AI to protect NZ birds

Shaun Ryan

THE
Cacophony
PROJECT

2040

THE Cacophony PROJECT

2040

My background

Artificial Intelligence
Deep learning

Founding CEO

SLI SYSTEMS

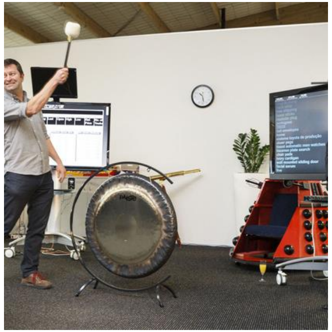

Governance

fleetpin

REDSEED

ARCHIPRO

Chair

THE Cacophony PROJECT

2040

Grant’s background

GlobalBrain

RealContacts™ word of mouth solutions™

eurekster

SLI SYSTEMS

YikeBike™

purepods™

History of the Cacophony Project

- Bought pest infected property
- Thought bird song might be increasing: how to measure?
- Decided he was the worlds worst trapper
- Saw an opportunity to apply technology to the problem

Cacophony development strategy

- Moore's Law – twice as good or half the price every 18 months
- Open source – collective intelligence
- Focus on engineering solutions not scientific discovery

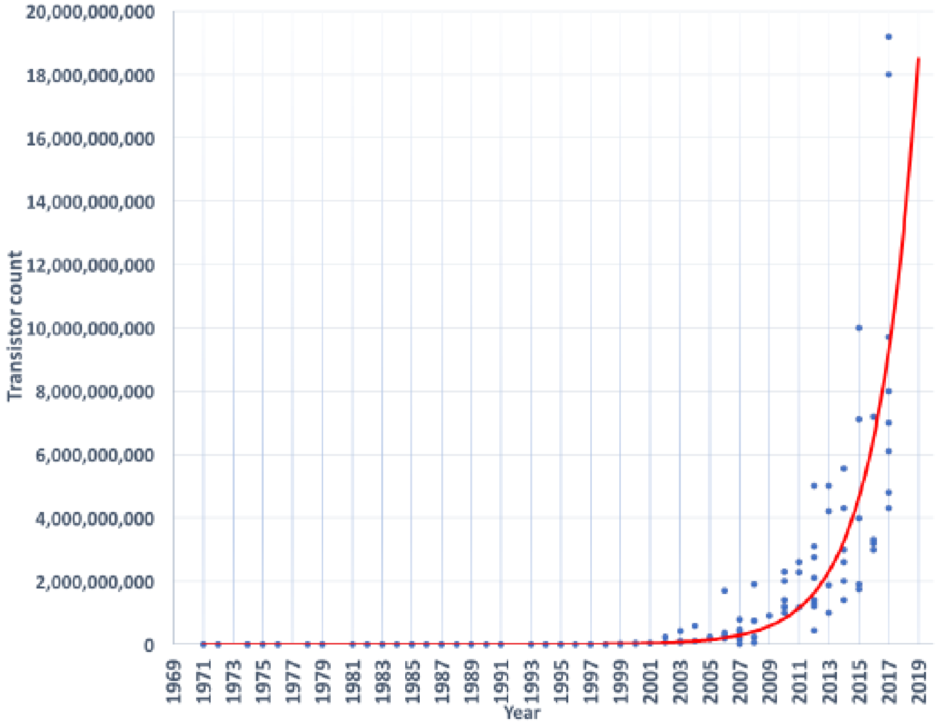

THE Cacophony PROJECT

A not for profit : a variety of contributors

2040

NEXT FOUNDATION
Investing in our land, our people

jasmine SOCIAL INVESTMENTS

Spark

CCL

PREDATOR FREE 2050

CONCENTRATE

ZIP ZERO INVASIVE PREDATORS

BALANCED

ANDERSONLLOYD LAWYERS

2040 - a social enterprise

www.2040.co.nz

The Mission of the Company is:

1. To use technology developed by The Cacophony Project Charitable Trust (CC54701) to eradicate all the predators of native birds in New Zealand by 2040;
2. To make the technology available to similar projects globally; and
3. To financially support The Cacophony Project.

The Cacophony Project

2040

Bird Monitor

You can't manage what you can't measure

- 1 minute recordings ~38 times/day
- With date, time, location
- Cacophony Index measures amount of bird song – calculated every 20s
- Bird species automatically identified (soon)
- Uploaded to the cloud
  - Can listen and export data to spreadsheet.
  - Will store up to 3 years with a memory card if no reception

Price: $419

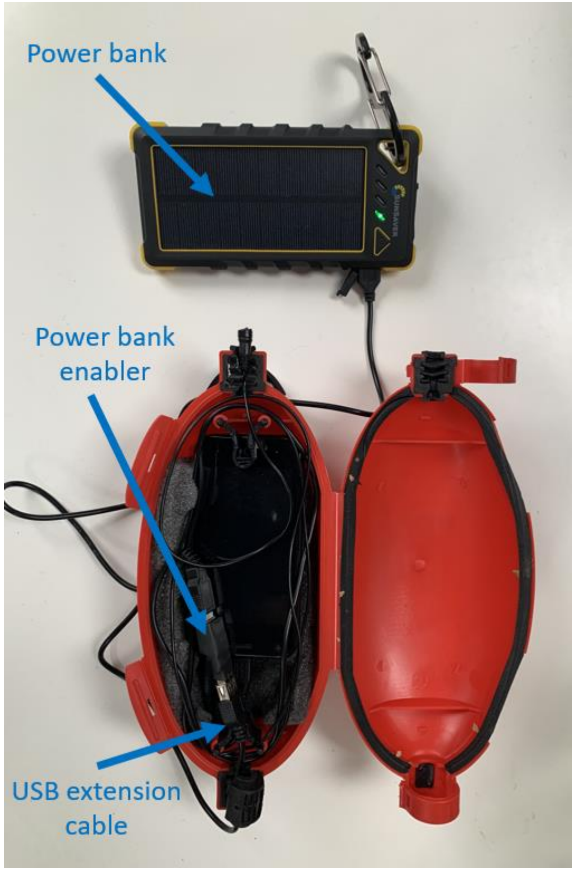

Accessing the recordings in the cloud

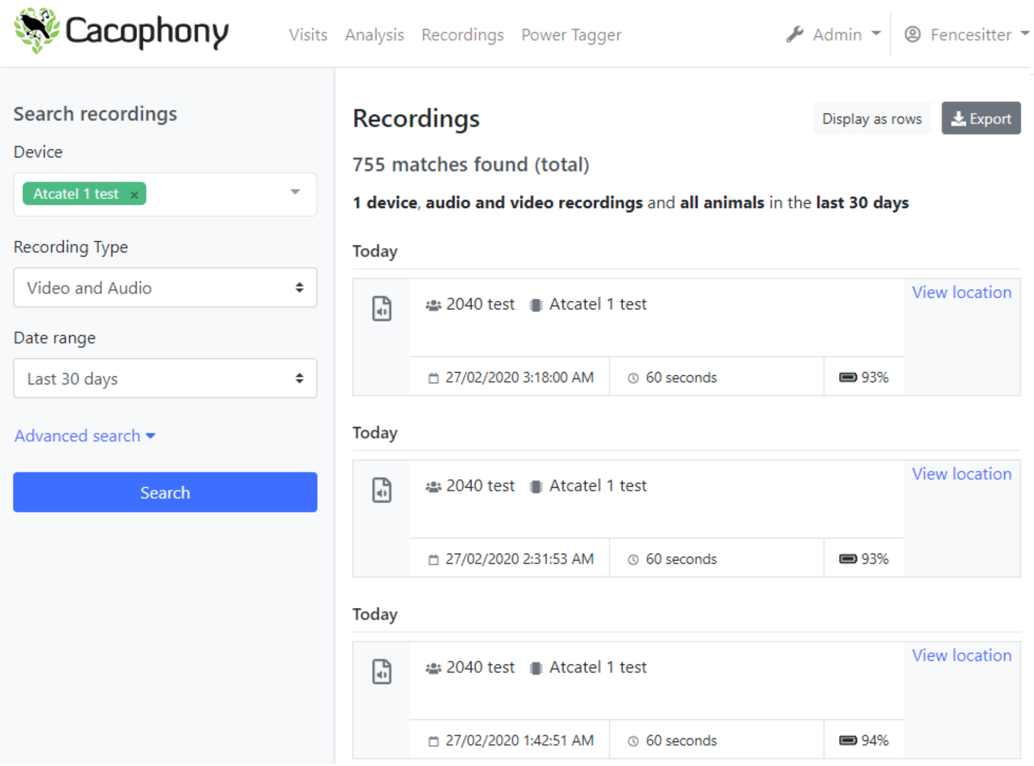

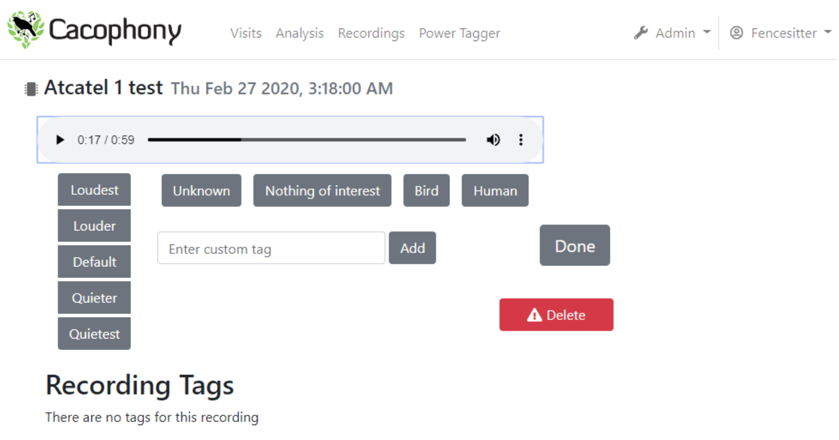

Listen, tag, view location & battery.
Delete & Export

Cacophony Index

U	V
Cacophony Index	
42.4;39.5;39.2	
65.5;57.3;53	
55.1;44.9;46.5	
46;43.8;43.2	
63.6;57.7;54.2	
45.2;42.9;38.9	
40.8;46.2;37.5	
56.2;62.2;53.2	
49.9;70.3;66.5	

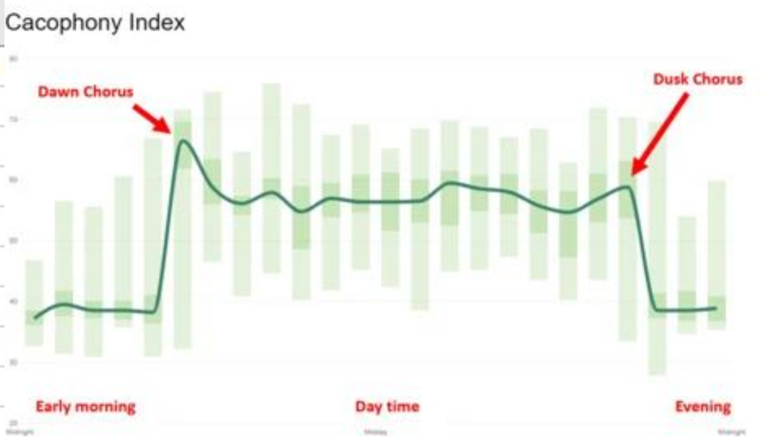

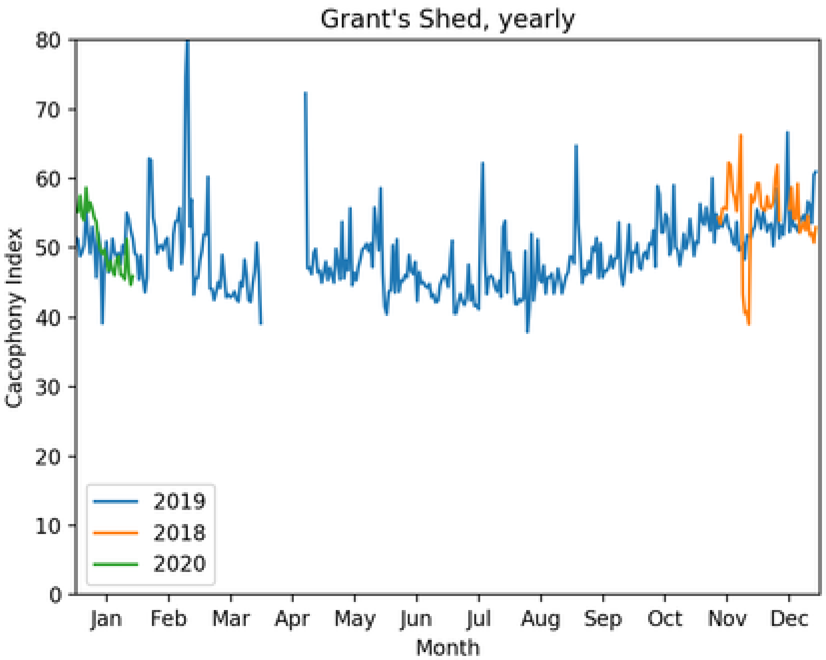

High Cacophony Index 69.6, 67.4, 69.1

Medium Cacophony Index 56.7 48.9 55.3

Bird classification open source contributers

Tim Hunt

Principal Lecturer, Wintec, Hamilton

tim.hunt@wintec.ac.nz

Dennis Sosnoski

Enterprise Software Consultant

dms@sosnoski.com

Audio Classification 1.0

- Previously developed a model for automatically identifying morepork/ruru calls using the Weka workbench for machine learning – from Waikato University (cs.waikato.ac.nz/ml/weka/)
- Method
  - Manually locate calls in recordings
  - Create spectrogram images of individual calls
  - Investigate which classifiers give the best performance
  - Used onset detection to find locations in recordings
    - Fed model with spectrograms of locations
  - Update model and repeat training/validation loop
- Result
  - Best results with a Random Forest classifier
  - But –false positive rate deemed too high (exact value depended on test data)

Training data

- Python app for creating training/validation and test data.
- Plays the recording and user to selects areas on spectrogram – labels & saves to database

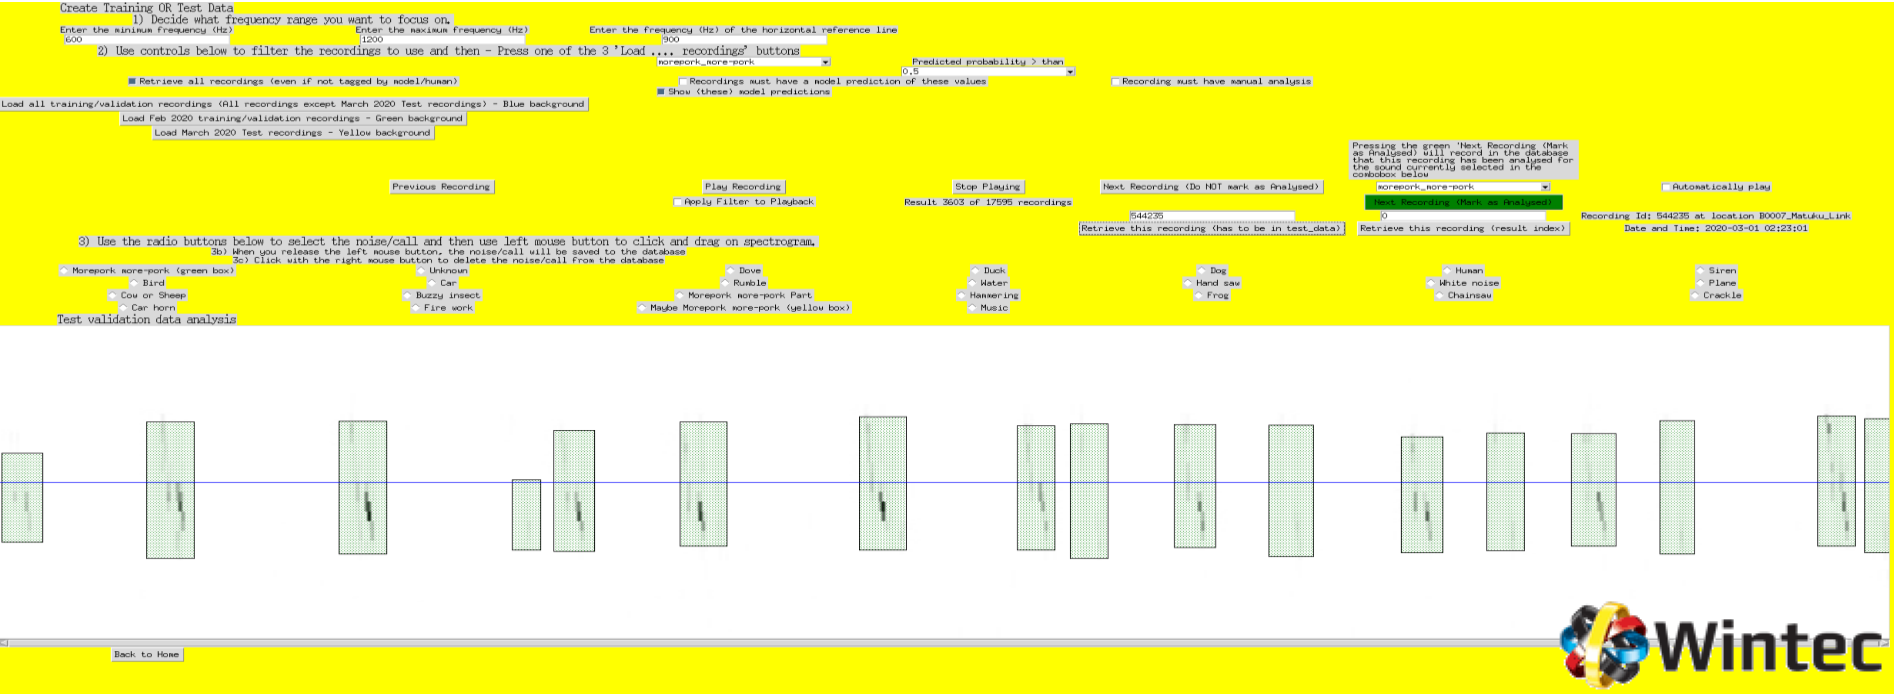

Audio Classification 2.0

- Keras/Tensorflow (tensorflow.org/)
  - Still using spectrogram images
  - Convolutional Neural Networks (CNNs)
- No longer using onset detection (As it missed a significant number of calls)
  - Now using a sliding window across recording
- Various model configurations investigated including:
  - Basic sequential model
  - VGGs, ResNets and DenseNets

Current Approach

- Deep CNN network (ResNet34 works well)
- Input is 60x60x1 spectrogram "image"
  - 60 frequency buckets for 600-1200 Hz
  - 60 time slices for 3 seconds of recording
  - Spectrogram dB power scaled, normalized
- Output is binary classifier for morepork or not

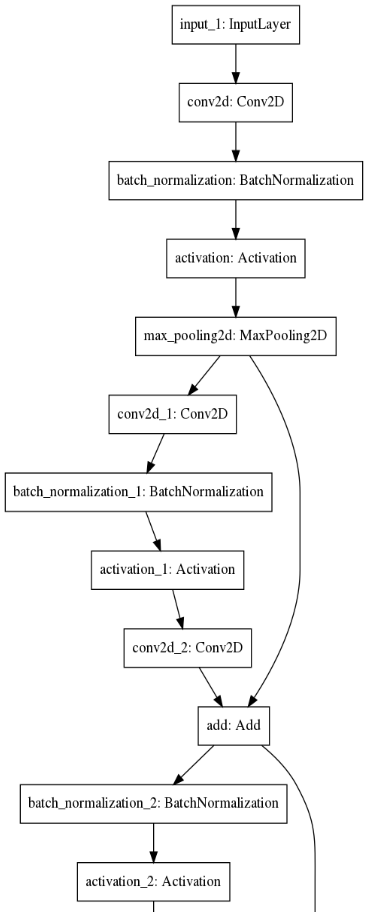

Bird Classification Results

- Training results measured on fixed sample subsets
  - Evaluating on fixed training data subset typically stabilizes at 1.0, with some dips
  - Validation (completely separate segments) stabilizes about .93 to .97
- Test results evaluated on sliding window in test data
  - Labels include “maybes”, which are skipped – not counted either way
  - Also skip counting where positive label only partially in window
  - Positives a (very) small fraction of total (944 out of 476298)
  - Typical result precision 0.8347, recall 0.6472, f-score 0.7291

Bird classification example

- Will be available in export soon
- Example finding morepork in Karaka Bay

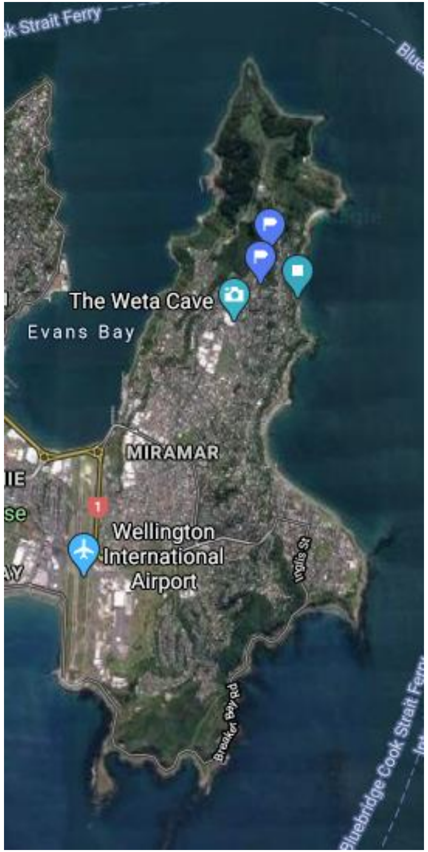

Thermal Camera

Grant started using trail cameras but found they were missing things.

- Camera too slow to come on
- Missing small predators

Viewing video time consuming

Experiments showed thermal cameras could overcome these.

Didn't want to develop our own camera – but there was nothing out there that could record autonomously, last for many days and upload to the cloud.

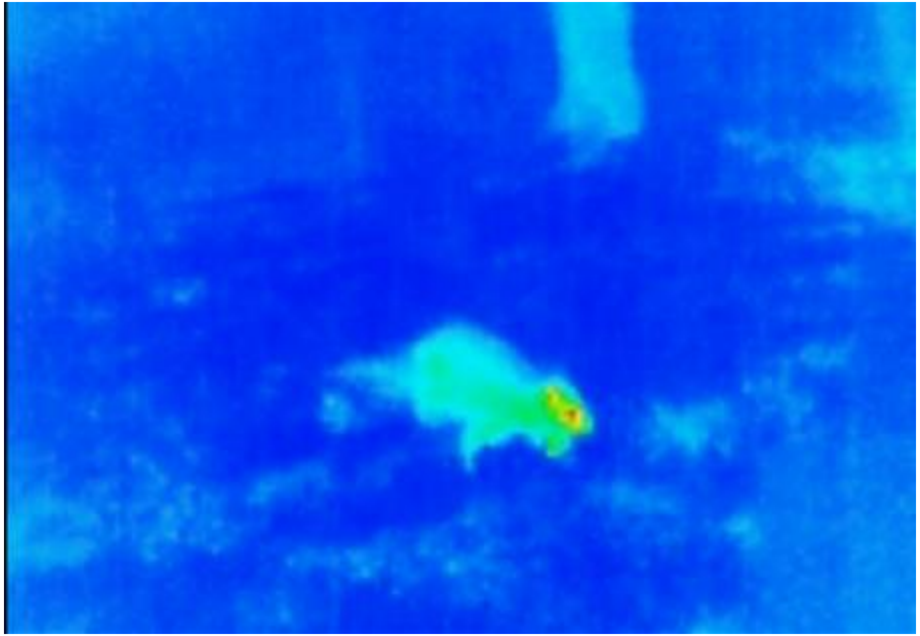

Thermal Camera

Thermal Camera and battery combo

~~$3,876.00~~ $3,669.00 Sale

- Automatically records movement
- Uploads to cloud
- Machine Vision algorithms detect predators
- Videos can be viewed, tagged, deleted, shared
- Consecutive videos of animals are converted into visits
- Stats can be viewed and pulled into a spreadsheet

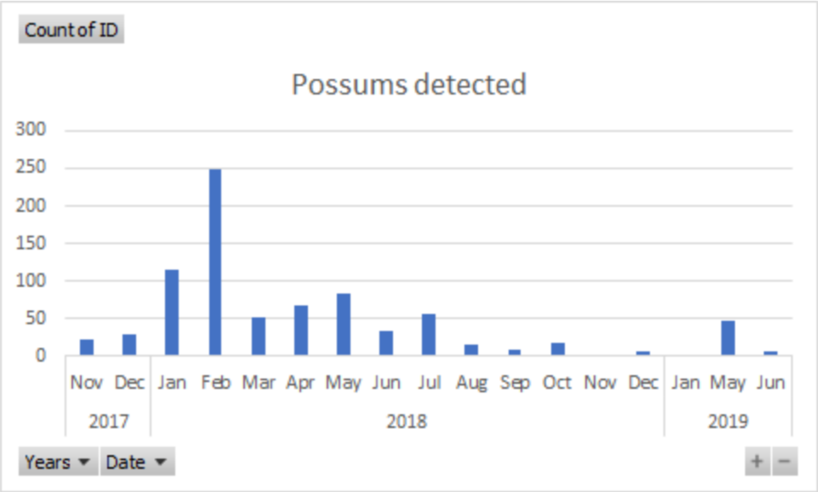

Data collection

- Camera detects motion and starts recording
- Videos are uploaded to our servers and processed
- Identify tracks
- Tagging interface – including Power

Students, volunteers, & team tag

Have >15K tagged videos

Demo of tagging interface.

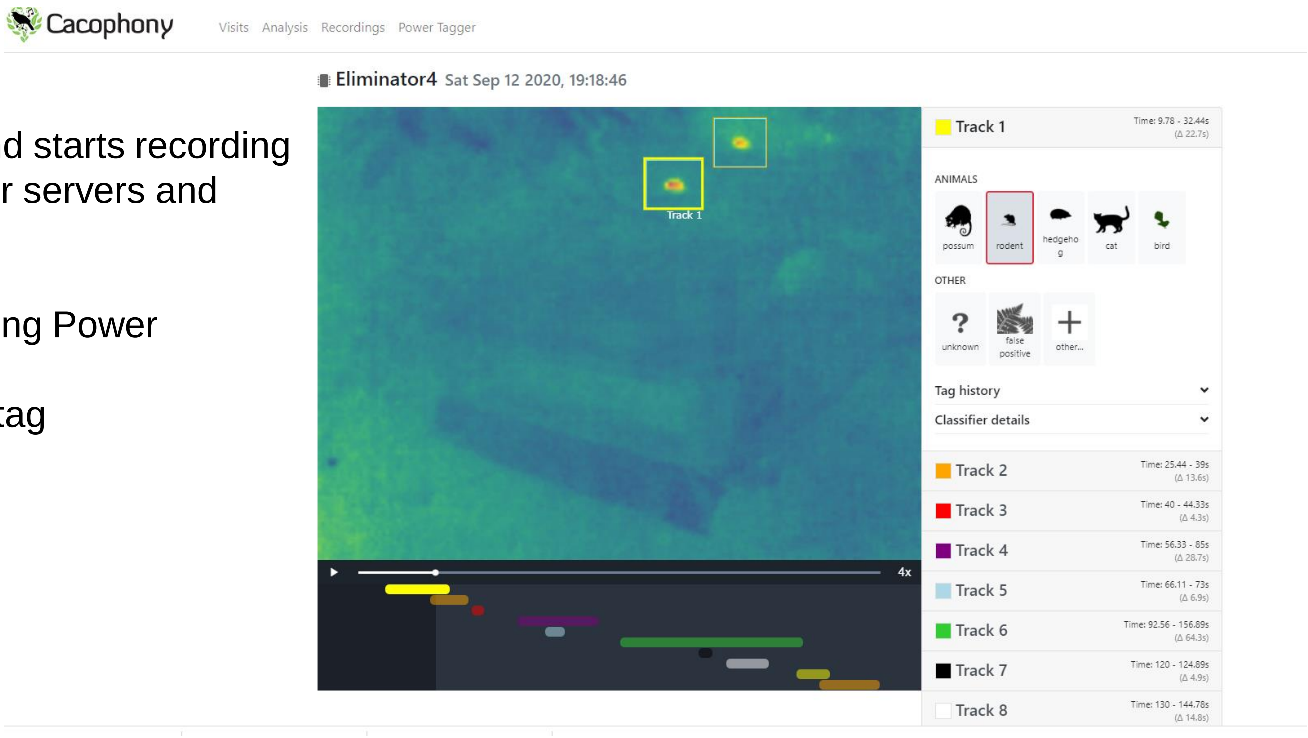

Machine Vision Models

- Matthew – student built first model over summer holidays – working last year
- Retrained the model as we got more data
- Added more animals
- Created a new model
- Created a wallaby/not wallaby model.
- Looking at retraining existing models like inception
- Tensorflow

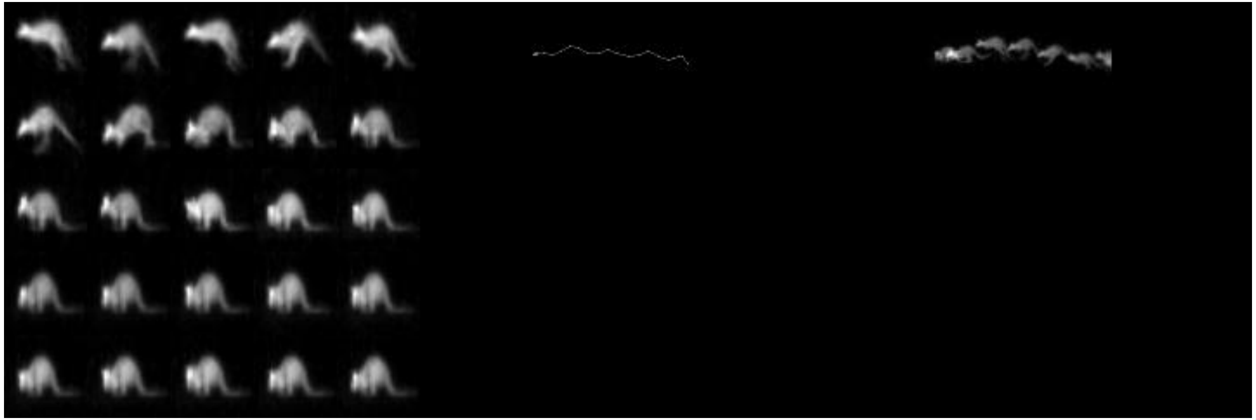

THE Cacophony PROJECT

2040

Example of machine vision videos

THE Cacophony PROJECT

2040

Thermal Camera Visits

Visit Summary Per Device

Pourewa camera

Animal	First Visit	Last Visit	Visits
rodent	03/06/2020 11:50:13 PM	03/09/2020 7:34:10 AM	8
bird	03/05/2020 7:33:17 AM	03/09/2020 7:21:32 AM	7
possum	03/06/2020 9:12:32 PM	03/09/2020 4:48:53 AM	10
cat	03/04/2020 1:56:34 AM	03/09/2020 2:38:58 AM	11
unknown	03/09/2020 12:47:35 AM	03/09/2020 12:47:12 AM	1
insect	03/07/2020 10:37:02 PM	03/08/2020 9:07:32 PM	3
mustelid	03/08/2020 4:59:58 AM	03/08/2020 4:59:54 AM	1
unidentified	03/06/2020 11:32:17 PM	03/07/2020 2:16:56 AM	3

Lincoln University Research: possums

	Trail Camera	AI heat camera
Number of recordings	5264	342
Average recording length (seconds)	10	16
Number of recordings to be analysed	5264	49
Minutes of recordings to analyse	877.33	13.07
Number of possum identifications	5	18
Minutes of analysis per detection	175.47	0.73

Better: >3x more detections
Faster: <1/200th analysis time
Cheaper: <6th cost/detection

Timestamp	Heat Camera	Trail camera	Chew card - camera	Chew card - other
16/11 10.02pm	X	X		
16/11 10.15pm	X			
16/11 10.18pm	X			
19/11 03.09am	X			
21/11 10.18pm	X	X		
24/11 12.10am	X	X		
24/11 10.32pm	X			
24/11 11.46pm	X			X
25/11 12.29am	X			
25/11 01.36am	X			
27/11 09.38am		X		
30/11 09.50pm	X			
01/12 05.17am	X			
02/12 02.34am	X			
03/12 03.08am	X			X
03/12 03.22am	X			
07/12 02.10am	X	X		X
12/12 09.58pm	X			
12/12 10.27pm	X			
Total detections	18	5	0	3

A bigger difference with rats

Rat results

	Thermal Camera	Trail Camera	Times better
Detections - close camera	19	1	19
Detections - far camera	19	4	5
Video length close camera (seconds)	1524	6	254
Video length camera (seconds)	1524	1	1524

The importance of scent trails

A COVID diversion

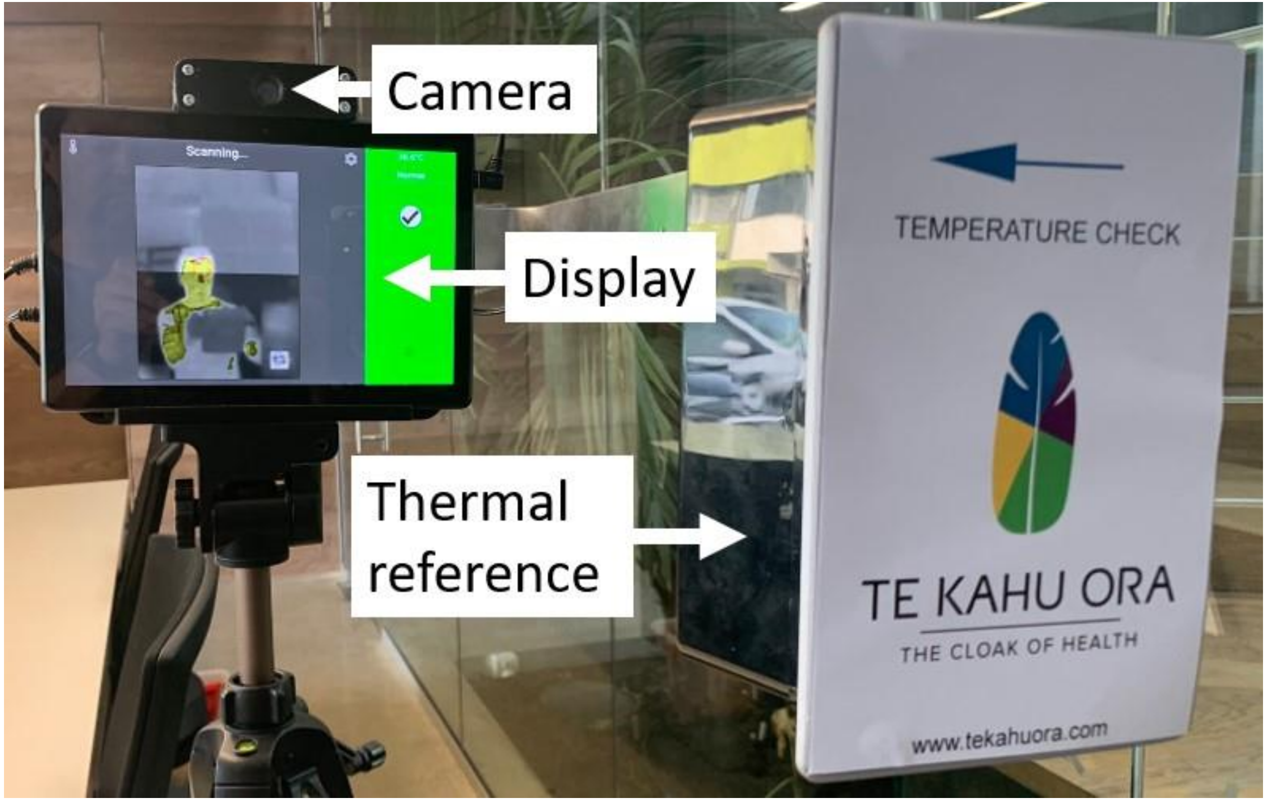

Potential for IT to improve trapping

	Attracting predators over a larger area	Higher catch rate	More predator types in same trap	Overall improvement
Potential	100 + times	100 + times	6 types	60,000 times better
Initial experiments	4	20-50 times	4 types	320 times better

Catch rates

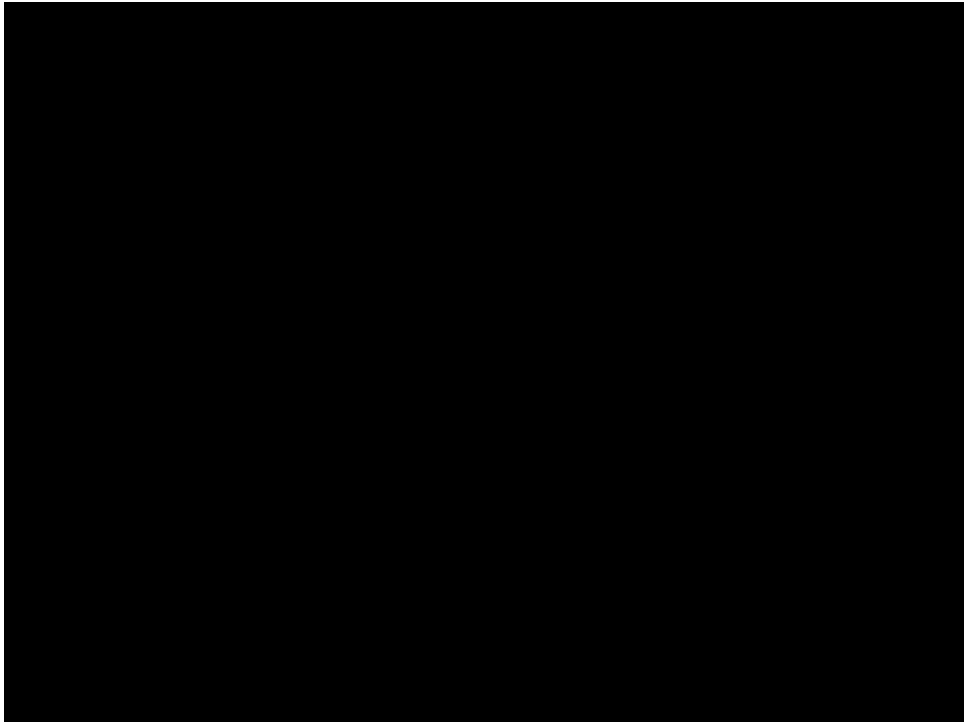

Overall catch rates

Trap type	Rats	Possums	Mustelids	Cats	Hedgehogs
Cacoph Open V4		35%		7%	56%
Timms		1%		1%	
NZ Autotrap		?			
Good Nature	0.9%	1%			
Trapinator		0.3%			
Live capture	6.0%	7%		3%	28%
Bait Station	0.6%				
Doc 200		NA			

High interaction rate makes a massive difference!

Variable Model inputs	Change numbers below to see how population changes	Low trap interaction rate
Trap interaction rate (chance that a predator triggers a trap in any given week in 10 hectare area)	9.00%	0.40%
Kill rate (for every triggering of trap what % die)	100%	100%
Number of elimination devices (auto resetting)	10	10
Population of predators	40	40
Fixed model inputs		
Area (hectares)	10	
Quick calculation that it looks about right		
Catches per year	40.0	35.0
Catches per trap per year	4.0	3.5

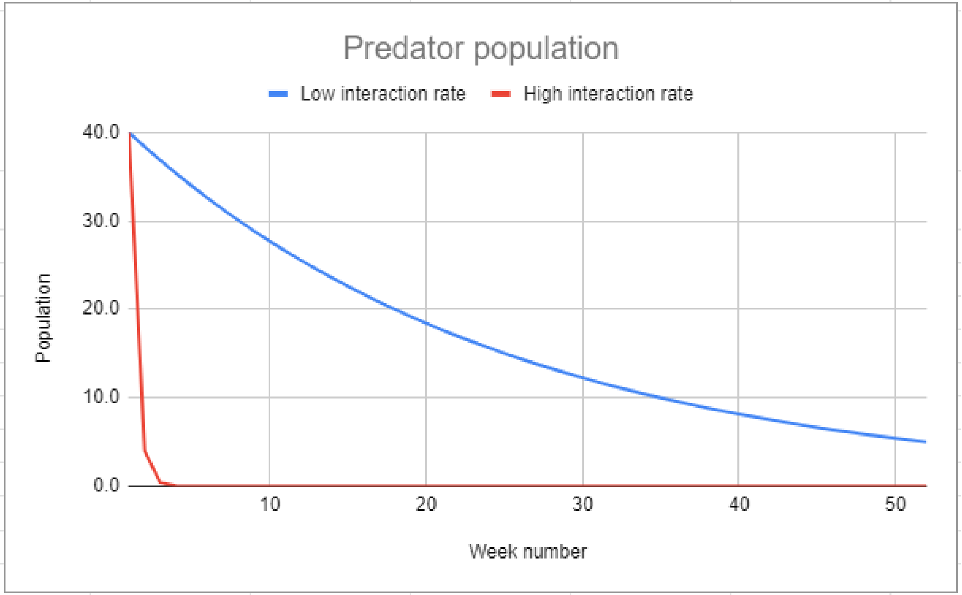

Our trap philosophy

- Open traps work have higher interaction rates
- Powered by AI can ensure only target species captured
- Live capture and hold can be used as a lure
- Auto kill and self resetting will reduce labour

3 Chamber trap

1. Initial capture – powered by AI – resets when animal in chamber 2
2. Holding chamber – access via a 1-way door – held for period of time
3. Death chamber – self resetting trap

The evolution of the Yike Bike

The evolution of the trap (so far)

The latest trap

Do we need a trap?

Poison Paintball

Looking for more contributors

- Happy to share data sets
- Opportunity to hone your neural network skills
- Several contributors have ended up with paying roles at Cacophony
- Is a great cause

THE Cacophony PROJECT

2040

Sign up for newsletters:

- www.cacophony.org.nz
- https://www.2040.co.nz/pages/2040-newsletter