

Listening to Nature: Techniques for Large-scale Monitoring of Ecosystems using Acoustics

Birgit Planitz, Paul Roe, Jiro Sumitomo, Michael Towsey, Ian Williamson, Jason Wimmer

QUT, Brisbane, Australian

{b.planitz, p.roe, j.sumitomo, m.towsey, i.williamson, j.wimmer}@qut.edu.au

INTRODUCTION

Climate change and human activity are subjecting the environment to unprecedented rates of change. Monitoring these changes is an immense task that demands new levels of automated monitoring and analysis. We propose the use of acoustics as a proxy for the time consuming auditing of fauna, especially for determining the presence/absence of species. Acoustic monitoring is deceptively simple; seemingly all that is required is a sound recorder. However there are many major challenges if acoustics are to be used for large scale monitoring of ecosystems. Key issues are scalability and automation. This paper discusses our approach to this important research problem. Our work is being undertaken in collaboration with ecologists interested both in identifying particular species and in general ecosystem health.

DATA COLLECTION

We have experimented with several systems for acoustic data collection, including acoustic data loggers and autonomous wireless sensors. Data loggers have the advantage of being simple and work well for short term monitoring. However they do not scale well due the need for frequent retrieval, depending on storage and battery capacity. Several commercial systems are available or, alternatively, loggers can be constructed from voice recorders such as Dictaphones, recording MP3 players or even Smartphones. Constructed sensors must be enclosed in a weatherproof container and need supplemental power from rechargeable batteries.

Autonomous wireless recording devices collect sound and analyse it locally or upload to a server, possibly through several intermediate machines. The challenge with autonomous devices is power consumption, especially if they are to operate for long periods without intervention. Processing and transmitting large volumes of data requires storage space (consider that a one minute MP3 recording occupies about one megabyte), CPU power, bandwidth and by implication energy. We prefer transmitting the data rather than local processing, to ensure all of it is available for subsequent analysis. We have successfully constructed and deployed sensors which utilise Smartphones to capture and transmit sound over a 3G network [1]. When combined with a 12W solar panel and 35Ah battery, we have maintained such sensors in remote locations for 6 months without human intervention.

All recordings are stored in a relational database, along with associated metadata, such as date, recording device and user ID. The recording protocol (duration and timing of recordings) is set by the ecologist according to his/her experimental requirement.

DATA ANALYSIS

While users must have the option of listening to any recording, in reality the huge amount of data collected demands automated analysis. In some cases, it may be sufficient to convert acoustic data to spectrograms in which the desired vocalisations can be quickly recognised by their spectral signatures. But even this simple time saving device may not be practical for large volumes of data. The ecologists with whom we work require two kinds of analytical output – summary statistics that point to, for example, environmental health and recognition of the presence, absence, density, distribution of one or more target species. To this end we have developed a three step sequence; energy analysis, event analysis and template analysis.

Energy analysis locates acoustic energy in frequency and time. When energy data are averaged over weeks, months or years, it is possible to derive indices that are useful indicators of environmental health [2]. Acoustic event analysis identifies energy events in a spectrogram that can be attributed to a single source. Event analysis provides complementary information to energy analysis, because the latter is confounded by proximity of a source to the sensor. Event analysis has a number of uses including: The event distribution at a location can be used to characterise its acoustic landscape and to make comparisons with other locations; The kinds of events detected in a recording can be used to select a shortlist of templates for more detailed analysis; and the properties of an event (frequency band, duration, energy distribution) can be used to determine broad categories of non-vocal events, such as wind and machine noise. Finally, template analysis adapts speech recognition and machine learning techniques to the recognition of specific calls, such as bird species calls or Koala bellowing.

AN ACOUSTIC WORKBENCH

We have developed a web interface, to support ecologists with the analysis and annotation of their recordings [3]. Annotation is achieved using a folksonomy approach. Events of interest can be labelled with newly created or existing tags. Tags can be used to find and group recordings. Recordings are simultaneously heard and viewed using a Microsoft Silverlight control. Figure 1 shows the sonogram of a koala bellow, with a portion selected for tagging. Tag information

(e.g. start and end time) can be saved to the database. The user interface also provides a mechanism whereby users can select a template to scan multiple recordings for presence of a pre-defined vocalisation.

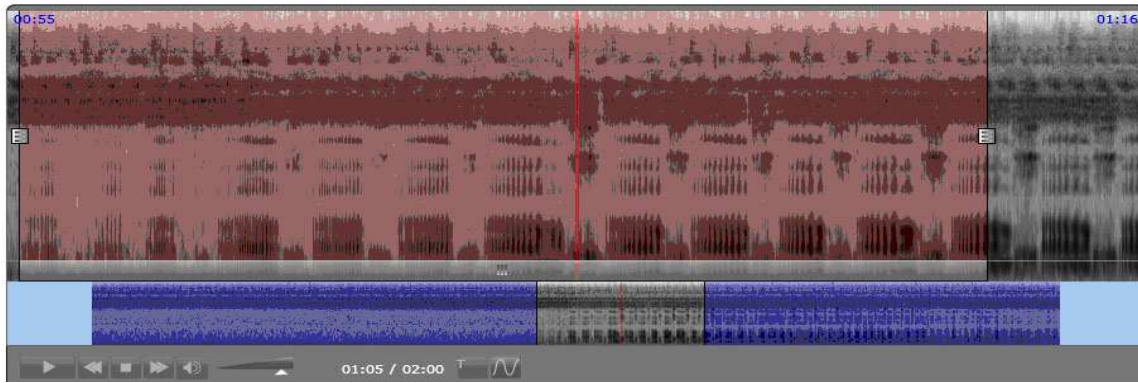


Figure 1. Koala bellow as it appears in the user interface.

RESULTS

Figure 2 displays the output of energy analysis (top) and event analysis (bottom) for a location on the edge of Brisbane city. The data were accumulated over a two week period and are displayed as the averages of one minute recordings taken every 30 minutes over a 24 hour cycle. The energy map is dominated by crickets (~8 kHz) and distant traffic (~500 Hz). However event analysis reveals a prominent morning and evening chorus and an abundance of bird calls in the mid frequency band. A comparison of these two acoustic maps illustrates the importance of having multiple views over the acoustic landscape.

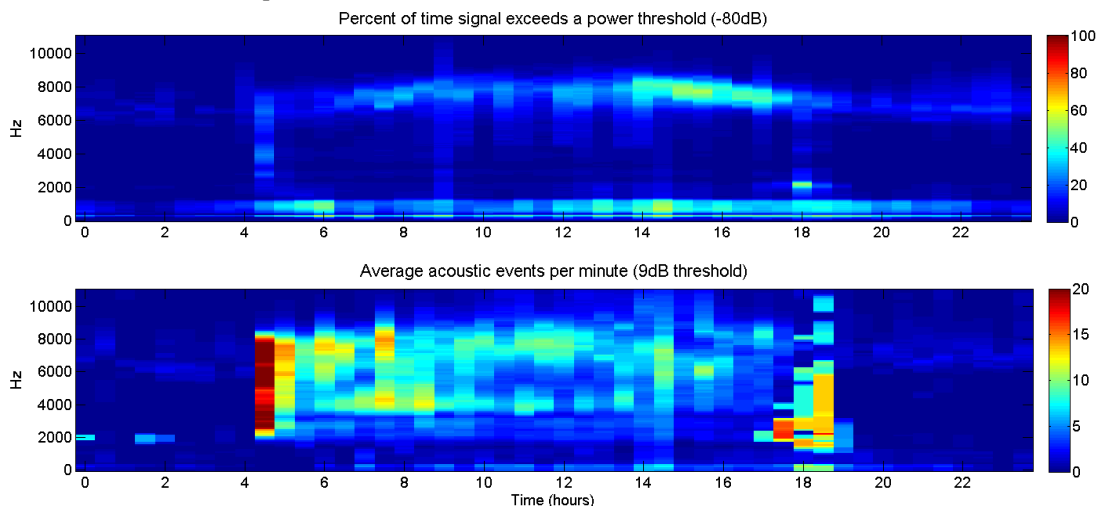


Figure 2: 24 hour maps of acoustic energy (top) and acoustic events (below).

Templates have been constructed for the recognition of the Currawong and the rare Lewins Rail. The former template achieves an accuracy of 94% on unseen test data. The latter (derived from the vocalisations of one bird) achieves only 60% accuracy due to confusion with similar bird calls in the locality but is nevertheless sufficiently accurate to have been of use in determining the distribution of this reclusive and difficult to study species.

We conclude that acoustics has much to contribute to the investigation of species distributions and ecosystem health in general.

REFERENCES

1. J. Zhang, P. Roe, R. Mason, B. Pham, “Archiving nature heartbeat using smart phones”, in *Movement-aware applications for sustainable mobility: technologies and approaches*, Editor: Monica Wachowicz, Technical University of Madrid, Spain, recommended for publication.
2. Joo, W., S. Gage, and S. Biswas. (2008) Analysis and interpretation of the “Heartbeat of the City” using acoustic signatures along an urban-rural gradient. In *the 93rd ESA, Annual Meeting, Milwaukee*. 88-155.
3. Mason, R., P. Roe, M. Towsey, J. Zhang, J. Gibson and S. Gage. (2008) Towards an Acoustic Environmental Observatory. In *Fourth IEEE International Conference on eScience*. 7-12 Dec 2008, 135-142, DOI:10.1109/eScience.2008.16.