

Metadata Type	Element/Attribute	Required or Recommended Value	Validation Recommendations
	sampleRate	96000, unless original format is DAT or CD, then sampleRate should match that of original object	If <sourceMD>/<aes:format> is "CD" or "DAT", then <techMD>/<aes:sampleRate> must have the same value as <sourceMD>/<aes:sampleRate>. Otherwise, <techMD>/<aes:sampleRate> must be 96000. Also, match to each file's JHOVE output to ensure that file does in fact have the properties indicated in the METS
	bitDepth	24, unless original format is DAT or CD, then bitDepth should match that of original object	If <sourceMD>/<aes:format> is "CD" or "DAT", then <techMD>/<aes:bitdepth> must have the same value as <sourceMD>/<aes:bitdepth>. Otherwise, <techMD>/<aes:bitDepth> must be 24. Also, match to each file's JHOVE output to ensure that file does in fact have the properties indicated in the METS

THE UK SOUNDMAP: AN AUDIO CROWDSOURCING EXPERIMENT

Richard Ranft, Head of Sound & Vision, The British Library
Paper presented at the IASA-AMIA 2010 Conference, Philadelphia, USA

Why a sound map?

Certain kinds of sound collections are amenable to presentation online using a map interface. Plotting the geographical position of recordings on a map gives a visually attractive and intuitive way to compare and group neighbouring recordings and may be a better alternative to text-based search and browsing. For example, many of the 3.5 million recordings in the British Library's sound collections contain geospatial data: locations where recordings were actually created, music or dialect origins, or place-names referenced to in the metadata. In 2001 the British Library launched *Listen to Nature*,¹⁹ a website showcasing 400 nature sounds from around the world from its existing wildlife sound collections. The site uses static web maps built by BL staff and is occasionally added to. In 2007 the Library created the *Sounds Familiar* British dialects map,²⁰ similarly designed and built in-house. Listeners may submit new dialect recordings of their own by mailing them on CD to the British Library. The recordings are then digitized and eventually added to the *Sounds Familiar* web map. Our *Archival Sound Recordings* website, launched in 2007, makes extensive use of interactive maps²¹ to show the locations of music, environment, oral history and dialect recordings that can be played directly from the map interface. Additional recordings are added from time to time to the maps. However, none of the aforementioned websites allow direct, real-time upload and interaction by contributing recordists. That has become a more recent phenomenon of a participatory culture, shown by the flourishing number of web maps that are updated by the general public, many of them using the Google Maps API.²² Examples include *Radio Aporee*,²³ launched in 2006, which allows anyone to upload geotagged environmental sounds worldwide under a Creative Commons License; and there are now numerous other websites of sound maps of Montreal, Berlin, Seoul, Barcelona and other cities.

An acoustic survey of Britain using geotagged sound samples

In 2009 British Library sound archive staff began tests for a new kind of field recording project to aggregate user-generated digital audio content using mobile phones. Named the UK SoundMap, the project represents a radical departure from the more traditional, curator-led professional archival practices we were used to. Hitherto, these involved a drawn-out sequence, beginning with acquisition, then professional accessioning, cataloguing, preservation and eventually, if funding and copyright terms allow, online access. The UK SoundMap instead uses an informal community of guided but untrained mobile phone users and field recordists to capture and describe their environmental sounds, then enable near-instant public sharing on a dedicated website: in effect, contributors as curator-publishers.

The initiative arose when one of our curators, Isobel Clouter, secured seed funding via the Noise Futures Network,²⁴ an interdisciplinary network of UK academics, urban planners, sound artists and noise pollution experts interested in Britain's acoustic landscape who need large audio datasets for research. The project aim is to create an acoustic snapshot of the UK in 2010-11, using low-cost technologies to archive a significant body of audio data, around 3,000 recordings, of sufficient quality so that it can be analysed by the Noise Futures Network and other researchers. Sounds that are likely to disappear from our environments, including everyday noises that are often overlooked in archival acquisition programmes, are being

19 www.bl.uk/listentonature

20 www.bl.uk/soundsfamiliar

21 <http://sounds.bl.uk/Maps.aspx>

22 See for example <http://googlemapsmania.blogspot.com/>

23 <http://aporee.org/maps/>

24 <http://www.noisefutures.org/>

captured and preserved. Sounds across the UK can be contrasted easily, and the project will concurrently capture the public's impression of what sounds are significant to them today. A parallel aim is to engage the general public with the project.

Use of mobile phone technology

The notion of using a pocket mobile device designed for voice communication to capture environmental sounds for a serious archival project would have seemed ridiculous just five years ago. Yet the ubiquity of phones means that nearly everyone now also carries an audio recorder. In the UK, the number of mobile phones exceeds the population figure, since many people own more than one handset. Global sales in 2010 of all kinds of mobile phones reached 0.4 billion. Furthermore, the audio quality achievable on mobile phones, while nowhere as good as a dedicated professional recorder, has nonetheless reached a point where they can be occasionally considered as suitable acquisition tool for a limited number of "mass observation" research projects. In effect, there are far more recording devices in existence today than ever before, and their owners carry them about everywhere. This opens up opportunities to collect very large numbers of audio samples. The challenge is to harness those opportunities to yield useful research data.

The solution for our project arrived in March 2009 with the launch of Audioboo,²⁵ a free mobile application that runs on Apple iPhone and on Android smart phones. Audioboo provides a simple interface for instant recording, tagging and uploading of sound clips to a web map using the built-in mobile microphone. To date 250,000 sound clips have been uploaded by 80,000 contributors worldwide, with the top contributors from UK, USA, Germany, and these can be listened to by anyone. There have been a total of 11 million listens to date (November 2010), an average 66,000 visits daily. Audioboo has been referred to as "Twitter without typing", "User generated BBC Radio 4" or "The YouTube of the spoken word" on account of its ease of use and social networking features. The phone's geolocation and the recording date and time are automatically registered. Users can optionally capture and add a photo and key in a description.

Immediately after recording, the audio, image and metadata are uploaded automatically and wirelessly to Amazon cloud servers. In the case of iPhone recordings, each audio clip is originally recorded as a losslessly compressed mono FLAC file at 16 bits, 22 kHz that is rendered into an MP3 for immediate web presentation. The maximum recording duration is 5 minutes, and the typical upload time is around the same as the recording duration (depending on connection speeds). So the interval between starting a 5-minute recording and publishing on the web can be as little as 10 minutes.

A web browser version of Audioboo provides an alternative to the Smartphone option. It supports recording via a microphone connected to a computer then upload to Audioboo's servers, as well as uploading of an existing audio file in a choice of lossy and lossless formats (WAV, FLAC, AIF, OGG, MP3, and AAC). As the browser upload method also accepts stereo files, it is often the choice for the more discerning field recordists to contribute higher quality recordings made with professional microphones and recorders, but of course it lacks the immediacy and simplicity of the mobile option.

25 <http://audioboo.fm/>

The power of crowdsourcing and its challenges

Crowdsourcing ("a distributed problem-solving and production model"²⁶) has been used for a wide variety of activities in support of heritage collections and scientific research,²⁷ including:

- metadata enrichment, e.g. basic labelling of photo and video collections;
- editing objects for web access, e.g. image cropping;
- improving or validating automatically extracted metadata, e.g. from digitised print and manuscript images;
- searching for extra-terrestrial intelligence.

There are many other examples²⁸ not without their disadvantages. In the case of the UK SoundMap, there are technical, legal and ethical risks arising from the 'publish first, archive later' crowdsourcing model. These include:

- poor sound quality, particularly wind noise and low quality recording equipment;
- deliberate or inadvertent contributions of inappropriate recordings (e.g. copyrighted music or spoken performances, invasions of privacy, derogatory or rude language);
- inconsistent and metadata quality and/or missing metadata;
- irrelevant recordings (e.g. outside the geographical scope or subject matter).

Little can be done with some of the technical limitations of using standard consumer mobile phones. Our own lab tests on an Apple iPhone microphone showed it has a pronounced treble emphasis at about 4 kHz, and the A-D converter has a very poor or even absent brick wall filter, resulting in severe aliasing for all signals above 10 kHz. An Android-based Smartphone, by comparison, had better anti-aliasing, but a high noise floor and poor low frequency sensitivity. These limitations do not reduce the overall value of the data for most comparative studies, but may make many of the recordings unsuitable for detailed acoustic analysis.

Other risks are mitigated by the following:

- all participants are bound by Audioboo terms and conditions;
- clear instructions are provided to all contributors;
- pre-publication moderation: all recordings tagged for the UK SoundMap are moderated by listening through and checking the metadata;
- post-publication notice and take-down procedures — probably not needed because all contributions are carefully moderated in advance of publication;
- giving advice on recording techniques, in particular reducing wind noise which can be obtrusive on outdoor recordings.

On balance, the solution chosen — crowdsourcing mobile phone contributions and a unique three-way mash up between Google Maps, Audioboo and the British Library's website — was a practical and cost-effective compromise to gather a large body of data in a relatively short space of time.

26 <http://en.wikipedia.org/wiki/Crowdsourcing>

27 Examples: (a) the SETI@home project uses the enormous processing power of many distributed personal computers to analyse radio telescope data. The programme runs when the screensaver is triggered (http://science.nasa.gov/science-news/science-at-nasa/1999/ast23may99_1/); (b) the Waisda? video labelling game is used to label TV programmes held at the Dutch Institute for Sound & Vision ("Using a Video Labelling Game in Audiovisual Archives", paper presented at IASA 2010 by J. Oomen, L. B. Baltussen & S. Limonard); (c) reCAPTCHA uses OCR images of the *New York Times* to prevent spambot submissions to websites. Human readers must interpret the images and consistent results are pooled; (d) *Transcribe Bentham*, (www.ucl.ac.uk/transcribe-bentham/) invites anyone to transcribe the unpublished manuscripts of British philosopher Jeremy Bentham (1748-1832).

28 See for example: Long, P (2010). *How can crowdsourcing help museums and what are the implications for future practice?* London University, Unpublished dissertation.

Results

Following extensive testing, the SoundMap was publicly trialled in July 2010 in and around Sheffield city, in South Yorkshire in England, as the start of a 12 month project, extending to UK-wide coverage the following month. Recordings were published on the British Library's SoundMap within 48 hours of the contributors tagging their Audioboo recording with a 'uksm' tag. As the number of map points increased above a 100, the performance of the standard Google Maps mash-up started to decrease, so the Audioboo data feed is now configured via Google Fusion Tables API — this method potentially allows the display of 10,000s of map points at no slowdown of web pages. The original FLAC files are obtained from Audioboo under license and added to the British Library's digital library system along with images and metadata, for permanent reference.

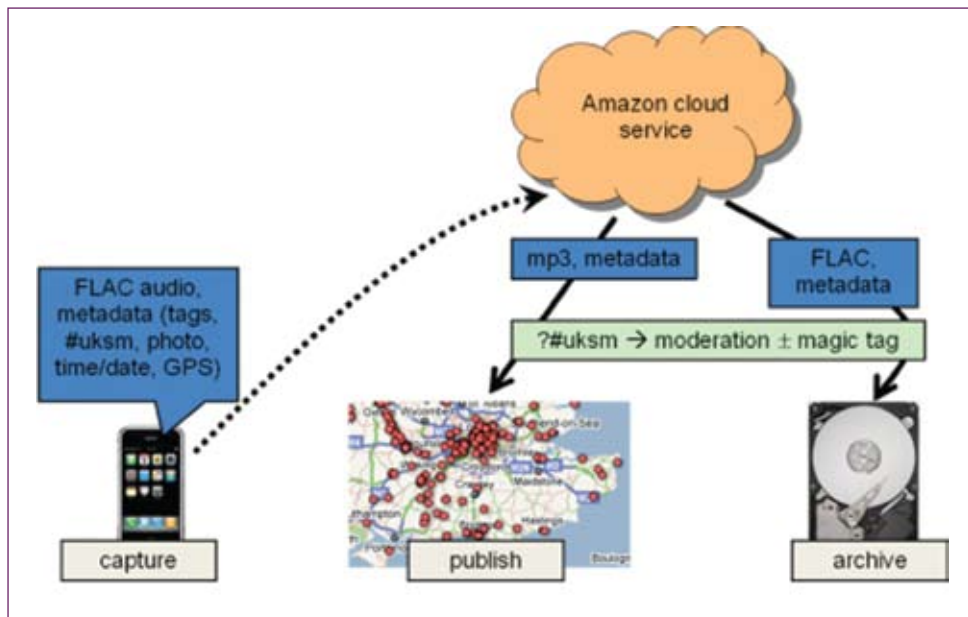


Figure 1. Schematic diagram of the recording-publishing-archiving processes

Between July-October 2010, 1,000 recordings were contributed from 260 recordists, comprising sounds of voices and direct human actions, amplified sounds and music, machinery (especially trains, buses), sounds of wind and water, and animals. 82% were made with mobiles, the remainder used dedicated recorders.

Public and media interest at the launch helped boost interest, and this was sustained through social networking sites. The feedback from online communities, the sector who inevitably were the most engaged with this project, was not always positive. "Whose brainchild was this? Will anyone really be interested in sounds originating from 2010? Don't we have anything better to do with taxpayers money?" wrote one blogger. But most were full of praise: "The increasingly innovative British Library" (BBC Technology blog), and a typical tweet was: "Can't describe the pure pleasure and pride in contributing to @UK_SoundMap".

As expected, the number of contributions declined after the summer, and is expected to pick up after the winter months and following the second planned publicity drive in spring 2011.




Figure 2. Number of accepted, rejected and 'pending' recordings, July-October

70 recordings (6% of total) have been rejected to date, due to:

Copyright violations (music/broadcast/performance)	36%
Poor quality (wind noise, low level)	22%
No geodata	19%
Obscenities, time wasters, advertising etc	19%
Recorded outside UK	3%

Conclusions

The project quickly proved its value in generating positive publicity for the British Library, creating a usable collection of environmental sounds, and demonstrating the potential of new technologies and the relatively simple methods used to bring them together. The success of the project has spawned a similar crowdsourcing project at the British Library. This new project seeks to map spoken English accents on a global scale. It is times to coincide with a temporary exhibition *Evolving Englishes*,²⁹ being held at the British Library. The English accents map requires users to recite a prescribed text so that contributions are more controlled and directly comparable.

Future challenges include extending similar community archiving projects to other kinds of sounds while ensuring that what is collected is of sufficient quality to have lasting research value. Today we cannot hear the sounds of say, Dickensian London streets, without travelling back in time. Yet in a century from now, we *shall* be able to listen to everyday sounds of today gathered in the early 21st century, thanks to the many contributors to the UK SoundMap and similar initiatives.

Acknowledgements

Thanks to Audioboo, The Noise Futures Network, Charlie Mydlarz (Salford University) and my colleagues Adrian Arthur, Chris Clark, Isobel Clouter, Andrew Pearson, David Penty, Ian Rawes and Chloe Titcombe.

[all web addresses correct as at 7 January 2011]

29 <http://www.bl.uk/evolvingenglish/>