Blogs (/blog) / mhellicar's blog (/blog/75) / How to find a Morepork

HOW TO FIND A MOREPORK

In this entry we are very pleased to welcome back Tim Hunt from the Centre for Information Technology at Wintec (https://www.wintec.ac.nz/). He has an update for us on his work in automatically detecting Morepork (Ruru) from audio recordings made by our bird monitors (https://www.2040.co.nz/collections/cacophonometer-bird-monitoring).

In a previous blog (https://www.2040.co.nz/blogs/news/first-morepork-automatically-identified (https://www.2040.co.nz/blogs/news/first-morepork-automatically-identified)), you may have read about how I started down the road of automatically detecting morepork calls. Since then I've made some further progress and thought I'd share the journey so far with you.

Here is an initial look at the numbers of morepork calls that we are starting to see automatically using the model that has been developed. Each colour is for a different recorder/location:

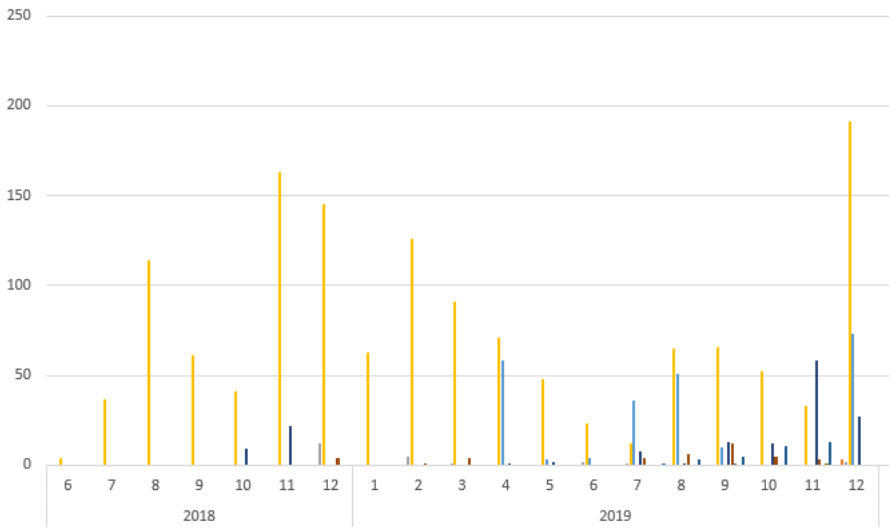

To see if we have found any moreporks on a particular recorder, you can scroll down through your recordings (see pic below) and see the red markers in the Tags column.

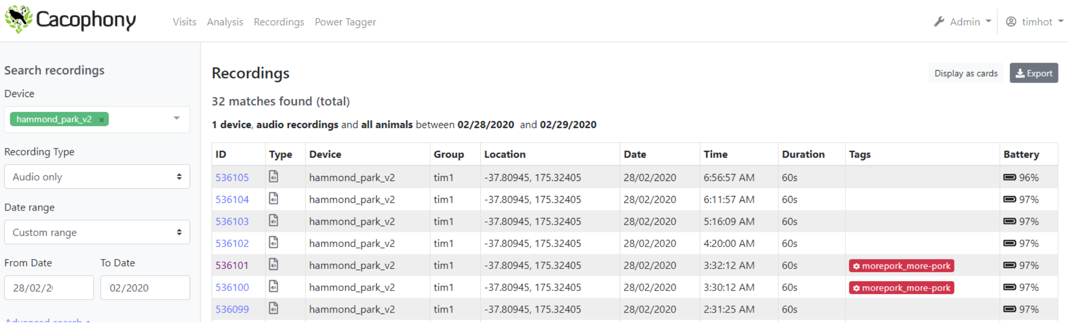

Search recordings

Device: hammond_park_v2 ×

Recording Type: Audio only

Date range: Custom range

From Date: 28/02/2

To Date: 02/2020

Recordings

32 matches found (total)

1 device, **audio recordings** and **all animals** between **02/28/2020** and **02/29/2020**

ID	Type	Device	Group	Location	Date	Time	Duration	Tags	Battery
536105		hammond_park_v2	tim1	-37.80945, 175.32405	28/02/2020	6:56:57 AM	60s		96%
536104		hammond_park_v2	tim1	-37.80945, 175.32405	28/02/2020	6:11:57 AM	60s		97%
536103		hammond_park_v2	tim1	-37.80945, 175.32405	28/02/2020	5:16:09 AM	60s		97%
536102		hammond_park_v2	tim1	-37.80945, 175.32405	28/02/2020	4:20:00 AM	60s		97%
536101		hammond_park_v2	tim1	-37.80945, 175.32405	28/02/2020	3:32:12 AM	60s	morepork_more-pork	97%
536100		hammond_park_v2	tim1	-37.80945, 175.32405	28/02/2020	3:30:12 AM	60s	morepork_more-pork	97%
536099		hammond_park_v2	tim1	-37.80945, 175.32405	28/02/2020	2:31:25 AM	60s		97%

Then click on the ID and you will see the tags listed as below. Note: currently recordings are only processed with the model on an ad-hoc basis.

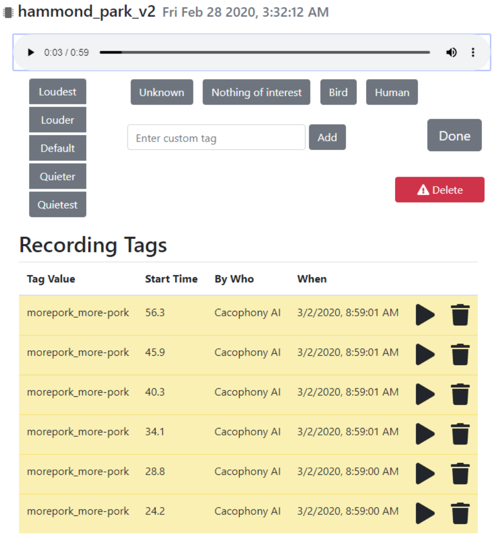

Recording Tags

Tag Value	Start Time	By Who	When
morepork_more-pork	56.3	Cacophony AI	3/2/2020, 8:59:01 AM
morepork_more-pork	45.9	Cacophony AI	3/2/2020, 8:59:01 AM
morepork_more-pork	40.3	Cacophony AI	3/2/2020, 8:59:01 AM
morepork_more-pork	34.1	Cacophony AI	3/2/2020, 8:59:01 AM
morepork_more-pork	28.8	Cacophony AI	3/2/2020, 8:59:00 AM
morepork_more-pork	24.2	Cacophony AI	3/2/2020, 8:59:00 AM

Background

Although I've been involved with IT for many years, and keep hearing about the wonders of machine learning, I hadn't personally created a machine learning application. It is a big and daunting area of knowledge and expertise that I just didn't have. So, as you do, I started watching youtube videos and found some great tutorials on using TensorfFow (https://www.tensorflow.org/ (https://www.tensorflow.org/)) the seemingly go to technique for machine learning. However, I was getting a nagging feeling that I wasn't really understanding what was going on and so was keeping an open mind about alternative ways of doing machine learning. I had also heard about a tool called Weka (https://www.cs.waikato.ac.nz/ml/weka/ (https://www.cs.waikato.ac.nz/ml/weka/)) which happens to have been created at The University of Waikato, right here in my hometown of Hamilton. I thought I should really have a look to see what it was capable of.

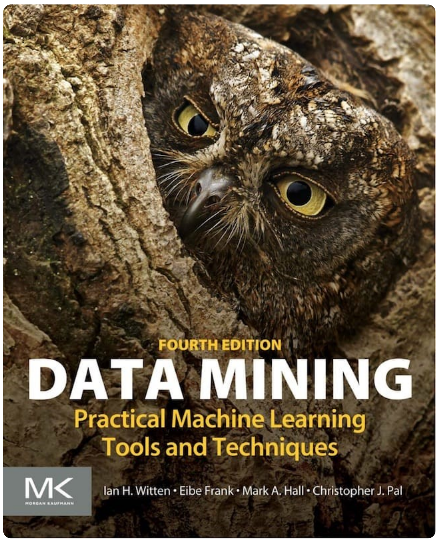
FOURTH EDITION
DATA MINING
Practical Machine Learning
Tools and Techniques
MK
Ian H. Witten • Eibe Frank • Mark A. Hall • Christopher J. Pal

Weka is a 'workbench for machine learning' and comes with a huge amount of helpful resources including a book that has been used my many people who are not necessary maths wizzes but wish to understand and implement machine learning solutions – it even has a picture of an owl 'hiding' in a tree, on the front cover – a good omen! I decided to take a closer look.

Machine learning basics – a simple understanding

Through watching the online tutorials (https://www.youtube.com/user/WekaMOOC (https://www.youtube.com/user/WekaMOOC)) and reading the book I slowly increased my knowledge of machine learning and how to use the Weka workbench. From the perspective of being able to recognize morepork calls, what we need to do is:

1) Find examples of morepork calls (see that previous blog (https://www.2040.co.nz/blogs/news/first-morepork-automatically-identified)),

2) Extract 'features' from the audio of each example (I'll come back to this later),

3) Give the (feature) examples to the machine learning algorithm (I just tried a few different algorithms available in Weka),

4) 'Ask' the algorithm to adjust its internal parameters (that's a bit of a black box) to find a set of parameters that work on deciding (classifying) if the sounds are morepork or not – this is known as learning – and the result is some code that we call a model,

5) Evaluate the model to see how well it works on sounds it hasn't yet heard.

Feature extraction

I kept going around in circles on how to do this. Weka requires the features to be input as a sequence of numbers or categories and I kept getting conceptually stuck on how to represent an audio signal in this format. Here is what a typical recording looks like when you view how the volume changes with time. Can you see the morepork? Unlikely!

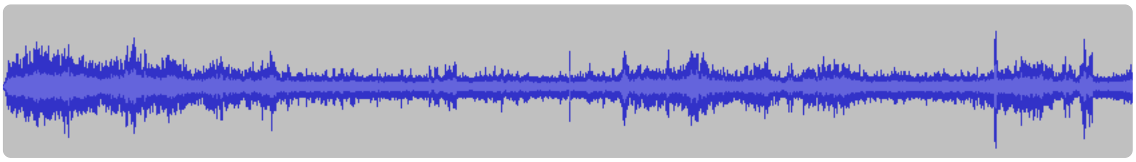

There are various published techniques and code for extracting features from audio but again I was concerned about losing information about how the sound changed with time and how to deal with the calls starting at different times. To cut a long story short, I ended up using a common technique of representing the sound as a spectrogram which indicates how the frequencies change with time. Can you see the morepork in this next image?

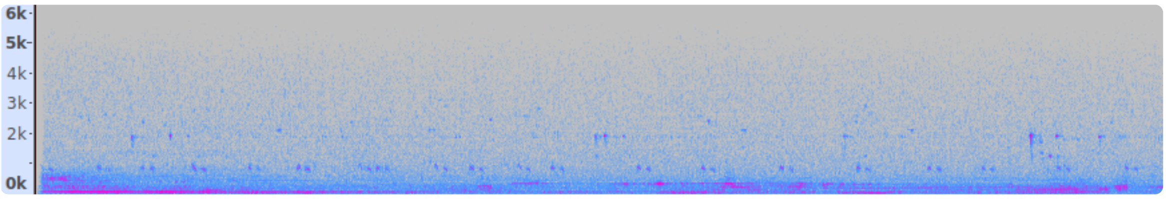

Maybe, but which of those blobs are morepork? Well as we know that morepork calls are in the frequency range of 700-1000Hz, let's look there.

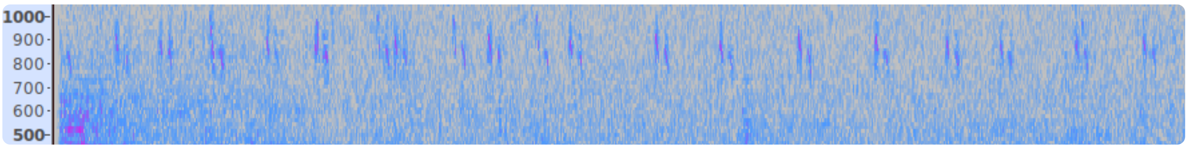

There we go – what you see is about 19 morepork calls in the minute recording.

To extract features, I 'cut-out' each of the example calls from that image (and many others) and feed them into Weka which can automatically create features from the images. Don't worry, I didn't cut them out by hand, but instead I used some modified code written by Chris Blackbourn that finds locations in an audio file where the volume increases. It is worth noting that even if there are very loud background noises in the recording, as long as the noises are not in the 700 – 1000 Hz frequency range, they will not interfere with the image of a very quiet morepork call.

I've heard that before?

Creating a model turned out to be a very iterative process. As I said before, once you have a model, you need to confirm that it is working – this can only be done by listening to the recordings and comparing what the model predicted with what it actually appears to be. I'm no bird expert, far from it, which is partly why I choose the classic sounding more-----pork call to start with. It soon became obvious that I found myself listening to the same recordings over and over again for each new model iteration. It was time for a database and some software to keep track of everything. I was also getting suspicious that the model might work better if all the non-morepork calls were actually categorized correctly rather than just as 'not-morepork' as I had been doing. Here is one of the interfaces from the Audio Manager application that I've developed to help with this process – you can get the code at https://github.com/TheCacophonyProject/audio_manager (https://github.com/TheCacophonyProject/audio_manager)

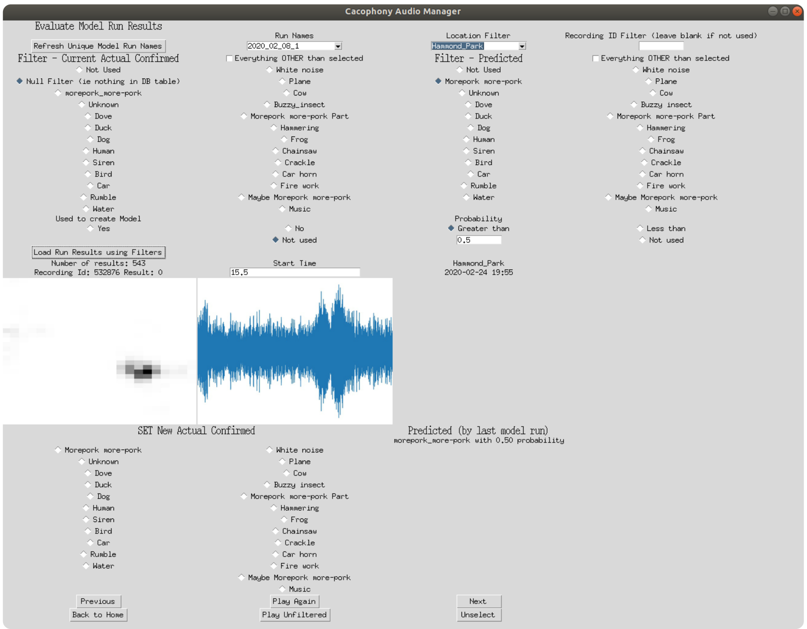

Run Names: 2020_02_08_1

Location Filter: Hammond_Park

Filter - Current Actual Confirmed: ● Null Filter (ie nothing in DB table)

Used to create Model: ● Not used

Filter - Predicted: ● Morepork more-pork

Probability: ● Greater than 0.5

Number of results: 543
Recording Id: 532876 Result: 0

Start Time: 15.5

Hammond_Park
2020-02-24 19:55

Predicted (by last model run)
morepork_more-pork with 0.50 probability

Results

For each audio clip that is fed into the model, it returns a predicted classification along with how sure (certain) it is of its prediction. The less certain the model is, the more likely it turns out to be wrong: that is, the model says it is a morepork when it isn't. The current model is currently far from perfect - however, it is giving enough information to be useful. In the following graph, the bottom blue dashed line shows that as the model certainty reduces (left to right) the number of errors increases.

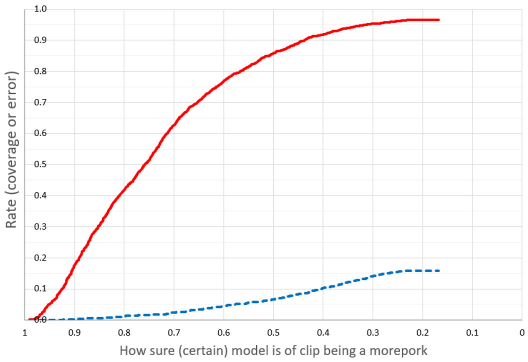

The top solid red line shows that the number of moreporks that we actually find, increases when we accept a higher number of errors. There is an obvious trade-off between locating all the morepork calls and having too many errors – known as false positives. The model is currently set to only 'tag' moreporks on the Cacophony interface if the model has a certainty of greater than 0.5 (50%).

Further work

- Currently not all the morepork calls in a recording are located for the model to analyze so the detection of these places of interest can be improved.
- The model is currently using the location of the Bird Monitor recorder as one of the input features to the model. This means that new Bird Monitors and those with few moreporks present may not perform as well as locations with many examples.
- There will certainly be better methods of creating the features – maybe I (or you) could use TensorFlow to create these features.

Tim Hunt

Many thanks to Tim for this article. Tim (and everyone at Cacophony) would love to hear any feedback you have.

You can leave a comment below and/or you can email us at blog@cacophony.org.nz (mailto:blog@cacophony.org.nz?subject=Ruru%20identification).

Publication Date:
Tuesday, 24 March 2020

1 Comment **The Cacophony Project** **Disqus' Privacy Policy** **Tim Hunt**

♡ **Recommend** Tweet Share **Sort by Best**

⚠ The Cacophony Project requires you to verify your email address before posting. Send verification email to timhot+14512153@gmail.com ✕

Join the discussion…

Anko Hanse • 13 days ago

Thank you for all your hard work. Had a quick look at all the moreporks identified by your tool in my recordings of the last few months. The AI seems to work very well!

Looking forward to automated recognition of many more different birds. But judging on your blog here, that is going to be an immense amount of work...

^ | ˅ • Reply • Share ›

✉ **Subscribe** **Add Disqus to your siteAdd DisqusAdd** ⚠ **Do Not Sell My Data**

mhellicar's blog (/blog/75)

Subscribe to our newsletter

email address

Subscribe

(https://twitter.com/birdsongnz) (https://www.facebook.com/cacophonyproject) (/contact) (https://drupal.org/u/lightweight) (https://github.com/TheCacophonyProject)

Hosting sponsored by Catalyst Cloud (https://catalystcloud.nz/)