VoceVista Video

User Manual and Reference Guide

VoceVista Video: User Manual and Reference Guide

Chapter 1. Introduction

VoceVista Video is a tool that helps you you understand the phenomenon of harmonic overtones of sounds intuitively and in a musical context. The program was originally developed for musicians, composers and singers who work with overtones in performance, education and musical theory. However, it should excite anyone who wants to analyze sound in a musical context.

1.1. Why a new software?

There are several good software applications for analyzing sound on the market, but we could not find any that help visualize the musical interpretation of the sound. Most spectrum analyzers are optimized for experts with a background in the physics of sound. Musicians and voice teachers are not normally trained to read physical data. Hence it is useful for musicians to relate the physical representation of the sound to the musical terms that they are familiar with.

Therefore we started to add visual aids to the spectrum analyzer that make the visual representation of frequencies accessible to musicians directly and intuitively. The addition of a musical staff and a piano keyboard already simplified things greatly. A further step to enhance the understandability of spectrograms was to switch from the linear frequency scale that is common in physics to the logarithmic scale that is more familiar to musicians (because it more closely models how the ear perceives pitch). The overtone sliders highlight the relationship between a fundamental and its overtones in a direct and intuitive way and helps simplify the understanding of the sometimes very complex spectrograms.

Many other programs are also not very practical to use. VoceVista Video has been refined through much feedback from people who use it daily to record their voice students or their instruments. It is not simply a tool for analyzing sounds: it also helps to collect and manage a large set of recordings.

1.2. What can you do with VoceVista Video?

As a software application for recording and exploring sound, especially the sound of the voice and of musical instruments, VoceVista Video helps visualize, measure and understand various aspects of your sound:

Pitch

- What pitch am I singing or playing?
- Am I in tune?
- How is my vibrato?

Timbre

- How strong are the different harmonics / overtones in my sound?
- How is my resonance?
- What is the relationship between the physiology of the voice, the physics of sound and the theory of music?

Change

- How does my sound change over time?
- How does my voice develop and what progress have I made?
- How do different recordings look and sound in comparison?

While the previous list includes most functions of VoceVista Video, it may not yet be obvious how this relates to specific tasks that you can perform with the program. Here are some more specific examples of things you can do with VoceVista Video:

- Record your voice, visualize it and listen to it.

- Measure the pitch of an instrument. Practice holding the pitch over an entire note.
- Analyze the harmonic structure of an instrument.
- Listen to and learn the harmonic series. Identify the overtones belonging to a specific fundamental. Practice singing the overtone scale.
- Transcribe the notes in a musical recording.
- Construct a musical scale through relating the constituent tones by their overtones.

1.3. Who is VoceVista Video for?

Singing teachers can use VoceVista Video like a mirror for the voice to explain to the student what he or she is doing.

Singers can use the program as a visual feedback aid to practice pitch and to support the development of vibrato and formants as well as other aspects of their vocal technique.

Voice therapists can use the program to monitor the progress of their client, and as another sensory feedback channel for certain exercises.

Instrument builders and **tuners** can easily measure and analyze their instruments with extreme precision and detail.

Instrument vendors can document the sound characteristics of their wares. This is especially useful for handmade instruments where each item is highly unique.

Overtone musicians can use VoceVista Video to study and improve their vocal or instrumental technique and to create or rehearse complex compositions, assisted by the visualization of their own overtones and by listening to changing fundamentals and harmonics.

Musical theorists find in VoceVista Video a new visual aid to illustrate the development of tuning systems and the interplay of natural overtones in classical harmony theory.

Choir conductors can recognize the role of overtones in chords and can use this knowledge to improve the brilliance of their choir and to achieve pure intonation. Recordings can be analyzed in order to systematically optimize the sound of specific passages. The timbre of individual passages can be adjusted to their current chord, which enables even amateur choirs to achieve a professional sound.

Composers that want to understand the role of overtones in their compositions can use VoceVista Video to learn the necessary background knowledge about overtone music. Very few professional musicians are familiar with the phenomenon of overtone singing, since this is not yet taught in most music schools. The software supports composers in relating *ambitus* (vocal range), vocal technique, and harmony of the singable overtones. Analyzing existing recordings gives insights into the sound and the expressive abilities of individual artists.

These are just a few examples of possible applications of VoceVista Video. Of course there are no limits to one's imagination (for example, the author analyzed the sound of dolphins to sing them as overtones).

1.4. How to use this manual

This manual can be used in several ways. If you already know what you want to do and simply need to know how certain things work, have a look through the Reference Guide. When you are using the program and need help with a settings page that is currently open, click on its Help button or press F1 to bring up the online help for this settings page.

To simply get yourself started, go to the Quickstart Guide and at the User Interface Overview.

1.4.1. Terminology used in this document

A menu choice is indicated with an arrow. File → New means: select *New* from the *File* menu.

User Interface Buttons are indicated like this: Press OK to continue.

Keyboard commands look like this: press F1 to open help.

The Glossary at the end contains definitions for many subject-specific terms used in this document.für

1.5. System Requirements

This section explains the minimum system requirements that you need to install and use VoceVista Video.

1.5.1. Operating System

macOS
: VoceVista Video requires *OS X 10.10* or later. We recommend *macOS High Sierra 10.13*.

Windows
: VoceVista Video requires *Microsoft Windows 7* or later. We recommend *Windows 10*. The program requires a 64 bit edition of Windows.

1.5.2. Computer

We recommend a recent computer with at least 4GB RAM and an Intel Core CPU. If you buy a new computer for audio work, select a model that operates as quietly as possible. A Solid State Drive (SSD) is highly recommended, as it eliminates hard drive noise.

1.5.3. Mouse

Although not essential, an external mouse with a scroll wheel can be helpful for using the program. If you use a laptop, you can use two-finger gestures on your touchpad for zooming in and out of the frequency and time range.

Zooming and scrolling can also be done with the 1 and 2 keys on your keyboard. When the mouse pointer is pointing at a scale or at the spectrogram, try tapping the 1 or 2 key to see what happens.

1.5.4. Microphone

To get started, the internal microphone of your computer is totally fine. If you want something better, get a USB microphone. The models from *Samson* are good for beginners and are inexpensive. For even better sound quality get an external USB audio interface together with a professional condenser microphone.

Another interesting option is a mobile recorder that can be used as a USB audio interface. These devices can provide excellent performance for their price category when used as a microphone and have the added benefit of being a stand-alone mobile recorder that is always at hand.

Position the microphone as far away from the computer as possible to minimize the pick up of noise from the computer, or get a silent computer. Be aware that many monitors also emit some sound, especially at lower brightness levels. Most of these ambient noises can be seen in the recording and should be avoided as much as possible.

Chapter 2. Quickstart Tutorial

This guide will show you how to record and visualize your voice or instrument with VoceVista Video. If you have no idea what to do with VoceVista Video, this is the place to start. It will also show you how to use the main user interface elements of the program. This quickstart guide is intended to be a short introduction. You can find more detailed explanations of everything in the Reference Guide.

> **Important**
>
> Some of the features shown in this guide (most notably the Frequency Filters) are only available in *VoceVista Video* and *VoceVista Video Pro*, but not in *Overtone Analyzer*. To follow this tutorial, you can download a free 30-day trial version of VoceVista Video at our download page [http://www.sygyt.com/en/select-download-version/].

> **Note**
>
> This guide is available in various formats:
>
> Online
>
> You can find an online version which you can view in your webbrowser at www.sygyt.com/en/quickstart [http://www.sygyt.com/en/quickstart].
>
> PDF
>
> You can download a PDF version of the entire manual at www.sygyt.com/en/support [http://www.sygyt.com/en/support].
>
> Video
>
> A video version of this tutorial is available at www.sygyt.com/en/videos/quickstart-tutorial [http://www.sygyt.com/en/videos/quickstart-tutorial].

2.1. Prerequisites

- You have installed VoceVista Video on your computer.
- You have a working microphone connected to your computer.

If you have not completed the above two steps, see System Requirements for more information about what you need.

2.2. Configuration

2.2.1. Apply Settings Profile

To prepare the program to make a recording, click on the profile list on the toolbar (the box that says “Select Profile”) and select the **Quickstart** profile:

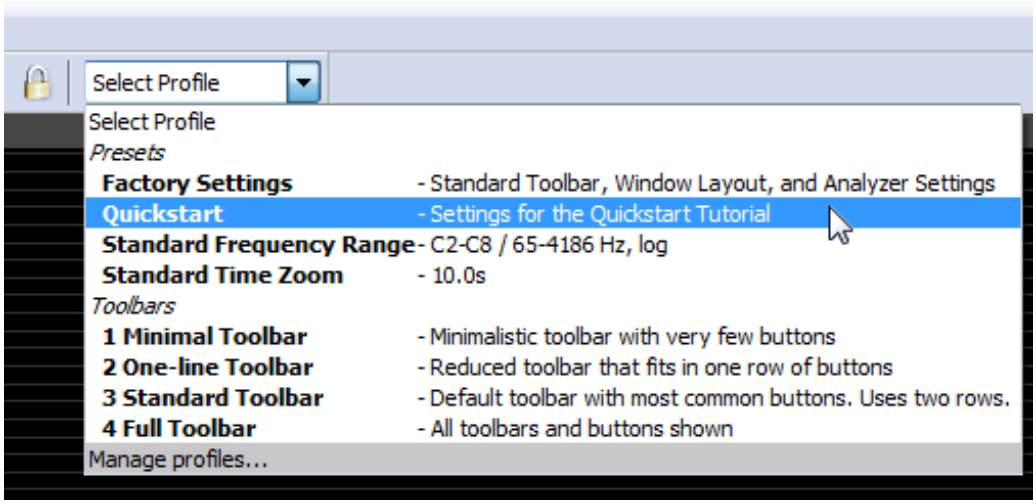

Figure 2.1. Selecting the Quickstart profile

Selecting a profile is a shortcut for applying all the settings stored in the profile. There is no difference between applying a profile and going through the program options and manually selecting the corresponding settings. Now the program should look like this:

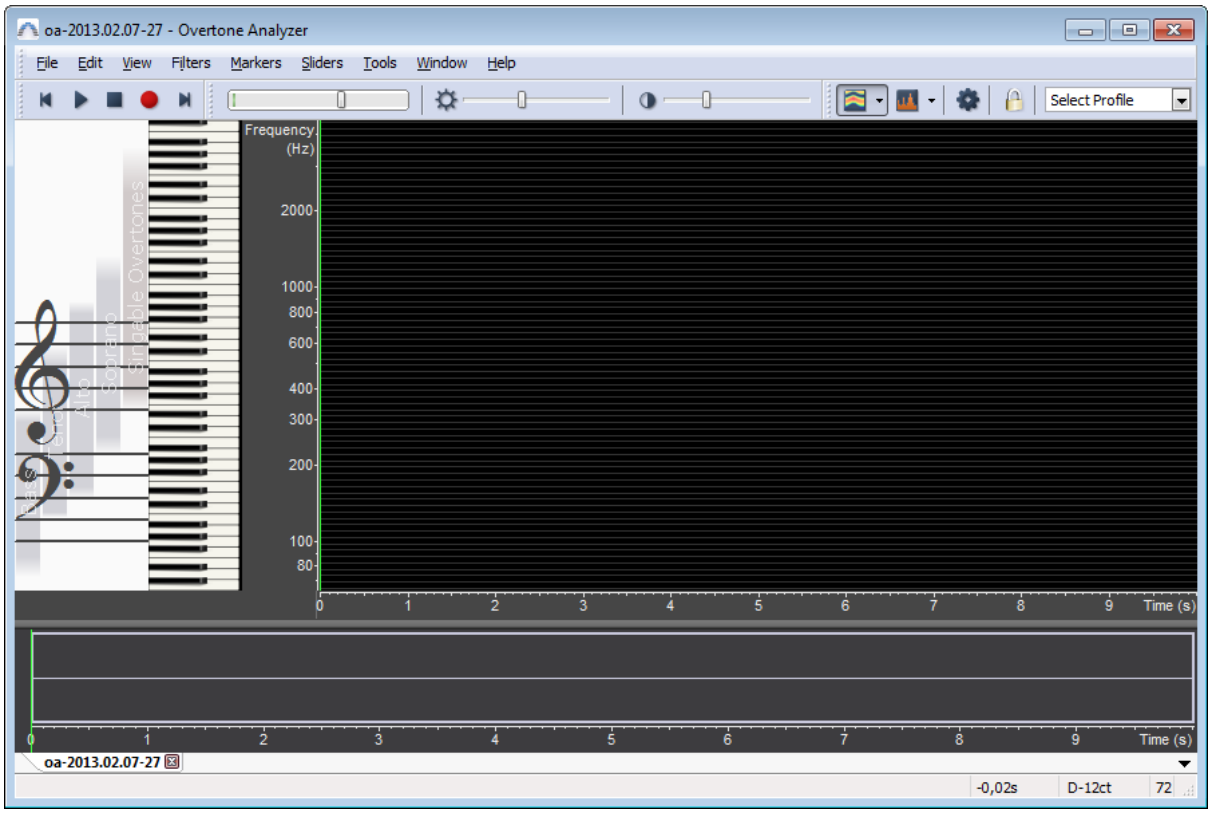

Figure 2.2. Program with Quickstart profile applied

Here you can see the empty Timeline on the bottom, the Staff View and Piano on the left and the empty Analyzer View in the main section.

Here is a closer look at the Toolbar:

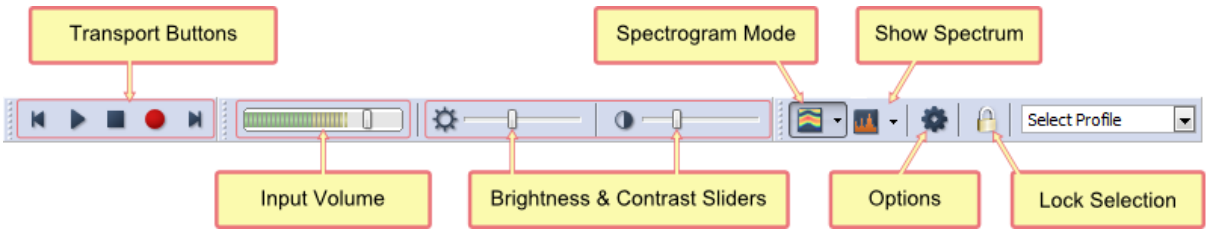

Figure 2.3. Toolbar buttons

2.2.2. Select Recording Source

Now we will check that you are recording from the correct microphone. **Right-click** on the Input Volume control on the Toolbar and make sure the correct microphone is selected. Here is an example (although this may look different on your computer due to the variety of audio devices available):

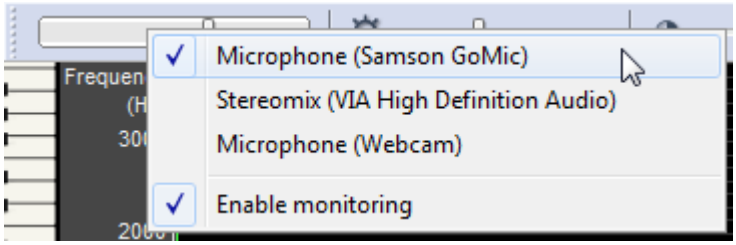

Figure 2.4. Right-click on Input Volume to select mic

2.2.3. Set Input Level

Right-click again on the Input Volume control and make sure that the last entry called Enable monitoring is checked. This will show the current input level even when you are not recording.

The Input Volume control also allows you to monitor and adjust the volume of the recorded sound. The slider that you can move sets the input volume of the microphone. The colored stripes in the background show the strength of the current signal. If you make a noise into your microphone, you should see some activity there. It is very important to adjust the recording volume correctly to prevent clipping. If the input volume reaches the the red area, reduce the input level, or increase the distance to your microphone. You should aim to keep the maximum volume of the recorded sound just at the upper end of the yellow area:

	Recording Volume too low: Analyzer Display will lack detail.
	Recording Volume too high: Analyzer will show clipping artefacts (this is worse than the volume being too low!).
	Recording Volume optimal: Signal uses most of the available dynamic range without clipping. Analyzer will show best amount of detail.

Make some test sounds into your microphone and adjust the recording level until it is getting a strong signal without clipping.

If there is no signal from the microphone, make sure it is connected correctly and that you have selected the right one. Also, some devices can only record in mono, while others can only record in stereo. You may have to click on the Options button and open the Recording Settings and experiment with different settings.

2.3. Recording

Now you are ready to record some sound. Click the Record button on the toolbar, or press Ctrl-Space to start a recording. During the recording keep an eye on the Input Volume control and make adjustments if the input level is too quiet or too loud.

Make some sounds. For example, sing the vowel 'Ah' for one breath on the same pitch. You could also click on a key of the piano keyboard, listen to it, and then try to sing the same pitch. Fill up half a screen and then press Space or click the Stop button to stop recording. This should look something like this:

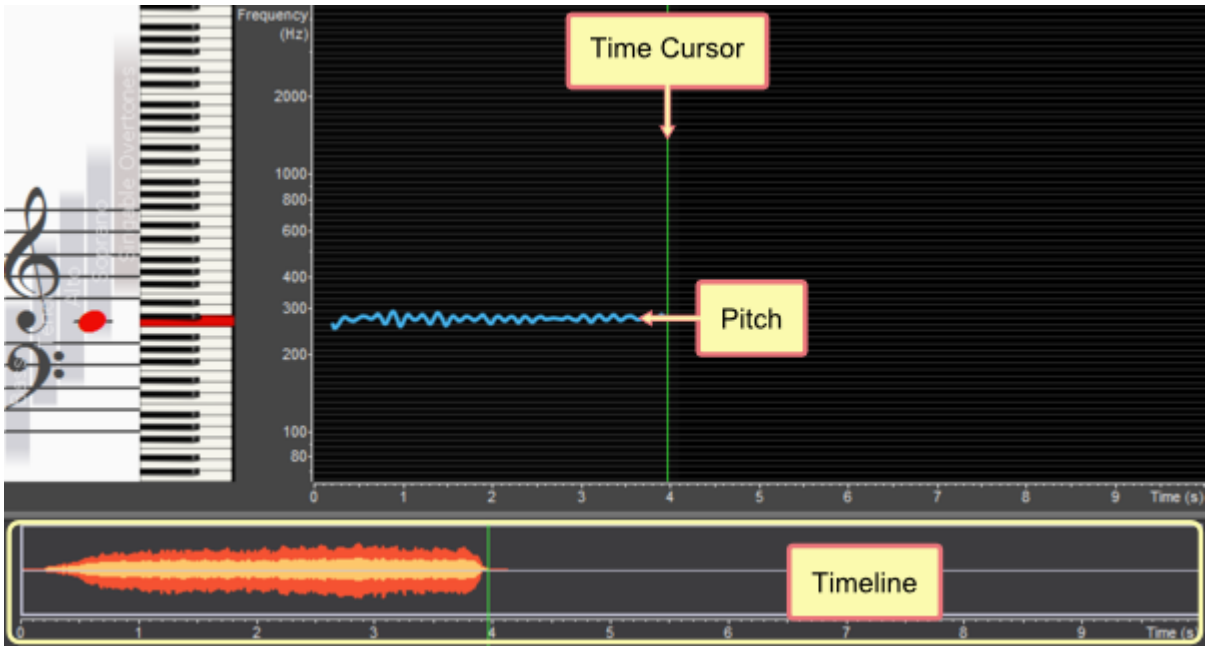

Figure 2.5. Singing a tone on the same pitch for one breath

The Timeline at the bottom shows the intensity (or volume) of your sound. The horizontal blue line in the middle is the pitch of your voice. The vertical green line is the Time Cursor. The Piano highlights the key that best matches the pitch of the recording at the time where the Time Cursor is, and the Staff View shows the same key as a musical note.

2.4. Loading sample file

To make this tutorial easier to follow, you can load the same audio file that was used to make the screenshots shown here, so that you will see exactly the same screens. Click on Help → Open sample file and select the file “vibrato_female.ogg”.

After you have opened the file, click on the profile list and select the profile **“Standard Time Zoom”** to ensure your view matches the following screenshots.

2.5. Zooming frequency range

Notice how the blue line is not straight, but is slightly moving up and down. This is due to the vibrato of the singing voice. Let's examine this a bit more closely. Move the mouse over the Frequency Scale (as shown in Figure 2.6) and rotate the scroll wheel of your mouse, or experiment with pressing the 1 and 2 buttons on your keyboard. This will zoom the frequency range in or out. Once you have zoomed in a bit, click on the Frequency Scale and drag it up or down with your mouse. This will scroll it. Zoom and scroll the frequency range until you see much more detail of the pitch line.

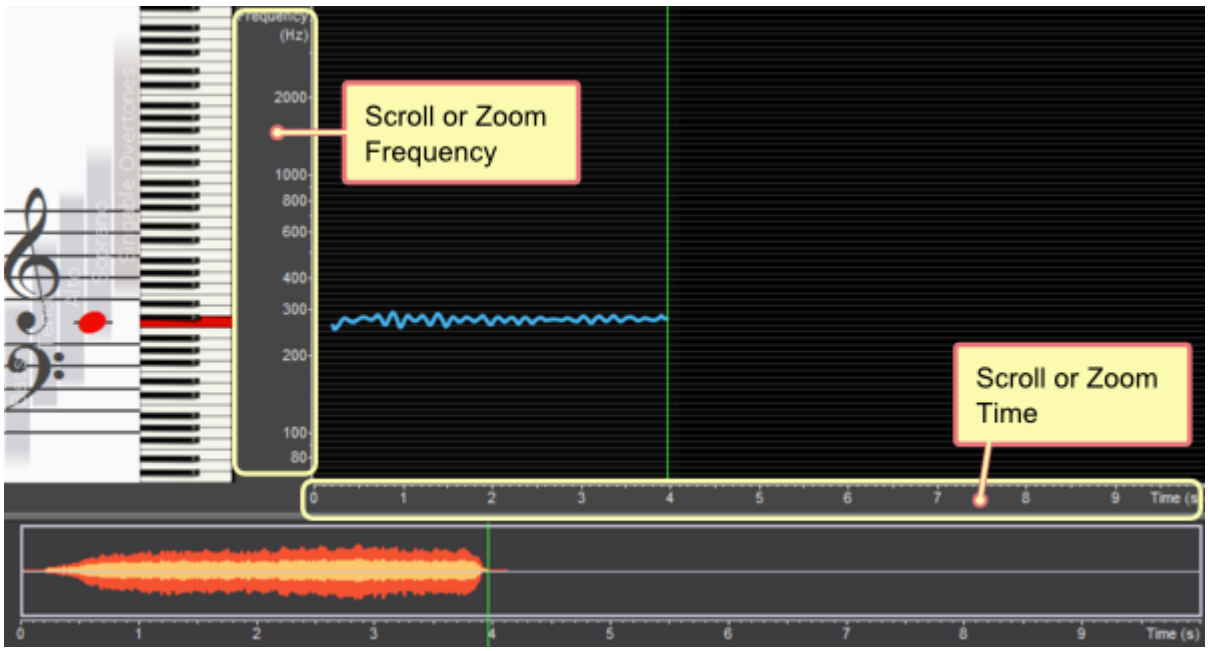

Figure 2.6. Using the mouse on the scales for scrolling and zooming

Now use your mouse on the Time Scale (also shown in Figure 2.6) and zoom into time a bit. So now you have zoomed the frequency range, and the time range, to show more detail on both scales. Your screen should now look similar to this:

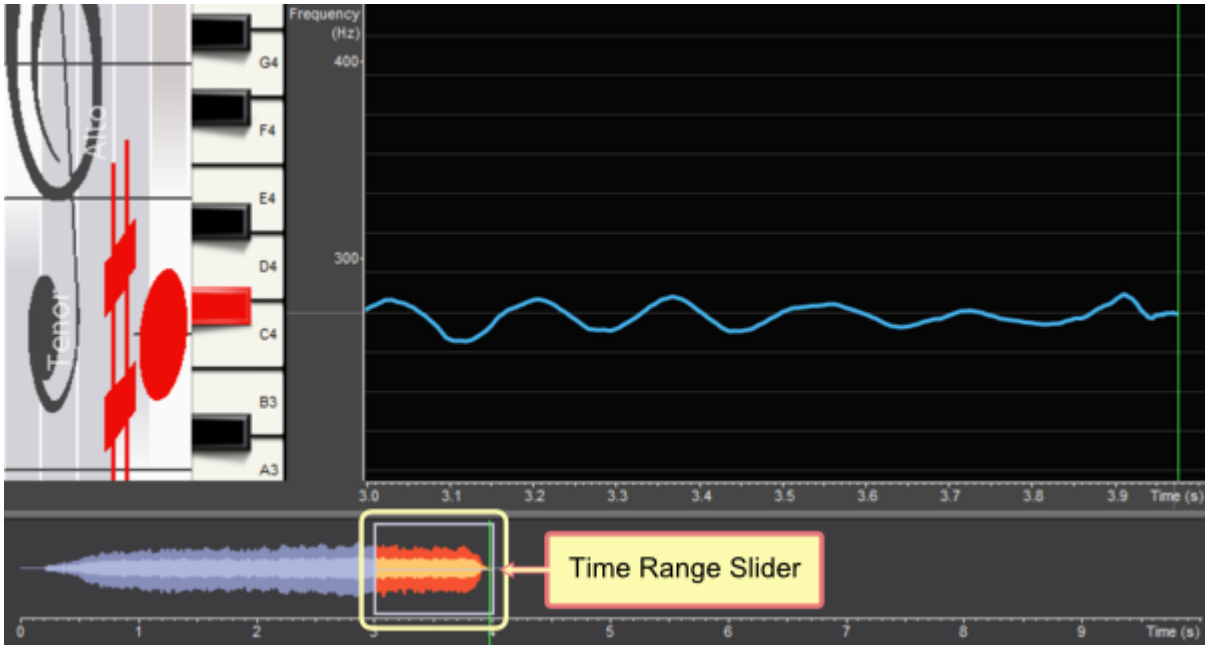

Figure 2.7. Zoomed-in detail of pitch

Notice how the Time Range Slider in Figure 2.7 is now much smaller than in the previous image, because the time range has been zoomed-in. You can drag the Time Range Slider with the mouse when you grab it in its middle, where the mouse cursor will change into a hand. Try this now. Or you can simply click anywhere on the Timeline to bring the Time Range Slider to that position.

2.6. Short-term view

We have already observed that the pitch line moves up and down. Now, with the detailed close-up, let's analyze the range of this pitch movement. Click on the Show Spectrum button (as shown on Figure 2.8) to bring up the Short-Term view:

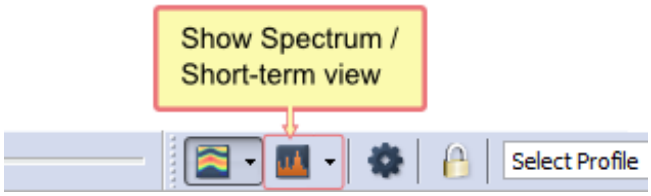

Figure 2.8. Click “Spectrum” button to bring up secondary view

Now the Analyzer View is split into two parts. On the left side is the long-term view, and on the right is the short-term view:

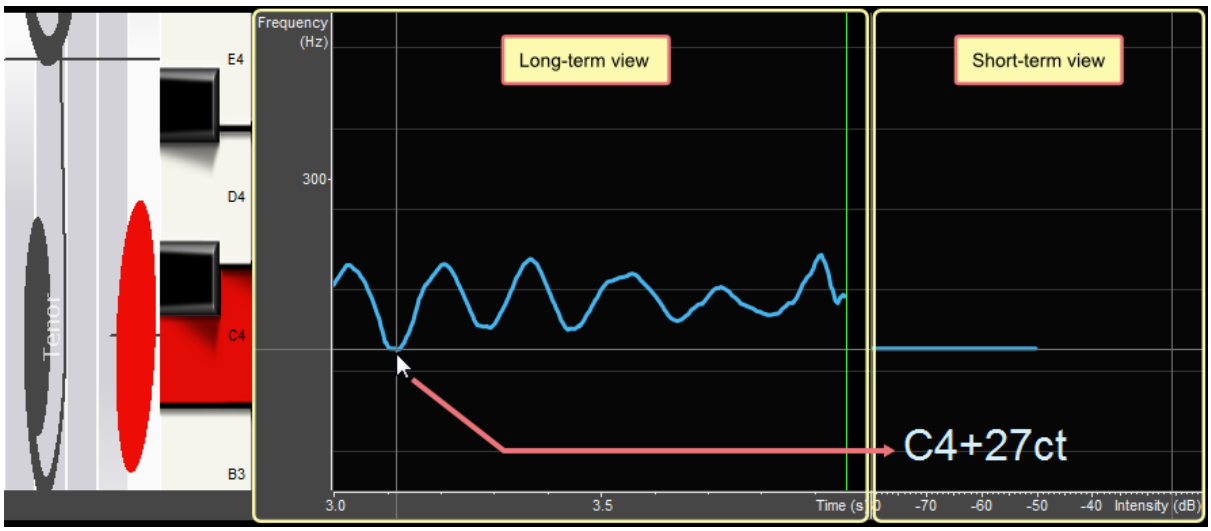

Figure 2.9. Long-term and short-term view

The two views are linked. If you move the mouse cursor over the long-term view on the left side, the short-term view on the right side will show a detailed close-up for the time at which you point. Right now we are displaying only pitch, so the short-term view shows a blue line with the exact pitch of the cursor location, and it displays the note name of this pitch in large letters.

Trace the pitch line on the left with the mouse cursor, and notice how the pitch display on the right changes. If the mouse cursor leaves the Long-term view, the short-term view will display details for the time location of the green Time Cursor line.

2.7. Play selection

In this recording, the singer is singing the note C#. Press the black C# key on the piano (highlighted in red in Figure 2.10), to hear the note from the piano. Now press the Play button on the toolbar, or press Space on your keyboard, to play back the recording. Listen to it and hear for yourself if the singer is singing in tune with the piano key.

Now let's select a part of the recording. Stop playback by pressing Space or by clicking on the Stop button on the toolbar. Now click on View → Zoom out fully to show the entire recording. Now click into the Analyzer View at the time position of 1.0 seconds and drag the mouse towards the time position of 3.0 seconds. Clicking and dragging in the Analyzer View will make a new selection which is marked by a white frame and a different background color. You can click and drag the borders of the selection to change its start and end position. Do this until you have selected the time range 1.0 to 3.0 seconds. This should look like this:

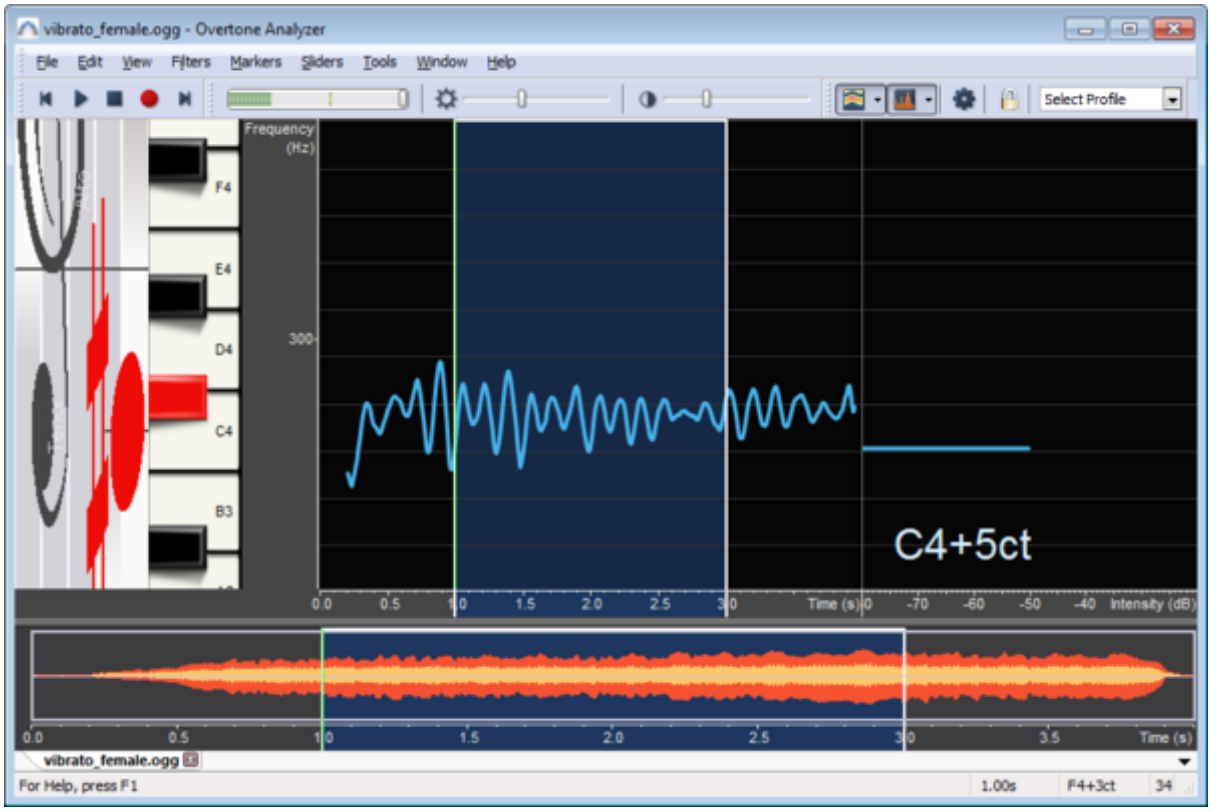

Figure 2.10. Select time range 1.0 to 3.0 seconds

Now start playback again. Notice how the selected time range is now played in a loop. You can switch looping on or off by clicking on Tools → Loop playback. If looping is turned off, pressing Space will play the selection just once, until you press it again.

If you click anywhere on the Timeline or the Analyzer View, the selection will disappear. If you want to keep a selection and be able to click into it during playback to hear specific parts, you can lock the selection by clicking on Edit → Lock selection or by clicking on the Lock selection button on the toolbar. Try this now. Make a selection, lock it, start playback, and then click at various points in time. Don't forget to unlock the selection when you are done with it.

2.8. Spectrogram

So far we've only looked at the pitch line, which shows the fundamental pitch of the recording. Now let's look at the Spectrogram, which shows how loud the recording is at each point of the frequency range. Click on Tools → Options or press the Options button on the toolbar to bring up the Options dialog, then click on “Analyzer View” on the left side to select the Analyzer View settings page:

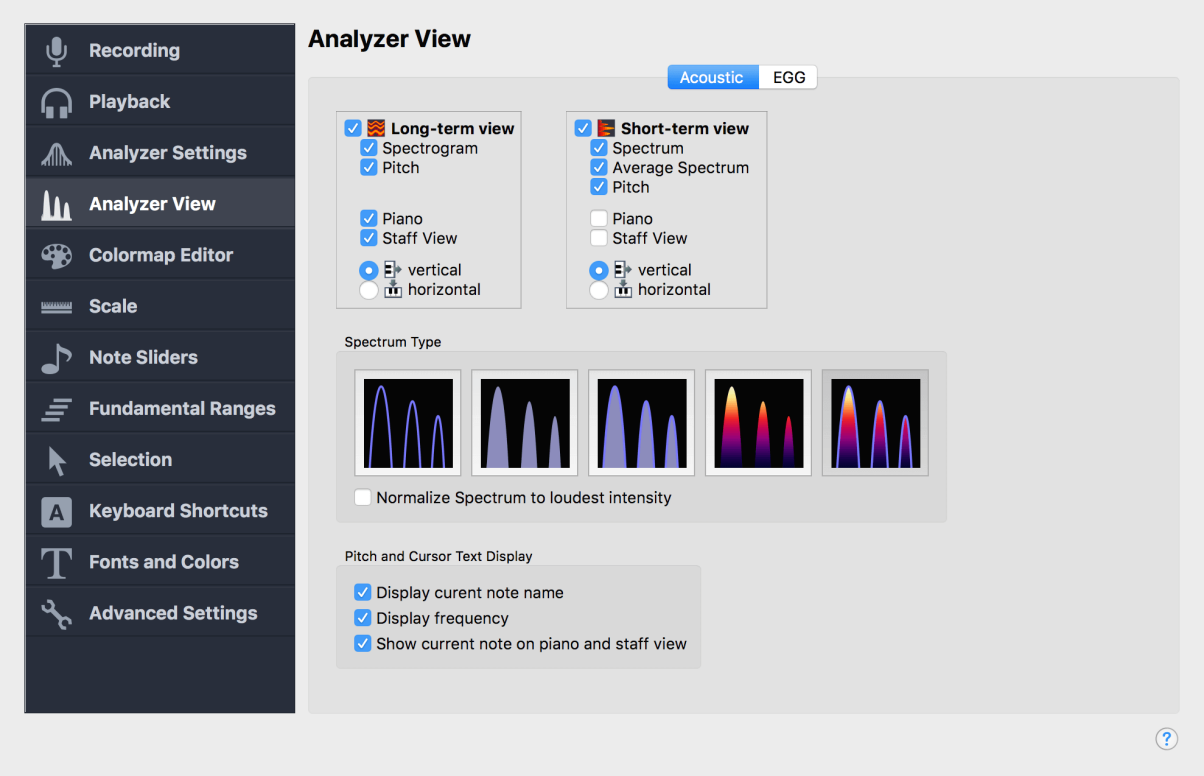

Figure 2.11. Analyzer View Settings

Check the Spectrogram and Spectrum options on the settings page so that it looks exactly like Figure 2.11. Then press on the OK button to close the settings page.

Now you should already see the Spectrogram and the Spectrum, but the screen might still be zoomed in very far from the previous step. On the toolbar, click on the Select Profile button and select the **"Standard Frequency Range"** profile. Then click on View → Zoom out fully. Now your screen should look similar to this:

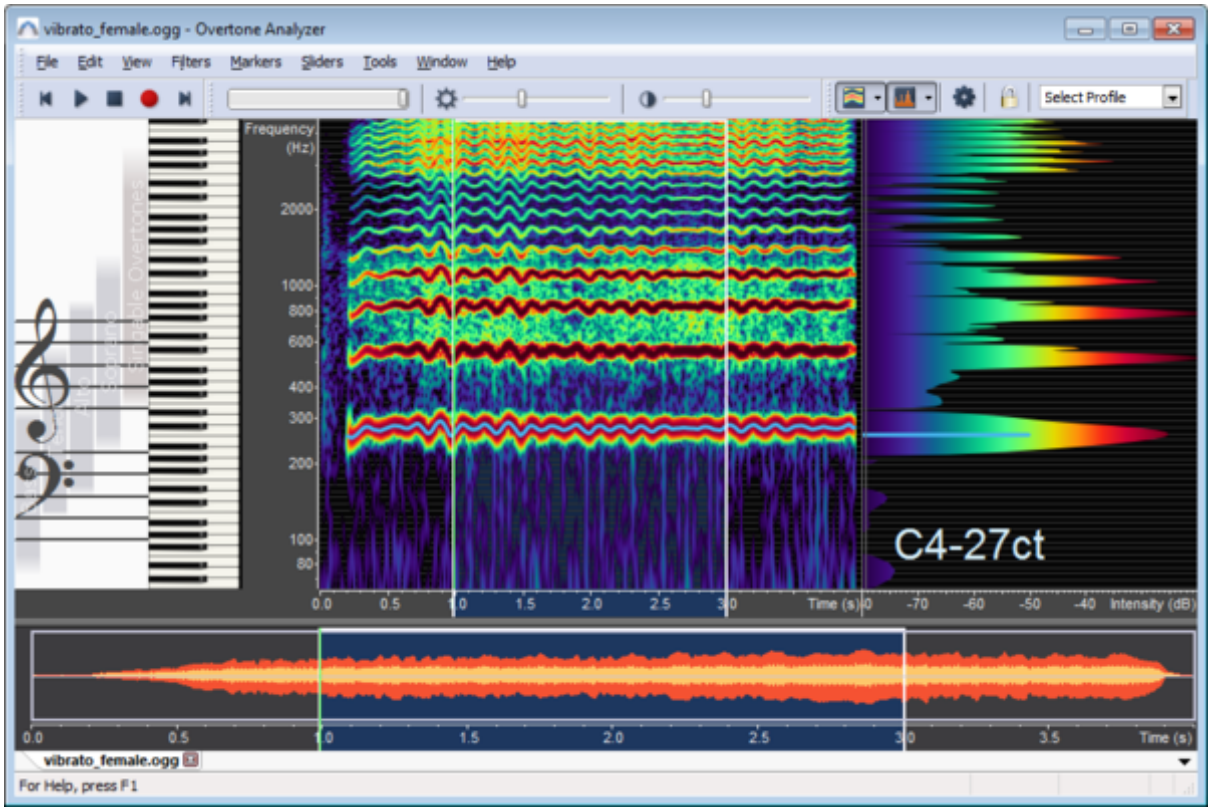

Figure 2.12. Displaying Spectrogram and Spectrum

Figure 2.12 once again shows the long-term and short-term views, but now they contain not only the pitch line, but also the Spectrogram (in the long-term view on the left side) and the Spectrum (in the short-term view on the right). You can adjust the amount of detail of Spectrogram and Spectrum by changing the brightness and contrast sliders on the toolbar (as shown on Figure 2.3).

The Spectrum and Spectrogram show the intensity (or volume) of the individual frequency components of a sound. To better understand what this means, let's use another feature of VoceVista Video: the Frequency Filter.

2.9. Frequency Filter

> **Note**
>
> Frequency Filters are only available in *VoceVista Video* and *VoceVista Video Pro*.

If you still have the selection from the previous steps, great. Otherwise, once again select the time range from 1 to 3 seconds, and lock the selection by clicking on Edit → Lock selection or by pressing the Lock selection toolbar button. Position the Time Cursor line in the middle by clicking into the long-term view at 2.0 seconds.

Now click on Filters → Add new frequency filter. Your screen should now look like this:

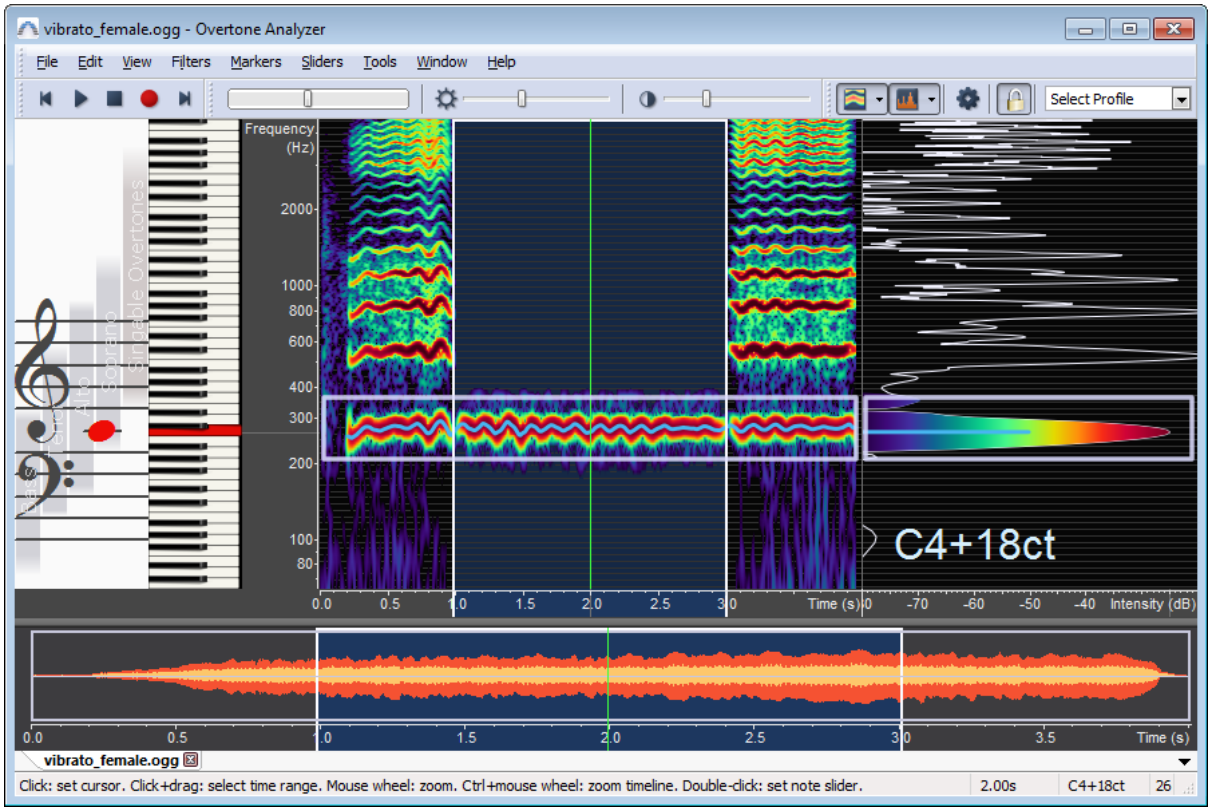

Figure 2.13. Adding a Frequency Filter

The grey frame between 200 and 400 Hertz is the Frequency Filter. It will remove all frequencies outside of its range. Enable looping. Start playback again and hear the effect of the filter. Then, while playback is running, move the filter up and down by dragging it with the mouse. The Spectrum on the right will show a white outline of the filtered parts, while the unfiltered (and therefore audible) part is shown in color. You can change the width of the filter by dragging the red handles that appear when the mouse cursor hovers over the filter.

Now click on Filters → Mute filtered frequencies to invert the filter. Instead of keeping only the frequencies inside of it, it now punches a hole in the spectrum by filtering out the frequencies inside of it. To see this more clearly, set the brightness slider on the toolbar to -10 dB, and set the contrast slider to 40 dB. Drag the Frequency Filter to cover the range from 400 to 700 Hz, so that it removes the second harmonic in the recording:

Figure 2.14. Subtractive Frequency Filter

Notice in Figure 2.14 how the Spectrum shows an outline of the second harmonic, while the other harmonics are in color. Once again, start playback, and move the filter up and down to hear its effect.

To remove the filter, click on the box on its top right corner.

2.10. Overtone Slider

The Frequency Filter enables you to listen to the individual harmonics (or overtones) in a recording, and the Overtone Slider is a tool that allows you to easily count the overtones and to see which note or frequency they correspond to. Make sure that you still have a selection. Then click on Sliders → Insert note slider at selection. A slider will appear at the fundamental pitch of the selected range.

Use the mouse to point at the slider, and drag up the red upwards-pointing triangle to expand the number of overtones shown:

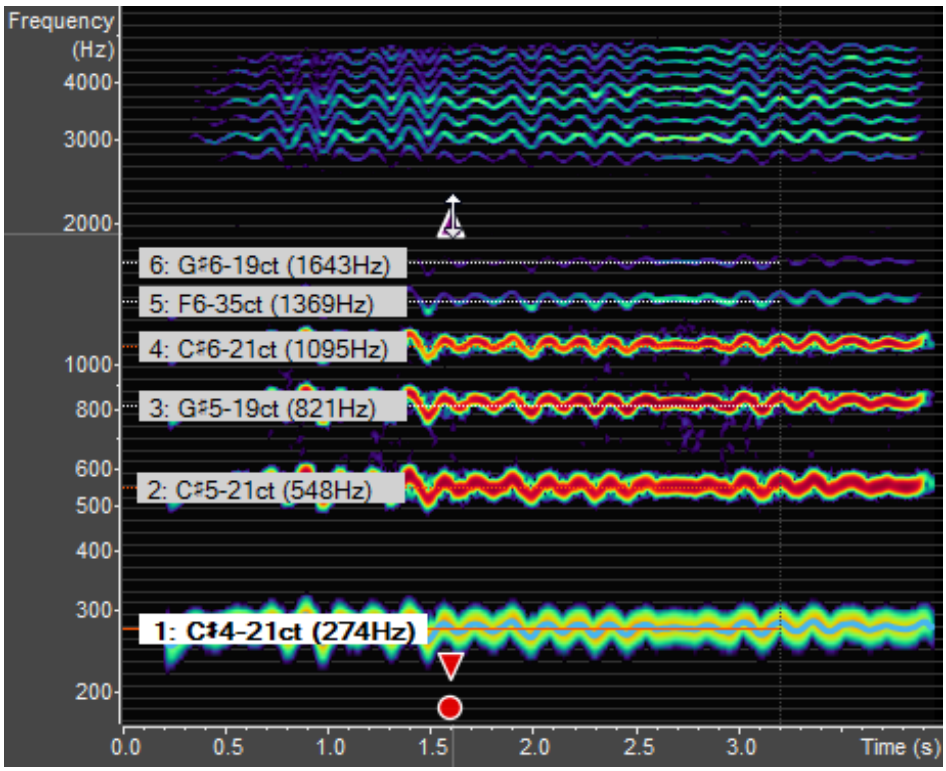

Figure 2.15. Adjusting the number of harmonics shown by the Overtone Slider

You can scroll down the Frequency Scale and zoom it slightly until you see enough detail of the Overtone Slider. The slider gives you a theoretical view of the overtones that belong to a particular sound. The labels of the slider lines show the number of the harmonic, its music note, and its frequency. The underlying Spectrogram shows the harmonics that are actually present in the recording. Click on the labels of the Overtone Slider and move the mouse up and down while holding the mouse button to hear the overtone scale belonging to the current fundamental pitch. Compare this to the actual overtones in the sound that you could hear with the Frequency Filter.

2.11. Summary

This concludes this quickstart guide. You have learned the basics of how to record and visualize sounds with VoceVista Video and had an introduction to the basic features and interface elements of the program. The main steps to remember are:

- Select the microphone to be used by right-clicking on the Input Volume control on the toolbar.
- Ctrl-Space starts the recording. Space will start playback or stop the current recording or playback.
- Always monitor the input volume during recording to prevent clipping.
- Use the mouse on the Frequency Scale and Time Scale for scrolling and zooming. Click & drag to scroll; use the mouse wheel or the 1 and 2 keys for zooming.
- The Long-term view shows a time range with the Spectrogram and/or fundamental pitch. The Short-term view shows a detailed close-up of the time where the cursor is in the Long-term view.
- Make a time selection by clicking and dragging on the Spectrogram. The selection will be played in a loop if looping is enabled. You can lock the selection to be able to click into it to position the playback cursor.
- The Spectrum shows the intensity (or volume) of the individual frequency components of a sound. The Spectrogram shows how the Spectrum changes over time.
- Adjust the brightness and contrast sliders on the toolbar to show the desired amount of detail in the Spectrogram and Spectrum.

- The Frequency Filter can filter out individual harmonics or frequency ranges and allows to listen to them individually, or to remove them.
- Use the Overtone Slider to count and hear the theoretical harmonics that belong to a particular tone.

Next Steps

For a more detailed description of the interface elements introduced here, have a look at the chapter called User Interface in the Reference Guide.

Chapter 3. Reference Guide

This section explains the user interface of VoceVista Video and the different program settings.

3.1. User Interface

3.1.1. Overview

Here is an overview of the main elements of the VoceVista Video user interface:

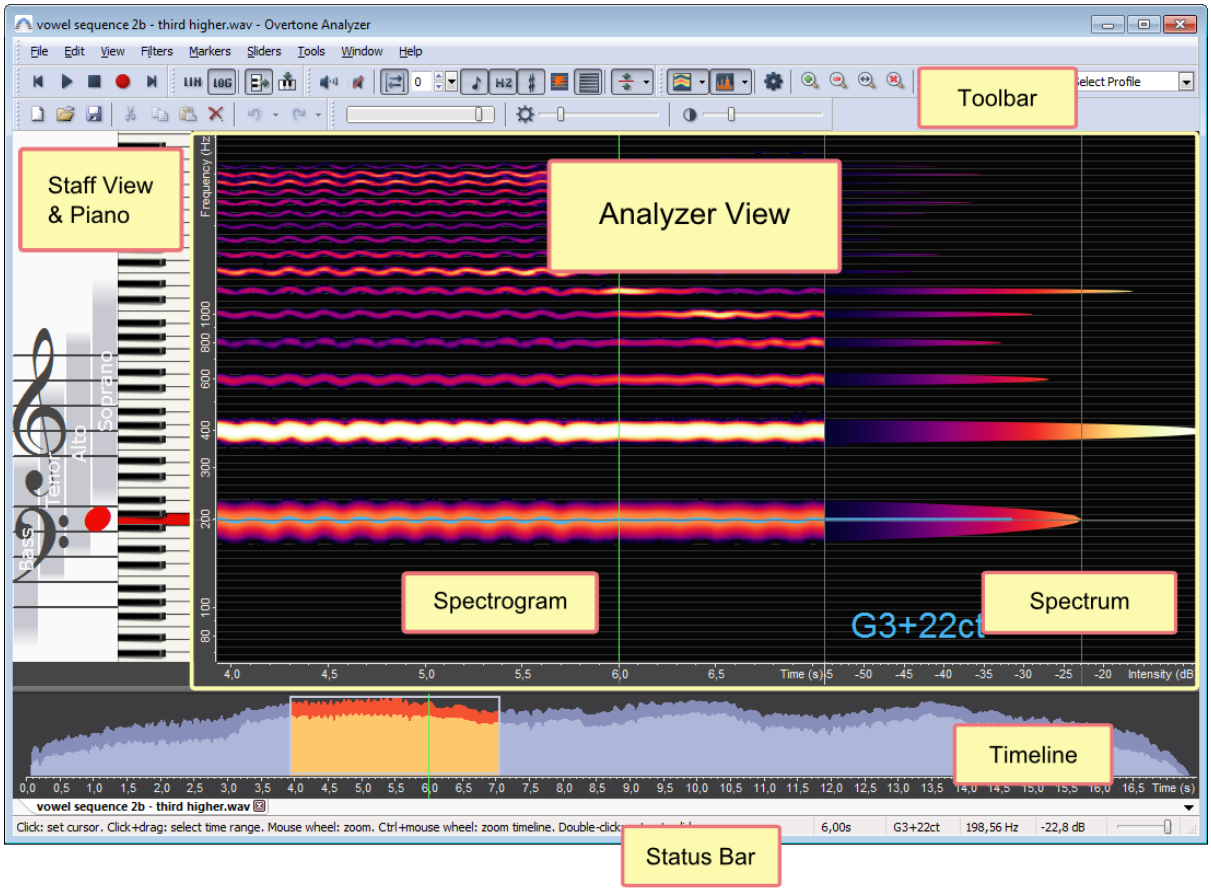

Figure 3.1. User Interface Overview

Toolbar

The Toolbar provides quick access to many functions of the program.

Staff View & Piano

The Piano is a playable keyboard. You can click on individual tones, or use Ctrl-Click to select entire chords and play them with the Play selected tones command.

The Staff View shows a musical staff. It also shows the current notes played on the piano. The colored areas in the background indicate the ambitus (the range from lowest to highest singable note) of the singing voices and of the singable overtones.

Analyzer View

This is the main view of VoceVista Video and contains output from the analyzer (such as spectrum, spectrogram, and pitch), as well as note/overtone sliders, rulers, markers, and frequency filters.

Timeline

The Timeline shows an overview of the entire recording or loaded audio file as a waveform. The Timeline also acts like a scrollbar and lets you navigate through a long recording by moving the current range slider with the mouse.

Status Bar

The status bar shows context help and status information. If the mouse cursor points at something that can be clicked or manipulated with the mouse, the status bar will show details about the specific operations that you can perform.

If the mouse cursor is over the Analyzer View, the Status Bar will show the time, frequency, note name and intensity of the recording at the current mouse cursor position. The Status Bar also has a slider for the main output volume.

3.1.2. Scrolling and Zooming

VoceVista Video is designed to make it easy to change parameters, and to see the effects of those changes immediately. The most important parameters that are changed frequently are the displayed **time range**, and the displayed **frequency range**. The other frequently changed values that affect the display are the **dynamic range** of the Spectrogram and Spectrum, and the dynamic range of the Waveform and Timeline.

There are various ways to change those values, but one of the fastest is to use the mouse to **drag** things with **the left button**, or to **zoom** the scales with the **mouse wheel**.

Mouse wheel emulation

If you don't have a mouse wheel, you can point at the scale or window to be zoomed and use the two-finger gesture for scrolling and zooming on your touchpad. Generally this consists of moving two fingers up or down on the touchpad. Another way to zoom and scroll is to use the 1 and 2 keys on your keyboard.

Interactive areas

The following table lists the mouse commands for scrolling and zooming, and where they can be used:

Command	**Effect**	**Can be used on**
Left click and drag	Scroll	any scale (Time Scale, Frequency Scale, Intensity Scale, Amplitude Scale)
Use mouse wheel	Zoom	any scale, Spectrogram
Use two-finger scroll gesture on touchpad		
Press the 1 and 2 keys		
Shift + mouse wheel	Scroll time range	Spectrogram
Ctrl + mouse wheel	Zoom the Timeline View (instead of zooming the time range of the Spectrogram)	Spectrogram, Timeline

Table 3.1. Commands for scrolling and zooming

The following image shows all areas that can be manipulated by left click or mouse wheel:

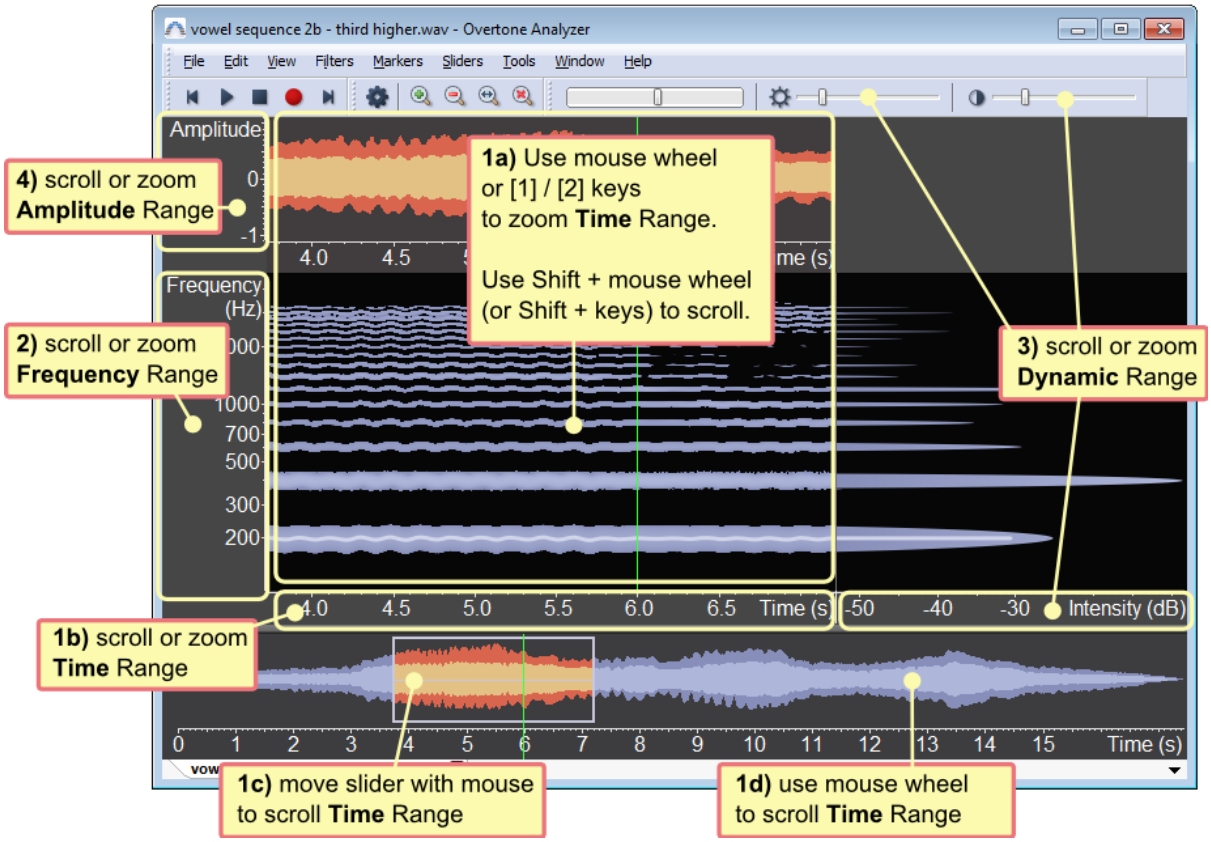

Figure 3.2. Scrollable and zoomable areas

1) Time Range

Because the time range of the Spectrogram is the most frequently changed range, there are many different ways to manipulate it. Try out the various ways of scrolling and zooming the time range with the mouse shown in Figure 3.2. Moving the Time Range Slider (**1c**) or clicking on the Timeline (**1d**) is for navigating quickly to any part of the recording, while scrolling the time scale (**1b**) is much more precise for small adjustments.

The time range can also be manipulated by the forward and backward buttons and the zoom buttons on the toolbar, or by the Zoom... entries in the View menu. You can also use Markers to bookmark specific parts of a recording and return to them quickly.

2) Frequency Range

You can change the frequency range by scrolling and zooming the frequency scale (**2**). You can also set precise values on the Scale settings page, or on the Scale Range toolbar (this toolbar is not shown by default and needs to be enabled).

3) Dynamic Range

The displayed dynamic range determines the range of intensities shown on spectrogram and spectrum. It can also be set on the Colormap editor. It can be changed by scrolling and zooming the spectrum intensity scale, or by dragging the brightness and contrast sliders on the toolbar (**3**).

4) Waveform / Timeline Amplitude Range

This setting determines the dynamic range shown on the Waveform and Timeline. You need to enable the Waveform in the View menu to see the amplitude scale.

5) Timeline zoom

The Timeline itself can also be zoomed by holding Ctrl while using the mouse wheel (or the mouse wheel emulation keys) on the Spectrogram or the Timeline. You can also enable the Timeline Zoom Slider in the View menu and use that to change the Timeline zoom.

To reset the zoom level of the Timeline, click on View → Zoom out fully.

3.1.3. Selection

In VoceVista Video, the concept of a selection extends to a time range of the recording, and to elements that can be selected individually, such as piano keys and note sliders:

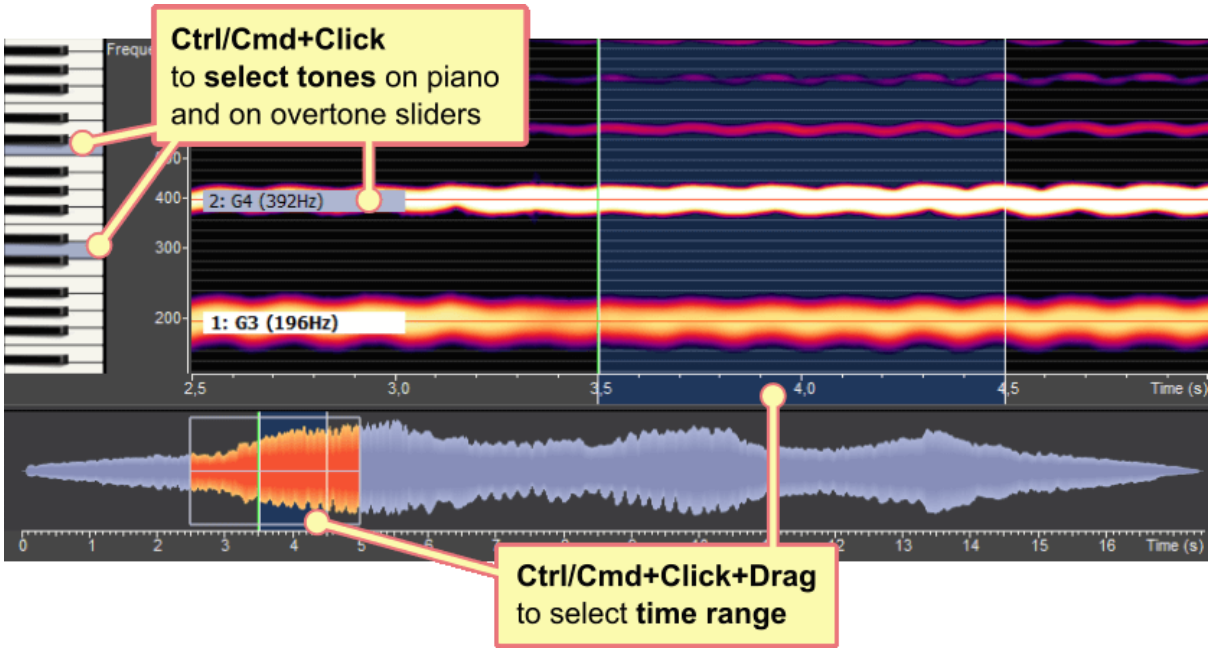

Figure 3.3. Selecting time range and notes

Using keyboard modifiers (Control / Command)

To change the selection with the mouse, you need to hold the Control key (on Windows), or the Command key (on Mac).

On the Selection Settings page, you can set the Mouse Click Action to change the selection with the simple click instead of requiring the Ctrl/Cmd key. This is the behavior used by Overtone Analyzer 4.

Selecting time range of a recording

Parts of the recording can be selected on the Timeline, the Waveform View, and on the Spectrogram. Simply click anywhere on these views, hold the left mouse button, and move the mouse to the end of the area that you want to select. The selected area will be marked with a different background color. You can click and drag the borders of an existing selection to make it longer or shorter.

Working with a time selection

Once you have selected a time range, you can press the Play button to play it repeatedly. You can choose if playback should be repeated automatically by clicking on Tools → Loop playback.

You can also click on View → Zoom to fit selection to set the visible time range to match the current selection.

Conversely, you can click on Edit → Select visible range to select the time range that is currently visible.

Use the command Edit → Set selection to enter precise numbers for the time range to be selected.

Further, you can apply edit commands such as Cut, Copy, Paste, etc. to the current selection.

Locking a selection

If you click anywhere on the Timeline or the Spectrogram, the playback cursor will jump to the position of the click, and the current selection will be reset. Sometimes you might want to keep the selection a bit longer, but still move the cursor within the selection to hear specific parts. To prevent the selection from being lost when you move the cursor, click on Edit → Lock selection. Now the selection can still be resized by moving its borders, but

clicking somewhere on the Timeline or Analyzer View won't destroy the existing selection. This is particularly useful when comparing files, and when working with Frequency Filters.

Resetting the selection

If the selection is not locked, you can simply click on the Timeline or Analyzer View to remove it. This will reposition the playback cursor and reset the selection. You can also clear the selection by clicking on Edit → Reset selection, or by pressing the ESC key.

Selecting Tones

You can select and deselect tones on the piano keyboard and on the Overtone Sliders by clicking on them while holding the Control key (on Windows), or the Command key (on Mac). To play those tones, click on Sliders → Play selected tones or press the Return key.

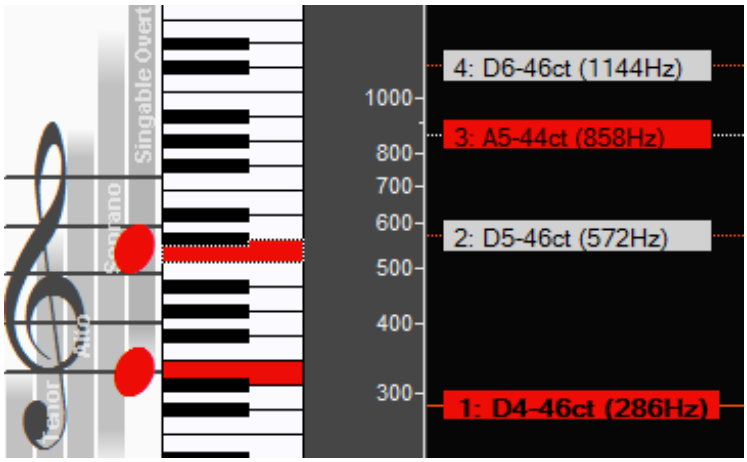

Figure 3.4. Playing selected tones

To make all tones unselected click on Sliders → Deselect all tones.

3.1.4. Timeline and Waveform

The Waveform and the Timeline are closely related and show the same information, which is how the intensity (or volume) of the recording changes over time. The difference is that the Timeline shows a larger time range than the Spectrogram, while the Waveform always shows the same time range as the Spectrogram. It can be useful to display both views when working with longer recordings, where the Timeline provides an overview over the whole recording, and the Waveform shows a detailed view of the current time range.

If the Waveform is hidden, you can show it by clicking on View → Waveform.

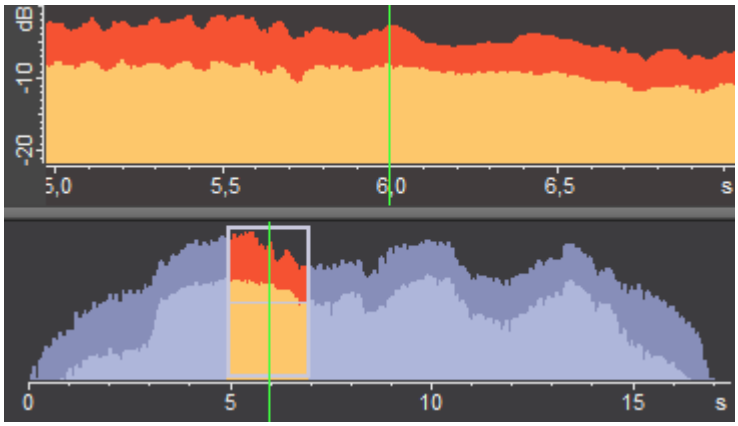

Figure 3.5. Waveform and Timeline

Figure 3.5 shows the Timeline at the bottom, and the Waveform at the top. The Timeline shows the whole recording, which is 17 seconds long. The **Time Range Slider** at the bottom highlights the range from 5.0 to 7.0 seconds: this is the range shown on the Waveform and on the Spectrogram (which is not shown in this image). The **Time Cursor** is at 6.0 seconds (this is the green vertical line).

Navigating with the Timeline

Since the Timeline shows an overview of the entire recording, you can use it to navigate to specific points in your recording. Simply click on it anywhere to jump to that position. You can also move the Time Range Slider with the mouse when you point at the middle of the slider, where the mouse cursor changes into a hand.

Waveform dynamic range

You can zoom the range of values shown on the Waveform and Timeline through the amplitude scale by pointing the mouse cursor at it and then using the mouse wheel, or by pressing the 1 and 2 keys. You can also **left-click** and **drag** the scale to scroll its range. This is explained in more detail in the chapter scrolling and zooming.

Timeline zoom

The Timeline itself can be zoomed by pointing at it and then holding Ctrl and using the mouse wheel. You can also show the Timeline Zoom Slider in the View menu and use that to zoom the Timeline.

Oscilloscope

When the Waveform is zoomed out to show a relatively long time range, each pixel on the screen shows the maximum amplitude for all the audio samples represented by the time range of that pixel. When you zoom in very close, the displayed time range becomes so small that individual audio samples become visible. In that case the Waveform looks like an oscilloscope:

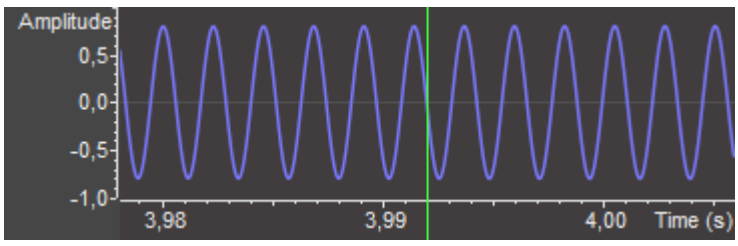

Figure 3.6. Fully zoomed-in Waveform as oscilloscope

Waveform display modes

The Waveform and Timeline can display the intensity of the recording in various ways. If you **right-click** on the Waveform or the Timeline, you can select one of these display modes:

Layout	Scale mode	Example
Symmetric	linear	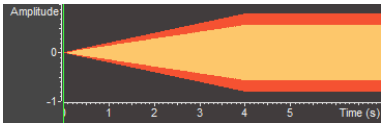
	log	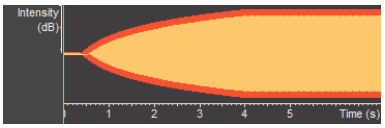
Skyline	linear	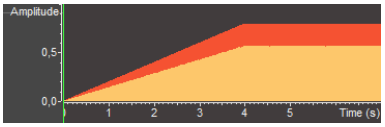
	log	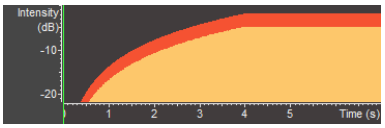

Table 3.2. The same recording in various waveform display modes

Table 3.2 illustrates the four display modes. The layout determines if the view is symmetrical or not. In the skyline mode, only the upper half of the scale range is shown.

The scale mode determines if the intensity is shown on a linear or a log scale. Depending on what you want to do, one mode may be more useful than the other. Notice how in the bottom row of Table 3.2 the beginning of the recording is not shown. That is because the log scale has no lower end (it would be at minus infinity), and therefore small values are hidden in this mode. On the other hand, the display of intensity in log mode is more similar to the way we hear volumes than in the linear mode.

Linear Amplitude	Log Intensity
Shows entire range from 0 to 1.	Shows intensity from loudest range (at 0 dB) towards some cutoff point, depending on scale zoom.
Very quiet passages are still visible.	Quiet passages may be hidden.
Does not accurately represent human perception and makes it hard to compare the volume of different sections.	Is a better model of human perception and allows for better volume comparisons.
Good for editing and for general navigation through a recording.	Good for analyzing the dynamic range of a recording and for inspecting variations in loudness.

Table 3.3. Waveform scale mode comparison

3.1.5. Analyzer View

This is the main view of VoceVista Video. The Analyzer View can show the spectrum, the spectrogram, the fundamental pitch, or any combination of these. You can select what to display on the toolbar or on the Analyzer View settings page.

This view can also display overtone sliders and other overlays. Surrounding the Analyzer View are the Frequency Scale, Time Scale, and Intensity Scale. They can be turned on and off in the View Menu.

Spectrum Display / Short Term View

The Spectrum Display shows the intensity of the frequency components of a sound at a specific point in time. It can also show the fundamental pitch of the sound at this time:

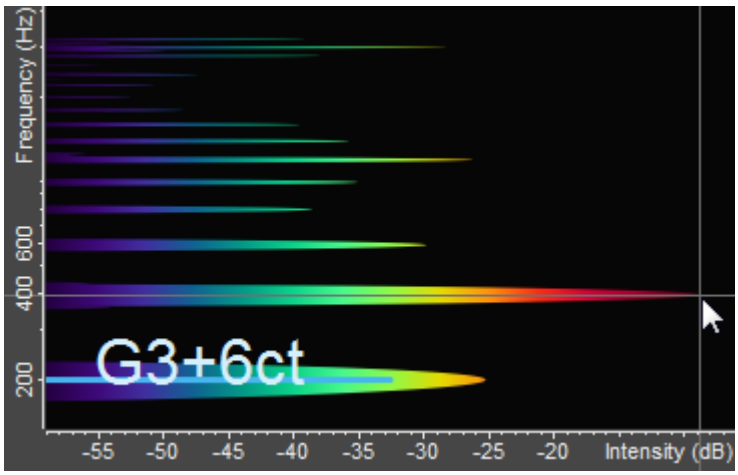

Figure 3.7. Frequency Spectrum and Fundamental Pitch

The frequency axis is shown in Hertz (cycles per second, abbreviated Hz), and the intensity axis is shown in decibel (dB). The loudest possible tone has 0 dB, and all tones quieter than that have a negative decibel value.

In Figure 3.7, the first spectral peak is at about 200 Hz and has an intensity of -25 dB. The second peak is louder. It has a frequency of about 400 Hz and an intensity of -9 dB. The cursor cross hair is snapped to the second peak at 400 Hz. The pitch display (the straight blue line) shows a fundamental pitch of about 200 Hz, which corresponds to the note “G3+6ct”.

The range of the intensity scale can be adjusted with the dynamic range slider on the toolbar. The range of the frequency scale can be adjusted on the Scale settings page, or on the Scale Range toolbar.

All scales can also be zoomed and scrolled with the mouse. To scroll the scale range, left click and drag a scale. To zoom it, point at the scale and use the mouse wheel. If you don't have a mouse wheel, you can also point the cursor at the scale and use the the 1 and 2 keys to zoom.

You can select different color schemes and display types for the spectrum display in the program options, in particular on the Analyzer View settings page and the Colormap Editor.

Spectrogram Display / Long Term View

The Spectrogram Display is very similar to the spectrum, but instead of a single spectrum, it shows a series of spectra over time:

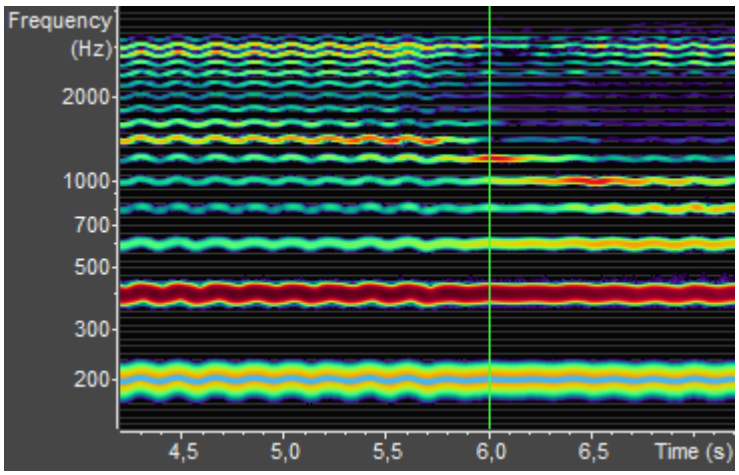

Figure 3.8. Spectrogram Display

The bottom scale now shows time instead of intensity, and the intensity is now represented by the color of each point on the display. Think of it as looking at a series of stacked spectra from the top.

In Figure 3.8, you see the spectrogram of a male singer. The fundamental pitch moves around 200 Hz. The periodic rise and fall of the pitch is the vibrato of the singer. The second harmonic at 400 Hz is much louder than the first. The higher harmonics between 600 Hz and 2000 Hz change over the time range shown. This is because the singer is changing the vowel he is singing. The singer's formant at 3000 Hz stays strong throughout.

Combined view: Spectrogram + Spectrum

It is possible to show spectrogram and spectrum side by side:

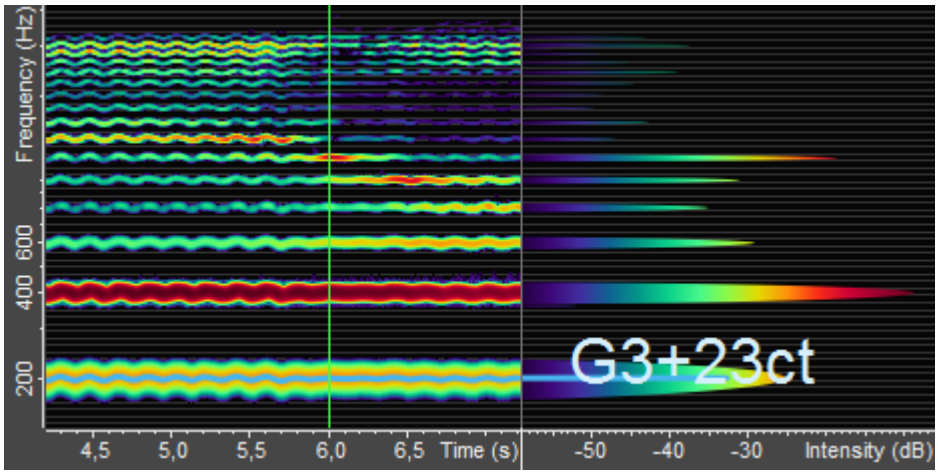

Figure 3.9. Spectrogram and Spectrum side-by-side

In this case the spectrogram on the left side shows how the frequency components change over time, and the spectrum on the right shows more details for a specific point in time. If you move the mouse cursor over the spectrogram, the right side will show the spectrum for the time pointed at by the mouse. If the mouse cursor is not over the spectrogram, the spectrum will show the frequency components for the time where the green cursor line is.

For example, in the image above, the cursor is at 6.0 seconds, so the spectrum shows a cross-section of the spectrogram at this time.

Pitch only display

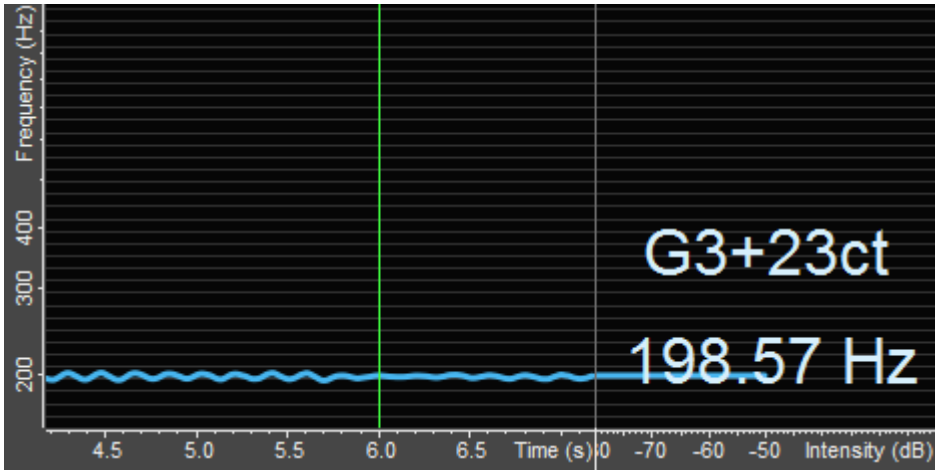

Figure 3.10. Pitch only display

It is possible to turn off spectrogram and spectrum and show only the fundamental pitch, as in Figure 3.10.

3.1.6. Sound Generator

The Sound Generator is a tool to explore fundamental aspects of acoustics and of voice science by means of reconstructing sounds from scratch. It also allows you to generate synthetic sounds for calibrating measurements.

The workflow is to adjust some settings, and then to click on the Generate button to add the new sound to the current recording. To get started, click on the Saved Profiles list on the right, pick one of the predefined profiles, and click on Generate. Try out the different predefined profiles to see what kind of sounds they generate.

Fundamental and Overtones

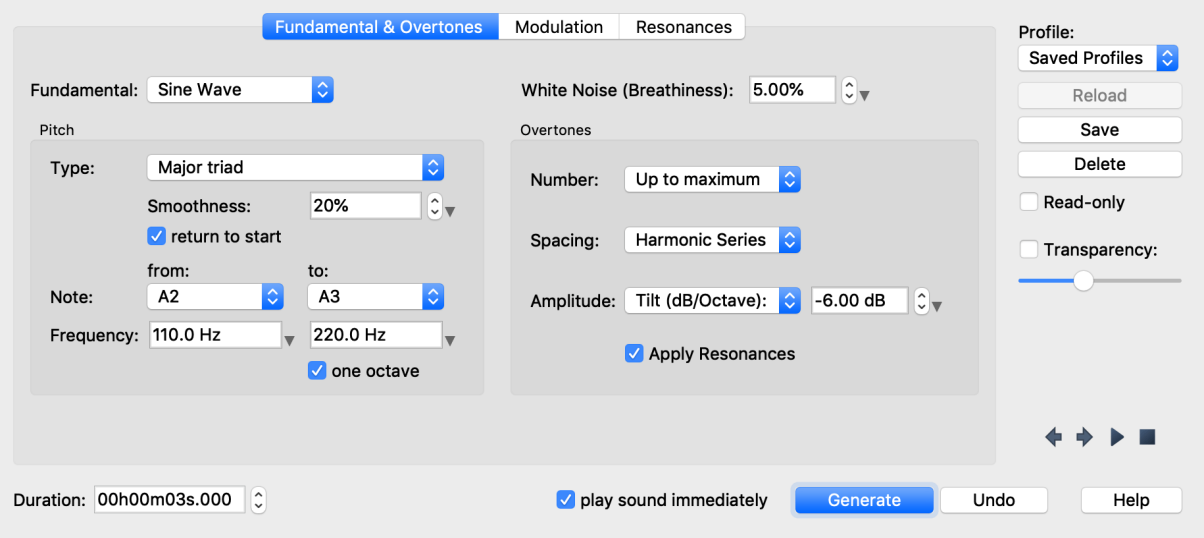

Figure 3.11. Sound Generator - Fundamental and Overtones

Fundamental

The fundamental is the basic waveform from which the generated sound is composed. In can be one of the following:

Sine Wave
: Your sound will be composed of one or more sine waves added together.

White Noise
: Your sound will be pure white noise. This fundamental sound type has no additional parameters.

Coherent Noise
: Coherent Noise is somewhere between sine waves and white noise. It is a waveform that is more irregular than a sine wave, but still has a frequency.

White Noise (Breathiness)

This setting allows to mix a percentage of pure white noise to the other fundamental waveform types, which can be used to simulate the breathiness of a voice, or the background noise of a recording.

Pitch

This section determines if the pitch stays constant over the duration of the sound, or if it moves, for example by forming a triad or a scale.

When the pitch type is not constant, the pitch will move from the from to the to note or frequency. You can check return to start to make the pitch move back to the starting note.

Type

Pitch type can be one of the following:

Constant Pitch
: The pitch will be constant over the entire duration of the sound.

Single Step
: The pitch will smoothly move from the start to the finish note.

Major / minor triad
: The pitch will form a major or minor triad.

Major / minor scale
: The pitch will form a major or minor scale.

Sweep (linear)
: The pitch form a straight interpolation along the linear scale.

Sweep (log)
: The pitch form a straight interpolation along the logarithmic scale.

Smoothness

This setting defines if the pitch changes gradually from one tone to the next (high smoothness), or with a more sharply defined step (low smoothness).

Overtones

This section determines the number, spacing, and intensity of the overtones of the sound, which will shape its timbre.

Number

The number of overtones can be anywhere from one (which means the sound only has the fundamental) up to the maximum that is determined by the sampling rate of the sound. In digital audio, a waveform can represent frequencies of up to half the sampling rate. For example, with a sampling rate of 44100 Hz, the highest frequency that can be represented is 22050 Hz. Therefore, overtones of up to a maximum of 22050 Hz can be included in the sound.

Specify amount
: Enter the amount of overtones. Overtones above the maximum frequency will be ignored.

Up to maximum
: Include all overtones below the maximum frequency, which is half the sampling rate.

Up to frequency
: Include all overtones below the entered frequency (or half the sampling rate, whichever is lower).

Spacing

The spacing of the overtones determines the distance or the ratio between successive overtones. This allows to create sounds that are based on the natural harmonic series, as well as other sounds that have different frequency ratios between the overtones.

Harmonic Series
: Each overtone will be a whole multiple of the fundamental. For example, if the fundamental is 100 HZ, the overtones will be 200 Hz, 300 Hz, 400 Hz, 500 Hz, etc. In this case the overtones will be the harmonics of the fundamental.

Octaves
: Each overtone will have twice the frequency of its predecessor. For a fundamental of 100 Hz, the overtones will be 200 Hz, 400 Hz, 800 Hz, 1600 Hz. In this case the overtones will all be octaves of the fundamental.

Multiply by
: Each overtone will have the frequency of its predecessor multiplied by the entered factor. If the factor is 2.0, this is the same as selecting Octaves.
In mathematical terms: `frequency(n) = factor*frequency(n-1)`

Fundamental factor
: Each overtone will have the frequency of the fundamental, multiplied by a factor and by the number of the overtone. If the factor is 1.0, this will generate the harmonic series. In other words, the factor sets the distance between consecutive overtones.

In mathematical terms: `frequency(n) = n*factor*fundamental`

Amplitude

This section determines the amplitude, or intensity, of successive overtones. The values are measured in decibel (dB). A value of -6 dB corresponds to a reduction of the intensity by half. If the dropoff by octave or overtone is zero, all overtones will have the same intensity. If it is negative, higher overtones will be quieter than the lower ones (this is usually the case in natural sounds).

Tilt (dB/Octave)
: The tilt sets the intensity drop per octave. With a value of -6 dB, each octave will be half the intensity of the previous octave.

db/Overtone
: This sets the intensity drop per overtone. Since the overtones follow a linear frequency scale, this setting will create a linear intensity dropoff (unlike the tilt, which will create a logarithmic dropoff).

Half intensity / Octave
: With this setting, the intensity will drop exactly in half with each octave, which is equivalent to having a tilt of -6,0206 dB per octave.

  Each overtone has an intensity that is one over N, where N is the frequency of the overtone divided by the frequency of the fundamental.

  In mathematical terms: `intensity(n) = 1/(frequency(n)/fundamental)`

Apply Resonances

When this is checked, the generated sound will be filtered through the resonances currently defined by the vowel chart and the Resonances part of the Sound Generator. This is intented to simulate sounds made by the human voice.

Modulation

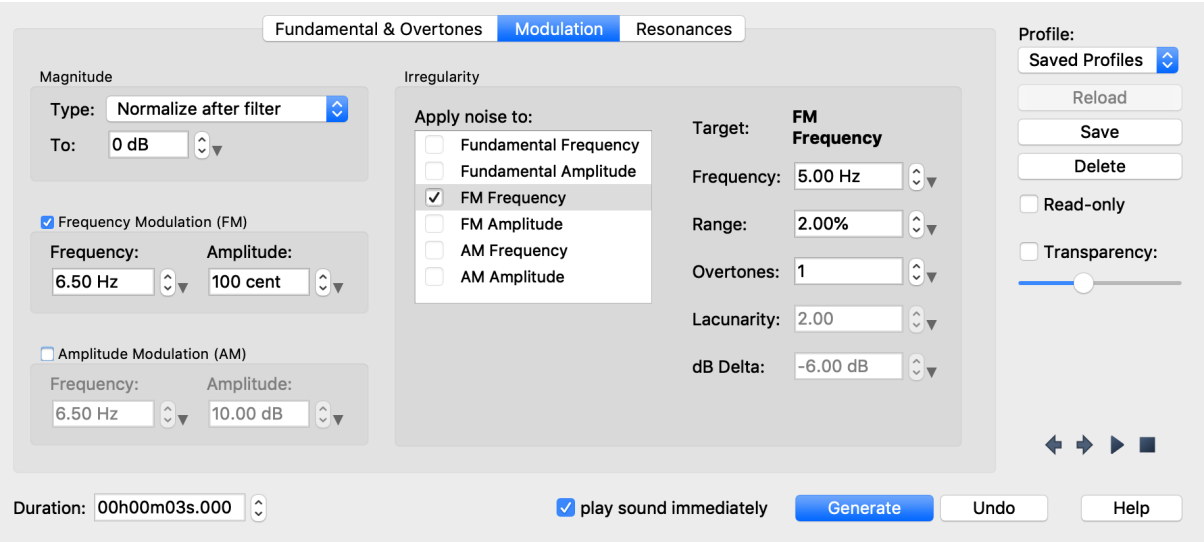

Figure 3.12. Sound Generator - Modulation

The modulation page offers several ways in which the generated sound can be changed. In particular, it has ways to affect (modulate) the intensity of the sound, and its frequency.

The words intensity, magnitude, amplitude, and loudness, all refer to the same thing, which is how loud the generated sound is. Each term has a slightly different meaning. When referring to individual points in time of a waveform, amplitude is the maximum deviation of the waveform away from its center, so it can be positive or negative. Magnitude is the absolute value of the amplitude. When referring to a longer waveform rather than individual points in time, amplitude and magnitude are often used synonymously, and in that case they both mean the absolute value of the maximum deviation from the center.

Intensity is a synonym for magnitude.

Amplitude, Magnitude, and Intensity can be measured in linear values, or in decibel. A value of 0 dB is the maximum intensity that can be represented in a digital waveform. It corresponds to a linear amplitude of 1.0. A value of -6 dB corresponds to a linear amplitude of 0.5, a value of -12 dB corresponds to a linear amplitude of 0.25, a value of -18 dB corresponds to a linear amplitude of 0.125, etc.

Loudness is a subjective measure that refers to the how loud a sound is perceived as by the listener. This depends on many factors such as the used speakers, their volume settings, and the ears of the listener. Therefore the software cannot say much about the perceived loudness of the sound, but it can indicate the relative intensity of the sound in qualitative terms, and it can indicate the mathematical magnitude of the sound.

The frequency affects the pitch of the generated sound.

Magnitude

This setting determines the overall intensity of the generated sound.

Normalize after filter
: The final intensity of the sound will be normalized to the entered value after the resonances have been applied.

Normalize before filter
: The generated sound will be normalized before it is filtered through the resonances. This can be used to compare the relative strength of different resonance settings. If the initial sound is too loud, this will generate clipping, so start with very low values, such as -40 dB.

Frequency Modulation (FM)

This setting enables a periodic change of the fundamental frequency (or pitch), for example to create a vibrato effect in singing.

FM Frequency
: The frequency of the pitch modulation. In singing, typical values of a singer's vibrato are around 6 Hz. That is the frequency with which the pitch of the fundamental will get faster and slower.

FM Amplitude
: This is the amount in cent by which the fundamental varies. A value of 100 cent will mean that the fundamental pitch will go up and down by one semitone.

Frequency Modulation (FM)

This setting enables a periodic change of the intensity of the generated sound.

AM Frequency
: The frequency of the intensity modulation.

AM Amplitude
: This is the amount in dB by which the magnitude of the generated sound varies. A value of 6 dB means that the intensity of the sound doubles with each oscillation.

Irregularity

This section allows to make each of the preceding aspects of the sound generation less regular by adding noise to it. This can prevent sounds that appear too "computerized" and unnatural by being too perfectly regular. The added noise is of the type "Coherent Noise" that can also be used as a fundamental waveform. It is not a truly random noise, but has a frequency and an amplitude.

To add noise to a parameter, select it in the `Apply noise to` field, and then define the properties of the noise generator on the right side. There is some overlap between the different ways of adding irregularity to a sound, because ultimately all parameters will either change its frequency or its amplitude. This can be done directly by adding noise to the fundamental, or indirectly, by adding irregularity to the FM or AM modulation.

Also, this section will add coherent noise to the generator, which is somewhat random, but still periodic. To add pure white noise, use the setting on the Fundamental and Overtones page.

Frequency
: The fundamental frequency of the noise generator. For human singers, a value of 5 Hz is a good starting point for many irregularities in the voice. You can also set values that change extremely slowly by entering frequencies such as 0.1 Hz, or even 0.01 Hz.

Range
: The range, in percent, by which the affected parameter will increase and decrease. If the target is a frequency modulation, this will affect the pitch, in cent. If the target is an amplitude modulation, the percentage value will affect the intensity.

Overtones
: The number of overtones for this noise generator. If it is set to 1, only there will be only one level of noise at the given frequency. With more than one overtone, higher frequencies of noise will be added.

Lacunarity
: This determines the spacing between overtones. More precisely, it is the frequency ratio of successive overtones. This setting is equivalent to the `Multiply by` setting for overtone spacing. A setting of 2.0 means that all overtones are octaves of the fundamental.

dB Delta
: The intensity reduction, in dB, for each successive overtone. This is equivalent to the `dB / Overtone` setting for the overtone amplitude.

Resonances

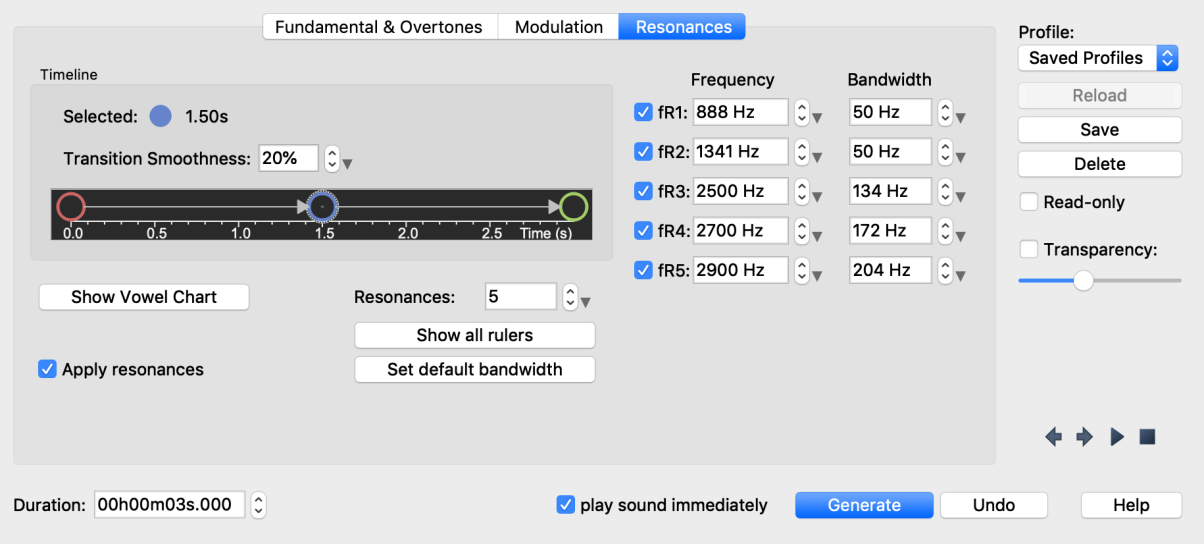

Figure 3.13. Sound Generator - Resonances

This section allows to pass the generated sound through a series of resonant filters akin to the vocal tract. Each filter has a frequency and a bandwidth that define the frequencies that this resonance amplifies / lets through.

This section of the Sound Generator is closely linked to the Vowel Chart and the Frequency Scale. You can create Resonance Path Nodes on the Vowel Chart, and you can change the frequencies of each resonance by moving the corresponding ruler on the Frequency Scale.

Simulate Vowel Adjustments

By default, the Generator will use the same resonances throughout the whole duration of the generated sound. However, the Sound Generator can also simulate how the resonances in the vocal tract change over time, for example when transitioning between different vowels, or when making more fine-grained adjustments to the current vowel when changing pitch to a different vocal register.

To get started with this, open the Vowel Chart, click on the  icon, and then click on the Vowel Chart to create two or more nodes.

Now you can see the created nodes on the Resonance Path Timeline. Click on a node to be edited, and then adjust the resonances for this point in time. When generating the sound, the resonances will change over time according to the defined path. When you click on any point in time on the timeline (but not on a node), the current values for the resonances at this point in time will be shown.

3.1.7. Frequency Filters

> **Note**
>
> Frequency Filters are only available in *VoceVista Video* and *VoceVista Video Pro*.

Frequency Filters are a powerful tool that allows you to listen to the different frequency ranges of a sound separately. This enables you, for example, to inspect individual overtones of a recording, or to reduce or amplify the singer's formant and hear the effect of this change.

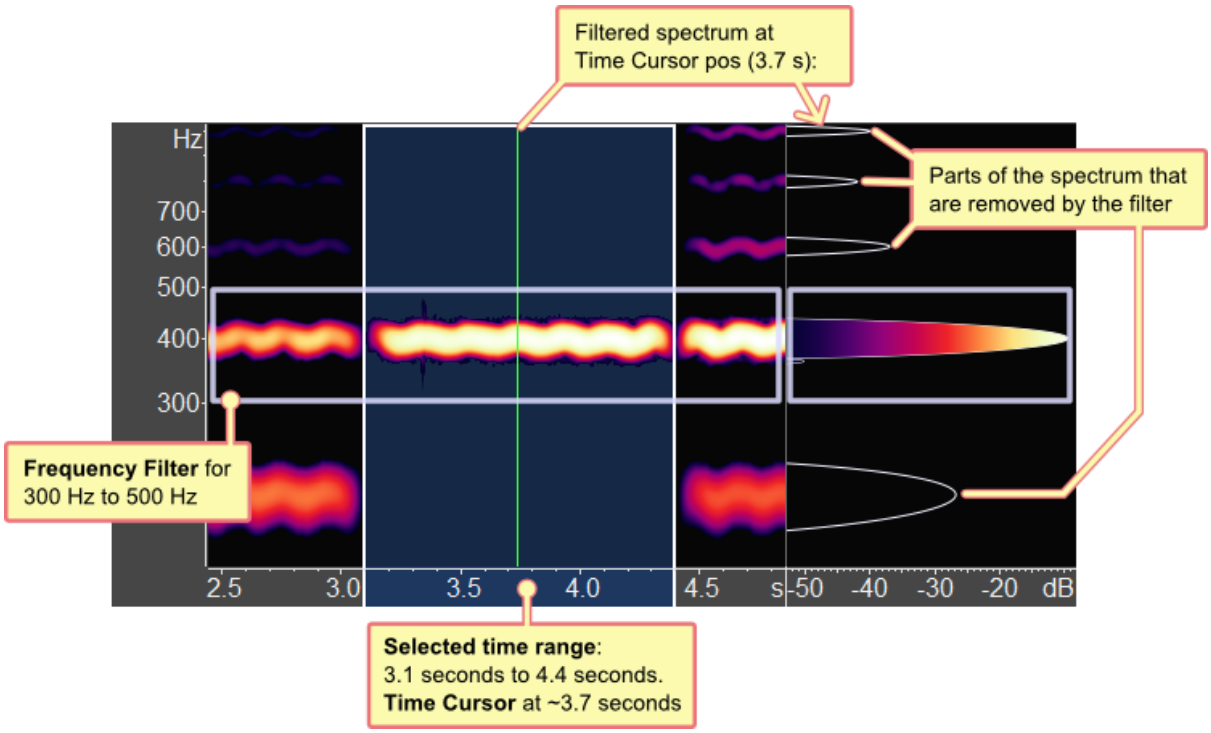

Figure 3.14. Using a Frequency Filter

Figure 3.14 shows a frequency filter that removes any frequencies outside of the range from 300 to 500 Hz. The left side of the image shows the spectrogram with a time selection. Only the selected time range is filtered.

The right side of the image shows the spectrum. The only part of the spectrum that is left by the filter is the second harmonic at 400 Hz. All other parts are removed, and will not be audible when the recording is played back. However, you can still see the outline of the removed parts.

Using Frequency Filters

Create a new filter by clicking on Filters → Add new frequency filter. If the new filter is not visible, scroll the Frequency Scale down or zoom it out until you see the filter.

You can move and resize the filter with the mouse. Here is an overview of the various controls:

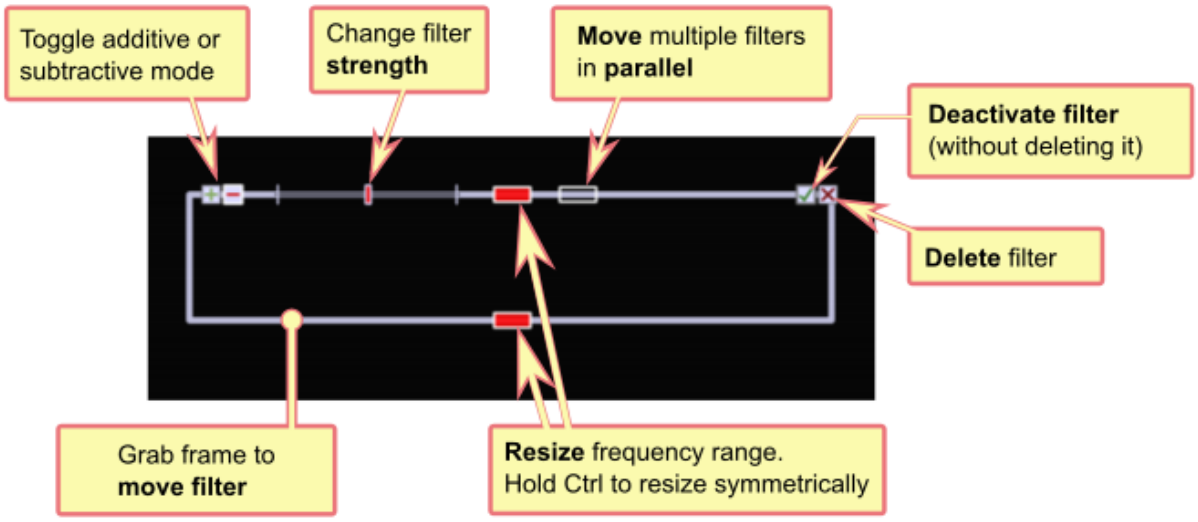

Figure 3.15. Filter Controls

Move filter

Simply grab the frame of the filter with the left mouse button to move it to a different frequency range.

Move multiple filters in parallel

If multiple filters exist, a new handle will appear next to the handle for resizing the frequency range. This new handle allows to move all existing filters in parallel as a group.

Resize filter

The width of the filter can be changed by moving the red resize handles in the middle of the filter. If you hold Ctrl while using the resize handles, the filter is resized symmetrically around its center.

Toggle additive or subtractive mode

The + and - buttons change the behavior of the filter. An additive filter (+ Mode) keeps the filtered frequencies and removes everything else. A subtractive filter (- Mode) removes the filtered frequencies and keeps everything else.

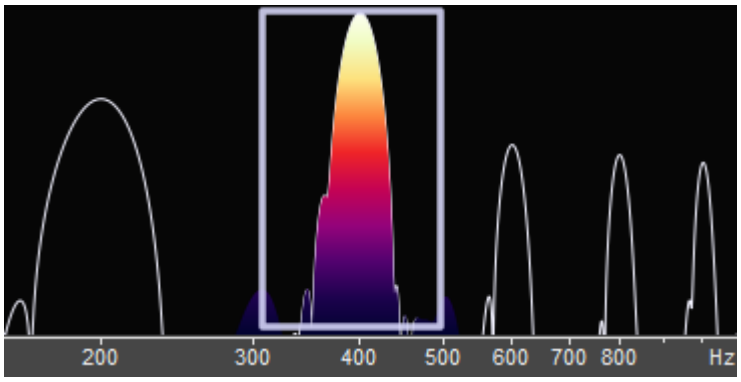

Figure 3.16. Additive filter (+ Mode)

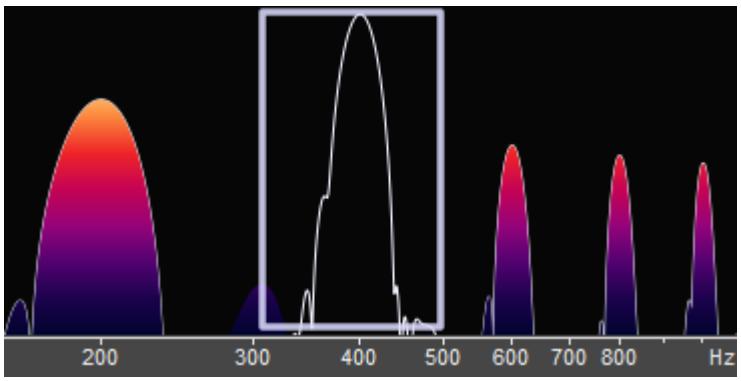

Figure 3.17. Subtractive filter (- Mode)

Change filter strength

By default, the filters are an all-or-nothing affair. This can be changed by moving the Filter Strength Slider (shown in Figure 3.15).

Deactivate filter

This button allows you to turn off a filter, so that it is still there, but has no effect. If you want to reactivate the filter, simply click the button again. This allows you to experiment with the effect of the filter.

Delete filter

When you are done with a filter, click this button to completely remove it.

Filters menu

The Filters menu has entries to add a new filter, and to toggle between the + and - modes. It also has the following entries:

Enable real time filter preview

If this option is checked, the effect of the filter is shown immediately, both on the spectrogram, and for playback. This means you can see and hear how your recording is affected by the filter. You can generally leave this on, but in some cases it may be helpful to turn it off. In particular, you might want to arrange filters and still see the unmodified spectrogram before turning the filters on.

Apply filters permanently

This command will apply the active filters to the current selection of your recording. In contrast to the real time preview, which does not change your sound file, this is a permanent alteration of the recording (although it can be undone by clicking on Edit → Undo).

3.1.8. Markers

> **Note**
> Markers are only available in *VoceVista Video Pro*.

Markers allow you to set bookmarks in files, and to add descriptions to them. This allows you to manage large collections of recordings and to find parts of interest again. Here is a screenshot of a recording with markers:

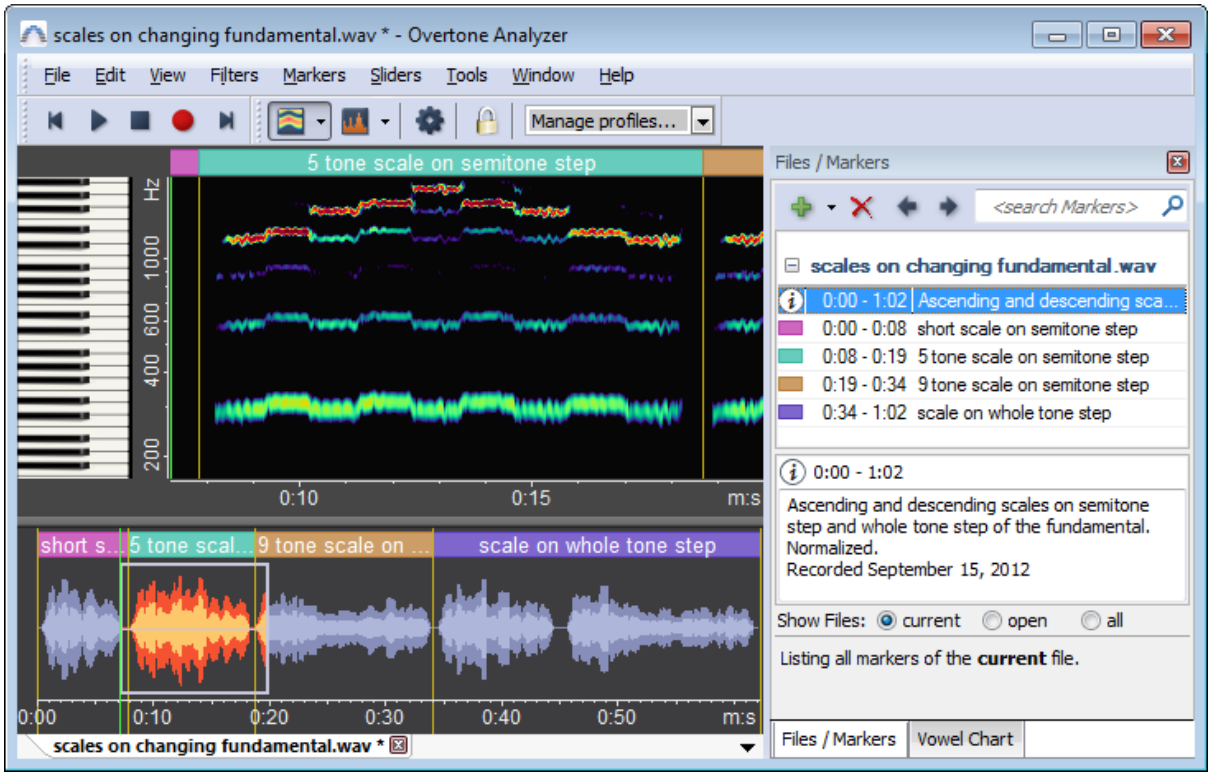

Figure 3.18. A recording with markers

Figure 3.18 shows an audio file that includes a number of separate sections. Each section has been marked and labeled with a marker. The “File and Marker View” on the right shows a list of all markers and can be used as a summary and a table of contents of the file.

Marker Types

File Description

The File Description is a special type of marker that is automatically added to every file. You can use it to add information that applies to the whole file, such as where it was recorded, when, and with whom. The File Description marker is always the first entry in the marker list of a file, and it has the icon.

Auto Markers

Every time you press record, record for a while, and then press stop, an Auto Marker is added to mark this last recorded segment. Auto Markers have a circle icon on the Marker List. Each segment is called a “take”, and takes are numbered.

By default, 0.5 seconds of silence are added after each take. This can be changed in the Advanced Settings.

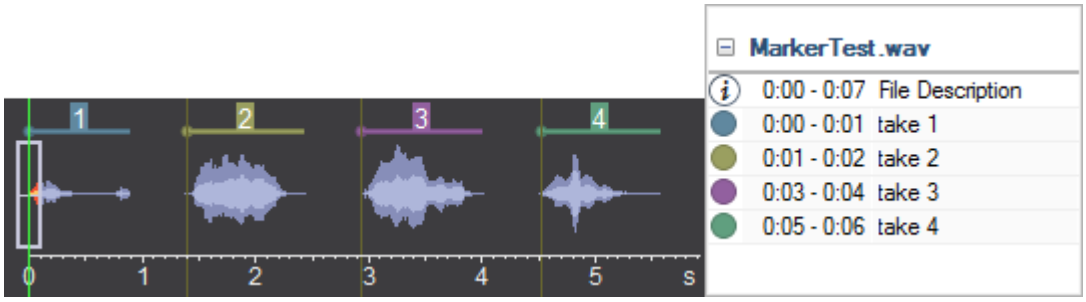

Figure 3.19. Auto Markers on the Timeline and the Marker List

Range Markers

A Range Marker marks a period of time with a beginning and and end. To add a new Range Marker, select a time period, and then click on Markers → Add marker. Range Markers have a rectangular icon on the Marker List.

Point Markers

A Point Marker marks a point in time. To add a new point marker, click somewhere on the Timeline or the Spectrogram to position the Time Cursor and reset the selection. Then click on Markers → Add marker. Point Markers have the triangle icon. They can also have the “Thumb up” or “Thumb down” icon. There is no difference between the Thumb Markers and regular Point Markers except for the icon.

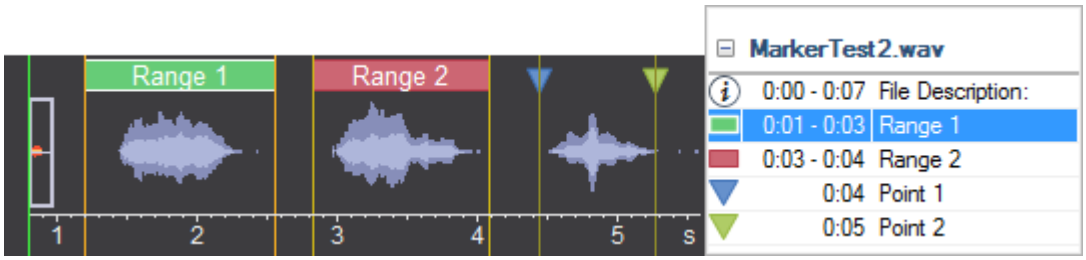

Figure 3.20. Range and Point Markers on the Timeline and the Marker List

Navigating with Markers

To jump back and forth between markers, click on Markers → Jump to previous marker or Jump to next marker. You can also use the Q and W keys for this. When you press Shift-Q or Shift-W, you will jump to the previous or next marker and zoom to it, so that the entire time range of the marker is visible and selected.

You can click on a marker on the Marker List to jump to it, or double-click on it to select its time range and zoom to it.

Editing Markers

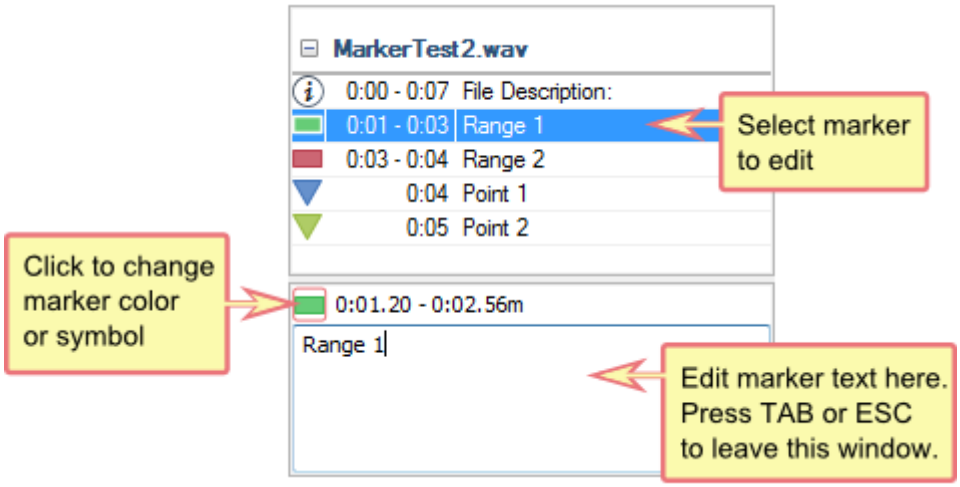

Figure 3.21. Editing Markers in the File / Marker List

You can edit the text and color of a marker in the lower half of the “File / Marker List” window . If this window is not visible, you can bring it up by clicking on Markers → Show File / Marker List.

To change the time range of a marker, simply move it with the mouse on the Timeline. You can also click and drag the edges of a Range Marker to change its start or end time.

Setting Marker Visibility

You can toggle the visibility of individual marker types with the Show markers and Hide markers commands in the Markers menu.

Searching Markers and Files

The search box at the top right corner of the “File / Marker List” window allows you to search for file names and for marker texts that contain specific words.

At the bottom of the window are three buttons that allow you to switch the scope of the search between searching all files, or just those that are currently loaded.

Recent file history

When no search term is entered, the File / Marker List shows a list of the most recently opened files in reverse order (so that the last file you opened comes first). Simply double-click on an entry in the list to open that file.

3.1.9. Overtone Sliders

An Overtone Slider is a visual tool that is laid over the spectrogram. A slider consists of lines that each represent a specific frequency. The distance between the lines corresponds to the harmonic series of a given fundamental. Overtone Sliders can be used to learn the harmonic series for each tone, and to analyze the harmonic content of a recording. They can also be used to transcribe the notes in a recording. Another use is to illustrate the role of overtones for composition and music theory. Overtone Sliders may also be called *Note Sliders*, depending on the context of use.

Main elements of Overtone Sliders

The following image shows an Overtone Slider with four harmonics, where the fundamental frequency is 110 Hz. On the piano, this would be an A. Each harmonic has a frequency that is a multiple of the fundamental. Therefore, the second harmonic has a frequency of 2*110 Hz = 220 Hz, the third one 330 Hz, and so on.

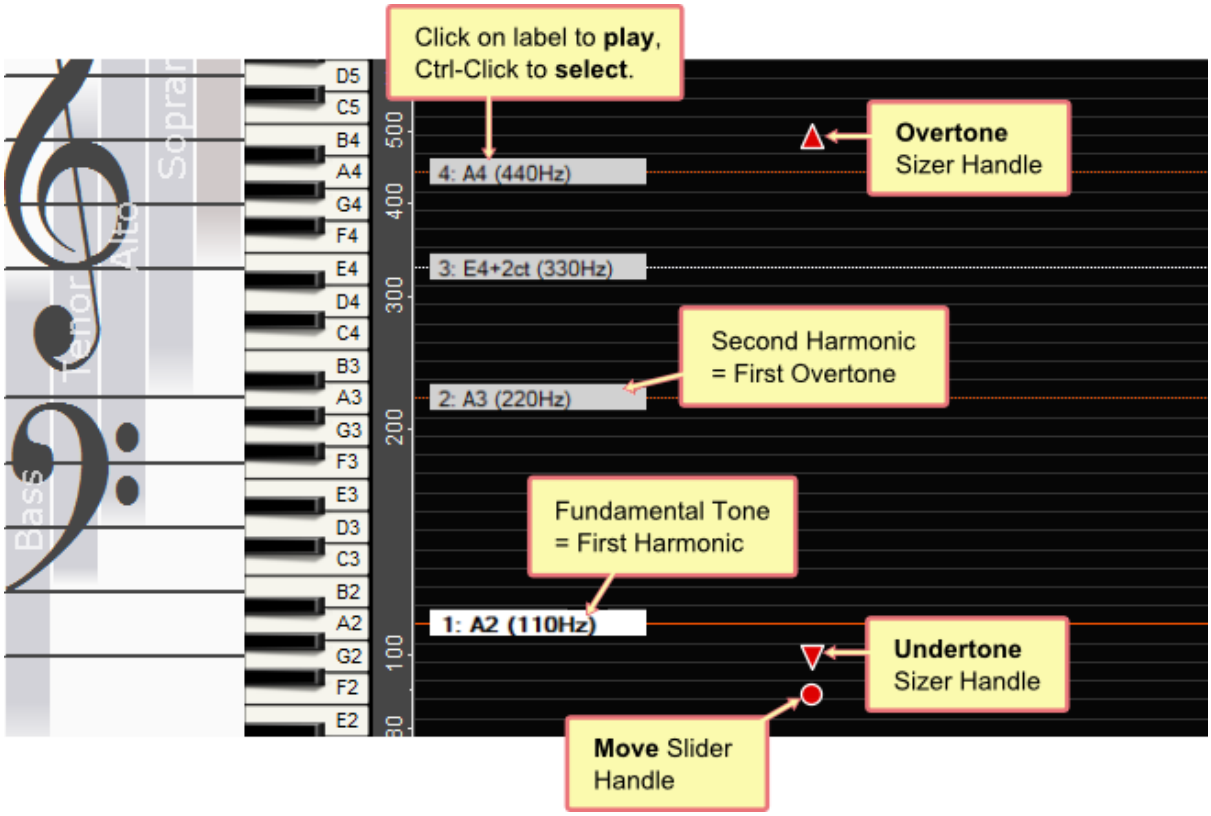

Figure 3.22. Main elements of Overtone Sliders

The image also shows the basic elements of the slider, including the controls that you can use to manipulate an Overtone Slider with the mouse:

Slider Labels

The labels show the number of the harmonic, and optionally the note name and frequency. This appearance can be adjusted on the Note Sliders settings page, and on the toolbar.

You can **click** on Labels to **play** them. The instrument used for playing them can be selected by the “Overtone Instrument” on the Playback Settings.

You can **select** one or more Labels with **Ctrl-Click**. Press on Sliders → Play selected tones (or press Enter to play them).

Slider Lines

Lines can be moved with the mouse. The lines have different colors:

1. Red - Octaves of the fundamental
2. White - Harmonics that are not octaves

3. Green - Harmonics that perfectly match the harmonic of another slider on the screen.

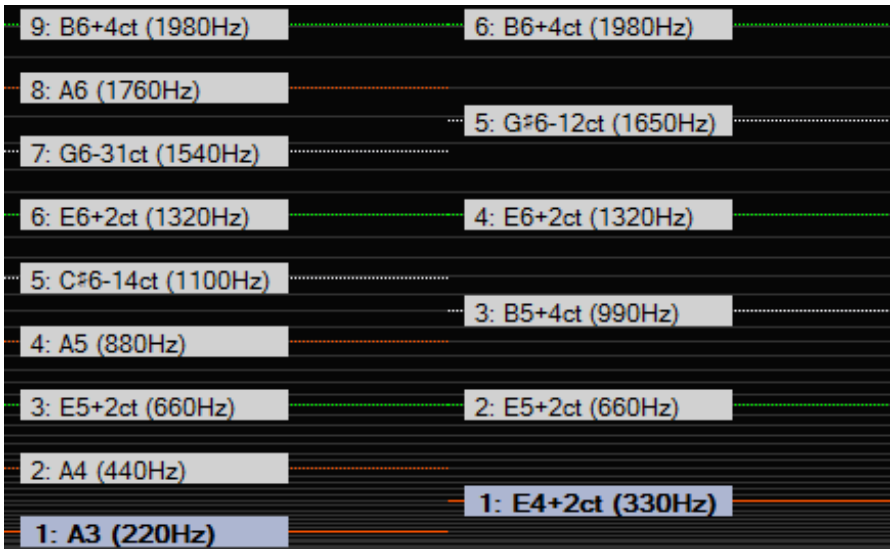

Figure 3.23. Two overtone sliders forming a perfect fifth

For example, Figure 3.23 shows two sliders whose interval forms a perfect fifth. The frequency of the second slider is exactly 3/2 times that of the first slider. Every third harmonic of the first slider matches every second harmonic of the second slider.

Overtone Sizer Handle
: The handles appear when you move the mouse over a slider. You can grab the handles with the mouse and move them. The Overtone Sizer Handle controls the number of harmonics (or overtones) shown for this slider.

Undertone Sizer Handle
: This handle controls the number of undertones shown for the corresponding slider.

Move Slider Handle
: This handle allows you to move the frequency or the time position of a slider. When the piano is vertical, moving this handle up or down has the same effect as moving a slider line with the mouse. Moving the slider left or right will leave the slider at the same frequency, but move its start and end position in time.

The Overtone Sliders Toolbar

Many aspects of sliders can be controlled through the Overtone Sliders Toolbar. If the toolbar is not visible, you can enable it by clicking on View → Toolbars → Overtone Sliders. The toolbar settings are explained on the Note Sliders settings page.

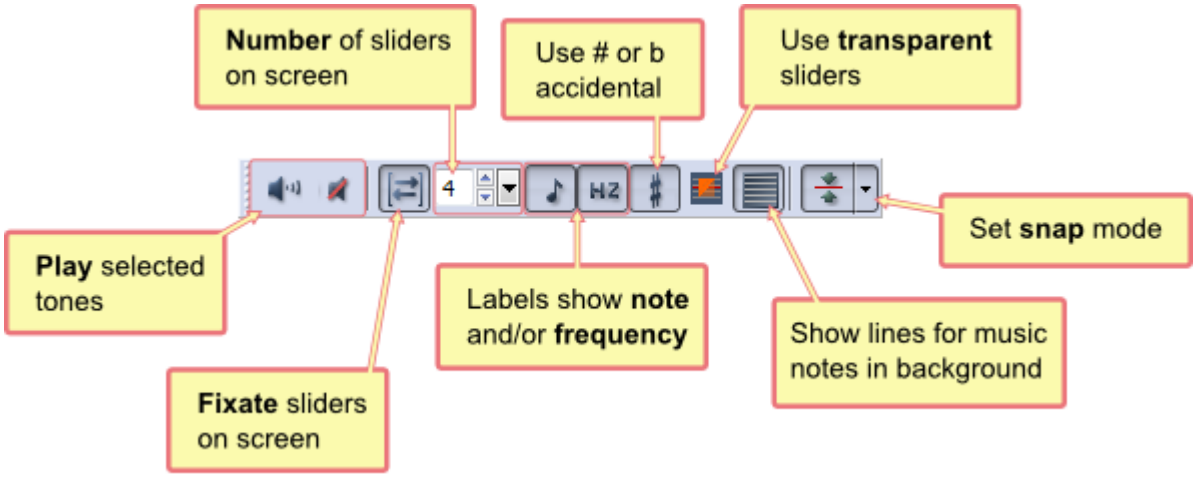

Figure 3.24. The Sliders Toolbar

Displaying Overtone Sliders

You can set the number of shown sliders on the toolbar, or on the Note Sliders settings page. There you can also choose to how display the labels of each slider, and whether to show the note name and the frequency value.

Manipulating Overtone Sliders

The following table lists the various mouse commands that you can use to manipulate overtone sliders. Refer to Figure 3.22 for the slider elements.

Mouse action	Effect
Change frequency	
Click + Drag Line	**Move line.** If snapping is enabled, the line will snap to the nearest tempered tone or the nearest spectral peak.
Shift + Click + Drag	**Move line in 1 cent increments.** This allows more precision than the normal way of moving.
Control + Shift + Click + Drag	**Move line in 1/4 cent increments.** This allows the highest amount of precision.
Double-click in empty space	**Move fundamental of slider** to that position. This is useful when the slider is outside the visible range.
Change frequency or time position	
Click + Drag Move Handle	**Move time position or frequency position** of slider, depending on the direction of the move.
Change number of displayed over- and undertones	
Move Overtone Sizer	**Change the number overtones** (or harmonics) that are shown for this slider.
Move Undertone Sizer	**Change the number of undertones** (or sub-harmonics) shown for this slider.
Double-click on Over- or Undertone Sizer	**Reset** number of over- or undertones to zero.
Snap line to reference frequency	
Alt + Click + Drag	**Move line, but with reversed snapping** behavior. For example, if snapping is enabled, holding Alt while moving the line will disable snapping, and vice versa.
Control + Click + Drag	**Move line and snap to nearest line** on the slider to the left, or the slider that currently has the focus. This is very useful to construct specific intervals. For example, to create a perfect fifth, snap the second harmonic of one slider to the third harmonic of another.
Control + Alt + Click + Drag	**Move line and snap to the nearest line** on the slider to the left, even if the focus is on a different slider. This is useful for quickly constructing scales.
Other commands	
Right-click on Label	Open **Context Menu** for this slider.

Table 3.4. Overtone Slider Mouse Commands

Selecting and Playing Slider Tones

The labels of the overtone sliders show information about the corresponding frequency and note, but they also serve as buttons that allow playing back the corresponding tone by clicking on them. This is explained in the section called “Selecting Tones”.

Once you have clicked on a slider label, that slider has the input focus and you can use the arrow keys on the keyboard to quickly move up and down the harmonic series.

Snapping Sliders

When moving sliders, it can be useful to snap the manipulated line to some reference. This may be a music note, a nearby peak in the spectrum, or another slider. You can control the current snapping behavior on the toolbar, on the slider settings page, and by using keyboard modifiers when moving a slider as explained in Table 3.4, “Overtone Slider Mouse Commands”.

Fixate sliders on screen or on Timeline

By default, sliders are fixated to the screen. When the underlying recording moves in time, the sliders stay at the same position on the screen. This allows you to use a small number of sliders as a reference, for example when practicing pitch, or when measuring the overtones in a sound.

The Sliders can also be fixated to the Timeline (when the option “Fixate note sliders on screen” is **not** checked). In that case each slider has a specific start and end time, like the notes in a piece of music. This can be used to transcribe the notes in a recording, or to construct new compositions or scales. In this mode the Timeline shows a **miniature view of all the note sliders** in a recording:

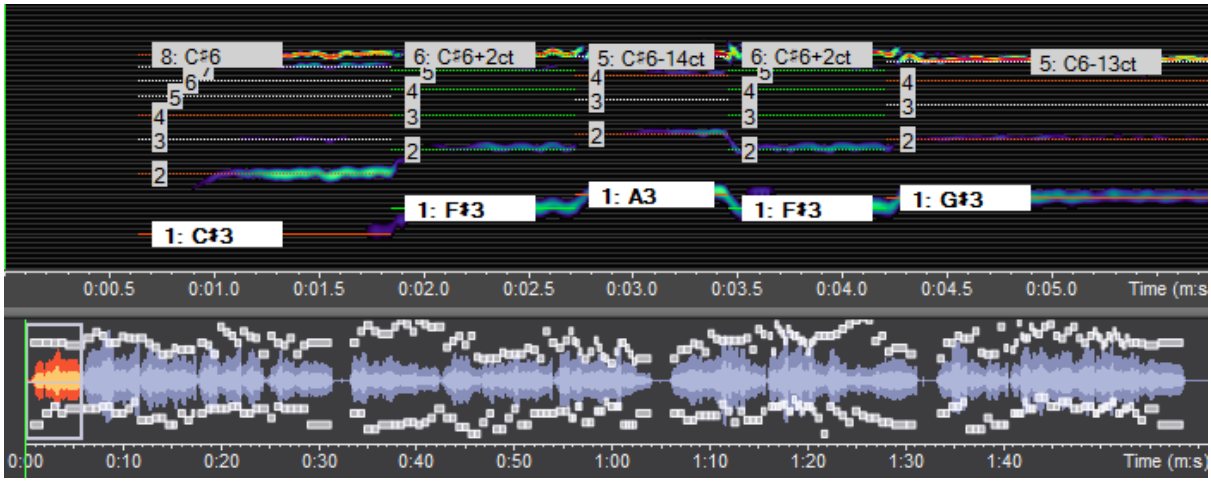

Figure 3.25. Note Sliders fixated to the Timeline

Figure 3.25 shows a recording that is about 2 minutes long. In the upper part of the image, the spectrogram shows the first five seconds. The lower part contains the Timeline with on overview of the entire piece, including a miniature view of the Note Sliders. The lower line shows the fundamentals, and the upper line shows the highest overtones for each note.

Editing Note Sliders

Insert note slider at selection

When the sliders are fixated to the screen, you can simply add more sliders by increasing the number of sliders on the slider settings page, or on the toolbar. This doesn't work when the sliders are fixed in time. Instead, you can select a period of time and then click on Sliders → Insert note slider at selection. The slider will be inserted for the selected time period, and it will have the frequency of the most common fundamental of this period. If you don't see the new slider, scroll the frequency scale or zoom it out.

Use Note Transcription Tool

To speed up the workflow of selecting a time range and inserting a new slider for consecutive notes, you can use the Note Transcription Tool in the Sliders menu. When the tool is activated, the mouse cursor will change to have a note symbol:

Now you can simply click on the boundaries of consecutive notes on the spectrogram, and the transcribe tool will insert a new note slider with every click.

To leave the transcribe tool, you can double-click, right-click, press ESCAPE, or press the keyboard shortcut for the transcription tool (which by default is T).

Delete selected note sliders
: This command will remove any sliders in the current selection without affecting the audio data in this time period.

Space sliders evenly across selection
: This command ensures that the sliders in the current selection all have the same width and are adjacent and non-overlapping.

Sort sliders by pitch
: This command may reorder the selected sliders such that they are sorted by their fundamental frequency, where the slider with the lowest fundamental will come first. This can be used, for example, during the construction of scales.

Use lowest undertone as fundamental
: If a slider has some undertones (in other words, if the Undertone Sizer Handle has been drawn down), this command will swap the lowest undertone and the fundamental.

Note Practice Mode

If you want to practice singing specific patterns or compositions, you can click on the sliders in one window to control the single slider in another. To understand how this works, first close all open windows by clicking on Windows → Close all. Then open two new windows by clicking on File → New twice. Arrange the windows by clicking on Windows → Tile horizontally. Now create three note sliders in the first window and one note slider in the second window. Move the sliders in the first window so that they have different frequencies. Then click on their labels. Note that the slider in the second window always jumps to the label on which you clicked in the first window.

When you load an existing recording in the first window, and have a new recording in the second window, the practice mode can be used to practice the pitch and timbre of individual notes in an existing piece one by one.

Loading and saving sliders

When a recording that contains Note Sliders is saved, the sliders are automatically saved in the file for most file formats. However, it is also possible to save the sliders into a separate file, for example to save the note track after a file has been transcribed. This is described in Section 3.3.3, “Loading and Saving Overtone Slider Layouts”.

3.1.10. Vowel Chart

Rulers are laid over of the spectrogram or waveform. They are a visual guide that highlights a specific frequency or amplitude.

The Vowel Chart shows the two frequencies that are needed to form specific vowels in various languages.

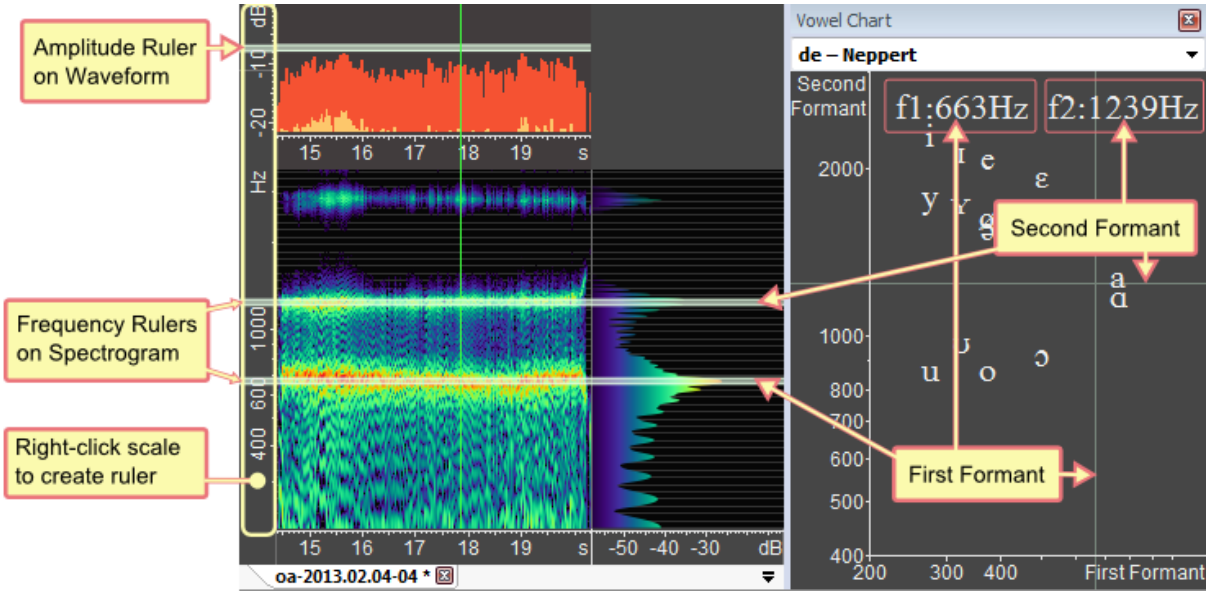

Figure 3.26. Rulers and the Vowel Chart

Frequency Rulers

Frequency Rulers mark a specific frequency. They can be created by right-clicking on the Frequency Scale (as shown in Figure 3.26) and selecting Insert new ruler here. To move a ruler, you can click and drag it on the Frequency Scale.

To delete a ruler, right-click on it on the Frequency Scale and then select Delete this ruler. You can also delete it by dragging it out of the visible range.

Amplitude Rulers

Amplitude Rulers work in the same way as Frequency Rulers, except that they are created and moved by clicking on the Amplitude Scale, and that they mark a specific amplitude instead of a frequency.

The Vowel Chart

The Vowel Chart can be used to explain how sounds of specific frequencies form recognizable vowels. This is based on the finding that vowels in many languages are created by amplifying two frequencies in the oral cavity. The chart shows the first frequency on the horizontal axis, and the second frequency on the vertical axis. On the chart are several symbols from the International Phonetic Alphabet (IPA). The location of the symbol gives the approximate first and second frequencies that form the corresponding vowel.

If the Vowel Chart is not visible, you can enable it by clicking on View → Vowel Chart. At the top of the chart is a list box that lets you select a language. There may be several entries for some languages because different researchers have identified different frequencies for each vowel. In any case, the two frequencies for each vowel are not exact. The chart just gives a general idea of the frequency range where a certain vowel may be heard.

Clicking on the Vowel Chart will move the lower two Frequency Rulers to the corresponding position. If necessary, two new Frequency Rulers will be created.

The frequency ranges and the scale mode of the Vowel Chart can be adjusted in the advanced settings, or by right-clicking on the Vowel Chart and selecting Edit frequency ranges.

Harmonic Vowels

The Vowel Chart can be configured to show the frequencies of Harmonic Vowels. These are Vowels whose first and second frequencies have a whole number ratio such as 2/1 or 3/2. To enable this display, right-click on the Vowel Chart and select Show harmonic vowels.

3.1.11. EGG (Electroglottograph) Display

The Electroglottograph, or short "EGG" is a physiological signal that allows us to follow the vibrations of the vocal folds that produce the primary sound at the glottis. A small high-frequency current runs between two electrodes that are held in place on either side of the larynx. The conductance between the electrodes increases by a small amount when the vocal folds make contact, initiating the closed phase of the glottis. The conductance decreases again as the glottis opens. These modulations in conductance give us the EGG signal, useful as a measure of contact between the vocal folds.

3.1.12. Toolbar

The Toolbar provides quick access to many frequently used functions of the program. If you point the mouse cursor on a toolbar button, a description of this button will pop up, and a more detailed description will appear at the bottom of the main window on the Status Bar.

Toolbar Icon Size

The size of the toolbar icons can be adjusted on the Toolbar Icons under Advanced Settings.:

Figure 3.27. Extra Small Toolbar Size

Figure 3.28. Small Toolbar Size

Figure 3.29. Normal Toolbar Size

Figure 3.30. Large Toolbar Size

Figure 3.31. Extra Large Toolbar Size

Toolbar Profiles

The Profile List on the toolbar contains several predefined toolbar arrangements that you can activate by clicking on one of the entries in the list.

The following images show the toolbars for the *VoceVista Video Pro* edition. In the other editions, not all icons will be visible.

Minimal Toolbar

Figure 3.32. Minimal Toolbar

A minimalistic toolbar for absolute beginners. Only contains buttons for recording and playback.

Single Row Toolbar

Figure 3.33. Single Row Toolbar

A toolbar with most standard buttons for recording, editing, and visualization. Fits one a single row of buttons to preserve screen real estate.

Standard Toolbar

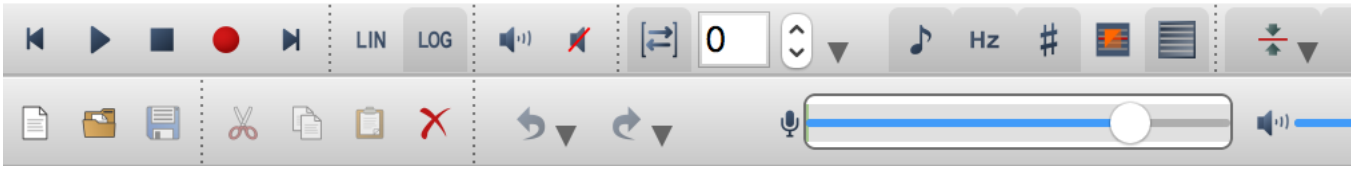

Figure 3.34. Standard Toolbar

A two-row toolbar with buttons for most operations, including Overtone Sliders, zooming, and window management. Recommended most users that have some experience in using the program.

Full Toolbar

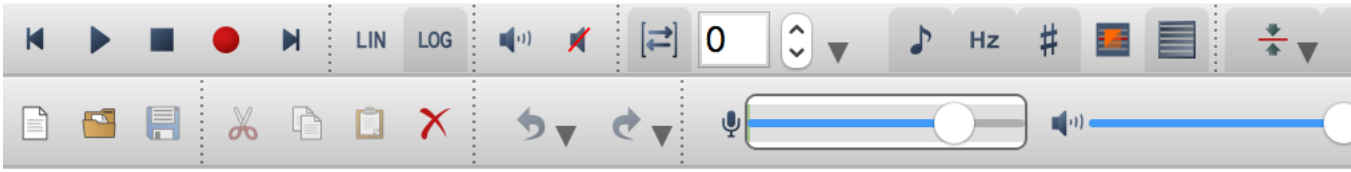

Figure 3.35. Full Toolbar

A two-row toolbar with most buttons that are available, including controls to set the frequency scale range.

3.2. Preferences

Most settings in VoceVista Video are controlled through the View Menu and the Preferences Window. The most common settings are also on the toolbar.

You can bring up the preferences by clicking on the preferences toolbar button, by clicking on VoceVista / Preferences on macOS, or Tools / Options on PC, or by pressing F10. Below are explanations for each different part of the options dialog.

3.2.1. Recording Settings

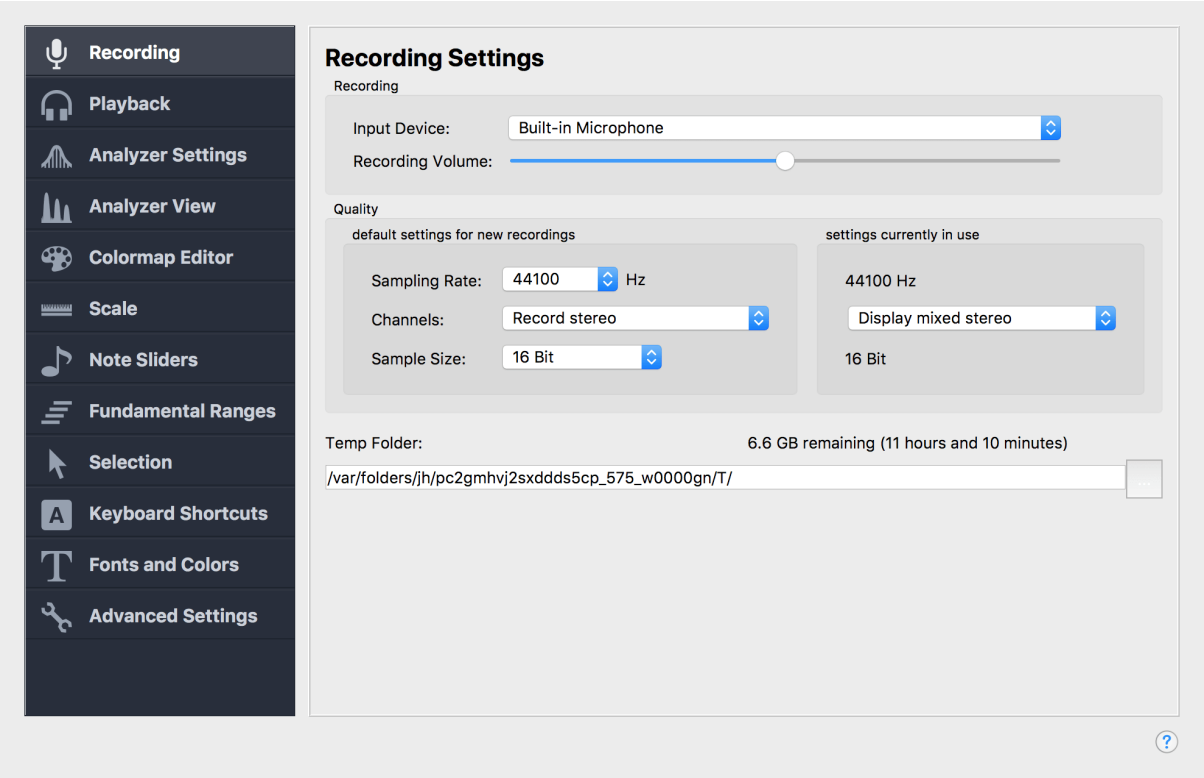

Figure 3.36. Recording Settings

On the Recording Settings page you can select the sound card from which to record, and you can adjust its volume. You can also adjust the format of the recorded data, and you can select which stereo channel is displayed from the active recording.

Input Device

This is the input line from which to record audio data when you press the record button. If you have more than one sound card installed, for example because you are using an external USB microphone, make sure that the correct one is selected here.

Recording Volume

This determines the sensitivity of your microphone or input line. It is very important to set this so that the full dynamic range is used without clipping to get a good recording. A clipped signal will cause artifacts in the spectrogram.

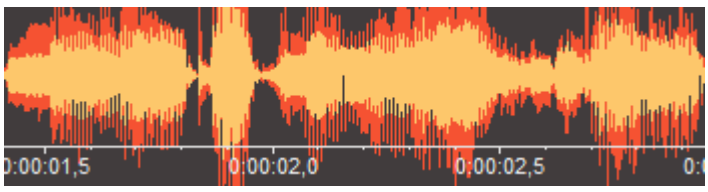

Figure 3.37. input level too high

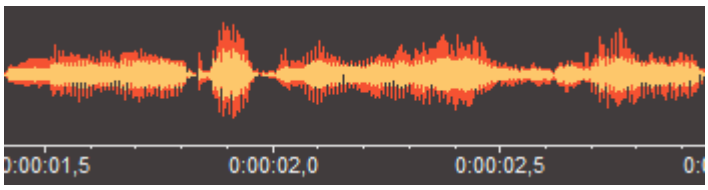

Figure 3.38. recording volume ok

Recording source and volume can also be set through the input level meter on the toolbar. You can set the recording volume by moving the slider, and you can select the input source by right-clicking on the input level meter:

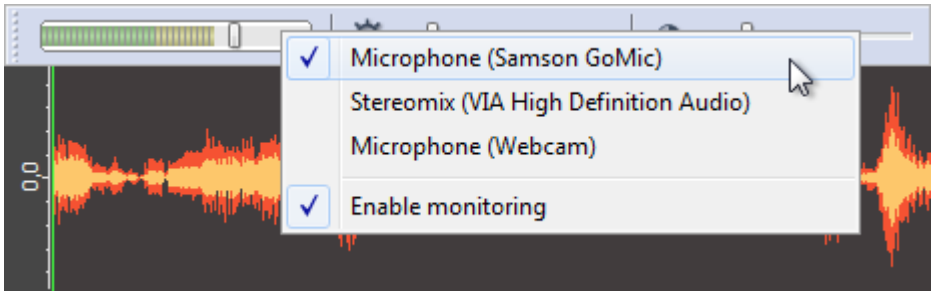

Figure 3.39. Right-clicking on input volume meter on toolbar

Quality

> **Note**
>
> These settings can only be changed on new files, before you have made any recording. The format of existing recordings cannot be changed. If you want to change these settings and they are disabled, click on File → New first.

Sampling Rate

The Sampling Rate determines the number of measurements (samples) per second recorded from the input source. For most purposes, a sampling rate of 44100 Hz is appropriate. When using the EGG, a sampling rate of 48000 Hz might give the best results.

Channels

The channel format allows you to select between stereo and mono recordings. On the left side, under *default settings for new recordings*, the channel format is set for new recordings. Changing this setting will not affect any files that are already open.

On the right side, under *settings currently in use*, you can change how the currently active file is displayed. This can also be changed on the toolbar.

Here are the various channel display options:

default settings for new recordings

Record mono
: Record a mono file with a single audio track.

Record stereo
: Record a stereo file with two audio channels, a left and a right channel.

Record left/right stereo channel
: This is useful when a single mono microphone is connected to the left or right channel of an audio interface. In that case, when recording "mono", the audio interface would mix the channel of the microphone with the empty channel that has nothing connected. By recording only the left or only the right channel, this mixing is avoided.

The recording will be saved as a mono file with a single channel.

Record Audio left; EGG right / EGG left; Audio right
: Record a stereo file with two channels, where one channel is interpreted and displayed as the audio channel recorded from the microphone, and the other channel is interpreted and displayed as the Electroglottograph (EGG).
This option is only available in VoceVista Video Pro.

settings currently in use

Display mono
: When a mono recording is active, the only available option is to display its single channel.

Display mixed stereo
: When a stereo recording is active, this setting will display both channels mixed together. So the spectrogram will show audio content from both channels.

Display left / right stereo channel
: With these settings, only the left or the right channel of a stereo recording will be displayed.

Display Audio left, EGG right / EGG left, Audio right
: Interpret one channel as audio data, and one channel as EGG.

Sample Size

The sample size influences the quality but also the storage requirements of the recorded sound. The default setting of 16 bit per sample corresponds to the quality of standard audio CDs and should be sufficient for most purposes. A higher bit depth of 24 or 32 bit can be useful when the dynamic range of the recorded signal is very high, and when it contains very quiet passages that should later be amplified.

Temp Folder

This is the folder where VoceVista Video will store data as you are recording it. Make sure that you have enough hard drive space left. A mono recording with 44100 samples per second and a sample size of 16 bit requires about 350 MB (half a CD) per hour.

3.2.2. Playback Settings

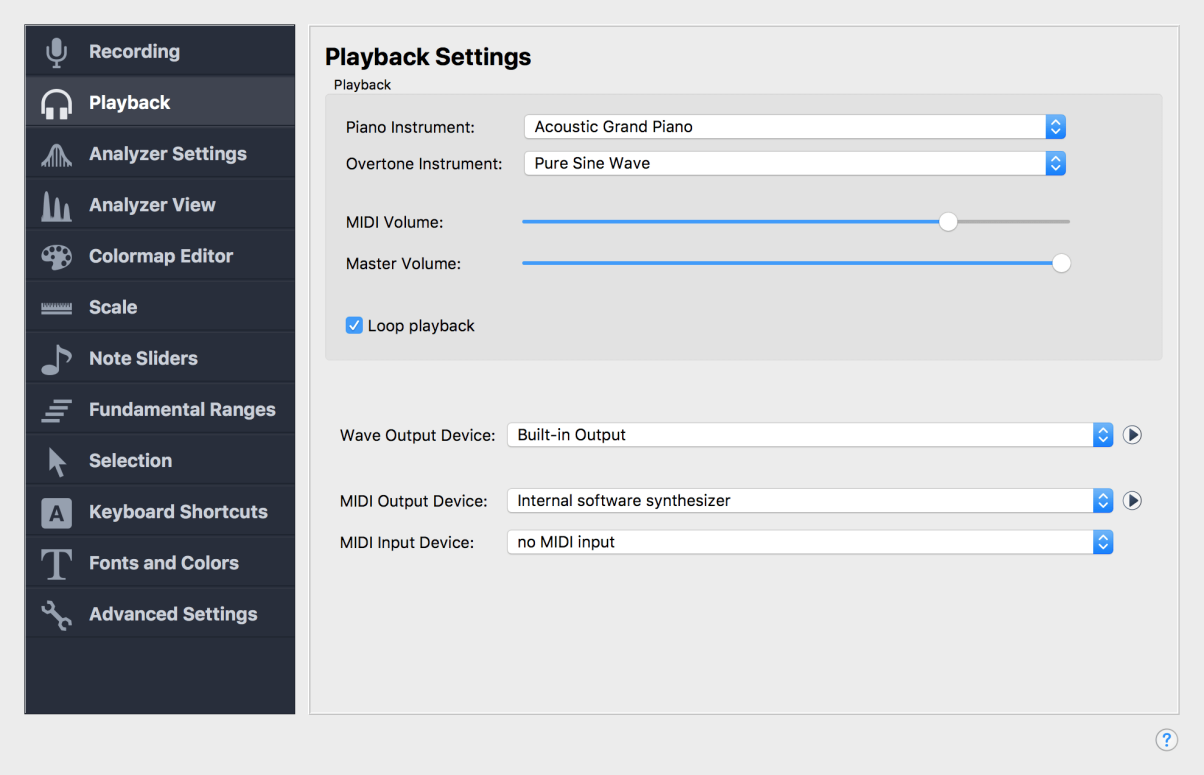

Figure 3.40. Playback Settings

Playback

Piano and Overtone Instrument

The two settings for Piano Instrument and Overtone Instrument control the instruments used for playing tones when you press keys on the piano, or labels on the Overtone Sliders. You have a choice between a large list of standard MIDI instruments, and between a sine wave generator that plays pure sine waves:

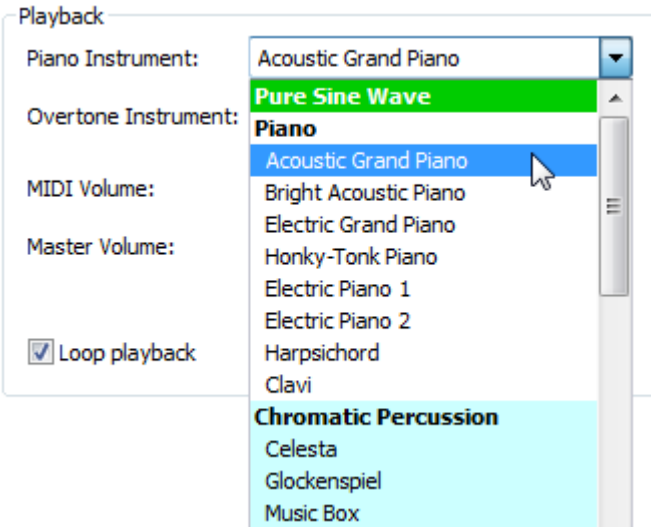

Figure 3.41. Piano and Overtone Instruments

Sine Wave Generator

The first entry in the list, with the label Pure Sine Wave, uses a tone generator that plays pure sine waves with high accuracy.

MIDI Instruments

All other entries use the standard MIDI instruments, and their sound depends on the quality of your sound card or of your connected keyboard. Their volume is controlled by the MIDI Volume slider.

The frequency of the MIDI Instruments may be less accurate than that of the sine wave generator.

MIDI Volume

This volume affects all sounds that are played through the MIDI Output device when playing piano keys or note sliders.

Master Volume

This slider affects the playback volume for all other sound output, including playing recorded audio content.

Loop playback

If this option is checked, playback of the current file or selection will be repeated in a loop. Otherwise, playback will stop when it reaches the end of the current selection or file.

Devices

Wave Output Device

This is the device for playing back the recorded sound.

Midi Output Device

This is the device for playing the sounds from the piano keyboard and the overtones (except for the "Pure Sine Wave" overtone sound, which is played through the wave output device).

Midi Input Device

If you have connected an external Midi keyboard, you can select it as input source here. Pressing a key on your external keyboard will have the same effect as clicking on the piano view with the mouse.

3.2.3. Analyzer Settings

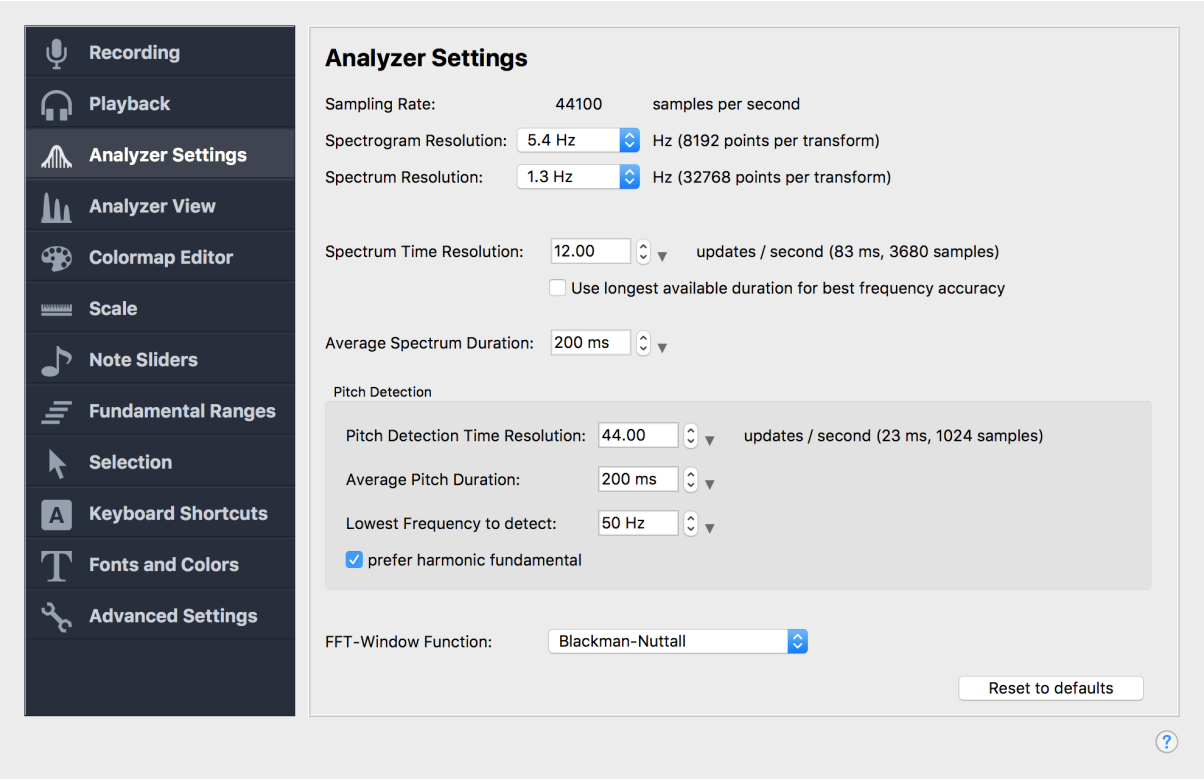

Figure 3.42. Analyzer Settings

The Analyzer Settings determine the frequency and time resolution and other aspects of the spectrum, spectrogram, and pitch analysis. In particular, the Time Resolution has a large influence on the kind of information that the analyzer will extract from a recording.

Sampling Rate

The Sampling Rate determines the number of measurements (samples) per second recorded from the input source. You can measure frequencies of up to half the sampling rate. For example, if the sampling rate is 11025 samples per second, then the analyzer can measure frequencies up to around 5500Hz. Standard music CDs have a sampling rate of 44100Hz. This should be a reasonable setting for most practical purposes, as it gives a good balance between frequency resolution and sound quality (however, see the note below).

> **Note**
>
> The sampling rate can be changed on the Recording Settings page.

VoceVista Video uses two separate mathematical methods to analyze sound. One method is the Fast Fourier Transform (FFT), which calculates the Spectrum. This gives the intensity of the individual frequency components of a sound. The other method is the detection of the fundamental pitch, which is completely separate from the FFT. The two methods are both dependent on the sampling rate of a recording. However, while lowering the sampling rate can increase the accuracy of the FFT in same circumstances, the pitch detection works best with sampling rates of 44100Hz or higher, and the EGG works best with 48000Hz.

Frequency Resolution (FFT Size)

Spectrogram Resolution

This is the Frequency Resolution for the Spectrogram on the Long Term View.

Spectrum Resolution

This is the Frequency Resolution for the Spectrum on the Short Term View.

The frequency resolution is the smallest difference between two frequencies that the analyzer can distinguish. Internally, this setting is stored as the size of the Fast Fourier Transform (FFT), which is the number of points that are computed for each update.

A higher FFT Size gives you more accuracy and shows more detail in the spectrum and spectrogram, but it also requires more processing power and may slow down your computer. In general, you should choose the highest setting that still gives you acceptable performance when moving the range slider on the Timeline.

Spectrum Time Resolution (updates per second)

This setting determines whether the Analyzer should be more accurate in the frequency or in the time domain. In other words, are you more interested in measuring the exact pitch, or in measuring the melody (the variations of the pitch over time)? Fewer updates per second will increase the accuracy of the pitch display but hide the melody, while more updates per second will cause the analyzer to show the melody more clearly, but with less accuracy in the pitch.

use longest available duration

Checking this box will automatically select the longest possible duration for the given frequency resolution at the expense of showing fewer details of the melody.

Average Spectrum Duration

The Spectrum on the Short Term View will be averaged over the given time period to avoid flicker and make it more stable. With a smaller duration, it will be more responsive to sudden changes, but flicker more, while a longer value will make it appear sluggish but more stable.

Pitch Detection

Pitch Detection Time Resolution

This is similar to the Time Resolution for the FFT, but is used by the algorithm for pitch detection, which is separate. The Time Resolution value determines the length of a recording segment that the pitch algorithm considers to find its fundamental pitch. Changing this value is for advanced users only, as the default value works best in most situations.

Average Pitch Duration

The current pitch will be averaged over this period of time during recording and playback to make it more stable and avoid flickering.

Lowest Frequency to detect

The pitch detection module will ignore frequencies below this value, which can avoid false detections when the recording has low-frequency noise.

prefer harmonic fundamental

If checked, the pitch detection module will attempt to find the pitch that is the fundamental of a harmonic sound. Otherwise it might pick the loudest harmonic, if that is much stronger than the fundamental.

FFT-Window Function

Due to the way the FFT method works, its input needs to be filtered through a window function to avoid incorrect results. A discussion of this can be found at en.wikipedia.org/wiki/Window_function [https://en.wikipedia.org/wiki/Window_function].

For VoceVista Video the default window function *Blackman-Nuttall* works very well and gives a useable dynamic range of about 100 decibels. It should not be changed unless you understand windows functions and their respective properties and trade-offs.

Reset to defaults

This button will restore the default Analyzer Settings.

3.2.4. Acoustic Analyzer View

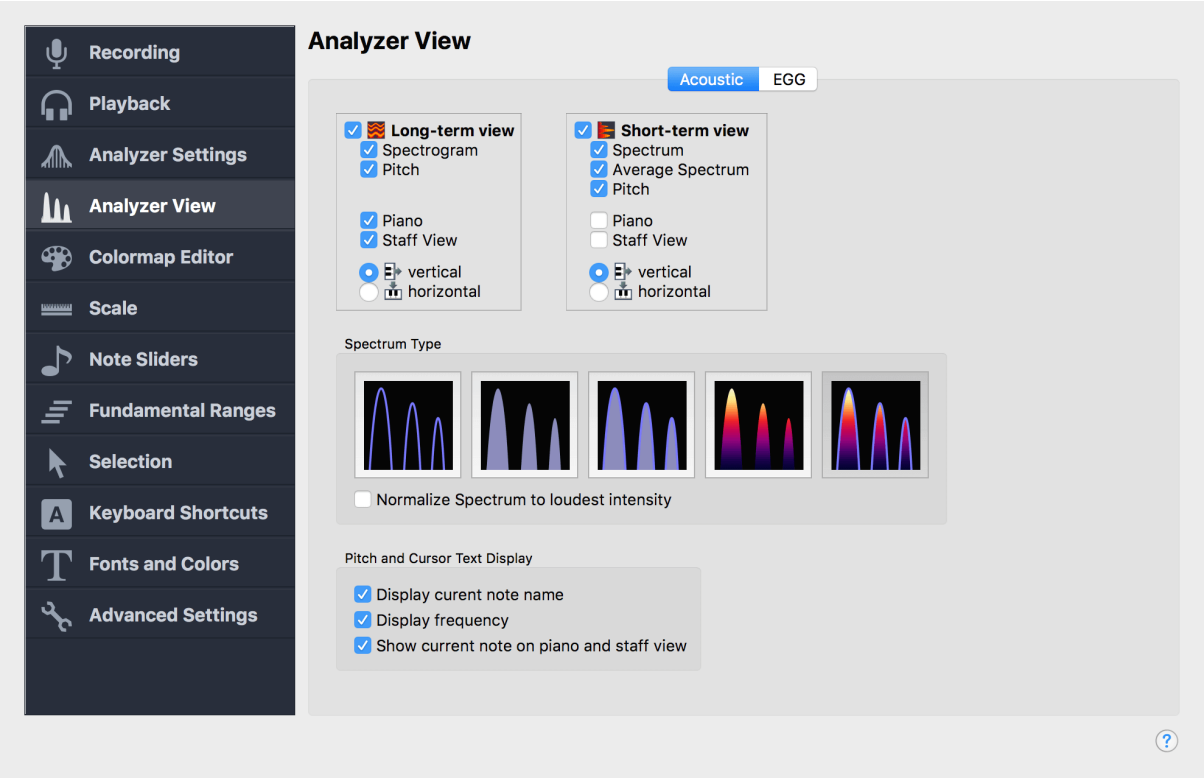

Figure 3.43. Acoustic Analyzer View

On this settings page you can choose the appearance and behaviour of the Analyzer View.

Analyzer Display

This selects the kind of analysis that you want to show. The main analyzer window can show a single view, or it can be split into a long-term view (for the spectrogram) and short-term view (for the spectrum). If both views are active, the cursor position in the long-term view determines the content of the short-term view.

The spectrum shows a graph of intensity per frequency. The spectrogram shows how the spectrum changes over time. Pitch shows the fundamental frequency. Showing neither spectrogram nor pitch can useful if you want to work with the Overtone Sliders alone.

Spectrum Type

Here you can select various options for the appearance of the spectrum.

Normalize Spectrum to loudest intensity

If this is checked, the spectrum will show the loudest spectral peak at 0dB and scale all other values accordingly. If it is not checked, the intensity of the spectral peaks will be shown as their actual value.

Pitch Display

These options allow you to customize the appearance of the pitch display. The current note always corresponds to the fundamental pitch at the current cursor position.

Display current note name
: If the short-term view is active, the name of the current note is displayed on top of the spectrum.

Display frequency
: If the short-term view is active, the frequency of the current pitch will be shown in Hz on top of the spectrum.

Show current note on piano and staff view
: If this is checked, the current note is highlighted on the piano, and shown as musical note on the staff view.

3.2.5. EGG Display Settings

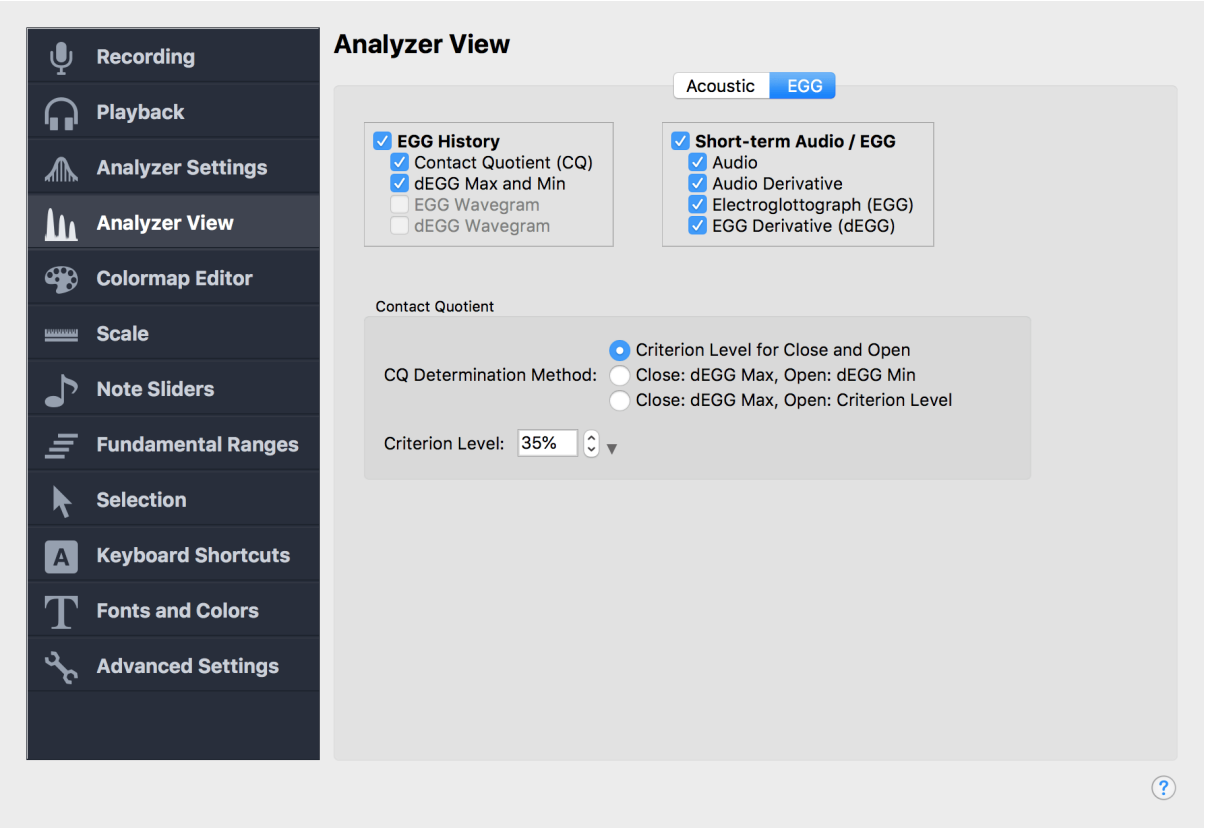

Figure 3.44. EGG Display Settings

On this settings page you can choose the appearance and behaviour of the Electroglottograph (EGG) display.

> **Important**
>
> The EGG features are only available in *VoceVista Video Pro*. Please view our products page [https://www.sygyt.com/en/products/] for a comparison of the different editions.

EGG History

The EGG History shows how the contact quotient and the dEGG Maximum and Minimum change over time. It is equivalent to the long-term analyzer display that shows the spectrogram for audio content.

Contact Quotient (CQ) History
: Select this option to display the Contact Quotient in the EGG History view.

dEGG Max and Min
: If selected, this will display the maximum and minimum of the derivative of the EGG (dEGG) in the EGG History.

EGG Wavegram
: This will show the EGG Wavegram in a future version of the software.

dEGG Wavegram
: This will show the dEGG Wavegram in a future version of the software.

Short-term Audio / EGG

The short-term EGG display shows a few periods of the EGG and the Audio signal. It is equivalent to the short-term view that shows the spectrum.

The short-term EGG View has an upper half for the audio signal, and a lower half for the EGG signal. You can select which of the following signals to display:

Audio
: This will display a few periods of the audio channel.

Audio Derivative
: If selected, the derivative of the audio signal will be displayed.

Electroglottograph (EGG)
: Display the EGG signal.

EGG Derivative (dEGG)
: Display the dEGG, which is the derivative of the EGG signal.

Contact Quotient

The *Contact Quotient (CQ)* is a measure of the amount of vocal fold contact. When looking at a period of the EGG signal (the amount of time until the signal repeats itself), a CQ of 50% would mean that the vocal folds are have contact for half the period, a CQ of 25% percent would mean that they have contact for one quarter of the period, and so on.

It is important to note that the EGG cannot give complete information about the vocal fold closure, it can only indicate that the contact of the vocal chords has somehow increased. By looking at the EGG alone, we cannot with certainty say if the vocal chords are completely closed, or if they are closed at all. That's why it is more accurate to use the term *Contact Quotient* instead of *Closed Quotient*.

There are several methods to calculate the Contact Quotient:

Criterion Level for Close and Open
: This method draws an imaginary line across the EGG signal and measures where the EGG signal crosses the line. The part where it is above the line is counted as closed, and the part below the line is counted as open. The location of the line is called the *Criterion Level (CL)*. The value of the Criterion Level is arbitrary, and can be changed by entering a value on this settings page, or by dragging the measurement line on the EGG display. Calculating the CQ with the Criterion Level is potentially less accurate than the other methods, but it is more robust than the other methods when the signal is noisy.

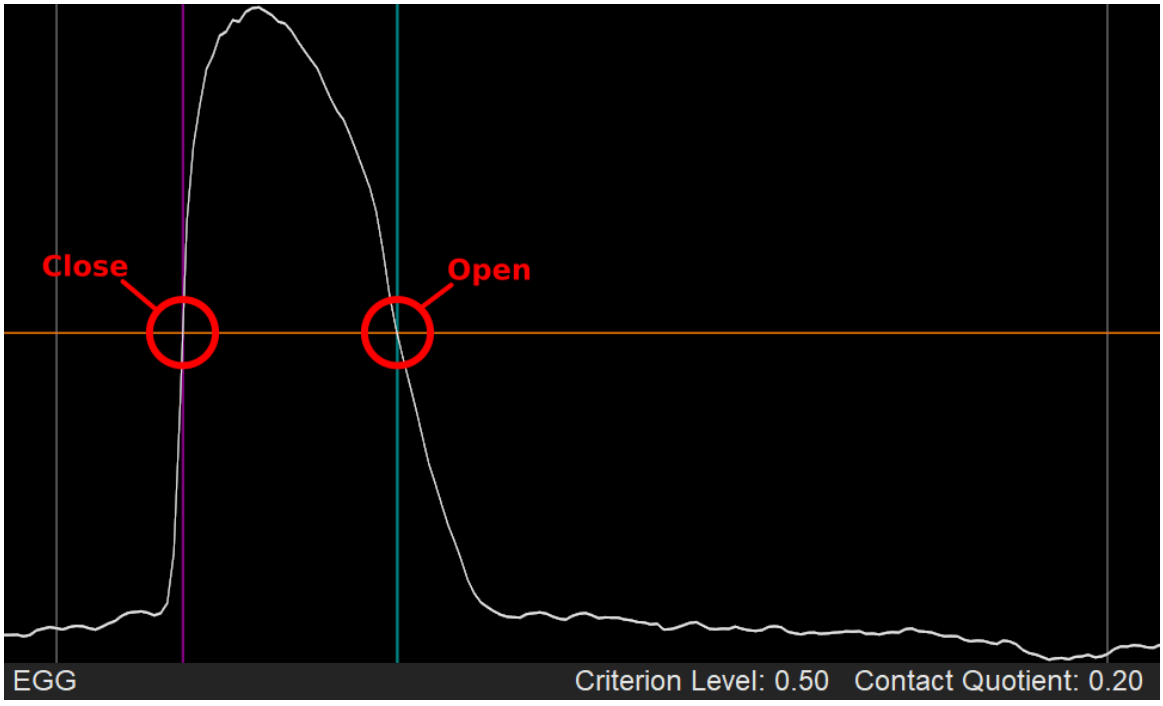

Figure 3.45. Using Criterion Level at 50% for estimating glottal contact period

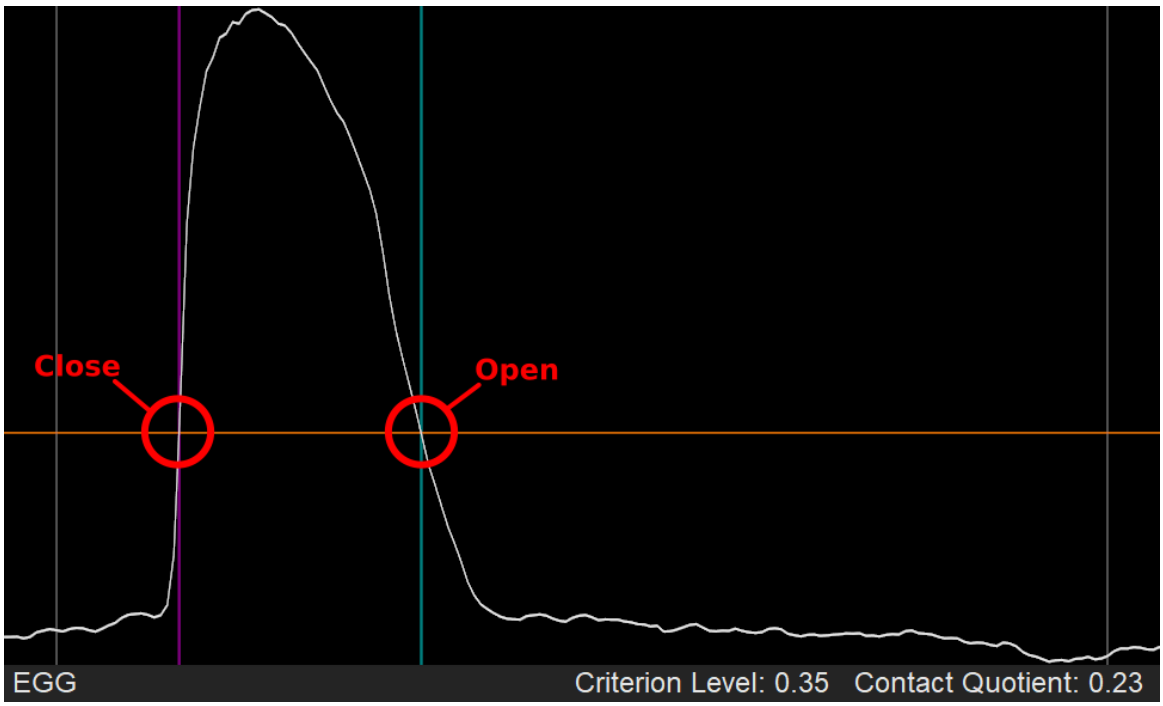

Figure 3.46. Using Criterion Level at 35% for estimating glottal contact period

Close: dEGG Max, Open: dEGG Min

This method looks at the derivative (the rate of change) of the EGG signal, the dEGG, and finds the maximum and minimum value in a period. The dEGG maximum is the point where the EGG signal rises most sharply, and the dEGG minimum is the point where the EGG has the fastest drop. In a clean signal, these points are most likely to indicate the vocal fold contact and de-contact events.

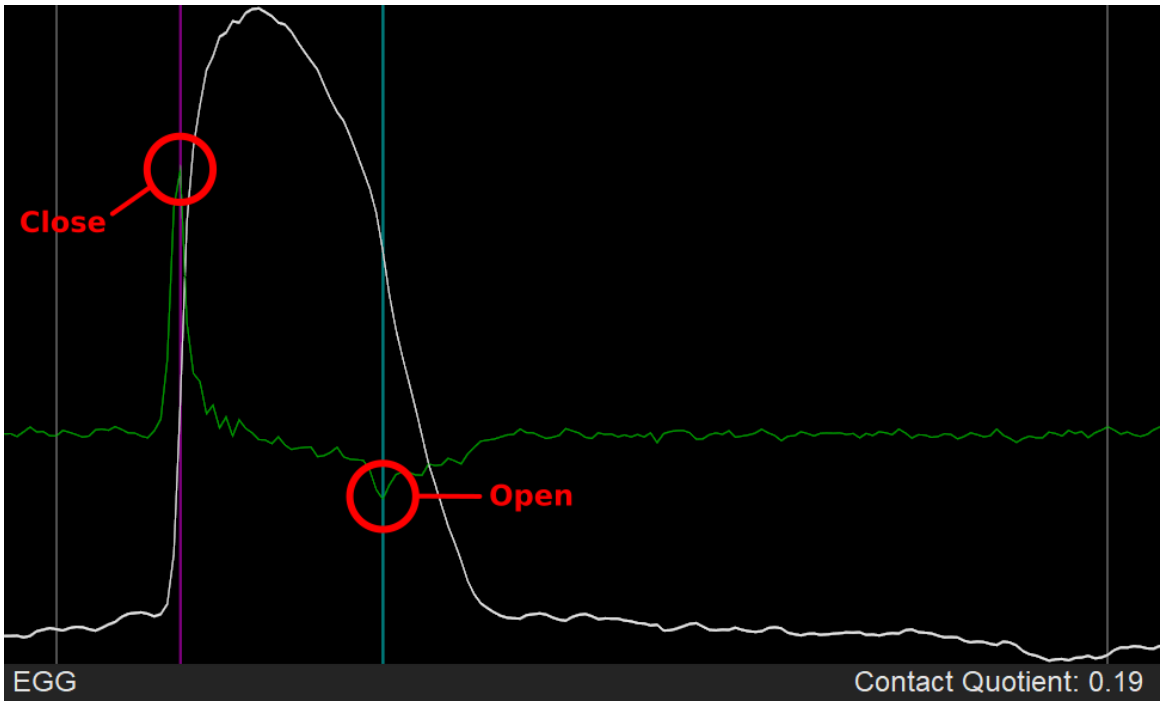

Figure 3.47. Using dEGG maximum for glottal closure, and dEGG minimum for opening

Close: dEGG Max, Open: Criterion Level

This method is a mix of the other two methods. Some signals have a good dEGG Maximum, while the dEGG Minimum is buried in noise. In that case it might work well to use the dEGG Maximum to mark the vocal fold

contact event, but to use the point where the signal drops below the criterion level to mark the de-contacting event.

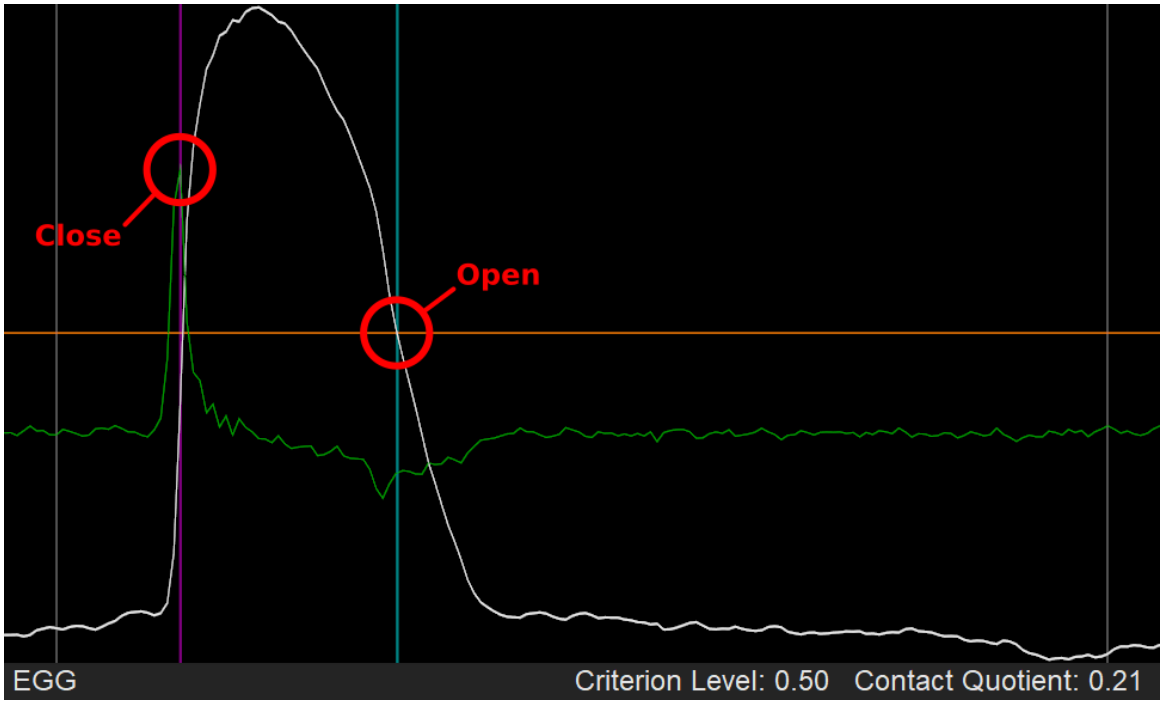

Figure 3.48. Using dEGG max for glottal closure, and Criterion Level at 50% for estimating glottal opening

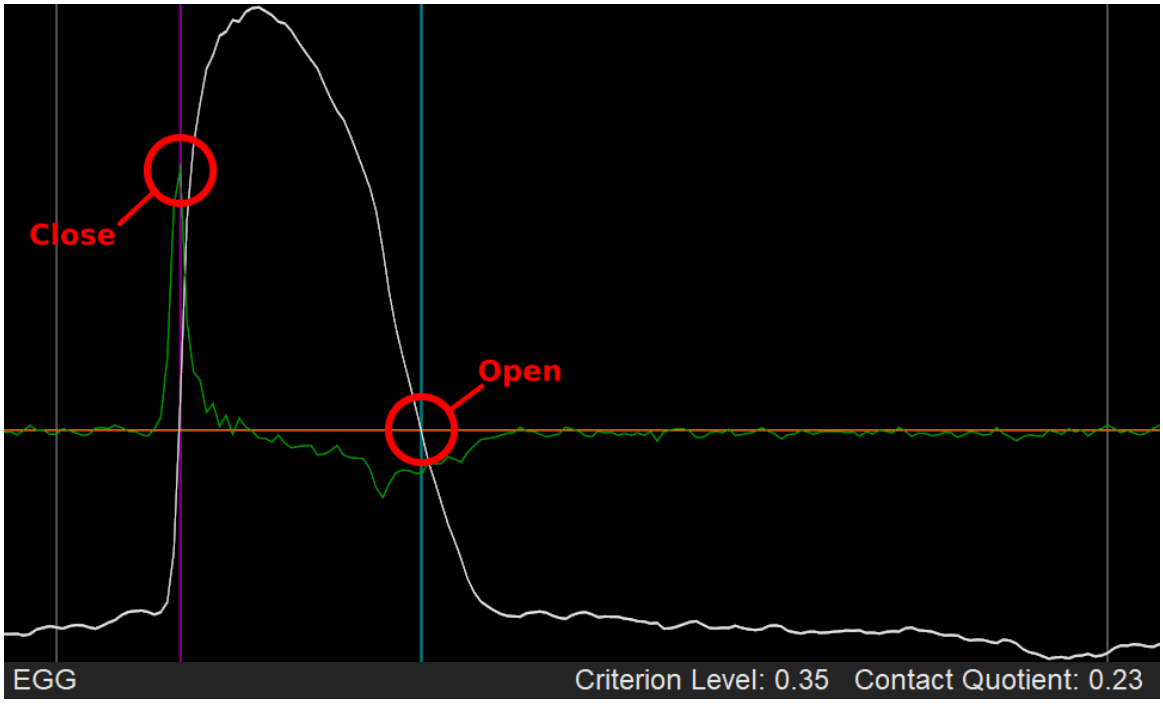

Figure 3.49. Using dEGG max for glottal closure, and Criterion Level at 35% for estimating glottal opening

3.2.6. Colormap Editor

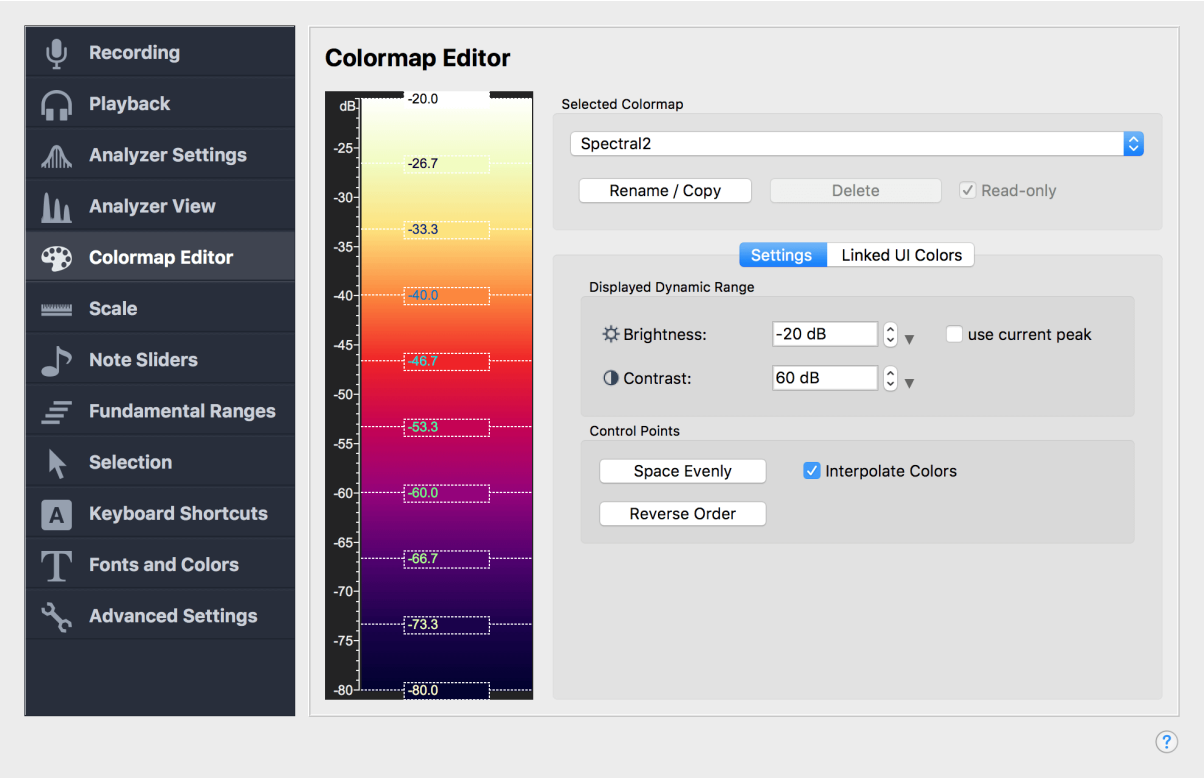

Figure 3.50. Colormap Editor

On this settings page you can choose among various colormaps to represent the intensities on the spectrum and spectrogram. You can also modify existing colormaps or create your own.

Control Point Editor

The colored field on the left shows which colors represent a given intensity level. Colors are interpolated between the existing control points. You can move and change the control points with the mouse, or create new control points by double-clicking into empty space:

Mouse action	Command
Click + Drag control point	Move control point
Shift + Click + Drag control point	Move control point with more precision
Double-click control point	Set color
Right-click control point	Bring up context menu to set color or delete point
Double-click empty space	Create new control point

Table 3.5. Colormap Editor Mouse Commands

Note: The lowest control point will always be set to the background color of the Analyzer View.

Space evenly

This button will arrange all control points such that they all have the same distance to each other. This can be used after you have changed the number of control points and want them to be evenly spread.

Interpolate colors

If this is checked, the colors between the control points will be smoothly interpolated. If the box is not checked, the color scheme will use hard color boundaries between the different control points (also look at the different but related setting Smooth Spectrogram Colors).

Here is an example of a spectrum with and without colormap interpolation:

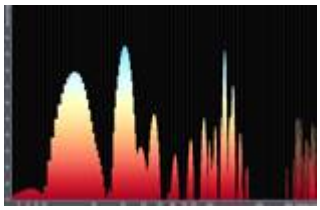

Figure 3.51. Spectrum with interpolated colormap

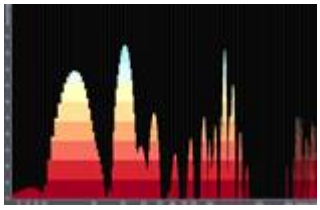

Figure 3.52. Spectrum with discrete colormap

Reverse order

Clicking this button will reverse the order of the control points, so that the colors for the highest intensity are now at the bottom, and vice versa. This will invert the appearance of the spectrogram and can be useful when trying out new color schemes.

Displayed Dynamic Range

This is the range of intensities represented by the colormap. The intensities are measured in decibels (dB). A value of 0dB corresponds to the loudest possible sine wave that can be represented by the selected audio format. With a recording bit depth of 16 bits per sample, it is theoretically possible to represent a dynamic range of 96dB. However, on the spectrogram it is often better to limit the shown dynamic range to reduce the amount of displayed noise. The dynamic range sliders can also be found on the toolbar and are useful for adjusting the displayed dynamic range on the fly to get the best contrast on the spectrogram. The best way to understand those sliders is to play with both of them.

Dynamic Range Top Level (Brightness)

This slider adjusts the top level (or saturation value) of the displayed dynamic range. In the displayed spectrogram or spectrum, all values with an intensity above the value of this slider will be displayed as maximum intensity.

Displayed Dynamic Range (Contrast)

This determines the width or spread of the displayed dynamic range. A large value will show more detail at the expense of clarity, and a small value will make the stronger intensity values stand out more.

Use current peak

If this option is enabled, the colormap top will be set to the loudest sample of the currently loaded audio file. Disable this option to set the top value manually.

Linked UI Colors

This tab will bring up a window similar to the Fonts and Colors editor. However, all changes made here are local to the current colormap. For example, the background color of the spektrum is normally determined by the global color scheme. With the 'Linked UI Colors' feature, you can change the background locally, for example to white. Similarly, you can adjust other UI element colors to match your preferred colormap.

3.2.7. Scale Settings

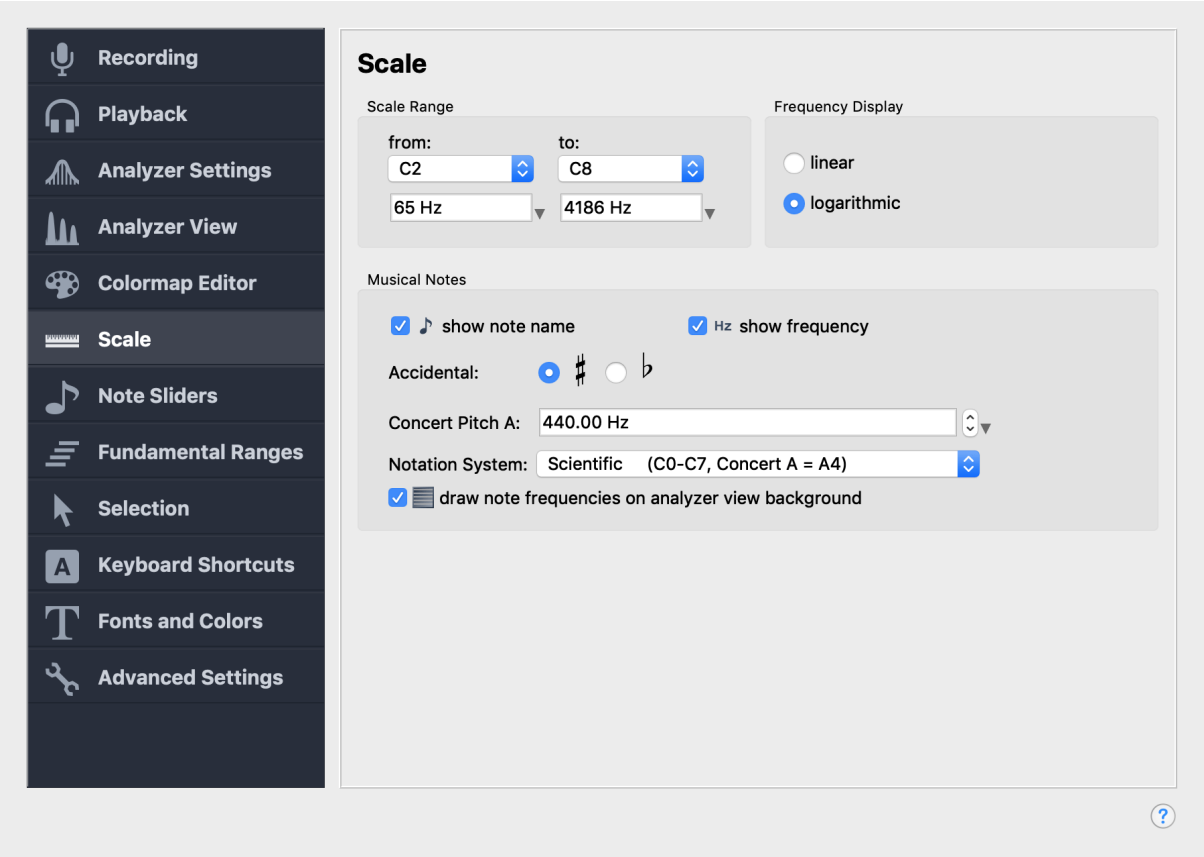

Figure 3.53. Scale Settings

On this settings page you can select range, type and orientation of the frequency scale. Most of those settings can also be found on the toolbar.

Frequency Scale

Scale Range

This is the lowest and highest note or the lowest and highest frequency that are displayed the piano and Analyzer Views. The scale range can also be changed by clicking and dragging the frequency scale in the main view, or by using the mouse wheel in that view.

Frequency Display

This determines if the frequency scale is linear or logarithmic.

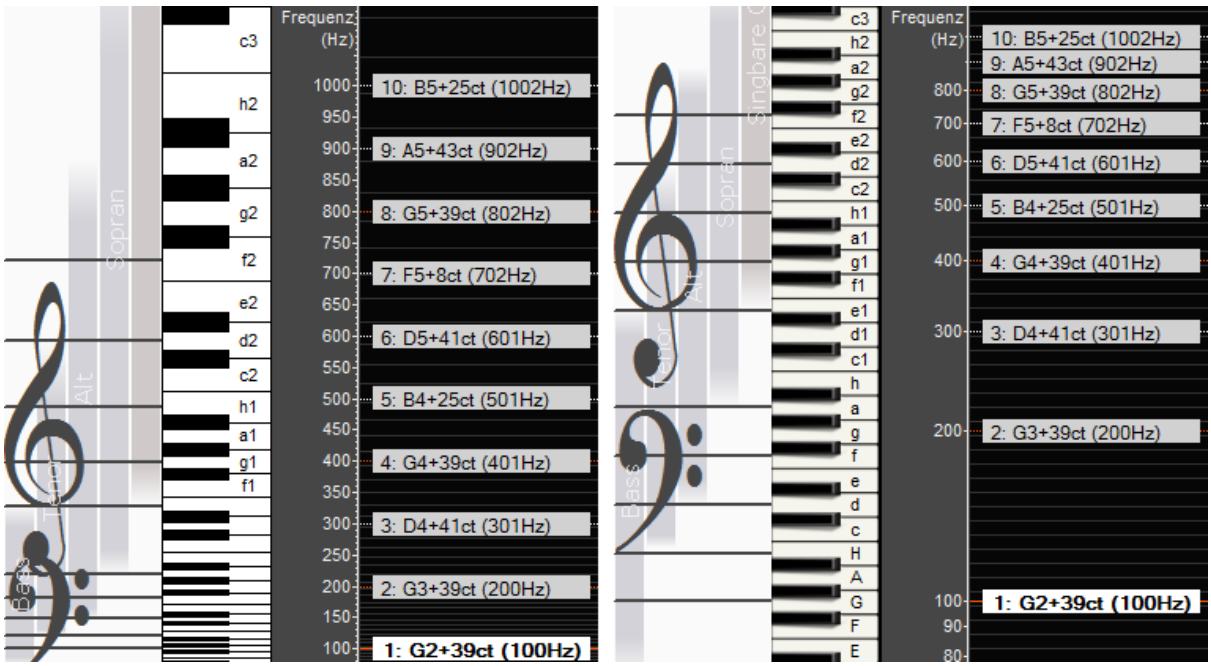

Figure 3.54. Linear and Logarithmic Frequency Scale

Figure 3.54 illustrates the difference between the linear (left side) and logarithmic (right side) frequency scale. Both sides show the same frequency range from about 80Hz to 1100Hz, which corresponds to the note range from E2 to C6.

When the scale is **linear**, consecutive harmonics have the same distance on the screen, but the piano keys get wider for higher frequencies.

With the **logarithmic** frequency scale, all piano keys have the same size, but higher harmonics are spaced more and more closely than the lower ones. The logarithmic scale represents how the ear perceives pitch. Towards the lower end of the scale, we can distinguish very small frequencies and hear them as different notes. On the upper end of the scale, each note covers a much wider frequency range. The logarithmic scale is therefore more suitable for analyzing music, as it shows much more detail in the lower note range.

Orientation

This setting simply determines whether the frequency scale is aligned horizontally or vertically.

Musical Notes

This group of settings controls how note names are shown on the piano keyboard and on the overtone sliders.

show note name

This will show the note name on note sliders on the pitch display.

show frequency

This will show the frequency in Hz on the labes of note sliders and on the pitch display.

Accidental

This determines if the notes on the black keys of the piano are displayed with the sharp or flat sign. You can also right-click into the Staff View to switch the accidental.

Concert Pitch A

Here you can set the frequency of the concert pitch that is used as a reference for all other tones. By default, the concert pitch A is 440Hz.

Notation System

You can choose between three different notation systems:

Helmholtz Notation
: This system is used mainly in German speaking countries. The notes on the piano range from A2 to C, and then from c-c5. Notice the use of both upper and lower case.

Scientific Notation
: This notation is more common in English speaking countries. The notes on the piano range from A0 to C8.

Piano Keyboard Notation
: In this system, the keys are numbered from 1 to 88 as they would appear on a real piano, counting white and black keys. This notation is mainly used by piano tuners.

Draw note frequencies in Analyzer View background

If this option is checked, thin lines will be drawn in the Analyzer View background that represent the frequencies of notes on the piano.

3.2.8. Note Slider Settings

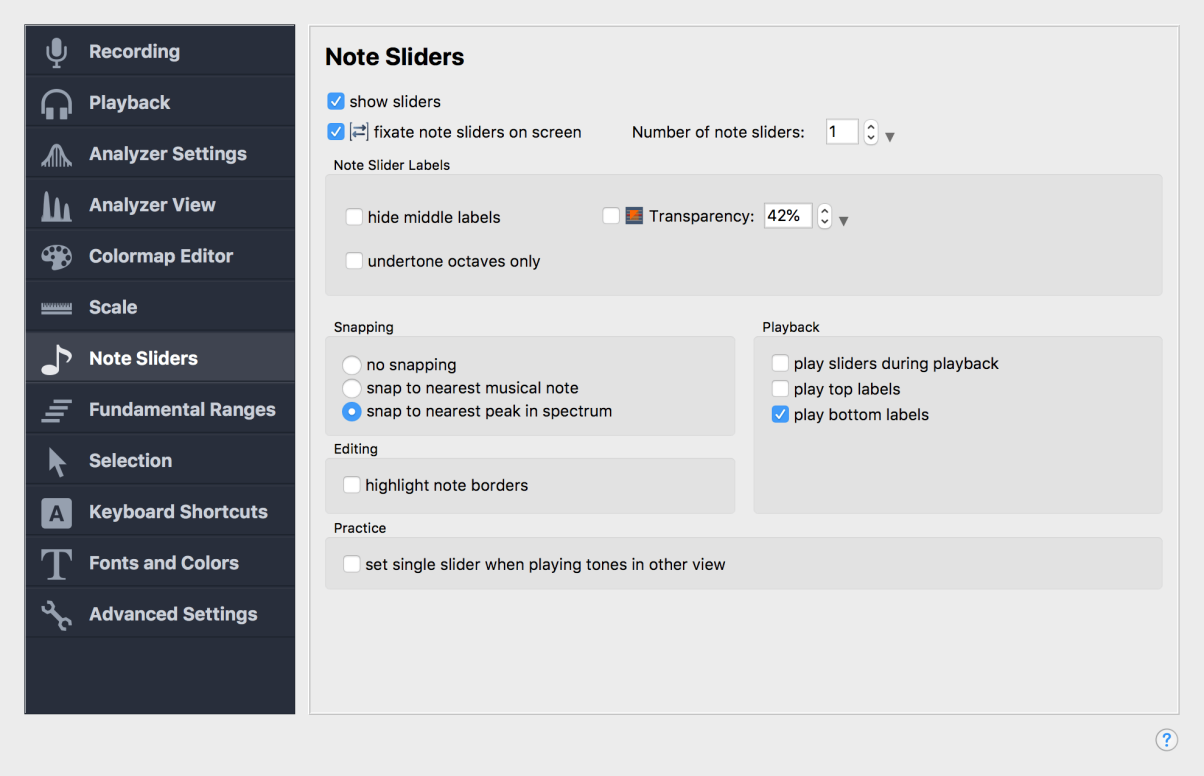

Figure 3.55. Note Slider Settings

On this settings page you can adjust the appearance and behavior of the Overtone Sliders (which are also called Note Sliders). Most of those settings can also be found on the toolbar.

Fixate note sliders on screen

If this setting is active, the note sliders are fixated on the screen and will stay in the same position regardless of how the time range of the Analyzer View is changed. In VoceVista Video Free Edition the sliders are always fixated on the screen. When the sliders are not fixated on the screen, each slider belongs to a specific time range of the recording, like the notes in a piece of music. This allows to represent an entire recording as a melody comprised of musical notes, which can also be exported as a MIDI file.

Number of Overtone Sliders

This is the number of distinct sliders shown. Each slider can have a separate fundamental frequency. The number of sliders can only be changed through this control if the sliders are fixated on the screen. Otherwise, sliders can be added and deleted through the commands Tools → Insert note slider at selection and Tools → Delete note selected sliders.

Note Slider Labels

Show note name

Show the name of a slider's note, such as “A3+12ct”. The second part of the name shows the difference in cent between the exact pitch of the slider, and the named note on the piano.

Show frequency

Show the pitch of a slider's frequency, such as “220 Hz”.

Hide middle labels

If this option is activated, the note name and frequency will only be shown for the fundamental, and for the highest displayed overtone, and the lowest displayed undertone of each slider.

Transparency

This option makes the sliders transparent so that the spectrogram underneath can still be seen. This can make it easier to position sliders accurately. The amount of transparency requires some experimentation for each situation. It can be adjusted here, but the transparency can also be adjusted through a slider on the status bar.

Undertone octaves only

This option will hide all undertones that are not an octave of the fundamental. This can be useful when constructing scales, where tones are used across many octaves.

Play top labels

Checking this option will cause the highest visible overtone of each slider to be played together with the fundamental, when the fundamental is played by clicking on it. This can be used to play overtone melodies together with the fundamental.

Snapping

When moving a slider along the frequency axis while this option is enabled, the slider will snap to the nearest piano note or spectral peak. This behavior can be reversed when the ALT key is pressed while moving a slider (for more keyboard modifiers, including the possibility to snap to other sliders, refer to Table 3.4, “Overtone Slider Mouse Commands”).

no snapping
: The sliders can be moved freely.

snap to nearest musical note
: The slider will snap to the frequency that represents the nearest note on the piano.

snap to nearest peak in spectrum
: The slider will snap to the nearest maximum in the spectrum. This allows to accurately measure the loudest frequencies in a recording.

Editing

highlight note borders

Checking this option will draw the note borders with high contrast, which can be useful when editing note sliders.

3.2.9. Fundamental Ranges

The Customize Fundamental Ranges Dialog lets you change the background of the Staff View, where the different ranges of normal singers and of overtone singers can be displayed.

Voice Range Profiles

You choose from several predefined profiles through the list box on the right:

SATB

Figure 3.56. SATB Voice Range Profile

This profile shows the normal four choir voice ranges Soprano, Alto, Tenor, and Bass.

Standard Voice Range Profile

Figure 3.57. The standard voice range profile

This profile shows the four choir voice ranges, and the range of overtones that are singable in a controlled manner (by consciously amplifying specific overtones).

Overtones NG / L

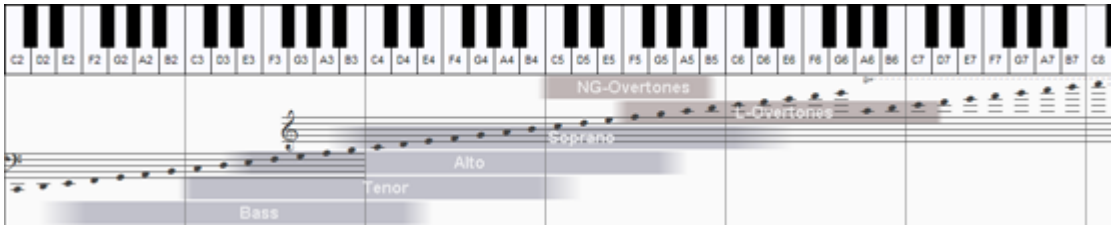

Figure 3.58. The Overtone Technique Profile

In this profile the singable overtones are divided by the technique used to produce them. The NG-Technique uses the entire oral cavity without the tongue and is more suitable for the lower overtones. The L Technique uses the tongue to divide the oral cavity into two separate chambers. This technique is used to produce the higher overtones.

Changing the frequency ranges

The predefined frequency ranges are only approximations, and there are different opinions about what constitutes the range of singable overtones, and the choir voices. You can select the ranges that you want to display by changing the start and end frequencies for each voice.

Display Singing Voice Ranges

This switch can be used to completely disable the display of the ranges when they are not needed, so that the Staff View will only show the music notes.

Inner and Outer values

The Outer value specify the point where the voice starts to fade in from the background color, and the Inner value is the point where it will reach the full color.

Color Type (Overtone or Voice)

This selects the color that will be used for this range. There are two different colors available: Overtone and Voice. Both can be customized on the Fonts and Colors settings page under Colors / Staff View / Range Overtone and Range Voice.

3.2.10. Selection and Pointer Settings

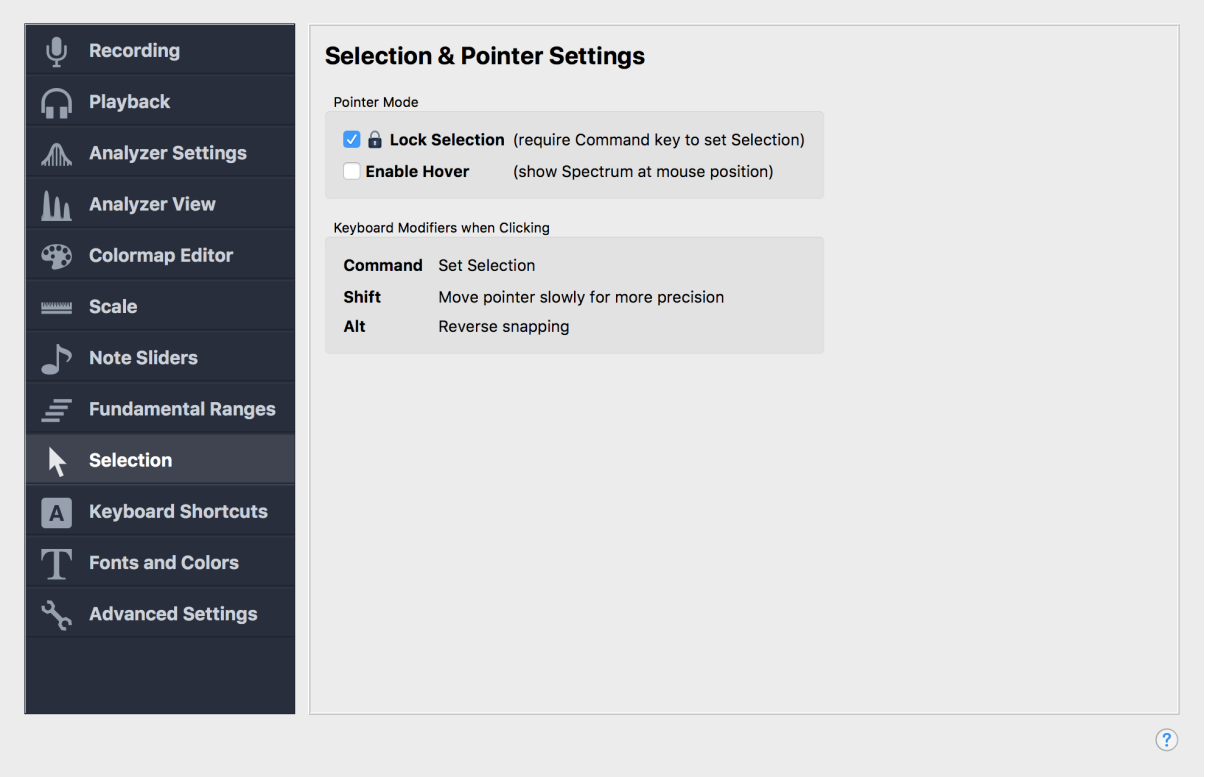

Figure 3.59. Selection and Pointer Settings

On this settings page you can choose the behaviour of the mouse pointer, and how it affects the cursor position and the Selection.

Pointer Mode

Lock Selection

Selection Locked: Ctrl/Cmd-Click sets Selection
: If the selection is locked, it cannot be changed or unset with a single mouse-click. To set the selection, press the Control key (on Windows) or the Command key (on Mac), and then click and drag on the Spectrogram or on a Time Scale.

Selection Unlocked: Click sets Selection
: With the selection unlocked, a simple mouse-click will reset it. To set the selection, click and drag on the Spectrogram or a Time Scale.

Enable Hover

By default, moving the cursor across the Spectrogram or a Time Scale will have no effect.

If the *Hover mode* is enabled, moving the mouse pointer across the Spectrogram, the Waveform or the Timeline will immediately set the Spectrum in the Short-Term View to that position of the recording.

Keyboard Modifiers when Clicking

This table lists the modifier keys that will affect the pointer behaviour when clicking.

Command / Control
: Holding the Command key (on Mac) or the Control key (on Windows) while clicking will always set the Selection, regardless whether the Selection is locked or not.

  Holding the Cmd/Ctrl key and then clicking on sliders or piano keys will select multiple elements.

Shift
: Holding the Shift key while clicking and dragging the pointer will make the pointer move much more slowly. This allows for greater precision when moving an element, making a selection, or scrolling a scale.

Alt
: Holding the Alt key while moving a Note Sliders will invert its snapping behaviour. If it is set to snap to musical notes or spectral peaks, it won't snap while holding the Alt key, and vice versa.

3.2.11. Keyboard Shortcuts

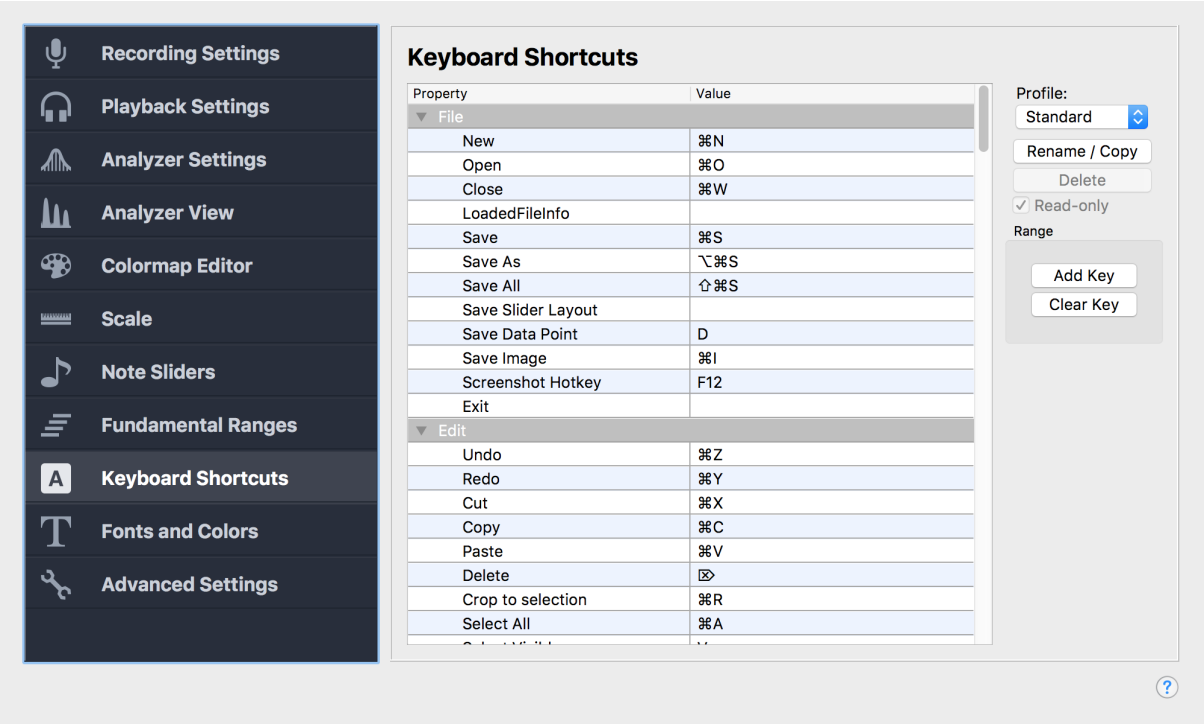

Figure 3.60. Keyboard shortcuts

On this settings page you can view and modify the keyboard shortcuts for all the different commands. To modify a shortcut, simply click on it and press the new key combination that you want to use for the selected command. Duplicate assignments will automatically be removed.

The Profile bar on the right allows you to save and retrieve different keyboard layouts. In addition, there are buttons to add and remove shortcuts to commands:

Add Key
: When a command is selected, pressing the Add Key button will add a new shortcut to this key, so that a single command can have multiple shortcuts.

Clear Key
: Clear the key for the selected command.

Here is an overview of the default shortcut assignments:

Command	**Shortcut**
File	
New	Ctrl + N
Open	Ctrl + O
Close	Ctrl + W
LoadedFileInfo	
Save	Ctrl + S
Save As	Ctrl + Alt + S
Save All	Ctrl + Shift + S
Save Slider Layout	
Save Data Point	D
Save Image	Ctrl + I

Command	Shortcut
Screenshot Hotkey	F12
Send Screenshot Email	Shift + F12
Exit	
Edit	
Undo	Ctrl + Z
Redo	Ctrl + Y
Cut	Ctrl + X
Copy	Ctrl + C
Paste	Ctrl + V
Delete	Num Del
Crop to selection	Ctrl + R
Select All	Ctrl + A
Select Visible	V
Set selection	S
Reset selection	Esc
Lock selection	L
Normalize	
FadeIn	
FadeOut	
Replace with silence	
Insert silence	
Reinsert silence	
View	
Zoom in current view	1
Zoom out current view	2
Zoom in	Shift + 1
Zoom out	Shift + 2
Zoom to fit selection	3
Zoom out fully	4
Center on reticule	5
Timeline	Ctrl + F2
Waveform	Ctrl + F3
Staff View	Ctrl + F4
Piano View	Ctrl + F5
Status Bar	Ctrl + F6
FileMarkerList	Ctrl + F7
VowelChart	Ctrl + F8
Waveform Time Scale	
Frequency Scale	
Spectrum Intensity Scale	Ctrl + F9

Command	Shortcut
Spectrogram Time Scale	Ctrl + F10
Reset Window Layout	
Set Language	
Toggle Language	Ctrl + Shift + L
Full Screen	Alt + Enter
Filters	
Add Filter	F
Enable Realtime Preview	
Apply Filters	
Filter Additively	
Filter Subtractively	
Markers	
Add and edit Marker	Ctrl + M
Add Marker	M
Add Range Marker	
Add Thumb up Marker	G
Add Thumb down Marker	B
Delete Marker	Ctrl + Num Del
Go To Previous Marker	Q
Go To Next Marker	W
Zoom To Previous Marker	Shift + Q
Zoom To Next Marker	Shift + W
Merge Markers	
Search Markers	Ctrl + F
Edit Current Marker	
Zoom to Current Marker	E
Activate Marker 1	
Activate Marker 2	
Activate Marker 3	
Activate Marker 4	
Activate Marker 5	
Activate Marker 6	
Activate Marker 7	
Activate Marker 8	
Activate Marker 9	
Activate Marker 10	
Tools	
Rewind to start	Home
Rewind page	Page Up
Rewind	Ctrl + Left

Command	Shortcut
Play and Stop	Space
Stop	
Record	Ctrl + Space
Record	R
Forward	Ctrl + Right
Forward page	Page Down
Forward to end	End
Replay last take	
Loop playback	
Play Selected Tones	Enter
Stop Playing Tones	Backspace
Deselect all Tones	Ctrl + Backspace
ProfileManager	
Options	F10
Volume	
Increase input volume	
Decrease input volume	
Increase output volume	
Decrease output volume	
Window	
NewWindowOnCurrentSelection	Ctrl + Shift + N
Tile horizontally	Ctrl + Shift + H
Tile vertically	Ctrl + Shift + V
Maximize current window	Ctrl + Shift + M
Close all	Ctrl + Shift + W
Close all but active	Ctrl + 0
Note Sliders	
Insert Note Slider	N
Use Note Transcription Tool	T
Fix Sliders on Screen	
Display	
Toggle Spectrum	
Toggle Spectrogram	
Toggle Pitch	
Disable Analyzer Display	7
Show Spectrogram	8
Show Spectrum	9
Show Spectrogram and Spectrum	0
Linear Scale	Ctrl + 1
Log Scale	Ctrl + 2

Command	Shortcut
Vertical Scale	Ctrl + 3
Horizontal Scale	Ctrl + 4
Show Notename	Ctrl + 5
Show Frequency	Ctrl + 6
Use sharp accidental	A
Transparent Sliders	Ctrl + 7
Snap to tempered	Ctrl + 8
Draw tempered lines	Ctrl + 9
Option Pages	
Audio Settings	F2
Device Options	F3
Analyzer Settings	F4
Analyzer View	F5
Colormap Editor	F6
Scale	F7
Sliders	F8
File Export	
Save As Wav	
Save As Flac	
Save As Ogg	
Save As Mp3	
Send Email As Ogg	
Send Email As Mp3	
Save Copy As Ogg	
Save Copy As Mp3	
Profiles	
1 Minimal Toolbar	
2 One-line Toolbar	
3 Default Toolbar	
4 Full Toolbar	
Quickstart	
Standard Frequency Range	
Standard Time Zoom	

Table 3.6. Keyboard Shortcuts

3.2.12. Fonts and Colors

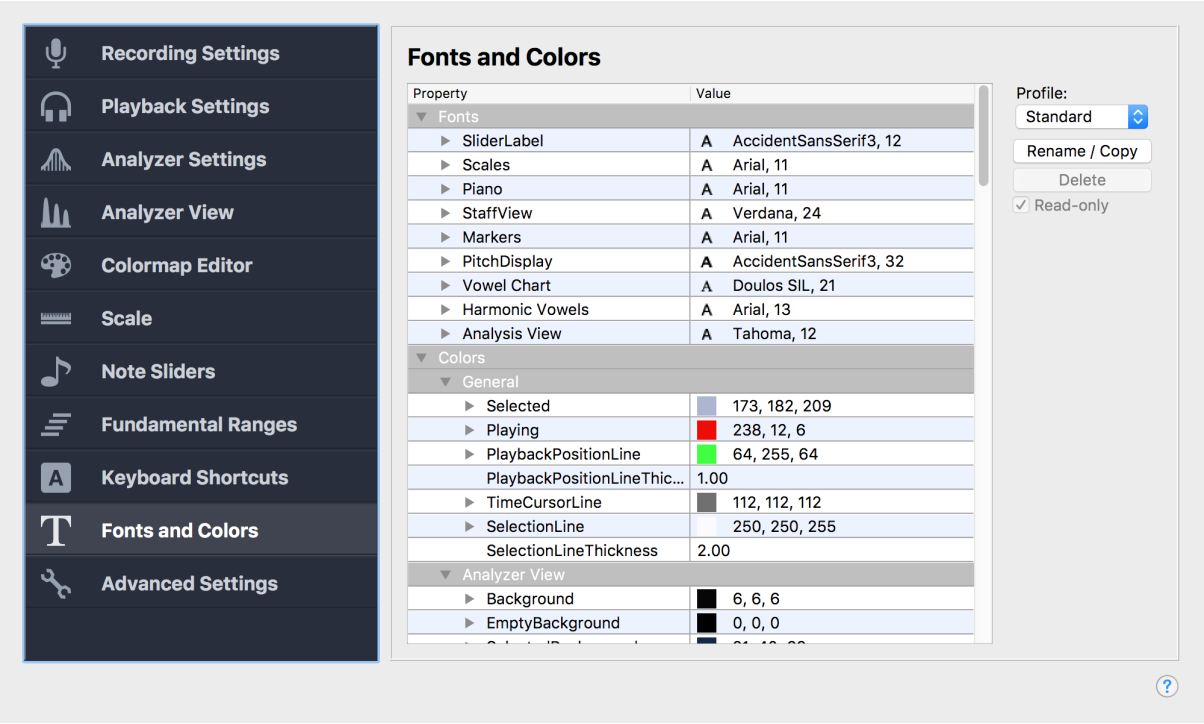

Figure 3.61. Fonts and Colors settings

On this page you can customize the fonts, colors, and line thicknesses for many screen elements. For example, you can adjust the colors of the Analyzer View background and of the overtone sliders to achieve the best contrast with specific spectrum colormaps.

You can manage several different color schemes through the profile bar on the right.

Note: On this page you define the global fonts and colors profile. If you want to, you can associated colors of specific UI elements with the colormaps of the spectrum by going to the Linked UI Colors section of the colormap editor. For example, there is a "White background" colormap that defines the spectrum background to be white, thereby overriding whatever background color is defined by the current color profile.

3.2.13. Advanced Settings

This page has various settings that are intended for advanced users to further customize the program:

System

Reopen last session at startup

If this is checked, VoceVista Video will remember which files were open when it was last closed, and open those files again on the next startup.

Number of uncompressed files to cache

When VoceVista Video loads compressed files such as Mp3 or Ogg files, it first converts them into an uncompressed file that will be stored in your temp folder. This allows much faster access for scrolling and editing the file, but it might use up a lot of hard drive space. This setting determines how many uncompressed files will be stored in the temp folder. If you open a file that has already been uncompressed and is still in the cache, VoceVista Video will skip the uncompression, which will be much faster.

Automatically check for updates

When this setting is checked, VoceVista Video will periodically check if a new version is available, which will require an internet connection.

Logging Level

This specifies the amount of detail that is written to the log file.

Language

Set the language in which the user interface of VoceVista Video is displayed. Changing this requires a restart of the app for all changes to take effect.

Display

Waveform

The Waveform settings control the appearance of the Waveform and the Timeline Views. They are also available by right-clicking on either of those views.

Waveform uses Decibel

If checked, the Waveform and the Timeline will display intensity on a logarithmic scale.

Use symmetric waveform

If checked, the Waveform and Timeline will display intensity symmetrically.

Scales

Use vertical scale labels

Determines if the text on all scales along the Y-Axis will be vertical or horizontal. Vertical text uses less screen space, but horizontal text is easier to read. You can also access this setting by right-clicking on the actual scale.

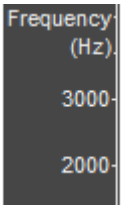	
Y-Scale with horizontal text	Scale with vertical text

Show only scale units

By default, the scales will display the name of the scale and the unit, such as “Time (s)” and “Frequency (Hz)”.

If this option is checked, the scales will only display their unit, such as “s” and “Hz”.

Markers

This settings control the visibility of Markers on different windows. This allows to further refine the Marker visibility in addition to setting the visibility of Markers by type.

Draw marker heads on timeline

If checked, the labels of Markers will be shown on the Timeline. If this option is not checked, only the boundaries of the Markers will be shown.

Show markers above the spectrogram

If checked, Markers will be drawn above the Spectrogram (the visibility for each marker type must also be enabled in the Markers menu).

Initial Time Range (in seconds)

This is the time range that will be set for new empty documents. By default it is 10 seconds.

Fixed Timeline Range (in minutes)

By default, the Timeline range starts with a range of 10 seconds for new recordings and expands as a recording grows longer. If the Timeline range is fixed by setting this option, it will stay the same length at all times.

Initial number of Overtone Sliders

Sets the number of visible overtone sliders for new empty documents, or for existing documents that are loaded for the first time.

Show Timeline on top

Shows the Timeline above the Spectrogramm instead of below it.

Spectrogram Scroll Mode

The Scroll Mode determines what happens when the time cursor reaches the end of the visible range during recording and playback.

Scrolling
: Keeps the time cursor at fixed position and scrolls the window to the left.

Paging
: Switches to a new empty page and then moves the time cursor across the page.

Decibel distance for background lines on spectrum

This setting allows background lines to be drawn on the Spectrum that mark different intensity levels. By default, these lines are draw 10 decibels apart from each other.

Decibel range of normalized spectrum

The dynamic range of intensities that is shown on the Spectrum when it is normalized.

Toolbar Icon Size

Allows to set the size of the Toolbar Icons to make them larger or smaller.

Recording

Insert silence after each take

This allows to specify an amount of silence that will be added to a file each time recording is stopped.

Automatically start recording when program starts

If this option is checked, the program will automatically start recording into a new empty file after it is started. This can also be achieved by starting the program with the `/record` command line parameter.

Clear recording after this many minutes

Setting this option can allow the program to run continuously without filling up the harddrive. For example, if this value is set to "10", the program will record and visualize sound for 10 minutes and then delete the current file and start over.

Signal Generator

The Signal Generator can be used to play back a test signal over a separate sound card. This feature is specifically intended to measure the resonance of the oral cavity by playing a sound into it, whose echo can then by analyzed.

Play signal while recording

If this option is checked, the test signal will always be played during recording.

Wave Output Device

If you have multiple sound cards, you can use a separate output device to play the test signal, while the recorded sound can be played back with the main speakers.

Frequency (Hz)

The frequency in Hz of the test signal. Right now the only available signal is a square wave at maximum amplitude.

Signal Phase Ratio

This ratio determines the width of the up and down phases of the square wave signal. A value of 0.5 creates a symmetric signal. A smaller value will create a shorter up spike followed by a longer down period.

Cursor Lines

When moving the mouse across the Analyzer View, helper lines appear to make it easier to read the current cursor position on all scales. These settings allow to individually control the visibility of those lines.

Show Playback Cursor Line

Show the green playback cursor line. Turning this off is mostly useful for taking screenshots where the playback cursor would be in the way.

Show Time Cursor Line

Show the time of the cursor location across all views that have a time scale.

Show Frequency Cursor Line

Show the frequency of the cursor location across all views.

Show Intensity Cursor Line

Show the intensity line for the cursor location on the Spectrum.

Show time range on crosshair

Mark the time ranges on the cursor that are used to calculate the spectrum and the pitch.

Vowel Chart

These settings control the frequency ranges and scale mode of the Vowel Chart.

Audio Performance

Enable WASAPI Audio API

This setting applies to Windows only, not to Mac. On Windows, the default Audio API to be used is DirectSound. WASAPI is an alternative, more low level API that may offer slightly better performance. It also allows to record the output of playback devices via Loopback Recording. For example, if WASAPI is enabled, you can record "What you hear", in other words, what is currently being played, if you select your current output device as recording source. The drawback of WASAPI is that, unlike DirectSound, it has no automatic sample rate conversion. For example, if the audio device is set to a sampling rate of 48kHz, you can only record from it with that sampling rate.

Recording Latency

This is the approximate time in milliseconds that passes from the moment a signal arrives in the microphone until it is displayed on the screen. A lower latency means that the screen is updated more frequently. Therefore this value should be as low as possible to ensure that the analyzer display is responsive. However, if the value is too low, your computer may be unable to process the incoming data quickly enough, and some parts of the recording my be lost. This will appear as clicks or gaps in the recording. If you have a fast computer, you can try to lower this value and see if the recording still appears correct.

File Playback Latency

This setting is similar to the Recording Latency, but for playback. A lower number will mean there is less delay between performing an action such as changing the playback position and hearing the result, but if the value is too low you may hear clicks or gaps during playback. If that is the case, raise the value until playback sounds smooth and continuous.

Tone Generator

These settings affect only the Sine Wave Generator that can be used to play overtones and piano keys. You can change the sampling rate and the playback latency. A lower latency value will make the program appear more responsive to changes of the pitch of a slider when you move it.

Input Level Meter

The Input Level Meter is the slider on the toolbar that controls input volume and shows the strength of the current input signal. Here you can change its parameters.

Enable monitoring

When the Input Level Meter is on, VoceVista Video will always record sound from your sound card, which could be undesirable in some cases, so you can turn it off here. You can also right-click on the Input Level Meter on the toolbar to enable or disable it.

Displayed range

This is the dynamic range shown on the Input Level Meter.

Max updates per second

This number determines how often the Input Level Meter is redrawn. You could lower this value if your computer is very slow, because every redraw uses up a small amount of CPU time.

Peak History

This is the amount of time (in seconds) that the peak display is using. For example, if the peak history is 3 seconds, the Level Meter will show the maximum value from the last 3 seconds.

3.3. Loading and Saving

This section explains how to load and save audio recordings and other file types such as images.

3.3.1. Loading and saving Audio Files

Loading Audio Files

If VoceVista Video is already running, click on File → Open to open an existing file. You can also drag files into the VoceVista Video window. In Windows Explorer, you can right-click on a file and select "Open with" and then chose VoceVista Video. You may have to click on "Choose program..." if VoceVista Video is not already listed.

Saving Audio Files

Click on File → Save to save a recording. On the File Save Dialog, you can select the file type on the bottom.

There are several keyboard shortcuts for saving files in a specific format.

Supported File Types

VoceVista Video can open most audio file formats. It can also open some video formats and load their audio stream.

For saving, VoceVista Video supports .wav, .ogg, .mp3, and .flac. Wav files are uncompressed. Those files require the most hard drive space, but they are also the fasted to open and save.

Lossless and lossy compression

The other three formats are compressed, so they all produce smaller files than .wav. For compressed formats, one distinguishes between lossy and lossless compression. Flac is a lossless format, which means that when a recording is saved as a .flac file and then loaded again, the reloaded data will be exactly what it was before. This perfect quality comes at a price, however, because the compression is usually not more than 50% compared to an uncompressed wav file.

Saving as Mp3 or Ogg

The formats .mp3 and .ogg use lossy compression. When a recording is saved in these formats and then loaded again, the restored data will be different from the original and depending on the strength of the compression, more or less detail will be lost. However, this enables much higher amounts of compression than the lossless formats. Files saved as .mp3 or .ogg are often 90% smaller than a .wav file with the same data. You can specify the amount of compression if you have selected one of those file types and press the Options button on the Save File Dialog.

Mp3 is the most popular format for compressed audio files. Ogg is a newer format that has been designed to replace mp3. It claims to have a better sound quality for files of the same size.

While you are editing a sound file, it is recommended to save it in the .wav format to avoid losing quality from compressing and uncompressing the file many times while it is opened and saved again. When the file is ready for archival, save it in .flac if you need perfect quality, or save it in .ogg if you want smaller files and can accept a certain loss in quality.

If you want to save files in the .mp3 format, you have to first install the Lame mp3 encoder dll.

Installing the LAME mp3 Encoder

To create Mp3 files with VoceVista Video, the LAME Encoder (lame_enc.dll) needs to be installed on your computer. Since Mp3 is a patented technology that requires an expensive license, we cannot distribute this file ourselves and recommend the use of the Ogg file format for creating compressed audio files. Ogg is a free alternative to Mp3 that also claims to have a better sound quality than Mp3 for compressed files of similar size (http://en.wikipedia.org/wiki/Ogg).

If you want to create Mp3 files, there are a number of sites where you can freely download the LAME Encoder. For example: www.rarewares.org [http://www.rarewares.org/mp3-lame-bundle.php].

If that link doesn't work, you could also go to the official LAME website [http://lame.sourceforge.net/links.php#Binaries] and look at their links section or try a web search.

> **Note**
>
> Please be careful when downloading software from the Internet. We are not in any way affiliated with rarewares.org, or any other site that distributes the LAME encoder, we have absolutely no control over what is offered there, and we cannot guarantee the safety of files offered by them. **Use files from these sites at your own risk.**

To use the LAME Encoder, put lame_enc.dll somewhere on your computer. When you want to save a file as mp3 in VoceVista Video, browse to the location of the lame_enc.dll when VoceVista Video asks for it during the saving of the file.

3.3.2. Saving Images and Screenshots

VoceVista Video can take screenshots from any part of the user interface and save them as image files.

Click on File → Export Screenshot to open the Export Screenshot Dialog:

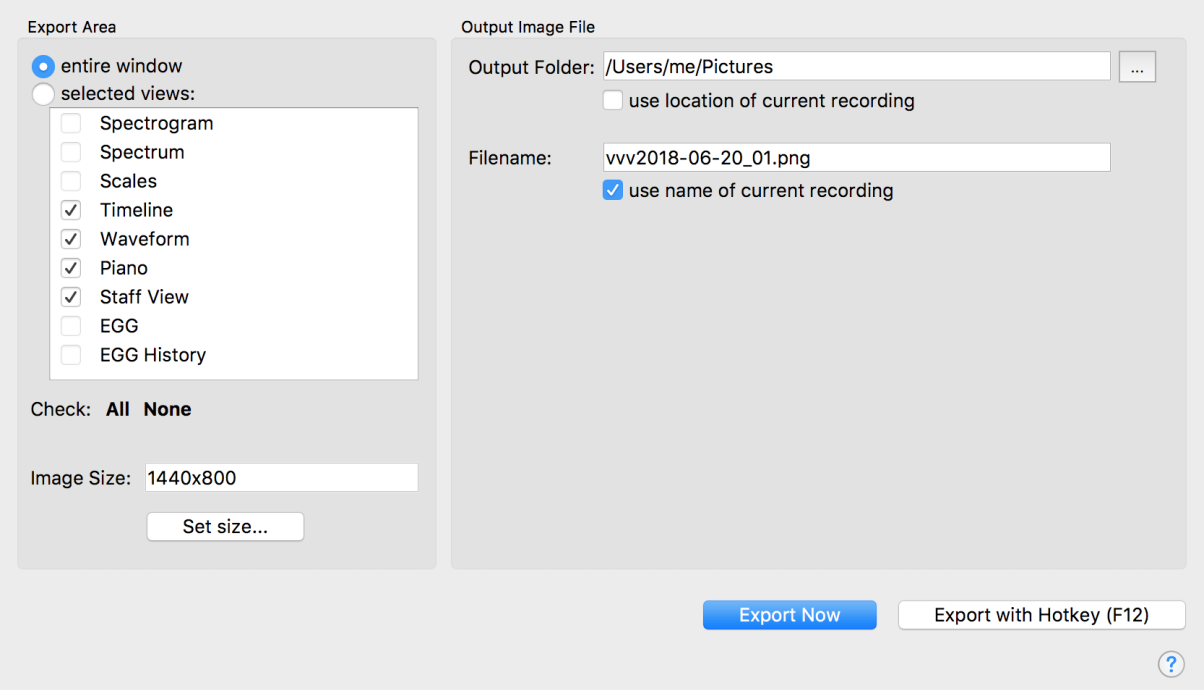

Figure 3.62. Export Screenshot Dialog

This window has two main sections. On the left side in the Export Area section you specify what you want to save, and on the right side in the Output Images section you specify where you want to save it, and under which name.

Export Area

In this section you control if the entire VoceVista Video window is saved, or if only parts of the interface should be saved. If you choose to save only selected views, check those views that should be included. VoceVista Video will export the smallest rectangle that includes all selected views, so in same cases that views may be implicitly included. You will see a yellow frame marking the selected area in the Overtone Analyzer window.

Image Size

This field shows the size in pixels of the saved screenshot. You can change the size by resizing the exported windows, by selecting different views to be exported, or by clicking on the Set Size button.

Set size

This button allows to set the size of the exported area by adjusting the size of the main window. First select the views to be exported, and then click on the Set Size button to enter the desired target size. VoceVista Video will attempt to resize the main window such that the exported views fit your requested size.

Output Images

This section determines the location and name of the exported image file.

Output Folder

Select a folder where the image will be saved. If you check the box use location of current recording, the image will be saved in the same location in which the current audio file is stored.

Filename

Enter the name of the current screenshot. If you check the box use name of current recording, the name of the image will be derived from the name of the current audio file. If you save multiple screenshots, a number will be appended to the filename so that each file has a unique name. VoceVista Video will not overwrite existing files.

send image as email

If this is checked, VoceVista Video will attempt to compose a new email in your email program with the saved screenshot as an attachment. You still have to specify the recipient and the other content of the email before you can decide to send it. Sending images as email only works if you use local email client (such as Microsoft Outlook, or Mozilla Thunderbird). It won't work with web-based email clients that you use through your browser (such as Gmail).

You can avoid having to explicitly check the send image as email box by using the Send Screenshot email keyboard shortcut (which by default is Shift + F12).

Export Now

This will save the selected output image immediately.

Export with Hotkey

This button will store the current image export settings and close the window. Now whenever you press the Save Image hotkey, a new screenshot will be saved into the output folder. By default, the hotkey is F12, but you can change that in the keyboard settings page. You can check the include mouse cursor option to include the mouse cursor in screenshots that are saved with the hotkey.

3.3.3. Loading and Saving Overtone Slider Layouts

If you have created a complex configuration of Overtone Sliders, you might want to save it so that it can be used again later. Overtone Analyzer will automatically remember the slider layouts for the last few audio recordings, but you can also explicitly save a slider layout so that you can use it for different audio files, or on their own.

Saving a slider layout

Click on Sliders → Save note slider layout and specify where you want to save the layout file. The file type ending is .oa (for **O**vertone **A**nalyzer Files).

Loading a slider layout

Click on Sliders → Load note slider layout to open an existing .oa file. This will replace your current slider layout.

Export note sliders as MIDI

Click on Sliders → Export note sliders as MIDI to save the current note slider layout into a MIDI file which can be played by many audio programs, and which can be imported as a note track into your music notation software.

At the bottom of the export dialog you can choose if you want to export just the fundamental notes, or the fundamentals and the overtones. The latter option is only relevant if you have transcribed or composed a piece with overtones, where the Overtone Sliders have a fundamental, and at least one overtone. In that case the fundamentals and overtones will be exported as two separate voices.

Specifying the default layout for new files

By default, when you create a new file, you will have a single overtone slider with a fundamental frequency of 220Hz (which is the note A3). If you simply want to change the number of displayed sliders, or turn sliders off for new documents, go to Options / Customize / Advanced Settings and change the Initial number of Overtone Sliders.

You can also save a slider layout as DefaultSliderLayout.oa into the folder where VoceVista Video is installed. After you restart the application, new documents will use this layout. Note: You still have to set the number of sliders to appear in the advanced options.

3.4. Editing Audio Files

This section explains how to perform editing operations such as cutting and copying, inserting silence, fading, and normalization. Most of those commands are found in the Edit Menu. Some also have toolbar buttons, and all of them can be accessed through customizable keyboard shortcuts.

3.4.1. Undo / Redo

The Undo and Redo commands revert or restore the last action that has been performed. The commands can be found in the Edit Menu and on the Toolbar. The toolbar buttons for Undo and Redo can show a history of all commands.

3.4.2. Cut / Copy

The clipboard commands Cut, Copy and Paste work mostly like you would expect from any other program. You can select a parts of your recording, or the entire file, and copy or cut it. Doing so will put the copied section on the clipboard, so that it can be pasted elsewhere.

3.4.3. Paste

Once you have copied or cut a piece of a recording, you can insert it with the paste command. If you simply place the cursor somewhere and paste, the audio data will be inserted at this position. If you make a selection and paste, the pasted audio data will replace the current selection.

> **Note**
>
> Pasting across files only works when the receiving file is an empty new file, or if the files have the same sampling rate and channel format. VoceVista Video does not perform any automatic file format conversions during paste.

3.4.4. Delete

This command deletes the current selection. If nothing is selected, the entire file is deleted. This can be useful if you want to quickly clear the screen when using VoceVista Video for practice.

If you deleted something by accident, always remember that you can use the Undo command to revert your action.

3.4.5. Crop to selection

This command does the opposite of Delete. Instead of deleting the selected part of a recording, it deletes those parts that are not selected, which can be used to trim unneeded parts before and after a recording.

3.4.6. Normalize

This command applies to the current selection or, if nothing is selected, to the entire file. Normalizing a recording will raise its amplitude such that the loudest part of the selection will use the largest possible value of the current file format. This can be used to raise the volume of very quiet recordings. It can also be useful to normalize recordings before sharing them with others to ensure a consistent volume across multiple files.

However, normalization cannot substitute a good recording that was made with a properly adjusted microphone. Normalizing is similar to increasing the brightness of a photograph that is too dark: no new information is added, the available information just becomes slightly easier to see or hear.

Also, since normalization uses the loudest part of a selection as a reference, loud noises in the recording will prevent successful normalization. Hence if loud noises are present, it may be necessary to delete them before normalization.

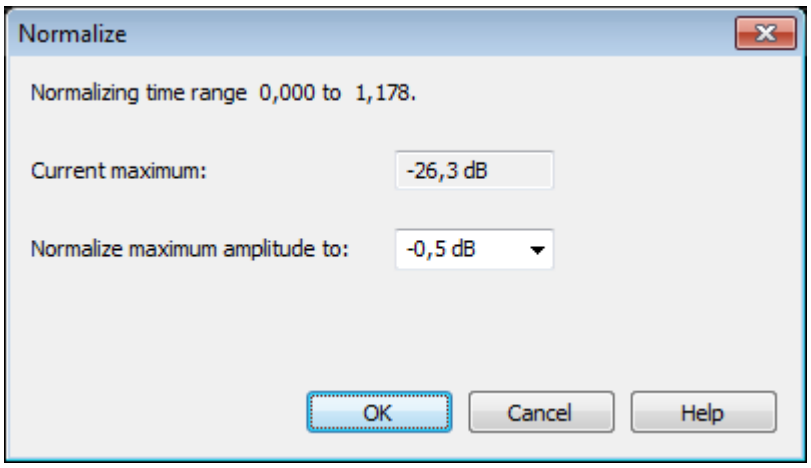

Figure 3.63. Normalize Dialog

Current maximum
: The largest intensity of the current selection.

Normalize maximum amplitude to
: This is the new value that the current maximum will be raised to. If the current maximum is already equal to or larger than this value, nothing will happen.

3.4.7. Replace with silence

This command replaces the current selection with silence. This is different from the Delete command in that the recording retains its length when Replace with silence is invoked, while with Delete, the recording is shortened.

3.4.8. Insert silence

This command will insert silence at the current cursor position. It will open a window where you can specify the duration of the silence that you wish to insert. If you later wish to quickly insert the same amount of silence again, you can press the shortcut for the Reinsert silence command.

3.4.9. Fade in / Fade out

These commands will change the volume of the current selection so that it gradually fades into zero. If you have copied and pasted part of a recording to a new file, the fade commands can be applied to the edges of this segment to avoid any loud pops or clicks at the beginning and end.

3.5. Profile Manager

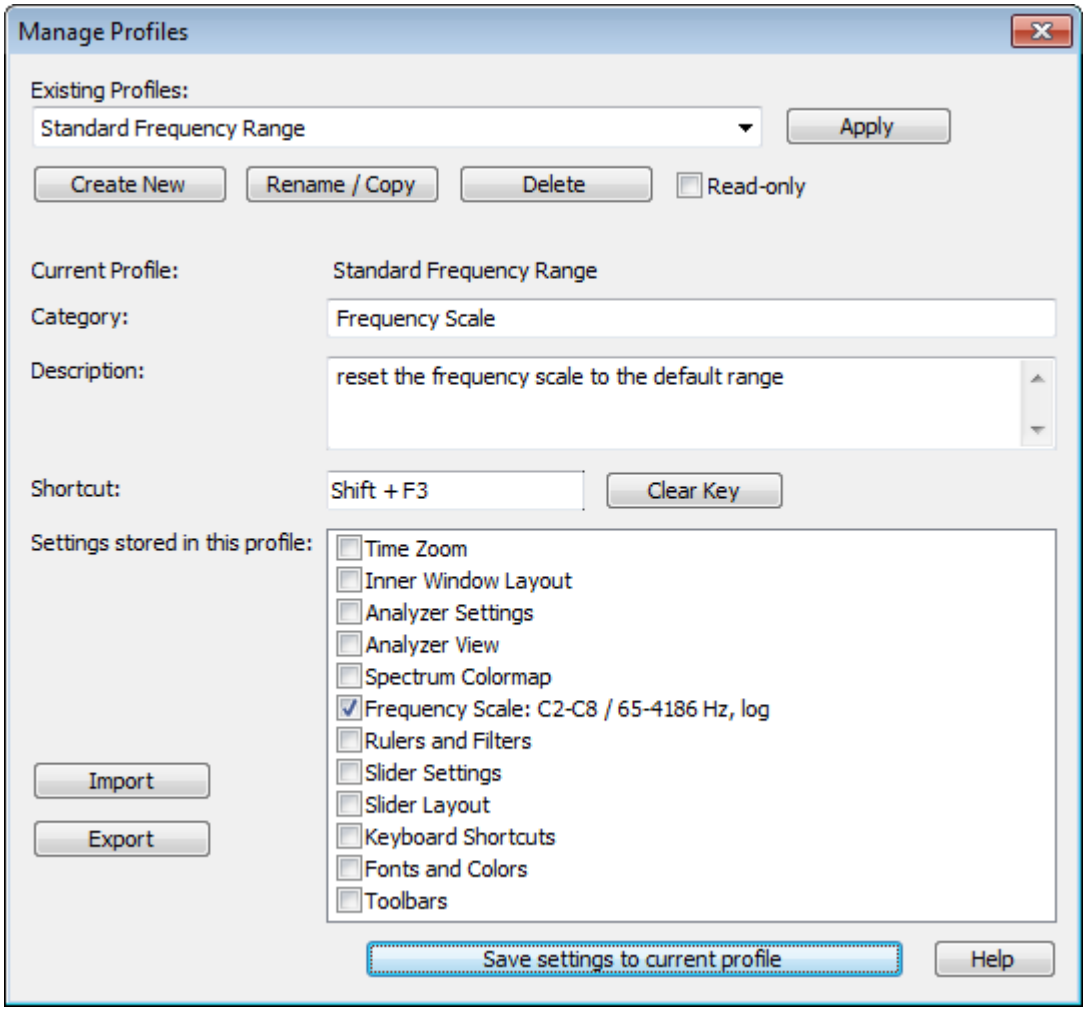

Figure 3.64. Profile Manager

> **Note**
> The Profile Manager is only available in *VoceVista Video Pro*.

The Profile Manager is a tool for advanced users that allows to save a group of settings for later use. This allows, for example, to load different predefined toolbar layouts for beginners, intermediate or advanced users. It also allows to share settings between users by exporting layouts from one computer and import them on another. Another use is to assign hotkeys to frequently used settings, for example to quickly return to a specific frequency range.

Each Profile can contain settings from one or more categories such as frequency scale settings, toolbar arrangement, fonts and colors, etc. When the Save settings to current profile button is pushed, the current value of the selected settings is stored into the current profile.

3.5.1. Workflow

Creating profiles

1. First, bring the program into a state that you want to save. Now click on Tools → Manage profiles to bring up the Profile Manager.

2. Press Create New to create a new profile, or press Rename / Copy to make a copy of the current profile. Choose a name for your profile that indicates its purpose.

3. Then select the categories that you want stored in this profile by clicking on the various entries in the category list (such as *Frequency Scale*, *Timezoom*, etc.).

4. Now press Save settings to current profile to store the current state of the program into your profile.

5. Optionally, you can also add a description for your profile, and assign a keyboard shortcut to it.

Using profiles

To apply a profile, select it from the list of existing profiles on the toolbar. This is the listbox that says Select profile. If the profile has a keyboard shortcut, you can press that to apply the profile. If you are in the Profile Manager, press Apply to activate the current profile.

Applying a profile means that the settings stored in the profile are applied to the program. This is the same as adjusting these settings by hand, and it is instantaneous. For example, the profile *Standard Frequency Range* sets the frequency scale to range from 65 Hz to 4186 Hz. When you apply this profile, the frequency scale is set to this range, as if you had manually set this range by hand. There is no concept of "being in" a specific profile. Instead, you are free to change the settings of the program, and to restore specific settings when you need them by applying a profile that contains the settings you want.

Sharing profiles

You can save the current profile to a file by pressing Export, and load this file on a different computer by pressing Import in the profile manager on that computer.

Loading profile files

Profile files have the .oapr file extension. You can open those files in VoceVista Video through the File menu, or by dragging the .oapr file on the application window.

Importing profiles from web

To make it easier to follow tutorials about VoceVista Video, you can import settings from a website that offers a profile through a profile link. Some browsers may show a security warning before opening the link.

3.5.2. Profile Categories

Here is an overview of the various categories of settings that profiles can contain, and that can be selected on the Profile Manager dialog:

Time Zoom

The time range of the Spectrogram and associated views. This can be changed by zooming.

Inner Window Layout

This category stores the layout, visibility and configuration of the Staff view, the Piano, the Analyzer View, the Waveform View, the Timeline, and the Scales (Frequency Scale, Time Scale, Intensity Scale, Amplitude Scale, etc.). This includes configuration of the Scales (for example if they show vertical labels), and of the Waveform and Timeline (for example if they are symmetric, and if they show decibel). It also includes Marker visibility settings.

Analyzer Settings

This category stores the settings from the Analyzer Settings page.

Analyzer View Settings

This category mostly contains the settings that are found on the Analyzer View settings page.

Spectrum Colormap

This category stores the current colormap of the spectrum, which can be selected on the ColorMap Editor.

Frequency Scale

This category contains the range of the frequency scale, and whether the scale is linear or logarithmic. These settings are found on the Scale Settings page.

Rulers and Filters

This category stores the current layout of rulers on the frequency and amplitude scales, and the current number and position of frequency filters.

Slider Settings

This category contains the settings for Overtone Sliders. These are found on the toolbar, and on the Note Sliders settings page.

Slider Layout

This category stores the current number and position of note sliders.

Keyboard Shortcuts

This category stores the current keyboard shortcuts, which can be set on the Keyboard Shortcuts settings page.

> **Note**
>
> If you assign a shortcut key to your profile, and you use profiles with different keyboard shortcut configurations, be aware that the profile hotkey must be separately assigned in all configurations.

Fonts and Colors

This stores the settings from the Fonts and Colors settings page.

Toolbars

Profiles with this category contain the layout of all toolbars and auxiliary windows such as the File / Markerview and the Vowel Chart. This is useful to save the layout of toolbars that have been customized.

3.5.3. Predefined Profiles

VoceVista Video comes with a set of predefined profiles that you can use to apply various standard settings:

Presets
: The profiles in the “Presets” category contain collections of various settings to give examples of how profiles can be used, and to provide a way to restore default settings of the program.

  Quickstart
  : Settings for the Quickstart Tutorial.

  Standard Frequency Range
  : Set the Frequency Scale to be logarithmic and to range from C2-C8 (this is 65 Hz to 4186 Hz).

  Standard Time Range
  : Set the Time Scale to have a range of 10.0 seconds.

Toolbars
: The profiles in this category apply to the Toolbars:

  Minimal Toolbar
  : A minimalistic toolbar for absolute beginners. Only contains buttons for recording and playback.

  Single Row Toolbar
  : A toolbar with most standard buttons for recording, editing, and visualization. Fits one a single row of buttons to preserve screen real estate.

  Standard Toolbar
  : A two-row toolbar with buttons for most operations, including Overtone Sliders, zooming, and window management. Recommended most users that have some experience in using the program.

  Full Toolbar
  : A two-row toolbar with most buttons that are available, including controls to set the frequency scale range.

Glossary

Ambitus	The ambitus of a voice is the range, or the distance, between the highest and the lowest singable note. The staff view can show the typical ambitus of speaking and singing voices and of singable overtones.
Amplitude	The amplitude is the maximum value of a signal over a given period of time. This correlates to the *intensity* and to the perceived loudness of a signal. It has no unit, but is scaled into the range [-1, 1], where -1 and 1 represent the largest values that a particular file format can encode.
Analyzer View	The Analyzer View is the central window in VoceVista Video and contains one or two sub-windows that can show the Spectrogram, the Spectrum, or both.
Auto Marker	An Auto Marker is a type of *Marker* that is automatically created for each recorded segment. In other words, every time you press record, and then stop, a new auto marker is created to mark the recorded time period.
Bit Depth	Same as *sample size*.
Cent	A cent is one hundredth of the distance between two notes on the piano, or between two semitones of the tempered scale. In other words, two consecutive keys on the piano (regardless if black or white) are 100 Cent apart. The cent is used to measure extremely small intervals. One octave is divided into 1200 Cent.
Clipping	Clipping is the effect when parts of the recorded audio signal are too loud to be represented by the used sample format, and are therefore cut off. For example, the audio format may be able to represent sample values between -1.0 and 1.0. If the incoming signal contains values larger than 1.0, they will all be set to 1.0, which causes a loss of information, and a distortion of the signal.
Decibel (dB)	The decibel is a logarithmic unit that indicates the ratio of an intensity relative to a reference level. When used to represent the *intensity* of an audio signal or of individual frequency components, the reference level is 0dB, which represents the loudest sound that can be encoded in a particular file format. A decibel value of 0dB equals an amplitude of 1. All intensities that are smaller than the loudest reference level have a negative decibel value. The available range depends on the bit depth of the file format. With 16 bit, the smallest intensity that can be represented is -90dB, and with 24 bit, it is approximately -140dB.
Dynamic Range	The dynamic range is the ratio between the largest and the smallest value that can be represented by a given format. The dynamic range is typically measured in *decibel*. In digital audio, common dynamic range values are 90db (for 16-bit audio), and 140db (for 24-bit audio).
EGG	Short for *Electroglottograph*.
Electroglottograph (EGG)	The Electroglottograph (or "EGG") is a small device that allows to estimate the closing and opening of the glottis [https://en.wikipedia.org/wiki/Glottis], the opening between the vocal chords.
Fast Fourier Transform (FFT)	The FFT is a mathematical process that converts a series of samples in the time domain (such as a digital audio recording) into a list of frequencies and their intensity.
FFT Window Function	The window function is a set of coefficients between 0 and 1 that are multiplied with a sequence of samples before taking the *FFT* of this sequence.

	The purpose of this is to reduce mathematical artifacts in the Spectrum arising from discontinuities between the beginning and the end of the signal.
File / Marker List	The File / Marker List is a window that lists the *Markers* of the current file. It can also show a list of recently used files, or a list of search results. Further, it allows to add and edit markers and marker descriptions.
File Description	The File Description is a special type of *Marker* that is automatically added to every file. Each file has a description, which is the first entry in the marker list, and which has the round information icon as symbol. It can be used to add a description to the file (such as what it contains, when it was recorded, where, with whom, and any other relevant information).
Filter	Short for *Frequency Filter*.
Formant	A formant is a resonance frequency in the vocal tract. The vocal tract has multiple resonance tones that will amplify sound with the frequency of that tone. The sound can come from the vocal chords, but it may also come from other sources. The literature on the voice does not always clearly distinguish between formants and *overtones*. Overtones are frequency components of a sound that may be amplified by the vocal tract if they match the frequency of a formant.
Frequency	The frequency is the number of cycles per second. The unit of frequency is the *Hertz (Hz)*. The frequency of a sound wave determines its *pitch*.
Frequency Filter	Frequency Filters are a tool to isolate individual parts of a recording in the frequency domain and make them louder or quieter. This allows, for example, to listen only to specific frequencies in a recording, or to take them away entirely.
Frequency Resolution	The frequency resolution of the *Spectrum* is the difference in Hz between two frequencies that the analyzer can distinguish. The frequency resolution can be set on the Analyzer Settings page. Smaller values show more detail in the Spectrum and Spectrogram, but they also require more processing power and can make the program slower.
Fundamental	For a tone that has multiple harmonic components, the fundamental tone is the frequency that forms the base of an overtone scale that contains all these harmonics. In most cases the fundamental is the pitch that a human listener will identify when hearing the tone.
Harmonic	Harmonic is another word for *overtone*, with one small difference: Harmonics are counted such that the fundamental is the first harmonic, while overtones are counted such that the first overtone is the second harmonic.
Harmonic Series	The harmonic series is the set of frequencies that are all integer multiples of a fundamental frequency.
Hertz (Hz)	Hertz is the unit of frequency to indicate the number of cycles per second of a periodic phenomenon. It is named after the German physicist Heinrich Hertz.
Intensity	The intensity is a measure of how loud or strong a signal is. The Waveform shows the intensity of the entire recording for each point in time, while the Spectrum shows the intensities of the individual frequency components. The intensity can be measured as *amplitude*, or in *decibel*. The intensity is not identical to the loudness of the whole signal or of the frequency components, because the human ear perceives different frequencies differently. For example, if two tones are played with the same intensity, one with 100Hz, and the other with 1000Hz, a human listener might hear one as

	louder than the other, even though they have the same amplitude when leaving the speaker. The intensity that VoceVista Video can show is therefore not the loudness experienced by a human listener, but the sound pressure level recorded by the microphone.
Lin	Short form of *linear*. Opposite of *logarithmic*. On a linear scale, numbers with the same distance have the same difference.
Log	Short form of *logarithmic*. Opposite of *linear*. A log scale can be useful to display numbers that range from very small to very large, especially values that represent quantities perceived by humans. On a log scale, numbers with the same distance to each other have the same ratio, whereas on a linear scale, numbers with the same distance have the same difference. The piano has a log scale. All octaves are the same distance apart, as each octave is a doubling of the frequency. If the piano is projected on a linear scale, the piano keys become progressively wider.
Long-term view	The long-term view is part of the *Analyzer View* and shows things that span a relatively long range of time, such as a Spectrogram, a melody, or a musical piece. The long-term view has a frequency scale and a time scale.
Marker	A marker marks a specific point in time, or a time range, in a recording. It can hold text to name and describe the area of interest. Markers can be used as searchable bookmarks to easily find specific points in a recording, and to add comments and notes. There are four types of markers: *Auto Markers*, *Range Markers*, *Point Markers*, and the *File Description*.
MIDI	Short for *Musical Instrument Digital Interface*, a standard protocol to encode messages to electronic instruments. In VoceVista Video, MIDI output is used to play the keys of the piano keyboard and of overtone sliders. It can be send to the standard MIDI synthesizer that is part of the operating system, or it can be send to external instruments connected to the computer. MIDI is also used as a file format to store a musical piece as a sequence of notes.
Mono	A mono recording has one channel, for example the input of a single microphone.
Note Slider	Same as *Overtone Slider*. The terms “Note Slider” and “Overtone Slider” are used synonymously, depending on the context.
Oscilloscope	A display that shows how a signal changes over time on a two dimensional graph, where one axis is time, and the other axis is the intensity of the signal. In VoceVista Video, an oscilloscope display can be shown by zooming in the *Waveform View* very far.
Overtone	An overtone is a tone that relates to a specific fundamental tone. Each overtone has a frequency that is a whole multiple of the fundamental frequency. For example, if the fundamental has a frequency of 100Hz, its overtones have 200Hz, 300Hz, etc. Also called *harmonic*, or *partial tone*.
Overtone Slider	Same as Note Slider. Overtone Sliders are a visual tool that is laid over the Spectrogram. Each slider represents a frequency. This can be interpreted as a music note, and it can be played as a sound. Sliders can be “drawn out” to show the overtones and undertones of the fundamental frequency. Sliders can be used to highlight a specific frequency or note, to illustrate principles of music theory and acoustics, or to transcribe a piece of music and show its notes.

Partial tone	Other word for *overtone*.
Pitch	Pitch is a perceptual property of a sound that corresponds to the frequency of a tone. Pitch allows to classify tones as “higher” or “lower”. Pitch is not a purely objective physical property because a human listener may perceive the pitch of a tone differently from its measurable fundamental frequency. However, in VoceVista Video, pitch and frequency of a sound are used mostly synonymously.
Playback Cursor	Other word for *Time Cursor*, especially during Playback.
Point Marker	A Point Marker is a type of *Marker* which marks a specific point in time and has no range.
Profiles	Profiles are a set of user settings that can be stored and retrieved. Profiles can contain most settings that can be changed by the user, such as the range of the frequency scale, the arrangement of toolbar buttons, or the display configuration. When a profile is saved, the current state of those settings is written into the profile. When the profile is later activated, all affected settings will be set to the value in the profile.
Range Marker	A Range Marker is a type of *Marker* that marks a period of time with a beginning and an end.
Ruler	A ruler is a visual aid that marks a specific frequency or amplitude. Over the Spectrogram, rulers are similar to *Overtone Sliders* in that they represent a frequency. However, contrary to sliders, rulers have no label, no overtones, and cannot be played. They are simply a visual tool.
Sample	A single measurement of sound pressure, or amplitude. In a digital recording, sound is stored as a sequence of numbers. A sound wave travels through the air and moves the membrane of a microphone. The microphone converts this mechanical movement into an electrical current, and the sound card reads out this current many times per second and stores each sample as a number that can be further processed by the computer.
Sample Size	The number of bits of each sample in a digital recording. Common values are 16, 24 and 32 bit. Larger values can represent a larger *dynamic range* of intensities.
Sampling Rate	The number of discrete measurements (or “samples”) per second stored in a digital audio recording. The sampling rate determines the frequency range that can be represented by an audio file. The highest representable frequency is half the sampling rate. For example, in a file with a sampling rate of 44100 Hz, the highest frequency that can be displayed in the Spectrum is 22050 Hz. Common values are 44100 samples per second for CD-Quality sound, or 48000, 96000 and 192000 samples per second for studio-quality sound.
Short-term view	The short-term view is part of the *Analyzer View* and shows things that span a relatively short range of time, such as a single Spectrum. The short-term view has a frequency scale and an intensity scale. However, the intensity scale only applies to the Spectrum, and not to the pitch value.
Spectrogram	The Spectrogram is a series of spectra. Whereas the Spectrum shows a single frequency-intensity diagram, the Spectrogram shows many such diagrams side-by side. Therefore, the Spectrogram is a two-dimensional diagram where one axis shows time, and the other shows the frequency. The intensity of each frequency at a specific point in time is now represented by the color of this point.

Spectrum	The Spectrum shows the strength of the individual frequency components in a piece of sound at a specific point in time. The Spectrum is a two-dimensional diagram, where one axis shows the frequency, and the other shows the intensity of each frequency.
Staff View	The Staff View shows a musical staff with treble and bass clefs. The location of the staff lines corresponds loosely to the location of the associated pitch on the frequency scale. When notes are played on the piano or the overtone sliders, they are shown as musical notes on the staff view.
Stereo	A stereo recording has two channels. To make a stereo recording, you need a recording device with two separate microphones. Stereo recordings are normally used to add “depth” to a recording by reproducing sound as a human listener would hear it with two ears. However, the two channels can also be used for different purposes, for example to record the sound from within an organ with one microphone, and the sound from the outside with another.
Time Cursor	Green line that indicates the time in the recording that is currently being played (or that will be played next). Also, when the Spectrogram and the Spectrum are both visible, the Time Cursor determines the time position of the Spectrum.
Time Range Slider	The Time Range Slider is a graphical interface element on the *Timeline View* that shows the current time range of the Spectrogram and the Waveform.
Time Resolution	The time resolution of the analyzer determines the length of a piece of a recording that the analyzer uses to calculate its Spectrum or pitch. A lower time resolution means that the analyzer can look at a longer piece of a recording. This will give more accuracy in the frequency domain at the expense of resolution in the time domain.
Timeline	The Timeline View shows an overview of the entire recording. It is similar to the *Waveform View*. The difference to the Waveform View is that the Timeline is zoomed out further than the Spectrogram and may show the whole recording, while the Waveform always shows the same time range as the Spectrogram.
Undertone	An undertone is a tone that relates to a specific fundamental tone. Each undertone has a frequency that is a whole ratio of the fundamental tone. So undertones follow the sequence 1/2, 1/3, 1/4, 1/5 etc. For example, if the fundamental has 100Hz, the undertones have the frequencies 50Hz, 33.33Hz, 25Hz, 20Hz, etc. Each undertone is a tone that has the reference tone as one of its overtones.
Vowel Chart	The vowel chart shows the first and second resonance frequencies of the oral cavity (sometimes called “Formants”) that are used in many languages to form a specific vowel. The chart is a two-dimensional diagram where one axis represents the first, and the other the second formant. The vowels are shown as symbols from the *International Phonetic Alphabet (IPA)*.
Waveform	The Waveform View shows the samples of a digital recording. When the displayed time range is very small (in other words, when the view is zoomed in very far), the individual samples are shown, as on an oscilloscope. When the view is zoomed out, each pixel shows an aggregate with the maximum and minimum values of the samples contained in the time range corresponding to this pixel. The values in the vertical middle of the Waveform show the “Root Mean Square (RMS)” of the signal.

Bibliography

[Miller2008] Donald Gray Miller. *Resonance in Singing*. Voice Building through Acoustic Feedback. Inside View Press. 2008. ISBN 978-0-9755307-5-7.

[Saus2004] Wolfgang Saus. *Oberton singen*. Das Geheimnis einer magischen Stimmkunst. Traumzeit-Verlag. 2004. ISBN 3-933825-36-9.

[Herbst2010] C. T. Herbst, W. T. Fitch, and J. G. Švec. *Electroglottographic wavegrams: a technique for visualizing vocal fold dynamics noninvasively*. *Journal of the Acoustical Society of America*. 128. 3070–3078. doi:10.1121/1.3493423 [https://doi.org/10.1121/1.3493423].