

Sonic Visualiser: An Open Source Application for Viewing, Analysing, and Annotating Music Audio Files

Chris Cannam
Centre for Digital Music
Queen Mary University of
London
chris.cannam@elec.qmul.ac.uk

Christian Landone
Centre for Digital Music
Queen Mary University of
London

Mark Sandler
Centre for Digital Music
Queen Mary University of
London
mark.sandler@elec.qmul.ac.uk

ABSTRACT

Sonic Visualiser is a friendly and flexible end-user desktop application for analysis, visualisation, and annotation of music audio files. Its stated goal is to be “the first program you reach for when want to study a musical recording rather than simply listen to it”. To this end, it has a user interface that resembles familiar audio editing applications, a set of useful standard visualisation facilities, and support for a plugin format for additional automated analysis methods.

Categories and Subject Descriptors

D0 [Software]: General

General Terms

Measurement, Human Factors, Experimentation

Keywords

Audio analysis, Visualisation, Recorded music, Open Source

1. RATIONALE

The development of Sonic Visualiser was initiated at the Centre for Digital Music, Queen Mary University of London, for two primary reasons. First, as a means of delivering useful outcomes from music informatics research carried out within the Centre to the general public and to researchers in related domains such as computational musicology. Second, to provide a visualisation platform that researchers within the Centre can use in their own work.

Sonic Visualiser therefore contains a number of high-quality general-purpose visualisation tools such as waveform and spectrogram views, which are designed to have useful default parameters for musical analysis. It also includes simple methods for annotating audio with features such as beat times, using keyboard, mouse, or an attached MIDI device. It is extensible with a number of automatic audio analysis methods using the Vamp plugin format (see <http://vamp-plugins.org/>). Finally, it provides import and export for

annotations in formats including Music Ontology RDF for use in Linked Data applications on the Semantic Web.

The application is written in C++ using Qt and saw its 1.0 release under the GNU GPL in 2007. Since then it has been continuously developed and published as source code and in binary form for major desktop platforms, and has been downloaded over 100,000 times via its website at <http://sonicvisualiser.org>.

Sonic Visualiser has been popular with users in a number of fields: for example, it is used for visualisation and analysis in musicological research by centres such as the AHRC Research Centre for the History and Analysis of Recorded Music, where some useful tutorial material about the application has also been published (see section 5).

2. SOME FEATURES OF THE PROGRAM

2.1 Built-in Visualisations

Sonic Visualiser displays a simple waveform view by default when an audio file is loaded. The view can be zoomed and scrolled using a number of navigation tools.

A configurable spectrogram is also built in, and can be called upon directly with three different sets of parameters by using the shortcuts found on the G, M, and K keys: G for a generic spectrogram covering the full frequency range with a linear frequency scale, M for a melodic-range spectrogram emphasising information in the range in which musical detail is most usually found, and K for a peak-frequency spectrogram in which an estimate is made of the exact frequency at each peak cell based on a phase difference calculation. Like all shortcuts in Sonic Visualiser, these can be discovered from the menu descriptions or through the Key and Mouse Reference window in the Help menu.

Sonic Visualiser uses a pane and layer paradigm to allow the user to stack up any number of different views of things on top of one another, all aligned in the time axis. It can show any number of different panes (distinct areas of the window, vertically stacked) and each pane can display any number of layers (overlaid data representations).

2.2 Annotation Layers

In addition to the built-in visualisations, Sonic Visualiser can display editable annotation layers of various types, such as time instants (moments with no associated value, such as beat locations), time-value plots, labels, and images. As an example, one might annotate points in a recording by adding a time instant layer and clicking with the mouse, or by using the keyboard to “tap” beats while the audio is playing.




Figure 1: Sonic Visualiser

2.3 Plugins

Sonic Visualiser supports the Vamp plugin format to provide additional analysis methods such as beat trackers, structural segmentation, key estimation and so on. Many of these are available for download from the Vamp plugin website at <http://vamp-plugins.org>; please see that website for more information on obtaining and installing plugins. Also supported are LADSPA effects plugins, which generate a new audio file from the original one which can then be used for analysis in the same way.

To use a plugin in Sonic Visualiser, find and select it on the Transform menu. The results will be displayed in a new layer of the appropriate type on the currently active pane. To display it in a separate pane, create the new pane (for example with the N key shortcut) before running the plugin.

The task of the plugin is to take audio input and parameters, and from them calculate values for display: it does not actually do any drawing or user interaction itself.

Sonic Visualiser also features a facility for searching the descriptions of plugins that are installed or that are known about but not yet installed: this can be reached through “Find a Transform” on the Transform menu.

3. AN EXAMPLE

Figure 1 shows Sonic Visualiser in use, with a stereo music audio file loaded.

In the topmost of the four panes the audio is shown as a waveform, overlaid with a colour gradient from a note onset detection likelihood function in an additional layer.

The second, larger pane contains one of Sonic Visualiser’s built-in visualisations, a spectrogram in peak-frequency estimation mode. Over this is shown the output of a Vamp plugin which performs key estimation, with a line and a label for each estimated key change.

The third pane contains the output of a chromagram Vamp plugin which calculates pitch chroma distribution, and in the bottom pane the two layers contain a spectral centroid plot and, in the coloured regions, a segmentation by timbre, both calculated by further Vamp plugins.

4. OBTAINING SONIC VISUALISER

Sonic Visualiser is distributed in binary form for Linux, OS/X, or Windows, from the project download page at <http://sonicvisualiser.org/download.html>.

It can also be compiled from the source code, although its extensive library dependencies can make it difficult to build on platforms other than Linux. The source code package is available from the same location as the binaries; compile instructions can be found in the INSTALL and INSTALL.ubuntu files in the package.

5. MORE DOCUMENTATION

There are several sources of tutorial and reference documentation:

- Example and demo videos, collected at <http://sonicvisualiser.org/videos.html>
- An in-depth tutorial about the use of Sonic Visualiser for musicological analysis, at http://www.charm.rhul.ac.uk/analysing/p9_0_1.html
- Further tutorials and other documentation collected at <http://sonicvisualiser.org/documentation.html>
- Reference manual and key shortcuts index available from the Help menu of the application itself.

6. ACKNOWLEDGEMENTS

Sonic Visualiser was initiated at the Centre for Digital Music, Queen Mary University of London. Some development was funded by the Engineering and Physical Sciences Research Council through the OMRAS2 project EP/E017614/1 and the SeMMA project GR/S84743/01; by the European Commission Information Society Technologies through the SIMAC project IST-FP6-507142 and the EASAIER project IST-FP6-033902; and by the Arts and Humanities Research Council through the Research Centre for the History and Analysis of Recorded Music.