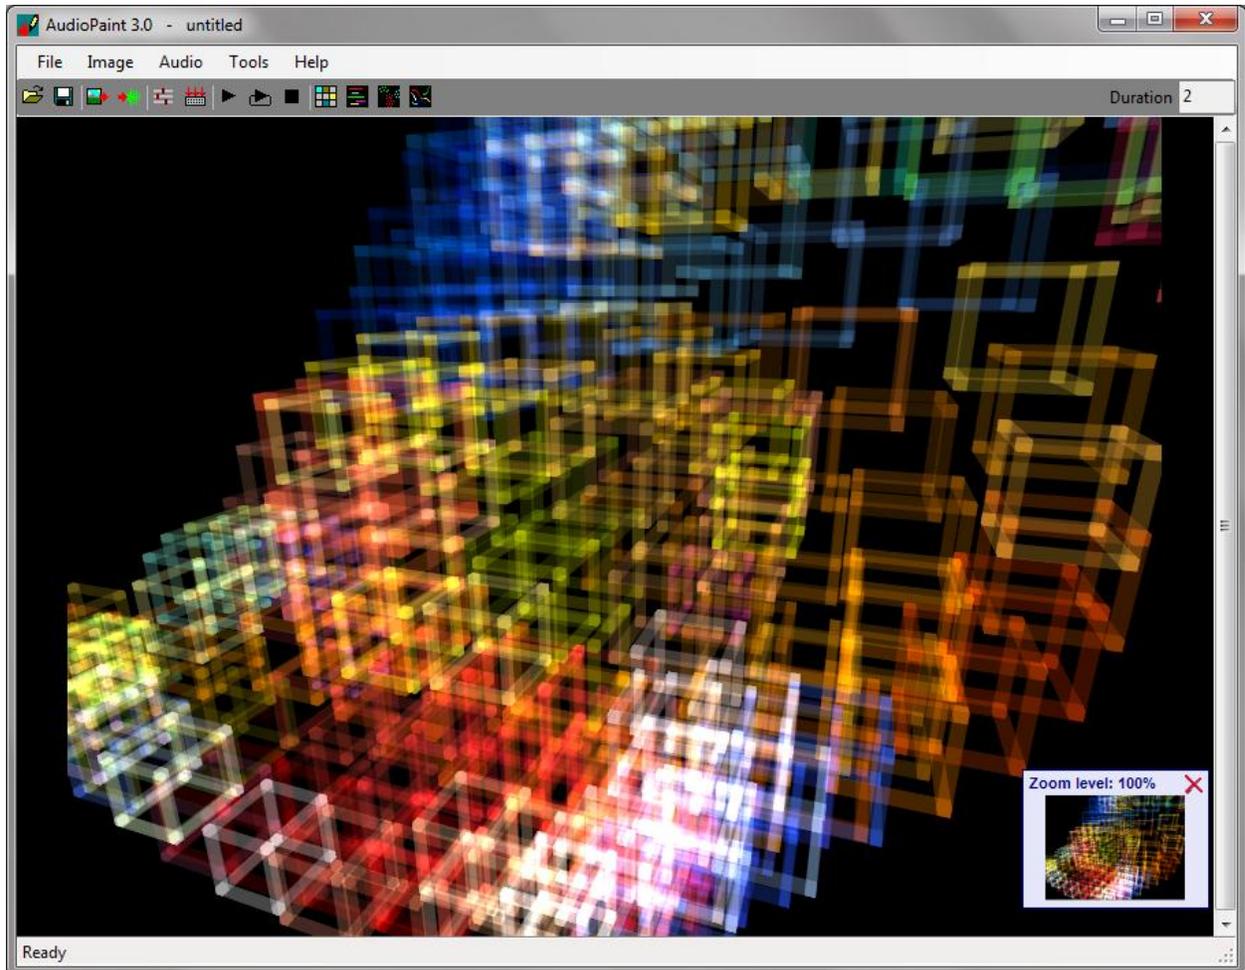


AudioPaint 3.0

User Documentation



This is a short documentation for AudioPaint 3.0. Make sure this is the latest version by checking the web page: <http://www.nicolasfournel.com/audiopaint.htm>

Please note that this program has been initially written for my personal use. Therefore, there is no warranty of any kind. Use it at your own risk!

AudioPaint (c) 2002-2011 Nicolas Fournel

Contents

1 - Introduction

- 1.1 - What is AudioPaint? *page 3*
- 1.2 - How does it work? *page 3*
- 1.3 - How to use it? *page 3*

2 - Reference

- 2.1 - The File Menu *page 4*
- 2.2 - The Audio Menu *page 5*
- 2.3 - The Image Menu *page 8*
- 2.4 - The Tools Menu *page 10*
- 2.5 - The Help Menu *page 16*
- 2.6 - The View Menu *page 16*

3 - Support

- 3.1 - Program History *page 17*
- 3.2 - Known issues *page 18*
- 3.3 - What's next? *page 18*

1 - Introduction

1.1 What is AudioPaint?

AudioPaint generates sounds from pictures. The program reads JPEG, GIF, PNG and BMP files and translates each pixel color and position into frequency, amplitude and pan information. AudioPaint can therefore be considered as a massive additive synthesizer.

1.2 How does it work?

A picture is actually processed as a big frequency / time grid. Each line of the picture is an oscillator, and the taller the picture is, the higher the frequency resolution is. While the vertical position of a pixel determines its frequency, its horizontal position corresponds to its time offset. By default, the color of a pixel is used to determine its pan, the red and green components controlling the amplitude of the left and right channels respectively (the brighter the color, the louder the sound), and the blue component is not used. The action of each component can be modified in the Routing section of the Audio Settings window. Starting with version 2.0, AudioPaint can also convert the color components into HSB values, and use hue, saturation and brightness instead of red, green and blue.

1.3 How to use it?

Depending on the picture, several hundreds (or even thousands) of oscillators can be generated simultaneously. Due to the large amount of data to be processed, AudioPaint is not real-time. Once you chose your picture and the parameters, a sound is generated. Then you can save it as a .WAV file for further processing in a sound editor such as Sound Forge / Wavelab etc.... or you can import it as an audio track in a sequencer (Sonar, Acid etc...).

Of course the sounds generated will vary depending on the type of picture used. I tested AudioPaint with pictures taken by the Hubble Space Telescope and available at the Hubble Heritage Gallery of Images: <http://heritage.stsci.edu/gallery/galindex.html>. The results are complex and fascinating futuristic soundscapes, especially with long durations. Generative graphics, which usually exhibits a lot of colors while having a geometric composition, work especially well. During development, I was using works from Dave Bollinger for example (<http://www.flickr.com/photos/davebollinger>).

But you can really try any type of picture: maps from Google are also a good starting point for your sonic exploration, as well as paintings of Picasso, Monet and other great artists. Starting with AudioPaint 2.0, it is even possible to acquire random (but suitable) pictures from the Internet directly from the program.

And of course, don't hesitate to create pictures especially for AudioPaint, in PhotoShop or a similar graphic package. By playing with geometrical shapes and colors, you can create very interesting sounds.

2 - Reference

2.1 The File Menu

Open Project...(Ctrl+O)

Opens an AudioPaint project. A project contains a path to the picture as well as the audio settings defined in the “Audio” menu. It is in XML format, which is easy to read, and other tools can import or export AudioPaint project files easily.

Save Project (Ctrl+S)

Saves the current AudioPaint project.

Save Project As...

Saves the current AudioPaint project under a new name.

Import Picture... (Ctrl+I)

Opens a picture file. AudioPaint recognizes the following formats: JPEG, BMP, GIF and PNG.

Export Picture:

Saves the current picture. This might be needed if you want to keep a picture that you generated with the commands from the “Tools” menu. Note that if you did so and want to save the project, AudioPaint will make you export the picture first. Pictures are saved in BMP.

Export Sound As... (Ctrl+E)

Saves the generated sound as a wave file, in 44.1 kHz, 16-bit, stereo.

Exit

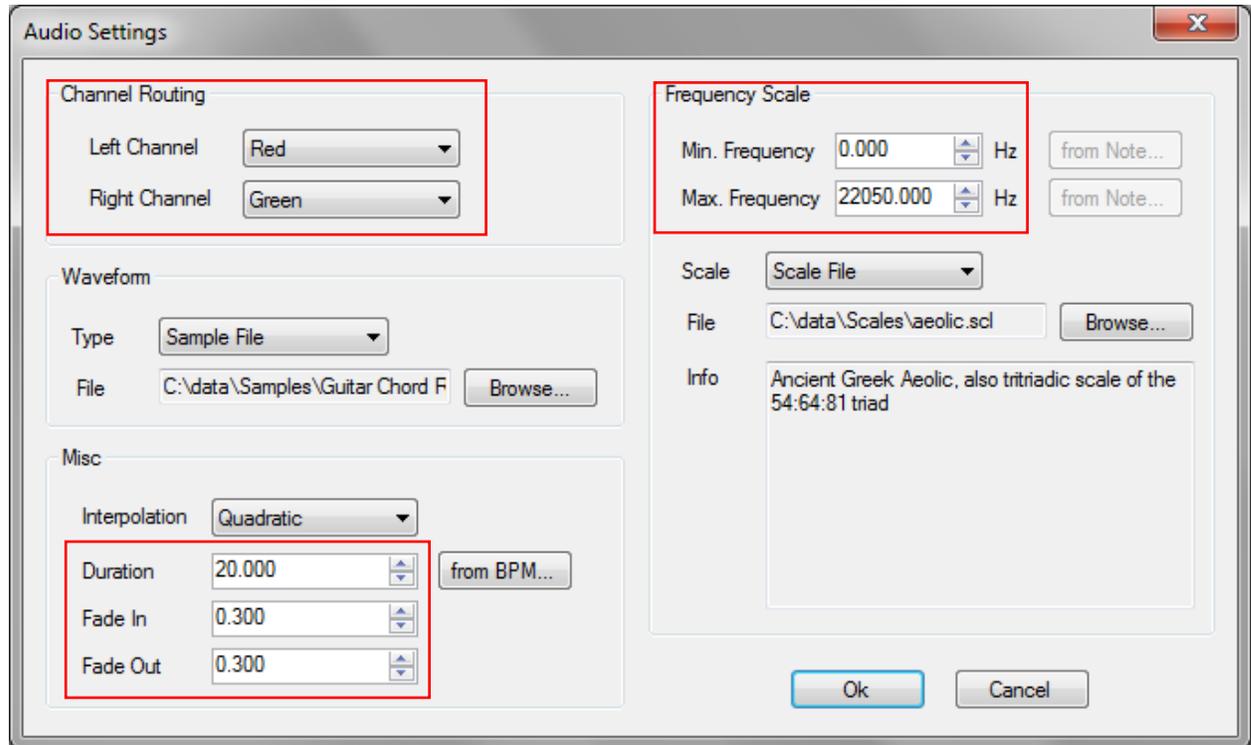
Quits the program after a confirmation message.

The File menu also includes a list of the most recently opened projects, for quick access. The number of project files listed can be selected in the program’s options (accessible from the “Tools” menu).

2.2 The Audio Menu

2.2.1 The Audio Settings window (Ctrl+A)

This is where you will configure the synthesis parameters used to generate the sound corresponding to the picture.



Channel Routing

In AudioPaint, each line of a picture corresponds to a running oscillator, and the color of a pixel determines the amplitude of this oscillator at a given time. The channel routing lets you select what color component (red, green, blue) or level (hue, saturation, brightness) will be used to generate the left and right signals. If you wish to keep a channel silent, simply select “None” for it.

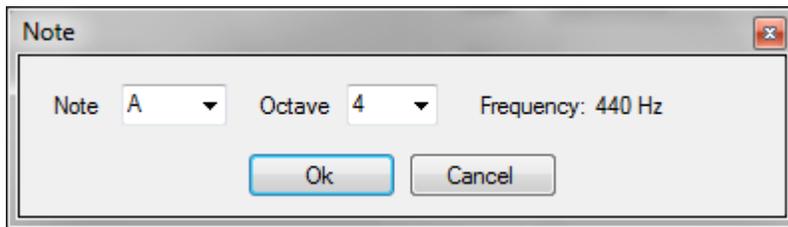
Waveform

AudioPaint uses a sine wave for the oscillators’ waveform by default. However it is also possible to use a custom waveform by specifying a sample file. To do so, simply select “sample file” for the type of the waveform with the combo box, and click on the browse button to select a wave file. The wave file can be in 8, 16, 24 or 32-bit and can be recorded at any sample rate. It can also have any number of channels, although they will be mixed down to a monophonic waveform.

Frequency scale

A frequency is associated with each line / oscillator of the picture, from the min. frequency at the bottom of the picture, to the max frequency at the top. Since AudioPaint generates sounds at 44100 Hz, the maximum frequency allowed is 22050 Hz, to avoid aliasing in accordance with Nyquist theorem. Note: at this point, when you select a sample as the waveform, nothing is done to prevent aliasing at the highest frequencies, so you may want to set the max frequency to a lower value.

Min and max frequencies can also be entered by selecting the corresponding musical note. To do that, simply click on the “from Note...” button, which will open the window pictured below:



The frequency scale determines the way the lines' frequencies are spread between the min and max frequencies. The scale can be linear or exponential. Starting with version 2.0, it is also possible to import any Scala file to specify the frequency scale. Scala is a powerful software tool for experimentation with musical tunings, such as just intonation scales, equal and historical temperaments, microtonal and macrotonal scales, and non-Western scales. It can be found at: <http://www.huygens-fokker.org/scala>. The Scala file format (extension .scl) for musical tunings has become a standard for exchange of scales, with an archive of over 3600 scales (<http://www.huygens-fokker.org/docs/scales.zip>). Using a scale file usually allows getting more harmonic and less noisy sounds. Since there are fewer available frequencies in this mode, they are repeated, from the min. frequency to the max.frequency, and then again and again, from the bottom of the screen to the top.

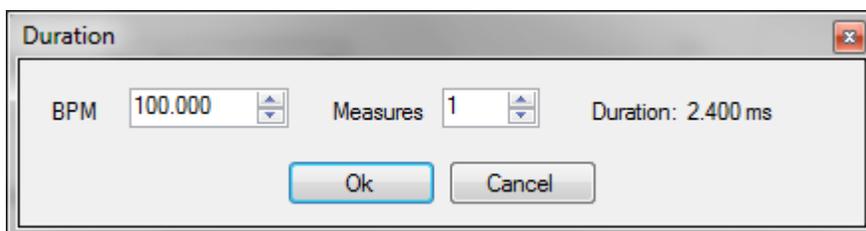
Interpolation

The type of interpolation (none, linear or quadratic) used to go from the amplitude of a pixel to the amplitude of its neighbor. By default, the interpolation is set to none. If you hear noise in the sound generated, switch to linear, or even better, quadratic interpolation. However, the better the interpolation is, the longer the calculation time will be.

Duration

This is the total duration of the sound to be generated, expressed in seconds. AudioPaint writes the sound file as it is computing it, so the sound does not need to be stored entirely in memory. Therefore this allows for very long pieces to be generated.

You can also specify the duration of the sound as a number of measures at a given BPM. To do that, press on the “from BPM...” button. The following window will appear:



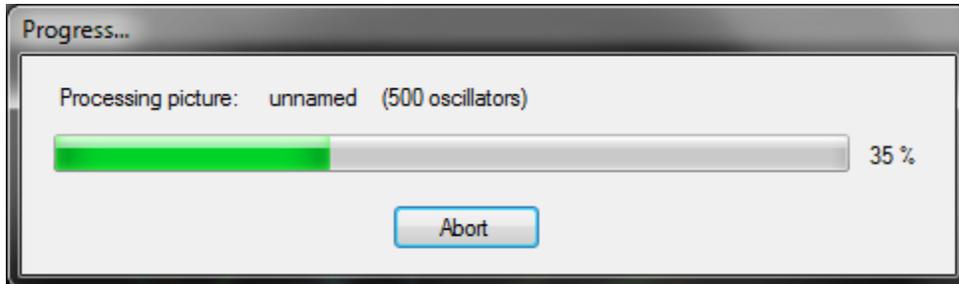
The duration can also be changed directly from the main window (in the top-right corner) for fast experimentation.

Fade In / Fade Out

AudioPaint can add a fade-in and a fade-out to the sound generated. Their duration can be set in seconds. Obviously the sum of their durations should be shorter or equal to the total duration of the sound.

2.2.2 Generate (Ctrl+G)

This is the command which generates a sound from the loaded picture, given the synthesis parameters defined by the user. A progress bar is displayed, and it is possible to abort the processing at any time.



2.2.3 Playback commands

The Audio menu also contains the playback commands:

Play (Space)

Play the last generated sound. Once AudioPaint has generated a sound, it saves it as a file named "audiopaint_temp.wav" in AudioPaint's root folder. This is this sound which will be played. If you change the audio settings but do not recompute the sound, the file will not be updated, and you will hear the last generated sound. This temporary file is deleted when you exit AudioPaint. The "Export Sound" command of the "File" menu, simply copies this file to the location of your choice.

Play in loop (Ctrl+Space)

Play the last generated sample in loop.

Stop (Ctrl+Shift+Space)

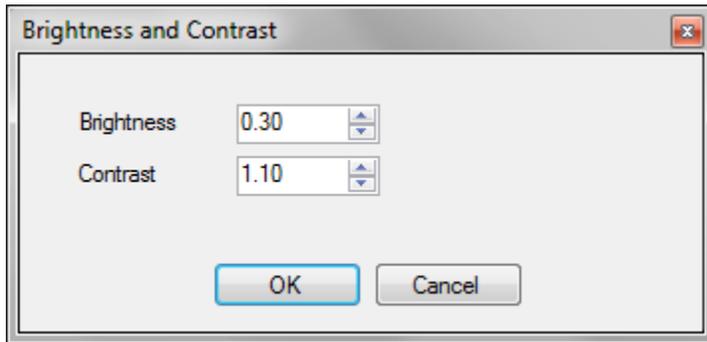
Stop the sample being played.

2.3 The Image Menu

The image menu offers many image processing functions that will alter the picture - and therefore the sound generated - in various ways.

The first set of functions allows you to *rotate the* picture or *flip* it. Thanks to these commands, you can *reverse a sound or switch the low and the high frequencies* for example.

The second set of functions lets you manipulate the colors of the picture: turning it into *grayscale* will ensure that the *left and right channels are identical*. The *negative* command will make *frequencies with low magnitudes more prominent and frequencies with high magnitudes more subtle*. *Red/green/blue component filters*, as well as the *brightness and contrast adjustments* are other ways to modify the frequency content of both audio channels.

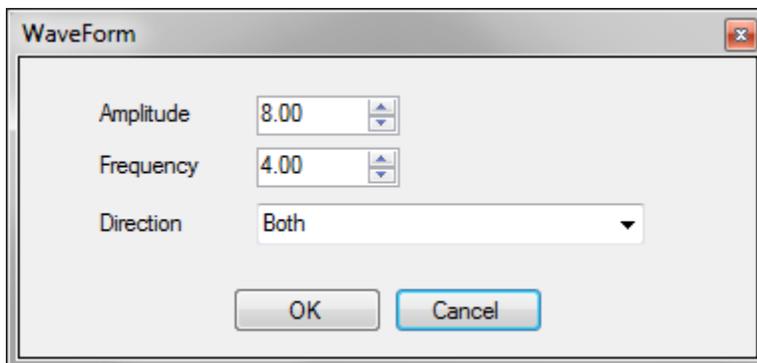


Next, a group of graphic filters is available to *sharpen or smooth* the picture, or to apply a *box blur* or *gaussian blur*. This helps making a *more defined / present sound* or on the contrary *to dilute it* a bit...

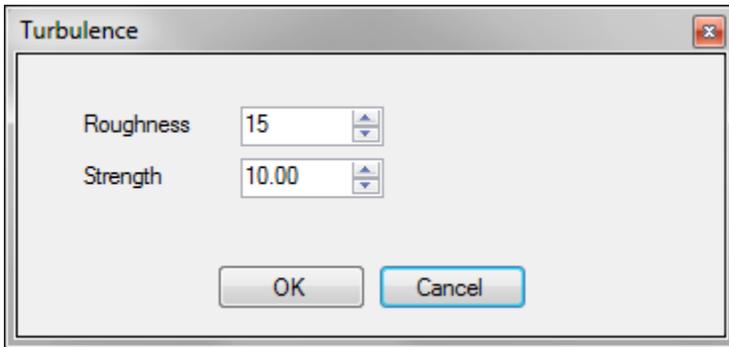
The *edge detection* and *emboss functions* can be useful to get an interesting sound from very busy images (hint: you can use *negative* after the *edge detection*), by making the *structure of the images more apparent*.

Dilate makes the pixels bigger and overlapping, while *Pixelate* divides the picture into large colored blocks. *Jitter* moves individual pixels randomly within a given radius.

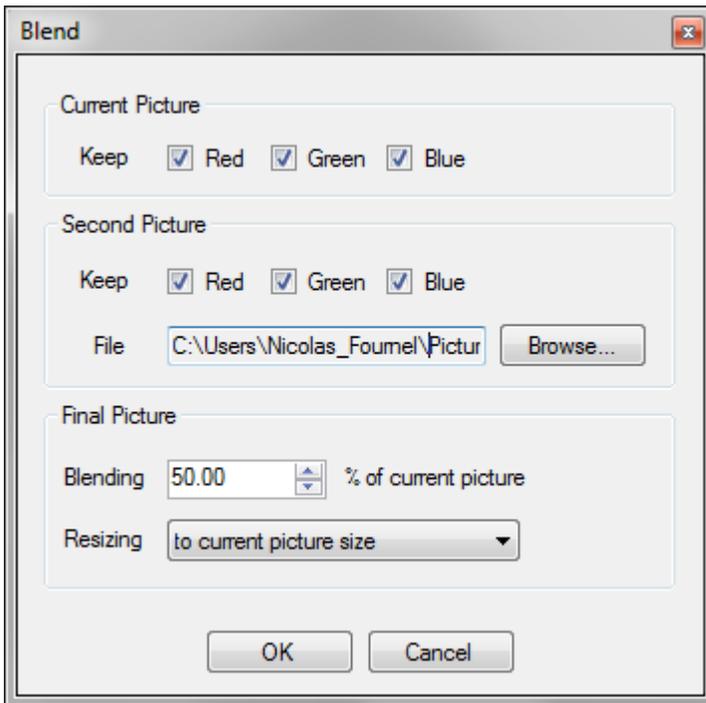
Wave will modulate the picture by a sine wave, vertically, horizontally or both (the amplitude and the frequency of the sine wave can be selected).



Turbulence modulates the picture as well but more randomly, based on a roughness and a strength parameters.



Finally, it is also possible to **blend two pictures**, which corresponds to **audio mixing**. The “level” of each image can be selected via the blending percentage, and it is possible to keep or discard any color component from both images. If the two pictures don’t have the same dimensions, it is possible to chose which one should be resized.

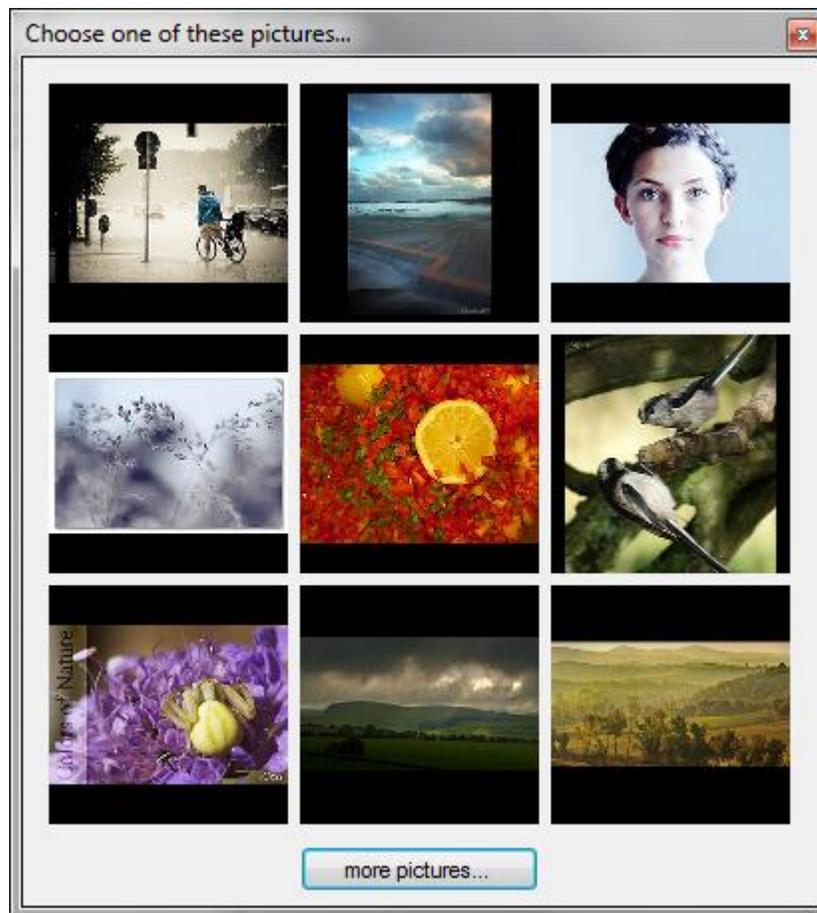


2.4 The Tools Menu

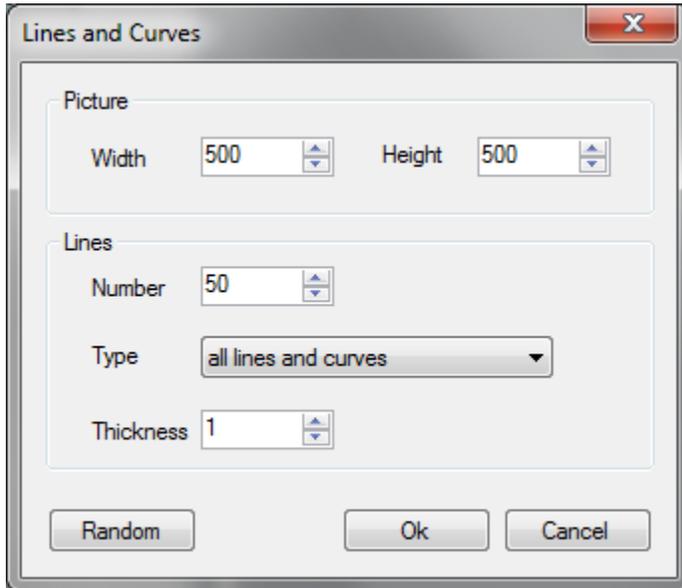
One of the most time-consuming tasks when using AudioPaint is to find or draw an interesting picture to analyze. The tools menu offers several functions to help you.

2.4.1 Random picture from the Web (Ctrl+W)

This command connects to the Internet and downloads 9 random pictures. Actually, these pictures are not totally random but have a size and a composition which makes them suitable for experimentation with AudioPaint. Once downloaded, AudioPaint displays thumbnails for these 9 pictures in a window, and you can select one of them by simply clicking on it. The pictures will be different each time you call this command. In addition, if you are not satisfied with the proposed pictures, you can click the “more pictures...” button which will download 9 new ones.



2.4.2 Lines & Curves (Ctrl+L)

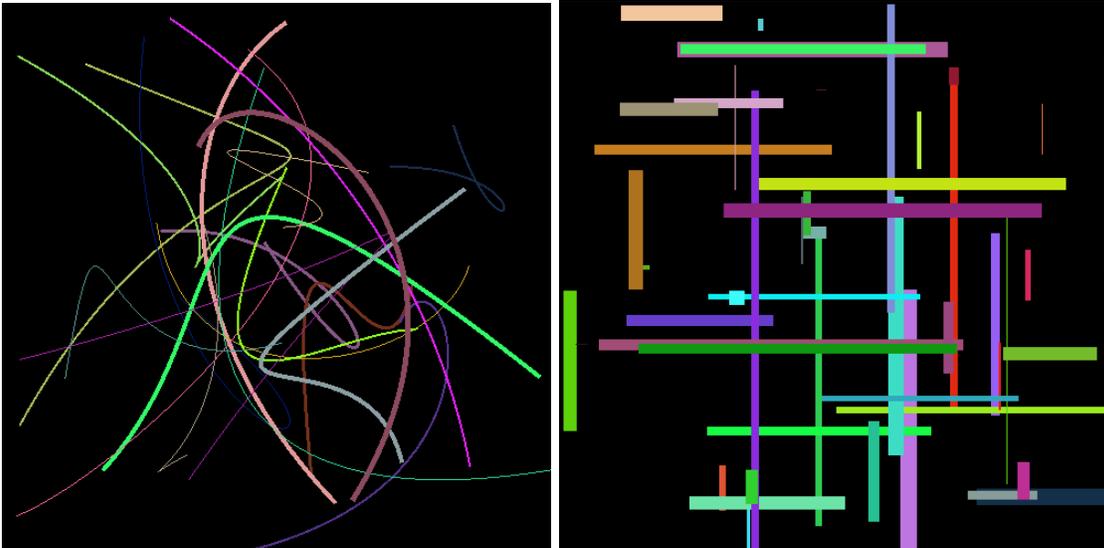


This command creates a picture filled with lines and / or curves. The size of the picture (width and height) can be selected. Horizontal lines will sound like sustained sine waves, while vertical lines will sound like chords or noise bursts. Curves will correspond to frequency sweeps.

It is also possible to draw combinations of horizontal and vertical lines, or all lines and curves at the same time. The total number of lines / curves drawn is user-definable. The thickness parameter specifies the maximum thickness of the lines or curves drawn. If it is set to 5, for example, all of the lines / curves will have a random thickness between 1 and 5.

The colors of the curves are always randomly selected. The “Random” button allows the setting of all the other parameters randomly.

Here are a couple of examples:



2.4.3 Clouds of points (Ctrl+C)

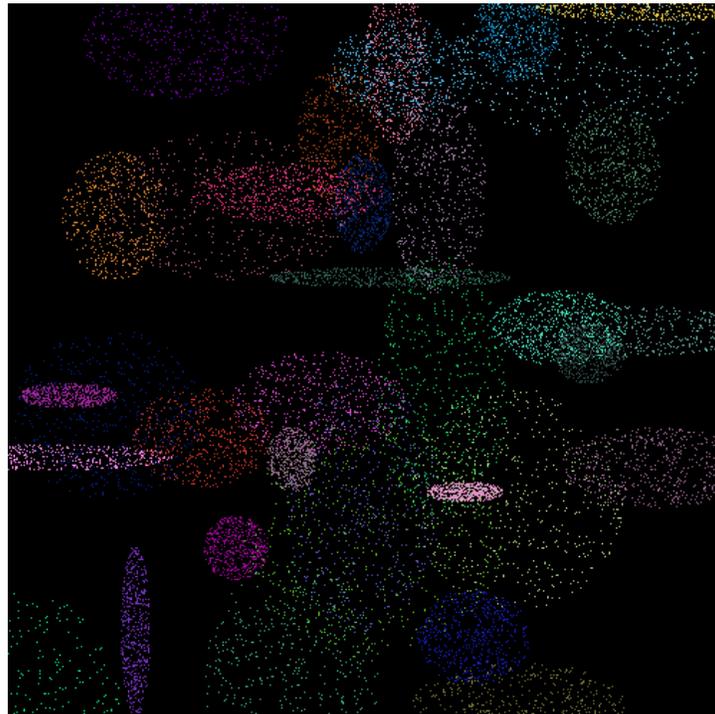
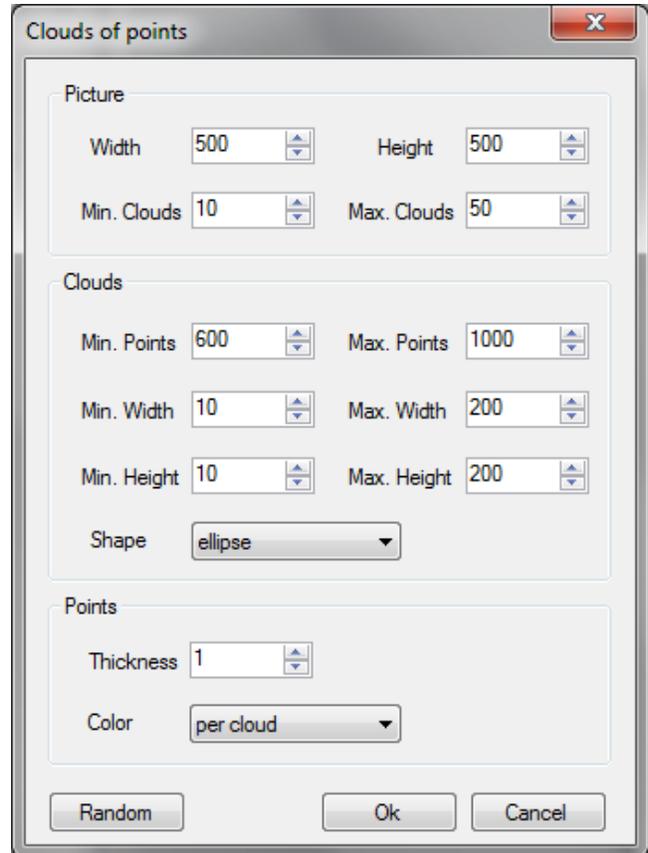
With this command, AudioPaint will create a random number of clouds of points (within the range you select). The clouds can have an ellipsoidal or rectangular shape. As for the Lines & Curves command, you will be able to choose the size of the picture to generate.

The number of points in a cloud, as well as its width and its height will also be randomly selected according to minimum and maximum values you enter.

Three color modes are available: each point can have a random color, or the points from the same cloud can share a single color, or all the points in the picture can have the same color.

This tool is especially useful to recreate granular sounds. It is very easy to generate rain, ocean or flowing water sounds with the pictures generated.

Here is an example:



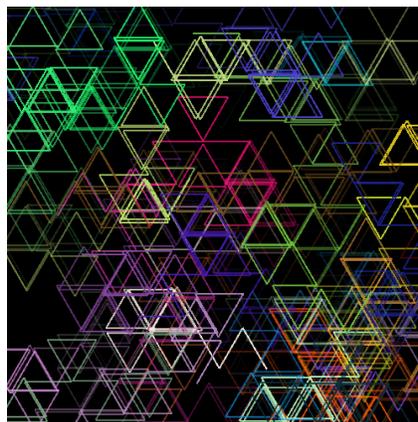
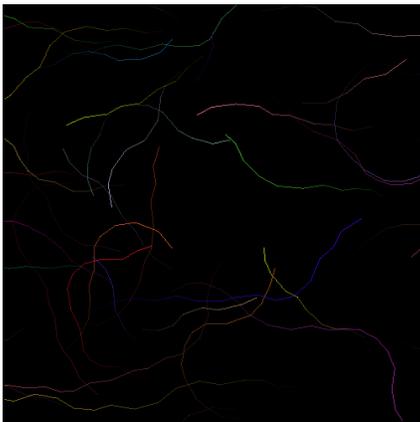
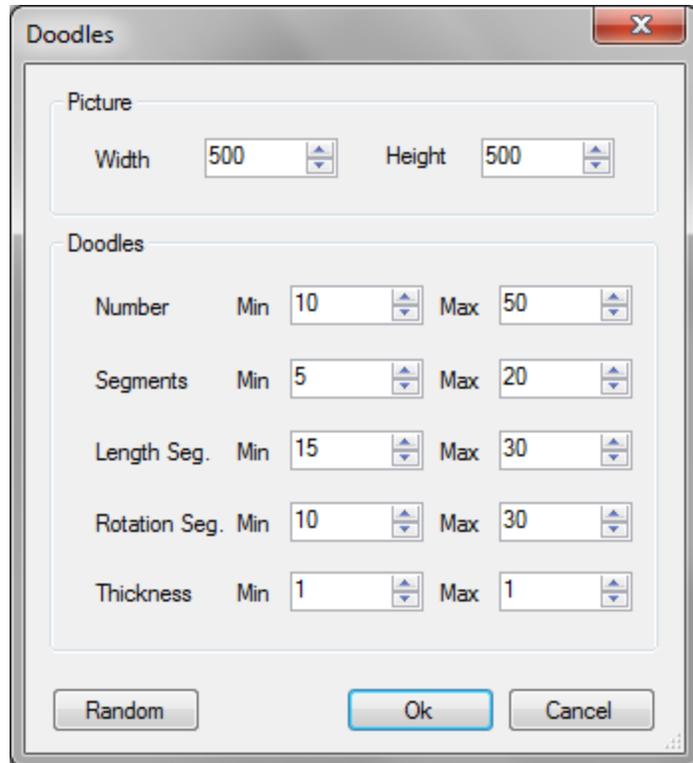
2.4.4 Doodles (Ctrl+D)

With this command, AudioPaint will draw various doodles on the screen. A doodle is a curve composed of a random number of line segments.

Each segment can have a different length and turn towards a different direction (both the length and the angle with the previous segment are chosen randomly between a minimum and a maximum value).

The thickness and the color of a doodle are chosen randomly too, with the first segment being opaque and the last one almost transparent, which makes for a more natural drawing and a better sounding result.

Depending on the settings, various types of pictures can be obtained:



2.4.5 The Batch Processor

This tool allows you to trigger the processing of a large number of pictures. It can be used for example to generate sound for every frame of a movie, or to process the same set of pictures using different audio settings each time.

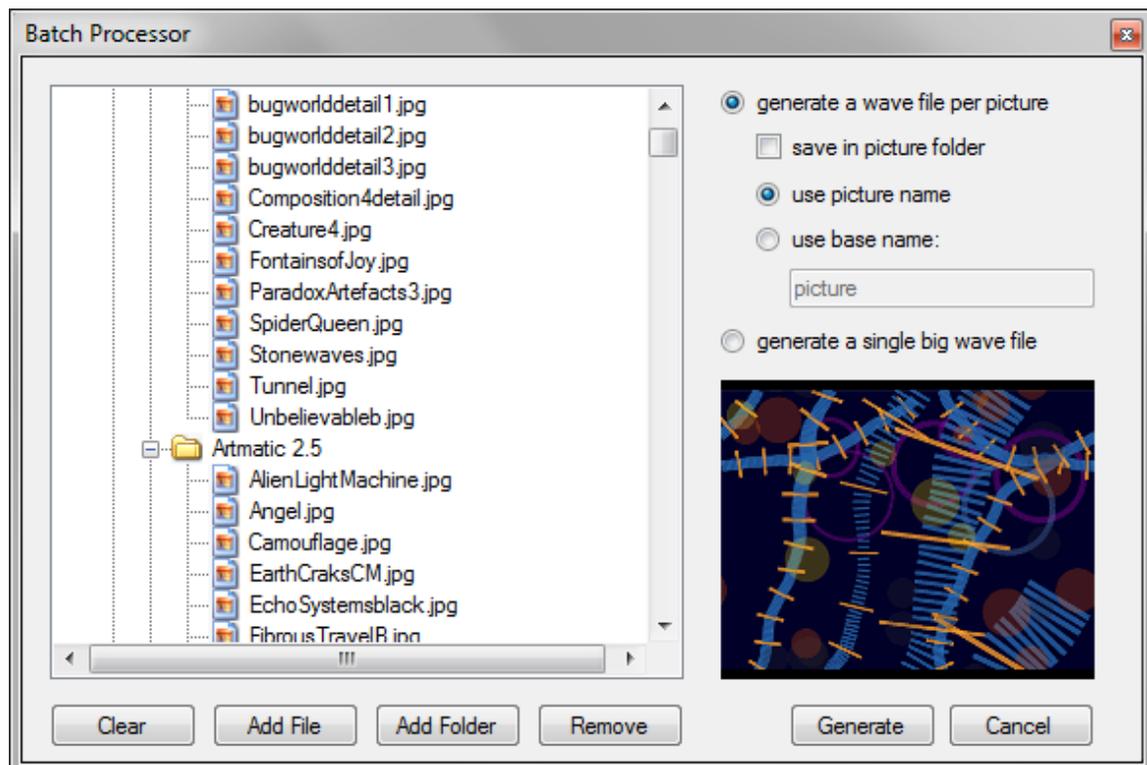
You can select one or more pictures with the "Add File" button, or add complete folders (including subfolders) with the "Add folder" button. Use the "Remove" button to remove either files or folders from the list. Adding folders which contain a large number of pictures can take time, as AudioPaint will first try to open every file to make sure it is a picture before actually adding it to the list.

Selecting the files can be time consuming, especially if they are in different folders on your hard disk. AudioPaint will keep track of the files you selected for batch processing. This means that you can exit the Batch Processor window without generating (with the "Cancel" button), change the audio settings or the general options, come back and still have all the files selected. If you want to start from scratch, simply press the "Clear" button.

If you decide to generate a wave file for each picture analyzed, you can choose to save it along with the picture (in the same folder). If not, the resulting wave files will be saved in the default sound folder (as specified in the "Options" window).

You can also choose how these wave files will be named. They can either use the same names as the pictures (the extension being replaced by ".wav"), or be based on a user-defined name, to which a numerical suffix will be added. For example, if you enter: "mysound", the wave files generated for your pictures will be named "mysound0.wav", "mysound1.wav" etc...

Instead of writing a wave file for each sound generated, you can stitch them together into a single big wave file. No fade in / out is currently applied. Therefore, if you stitch sounds coming from random pictures together, you may experience clicking.



2.4.6 The Options window

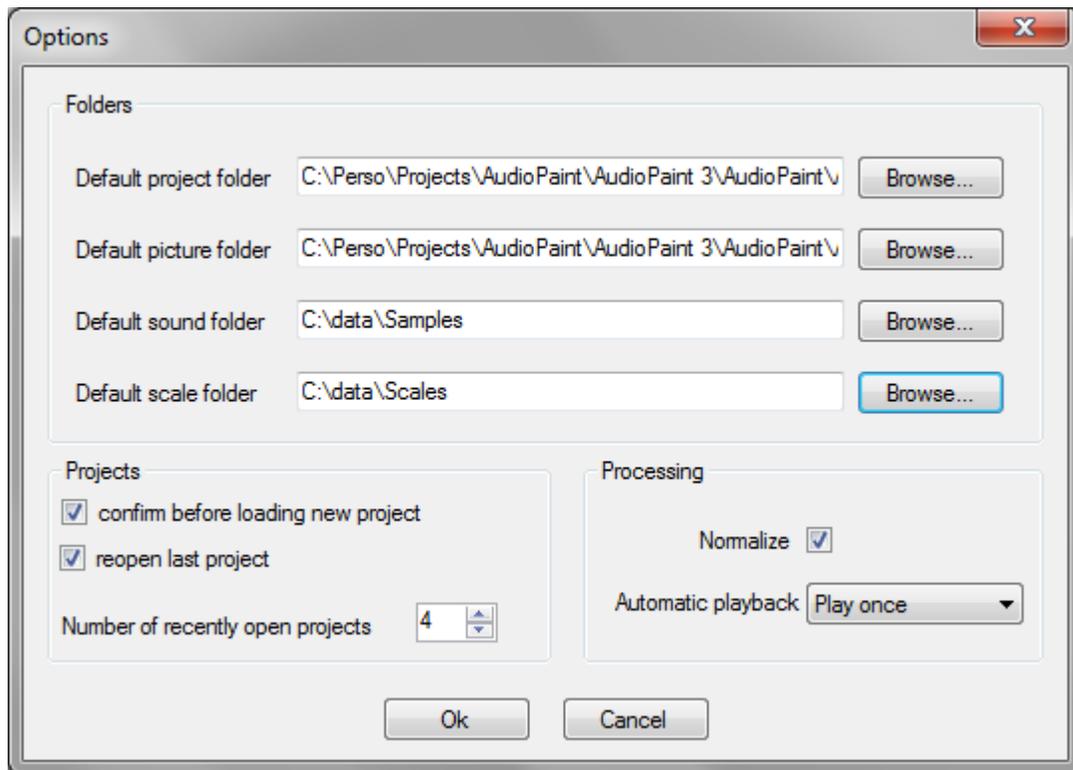
The Options window lets you select **the default folders for the project, pictures, sounds and scales.** These folders are used when you want to load / save a file for the first time. Then, the last folder where you loaded / saved a file will be used instead.

Post-processing options can also be set here:

- the playback mode determines how the generated sound will be played once it has been computed (no playback, played once, or played in loop).
- the normalize box should usually be checked. It ensures that the sound will have sufficient amplitude, but will not be clipping.

Finally, a set of options lets you configure the way you load projects. You can specify if you want a confirmation message to pop up before loading a new project, if you want to reopen the last project automatically when you launch AudioPaint, and how many projects should go in the most recently used list.

The parameters in the Options window are saved in the registry and will be recalled the next time you launch AudioPaint.



2.5 Help Menu

The “About AudioPaint” command displays the program’s version number, copyright and information, including the modules / dlls it is using.

The “Web Site” command opens your browser and connects to www.nicolasfournel.com where you can check the latest version of AudioPaint and other tools.

Finally, the “Documentation” command displays this PDF file.

2.6 View Menu

Sometimes the picture you import will not fit in the main window. In that case, the cursor will be set to a magnifier. If you click on it, you will be able to zoom in on the picture. A smaller version of the picture will also be displayed at the bottom of your window. It will let you see what part of the picture is currently displayed in the main window, and will allow you to scroll in it.

When you are in zoom mode, you can also access the View popup menu (with a right click), which offers the following options: 25%...100% (displays the picture at xx% of its original size), show whole, and fit width.

3 - Support

3.1 History

Version 3.0 – 05/10/2011

- added installer
- better support of wave files (now accepts wave files in 8-16-24-32 bits with various channel numbers)
- added "more pictures..." button to the random web pictures tool
- added Doodle generator in Tools menu
- added quick way to change duration of the sound generated (in the main window, without opening the "Audio Settings" window)
- added fade-in / fade-out audio settings
- added Image menu, including:
 - rotate / flip images
 - color manipulation: grayscale, negative, red/green/blue filters, brightness and contrast adjustment
 - sharpen, smooth, box and gaussian blur
 - edge detection, emboss
 - dilate, jitter, pixelate, turbulence, wave
 - blending of two pictures, including selection of RGB channels
- fixed "random web picture" after a change in the host broke the download code

Version 2.1 – 03/12/2008

- allows to enter the sound duration in BPM form
- allows to specify the min and max frequencies as notes
- displays scale file information
- added batch processor

Version 2.0 – 12/28/2007

- entirely rewritten (from C++/MFC to C#/.Net)
- audio engine optimized for speed
- no more limitation on the duration of the sound generated
- processing is now a background task
- processing can be cancelled
- can import Scala files for the frequency scale
- can generate a sound from hue, saturation and brightness in addition to red, green and blue
- generative algorithms enable the creation of pictures composed of lines, curves or clouds of points
- new function automatically finds suitable pictures on the web and import them in AudioPaint
- projects saved in XML for easy import / export to other tools
- floating window for better visualization while scrolling
- pictures dropped on the main window will automatically be imported
- PDF documentation

version 1.3 - 12/28/2002

- samples can be previewed in the file requester (Parameters window)
- stereo samples are now accepted for waveforms
- fixed a bug when user tried to read an invalid picture file

version 1.2 - 11/22/2002

- the max duration of the sound generated has been extended
- fixed a bug with very short sample files
- a more convenient documentation (help file)

version 1.1 - 11/17/2002

- project files (store the picture name + audio parameters)
- "play loop" and "stop" commands / buttons
- routing of red/green/blue component to left / right channel
- use of samples as waveforms for the synthesis
- new user's preferences window including: default folders, playback mode and normalization

version 1.0 - 11/03/2002

First release.

3.2 Known issues

No known issues at this point.

3.3 What's next?

The current version is 3.0. The following features could find their way into the next version(s):

- a component (blue by default) could be used to control an effect such as a waveshaper.
- integrated delay / reverb effects
- new interpolation modes for the wavetables
- new synthetic waveforms (saw, pulse...)
- more generative graphic functions
- processing of movie files (.avi, .wmv etc...)
- import a sound file as a picture by using FFT for spectral analysis
- display of frequency / volume / pan information when the mouse is over a pixel
- multicore support for faster processing
- picture browser
- drawing directly on the canvas with various brushes