

# Identifying nocturnal bird calls

**Presentation at the Department of Conservation,  
Christchurch**

Douglas Bagnall and Edward Abraham

5 June, 2014

Executive summary: not yet useful for Tier - 1 monitoring

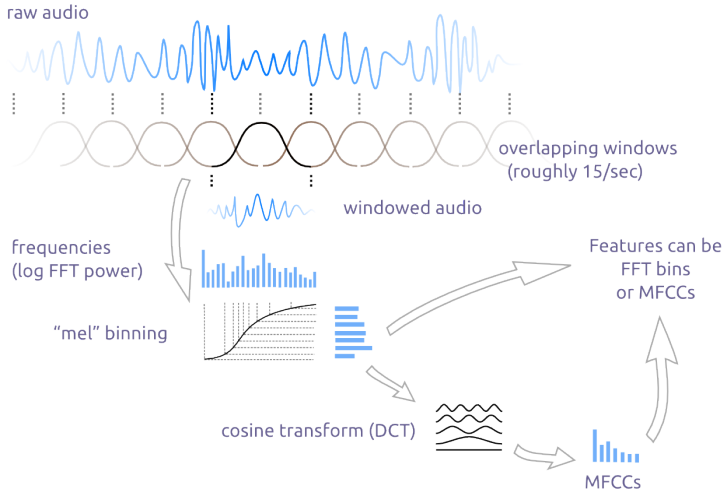


## Goals


- Identify kiwi, morepork, and weka in nocturnal recordings
- Allow recordings to be ignored that are unlikely to contain calls to reduce the effort needed to score calls
- Facilitate consistent, automated monitoring of acoustic data from around New Zealand

# Recurrent Neural Networks

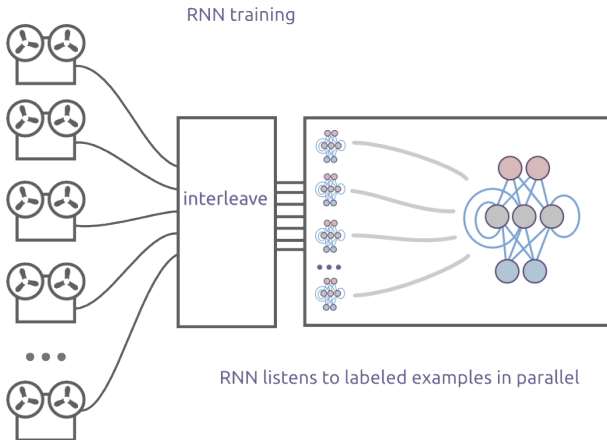
# Audio processing




# Recurrent Neural Networks




# Parallel processing



# An example prediction



**RFPT-LPC-2011-11-26T13:45:03Z-540-60.wav**



# Identifying kiwi



## Training data

- Provided with 2439 15 - minute audio files from the Tier - 1 data set
- First calls of each bird species in each file labelled using Freebird
- A total of 1327, 183, and 103 files containing kiwi, morepork, and weka, respectively

### Files with kiwi in the Tier - 1 training set

species	2011-12	2012-13
brown/ tokoeka	78	34
great spot	28	21
little spot	0	3
<i>spp</i>	0	19
total	106	77


# Training data

- Requires a well - labelled training set
- Current Tier - 1 protocol not ideal for three reasons
  - ① For kiwi and weka, there were insufficient examples in the training data
  - ② not all calls are labelled
  - ③ time bounding of calls isn't precise
- Carried out our own labelling of morepork calls
- Used data from the Rimutaka Forest Park Trust for kiwi


## Rimutaka kiwi

- Data from the Rimutaka Forest Park Trust
- 20 444 1 - minute clips
- Half of the clips with high energy in the kiwi frequency
- Half of the clips randomly sampled from the remaining 600 000 clips
- Added in 6870 1 - minute kiwi - less clips from the Tier - 1 set
- Held out 2500 clips as a test set


# A successful prediction




**RFPT-LPC-2011-11-26T13:45:03Z-540-60.wav**




# No kiwi here




**RFPT-LPA-2011-12-25T16:45:02Z-120-60.wav**




# This tūi might be a kiwi




**RFPT-SG2-2012-03-16T22:45:03Z-660-60.wav**



## And it didn't find this call




**RFPT-LPB-2011-11-19T15:00:02Z-600-60.wav**



# Kiwi RNN applied to test data

An AUC of 96%




	Score below 0.4	Score above 0.4	Correct (%)
No kiwi	2168	187	92.0
kiwi	12	133	91.7



## Kiwi RNN applied to Tier-1 data

An AUC of 71%

Not useful for discriminating kiwi in the Tier-1 data set




	Score below 0.4	Score above 0.4	Correct (%)
No kiwi	1907	338	84.9
kiwi	82	71	46.4

# What has gone wrong?


- Multiple kiwi species in Tier - 1 data set
- Greater diversity of background sounds
- More possibility of mistakes in 15 - minute data

# Difficult to distinguish kiwi and weka

Weka




Weka sounding like a kiwi



Labelled in Freebird as kiwi - weka duet



Labelled in Freebird as weka




# Identifying morepork

## Morepork calls


- Use the first minute of each Tier - 1 file, and the rest of 'interesting' files
- Count any morepork call type (ruru, quee, etc.) as a morepork
- Extend data by changing the levels, and blending known morepork with a range of background noise
- A total of 16 146 labelled minutes, 7938 with morepork
- A total of 26 651 calling periods

# Finding morepork calls



# Morepork RNN applied to test data

An AUC of 88%




	Score below 0.2	Score above 0.2	Correct (%)
No morepork	1703	281	85.8
Morepork	151	365	70.7

## Morepork RNN applied to Tier-1 data

An AUC of 74%

Not useful for discriminating morepork in the Tier - 1 data set



	Score below 0.2	Score above 0.2	Correct (%)
No morepork	877	398	68.8
Morepork	393	730	65.0



## Morepork RNN

- Classifier not as accurate on 1 - minute training hold - out clips as kiwi
- Morepork are harder as the individual calls are shorter
- Perhaps there are difficulties with the diversity of calls, and wide variation in intensities
- Performance degrades as the interval is extended to 15 minutes

## Summary

- Recurrent Neural Networks not yet suitable for automating Tier - 1 acoustic monitoring
- To improve would require specialised training data (calls well - located in time, and with large numbers of cases)
- May need other modelling methods (e.g., Random Forests) to go from continuous score of the RNN to a classification of the audio file
- Positively, the RNNs will be useful for finding infrequent kiwi calls at sites similar to the Rimutakas