array,3,0)` returns `[1,2,3]`
- `get_batch(array,3,3)` returns `[7]`

---

## CNN Architectures

---

Module to initialize PyTorch CNN architectures with custom output shape

This module allows the use of several built-in CNN architectures from PyTorch. The architecture refers to the specific layers and layer input/output shapes (including convolution sizes and strides, etc) - such as the ResNet18 or Inception V3 architecture.

We provide wrappers which modify the output layer to the desired shape (to match the number of classes). The way to change the output layer shape depends on the architecture, which is why we need a wrapper for each one. This code is based on [pytorch.org/tutorials/beginner/finetuning\\_torchvision\\_models\\_tutorial.html](https://pytorch.org/tutorials/beginner/finetuning_torchvision_models_tutorial.html)

To use these wrappers, for example, if your model has 10 output classes, write

```
my_arch=resnet18(10)
```

Then you can initialize a model object from *opensoundscape.torch.models.cnn* with your architecture:

```
model=CNN(my_arch,classes,sample_duration)
```

or override an existing model's architecture:

```
model.network = my_arch
```

Note: the InceptionV3 architecture must be used differently than other architectures - the easiest way is to simply use the InceptionV3 class in *opensoundscape.torch.models.cnn*.

```
opensoundscape.ml.cnn_architectures.alexnet(num_classes,  
  freeze_feature_extractor=False,  
  weights='DEFAULT', num_channels=3)
```

Wrapper for AlexNet architecture

input size = 224

### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained

- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if ‘DEFAULT’, model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.change_conv2d_channels (conv2d,
  num_channels=3,
  reuse_weights=True)
```

Modify the number of input channels for a pytorch CNN

This function changes the input shape of a torch.nn.Conv2D layer to accommodate a different number of channels. It attempts to retain weights in the following manner: - If num\_channels is less than the original, it will average weights across the original channels and apply them to all new channels. - if num\_channels is greater than the original, it will cycle through the original channels, copying them to the new channels

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape
- **reuse\_weights** – if True (default), averages (if num\_channels<original)
- **cycles through** (*or*) – and adds them to the new Conv2D

```
opensoundscape.ml.cnn_architectures.change_fc_output_size (fc, num_classes)
```

Modify the number of output nodes of a fully connected layer

#### Parameters

- **fc** – the fully connected layer of the model that should be modified
- **num\_classes** – number of output nodes for the new fc

```
opensoundscape.ml.cnn_architectures.densenet121 (num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT',
  num_channels=3)
```

Wrapper for densenet121 architecture

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if ‘DEFAULT’, model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.efficientnet_b0 (num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT',
  num_channels=3)
```

Wrapper for efficientnet\_b0 architecture

**Parameters**

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.efficientnet_b4 (num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT',
  num_channels=3)
```

Wrapper for efficientnet\_b4 architecture

**Parameters**

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.efficientnet_widese_b0 (num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT',
  num_channels=3)
```

Wrapper for efficientnet\_widese\_b0 architecture

**Parameters**

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.efficientnet_widese_b4 (num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT',
  num_channels=3)
```

Wrapper for efficientnet\_widese\_b4 architecture

**Parameters**

- **num\_classes** – number of output nodes for the final layer

- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

`opensoundscape.ml.cnn_architectures.freeze_params(model)`  
 remove gradients (aka freeze) all model parameters

`opensoundscape.ml.cnn_architectures.inception_v3(num_classes,  
 freeze_feature_extractor=False,  
 weights='DEFAULT',  
 num_channels=3)`

Wrapper for Inception v3 architecture

Input: 229x229

WARNING: expects (299,299) sized images and has auxiliary output. See InceptionV3 class in *opensoundscape.torch.models.cnn* for use.

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

`opensoundscape.ml.cnn_architectures.list_architectures()`  
 return list of available architecture keyword strings

`opensoundscape.ml.cnn_architectures.register_arch(func)`  
 add architecture to ARCH\_DICT

`opensoundscape.ml.cnn_architectures.resnet101(num_classes,  
 freeze_feature_extractor=False,  
 weights='DEFAULT', num_channels=3)`

Wrapper for ResNet101 architecture

input\_size = 224

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.resnet152(num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT', num_channels=3)
```

Wrapper for ResNet152 architecture

input\_size = 224

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.resnet18(num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT', num_channels=3)
```

Wrapper for ResNet18 architecture

input\_size = 224

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.resnet34(num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT', num_channels=3)
```

Wrapper for ResNet34 architecture

input\_size = 224

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.resnet50 (num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT', num_channels=3)
```

Wrapper for ResNet50 architecture

input\_size = 224

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.squeezenet1_0 (num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT',
  num_channels=3)
```

Wrapper for squeezenet architecture

input size = 224

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>
- **num\_channels** – specify channels in input sample, eg [channels h,w] sample shape

```
opensoundscape.ml.cnn_architectures.vgg11_bn (num_classes,
  freeze_feature_extractor=False,
  weights='DEFAULT', num_channels=3)
```

Wrapper for vgg11 architecture

input size = 224

#### Parameters

- **num\_classes** – number of output nodes for the final layer
- **freeze\_feature\_extractor** – if False (default), entire network will have gradients and can train if True, feature block is frozen and only final layer is trained
- **weights** – string containing version name of the pre-trained classification weights to use for this architecture. if 'DEFAULT', model is loaded with best available weights (note that these may change across versions). Pre-trained weights available for each architecture are listed at <https://pytorch.org/vision/stable/models.html>



---

## Logging with WandB (Weights and Biases)

---

helpers for integrating with WandB and exporting content

```
opensoundscape.logging.wandb_table(dataset, n=None, classes_to_extract=(), random_state=None, raise_exceptions=False, drop_labels=False)
```

Generate a wandb Table visualizing n random samples from a sample\_df

### Parameters

- **dataset** – object to generate samples, eg AudioFileDataset or AudioSplittingDataset
- **n** – number of samples to generate (randomly selected from df) - if None, does not sub-sample or change order
- **bypass\_augmentations** – if True, augmentations in Preprocessor are skipped
- **classes\_to\_extract** – tuple of classes - will create columns containing the scores/labels
- **random\_state** – default None; if integer provided, used for reproducible random sample
- **drop\_labels** – if True, does not include ‘label’ column in Table

Returns: a W&B Table of preprocessed samples with labels and playable audio



tools for subsetting and resampling collections

```
opensoundscape.data_selection.resample(df, n_samples_per_class, upsample=True,
                                       downsample=True, with_replace=False, random_state=None)
```

resample a one-hot encoded label df for a target n\_samples\_per\_class

### Parameters

- **df** – dataframe with one-hot encoded labels: columns are classes, index is sample name/path
- **n\_samples\_per\_class** – target number of samples per class
- **upsample** – if True, duplicate samples for classes with <n samples to get to n samples
- **downsample** – if True, randomly sample classes with >n samples to get to n samples
- **with\_replace** – flag to enable sampling of the same row more than once, default False
- **random\_state** – passed to np.random calls. If None, random state is not fixed.

Note: The algorithm assumes that the label df is single-label. If the label df is multi-label, some classes can end up over-represented.

Note 2: The resulting df will have samples ordered by class label, even if the input df had samples in a random order.

```
opensoundscape.data_selection.upsample(input_df, label_column='Labels',
                                       with_replace=False, random_state=None)
```

Given an input DataFrame of categorical labels, upsample to maximum value

Upsampling removes the class imbalance in your dataset. Rows for each label are repeated up to *max\_count // rows*. Then, we randomly sample the rows to fill up to *max\_count*.

The input df is NOT one-hot encoded in this case, but instead contains categorical labels in a specified label\_columns

### Parameters

- **input\_df** – A DataFrame to upsample
- **label\_column** – The column to draw unique labels from
- **flag to enable sampling of the same row more than once, default False** (*with\_replace*) –
- **random\_state** – Set the random\_state during sampling

**Returns** An upsampled DataFrame

**Return type** df

Preprocessors: `pd.Series` child with an action sequence & forward method

**class** `opensoundscape.ml.datasets.AudioFileDataset` (*samples*, *preprocessor*, *by-pass\_augmentations=False*)

Base class for audio datasets with OpenSoundscape (use in place of torch Dataset)

Custom Dataset classes should subclass this class or its children.

Datasets in OpenSoundscape contain a Preprocessor object which is responsible for the procedure of generating a sample for a given input. The DataLoader handles a dataframe of samples (and potentially labels) and uses a Preprocessor to generate samples from them.

### Parameters

- **samples** – the files to generate predictions for. Can be: - a dataframe with index containing audio paths, OR - a dataframe with multi-index of (path,start\_time,end\_time) per clip, OR - a list or `np.ndarray` of audio file paths

### Notes for input dataframe:

- df must have audio paths in the index.
- If label\_df has labels, the class names should be the columns, and

### the values of each row should be 0 or 1.

- If data does not have labels, label\_df will have no columns

- **preprocessor** – an object of BasePreprocessor or its children which defines the operations to perform on input samples

**Returns** sample (AudioSample object)

**Raises** PreprocessingError if exception is raised during `__getitem__`

### Effects:

**self.invalid\_samples** will contain a set of paths that did not successfully produce a list of clips with start/end times, if `split_files_into_clips=True`

**class\_counts** ()  
count number of each label

**head** (*n=5*)  
out-of-place copy of first *n* samples  
performs `df.head(n)` on `self.label_df`

**Parameters**

- **n** – number of first samples to return, see `pandas.DataFrame.head()`
- [**default** – 5]

**Returns** a new dataset object

**sample** (*\*\*kwargs*)  
out-of-place random sample  
creates copy of object with *n* rows randomly sampled from `label_df`  
Args: see `pandas.DataFrame.sample()`

**Returns** a new dataset object

**class** `opensoundscape.ml.datasets.AudioSplittingDataset` (*samples*, *preprocessor*,  
*overlap\_fraction=0*, *final\_clip=None*)

class to load clips of longer files rather than one sample per file

Internally creates even-lengthed clips split from long audio files.

If file labels are provided, applies copied labels to all clips from a file

NOTE: If you've already created a dataframe with clip start and end times, you can use `AudioFileDataset`. This class is only necessary if you wish to automatically split longer files into clips (providing only the file paths).

**Parameters** `AudioFileDataset` and `make_clip_df` (*see*) –

---

## CAM (Class Activation Maps)

---

Class activation maps (CAM) for OpenSoundscape models

```
class opensoundscape.ml.cam.CAM (base_image, activation_maps=None, gbp_maps=None)
    Object to hold and view Class Activation Maps, including guided backprop

    Stores activation maps as .activation_maps, and guided backprop as .gbp_cams
    each is a Series indexed by class

    #TODO: implement plotting multiple classes, each a different color basically, create greyscale images, then
    convert each one to a different color from color cyler getting transparency right might be tricky though

    plot (target_class=None, mode='activation', show_base=True, alpha=0.5, cmap='inferno', interpola-
        tion='bilinear', figsize=None, plt_show=True, save_path=None)
        Plot the activation map, guided back propogation, or their product :param target_class: which class's maps
        to visualize

        • must be in the index of self.gbp_map / self.activation_maps
        • note that the class None is created when classes are not specified

    during CNN.generate_cams() [default: None]
```

### Parameters

- **mode** – str selecting which maps to visualize, one of: 'activation' [default]: overlay activation map 'backprop': overlay guided back propogation result 'backprop\_and\_activation': overlay product of both maps None: do not overlay anything
- **show\_base** – if False, does not plot the image of the original sample [default: True]
- **alpha** – opacity of the activation map overlap [default: 0.5]
- **cmap** – matplotlib colormap for the activation map [default: 'jet']
- **interpolation** – the interpolation method for the activation map [default: bilinear] see matplotlib.pyplot.imshow()
- **figsize** – the figure size for the plot [default: None]

- **plt\_show** – if True, runs plt.show() [default: True]
- **save\_path** – path to save image to [default: None does not save file]

**Returns** (fig, ax) of matplotlib figure

Note: if base\_image does not have 3 channels, channels are averaged then copied across 3 RGB channels to create a greyscale image



loss function classes to use with opensoundscape models

```
class opensoundscape.ml.loss.BCEWithLogitsLoss_hot
    use pytorch's nn.BCEWithLogitsLoss for one-hot labels by simply converting y from long to float

class opensoundscape.ml.loss.CrossEntropyLoss_hot
    use pytorch's nn.CrossEntropyLoss for one-hot labels by converting labels from 1-hot to integer labels
    throws a ValueError if labels are not one-hot

class opensoundscape.ml.loss.ResampleLoss (class_freq,                reduction='mean',
   loss_weight=1.0)

opensoundscape.ml.loss.binary_cross_entropy (pred,    label,    weight=None,    reduc-
   tion='mean', avg_factor=None)

    helper function for BCE loss in ResampleLoss class

opensoundscape.ml.loss.reduce_loss (loss, reduction)
    Reduce loss as specified.
```

#### Parameters

- **loss** (*Tensor*) – Elementwise loss tensor.
- **reduction** (*str*) – Options are “none”, “mean” and “sum”.

**Returns** Reduced loss tensor.

**Return type** Tensor

```
opensoundscape.ml.loss.weight_reduce_loss (loss,    weight=None,    reduction='mean',
   avg_factor=None)
```

Apply element-wise weight and reduce loss.

#### Parameters

- **loss** (*Tensor*) – Element-wise loss.
- **weight** (*Tensor*) – Element-wise weights.
- **reduction** (*str*) – Same as built-in losses of PyTorch.

- **avg\_factor** (*float*) – Average factor when computing the mean of losses.

**Returns** Processed loss values.

**Return type** Tensor

Safe Dataset

---

Dataset wrapper to handle errors gracefully in Preprocessor classes

A SafeDataset handles errors in a potentially misleading way: If an error is raised while trying to load a sample, the SafeDataset will instead load a different sample. The indices of any samples that failed to load will be stored in `._invalid_indices`.

The behavior may be desirable for training a model, but could cause silent errors when predicting a model (replacing a bad file with a different file), and you should always be careful to check for `._invalid_indices` after using a SafeDataset.

based on an implementation by @msamogh in nonechucks ([github.com/msamogh/nonechucks/](https://github.com/msamogh/nonechucks/))

```
class opensoundscape.ml.safe_dataset.SafeDataset (dataset,    invalid_sample_behavior,
  eager_eval=False)
```

A wrapper for a Dataset that handles errors when loading samples

WARNING: When iterating, will skip the failed sample, but when using within a DataLoader, finds the next good sample and uses it for the current index (see `__getitem__`).

Note that this class does not subclass DataSet. Instead, it contains a `.dataset` attribute that is a DataSet (or a Preprocessor, which subclasses DataSet).

#### Parameters

- **dataset** – a torch Dataset instance or child such as a Preprocessor
- **eager\_eval** – If True, checks if every file is able to be loaded during initialization (logs `_valid_indices` and `_invalid_indices`)

Attributes: `_valid_indices` and `_invalid_indices` can be accessed later to check which samples raised Exceptions. `_invalid_samples` is a set of all index values for samples that raised Exceptions.

`__build_index()`

tries to load each sample, logs `_valid_indices` and `_invalid_indices`

`__getitem__(index)`

If loading an index fails, keeps trying the next index until success

`__safe_get_item()`

Tries to load a sample, returns None if error occurs

**is\_index\_built**

Returns True if all indices of the original dataset have been classified into `_valid_indices` or `_invalid_indices`.

**num\_samples\_examined**

count of `_valid_indices` + `_invalid_indices`

**report** (*log=None*)

write `_invalid_samples` to log file, give warning, & return `_invalid_samples`

Class for holding information on a single sample

```
class opensoundscape.sample.AudioSample (source, start_time=None, duration=None, labels=None, trace=None)
```

A class containing information about a single audio sample

`self.preprocessing_exception` is initialized as `None` and will contain the exception raised during preprocessing if any exception occurs

**categorical\_labels**

list of indices with value==1 in `self.labels`

**end\_time**

calculate sample end time as `start_time + duration`

**classmethod from\_series** (*labels\_series*)

initialize `AudioSample` from a pandas Series (optionally containing labels)

- if series name (dataframe index) is tuple, extracts ['file', 'start\_time', 'end\_time']

these values to (`source`, `start_time`, `duration=end_time-start_time`) - otherwise, series name extracted as `source`; `start_time` and `duration` will be none

Extracts `source` (file), `start_time`, and `end_time` from multi-index `pd.Series` (one row of a `pd.DataFrame` with multi index ['file', 'start\_time', 'end\_time']). The argument *series* is saved as `self.labels`. Creates an `AudioSample` object.

**Parameters labels** – a `pd.Series` with name = file path or ['file', 'start\_time', 'end\_time'] and index as classes with 0/1 values as labels. Labels can have no values (just a name) if sample does not have labels.

```
class opensoundscape.sample.Sample (data=None)
```

Class for holding information on a single sample

a `Sample` in `OpenSoundscape` contains information about a single sample, for instance its data and labels

Subclass this class to create `Samples` of specific types

`opensoundscape.sample.collate_samples` (*samples*)

generate batched tensors of data and labels (in a dictionary)

returns collated samples: a dictionary with keys “samples” and “labels”

assumes that `s.data` is a Tensor and `s.labels` is a list/array for each sample `S`

#### Parameters

- **samples** – iterable of AudioSample objects (or other objects)
- **attributes .data as Tensor and .labels as list/array** (*with*)  
–

#### Returns

dictionary of { “samples”:batched tensor of samples, “labels”: batched tensor of labels,  
}

classes for strategically sampling within a DataLoader

**class** opensoundscape.ml.sampling.**ClassAwareSampler** (*labels, num\_samples\_cls=1*)

In each batch of samples, pick a limited number of classes to include and give even representation to each class

**class** opensoundscape.ml.sampling.**ImbalancedDatasetSampler** (*dataset, indices=None, num\_samples=None, call-back\_get\_label=None*)

Samples elements randomly from a given list of indices for imbalanced dataset :param indices: a list of indices

:type indices: list, optional :param num\_samples: number of samples to draw :type num\_samples: int, optional

:param callback\_get\_label func: a callback-like function which takes two arguments:

dataset and index

Based on Imbalanced Dataset Sampling by davinnovation (<https://github.com/ufoym/imbalanced-dataset-sampler>)





`opensoundscape.metrics.multi_target_metrics(targets, scores, class_names, threshold)`  
generate various metrics for a set of scores and labels (targets)

**Parameters**

- **targets** – 0/1 labels in 2d array
- **scores** – continuous values in 2d array
- **class\_names** – list of strings
- **threshold** – scores  $\geq$  threshold result in prediction of 1, while scores  $<$  threshold result in prediction of 0

**Returns** dictionary of various overall and per-class metrics

**Return type** metrics\_dict

`opensoundscape.metrics.predict_multi_target_labels(scores, threshold)`

Generate boolean multi-target predicted labels from continuous scores

For each sample, each class score is compared to a threshold. Any class can be predicted 1 or 0, independent of other classes.

This function internally uses torch.Tensors to optimize performance

Note: threshold can be a single value or list of per-class thresholds

**Parameters**

- **scores** – 2d np.array, 2d list, 2d torch.Tensor, or pd.DataFrame containing continuous scores
- **threshold** – a number or list of numbers with a threshold for each class - if a single number, used as a threshold for all classes (columns) - if a list, length should match number of columns in scores. Each value in the list will be used as a threshold for each respective class (column).

Returns: 1/0 values with 1 if score exceeded threshold and 0 otherwise

See also: `predict_single_target_labels`

`opensoundscape.metrics.predict_single_target_labels(scores)`

Generate boolean single target predicted labels from continuous scores

For each row, the single highest scoring class will be labeled 1 and all other classes will be labeled 0.

This function internally uses `torch.Tensors` to optimize performance

**Parameters** `scores` – 2d `np.array`, 2d list, 2d `torch.Tensor`, or `pd.DataFrame` containing continuous scores

Returns: boolean value where each row has 1 for the highest scoring class and 0 for all other classes. Returns same datatype as input.

See also: `predict_multi_target_labels`

`opensoundscape.metrics.single_target_metrics(targets, scores)`

generate various metrics for a set of scores and labels (targets)

Predicts 1 for the highest scoring class per sample and 0 for all other classes.

**Parameters**

- **targets** – 0/1 labels in 2d array
- **scores** – continuous values in 2d array

**Returns** dictionary of various overall and per-class metrics

**Return type** `metrics_dict`

---

### Image Augmentation

---

Transforms and augmentations for PIL.Images

`opensoundscape.preprocess.img_augment.time_split` (*img*, *seed=None*)

Given a PIL.Image, split into left/right parts and swap

Randomly chooses the slicing location For example, if *h* chosen

**abcdefghijklmnop** ^

hijklmnop + abcdefg

**Parameters** *img* – A PIL.Image

**Returns** A PIL.Image



Actions for augmentation and preprocessing pipelines

This module contains Action classes which act as the elements in Preprocessor pipelines. Action classes have `go()`, `on()`, `off()`, and `set()` methods. They take a single sample of a specific type and return the transformed or augmented sample, which may or may not be the same type as the original.

See the preprocessor module and Preprocessing tutorial for details on how to use and create your own actions.

**class** `opensoundscape.preprocess.actions.Action` (*fn, is\_augmentation=False, \*\*kwargs*)  
 Action class for an arbitrary function

The function must take the sample as the first argument

Note that this allows two use cases: (A) regular function that takes an input object as first argument

eg. `Audio.from_file(path, **kwargs)`

(B) method of a class, which takes ‘self’ as the first argument, eg. `Spectrogram.bandpass(self, **kwargs)`

Other arguments are an arbitrary list of kwargs.

**class** `opensoundscape.preprocess.actions.AudioClipLoader` (*\*\*kwargs*)  
 Action to load clips from an audio file

Loads an audio file or part of a file to an Audio object. Will load entire audio file if `_start_time` and `_end_time` are None. If `_start_time` and `_end_time` are provided, loads the audio only in the specified interval.

see `Audio.from_file()` for documentation.

**Parameters** `Audio.from_file()` (*see*) –

**class** `opensoundscape.preprocess.actions.AudioTrim` (*\*\*kwargs*)  
 Action to trim/extend audio to desired length

**Parameters** `actions.trim_audio` (*see*) –

**class** `opensoundscape.preprocess.actions.BaseAction`  
 Parent class for all Actions (used in Preprocessor pipelines)

New actions should subclass this class.

Subclasses should set `self.requires_labels = True` if `go()` expects (X,y) instead of (X). y is a row of a dataframe (a `pd.Series`) with index (.name) = original file path, columns=class names, values=labels (0,1). X is the sample, and can be of various types (path, Audio, Spectrogram, Tensor, etc). See `Overlay` for an example of an Action that uses labels.

**set** ( *\*\*kwargs* )  
only allow keys that exist in self.params

**class** `opensoundscape.preprocess.actions.Overlay` (*is\_augmentation=True, \*\*kwargs*)  
Action Class for augmentation that overlays samples on eachother

**Required Args:** `overlay_df`: dataframe of audio files (index) and labels to use for overlay `update_labels` (bool): if True, labels of sample are updated to include

labels of overlaid sample

See `overlay()` for other arguments and default values.

**class** `opensoundscape.preprocess.actions.SpectrogramToTensor` (*colormap=None, invert=False*)

Action to create Tensor of desired shape from Spectrogram

calls `.to_image` on `sample.data`, which should be type `Spectrogram`

exposes `invert` argument in `self.params`

**go** (*sample*)  
converts `sample.data` from `Spectrogram` to `Tensor`

`opensoundscape.preprocess.actions.audio_add_noise` (*audio, noise\_dB=-30, signal\_dB=0, color='white'*)

Generates noise and adds to audio object

#### Parameters

- **audio** – an Audio object
- **noise\_dB** – number or range: dBFS of noise signal generated - if number, creates noise with `dB` dBFS level - if (min,max) tuple, chooses noise `dBFS` randomly from range with a uniform distribution
- **signal\_dB** – dB (decibels) gain to apply to the incoming Audio before mixing with noise [default: -3 dB] - like `noise_dB`, can specify (min,max) tuple to use random uniform choice in range

Returns: Audio object with noise added

`opensoundscape.preprocess.actions.audio_random_gain` (*audio, dB\_range=(-30, 0), clip\_range=(-1, 1)*)

Applies a randomly selected gain level to an Audio object

Gain is selected from a uniform distribution in the range `dB_range`

#### Parameters

- **audio** – an Audio object
- **dB\_range** – (min,max) decibels of gain to apply - dB gain applied is chosen from a uniform random distribution in this range

Returns: Audio object with gain applied

`opensoundscape.preprocess.actions.frequency_mask` (*tensor*, *max\_masks=3*,  
*max\_width=0.2*)

add random horizontal bars over Tensor

#### Parameters

- **tensor** – input Torch.tensor sample
- **max\_masks** – max number of horizontal bars [default: 3]
- **max\_width** – maximum size of horizontal bars as fraction of sample height

**Returns** augmented tensor

`opensoundscape.preprocess.actions.image_to_tensor` (*img*, *greyscale=False*)

Convert PIL image to RGB or greyscale Tensor (PIL.Image -> Tensor)

convert PIL.Image w/range [0,255] to torch Tensor w/range [0,1]

#### Parameters

- **img** – PIL.Image
- **greyscale** – if False, converts image to RGB (3 channels). If True, converts image to one channel.

`opensoundscape.preprocess.actions.overlay` (*sample*, *overlay\_df*, *update\_labels*,  
*overlay\_class=None*, *overlay\_prob=1*,  
*max\_overlay\_num=1*, *overlay\_weight=0.5*)

iteratively overlay 2d samples on top of eachother

Overlays (blends) image-like samples from *overlay\_df* on top of the sample with probability *overlay\_prob* until stopping condition. If necessary, trims overlay audio to the length of the input audio.

**Overlays can be used in a few general ways:**

1. a separate df where any file can be overlayed (*overlay\_class=None*)
2. **same df as training, where the overlay class is “different” ie**, does not contain overlapping labels with the original sample
3. **same df as training, where samples from a specific class are used** for overlays

#### Parameters

- **sample** – AudioSample with .labels: labels of the original sample and .preprocessor: the preprocessing pipeline
- **overlay\_df** – a labels dataframe with audio files as the index and classes as columns
- **update\_labels** – if True, add overlayed sample’s labels to original sample
- **overlay\_class** – how to choose files from *overlay\_df* to overlay Options [default: “different”]: None - Randomly select any file from *overlay\_df* “different” - Select a random file from *overlay\_df* containing none  
of the classes this file contains  
specific class name - always choose files from this class
- **overlay\_prob** – the probability of applying each subsequent overlay
- **max\_overlay\_num** – the maximum number of samples to overlay on original - for example, if *overlay\_prob* = 0.5 and *max\_overlay\_num*=2,  
1/2 of samples will recieve 1 overlay and 1/4 will recieve an additional second overlay

- **overlay\_weight** – a float > 0 and < 1, or a list of 2 floats [min, max] between which the weight will be randomly chosen. e.g. [0.1,0.7] An overlay\_weight <0.5 means more emphasis on original sample.

**Returns** overlaid sample, (possibly updated) labels

`opensoundscape.preprocess.actions.scale_tensor(tensor, input_mean=0.5, input_std=0.5)`

linear scaling of tensor values using torch.transforms.Normalize

(Tensor->Tensor)

WARNING: This does not perform per-image normalization. Instead, it takes as arguments a fixed u and s, ie for the entire dataset, and performs  $X=(X-\text{input\_mean})/\text{input\_std}$ .

#### Parameters

- **input\_mean** – mean of input sample pixels (average across dataset)
- **input\_std** – standard deviation of input sample pixels (average across dataset)
- **are NOT the target mu and sd, but the original mu and sd of img (these) –**
- **which the output will have mu=0, std=1) (for) –**

**Returns** modified tensor

`opensoundscape.preprocess.actions.tensor_add_noise(tensor, std=1)`

Add gaussian noise to sample (Tensor -> Tensor)

**Parameters** **std** – standard deviation for Gaussian noise [default: 1]

Note: be aware that scaling before/after this action will change the effect of a fixed stdev Gaussian noise

`opensoundscape.preprocess.actions.time_mask(tensor, max_masks=3, max_width=0.2)`

add random vertical bars over sample (Tensor -> Tensor)

#### Parameters

- **tensor** – input Torch.tensor sample
- **max\_masks** – maximum number of vertical bars [default: 3]
- **max\_width** – maximum size of bars as fraction of sample width

**Returns** augmented tensor

`opensoundscape.preprocess.actions.torch_color_jitter(tensor, brightness=0.3, contrast=0.3, saturation=0.3, hue=0)`

Wraps torchvision.transforms.ColorJitter

(Tensor -> Tensor) or (PIL Img -> PIL Img)

#### Parameters

- **tensor** – input sample
- **brightness=0.3** –
- **contrast=0.3** –
- **saturation=0.3** –
- **hue=0** –

**Returns** modified tensor



`opensoundscape.preprocess.actions.torch_random_affine` (*tensor*, *degrees=0*, *translate=(0.3, 0.1)*, *fill=0*)

Wraps for torchvision.transforms.RandomAffine

(Tensor -> Tensor) or (PIL Img -> PIL Img)

#### Parameters

- **tensor** – torch.Tensor input sample
- **= 0** (*degrees*) –
- **=** (*translate*) –
- **= 0-255, duplicated across channels** (*fill*) –

**Returns** modified tensor

Note: If applying per-image normalization, we recommend applying RandomAffine after image normalization. In this case, an intermediate gray value is ~0. If normalization is applied after RandomAffine on a PIL image, use an intermediate fill color such as (122,122,122).

`opensoundscape.preprocess.actions.trim_audio` (*sample*, *extend=True*, *random\_trim=False*, *tol=1e-05*)

trim audio clips (Audio -> Audio)

Trims an audio file to desired length Allows audio to be trimmed from start or from a random time Optionally extends audio shorter than clip\_length with silence

#### Parameters

- **sample** – AudioSample with .data=Audio object, .duration as sample duration
- **extend** – if True, clips shorter than sample.duration are extended with silence to required length
- **random\_trim** – if True, chooses a random segment of length sample.duration from the input audio. If False, the file is trimmed from 0 seconds to sample.duration seconds.
- **tol** – tolerance for considering a clip to be of the correct length (sec)

**Returns** trimmed audio



---

## Preprocessors

---

Preprocessor classes: tools for preparing and augmenting audio samples

**class** `opensoundscape.preprocess.preprocessors.BasePreprocessor` (*sample\_duration=None*)  
 Class for defining an ordered set of Actions and a way to run them

Custom Preprocessor classes should subclass this class or its children

Preprocessors have one job: to transform samples from some input (eg a file path) to some output (eg an AudioSample with .data as torch.Tensor) using a specific procedure defined by the .pipeline attribute. The procedure consists of Actions ordered by the Preprocessor's .pipeline. Preprocessors have a forward() method which sequentially applies the Actions in the pipeline to produce a sample.

### Parameters

- **action\_dict** – dictionary of name:Action actions to perform sequentially
- **sample\_duration** – length of audio samples to generate (seconds)

**forward** (*sample*, *break\_on\_type=None*, *break\_on\_key=None*, *bypass\_augmentations=False*, *trace=False*)

perform actions in self.pipeline on a sample (until a break point)

Actions with .bypass = True are skipped. Actions with .is\_augmentation = True can be skipped by passing bypass\_augmentations=True.

### Parameters

- **sample** – either: - pd.Series with file path as index (.name) and labels - OR a file path as pathlib.Path or string
- **break\_on\_type** – if not None, the pipeline will be stopped when it reaches an Action of this class. The matching action is not performed.
- **break\_on\_key** – if not None, the pipeline will be stopped when it reaches an Action whose index equals this value. The matching action is not performed.
- **clip\_times** – can be either - None: the file is treated as a single sample - dictionary {"start\_time":float,"end\_time":float}:

the start and end time of clip in audio

- **bypass\_augmentations** – if True, actions with `.is_augmentatino=True` are skipped
- **trace** (*boolean – default False*) – if True, saves the output of each pipeline step in the `sample_info` output argument - should be utilized for analysis/debugging on samples of interest

**Returns** sample (instance of AudioSample class)

**insert\_action** (*action\_index, action, after\_key=None, before\_key=None*)

insert an action in specific specific position

This is an in-place operation

Inserts a new action before or after a specific key. If `after_key` and `before_key` are both None, action is appended to the end of the index.

**Parameters**

- **action\_index** – string key for new action in index
- **action** – the action object, must be subclass of BaseAction
- **after\_key** – insert the action immediately after this key in index
- **before\_key** – insert the action immediately before this key in index Note: only one of (`after_key`, `before_key`) can be specified

**remove\_action** (*action\_index*)

alias for `self.drop(..., inplace=True)`, removes an action

This is an in-place operation

**Parameters** **action\_index** – index of action to remove

**class** opensoundscape.preprocess.preprocessors.**SpectrogramPreprocessor** (*sample\_duration, over-  
lay\_df=None, out\_shape=(224, 224, 3)*)

Child of BasePreprocessor that creates spectrogram Tensors w/augmentation

loads audio, creates spectrogram, performs augmentations, creates tensor

by default, does not resample audio, but bandpasses to 0-11.025 kHz (to ensure all outputs have same scale in y-axis) can change with `.pipeline.bandpass.set(min_f=,max_f=)`

during prediction, will load clips from long audio files rather than entire audio files.

**Parameters**

- **sample\_duration** – length in seconds of audio samples generated If not None, longer clips trimmed to this length. By default, shorter clips will be extended (modify `random_trim_audio` and `trim_audio` to change behavior).
- **overlay\_df** – if not None, will include an overlay action drawing samples from this df
- **out\_shape** – output shape of tensor h,w,channels [default: (224,224,3)]

Utilities for preprocessing

**exception** `opensoundscape.preprocess.utils.PreprocessingError`

Custom exception indicating that a Preprocessor pipeline failed

`opensoundscape.preprocess.utils.get_args(func)`

get list of arguments and default values from a function

`opensoundscape.preprocess.utils.get_reqd_args(func)`

get list of required arguments and default values from a function

`opensoundscape.preprocess.utils.show_tensor(tensor, channel=None, transform_from_zero_centered=True, invert=True)`

helper function for displaying a sample as an image

### Parameters

- **tensor** – torch.Tensor of shape [c,w,h] with values centered around zero
- **channel** – specify an integer to plot only one channel, otherwise will attempt to plot all channels
- **transform\_from\_zero\_centered** – if True, transforms values from [-1,1] to [0,1]
- **invert** – if true, flips value range via  $x=1-x$

`opensoundscape.preprocess.utils.show_tensor_grid(tensors, columns, channel=None, transform_from_zero_centered=True, invert=True, labels=None)`

create image of nxn tensors

### Parameters

- **tensors** – list of samples
- **columns** – number of columns in grid

- **labels** – title of each subplot
- **other args**, see `show_tensor()` (*for*) –

Augmentations and transforms for torch.Tensors

`opensoundscape.preprocess.tensor_augment.freq_mask(spec, F=30, max_masks=3, replace_with_zero=False)`  
draws horizontal bars over the image

### Parameters

- **spec** – a torch.Tensor representing a spectrogram
- **F** – maximum frequency-width of bars in pixels
- **max\_masks** – maximum number of bars to draw
- **replace\_with\_zero** – if True, bars are 0s, otherwise, mean img value

**Returns** Augmented tensor

`opensoundscape.preprocess.tensor_augment.time_mask(spec, T=40, max_masks=3, replace_with_zero=False)`  
draws vertical bars over the image

### Parameters

- **spec** – a torch.Tensor representing a spectrogram
- **T** – maximum time-width of bars in pixels
- **max\_masks** – maximum number of bars to draw
- **replace\_with\_zero** – if True, bars are 0s, otherwise, mean img value

**Returns** Augmented tensor





Detect periodic vocalizations with RIBBIT

This module provides functionality to search audio for periodically fluctuating vocalizations.

`opensoundscape.ribbit.calculate_pulse_score` (*amplitude*, *amplitude\_sample\_rate*,  
*pulse\_rate\_range*, *plot=False*, *nfft=1024*)

Search for amplitude pulsing in an audio signal in a range of pulse repetition rates (PRR)

scores an audio amplitude signal by highest value of power spectral density in the PRR range

#### Parameters

- **amplitude** – a time series of the audio signal’s amplitude (for instance a smoothed raw audio signal)
- **amplitude\_sample\_rate** – sample rate in Hz of amplitude signal, normally ~20-200 Hz
- **pulse\_rate\_range** – [min, max] values for amplitude modulation in Hz
- **plot=False** – if True, creates a plot visualizing the power spectral density
- **nfft=1024** – controls the resolution of the power spectral density (see `scipy.signal.welch`)

**Returns** pulse rate score for this audio segment (float)

`opensoundscape.ribbit.ribbit` (*spectrogram*, *signal\_band*, *pulse\_rate\_range*, *clip\_duration*,  
*clip\_overlap=0*, *final\_clip=None*, *noise\_bands=None*, *plot=False*)

Run RIBBIT detector to search for periodic calls in audio

Searches for periodic energy fluctuations at specific repetition rates and frequencies.

#### Parameters

- **spectrogram** – `opensoundscape.Spectrogram` object of an audio file
- **signal\_band** – [min, max] frequency range of the target species, in Hz
- **pulse\_rate\_range** – [min,max] pulses per second for the target species

- **clip\_duration** – the length of audio (in seconds) to analyze at one time - each clip is analyzed independently and receives a ribbit score
- **clip\_overlap** (*float*) – overlap between consecutive clips (sec)
- **final\_clip** (*str*) – behavior if final clip is less than clip\_duration seconds long. By default, discards remaining audio if less than clip\_duration seconds long [default: None]. Options: - None: Discard the remainder (do not make a clip) - “remainder”: Use only remainder of Audio (final clip will be shorter than clip\_duration)
- **“full”**: Increase overlap with previous clip to yield a clip with clip\_duration length

Note that the “extend” option is not supported for RIBBIT.

- **noise\_bands** – list of frequency ranges to subtract from the signal\_band For instance: [ [min1,max1] , [min2,max2] ] - if *None*, no noise bands are used - default: None
- **plot=False** – if True, plot the power spectral density for each clip

**Returns** DataFrame with columns ['start\_time', 'end\_time', 'score'], with a row for each clip.

## Notes

\_\_PARAMETERS\_\_ RIBBIT requires the user to select a set of parameters that describe the target vocalization. Here is some detailed advice on how to use these parameters.

**Signal Band:** The signal band is the frequency range where RIBBIT looks for the target species. It is best to pick a narrow signal band if possible, so that the model focuses on a specific part of the spectrogram and has less potential to include erroneous sounds.

**Noise Bands:** Optionally, users can specify other frequency ranges called noise bands. Sounds in the *noise\_bands* are *subtracted* from the *signal\_band*. Noise bands help the model filter out erroneous sounds from the recordings, which could include confusion species, background noise, and popping/clicking of the microphone due to rain, wind, or digital errors. It's usually good to include one noise band for very low frequencies – this specifically eliminates popping and clicking from being registered as a vocalization. It's also good to specify noise bands that target confusion species. Another approach is to specify two narrow *noise\_bands* that are directly above and below the *signal\_band*.

**Pulse Rate Range:** This parameters specifies the minimum and maximum pulse rate (the number of pulses per second, also known as pulse repetition rate) RIBBIT should look for to find the focal species. For example, choosing *pulse\_rate\_range* = [10, 20] means that RIBBIT should look for pulses no slower than 10 pulses per second and no faster than 20 pulses per second.

**Clip Duration:** The *clip\_duration* parameter tells RIBBIT how many seconds of audio to analyze at one time. Generally, you should choose a *clip\_length* that is similar to the length of the target species vocalization, or a little bit longer. For very slowly pulsing vocalizations, choose a longer window so that at least 5 pulses can occur in one window (0.5 pulses per second -> 10 second window). Typical values for are 0.3 to 10 seconds. Also, *clip\_overlap* can be used for overlap between sequential clips. This is more computationally expensive but will be more likely to center a target sound in the clip (with zero overlap, the target sound may be split up between adjacent clips).

**Plot:** We can choose to show the power spectrum of pulse repetition rate for each window by setting *plot=True*. The default is not to show these plots (*plot=False*).

\_\_ALGORITHM\_\_ This is the procedure RIBBIT follows: divide the audio into segments of length clip\_duration for each clip:

- calculate time series of energy in signal band (signal\_band) and subtract noise band energies (noise\_bands) - calculate power spectral density of the amplitude time series - score the file based on the max value of power spectral density in the pulse rate range



Signal Processing

---

Signal processing tools for feature extraction and more

```
opensoundscape.signal_processing.cwt_peaks(audio, center_frequency, wavelet='morl',
   peak_threshold=0.2, peak_separation=None,
   plot=False)
```

compute a cwt, post-process, then extract peaks

Performs a continuous wavelet transform (cwt) on an audio signal at a single frequency. It then squares, smooths, and normalizes the signal. Finally, it detects peaks in the resulting signal and returns the times and magnitudes of detected peaks. It is used as a feature extractor for Ruffed Grouse drumming detection.

#### Parameters

- **audio** – an Audio object
- **center\_frequency** – the target frequency to extract peaks from
- **wavelet** – (str) name of a pywt wavelet, eg ‘morl’ (see pywt docs)
- **peak\_threshold** – minimum height of peaks - if None, no minimum peak height - see “height” argument to `scipy.signal.find_peaks`
- **peak\_separation** – minimum time between detected peaks, in seconds - if None, no minimum distance - see “distance” argument to `scipy.signal.find_peaks`

**Returns** list of times (from beginning of signal) of each peak `peak_levels`: list of magnitudes of each detected peak

**Return type** `peak_times`

---

**Note:** consider downsampling audio to reduce computational cost. Audio must have sample rate of at least 2x target frequency.

---

```
opensoundscape.signal_processing.detect_peak_sequence_cwt(audio,
   sample_rate=400,
   window_len=60,
   center_frequency=50,
   wavelet='morl',
   peak_threshold=0.2,
   peak_separation=0.0375,
   dt_range=(0.05, 0.8),
   dy_range=(-0.2, 0),
   d2y_range=(-0.05,
   0.15),
   max_skip=3,
   duration_range=(1,
   15),
   points_range=(9,
   100),
   plot=False)
```

Use a continuous wavelet transform to detect accelerating sequences

This function creates a continuous wavelet transform (cwt) feature and searches for accelerating sequences of peaks in the feature. It was developed to detect Ruffed Grouse drumming events in audio signals. Default parameters are tuned for Ruffed Grouse drumming detection.

Analysis is performed on analysis windows of fixed length without overlap. Detections from each analysis window across the audio file are aggregated.

#### Parameters

- **audio** – Audio object
- **sample\_rate=400** – resample audio to this sample rate (Hz)
- **window\_len=60** – length of analysis window (sec)
- **center\_frequency=50** – target audio frequency of cwt
- **wavelet='morl'** – (str) pywt wavelet name (see pywavelets docs)
- **peak\_threshold=0.2** – height threshold (0-1) for peaks in normalized signal
- **peak\_separation=15/400** – min separation (sec) for peak finding
- **dt\_range=(0.05, 0.8)** – sequence detection point-to-point criterion 1 - Note: the upper limit is also used as sequence termination criterion 2
- **dy\_range=(-0.2, 0)** – sequence detection point-to-point criterion 2
- **d2y\_range=(-0.05, 0.15)** – sequence detection point-to-point criterion 3
- **max\_skip=3** – sequence termination criterion 1: max sequential invalid points
- **duration\_range=(1, 15)** – sequence criterion 1: length (sec) of sequence
- **points\_range=(9, 100)** – sequence criterion 2: num points in sequence
- **plot=False** – if True, plot peaks and detected sequences with pyplot

**Returns** dataframe summarizing detected sequences

Note: for Ruffed Grouse drumming, which is very low pitched, audio is resampled to 400 Hz. This greatly increases the efficiency of the cwt, but will only detect frequencies up to  $400/2=200$ Hz. Generally, choose a resample frequency as low as possible but  $\geq 2$ x the target frequency

Note: the cwt signal is normalized on each analysis window, so changing the analysis window size can change the detection results.

Note: if there is an incomplete window remaining at the end of the audio file, it is discarded (not analyzed).

```
opensoundscape.signal_processing.find_accel_sequences (t, dt_range=(0.05, 0.8), dy_range=(-0.2, 0), d2y_range=(-0.05, 0.15), max_skip=3, duration_range=(1, 15), points_range=(5, 100))
```

detect accelerating/decelerating sequences in time series

developed for detecting Ruffed Grouse drumming events in a series of peaks extracted from cwt signal

The algorithm computes the forward difference of  $t$ ,  $y(t)$ . It iterates through the  $[y(t), t]$  points searching for sequences of points that meet a set of conditions. It begins with an empty candidate sequence.

“Point-to-point criteria”: Valid ranges for  $dt$ ,  $dy$ , and  $d2y$  are checked for each subsequent point and are based on previous points in the candidate sequence. If they are met, the point is added to the candidate sequence.

“Continuation criteria”: Conditions for  $max\_skip$  and the upper bound of  $dt$  are used to determine when a sequence should be terminated.

- $max\_skip$ : max number of sequential invalid points before terminating
- $dt \leq dt\_range[1]$ : if  $dt$  is long, sequence should be broken

“Sequence criteria”: When a sequence is terminated, it is evaluated on conditions for  $duration\_range$  and  $points\_range$ . If it meets these conditions, it is saved as a detected sequence.

- $duration\_range$ : length of sequence in seconds from first to last point
- $points\_range$ : number of points included in sequence

When a sequence is terminated, the search continues with the next point and an empty sequence.

#### Parameters

- $t$  – (list or `np.array`) times of all detected peaks (seconds)
- $dt\_range = (0.05, 0.8)$  – valid values for  $t(i) - t(i-1)$
- $dy\_range = (-0.2, 0)$  – valid values for change in  $y$  (grouse: difference in time between consecutive beats should decrease)
- $d2y\_range = (-0.05, 15)$  – limit change in  $dy$ : should not show large decrease (sharp curve downward on  $y$  vs  $t$  plot)
- $max\_skip=3$  – max invalid points between valid points for a sequence (grouse: should not have many noisy points between beats)
- $duration\_range = (1, 15)$  – total duration of sequence (sec)
- $points\_range = (9, 100)$  – total number of points in sequence

**Returns** lists of  $t$  and  $y$  for each detected sequence

**Return type** `sequences_t, sequences_y`

```
opensoundscape.signal_processing.frequency2scale (frequency, wavelet, sample_rate)
```

determine appropriate wavelet scale for desired center frequency

#### Parameters

- **frequency** – desired center frequency of wavelet in Hz (1/seconds)
- **wavelet** – (str) name of pywt wavelet, eg ‘morl’ for Morlet
- **sample\_rate** – sample rate in Hz (1/seconds)

**Returns** (float) scale parameter for `pywt.cwt()` to extract desired frequency

**Return type** scale

Note: this function is not exactly an inverse of `pywt.scale2frequency()`, because that function returns frequency in sample-units (cycles/sample) rather than frequency in Hz (cycles/second). In other words, `frequency_hz = pywt.scale2frequency(w,scale)*sr`.

`opensoundscape.signal_processing.gcc(x, y, cc_filter='phat', epsilon=0.001, radius=None)`

Generalized cross correlation of two signals

Computes a generalized cross correlation in frequency response.

The generalized cross correlation algorithm described in Knapp and Carter [1].

In the case of `cc_filter='cc'`, gcc simplifies to cross correlation and is equivalent to `scipy.signal.correlate` and `numpy.correlate`.

code adapted from [github.com/axeber01/ngcc](https://github.com/axeber01/ngcc)

**Parameters**

- **x** – 1d numpy array of audio samples
- **y** – 1d numpy array of audio samples
- **cc\_filter** – which filter to use in the gcc. 'phat' - Phase transform. Default. 'roth' - 'scot' - Smoothed Coherence Transform, 'ht' - Hannan and Thomson 'cc' - normal cross correlation with no filter
- **epsilon** – small value used to ensure denominator when applying a filter is non-zero.

**Returns** 1d numpy array of gcc values

**Return type** gcc

see also: `tdoa()` uses this function to estimate time delay between two signals

[1] Knapp, C.H. and Carter, G.C (1976) The Generalized Correlation Method for Estimation of Time Delay. IEEE Trans. Acoust. Speech Signal Process, 24, 320-327. <http://dx.doi.org/10.1109/TASSP.1976.1162830>

`opensoundscape.signal_processing.tdoa(signal, reference_signal, cc_filter='phat', sample_rate=1, return_max=False, max_delay=None)`

estimate time difference of arrival between two signals

estimates time delay by finding the maximum of the generalized cross correlation (gcc) of two signals. The two signals are discrete-time series with the same sample rate.

For example, if the signal arrives 2.5 seconds `_after_` the reference signal, returns 2.5; if it arrives 0.5 seconds `_before_` the reference signal, returns -0.5.

**Parameters**

- **reference\_signal** (*signal*,) – np.arrays or lists containing each signal
- **cc\_filter** – see `gcc()`
- **sample\_rate** – sample rate (Hz) of signals; both signals must have same sample rate
- **return\_max** – if True, returns the maximum value of the generalized cross correlation
- **max\_delay** – maximum possible tdoa. Cross-correlations that correspond to time delays outside of this range are ignored. For example, if `max_delay=0.5`, the tdoa returned will be the delay between -0.5 and +0.5 seconds, that maximizes the cross-correlation. This is useful if you know the maximum possible delay between the two signals, and want to ignore any tdoas outside of that range. e.g. if receivers are 100m apart, and the speed of sound is 340m/s, then the maximum possible delay is 0.294 seconds.



### Returns

estimated delay from reference signal to signal, in seconds (note that default sampling rate is 1.0 samples/second)

if `return_max` is `True`, returns a second value, the maximum value of the result of generalized cross correlation

See also: `gcc()` if you want the raw output of generalized cross correlation

`opensoundscape.signal_processing.thresholded_event_durations` (*x*, *threshold*, *normalize=False*, *sample\_rate=1*)

Detect positions and durations of events over threshold in 1D signal

This function takes a 1D numeric vector and searches for segments that are continuously greater than a threshold value. The input signal can optionally be normalized, and if a sample rate is provided the start positions will be in the units of  $[sr]^{-1}$  (ie if *sr* is Hz, start positions will be in seconds).

### Parameters

- **x** – 1d input signal, a vector of numeric values
- **threshold** – minimum value of signal to be a detection
- **normalize** – if `True`, performs  $x=x/\max(x)$
- **sample\_rate** – sample rate of input signal

**Returns** start time of each detected event durations: duration (# samples/sr) of each detected event

**Return type** start\_times



a set of utilities for converting between scientific and common names of bird species in different naming systems (xeno canto and bird net)

`opensoundscape.taxa.bn_common_to_sci` (*common*)

convert bird net common name (ignoring dashes, spaces, case) to lowercase-hyphenated name

`opensoundscape.taxa.common_to_sci` (*common*)

convert bird net common name to scientific name as lowercase-hyphenated

Ignores dashes, spaces, case

`opensoundscape.taxa.get_species_list` ()

returns a list of scientific-names (lowercase-hyphenated) of species

`opensoundscape.taxa.sci_to_bn_common` (*scientific*)

convert scientific lowercase-hyphenated name to birdnet common name as lowercasenospaces

`opensoundscape.taxa.sci_to_xc_common` (*scientific*)

convert scientific lowercase-hyphenated name to xeno-canto lowercasenospaces common name

`opensoundscape.taxa.xc_common_to_sci` (*common*)

convert xeno-canto common name (ignoring dashes/spaces/case) to lowercase-hyphenated name



## Localization

Tools for localizing audio events from synchronized recording arrays

**exception** `opensoundscape.localization.InsufficientReceiversError`  
 raised when there are not enough receivers to localize an event

**class** `opensoundscape.localization.SpatialEvent` (*receiver\_files*, *receiver\_positions*,  
*start\_time=0*, *duration=None*,  
*class\_name=None*, *band-*  
*pass\_range=None*, *cc\_threshold=None*,  
*max\_delay=None*)

Class that estimates the position of a single sound event

Uses receiver positions and time-of-arrival of sounds to estimate sound source position

**calculate\_distance\_residuals** (*speed\_of\_sound=343*)  
 calculate distance residuals for each receiver from tdoa localization

The residual represents the discrepancy between (difference in distance of each receiver to estimated position) and (observed tdoa), and has units of meters.

**Parameters** `speed_of_sound` – speed of sound in m/s

**Returns** array of residuals (units are meters), one for each receiver

**Effects:** stores the returned residuals as `self.distance_residuals`

**calculate\_residual\_rms** (*speed\_of\_sound=343*)  
 calculate the root mean square distance residual from tdoa localization

**Parameters** `speed_of_sound` – speed of sound in meters per second

**Returns**

root mean square value of residuals, in meters - returns None if `self.tdoas` or `self.position_estimate`

are None

See also: *SpatialEvent.calculate\_distance\_residuals()*

**estimate\_delays** (*bandpass\_range=None, cc\_filter='phat', max\_delay=None*)  
 estimate time delay of event relative to receiver\_files[0] with gcc

**Performs Generalized Cross Correlation of each file against the first**, extracting the segment of audio of length self.duration at self.start\_time

Assumes audio files are synchronized such that they start at the same time

#### Parameters

- **bandpass\_range** – bandpass audio to [low, high] frequencies in Hz before cross correlation; if None, defaults to self.bandpass\_range
- **cc\_filter** – filter for generalized cross correlation, see opensoundscape.signal\_processing.gcc()
- **max\_delay** – only consider values in +/- this range (seconds) for possible delay (see opensoundscape.signal\_processing.tdoa()) - default None allows any delay time

#### Returns

list of time delays, list of max cross correlation values

each list is the same length as self.receiver\_files, and each value corresponds to the cross correlation of one file relative to the first file (self.receiver\_files[0])

**Effects:** sets self.tdoas and self.cc\_maxs with the same values as those returned

**estimate\_location** (*algorithm='gillette', cc\_threshold=None, min\_n\_receivers=3, speed\_of\_sound=343*)  
 estimate spatial location of this event

uses self.tdoas and self.receiver\_positions to estimate event location

Note: if self.tdoas or self.receiver\_positions is None, first calls self.estimate\_delays() to estimate these values

Localization is performed in 2d or 3d according to the dimensions of self.receiver\_positions (x,y) or (x,y,z)

#### Parameters

- **algorithm** – 'gillette' or 'soundfinder', see localization.localize()
- **cc\_threshold** – see SpatialEvent documentation
- **min\_n\_receivers** – if number of receivers with cross correlation exceeding the threshold is fewer than this, raises InsufficientReceiversError instead of estimating a spatial position

**Returns** meters)

**Return type** position estimate as cartesian coordinates (x,y) or (x,y,z) (units

**Raises** *InsufficientReceiversError* if the number of receivers with cross correlation – maximums exceeding *cc\_threshold* is less than *min\_n\_receivers*

**Effects:** sets the value of self.position\_estimate to the same value as the returned position

**class** opensoundscape.localization.**SynchronizedRecorderArray** (*file\_coords*)  
 Class with utilities for localizing sound events from array of recorders

### **localize\_detections()**

Attempt to localize a sound event for each detection of each class. First, creates candidate events with: `create_candidate_events()`

Create `SpatialEvent` objects for all simultaneous, spatially clustered detections of a class

Then, attempts to localize each candidate event via time delay of arrival information: For each candidate event:

- calculate relative time of arrival with generalized cross correlation (`event.estimate_delays()`)
- **if enough cross correlation values exceed a threshold, attempt to localize the event** using the time delays and spatial positions of each receiver with `event.estimate_location()`
- **if the residual distance rms value is below a cutoff threshold, consider the event** to be successfully localized

### **create\_candidate\_events** (*detections, min\_n\_receivers, max\_receiver\_dist*)

Takes the detections dictionary and groups detections that are within *max\_receiver\_dist* of each other.  
:param detections: a dictionary of detections, with multi-index (file,start\_time,end\_time), and

one column per class with 0/1 values for non-detection/detection The times in the index imply the same real world time across all files: eg 0 seconds assumes that the audio files all started at the same time, not on different dates/times

#### **Parameters**

- **min\_n\_receivers** – if fewer nearby receivers have a simultaneous detection, do not create candidate event
- **max\_receiver\_dist** – the maximum distance between recorders to consider a detection as a single event

**Returns** a list of `SpatialEvent` objects to attempt to localize

**localize\_detections** (*detections, max\_receiver\_dist, min\_n\_receivers=3, localization\_algorithm='gillette', cc\_threshold=0, cc\_filter='phat', max\_delay=None, bandpass\_ranges=None, residual\_threshold=<sphinx.ext.autodoc.importer.\_MockObject object>*)

Attempt to localize positions for all detections

The user provides a table of class detections from each recorder with timestamps. The object's `self.file_coords` dataframe contains a table listing the spatial location of the recorder for each unique audio file in the table of detections. The audio recordings must be synchronized such that timestamps from each recording correspond to the exact same real-world time.

Localization of sound events proceeds in four steps:

1. Grouping of detections into candidate events (`self.create_candidate_events()`):

Simultaneous and spatially clustered detections of a class are selected as targets for localization of a single real-world sound event.

For each detection of a species, the grouping algorithm treats the receiver with the detection as a “reference receiver”, then selects all detections of the species at the same time and within *max\_receiver\_dist* of the reference receiver (the “surrounding detections”). This selected group of simultaneous, spatially-clustered detections of a class becomes one “candidate event” for subsequent localization.

If the number of recorders in the candidate event is fewer than *min\_n\_receivers*, the candidate event is discarded.

This step creates a highly redundant set of candidate events to localize, because each detection is treated separately with its recorder as the ‘reference recorder’. Thus, the localized events created by this algorithm may contain multiple instances representing the same real-world sound event.

## 2. Estimate time delays with cross correlation:

For each candidate event, the time delay between the reference receiver’s detection and the surrounding recorders’ detections is estimated through generalized cross correlation.

If `bandpass_ranges` are provided, cross correlation is performed on audio that has been bandpassed to class-specific low and high frequencies.

If the max value of the cross correlation is below `cc_threshold`, the corresponding time delay is discarded and not used during localization. This provides a way of filtering out undesired time delays that do not correspond to two recordings of the same sound event.

If the number of estimated time delays in the candidate event is fewer than `min_n_receivers` after filtering by cross correlation threshold, the candidate event is discarded.

## 3. Estimate positions

The position of the event is estimated based on the positions and time delays of each detection.

Position estimation from the positions and time delays at a set of receivers is performed using one of two algorithms, described in *localization\_algorithm* below.

## 4. Filter by spatial residual error

The residual errors represent discrepancies between (a) time of arrival of the event at a receiver and (b) distance from receiver to estimated position.

Estimated positions are discarded if the root mean squared spatial residual is greater than `residual_rms_threshold`

### Parameters

- **detections** – a dictionary of detections, with multi-index (file,start\_time,end\_time), and one column per class with 0/1 values for non-detection/detection. The times in the index imply the same real world time across all files: eg 0 seconds assumes that the audio files all started at the same time, not on different dates/times
- **max\_receiver\_dist** – float (meters) Radius around a recorder in which to use other recorders for localizing an event. Simultaneous detections at receivers within this distance (meters) of a receiver with a detection will be used to attempt to localize the event.
- **min\_n\_receivers** – int Minimum number of receivers that must detect an event for it to be localized [default: 3]
- **localization\_algorithm** – str, optional algorithm to use for estimating the position of a sound event from the positions and time delays of a set of detections. [Default: ‘gillette’] Options:
  - ‘gillette’: linear closed-form algorithm of Gillette and Silverman 2008 [1]
  - ‘soundfinder’: GPS position algorithm of Wilson et al. 2014 [2]
- **cc\_threshold** – float, optional Threshold for cross correlation: if the max value of the cross correlation is below this value, the corresponding time delay is discarded and not used during localization. Default of 0 does not discard any delays.



- **cc\_filter** – str, optional Filter to use for generalized cross correlation. See `signalprocessing.gcc` function for options. Default is “phat”.
- **max\_delay** – float, optional Maximum absolute value of time delay estimated during cross correlation of two signals For instance, 0.2 means that cross correlation will be maximized in the range of delays between -0.2 to 0.2 seconds. Default: None does not restrict the range, finding delay that maximizes cross correlation
- **bandpass\_ranges** – dict, optional Dictionary of form {“class name”: [low\_f, high\_f]} for audio bandpass filtering during cross correlation. [Default: None] does not bandpass audio. Bandpassing audio to the frequency range of the relevant sound is recommended for best cross correlation results.
- **residual\_threshold** – discard localized events if the root mean squared residual exceeds this value (distance in meters) [default: np.inf does not filter out any events by residual]

**Returns** list of localized events, list of un-localized events events are of class `SpatialEvent`

**Return type** 2 lists

[1] M. D. Gillette and H. F. Silverman, “A Linear Closed-Form Algorithm for Source Localization From Time-Differences of Arrival,” *IEEE Signal Processing Letters*

[2] Wilson, David R., Matthew Battiston, John Brzustowski, and Daniel J. Mennill. “Sound Finder: A New Software Approach for Localizing Animals Recorded with a Microphone Array.” *Bioacoustics* 23, no. 2 (May 4, 2014): 99–112. <https://doi.org/10.1080/09524622.2013.827588>.

**make\_nearby\_files\_dict** (*r\_max*)

create dictionary listing nearby files for each file

pre-generate a dictionary listing all close files for each audio file dictionary will have a key for each audio file, and value listing all other receivers within *r\_max* of that receiver

eg {ARU\_0.mp3: [ARU\_1.mp3, ARU\_2.mp3...], ARU\_1... }

Note: could manually create this dictionary to only list `_simultaneous_` nearby files if the detection dataframe contains files from different times

**The returned dictionary is used in `create_candidate_events` as a look-up table for** recordings nearby a detection in any given file

**Parameters** *r\_max* – maximum distance from each recorder in which to include other recorders in the list of ‘nearby recorders’, in meters

**Returns** dictionary with keys for each file and values = list of nearby recordings

`opensoundscape.localization.calc_speed_of_sound` (*temperature=20*)

Calculate speed of sound in air, in meters per second

Calculate speed of sound for a given temperature in Celsius (Humidity has a negligible effect on speed of sound and so this functionality is not implemented)

**Parameters** *temperature* – ambient air temperature in Celsius

**Returns** the speed of sound in air in meters per second

`opensoundscape.localization.calculate_tdoa_residuals` (*receiver\_positions*,  
*tdoas*, *position\_estimate*,  
*speed\_of\_sound*)

Calculate the residual distances of the TDOA localization algorithm

The residual represents the discrepancy between (difference in distance of each receiver to estimated position) and (observed tdoa), and has units of meters. Residuals are calculated as follows:

**expected = calculated time difference of arrival between reference and** another receiver, based on the positions of the receivers and estimated event position

**observed** = observed tdoas provided to localization algorithm

**residual time** = expected - observed (in seconds)

**residual distance** = speed of sound \* residual time (in meters)

#### Parameters

- **receiver\_position** – The list of coordinates (in m) of each receiver, as [x,y] for 2d or [x,y,z] for 3d.
- **tdoas** – List of time delays of arrival for the sound at each receiver, relative to the first receiver in the list (tdoas[0] should be 0)
- **position\_estimate** – The estimated position of the sound, as (x,y) or (x,y,z) in meters
- **speed\_of\_sound** – The speed of sound in m/s

**Returns** np.array containing the residuals in units of meters, one per receiver

`opensoundscape.localization.gillette_localize(receiver_positions, arrival_times, speed_of_sound=343)`

Uses the Gillette and Silverman [1] localization algorithm to localize a sound event from a set of TDOAs. :param receiver\_positions: a list of [x,y] or [x,y,z] positions for each receiver

Positions should be in meters, e.g., the UTM coordinate system.

#### Parameters

- **arrival\_times** – a list of TDOA times (arrival times) for each receiver The times should be in seconds.
- **speed\_of\_sound** – speed of sound in m/s

**Returns** a tuple of (x,y,z) coordinates of the sound source

**Return type** coords

Algorithm from: [1] M. D. Gillette and H. F. Silverman, “A Linear Closed-Form Algorithm for Source Localization From Time-Differences of Arrival,” IEEE Signal Processing Letters

`opensoundscape.localization.localize(receiver_positions, tdoas, algorithm, speed_of_sound=343)`

Perform TDOA localization on a sound event. :param receiver\_positions: a list of [x,y,z] positions for each receiver

Positions should be in meters, e.g., the UTM coordinate system.

#### Parameters

- **tdoas** – a list of TDOA times (onset times) for each recorder The times should be in seconds.
- **speed\_of\_sound** – speed of sound in m/s
- **algorithm** – the algorithm to use for localization Options: ‘soundfinder’, ‘gillette’

**Returns** The estimated source position in meters.

`opensoundscape.localization.lorentz_ip(u, v=None)`

Compute Lorentz inner product of two vectors

For vectors  $u$  and  $v$ , the Lorentz inner product for 3-dimensional case is defined as

$$u[0]*v[0] + u[1]*v[1] + u[2]*v[2] - u[3]*v[3]$$

Or, for 2-dimensional case as

$$u[0]*v[0] + u[1]*v[1] - u[2]*v[2]$$

#### Parameters

- **u** – vector with shape either (3,) or (4,)
- **v** – vector with same shape as **u**; if None (default), sets  $v = u$

**Returns** value of Lorentz IP

**Return type** float

`opensoundscape.localization.soundfinder_localize(receiver_positions, arrival_times, speed_of_sound=343, invert_alg='gps', center=True, pseudo=True)`

Use the soundfinder algorithm to perform TDOA localization on a sound event. Localize a sound event given relative arrival times at multiple receivers. This function implements a localization algorithm from the equations described in [1]. Localization can be performed in a global coordinate system in meters (i.e., UTM), or relative to recorder positions in meters.

This implementation follows [2] with corresponding variable names.

#### Parameters

- **receiver\_positions** – a list of [x,y,z] positions for each receiver. Positions should be in meters, e.g., the UTM coordinate system.
- **arrival\_times** – a list of TDOA times (onset times) for each recorder. The times should be in seconds.
- **of sound (speed)** – speed of sound in m/s
- **invert\_alg** – what inversion algorithm to use (only 'gps' is implemented)
- **center** – whether to center recorders before computing localization result. Computes localization relative to centered plot, then translates solution back to original recorder locations. (For behavior of original Sound Finder, use True)
- **pseudo** – whether to use the pseudorange error (True) or sum of squares discrepancy (False) to pick the solution to return (For behavior of original Sound Finder, use False. However, in initial tests, pseudorange error appears to perform better.)

**Returns** The solution (x,y,z) in meters.

[1] Wilson, David R., Matthew Battiston, John Brzustowski, and Daniel J. Mennill. "Sound Finder: A New Software Approach for Localizing Animals Recorded with a Microphone Array." *Bioacoustics* 23, no. 2 (May 4, 2014): 99–112. <https://doi.org/10.1080/09524622.2013.827588>.

[2] Global Positioning Systems handout, 2002 <http://web.archive.org/web/20110719232148/http://www.macalester.edu/~halverson/math36/GPS.pdf>

`opensoundscape.localization.travel_time` (*source, receiver, speed\_of\_sound*)

Calculate time required for sound to travel from a source to a receiver

**Parameters**

- **source** – cartesian position [x,y] or [x,y,z] of sound source, in meters
- **receiver** – cartesian position [x,y] or [x,y,z] of sound receiver, in meters
- **speed\_of\_sound** – speed of sound in m/s

**Returns** time in seconds for sound to travel from source to receiver

Utilities for opensoundscape

**exception** `opensoundscape.utils.GetDurationError`  
 raised if `librosa.get_duration(path=f)` causes an error

`opensoundscape.utils.binarize(x, threshold)`  
 return a list of 0, 1 by thresholding vector `x`

`opensoundscape.utils.generate_clip_times_df(full_duration, clip_duration, clip_overlap=0, final_clip=None, rounding_precision=10)`  
 generate start and end times for even-lengthed clips

The behavior for incomplete final clips at the end of the `full_duration` depends on the `final_clip` parameter.

This function only creates a dataframe with start and end times, it does not perform any actual trimming of audio or other objects.

#### Parameters

- **full\_duration** – The amount of time (seconds) to split into clips
- **clip\_duration** (*float*) – The duration in seconds of the clips
- **clip\_overlap** (*float*) – The overlap of the clips in seconds [default: 0]
- **final\_clip** (*str*) – Behavior if `final_clip` is less than `clip_duration` seconds long. By default, discards remaining time if less than `clip_duration` seconds long [default: None]. Options:
  - None: Discard the remainder (do not make a clip)
  - "extend": Extend the final clip beyond `full_duration` to reach `clip_duration` length
  - "remainder": Use only remainder of `full_duration` (final clip will be shorter than `clip_duration`)
  - "full": Increase overlap with previous clip to yield a clip with `clip_duration` length.

Note: returns entire original audio if it is shorter than `clip_duration`

- **rounding\_precision** (*int* or *None*) – number of decimals to round start/end times to - pass None to skip rounding

**Returns** DataFrame with columns for ‘start\_time’ and ‘end\_time’ of each clip

**Return type** clip\_df

`opensoundscape.utils.hex_to_time(s)`

convert a hexadecimal, Unix time string to a datetime timestamp in utc

Example usage: ““ # Get the UTC timestamp t = hex\_to\_time(‘5F16A04E’)

# Convert it to a desired timezone my\_timezone = pytz.timezone(“US/Mountain”) t = t.astimezone(my\_timezone) ““

**Parameters** *s* (*string*) – hexadecimal Unix epoch time string, e.g. ‘5F16A04E’

**Returns** datetime.datetime object representing the date and time in UTC

`opensoundscape.utils.identity(x)`

return the input unchanged

`opensoundscape.utils.inrange(x, r)`

return true if x is in range [r[0],r1] (inclusive)

`opensoundscape.utils.isNan(x)`

check for nan by equating x to itself

`opensoundscape.utils.jitter(x, width, distribution=‘gaussian’)`

Jitter (add random noise to) each value of x

**Parameters**

- **x** – scalar, array, or nd-array of numeric type
- **width** – multiplier for random variable (stdev for ‘gaussian’ or r for ‘uniform’)
- **distribution** – ‘gaussian’ (default) or ‘uniform’ if ‘gaussian’: draw jitter from gaussian with mu = 0, std = width if ‘uniform’: draw jitter from uniform on [-width, width]

**Returns** x + random jitter

**Return type** jittered\_x

`opensoundscape.utils.linear_scale(array, in_range=(0, 1), out_range=(0, 255))`

Translate from range in\_range to out\_range

**Inputs:** in\_range: The starting range [default: (0, 1)] out\_range: The output range [default: (0, 255)]

**Outputs:** new\_array: A translated array

`opensoundscape.utils.load_metadata(path, raise_exceptions=False)`

use soundfile to load metadata from WAV or AIFF file

**Parameters**

- **path** – file path to WAV or AIFF file
- **raise\_exceptions** – if True, raises exception, if False returns None if exception occurs [default: False]

**Returns** dictionary containing audio file metadata

`opensoundscape.utils.make_clip_df(files, clip_duration, clip_overlap=0, final_clip=None, return_invalid_samples=False, raise_exceptions=False)`

generate df of fixed-length clip start/end times for a set of files

Used internally to prepare a dataframe listing clips of longer audio files

This function creates a single dataframe with audio files as the index and columns: 'start\_time', 'end\_time'. It will list clips of a fixed duration from the beginning to end of each audio file.

Note: if a label dataframe is passed as *files*, the labels for each file will be copied to all clips having the corresponding file. If the label dataframe contains multiple rows for a single file, the labels in the *\_first\_* row containing the file path are used as labels for resulting clips.

#### Parameters

- **files** – list of audio file paths, or dataframe with file path as index - if dataframe, columns represent classes and values represent class labels. Labels for a file will be copied to all clips belonging to that file in the returned clip dataframe.
- **clip\_duration** (*float*) – see `generate_clip_times_df`
- **clip\_overlap** (*float*) – see `generate_clip_times_df`
- **final\_clip** (*str*) – see `generate_clip_times_df`
- **return\_invalid\_samples** (*bool*) – if True, returns additional value, a list of samples that caused exceptions
- **raise\_exceptions** (*bool*) – if True, if exceptions are raised when attempting to check the duration of an audio file, the exception will be raised. If False [default], adds a row to the dataframe with `np.nan` for 'start\_time' and 'end\_time' for that file path.

#### Returns

**dataframe multi-index ('file','start\_time','end\_time')**

- if files is a dataframe, will contain same columns as files
- otherwise, will have no columns

if `return_invalid_samples==True`, returns (clip\_df, invalid\_samples)

**Return type** clip\_df

**Note: default behavior for raise\_exceptions is the following:** if an exception is raised (for instance, trying to get the duration of the file), the dataframe will have one row with `np.nan` for 'start\_time' and 'end\_time' for that file path.

`opensoundscape.utils.min_max_scale(array, feature_range=(0, 1))`  
rescale vaues in an a array linearly to feature\_range

`opensoundscape.utils.overlap(r1, r2)`  
“calculate the amount of overlap between two real-numbered ranges  
ranges must be [low,high] where low <= high

`opensoundscape.utils.overlap_fraction(r1, r2)`  
“calculate the fraction of r1 (low, high range) that overlaps with r2

`opensoundscape.utils.rescale_features(X, rescaling_vector=None)`  
rescale all features by dividing by the max value for each feature

optionally provide the rescaling vector (1xlen(X) np.array), so that you can rescale a new dataset consistently with an old one

returns rescaled feature set and rescaling vector

`opensoundscape.utils.sigmoid(x)`  
sigmoid function





## CHAPTER 37

---

Index

---



## CHAPTER 38

---

### Modules

---

- modindex



### O

- `opensoundscape.annotations`, 143
- `opensoundscape.aru`, 149
- `opensoundscape.data_selection`, 161
- `opensoundscape.localization`, 203
- `opensoundscape.logging`, 159
- `opensoundscape.metrics`, 175
- `opensoundscape.ml.cam`, 165
- `opensoundscape.ml.cnn`, 135
- `opensoundscape.ml.cnn_architectures`, 153
- `opensoundscape.ml.datasets`, 163
- `opensoundscape.ml.loss`, 167
- `opensoundscape.ml.safe_dataset`, 169
- `opensoundscape.ml.sampling`, 173
- `opensoundscape.ml.utils`, 151
- `opensoundscape.preprocess.actions`, 179
- `opensoundscape.preprocess.img_augment`, 177
- `opensoundscape.preprocess.preprocessors`, 185
- `opensoundscape.preprocess.tensor_augment`, 189
- `opensoundscape.preprocess.utils`, 187
- `opensoundscape.ribbit`, 191
- `opensoundscape.sample`, 171
- `opensoundscape.signal_processing`, 195
- `opensoundscape.spectrogram`, 129
- `opensoundscape.taxa`, 201
- `opensoundscape.utils`, 211



## Symbols

`__getitem__()` (*opensoundscape.ml.safe\_dataset.SafeDataset* method), 169

`_build_index()` (*opensoundscape.ml.safe\_dataset.SafeDataset* method), 169

`_safe_get_item()` (*opensoundscape.ml.safe\_dataset.SafeDataset* method), 169

## A

*Action* (class in *opensoundscape.preprocess.actions*), 179

*alexnet()* (in module *opensoundscape.ml.cnn\_architectures*), 153

*amplitude()* (*opensoundscape.spectrogram.Spectrogram* method), 131

*apply\_activation\_layer()* (in module *opensoundscape.ml.utils*), 151

*audio\_add\_noise()* (in module *opensoundscape.preprocess.actions*), 180

*audio\_random\_gain()* (in module *opensoundscape.preprocess.actions*), 180

*audio\_sample\_rate* (*opensoundscape.spectrogram.Spectrogram* attribute), 131

*AudioClipLoader* (class in *opensoundscape.preprocess.actions*), 179

*AudioFileDataset* (class in *opensoundscape.ml.datasets*), 163

*audiomoth\_start\_time()* (in module *opensoundscape.aru*), 149

*AudioSample* (class in *opensoundscape.sample*), 171

*AudioSplittingDataset* (class in *opensoundscape.ml.datasets*), 164

*AudioTrim* (class in *opensoundscape.preprocess.actions*), 179

## B

*bandpass()* (*opensoundscape.annotations.BoxedAnnotations* method), 143

*bandpass()* (*opensoundscape.spectrogram.Spectrogram* method), 131

*BaseAction* (class in *opensoundscape.preprocess.actions*), 179

*BaseModule* (class in *opensoundscape.ml.utils*), 151

*BasePreprocessor* (class in *opensoundscape.preprocess.preprocessors*), 185

*BCEWithLogitsLoss\_hot* (class in *opensoundscape.ml.loss*), 167

*binarize()* (in module *opensoundscape.utils*), 211

*binary\_cross\_entropy()* (in module *opensoundscape.ml.loss*), 167

*bn\_common\_to\_sci()* (in module *opensoundscape.taxa*), 201

*BoxedAnnotations* (class in *opensoundscape.annotations*), 143

## C

*calc\_speed\_of\_sound()* (in module *opensoundscape.localization*), 207

*calculate\_distance\_residuals()* (*opensoundscape.localization.SpatialEvent* method), 203

*calculate\_pulse\_score()* (in module *opensoundscape.ribbit*), 191

*calculate\_residual\_rms()* (*opensoundscape.localization.SpatialEvent* method), 203

*calculate\_tdoa\_residuals()* (in module *opensoundscape.localization*), 207

*CAM* (class in *opensoundscape.ml.cam*), 165

*cas\_data\_loader()* (in module *opensoundscape.ml.utils*), 151

*categorical\_labels* (*opensoundscape* attribute), 151

*scape.sample.AudioSample* attribute), 171  
*categorical\_to\_one\_hot()* (in module *opensoundscape.annotations*), 147  
*change\_conv2d\_channels()* (in module *opensoundscape.ml.cnn\_architectures*), 154  
*change\_fc\_output\_size()* (in module *opensoundscape.ml.cnn\_architectures*), 154  
*class\_counts()* (*opensoundscape.ml.datasets.AudioFileDataset* method), 163  
*ClassAwareSampler* (class in *opensoundscape.ml.sampling*), 173  
*CNN* (class in *opensoundscape.ml.cnn*), 135  
*collate\_samples()* (in module *opensoundscape.sample*), 171  
*common\_to\_sci()* (in module *opensoundscape.taxa*), 201  
*convert\_labels()* (*opensoundscape.annotations.BoxedAnnotations* method), 143  
*create\_candidate\_events()* (*opensoundscape.localization.SynchronizedRecorderArray* method), 205  
*CrossEntropyLoss\_hot* (class in *opensoundscape.ml.loss*), 167  
*cwt\_peaks()* (in module *opensoundscape.signal\_processing*), 195

## D

*densenet121()* (in module *opensoundscape.ml.cnn\_architectures*), 154  
*detect\_peak\_sequence\_cwt()* (in module *opensoundscape.signal\_processing*), 195  
*diff()* (in module *opensoundscape.annotations*), 147  
*duration* (*opensoundscape.spectrogram.Spectrogram* attribute), 131

## E

*efficientnet\_b0()* (in module *opensoundscape.ml.cnn\_architectures*), 154  
*efficientnet\_b4()* (in module *opensoundscape.ml.cnn\_architectures*), 155  
*efficientnet\_widese\_b0()* (in module *opensoundscape.ml.cnn\_architectures*), 155  
*efficientnet\_widese\_b4()* (in module *opensoundscape.ml.cnn\_architectures*), 155  
*end\_time* (*opensoundscape.sample.AudioSample* attribute), 171  
*estimate\_delays()* (*opensoundscape.localization.SpatialEvent* method), 204  
*estimate\_location()* (*opensoundscape.localization.SpatialEvent* method), 204

*eval()* (*opensoundscape.ml.cnn.CNN* method), 135

## F

*find\_accel\_sequences()* (in module *opensoundscape.signal\_processing*), 196  
*forward()* (*opensoundscape.preprocess.preprocessors.BasePreprocessor* method), 185  
*freeze\_params()* (in module *opensoundscape.ml.cnn\_architectures*), 156  
*freq\_mask()* (in module *opensoundscape.preprocess.tensor\_augment*), 189  
*frequencies* (*opensoundscape.spectrogram.Spectrogram* attribute), 130  
*frequency2scale()* (in module *opensoundscape.signal\_processing*), 197  
*frequency\_mask()* (in module *opensoundscape.preprocess.actions*), 180  
*from\_audio()* (*opensoundscape.spectrogram.MelSpectrogram* class method), 129  
*from\_audio()* (*opensoundscape.spectrogram.Spectrogram* class method), 131  
*from\_raven\_files()* (*opensoundscape.annotations.BoxedAnnotations* class method), 144  
*from\_series()* (*opensoundscape.sample.AudioSample* class method), 171  
*from\_torch\_dict()* (*opensoundscape.ml.cnn.CNN* class method), 136  
*from\_torch\_dict()* (*opensoundscape.ml.cnn.InceptionV3* class method), 140

## G

*gcc()* (in module *opensoundscape.signal\_processing*), 198  
*generate\_cams()* (*opensoundscape.ml.cnn.CNN* method), 136  
*generate\_clip\_times\_df()* (in module *opensoundscape.utils*), 211  
*get\_args()* (in module *opensoundscape.preprocess.utils*), 187  
*get\_batch()* (in module *opensoundscape.ml.utils*), 151  
*get\_reqd\_args()* (in module *opensoundscape.preprocess.utils*), 187  
*get\_species\_list()* (in module *opensoundscape.taxa*), 201  
*GetDurationError*, 211



- `gillette_localize()` (in module `opensoundscape.localization`), 208
- `global_one_hot_labels()` (`opensoundscape.annotations.BoxedAnnotations` method), 144
- `go()` (`opensoundscape.preprocess.actions.SpectrogramToTensor` method), 180
- ## H
- `head()` (`opensoundscape.ml.datasets.AudioFileDataset` method), 164
- `hex_to_time()` (in module `opensoundscape.utils`), 212
- ## I
- `identity()` (in module `opensoundscape.utils`), 212
- `image_to_tensor()` (in module `opensoundscape.preprocess.actions`), 181
- `ImbalancedDatasetSampler` (class in `opensoundscape.ml.sampling`), 173
- `inception_v3()` (in module `opensoundscape.ml.cnn_architectures`), 156
- `InceptionV3` (class in `opensoundscape.ml.cnn`), 140
- `inrange()` (in module `opensoundscape.utils`), 212
- `insert_action()` (`opensoundscape.preprocess.preprocessors.BasePreprocessor` method), 186
- `InsufficientReceiversError`, 203
- `is_index_built` (`opensoundscape.ml.safe_dataset.SafeDataset` attribute), 169
- `isNan()` (in module `opensoundscape.utils`), 212
- ## J
- `jitter()` (in module `opensoundscape.utils`), 212
- ## L
- `limit_db_range()` (`opensoundscape.spectrogram.Spectrogram` method), 132
- `linear_scale()` (in module `opensoundscape.utils`), 212
- `linear_scale()` (`opensoundscape.spectrogram.Spectrogram` method), 132
- `list_architectures()` (in module `opensoundscape.ml.cnn_architectures`), 156
- `load_metadata()` (in module `opensoundscape.utils`), 212
- `load_model()` (in module `opensoundscape.ml.cnn`), 140
- `load_outdated_model()` (in module `opensoundscape.ml.cnn`), 140
- `load_weights()` (`opensoundscape.ml.cnn.CNN` method), 137
- `localize()` (in module `opensoundscape.localization`), 208
- `localize_detections()` (`opensoundscape.localization.SynchronizedRecorderArray` method), 204, 205
- `lorentz_ip()` (in module `opensoundscape.localization`), 209
- ## M
- `make_clip_df()` (in module `opensoundscape.utils`), 212
- `make_nearby_files_dict()` (`opensoundscape.localization.SynchronizedRecorderArray` method), 207
- `MelSpectrogram` (class in `opensoundscape.spectrogram`), 129
- `min_max_scale()` (in module `opensoundscape.utils`), 213
- `min_max_scale()` (`opensoundscape.spectrogram.Spectrogram` method), 132
- `multi_target_metrics()` (in module `opensoundscape.metrics`), 175
- ## N
- `net_amplitude()` (`opensoundscape.spectrogram.Spectrogram` method), 133
- `num_samples_examined` (`opensoundscape.ml.safe_dataset.SafeDataset` attribute), 170
- ## O
- `one_hot_clip_labels()` (`opensoundscape.annotations.BoxedAnnotations` method), 144
- `one_hot_labels_like()` (`opensoundscape.annotations.BoxedAnnotations` method), 145
- `one_hot_labels_on_time_interval()` (in module `opensoundscape.annotations`), 147
- `one_hot_to_categorical()` (in module `opensoundscape.annotations`), 147
- `opensoundscape.annotations` (module), 143
- `opensoundscape.aru` (module), 149
- `opensoundscape.data_selection` (module), 161
- `opensoundscape.localization` (module), 203
- `opensoundscape.logging` (module), 159
- `opensoundscape.metrics` (module), 175
- `opensoundscape.ml.cam` (module), 165
- `opensoundscape.ml.cnn` (module), 135

opensoundscape.ml.cnn\_architectures  
(module), 153

opensoundscape.ml.datasets (module), 163

opensoundscape.ml.loss (module), 167

opensoundscape.ml.safe\_dataset (module),  
169

opensoundscape.ml.sampling (module), 173

opensoundscape.ml.utils (module), 151

opensoundscape.preprocess.actions (mod-  
ule), 179

opensoundscape.preprocess.img\_augment  
(module), 177

opensoundscape.preprocess.preprocessors  
(module), 185

opensoundscape.preprocess.tensor\_augment  
(module), 189

opensoundscape.preprocess.utils (module),  
187

opensoundscape.ribbit (module), 191

opensoundscape.sample (module), 171

opensoundscape.signal\_processing (mod-  
ule), 195

opensoundscape.spectrogram (module), 129

opensoundscape.taxa (module), 201

opensoundscape.utils (module), 211

overlap() (in module opensoundscape.utils), 213

overlap\_fraction() (in module opensound-  
scape.utils), 213

overlap\_samples (opensound-  
scape.spectrogram.Spectrogram attribute),  
131

Overlay (class in opensoundscape.preprocess.actions),  
180

overlay() (in module opensound-  
scape.preprocess.actions), 181

## P

parse\_audiomoth\_metadata() (in module open-  
soundscape.aru), 149

parse\_audiomoth\_metadata\_from\_path() (in  
module opensoundscape.aru), 150

plot() (opensoundscape.ml.cam.CAM method), 165

plot() (opensoundscape.spectrogram.MelSpectrogram  
method), 130

plot() (opensoundscape.spectrogram.Spectrogram  
method), 133

predict() (opensoundscape.ml.cnn.CNN method),  
137

predict\_multi\_target\_labels() (in module  
opensoundscape.metrics), 175

predict\_single\_target\_labels() (in module  
opensoundscape.metrics), 176

PreprocessingError, 187

## R

reduce\_loss() (in module opensoundscape.ml.loss),  
167

register\_arch() (in module opensound-  
scape.ml.cnn\_architectures), 156

remove\_action() (opensound-  
scape.preprocess.preprocessors.BasePreprocessor  
method), 186

report() (opensound-  
scape.ml.safe\_dataset.SafeDataset method),  
170

resample() (in module opensound-  
scape.data\_selection), 161

ResampleLoss (class in opensoundscape.ml.loss), 167

rescale\_features() (in module opensound-  
scape.utils), 213

resnet101() (in module opensound-  
scape.ml.cnn\_architectures), 156

resnet152() (in module opensound-  
scape.ml.cnn\_architectures), 156

resnet18() (in module opensound-  
scape.ml.cnn\_architectures), 157

resnet34() (in module opensound-  
scape.ml.cnn\_architectures), 157

resnet50() (in module opensound-  
scape.ml.cnn\_architectures), 157

ribbit() (in module opensoundscape.ribbit), 191

## S

SafeDataset (class in opensound-  
scape.ml.safe\_dataset), 169

Sample (class in opensoundscape.sample), 171

sample() (opensound-  
scape.ml.datasets.AudioFileDataset method),  
164

save() (opensoundscape.ml.cnn.CNN method), 138

save\_torch\_dict() (opensoundscape.ml.cnn.CNN  
method), 138

save\_weights() (opensoundscape.ml.cnn.CNN  
method), 138

scale\_tensor() (in module opensound-  
scape.preprocess.actions), 182

scaling (opensoundscape.spectrogram.Spectrogram  
attribute), 131

sci\_to\_bn\_common() (in module opensound-  
scape.taxa), 201

sci\_to\_xc\_common() (in module opensound-  
scape.taxa), 201

separate\_resnet\_feat\_clf() (in module open-  
soundscape.ml.cnn), 141

set() (opensoundscape.preprocess.actions.BaseAction  
method), 180

show\_tensor() (in module opensound-  
scape.preprocess.utils), 187

`show_tensor_grid()` (in module `opensoundscape.preprocess.utils`), 187  
`sigmoid()` (in module `opensoundscape.utils`), 213  
`single_target_metrics()` (in module `opensoundscape.metrics`), 176  
`soundfinder_localize()` (in module `opensoundscape.localization`), 209  
`SpatialEvent` (class in `opensoundscape.localization`), 203  
`Spectrogram` (class in `opensoundscape.spectrogram`), 130  
`SpectrogramPreprocessor` (class in `opensoundscape.preprocess.preprocessors`), 186  
`SpectrogramToTensor` (class in `opensoundscape.preprocess.actions`), 180  
`squeezenet1_0()` (in module `opensoundscape.ml.cnn_architectures`), 158  
`subset()` (`opensoundscape.annotations.BoxedAnnotations` method), 146  
`SynchronizedRecorderArray` (class in `opensoundscape.localization`), 204

## T

`tdoa()` (in module `opensoundscape.signal_processing`), 198  
`tensor_add_noise()` (in module `opensoundscape.preprocess.actions`), 182  
`thresholded_event_durations()` (in module `opensoundscape.signal_processing`), 199  
`time_mask()` (in module `opensoundscape.preprocess.actions`), 182  
`time_mask()` (in module `opensoundscape.preprocess.tensor_augment`), 189  
`time_split()` (in module `opensoundscape.preprocess.img_augment`), 177  
`to_image()` (`opensoundscape.spectrogram.Spectrogram` method), 133  
`to_raven_files()` (`opensoundscape.annotations.BoxedAnnotations` method), 146  
`torch_color_jitter()` (in module `opensoundscape.preprocess.actions`), 182  
`torch_random_affine()` (in module `opensoundscape.preprocess.actions`), 182  
`train()` (`opensoundscape.ml.cnn.CNN` method), 139  
`travel_time()` (in module `opensoundscape.localization`), 209  
`trim()` (`opensoundscape.annotations.BoxedAnnotations` method), 146  
`trim()` (`opensoundscape.spectrogram.Spectrogram` method), 133  
`trim_audio()` (in module `opensoundscape.preprocess.actions`), 183

## U

`unique_labels()` (`opensoundscape.annotations.BoxedAnnotations` method), 146  
`upsample()` (in module `opensoundscape.data_selection`), 161  
`use_resample_loss()` (in module `opensoundscape.ml.cnn`), 141

## V

`vgg11_bn()` (in module `opensoundscape.ml.cnn_architectures`), 158

## W

`wandb_table()` (in module `opensoundscape.logging`), 159  
`weight_reduce_loss()` (in module `opensoundscape.ml.loss`), 167  
`window_length` (`opensoundscape.spectrogram.Spectrogram` attribute), 134  
`window_start_times` (`opensoundscape.spectrogram.Spectrogram` attribute), 134  
`window_step` (`opensoundscape.spectrogram.Spectrogram` attribute), 134  
`window_type` (`opensoundscape.spectrogram.Spectrogram` attribute), 131

## X

`xc_common_to_sci()` (in module `opensoundscape.taxa`), 201