CacophonyViz : Visualisation of Birdsong Derived Ecological Health Indicators

Mark T.H. Nikora
Waikato Institute of Technology (Wintec), New Zealand
mark.nikora@wintec.ac.nz

Tim D. Hunt
Waikato Institute of Technology (Wintec), New Zealand
tim.hunt@wintec.ac.nz

Grant Ryan
The Cacophony Project, New Zealand
grant@hunchcruncher.com

Abstract

The purpose of this work was to create an easy to interpret visualisation of a simple index that represents the quantity and quality of bird life in New Zealand. The index was calculated from an algorithm that assigned various weights to each species of bird.

This work is important as it forms a part of the ongoing work by the Cacophony Project which aims to eradicate pests that currently destroy New Zealand native birds and their habitat. The map will be used to promote the Cacophony project to a wide public audience and encourage their participation by giving relevant feedback on the effects of intervention such as planting and trapping in their communities.

The Design Science methodology guided this work through the creation of a series of prototypes that through their evaluation built on lessons learnt at each stage resulting in a final artifact that successfully displayed the index at various locations across a map of New Zealand.

It is concluded that the artifact is ready and suitable for deployment once the availability of real data from the automatic analysis of audio recordings from multiple locations becomes available.

Keywords

The cacophony project, data visualisation, ecological health indicators, geographical information systems

1. Introduction

New Zealand's bird species and populations are under continual threat from predators, most of which were introduced by the arrival of humans at different time periods. The Kiore, or pacific rat was introduced by the first Polynesian settlers, ship rats and Norway rats arrived with the first Europeans. Stoats and cats were brought to control rabbit and rat populations (Brockiem, 2015). Acclimatisation societies were established by early European settlers for the importation of non-native plant and animal species to help new arrivals to settle by providing them with reminders of home (Troup, 2005). No real consideration of the impacts of such actions was considered, and several introduced species became pests.

A report on the conservation status of bird species in 2016 demonstrates the impact of such predators on bird populations, and the magnitude of the current situation. Of 426 living bird species, 16.7% were assessed as threatened and 25.1% at risk. (Department of Conservation, 2017).

The Cacophony project (Moore's law for New Zealand birds, 2016) focuses on the application of technology as a means of improving this situation through monitoring birdsong and improving trapping mechanisms. The purpose of this study is to apply data visualisation to visually communicate the Cacophony projects efforts. The goal is to develop a visualisation prototype that could provide answers to the following - What impact has the application of technology had within a geographical area? Has the situation improved and can this be attributed to the Cacophony project? Can further improvements in technology usage be recommended based on what the visualisation is telling us?

In July 2016 the New Zealand Government announced its intention to make New Zealand predator-free by 2050 (Kirk, S., 2016) by eradicating all pests that threaten New Zealand native birds. In alignment with this government directive, the Cacophony project represents an open source collaborative effort to monitor bird-life and reduce predator population through the innovative application of technology. The project assumes that monitoring of sounds such as birdsong will prove to be a good ecological health indicator and proposes a Cacophony Index as a measure of how effective conservation efforts are in a geographical location.

Currently the project has produced and is trailing Cacophonometers (see figure 1) developed using a combination of solar and cheap mobile phone technology for data collection in the wild (The Cacophony project: summary of the development so far, 2016). Other devices are being considered for visual data capture of predators at trap locations. The project is also looking at the creation of intelligent traps that could have a significant impact on eradication of predators. The current thinking around this is that the trap will listen to and identify the type of predator in the trap vicinity and then use a digital lure such as playing the sound of a female of the species to draw the predator to the trap.

Figure 1. Cacophonometer composed of smart phone, water proof case and solar panel

In this study we consider various data visualisation techniques and use this to inform the development of a data visualisation prototype (CacophonyViz). In the next section we examine the type of data currently being collected by the Cacophony project as well as an overview of relevant geospatial data visualisation techniques. Section 4 Analysis and Design describes the requirements for CacophonyViz and Section 5 Implementation and Development describes the process that was followed in its development. Finally an evaluation of the work completed thus far is given along with a summary of the key findings of this study and opportunities for future work.

2. Related Work

As the Cacophony project is still in its infancy, the raw data collected currently exists in the form of short audio recordings taken at set intervals by Cacophonometers. Predator data may come from intelligent traps and identification of predators via video camera and machine learning. Another sub-project is currently looking at the application of machine learning to identify bird species and numbers from audio. The expectation is that recordings will be automatically analysed to extract the number and type of bird calls in each recording. Table 1 shows a weighting that could be applied to each bird type to reflect the 'quality' of the bird. An index will then be calculated for each location based on the number of each bird call and scaled for the percentage of time that recordings occur.

Bird Species	Weighting
Kiwi	10
Kereru	9
Kea	9
Weka	8
Bellbird	8
Fantail	7

2.1 Geospatial data visualisation techniques

The raw data collected from the network of Cacophonometers is location-oriented and so it is appropriate to review a range of existing applicable geospatial data visualisation techniques that can facilitate display, selection and filtering.

Chloropleth Map

A map segmented by geographic regions using colour shading progression to visually represent a variable under consideration (Kirk, A, 2016). Various colour progression schemes can be used besides a basic single hue progression. Examples are bi-polar colour progression used to represent negative and positive values or values around some central tendency. One problem with this technique is that the viewer may assume the entire geographical region has the same value however this may prove not to be the case under further examination. Another is that size of regions may give undue visual emphasis in contrast to colour.

Figure 2. Example of chloropleth map with US states. (Severino, 2017).
From the data visualization catalogue.
Retrieved from https://datavizcatalogue.com/methods/choropleth.html. Reprinted with permission.

Cluster Map

A cluster map represents dense pockets of data points using a single cluster point. The cluster point can either be relatively sized or it can contain the number of data points in the cluster. Clusters allow for drill down interaction to a lower level of detail.
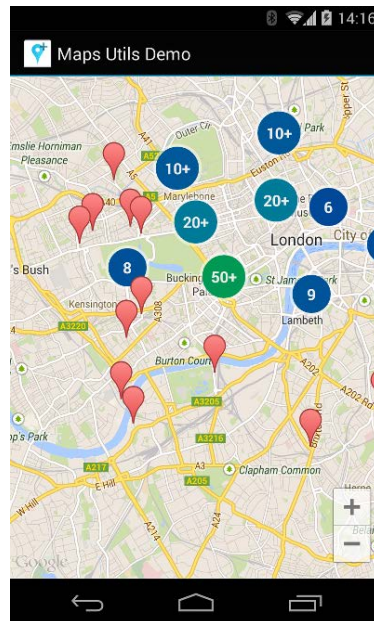

Figure 3. Cluster map. (Google, 2018). From Google Maps Documentation, Retrieved from https://developers.google.com/maps/documentation/android-sdk/utility/marker-clustering.

Hexagonal Binning

Hexagonal binning creates a uniform grid on the map with regular hexagons, allowing data to be aggregated into bins which can then be colour shaded to express the aggregation. Hexagonal shapes are selected as they are more similar to a circle than a square and offer more efficient data aggregation while still being able to tessellate (Nelli, 2014).
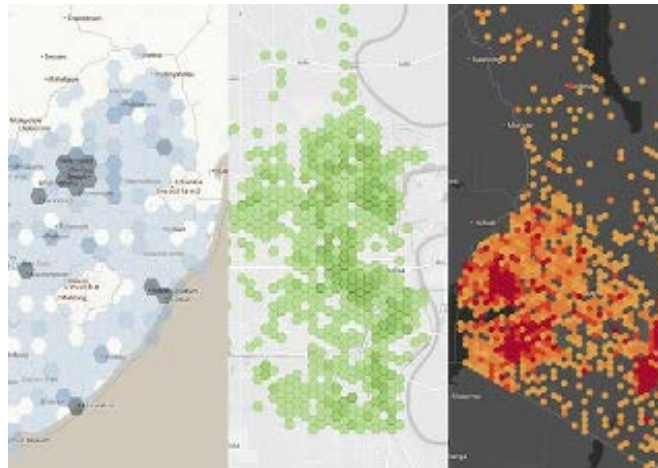
Figure 4. Three examples of hexagonal binning. (Smith, 2012). From mapbox blog, Retrieved from https://blog.mapbox.com/binning-an-alternative-to-point-maps-2cfc7b01d2ed

Animated Geospatial Techniques

A. Kirk (2016) supports the use of geospatial animation through plotting timestamped geographical data on a map to reveal the shifts and changes in the data. Such animation can tell a compelling data driven story. However, he also advises of the need to consider the type of changes that exist within the data and that such animated sequences are generally used to illustrate progressive or systematic changes as opposed to chaotic changes (Kirk, A., 2011). Interactive animation could be valuable in demonstrating the hoped-for changes in the Cacophony Index, proving the effectiveness of intelligent trapping techniques and monitoring.

3. Analysis and Design

The goal of this project is to develop a data visualisation prototype that presents to the end user the current status of the Cacophony project's conservation efforts. This is done by visually presenting Cacophony Index measurements (see section 2 related work) calculated from raw data collected by strategically placed Cacophonometers. After discussion with Cacophony project researchers an initial set of requirements were developed as follows.

1. Zoomable map of New Zealand showing geolocated Cacophony Index measurement plots with a 1 km grid resolution. This is to avoid showing actual GPS location and having equipment stolen.

2. Allow the user to select single or multiple plots.

3. Display current Cacophony Index value for each plot. For the purposes of this study the Cacophony Index is taken as a value in the range 0.0 to 1.0 with two decimal place accuracy.

4. Display Cacophony Index average for specific time-period such as monthly average.

5. Display Cacophony Index average for specific time-period and for time of day.

6. Show two overlays on a map, one for Cacophony Index and another for predator plots as equipment for monitoring predators will not be in same location as Cacophonometers.

7. Retrieve processed data for the visualisation from Full Noise. Full Noise is a sub-project focusing on a system that allows for uploading, storing and viewing of raw data from Cacophonometers via a basic web interface (Ryan-Pears, 2016).

4. Implementation and Development

The method followed in this research has been design science methodology (DSRM). Design science focuses on the design, creation and evaluation of an artefact as a driving mechanism for research and new knowledge creation (Hevner, March, Park, & Ram, 2004). Seven design science guidelines as outlined by Hevner et al (2004) provided direction for application of DSRM. These are: 1) Design as an artifact, 2) Problem relevance 3) Design evaluation, 4) Research contributions, 5) Research rigor, 6) Design as a search process, 7) communication of research.

Design as an artifact
The approach taken here was to create a software artifact for data visualisation. The artifact was developed as several initial low fidelity prototypes for exploratory purposes to further solidify ideas on visualisation. High fidelity prototypes were then developed leading to a final version.

Problem relevance
The relevance of any design science research effort is with respect to the community within which it operates. The visualisation problem addressed by this work has clear relevance to the community of contributors to the cacophony project and provides a visualisation of cacophony project efforts. The community was exposed to the work by private hosted versions of the visualisation on cacophony project servers, access to a github respository with all prototypes, and email correspondence.

Design Evaluation
Eliciting user feedback and testing to requirements were the two main forms of artifact evaluation used. This was done for each prototype (see section 4.3).

Design as a search process
Having defined the problem, reaching a solution involves a search process that considers application of applicable technology. Sections 4.1 and 4.3 describe this.

Research contributions
The design science contributions of this research are the development of a data visualisation prototype fit for visualisation of cacophony index data.

Communication of research
Design science research must be presented both to technology-oriented as well as management oriented audiences. The technology oriented audience is viewed as the cacophony project community and the audience for this research paper once published. The management oriented audience is seen as the cacophony project leader.

4.1 Technology selection

Technology selection was driven by the following requirements.

1. The ability to adjust to end user device display settings (tablet and desktop).

2. The geospatial nature of the project requiring maps with overlays as part of the visualisation.

3. Compliance with web standards to ensure accurate visual presentation across devices.

4. Cross browser compliance on major browser platforms - Google Chrome, Mozilla Firefox, Microsoft Edge, macOS Safari.

5. Open source non-commercial licensing restrictions.

Leaflet

Leaflet is an open source javascript web mapping library for constructing mobile friendly interactive maps (Agafonkin, 2017). It is comparable to the Google maps javascript API. Based on the API documentation, the API itself is small and focused on map-based functions, however Leaflet does not provide map data. The developer needs to access this from a tile server web service such as Mapzen, Mapbox or Open Street Maps.

D3.js

D3.js is a javascript data visualisation library that uses HTML5 technologies such as Hypertext Markup Language, Scalable Vector Graphics, and Cascading Style Sheets (CSS) in the creation of charts and graphs (Bostock, M). The advantage of this library is its ability to create advanced visualisations for complex data sets.

Bootstrap

Bootstrap is an open source front end framework focusing on developing responsive user interfaces that allow for a user interface to adjust to the display capabilities of the device it runs on (Bootstrap, n.d). Bootstrap incorporates a responsive grid for managing placement of content and resizing of layout. Resizing and rearrangement of layout is triggered by breakpoints that use CSS media queries. It also provides useful user interface components and can be extended further using JQuery.

JQuery

JQuery is a small javascript library having a lightweight footprint of 32kB after minification and gzip compression (JQuery, n.d). The library streamlines development by providing several convenience methods for manipulation of the web browser Document Object Model and is cross browser compliant.

Google Maps

Google maps represents the most accurate map resource with a very mature API with versions for javascript. It supports marker clustering and data visualisation methods (Google maps javascript API, n.d). Issues are the potential visibility of the API key within the application. This can be mitigated by limiting API usage to a

specific domain hosting the visualisation. As this is a commercial API it does not expose itself for modification at API code level in comparison with open source libraries. Request quotas are imposed with a complimentary free quota given with charging above free quota limits.

Mapzen

Mapzen is a vector tile service that delivers vector-based tiles to your application in several commonly used formats such as GeoJSON as opposed to predrawn bitmap image tiles (Mapzen javascript SDK n.d.). The vector tiles describe what can be drawn such as buildings, roads etc, leaving the selection up to the application by using tools such as scalable vector graphics (SVG), web graphics library (WebGL) or data visualisation libraries like D3.js. Generous free quotas are available on data requests with charging above quota limits.

Selection

Initially D3.js was chosen over Leaflet as Leaflet had a narrow focus whereas D3.js offered a far broader range of visualisation capabilities that can be called on in future development. Google maps javascript API was chosen as it represented the most mature API. Bootstrap and JQuery were chosen to provide DOM convenience methods and the responsive features needed to cater for different browsers and display device settings.

4.2 Development Environment

A private source code repository for CacophonyViz was created. Development tools used included the Atom hackable text editor which is a package-based editor that can be extended via package installation. The Atom Live Server package was used to quickly view and test prototypes using a live reload feature that automatically refreshes the browser on code edits. The package uses a basic HTTP server implementation via Node JS.

4.3 Prototyping Process

Several exploratory prototypes were developed to examine the Google maps javascript API, D3.js and alternatives such as D3.js with GeoJSON prior to evolutionary prototype development. Prototype 1 shown below in figure 5 was constructed using D3.js with a small data set. The aim was to show Cacophony Index readings on a gradient scale from light green (1.0) to dark green (0.0). Google marker customisation did not facilitate this, so readings were passed to D3.js and a custom overlay added to the map with gradient plots created using a colour interpolation function.

Prototype 2 was designed to incorporate a larger data set. A data generator program was developed that generated random data points constrained by region bounding boxes as shown in Figure 6. Prototype 3 represented a shift in technology direction as clustering of data points also required the ability to customize markers for colour gradients. Google marker and Leaflet marker customisation were examined with Leaflet providing the greater capability. Figure 7 shows Leaflet marker clustering. Clicking on a cluster will drill down to a view that shows the individual markers for the cluster. Figure 8 shows overriding of the marker rendering process to show Cacophony Index average and data point count.
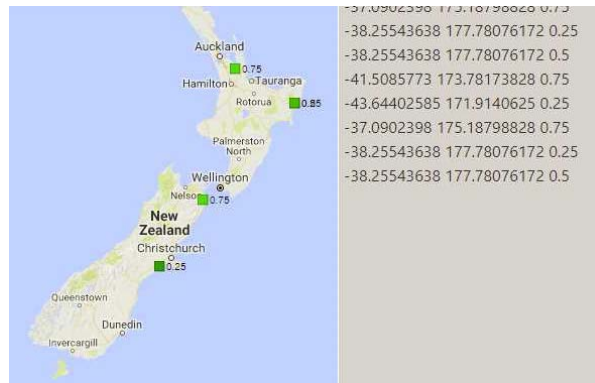
Figure 5. Prototype 1 – Gradient scale colour plot of Cacophony Index data using Google maps and D3.js with small data set
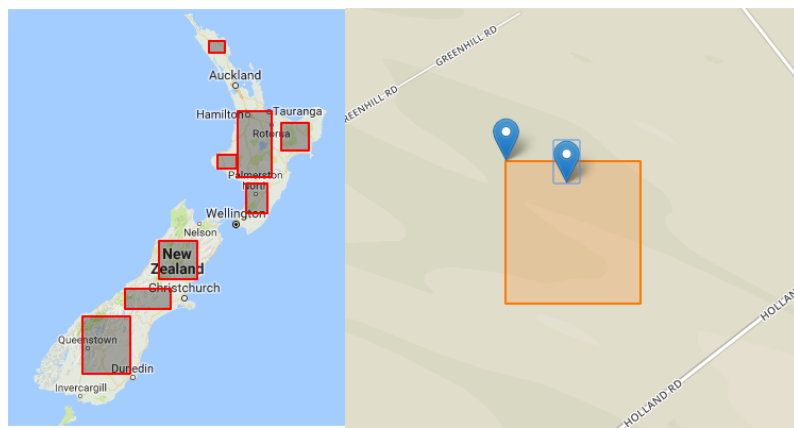

Figure 6. Prototype 2 – Data generator to create large random sample data set within region bounding boxes with debugging at 1 km resolution
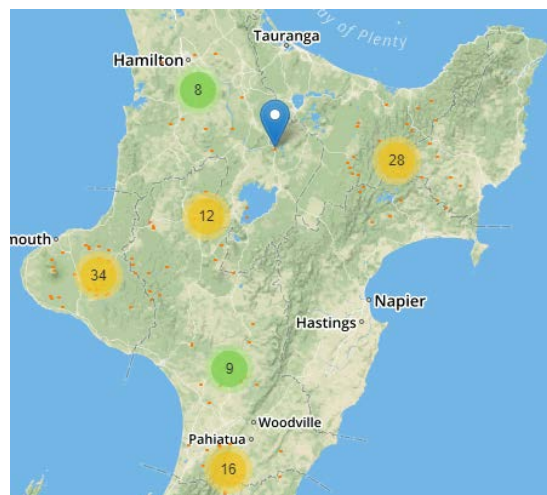

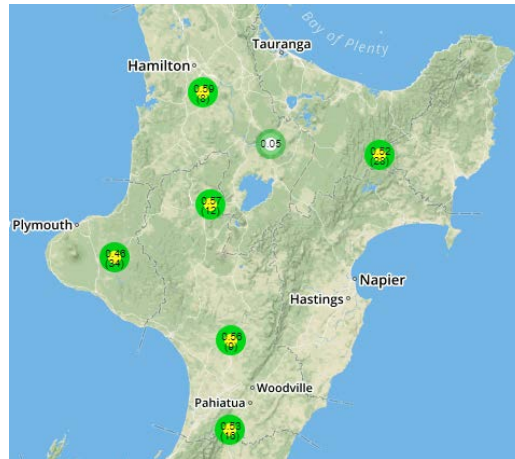Figure 7. Prototype 3 – Marker clustering using Leaflet and Mapbox Streets

Figure 8. Prototype 4 –Customised marker clustering showing Cacophony Index average and data point count

4.4 Prototype Mock Data

Data generated by prototype 2 was stored in a JSON file as an array of reading objects in the following format.

- Cacophonometer id
- Reading id
- GPS location for reading
- Date and time of reading
- Cacophony Index

5. Evaluation

Fry (2008) describes a process for visualisation design and development composed of the stages of Acquire, Parse, Filter, Mine, Represent, Refine and Interact. Fry states that "one should not become a slave to this process" suggesting the need for flexibility and indeed the approach here has departed from this in certain cases. The Acquire and Parse stages focus on obtaining source data and providing structure to it. In this study data required for visualisation does not currently exist yet so rudimentary mock data as outlined in section 5.4 has been created to drive the visualisation. Future prototypes may require a suitable database to provide multidimensional data with pre-calculated aggregations as well as base data to drive filtering requirements.

Refine and Interact focus on improving the visualisation to make it clearer and more visually engaging, along with methods for manipulation and control of the visualisation. Early stage prototype informal evaluations via email have been anecdotal at best with comments of '"that's great"' but no real user feedback to guide improvement. However later live prototypes privately hosted on cacophony servers have drawn better quality feedback. Such feedback has further helped to refine the drill down interaction and cacophony index plot presentation to better make use of preattentive processing.

The 1 km grid resolution requirement for the visualisation requires that 1km be measured in degrees of latitude and longitude. The issue with this conversion is that the earth is an ellipsoid and not a true sphere. Conversion results vary at the equator when compared to poles. One degree of latitude = 69 miles (111 kilometres) approximately, at the equator this results in 68.703 miles (110.567 km) and at the poles 69.407 (11.6999 km). For our visualisation the 1 km resolution requirement does not need to have a high degree of accuracy as the requirement is to simply mask the location of instrumentation equipment such as Cacophonometers.

Prototype 1 used D3.js to create an overlay showing color gradient plots. Further work on this prototype showed that gradient plot size needed to be adjusted in conjunction with a change in zoom level or plots needed to be clustered, which would require further coding in D3.js. Considering that map based visualisations are common it was assumed that nearly equivalent functionality would be found in a map API, and tests revealed that Leaflet provided it which triggered a change in technology from Google maps and D3.JS to Leaflet and Mapbox.

5.1 Incomplete Work

Prototype growth required modifications in the code and the application development platform. All prototypes had been contained within a single HTML5 project which did not allow for javascript to be easily used in a modular manner. Application code was refactored, and code library dependencies were added as node package manager modules (NPM, 2017). A grunt script was created to allow for automated building of application code for both development and for a production environment (Grunt, 2017). Production builds require minification to reduce overall size of code transmitted to the browser.

The prototype does not draw on processed data from the Full Noise sub-project as the machine learning research required to produce the processed data is still ongoing. While two overlays on a single map were initially specified, this requirement has been modified after feedback from Cacophony project participants and two separate map visualisations are now required. While only the Cacophony Index map visualisation exists within the CacophonyViz prototype, it is expected that creation of the map visualisation for predators will now be trivial.

6. Conclusions

The work completed thus far has provided a prototype that fulfills requirements for displaying Cacophony Index data collected from Cacophonometers. Visual clustering techniques have been used with customised interaction and display. Technology such as D3.js, Google Maps and Leaflet have been evaluated and tested for their suitability during prototype construction. A data generator has been created to produce large data sets in JSON format. Insights into the structure and delivery of data have been gained.

Opportunities for future work exist such as investigation into and implementation of animation techniques that are appropriate for the potentially chaotic nature of the data. Animation of a colour gradient in conjunction with movement may be better suited to showing changes in data over time. Examination and implementation of suitable interaction methods for filtering data by time and by bird species can be

conducted. For appealing to a diverse audience consideration of Matauranga Maori can be examined with the possibility of incorporating Maori knowledge on bird species. User interface issues still need to be addressed to have CacophonyViz run on various device and display types. Based on the experiences of this research the direction for the final version of CacophonyViz will be to have a cohesive suite of data visualisations and so further prototype development should yield opportunities for the application of additional visualisation techniques.

References

Agafonkin, V. (2017). *Leaflet: an open-source javascript library for mobile-friendly interactive map*s. Retrieved from http://leaflet.com

Bootstrap. (n.d.). Retrieved September 3, 2017, Retrieved from http://getbootstrap.com/

Bostock M. (2017). *D3 data-driven documents*. Retrieved from https://d3js.org/

Brockiem B. (2015). *Introduced animal pests*. Retrieved from https://www.teara.govt.nz/en/introduced-animal-pests

Fry, B. (2008). *Visualizing Data*. Sebestapol, CA: O'Reilly Media.

Grunt the javascript task runner. (2017). Retrieved August 21, 2017, from https://gruntjs.com/

Google. (2018). *Cluster map* [image] Retrieved from https://developers.google.com/maps/documentation/android-sdk/utility/marker-clustering. Used under Creative Commons Licence Attribution 3.0 Unported.

Google maps Javascript API. (n.d.). Retrieved August 21, 2017, from https://developers.google.com/maps/documentation/javascript

Severino, R. (2017). *chloropleth map* [image]. Retrieved from https://datavizcatalogue.com/methods/choropleth.html

Hevner, A.R., March, S. T., Park, J., Ram, S. (2004). Design Science in Information Systems Research. *Management Information Systems Quarterly*, *Volume 28(1)*, 75-105.

JQuery. (n.d.). Retrieved September 3, 2017, Retrieved from https://jquery.com/

Kirk, A. (2016). *Data Visualisation: A Handbook for Data Driven Design*. Los Angeles: Sage Publishing.

Kirk, A. (2011). *Visualisation deconstructed why animated geospatial data works*. Retrieved from http://radar.oreilly.com/2011/10/animated-geo-data.html

Kirk, S. (2016). *Government sets target to make New Zealand 'predator-free' by 2050'*. Retrieved July 27, 2017, from http://www.stu.co.nz/environment/82454116/government-sets-target-to-make-new-zealand-predatorfree-by-2050

Mapzen Javascript SDK. (2017). Retrieved August 21, 2017, from https://mapzen.com/documentation/mapzen-js/

Moore's Law for New Zealand Birds. (2016). Retrieved from https://cacophony.org.nz/

Nelli, F. (2014). *Hexagonal binning a new method of visualization for data analysis*. Retrieved July 26, 2017, from http://www.meccanismocomplesso.org/en/hexagonal-binning/

NPM. (2017). Retrieved from https://docs.npmjs.com/getting-started/what-is-npm

Department of Conservation. (2017). *Conservation status of New Zealand birds*, 2016. Retrieved from http://www.doc.govt.nz/Documents/science-and-technical/nztcs19entire.pdf

Ryan-Pears, C. (2017). *Full noise*. Retrieved July 20, 2017, from https://github.com/TheCacophonyProject/FullNoise

Severino, R. (2017). *Chloropleth map* [image]. Retrieved from https://datavizcatalogue.com/methods/choropleth.html

Smith, N. (2012). *Three examples of hexagonal binning* [image]. Retrieved from https://blog.mapbox.com/binning-an-alternative-to-point-maps-2cfc7b01d2ed

Troup. C. (2008). *Introduced land birds - introducing birds*. Retrieved from https://www.teara.govt.nz/en/introduced-land-birds/page-1

The Cacophony Project: summary of the development so far. (2017). Retrieved from https://cacophony.org.nz/summary-development-project-so-far