Future Generation Computer Systems 29 (2013) 560–568

Contents lists available at SciVerse ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Analysing environmental acoustic data through collaboration and automation

Jason Wimmer *, Michael Towsey, Birgit Planitz, Ian Williamson, Paul Roe

Faculty of Science and Technology, Gardens Point Campus, Queensland University of Technology, Brisbane, Queensland, 4000, Australia

ARTICLE INFO

Article history:
Received 12 March 2011
Received in revised form
31 January 2012
Accepted 1 March 2012
Available online 5 March 2012

Keywords:
Sensors
Acoustic sensing
Data analysis
Biodiversity

ABSTRACT

Monitoring environmental health is becoming increasingly important as human activity and climate change place greater pressure on global biodiversity. Acoustic sensors provide the ability to collect data passively, objectively and continuously across large areas for extended periods. While these factors make acoustic sensors attractive as autonomous data collectors, there are significant issues associated with large-scale data manipulation and analysis. We present our current research into techniques for analysing large volumes of acoustic data efficiently. We provide an overview of a novel online acoustic environmental workbench and discuss a number of approaches to scaling analysis of acoustic data; online collaboration, manual, automatic and human-in-the loop analysis.

1. Introduction

Monitoring environmental health is becoming increasingly important as human activity and climate change place greater pressures on global biodiversity. Protecting biodiversity and developing effective conservation strategies requires a thorough understanding of natural systems, the relationship between organisms and environment and the effects of climate change [1]. This understanding is traditionally derived from field observations using manual methods such as fauna and vegetation surveys [2]. While manual fauna survey methods can provide an accurate measure of species richness they are resource intensive and therefore limited in their ability to provide the large scale spatiotemporal observations required to monitor the effects of environmental change [3,4]. In this context, there is a need to provide scientists with technology and tools to rapidly collect and analyse environmental data on a large scale [5,6].

Acoustic sensors have the potential to increase the scale of ecological research by providing ecologists with acoustic environmental 'observations' simultaneously from numerous sites over extended periods of time. This delivers far more information, more rapidly than traditional manual methods [7,8]. There are limitations to the use of acoustic sensor technology however. Most obviously, acoustic sensors are typically confined to species with audible and predictable vocalisations such as some amphibian species, insect and avian species (with some notable exceptions, for example bat species [9]). Acoustic sensors are subject to extraneous noise such as wind and rain [10]. They also produce large volumes of complex data, which must be analysed to derive detailed species information. It is the analysis of large volumes of acoustic sensor data which this research seeks to address, through the use of online collaboration and automation tools.

Analysis of acoustic sensor data is a complex task. Acoustic sensors generate large quantities of raw acoustic data which must be stored, analysed and summarised. For example, traditional avian point counts may involve ecologists making ten-minute observations at dawn, noon and dusk over a period of five days at a single site. At 2.5 h, the total observation time for a short-term manual survey is a fraction of the potential 120 h of a continuous automated acoustic sensor recording over the same period, at the same site. At long-term scales, even scheduled recordings (e.g. five minute recordings every 30 min) provide ecologists with significantly more data than manually collected long term surveys. Detecting specific species in large volumes of acoustic data is a daunting task given factors such as varying levels of background noise, variation in species vocalisations and overlapping vocalisations. Because of this complexity, a 'one size fits all' automated approach to analysis of environmental acoustic sensor data is currently infeasible.

This paper describes a novel online Acoustic Environmental Workbench, which addresses some of the challenges of manipulating and analysing large volumes of acoustic data through collaboration and human-in-the-loop semi-automation. The workbench is a web-based application that includes data upload, storage, management, playback, analysis and annotation tools all of which enable users to work collaboratively to scale acoustic analysis tasks.

* Correspondence to: QUT, 126 Margaret St, Brisbane, Queensland, 4000, Australia. Tel.: +61 7 3138 2000; fax: +61 7 3138 9390.

E-mail addresses: j.wimmer@qut.edu.au (J. Wimmer), m.towsey@qut.edu.au (M. Towsey), b.planitz@qut.edu.au (B. Planitz), i.williamson@qut.edu.au (I. Williamson), p.roe@qut.edu.au (P. Roe).

0167-739X/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2012.03.004

In Section 2 of this paper, we outline the basic architecture of our system. In Section 3, we describe our analysis techniques. Section 4 describes its implementation and Section 6 discusses the results of our implementation and future work.

2. Online environmental workbench

Part of our ongoing research has been to compare acoustic sensors with traditional manual fauna survey methods. This required us to work closely with ecologists to manually analyse large volumes of data (over 400 h) to identify vocal species for comparison with tradition field survey results. Performing this analysis identified the need to provide ecologists with a framework, which facilitates close interaction with acoustic data, and the ability to work collaboratively with other scientists. The result of this collaboration is an environmental acoustic workbench for the analysis of acoustic sensor data. The workbench is a collection of online tools, which allow users to visualise and hear recordings to identify individual species and record their analysis results. The following is the core workbench functionality we have implemented to achieve this:

- Acoustic data upload and storage.
- Acoustic data organisation and structure.
- Recording playback and visualisation.
- Recording analysis and annotation.
- Discussion and review facility.

We describe these core functions in turn.

2.1. *Acoustic data upload and storage*

Acoustic recording devices are widespread and capable of recording in many different formats (e.g. MP3, WAV etc.). The acoustic workbench provides web-based access to recordings collected from a variety of sources including, but not limited to, networked sensors and standalone data loggers such as commercially available MP3 recorders. Acoustic data in either MP3 or WAV format may be uploaded from any device capable of generating files in these formats.

All acoustic sensor data is uploaded to a centralised, online repository. This centralised approach provides a number of advantages:

- Online access and collaboration: multiple users have access to the same data and same analysis tools, enabling users to collaborate on analysis tasks.
- Data retention: all raw data is retained to allow future analysis as techniques improve, to enable long-term comparisons of historical data and to verify analyses.
- Data security and backup: all data is stored securely with regular backups and recovery facilities to prevent data loss.
- Data provenance and context retention (metadata): key experimental design details are retained to ensure accurate comparisons between data sets.

In this case however, there are a number of drawbacks to data centralisation. Most notably, accessing large volumes of acoustic data via the internet requires relatively high speed internet access and sufficient download quota. We have found that many of our users do not have access to high-speed internet. We have therefore implemented a distributed system whereby raw acoustic data is installed to user's machines and accessed by our Silverlight audio player utilising Isolated Storage. Species annotation data is still stored in our centralised database, however data transfer is reduced by a factor of four.

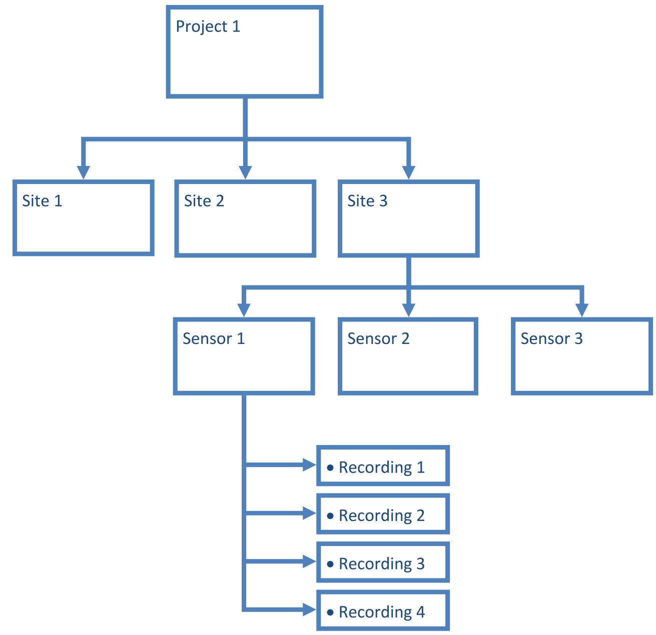

Fig. 1. Workbench data organisation and structure.

2.2. *Acoustic data organisation and structure*

The acoustic workbench allows users to browse and manipulate data in a logical, structured manner. Acoustic data are structured on a hierarchical model of Projects, Sites and Recordings. Projects are the top level. A project can represent any logical collection of experiments or studies and can be shared with other users. Each project consists of a collection of Sites. Sites are physical locations (identified by GPS coordinates), with sensors deployed at each site. Sensors are physical recording devices whose details are stored to ensure retention of experimental design details. Recordings are the raw acoustic data collected from sensor devices in the field and uploaded to the website. Fig. 1 illustrates the workbench data organisation and structure.

Users are granted role-based permissions on a project-by-project basis. These control the level of access to data and analysis tasks. Access levels include:

- None (default): user has no access to any data or any function in the project.
- Read Only: user can view/play acoustic data, can annotate spectra, but cannot upload data and cannot perform analysis tasks.
- Full: user can view/play acoustic data, can annotate spectra, can upload data and can perform analysis tasks.

These access levels allow collaborative tagging, review and discussions of tagging as required for the semi-automated analyses of data.

2.3. *Recording playback and visualisation*

Recordings can be played online using a custom-developed Microsoft Silverlight audio playback tool developed for the workbench. The playback tool plays audio and displays a spectrogram which allows the user to visualise and hear audio simultaneously. Long recordings are split into fixed-length segments (with configurable lengths) which are loaded dynamically as the player reaches the end of each segment. For example, a continuous 24 h recording can be divided into 240 six-minute segments. This allows the user to start listening without waiting for the entire 24 h recording