

VoceVista User Guide

Revision 162, Built 2026-05-14

Table of Contents

1. Introduction	1
2. Quickstart Tutorial	7
3. Reference Guide	25
3.1. User Interface	25
3.1.1. Overview	25
3.1.2. Scrolling and Zooming	26
3.1.3. Selecting Time Range and User Interface Elements	30
3.1.4. Timeline and Waveform	33
3.1.5. Analyzer View	35
3.1.6. Frequency Filters	37
3.1.7. Note Sliders	40
3.1.8. Vowel Chart	45
3.1.9. Long-Term Average Spectrum	47
3.1.10. Markers	50
3.1.11. Statistics Pane	55
3.1.12. Vibrato View	56
3.1.13. EGG (Electroglottograph)	57
3.1.14. Toolbar	58
3.2. Settings	58
3.2.1. Recording Settings	58
3.2.2. Playback Settings	62
3.2.3. Analyzer Settings	64
3.2.4. Acoustic Analyzer View	69
3.2.5. Spectral Peaks Settings	72
3.2.6. Vibrato Display Settings	74
3.2.7. EGG Display Settings	75
3.2.8. Colormap Editor	80
3.2.9. Scale Settings	83
3.2.10. Note Slider Settings	86
3.2.11. Fundamental Ranges	89
3.2.12. Selection and Pointer	91
3.2.13. Keyboard Shortcuts	94
3.2.14. Fonts and Colors	99
3.2.15. Advanced Settings	100
3.3. Loading and Saving	108
3.3.1. Loading and saving Audio Files	109
3.3.2. Export Screenshots	110
3.3.3. Export Data	111

3.3.4. Export Sound Generator Profiles	115
3.3.5. Loading and Saving Note Slider Layouts	115
3.4. Editing Audio Files	115
3.4.1. Undo / Redo	116
3.4.2. Cut / Copy	116
3.4.3. Paste	116
3.4.4. Delete	116
3.4.5. Crop to selection	116
3.4.6. Replace with silence	116
3.4.7. Insert silence	116
3.4.8. Fade in / Fade out	117
3.4.9. Normalize	117
3.4.10. Change File Format	117
3.5. Sound Generator	118
3.5.1. Fundamental and Overtones	118
3.5.2. Modulation	120
3.5.3. Resonances	123
3.5.4. Manage Sound Generator Profiles	124
3.6. Profile Manager	124
3.7. Customizing the User Interface	128
4. Theory and Concepts	130
4.1. Waveform	131
4.2. Frequency	133
4.3. Intensity	135
4.4. Spectrum	137
4.5. Spectrogram	138
4.6. Pitch / Fundamental Frequency	139
4.7. Harmonics	140
4.8. Subharmonics	141
4.9. Resonances and Formants	142
4.10. Smoothed Curves and the Spectral Envelope	143
4.11. Long-Term Average Spectrum	144
4.12. Vibrato	145
4.13. EGG (Electroglottograph)	146
5. Troubleshooting	149
5.1. Fix Recording	149
5.1.1. Program is not recording anything	149
5.1.2. Recorded sound is distorted	149
5.2. Contact Us for Support	149
Glossary	150
Bibliography	160

Chapter 1. Introduction

VoceVista shows you what your sound is made of. It is a high-quality, high-resolution real-time audio spectrum analyzer wrapped in a musical interface, designed to help singers, instrumentalists, teachers and anyone curious about sound to see, hear and understand the structure of what they sing, play or record.

Two questions, one tool

As a singer or instrumentalist, almost everything you work on can be traced back to two questions:

1. **Which note am I singing or playing?**

What is the fundamental pitch, and does it match the melody I want to produce?

2. **How am I singing or playing this note?**

What is the tone color, the **timbre** — and therefore: which **harmonics** make up my sound, and in what proportions?

VoceVista is built to answer both questions, in real time and side by side. You can watch the pitch of every note as you sing or play it, on a musical staff and a piano keyboard, and at the same time see the full harmonic structure of your sound on a spectrogram, with musically meaningful axes and overtone guides.

A precise spectrum analyzer for any sound

At its core, VoceVista is a precise, real-time audio spectrum analyzer. While it has been carefully tailored for musical use, it is not limited to voice or instruments: you can use it to analyze any kind of sound — singing bowls, bird song, machine noise, room acoustics, synthesized tones, or recordings you load from disk. The same display that shows you a singer's vibrato can show you the partials of a bell or the resonances of a room.

Understanding sound, music and the voice

VoceVista is also a tool for **understanding**. It makes the interplay between sound, music and our perception of music tangible: how harmonics combine to form a timbre, how musical intervals relate to the natural overtone series, and how the ear turns physical vibrations into the experience of pitch and tone color.

For the human voice, this becomes particularly powerful. VoceVista lets you observe and explore the different layers of voice production:

- the **harmonics in the source sound** produced at the **vocal folds**,
- the **resonances of the vocal tract** that shape this raw sound,
- the resulting **formants** in the radiated sound, and
- the role those formants play in the perception of **vowels** and vocal color.

Seeing these layers at work — and being able to change them with your own voice in real time —

turns abstract theory into something you can hear, see, and practice with.

Measure, understand, learn

VoceVista supports a simple but powerful workflow cycle that brings these capabilities together:

Measure

Record and analyze sound with high precision. Track pitch, vibrato, intensity, harmonics, formants and more — for a single note or for entire pieces — and store recordings for later comparison.

Understand

Use the visualizations, theory chapters and built-in [Sound Generator](#) to make sense of what you measure: see how a sound is built up, how harmonics and formants behave, and how musical concepts map to the physics of sound.

Learn & practice

Apply that understanding to your own playing or singing. Use VoceVista as a visual mirror for daily practice, follow your progress over time, and refine specific aspects of your technique with direct, immediate feedback.

These three aspects go hand in hand, and you can use each of them on its own. You can simply measure a sound, you can use the program to study the theory of harmonics and formants, or you can use it as a practice companion — or all three together, as a continuous loop.

What you can do with VoceVista

To make this concrete, here are some of the things VoceVista helps you do every day:

- Record your voice or instrument, visualize it and listen back to it.
- Measure the pitch of a sustained note and practice holding it steadily over its full duration.
- Visualize and shape your [vibrato](#).
- Analyze the harmonic structure of a voice or an instrument.
- Compare different singers, instruments, takes or vowels side by side.
- Identify and isolate individual overtones of a fundamental, and practice singing the harmonic series.
- Observe vocal tract resonances and formants while exploring different vowels.
- Transcribe the notes in a recording.
- Construct and explore musical scales through the natural relationships between their tones.
- Analyze any other sound that interests you — instruments, environments, acoustic experiments.

Who is VoceVista for?

Singers can use the program as a visual feedback aid to practice pitch and [intonation](#), to develop and refine vibrato, formants and vowel color, and to work on any other aspect of their vocal technique.

Singing teachers can use VoceVista like a mirror for the voice, to show students what they are actually doing and to make subtle changes in tone, resonance and pitch immediately visible.

Voice therapists can use the program to monitor the progress of their clients and as an additional sensory feedback channel for specific exercises.

Instrumentalists can use VoceVista to check intonation, analyze their tone, compare different instruments, mouthpieces, reeds or setups, and follow how their sound changes with technique.

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

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

Choir conductors can analyze blend, intonation and timbre across the ensemble, identify problem passages in recordings, and work on chord purity and a more uniform vowel color — helping even amateur choirs achieve a polished, professional sound.

Composers and arrangers can use VoceVista to study how voices and instruments produce their characteristic colors and how those colors combine, and to ground choices about register, scoring and timbre in what is actually happening in the sound.

Musical theorists find in VoceVista a visual aid for exploring intervals, scales and tuning systems, and for relating classical harmony to the natural relationships between tones.

Overtone musicians can use VoceVista to study and improve their technique and to create or rehearse complex compositions, with the visualization of their own harmonics as direct feedback.

These are just a few examples of who works with VoceVista; in practice, anyone curious about sound and how it is built up will find something useful here.

Overview

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 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](#).

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: Click **[OK]** to continue means: click on the button labelled **OK**.

Keyboard commands look like this: press **F12** to take a screenshot.

The [Glossary](#) at the end contains definitions for many subject-specific terms used in this document.

Installation

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

Operating System

macOS

VoceVista for macOS requires *macOS 11* or later. It works on both Intel and Apple Silicon chips.

Windows

VoceVista for Windows requires *Microsoft Windows 10* or later. It also works on Windows 11.

If you have an older computer which doesn't run the required OS version, you can download an older version of VoceVista that works with your system at <https://www.vocevista.com/en/download-old/>.

Computer

Most modern computers can run VoceVista without any problems, but a faster machine with more memory and hard-drive space will run the program more smoothly. Some lower-end models have slower graphics cards, which may result in poor performance and lower image quality when rendering spectrograms.

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.

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.

Installation

You can download the latest version of the software on our [download page](#). Clicking on the download button will download the installer file to your computer. If you cannot find it, open a file browser and locate your "Downloads" folder, where the VoceVista installer should be the newest item. Double-click on it to open it.

macOS

On macOS, once you have downloaded and opened the installer file, you will see a window like this:




Figure 1. macOS Installation window

In here, drag the "VoceVista" icon onto the Applications folder.

Now open a Finder window. On the left side, click on the Applications folder and confirm that the VoceVista icon has been copied into that folder. Now, under "Locations", find the entry called "Install VoceVista" and click on the "eject" symbol next to it (marked with a red circle in the image below) to close the installer.




Figure 2. Ejecting the installer

Now you can start VoceVista by clicking on its icon in the Applications folder or the Launchpad.

Windows

On Windows, once you have downloaded the setup program, double-click on it to open it, and then you should see a window like this:




Figure 3. setup program on Windows

Simply keep clicking on the **[Next]** button until the installer is finished.

Chapter 2. Quickstart Tutorial

This guide will show you how to record and visualize your voice or instrument with VoceVista. If you have no idea what to do with VoceVista, this is the place to start. This tutorial is intended to be a short introduction that gives you an overview of the elements of the program. You can find more detailed explanations of everything in the [Reference Guide](#).

NOTE

Some features require VoceVista Video or higher

Some of the features shown in this guide (most notably the Frequency Filters) are only available in *VoceVista Video* and higher. To follow this tutorial, you can download a free 30-day trial version of *VoceVista Video Pro* at [our download page](#).

This guide is available in various formats:

Online

You can find an online version which you can view in your webbrowser at www.vocevista.com/en/documentation/quickstart.

NOTE

PDF

You can download a PDF version of the entire manual at www.vocevista.com/en/support.

Video

A video version of this tutorial is available at www.vocevista.com/en/videos/quickstart-tutorial.

Prerequisites

- You have installed VoceVista on your computer.

Our software comes in multiple editions with different feature sets. You can actually follow this tutorial in all three editions. The only section that requires at least *VoceVista Video* is the section about Frequency Filters. Apart from that, some screenshots show sections of the recording labelled with [Markers](#). This feature is only available in *VoceVista Video Pro*, so if you want follow the tutorial and see exactly the same screenshots, use that edition.

- You have a working microphone connected to your computer, or you use a laptop with a built-in microphone.

If you have not completed the above two steps, see [System Requirements](#) for more information about what you need, and [Installation](#) for help with installing the software. This tutorial was made with VoceVista version 5.8.0. If you are using an older or newer version of the software, some things might look and work slightly differently.

Configuration

Apply Settings Profile

First let's configure the program to a predefined state, so that you can follow this tutorial and see the same things as shown here. Click on the profile list on the toolbar (the box that says "Select Profile") and select the **Quickstart** profile:




Figure 4. Selecting the Quickstart profile

A profile is a set of saved settings. Selecting a profile is a shortcut for applying all the settings stored in that 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 5. Program with Quickstart profile applied

In this screenshot you can see the empty Timeline on the bottom, the Staff View and Piano on the left, the empty Analyzer View in the main section, and the Toolbar on the top.

Here is a closer look at the Toolbar. This is a simplified toolbar profile with only the most important

elements. As you get more proficient with VoceVista, you can switch the toolbar by clicking on the profile selector and choosing another toolbar profile.




Figure 6. Toolbar buttons

Recording

Select Recording Source

Now we will check that you are recording from the correct microphone. **Point** at the Input Volume control on the toolbar and wait until a tooltip appears. This will show the active input device from which audio will be recorded:




Figure 7. Point at the Input Volume to see active mic

To see a list of all available audio input devices, open the [Recording Settings](#). To do so, click on the **Settings** button on the toolbar ⚙️, and then select **Recording** on the left. Now you can click on the list of input devices and select the one you want to record from:




Figure 8. Selecting input device on Recording Settings

Another way to quickly select the active audio device for recording is to **right-click** on the Input

Volume control on the toolbar and make sure the correct microphone is selected:




Figure 9. Right-click on the Input Volume to select input source

Start recording

Now you are ready to record some sound. Click the Record button on the toolbar , or press **Ctrl** + **Space** or **R** to start recording.

Grant microphone access

If you are running VoceVista on macOS for the first time, there may be a security prompt asking whether the program is allowed to access the microphone. Click on **[Allow]** to grant microphone access:







Figure 10. Allow microphone access on macOS

If microphone access hasn't been granted, recording will not work. If that is the case, you can open the macOS System Settings to grant access later under "Privacy & Security". On Windows, there is a similar setting under "PC Settings / Privacy". For more details, check our [guide for troubleshooting audio input](#).

Set Input Level

If recording has started successfully, the green time cursor line should start moving to the right. Make some noises and sounds into the microphone and watch how the Timeline at the bottom of the screen will fill up with content.

During the recording keep an eye on the Input Volume control and the Timeline and adjust the Input Volume Slider or the distance to the microphone if the recording is too quiet or too loud.

		Recording Volume too low: Analyzer Display will lack detail, and recording is barely audible during playback. Move the input volume slider to the right, or move closer to your microphone.
		Recording Volume too high: Analyzer will show clipping artefacts (this is worse than the volume being too low!). Recording sounds screechy and distorted during playback and is unusable. Move the input volume slider to the left, or increase your distance to the microphone.
		Recording Volume optimal: Signal uses most of the available dynamic range without clipping. Analyzer will show best amount of detail. Sounds good during playback.

Make some sounds into your microphone and adjust the recording level until it is getting a strong signal without clipping. This means that the background of the input level meter should never reach the red zone, and the content in the Timeline and/or Waveform should never touch the upper and lower edges of the view.

If you are recording singing, be aware that the dynamic range of a trained singer can be extreme. This means that singing can get very loud, and the input volume might need to be lower than you think to prevent clipping and distortion of the recording when the singer is singing high notes at high volume. In those cases, besides lowering the input volume, it might help to increase the distance to the mic a bit, or to use a better microphone.

Sing a tone

Now let's try some singing. For example, 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 on the

toolbar to stop recording. This should look something like this:




Figure 11. Singing a tone on the same pitch

The Timeline at the bottom shows the intensity (or volume) of your sound as a waveform display. 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.

- Now try to play some other notes on the piano keyboard, and record your voice to see if you can sing the matching pitch.
- Then try to sing or whistle some melodies, and observe how the blue pitch line and the notes on the piano and staff view respond to your sound.

Scrolling and Zooming

While recording some sound, you have seen how the horizontal axis shows the time of the recording, and the vertical axis shows the frequency of the pitch at any given time.

Once your recording gets longer than a few seconds, you might want to choose which part of the recording you currently see. This can be done with scrolling and zooming the visible time range. Also, depending on the content that you are currently working with, you might want to change the visible frequency range. In this section we will do both.

Loading sample file

Let's load a sample file so that you can follow this tutorial and see exactly the same screens as shown here. Click on **Help** > **Open sample file** and select the file "A01_male_intervals.flac".

After you have opened the file, click on the profile list and select the "**Quickstart**" profile again to ensure your view matches the following screenshots. Remember that you can always restore the **Quickstart** or **Factory Settings** profiles if you want to return the program to the initial state. Your screen should now look like this:




Figure 12. Loading sample file A01

First, let's listen to the file. Click on the play button on the toolbar ► or press Space to start playback. You can adjust the loudness on the output volume slider on the toolbar. Listen to the recording and then press Space again when you are done to stop playback.

In this sample, the singer is singing a triad consisting of three different musical notes. Let's zoom in to see precisely which notes those are, and how accurately the singer was singing them.

Zoom in time

Let's examine the first sung note more closely.

Press and hold the button of your trackpad or mouse, and move the pointer over the blue pitch line. Notice how this moves the green time cursor line. The current time position is displayed on the status bar at the bottom of the window, along with details about the current pitch frequency and musical note at that point in time.

Move the green time cursor line to be exactly at the time point "1.0 seconds". Now click on the menu **View > Zoom in** or press the 1 key to zoom in the current time range. Do this one more time. Now you should see something like this:




Figure 13. Time zoomed to first note

Besides using the menu command to zoom in, you can also move the pointer over the main window and use the two-finger scroll gesture on your trackpad, or the wheel on your mouse, to zoom in and out.

Zoom in frequency

We have zoomed the time scale so that the first sung note is stretched out over most of the visible time range. Now we want to zoom in the frequency range to see more details about the sung pitch on the frequency axis.

Move the pointer over the Frequency Scale just to the right of the piano keyboard. When pointing at the Frequency Scale, a grey horizontal line appears as a visual aid. Move the pointer so that the grey line is approximately matching the blue pitch line, and then use the two-finger scroll gesture on your trackpad, or rotate the wheel on your mouse, or press the 1 and 2 keys on your keyboard, to zoom in the visible frequency range to go approximately from notes A2 to A3 on the Piano, or 110 Hz to 220 Hz on the Frequency Scale:




Figure 14. Frequency range zoomed to A2 - A3

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. Also, it looks like the pitch is rising slightly throughout the duration of the note. VoceVista offers many ways to analyze this in more detail, but for now we will continue with the fundamentals of using the program.

Scroll visible time range

Now let's look at other parts of the recording. Move the pointer over the center of the Time-Range Slider on the Timeline until it changes into a hand. Then click and drag this slider with the mouse until the second note has been scrolled into view:




Figure 15. Move the Time Range Slider with the mouse

Another way to scroll time is simply to click anywhere on the Timeline to move the current range to that position. Try that now: Move the pointer over the Timeline, at around 4.5 seconds, and make a single click. This should move the time range to cover the third note. However, now the pitch of the third note is no longer in view. So point at the Frequency Scale and drag it down a bit until the pitch line is in view again:




Figure 16. Moving third note into view

So far we have set the visible ranges on the time and frequency axes by simply scrolling and zooming those axes with the mouse, and by interacting with the Timeline. If you want to set a specific frequency range, you can use the [Frequency Scale Settings](#). To set a specific time range, you can click on the menu **Edit > Set Selection** to select a specific time range, and then click on **View > Zoom to Fit Selection**.

Play selection

In the last section we have played back the entire recording. But let's say that we are interested in a specific part of the recording and want to repeat it a few times in a loop. To do that, you can select a time range by holding the Command key (on Mac) or the Control key (on PC), and then clicking and dragging the pointer over the time range that you want to select.

Do this now: Hold Command (or Control on PC), and then click on the pitch line at 3.5 seconds, and while holding the mouse button, drag the pointer to 5.5 seconds, and then release it. This should select the time range that encompasses the third note:




Figure 17. Select time range 3.5 to 5.5 seconds

Now you can press Space to play the current selection. If the loop button on the toolbar is pressed, the selection will repeat until you stop playback. Try to move the edges of the selection with the mouse. This can be done even during playback.

Selecting a time range is not just for looping playback, but also for using this range for further analysis and averaging, and to apply [edit commands](#).

To reset the selected time range, hold Command / Control and single-click anywhere in the main window. Alternatively, you can also invoke the menu **Edit > Reset selection**.

Spectrogram & Spectrum

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 the **Settings** button on the toolbar , and then select **Analyzer View** on the left, and **Acoustic** on the top, to open the [Analyzer View settings page](#):




Figure 18. Analyzer View Settings

Check the [**Spectrogram**] option on the left column, just under [**Long-term view**]. Then close the settings window.

Now you should already see the Spectrogram, but the screen might still be zoomed in very far from the previous step. On the toolbar, click on [**Settings Profiles**] and select the **Standard Frequency Range** profile. Then click on **View** > **Zoom out fully**. Now your screen should look similar to this:




Figure 19. Displaying Spectrogram

The colors on the spectrogram indicate the intensity of the sound for each frequency. In other words, the spectrogram shows which frequencies are in the sound, and how loud they are. To see this in more detail at any given point in time, we can enable the **Short-Term view**, and in particular, the **Spectrum**. Reopen the [Analyzer View settings page](#) and check the [**Short-Term view**] and [**Spectrum**] options.

Now the screen is split into two parts. The left side contains the **Long-Term view**, which shows how the recording changes over time. The right side contains the **Short-Term view**, which shows details of the recording at the specific point in time where the green time cursor line is located. Move the pointer around over the spectrum with the mouse button held down, and observe how the spectrum changes as you move the cursor to different time positions. Now move the green cursor line to 8.0 seconds.

Move the brightness slider on the toolbar (as shown on [Figure 6, “Toolbar buttons”](#)) and see how the dynamic range of the spectrum and spectrogram changes in response to that. Adjust the brightness slider so that all peaks on the spectrum are fully visible. Then point at some of the spectral peaks on the right side and see how the text display on the Short-Term view shows the current frequency, nearest musical note, and intensity, for each peak in the spectrum:




Figure 20. Pointing at spectral peak to see its frequency and intensity

So now you can see that the sound in our recording contains many individual frequency components that have different levels of intensity. These are the harmonics. To better understand what this means, let's use another feature of VoceVista: the Frequency Filter.

Frequency Filter

NOTE

Available in VoceVista Video and higher

Frequency Filters are only available in *VoceVista Video* and higher.

Frequency Filters are a tool that allows to pick out specific frequency ranges and listen to them in isolation, or to change their intensity to make them louder or softer. This helps to understand the spectral composition of the recorded sound.

Continuing with the same sample file ("A01_male_intervals.flac"), select the time range that encompasses the voiced part of the last sung note (hence select the range from 7.5 s to 9.5 s), and

center the view on this range. Adjust the visible frequency range so that the fundamental pitch of the note is at the bottom of the screen, and the part around 3000 Hz can be clearly seen (hence set the frequency range to something like 80 Hz - 4000 Hz).

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




Figure 21. Adding a Frequency Filter

The grey frame at 100 Hertz is the Frequency Filter. It will remove all frequencies outside of its range. Check that the **loop** icon on the toolbar is pressed, and then start playback again and hear the effect of the filter. Then, while playback is running, move the filter up and down by dragging its grey frame with the mouse. You can change the width of the filter by dragging the red handles that appear when the mouse cursor hovers over the filter.

The spectrum on the right now shows a white outline of the original spectrum. The part in color is all that is left by the Frequency Filter. With this, you can listen to the individual harmonics or overtones in your sound.

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 in its range. To see this more clearly, set the brightness slider on the toolbar to -18 dB, and set the contrast slider to 60 dB. Drag the Frequency Filter to cover the range from 200 to 700 Hz, so that it removes the second, third, fourth and fifth harmonics in the recording:




Figure 22. Subtractive Frequency Filter

Notice in Figure 22, “Subtractive Frequency Filter” how the Spectrum shows an outline of the second to fifth harmonics, 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 **[close]** button on its top right corner.

Note Slider

The Frequency Filter enables you to listen to the individual harmonics (or overtones) in a recording. A Note Slider is a complementary tool: it gives you a visual reference for a specific note or frequency, lets you count the harmonics that theoretically belong to that fundamental, and shows which note name and frequency they correspond to.

NOTE Note Sliders were called *Overtone Sliders* in earlier versions.

To add Note Sliders, first click on **[Settings Profiles]** on the toolbar and select the **Single Row Toolbar** profile. This will introduce a few new items on the toolbar. The two we are interested in now are the edit field for the number of Note Sliders, and the menu button to set the snap mode:




Figure 23. Toolbar controls for Note Sliders

Each slider has a fundamental frequency and could be interpreted as a musical note, and it can show the over- and undertones corresponding to this frequency.

The **Snap Mode** determines what happens when you move the slider. If snapping is off, it will move freely. If snapping is enabled, the slider will snap to the nearest musical note on the piano, or the nearest peak in the spectrum. Snapping to a musical note would be used if you want to set the note that you want to sing, and snapping to the spectral peak would allow you to measure what you actually sung.

Set the [**Number of Sliders**] edit field on the toolbar to "1". This will add a single slider overlay on the screen. Click on the red line of the slider and drag it up and down.

Click on the icon of the [**Snap mode**] button to enable snapping. Then click on the down arrow on the right side of the button to open the button menu and switch between snapping to musical notes and spectral peaks. Observe how these actions change the behaviour of the slider line when you move it up and down.

Now position the green time cursor line in the middle of Note 5 (at around 8.5 s) to set the spectrum to that position. Then move the slider on the Short-Term view to align with the lowest harmonic. Now point at the right upward facing triangle of the slider and drag it all the way up. Your Short-Term view should now look something like this:




Figure 24. Adjusting the number of harmonics shown by the Note Slider

You can scroll down the Frequency Scale and zoom it slightly until you see enough detail of the Note 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 Note 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.

Summary

This concludes this quickstart guide. You have learned the basics of how to record and visualize sounds with VoceVista 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.
- Press **Ctrl** + **Space** or **R** on your keyboard to start 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 pointer 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 while pressing the **Command** key on Mac, or the **Control** key on PC. The selection will be played in a loop if looping is enabled.
- 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 a Note 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 [Section 3.1, “User Interface”](#) in the Reference Guide.

Chapter 3. Reference Guide

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

3.1. User Interface

3.1.1. Overview

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




Figure 25. 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 **Command** + **Click** on a Mac or **Control** + **Click** on a PC 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 and contains output from the analyzer (such as spectrum, spectrogram, and pitch), as well as Note Sliders, rulers, markers, and frequency filters.

It is divided into two parts, the **Long-Term view** on the left, and the **Short-Term view** on the right side. The Long-Term view shows how the frequency content in the recording changes over time. The Short-term view shows the intensities of all frequencies in the recording at a specific point in time (which is at the location of the green time cursor line).

The Long-Term view can optionally show additional sub-views such as the Waveform, Vibrato Rate and Extent, and the EGG (Electroglottograph) History.

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.

Vowel Chart

The Vowel Chart is one of several side panels that don't belong to a specific file, but are part of the main window. The active side panel can be switched by clicking on the tabs underneath the vowel chart. The panels can also be detached or rearranged.

The Vowel Chart is a tool to visualize and listen to specific frequency pairs that are associated with certain vowels. You can click on the chart to hear the corresponding vowel chart to understand better what this means.

Open file tabs

Just like the tabs in a web browser, these tabs allow you to switch between open files in VoceVista. You can also right-click on a tab so see more actions, for example to arrange open files side-by-side.

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 enables you to easily change parameters directly from the user interface and see those effects immediately. Some of the most important ones you can change are the:

- time range
- frequency range
- the dynamic range of the spectrogram and spectrum
- the dynamic range of the waveform and timeline

The fastest two methods to change these parameters are to hover your mouse over an item and either:

- scroll (shift its range) by clicking and dragging
- zoom (expand or compress its range) by using one of the following:
 - a zoom gesture
 - the mouse wheel
 - the **1** and **2** keys

Note: If you don't have a mouse wheel, you can hover your mouse over an item and use the two-finger gesture for scrolling and zooming on your touchpad. For most computers, put two fingers on the touchpad and move them both up or down.

Interactive areas

The following image shows all areas that you can scroll or zoom:




Figure 26. Scrollable and zoomable areas

1) Time Range

You can easily change the visible time range using the following methods:

1a Spectrogram

- a zoom gesture
- the mouse wheel
- the **1** and **2** keys

1b Time range (use to select with greater precision)

- click and drag to scroll
- a zoom gesture
- the mouse wheel
- the **1** and **2** keys

1c Timeline selection (use to navigate quickly)

- click and drag right/left edges to zoom time range
- click and drag selection to scroll time range
- scroll the timeline using:
 - a zoom gesture
 - the mouse wheel
 - the **1** and **2** keys

1d Timeline (use to navigate quickly)

- click anywhere outside a timeline selection to scroll there
- scroll the timeline using:
 - a zoom gesture
 - the mouse wheel
 - the **1** and **2** keys

You can also zoom using the menu by clicking **View** and then selecting one of the four options at the top (**Zoom In**, **Zoom Out**, **Zoom to Fit Selection**, or **Zoom Out Fully**) or pressing the number that corresponds with each.

Adding [Markers](#) to a file enables you to bookmark specific parts of a recording so that they are easy to find in the future. You can then move between markers by clicking **Markers > Jump to next marker** or **Jump to previous marker**.

2) Frequency Range

You can easily scroll or zoom the frequency scale. Hover over the frequency values and:

- Scroll by clicking and dragging the frequency values
- Zoom by using a zoom gesture or pressing the **1** and **2** keys

Set precise values for this scale on the [Frequency 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. Access the [Colormap editor](#) on the **Settings** dialog box to change all of the values for the dynamic range.

To quickly change basic settings for the dynamic range:

- Hover over the Intensity scale and
 - Scroll by clicking and dragging the frequency values
 - Zoom by using a zoom gesture or pressing the **1** and **2** keys
- Drag the brightness and contrast sliders on the toolbar

4) Waveform / Timeline Amplitude Range

If enabled (**View > Waveform**), the amplitude range appears above the spectrogram and displays the waveform.

When you hover over the waveform, the following zooming actions are available similar to the spectrogram.

- a zoom gesture
- the mouse wheel
- the **1** and **2** keys

When you hover over the amplitude range to the left of the waveform, the following actions enable you to change the intensity of the displayed waveform.

- a zoom gesture
- the mouse wheel
- the **1** and **2** keys

5) Timeline zoom

You can zoom the Timeline by hovering over the timeline, pressing Command (Ctrl for Windows) and then using any of the following actions:

- a zoom gesture
- the mouse wheel
- the **1** and **2** keys

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 **View > Zoom out fully**, or press the **4** key.

Commands

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

Table 1. Commands for scrolling and zooming

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

3.1.3. Selecting Time Range and User Interface Elements

VoceVista has different options that enable you to select:

- time range of a recording
- some elements on the user interface such as multiple piano keys and note sliders

Learning how to select/de-select these items, especially the time range of a recording, is extremely important as it enables you to precisely analyze a range, loop it, apply filters, and other use other application tools.

The following graphic shows a file with a portion selected (highlighted background) for further analysis/manipulation.




Figure 27. Selecting time range and notes

Selection Options Using Mouse

There are two main options for using your mouse to select/de-select time ranges in a file.

In the spectrogram, waveform or timeline:

- Press the **Command** key (on Mac), or the **Control** key (on Windows) (default) while clicking and dragging your mouse to select or de-select. (default setting) This setting is better if you are making a selection and then using different methods to further analyze the recording because you won't lose your selection if you click somewhere else.
- Click and drag your mouse to make a selection (change option in settings). Click anywhere in the spectrogram, waveform or timeline to de-select. Some workflows are much easier and quicker if you can click and drag on the user interface to select part of a file. To enable this behavior, go to [Selection Settings](#) page and check **Lock Selection**.

Regardless of which option you have enabled, once you have selected part of the file, you can click and drag the left and right borders of an existing selection to make it longer or shorter. You can tell which area is selected because the background color changes in the waveform, spectrogram and timeline.

Selection Options Using Menus

You can also use menus and associated keyboard shortcuts to quickly and precisely select or view a

time range:

- Click **Edit** › **Select All** to select the entire range
- Click **Edit** › **Select Visible Range** to select the entire range that is visible
- Click **Edit** › **Set Selection** to more precisely select a range using time values
- Click **Edit** › **Lock Selection** to lock the range that you have selected (see below)
- Click **View** › **Zoom to fit selection** to set the visible time range to match the current selection

Working with a Time Selection

Once you have selected a time range, you have a number of options to view and work with it:

- Press the **[Play]** button to play it repeatedly. You can configure the application to play the selection once by de-selecting **Tools** › **Loop playback** on the **Playback Settings** of the **Settings** window. When you select this option, you can click the **Loop** button on the main toolbar to have the selection repeat.
- Cut, copy, paste, etc. the current selection
- Refine your selection by selecting **Edit** › **Set Selection** to enter exact values for the range.




Figure 28. Set Selection Dialog Box

Locking a Selection

If you click anywhere on the timeline or the spectrogram, the playback cursor will jump to that position, and the current selection will be reset. Sometimes you might want to maintain a selection, but still move the cursor within the selection to hear specific parts. This is particularly useful when comparing files, and when working with [Frequency Filters](#).

To lock the selection, click **Edit** › **Lock selection**. You can still resize the selection by clicking and dragging its borders, but you won't lose your selection if you click somewhere on the Timeline or Analyzer View.

Resetting the Selection

If the selection is not locked, you can simply click on the Timeline or Analyzer View to de-select 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 Note Sliders by clicking on them while holding the Control key (on Windows), or the Command key (on Mac). To play those tones, click on **Notes › Play selected tones** or press the Return key.

To deselect all tones, click **Notes › 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 29. Waveform and Timeline

Waveform and Timeline 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 30. 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:

Table 2. The same recording in various waveform display modes

Layout	Scale mode	Example
Symmetric	linear	
log		
Skyline	linear	
log		

[The same recording in various waveform display modes](#) 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 [The same recording in various waveform display modes](#) 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.

Table 3. Waveform scale mode comparison

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.

3.1.5. Analyzer View

This is the main view of VoceVista. 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 Note 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.

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:

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 the *Spectrogram Display* 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.

Spectrogram

Pitch

Mean Pitch

Short-Term view

Spectrum

Smoothed Curve

Long-Term Average Spectra

Unfiltered Spectrum

Spectrum Display

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 31. 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 [Frequency Spectrum and Fundamental Pitch](#), 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](#).

Combined view: Spectrogram + Spectrum

It is possible to show 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

It is possible to turn off spectrogram and spectrum and show only the fundamental pitch, as in the *Pitch only display*.

3.1.6. Frequency Filters

NOTE

Available in VoceVista Video and higher

Frequency Filters are only available in *VoceVista Video* and higher.

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 32. Using a Frequency Filter

[Using a Frequency Filter](#) 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 33. 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 34. Additive filter (+ Mode)




Figure 35. 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 [Filter Controls](#)).

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**).

3.1.7. Note Sliders

A Note Slider is a visual tool that is laid over the spectrogram. At its core, each slider has a fundamental frequency and, optionally, a start and end time — much like a note in a MIDI sequence. A slider can be played as a sound, and it can also be drawn out to show the harmonics (overtones) and subharmonics (undertones) of its fundamental. Note Sliders can be used as a visual reference for a specific pitch, to practice or transcribe the notes in a recording, to learn the harmonic series for each tone, to analyze the harmonic content of a recording, and to illustrate the role of overtones in composition and music theory.

NOTE

Note Sliders were called *Overtone Sliders* in earlier versions, and the two terms still appear interchangeably in some places. They are sometimes also referred to simply as *Notes*.

Main elements of a Note Slider

The following image shows a Note 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 $2 \times 110 \text{ Hz} = 220 \text{ Hz}$, the third one 330 Hz, and so on.

The image also shows the basic elements of the slider, including the controls that you can use to manipulate a Note 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 **Notes › Play selected tones** (or press Enter to play them).

Slider Lines

Lines can be moved with the mouse. The lines have different colors: + [arabic]

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. + + For example, [Two Note Sliders forming a perfect fifth](#) shows two sliders whose interval forms a perfect fifth. The frequency of the second slider is exactly $\frac{3}{2}$ 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 Note Sliders Toolbar

Many aspects of sliders can be controlled through the Note Sliders Toolbar. If the toolbar is not visible, you can enable it by clicking on **View > Toolbars > Note Sliders**. The toolbar settings are explained on the [Note Sliders settings page](#).




Figure 36. The Sliders Toolbar

Displaying Note 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 Note Sliders

The following table lists the various mouse commands that you can use to manipulate Note Sliders. Refer to [Main elements of a Note Slider](#) for the slider elements.

Table 4. Note Slider Mouse Commands

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.

Mouse action	Effect
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.

Selecting and Playing Slider Tones

The labels of the Note 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 ???.

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 [Note 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 37. Note Sliders fixated to the Timeline

Note Sliders fixated to the Timeline 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 an 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 **Notes > 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 Notes 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 [???](#).

3.1.8. Vowel Chart

NOTE

Available in VoceVista Video and higher

The Vowel Chart is only available in *VoceVista Video* and higher.

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 38. 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 [Rulers and the Vowel Chart](#)) 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 (the first two [formants](#)) 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 menu:[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 menu:[Show harmonic vowels].

3.1.9. Long-Term Average Spectrum

NOTE

Available in VoceVista Video Pro

The Long-Term Average Spectrum is only available in *VoceVista Video Pro*.

This page is about using the LTAS through the selection and through markers. Its theoretical background is explained at [LTAS background](#).

There may be times when you are analyzing a file or a selection of a file (or two files) that you want to have a longer term view of what is happening in the signal rather than just the point in time you get when you click in the spectrogram.

The long-term average spectrum (LTAS) feature enables you to select part of a file, or the complete file, and let the application determine what the average values are that appear in the spectrogram. This can be useful in many ways giving you the ability to:

- view a more stable signal over time
- better identify the [resonance](#) frequencies that make up a signal — for example, the [singer's formant](#) cluster around 2.5 - 3.5 kHz
- compare sections of a file
- compare sections of different files




Figure 39. Long-term Average Spectrum

The following image shows how you can display the LTAS for two different sections of the same file, or two different files, to easily analyze differences. Note how the top image (vocal fry of the letter "i") has much stronger resonances around 260Hz, 2000Hz and 3000Hz while the bottom (vocal fry of the letter "a") has its two strongest signals at around 700Hz 1150Hz.




Figure 40. Long-term Average Spectrum Comparing Sections of a File

View the long-term average spectrum:

1. Select a portion of the file in the spectrogram or press Command/Control + A to select the entire file. You can adjust your selection by clicking and dragging the borders or selecting **Edit > Set selection** to enter precise time values.
2. Click the LTAS button on the toolbar OR



3. Select **Markers > Set average spectrum to selection**

Reset the long-term average spectrum:

When you no longer want to view the LTAS, select **Markers > Reset average spectrum**.

Note If you have selected the option to synchronize the LTAS with the selection (see below), you do not have to perform this step. Simply click anywhere in the spectrogram.

Synchronize the long-term average spectrum with selection:

By default, the LTAS is not synchronized with what you have selected in the spectrogram. Every time you change your selection, you have to press the LTAS button or select **Markers > Set**

average spectrum to selection to update the LTAS.

You can change this setting so that the application updates the LTAS every time you modify what you have selected in the spectrogram. You have two options to set this:

- Select **Markers > Sync average spectrum with selection** OR
- On the **Settings** window, select **Pointer & Selection** and then select/de-select **Sync Average Spectrum with Selection**

3.1.10. Markers

NOTE

Available in VoceVista Video Pro

Markers are only available in *VoceVista Video Pro*.

Markers enable you to bookmark sections within files and label them. This makes it easy to:

- mark and find sections of long recordings
- organize large collections of recordings by using labels and point markers
- move quickly to marked sections within a recording
- analyze and compare marked sections in a recording




Figure 41. A recording with markers

Figure 41, “A recording with markers” is an example of a recording with markers that have been

added and labelled to identify different sections within the file. The labels appear above the spectrogram and in the **Files / Markers** list on the right side of the user interface. This list of markers lets you easily see the contents of the recording and quickly move to that section, as well as analyze and compare marked sections in a recording.

Marker Types

! File Description

The application automatically adds a File Description marker to every recording. You can use it to add any information about the recording. This marker always appears first in the marker list of a file, and it has the ! icon.

To add information to this marker, select it in the **Files/Markers** list, and then enter text in the box that appears below the list of markers.

● Auto Markers

An Auto Marker is created every time you press record, record for a while, and then press stop.

Auto Markers have a circle icon ● on the Marker List. Each segment is called a “take,” and takes are numbered. You can change the take descriptions later by selecting them and editing the text that appears in the box below the list of markers (see below).

By default, 0.5 seconds of silence are added after each take. You can change this value in the [Advanced Settings](#).




Figure 42. Auto Markers on the Timeline and the Marker List

■ Range Markers

You can add Range Markers to identify sections within a file. This can be useful when you have a long recording and want to be able to easily find that section later. You can also move easily

between sections you have marked by selecting them in the Marker List or compare sections by checking their associated boxes in the Marker List.

To add a new Range Marker, [select a time period](#), and then click **Markers > Add marker**. Range Markers appear with a rectangular icon  on the Marker List. The marked section also has a rectangle that appears above the signal in the spectrogram and waveform that matches the Range Marker color and description.




Figure 43. Range Markers on the Spectrogram, Timeline and the Marker List

▼ Point Markers

Use Point Markers to mark a point in time in a recording. To add a new point marker, click the Timeline or Spectrogram to position the Time Cursor and reset the selection. Then click **Markers > Add marker** (or right-click the Spectrogram and select **Marker**).

Standard Point Markers have the blue triangle icon. They can also have the “Thumb up” (green) or “Thumb down” (red) icon. There is no difference between the Thumb Markers and regular Point Markers except for the icon.




Figure 44. Range and Point Markers on the Timeline and the Marker List

Navigating with Markers

Markers make it easy to move between marked sections of a recording.

You can move to the next or previous marker through the menu. Click **Markers > Jump to previous marker** or **Jump to next marker**. You can also use the Q and W keys instead. If you press **Shift+Q** or **Shift+W**, you 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 also click any marker in the Marker List to jump to it, or double-click it to select its time range and zoom to it.

Editing Markers




Figure 45. Editing Markers in the Files / Markers List

You can edit the text and color of a marker in the lower half of the **File / Marker** list. If this window is not visible, click **Markers > Show File / Marker List**.

To change the time range of a marker:

- Click and drag the marker label that appears in the Spectrogram or Timeline and then double-click the marker label
- Click and drag the edges of a Range Marker to your (then click and drag the marker label to match)

Setting Marker Visibility

You can toggle the visibility of individual marker types by selecting one of the **Show auto/point/range markers** or **Hide markers** commands from the **Markers** menu.

Searching Markers and Files

To search for file names and marker text, enter a term into the <search Markers> box at the top right corner of the **Files / Markers** list.

There are three radio buttons at the bottom of the **Files / Markers** list that enable you to search:

- the current file
- all open files
- all application files on your local hard drive

Recent file history

If you haven't entered a search term, the **Files / Markers** list shows a list of the most recently opened files in reverse order (the last file you opened appears first). Simply double-click an entry in the list to open that file.

3.1.11. Statistics Pane

NOTE

Available in VoceVista Video Pro

The Statistics Pane is only available in *VoceVista Video Pro*.

The **Statistics** pane enables you to view and export the following types of numerical data about your recordings:

- pitch (frequency Hz)
- vibrato (rate and extent)
- harmonics (number, note, frequency Hz, intensity dB)
- spectral peaks (number, note, frequency Hz, intensity dB)
- EGG (contact quotient, dEGG max, dEGG min)

You can view this data for a point in time, or, if you have created a marker and selected it in the **Files / Markers** window, you can view the mean values for each of these properties for that period of time.

The following image contains the information about the harmonics at a point in time in the recording displayed in the **Statistics** pane.




Figure 46. Statistics Pane

To view the statistics for a recording:

1. Select a point in time in the spectrogram or select a marker on the **Files/Markers** pane.
2. Select **View > Statistics** to view the **Statistics** pane, if not visible.
3. Select one or more of the statistics you want to view from the dropdown checkbox that appears near the top of the **Statistics** pane (pitch, vibrato, harmonics, spectral peaks, EGG).
4. Either (optional step)
 - Click the **Copy statistics result to clipboard** icon at the bottom of the **Statistics** pane and paste the data into another application
 - OR
 - Click the **Open export data dialog** icon at the bottom of the **Statistics** pane to open a dialog box where you can specify in greater detail the data which will be exported as a file in Microsoft Excel format.

3.1.12. Vibrato View

NOTE

Available in VoceVista Video Pro

The Vibrato View is only available in *VoceVista Video Pro*.




Figure 47. VibratoView

The Vibrato view option lets show and hide the vibrato view pane. If the file has a measurable **vibrato**, this pane will display the vibrato rate and/or extent, depending upon whether either is selected.




Figure 48. VibratoView

When enabled, the Vibrato view appears in a separate pane and displays the vibrato rate and/or extent using the settings you have selected on the dialog box.

3.1.13. EGG (Electroglottograph)

NOTE

Available in VoceVista Video Pro

The EGG display is only available in *VoceVista Video Pro*.

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.

3.1.14. 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 Icon Size](#) under Advanced Settings.

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

image::images/ref-ui-toolbar/toolbar-small.png[width=153] + A minimalistic toolbar for absolute beginners. Only contains buttons for recording and playback.

Single Row Toolbar

image::images/ref-ui-toolbar/toolbar_single_row.png[width=1190] + 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

image::images/ref-ui-toolbar/toolbar_standard.png[width=1190] + 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

image::images/ref-ui-toolbar/toolbar_full.png[width=1190] + A two-row toolbar with most buttons that are available, including controls to set the frequency scale range.

3.2. Settings

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

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

3.2.1. Recording Settings




Figure 49. Recording Settings

The Recording Settings page controls how VoceVista captures audio from your sound card: which device to record from, the input level, the audio format, and where the data is stored on disk.

Input Device

The audio line from which to record. If you have more than one sound card installed — for example, when using an external USB microphone — make sure the correct one is selected here.

First Channel

Selects which physical input channel of a multi-channel audio device is used as the first channel of the recording. This control is only shown for devices with more than two inputs.

For example, if your audio interface has eight inputs and you are recording in mono, this lets you choose whether channel 1, 2, 3, ... or 8 is the source. When recording in stereo, the selected channel becomes the left channel and the next channel becomes the right.

Recording Volume

Sensitivity of the selected microphone or input line. Adjust the slider so the level meter peaks high but never clips — clipping introduces artifacts in the spectrogram.

CAUTION

A clipped signal cannot be repaired in software. If the meter's clip indicator lights up while recording, lower this volume and re-record.

Recording source and volume can also be set through the input level meter on the toolbar. Drag the slider to change the volume, or right-click the meter to switch input source:




Figure 50. Right-clicking on the input volume meter on the toolbar

Enable ASIO (Windows only)

When checked, VoceVista uses ASIO drivers for audio input and output. ASIO offers lower latency than the default Windows audio backend, but it requires ASIO drivers provided by your audio interface manufacturer. Leave this off if you do not have an ASIO-capable device.

This control is hidden on macOS and Linux, where Core Audio and ALSA already provide low-latency audio.

Quality

NOTE

The settings under *default settings for new recordings* can only be changed before any audio has been recorded into the current file. The format of an existing recording cannot be retroactively changed. To choose a new format, click **File** > **New** first.

Sampling Rate

The number of amplitude measurements (samples) per second taken from the recording source. The sampling rate determines the highest frequency that can be measured: a digital recording cannot accurately resolve frequencies above half its sampling rate. For example, at 11025 samples per second the analyzer can measure up to around 5500 Hz.

For most purposes a sampling rate of 44100 Hz is appropriate. When using the EGG, 48000 Hz often gives the best results.

Channels

Determines whether new recordings are stored internally as mono or stereo, and how the channels of the active recording are displayed. The left side of this section, *default settings for new recordings*, applies to recordings made from now on. Changing it does not affect files that are already open. The right side, *settings currently in use*, controls how the active file is displayed and can also be changed from the toolbar.

===== 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

Useful when a single mono microphone is connected to one channel of a stereo audio interface. Recording "mono" in that situation would mix the live channel with the empty channel that has nothing connected. Recording only the left or only the right channel avoids that mixing. + The result is saved as a mono file with a single channel.

Record Audio left; EGG right / EGG left; Audio right

Record a stereo file in which one channel is interpreted as the audio signal from the microphone and the other as the Electroglottograph (EGG). + Available in *VoceVista Video Pro*.

===== settings currently in use

Display mono

When a mono recording is active, only its single channel is available to display.

Display mixed stereo

Display both channels of a stereo recording mixed together. The spectrogram then shows audio content from both channels.

Display left / right stereo channel

Display only the left or right channel of a stereo recording.

Display Audio left, EGG right / EGG left, Audio right

Interpret one channel as audio data and the other as EGG.

Sample Size

The number of bits used per sample. A higher bit depth increases dynamic range and storage size.

- **16 bit** — standard CD quality, sufficient for direct analysis and the smallest on disk.
- **24 bit / 32 bit** — useful when the recording will be edited or amplified later, or when the input has a very high dynamic range and contains quiet passages that need to be amplified after the fact.

Temp Folder

The folder where VoceVista stores incoming audio data while you record. Make sure the chosen drive has enough free space.

The label above the path field shows how much free space is available on that drive and how long that translates to at the current recording format. As a reference, a mono recording at 44100 Hz / 16 bit needs about 350 MB per hour (roughly half a CD).

Click the ... button next to the path to browse for a different folder.

3.2.2. Playback Settings




Figure 51. Playback Settings

The Playback Settings page controls the instruments and volumes used for audio playback in VoceVista, and selects the wave and MIDI devices through which sound is sent.

Playback

This section sets the instruments used by the piano and the note sliders, and the individual volumes for each playback source. Each volume slider is independent, and the Master Volume slider scales all of them together.

Piano Instrument

Instrument for piano keys and for the **fundamental** of each note slider. You can choose the **Sine Wave Generator** (the first entry in the list) or any of the standard MIDI instruments:




Figure 52. Piano and Overtone Instruments

The Sine Wave Generator plays pure sine waves with high frequency accuracy. MIDI instruments play through the MIDI output device and may have less precise pitch, but provide a richer timbre.

The slider next to this combo box sets the volume of the piano / fundamental playback.

Harmonic Instrument

Instrument for the **harmonics and subharmonics** of the note sliders. Same instrument list as above (Sine Wave Generator or MIDI). The slider next to this combo sets the volume for harmonic and subharmonic playback.

Using the Sine Wave Generator for harmonics is recommended for accurate overtone tuning practice; MIDI instruments can sound richer but their overtones may not align cleanly with the fundamental.

Recording

Volume for playing back recorded audio content (the audio that was captured into the file). Adjust this slider relative to the piano and harmonic volumes when you want the recording louder or quieter against the played note sliders or piano keys.

Vowel Chart / Generator

Volume for the live audio produced by the Vowel Chart and the standalone Sound Generator. These tools synthesize sound on the fly as you move the mouse over them; this slider controls how loud that synthesized sound plays.

Master Volume

Master volume for all audio playback. This slider scales every other volume on this page (Piano, Harmonic, Recording, Vowel Chart) by the same factor, so you can quickly raise or lower overall output without disturbing the relative balance you have set up.

Loop playback

When checked, playback of the current selection — or of the whole file if there is no selection —

repeats in a loop. When unchecked, playback stops at the end of the selection or file.

Playback Speed

Slows down or speeds up playback of recorded audio. A value of 1.000 plays at the original speed; values below 1 slow it down, values above 1 speed it up. The **Reset** button next to the field returns the value to 1.000.

===== use pitch compensation

When enabled, the pitch is preserved when playback speed is changed. When disabled, slowing playback also lowers the pitch (and speeding it up raises it).

Pitch-compensated playback is useful for studying a fast passage at half-speed without transposing it. Uncompensated playback can be useful to transpose a recording up or down by changing the speed.

Devices

This section selects the audio devices through which VoceVista sends sound. Each device combo has a circular **Test** button to its right that plays a short sample so you can verify the device is working.

Wave Output Device

The audio device used to play back recordings and the Sine Wave Generator. If you have multiple sound cards or audio interfaces, make sure the desired output is selected here.

Click the test button to the right of the combo to play a short test sound through the selected device.

MIDI Output Device

The MIDI device used to play the piano and note slider sounds when a MIDI instrument (rather than the Sine Wave Generator) is selected above. If your system has no MIDI synthesizer, the choices are limited to whatever the operating system provides.

Press and hold the test button to the right to hear a short note from the currently selected MIDI instrument.

MIDI Input Device

If you have an external MIDI keyboard connected, select it here. Pressing keys on the external keyboard then has the same effect as clicking on the on-screen piano view: the corresponding note is played through the MIDI output device.

Leave this set to **No Input Device** if you do not use an external MIDI keyboard.

3.2.3. Analyzer Settings




Figure 53. Analyzer Settings

The Analyzer Settings page controls how VoceVista extracts information from the audio signal. It governs the frequency and time resolution of the spectrum and spectrogram, the parameters of the pitch detection algorithm, and the profile that bundles all of these settings together.

VoceVista runs two independent analyses on the signal. The Fast Fourier Transform (FFT) computes the **Spectrum** — the intensity of each frequency component. A separate pitch detector computes the **Fundamental Pitch**. The two methods are fully independent and have their own controls on this page.

NOTE

The sampling rate is set on the [Recording Settings](#) page. Most values on this page — frequency resolutions, time windows, pitch limits — depend on it: lowering the sampling rate can sharpen FFT frequency resolution, while the pitch detector works best at 44100 Hz or higher (and the EGG at 48000 Hz).

Spectral Analysis

The Spectral Analysis section configures the FFT and the time window over which it runs.

Frequency Resolution

The frequency resolution is the smallest difference (in Hz) between two frequencies that the analyzer can distinguish. Internally this is controlled by the **FFT size** — the number of points fed to each transform — but the dialog presents it as a frequency value in Hertz.

===== Spectrogram Resolution

The frequency resolution used by the spectrogram on the Long-Term view. The default is comfortable on a modern computer; very high values (below a few Hz) can become sluggish when scrolling, so lower it if dragging the time-range slider starts to feel laggy.

===== Spectrum Resolution

The frequency resolution used by the spectrum on the Short-Term view. The spectrum can run at a higher resolution than the spectrogram with little extra cost, because the spectrogram recomputes once per pixel column while the spectrum recomputes only at the cursor position.

Spectrum Time Window

The time window determines how much audio is folded into each FFT update — and therefore the trade-off between time resolution and frequency accuracy. A longer window improves frequency accuracy; a shorter window shows faster temporal changes.

The dropdown selects how you specify the window:

Updates per second

How many spectrum updates per second (e.g. 12). Larger values mean a shorter window and better time resolution.

Milliseconds

The duration of the analysis window in milliseconds (e.g. 83 ms).

Maximum Frequency Resolution

Automatically use the longest possible window for the selected frequency resolution. This pins the analyzer to the highest frequency accuracy your FFT size allows.

For example, with a 44100 Hz sampling rate and an FFT size of 8192 points, the maximum window is approximately 185.8 ms — about 5.4 updates per second.

===== link time windows

When checked, the spectrogram and the short-term spectrum share the same time window. Uncheck to set the spectrum time window independently from the spectrogram, which is useful when you want high temporal precision in the Short-Term view while keeping the spectrogram at a coarser window.

Spectrum Averaging

The short-term spectrum is averaged over this sliding time window during recording and playback to reduce flickering. Shorter durations make the spectrum more responsive to sudden changes but flicker more; longer durations are more stable but slower to react.

LTAS Overlap

The Long-Term Average Spectrum (LTAS) is computed by averaging many spectral slices taken across a selected time range. This control sets how many slices per second are taken. Higher values yield smoother LTAS results at the cost of more computation.

FFT Window Function

The FFT method requires its input to be filtered through a **window function** to avoid spectral leakage. A discussion can be found at en.wikipedia.org/wiki/Window_function.

For VoceVista the default **Dolph-Chebyshev** window works very well and gives a usable dynamic range of about 100 decibels. Do not change it unless you understand window functions and their respective trade-offs.

Pitch Detection

The Pitch Detection section configures the pitch algorithm. It runs in parallel with the FFT and produces the fundamental-pitch curve shown above the spectrogram.

Pitch Detection Time Window

The length of the audio segment fed to the pitch algorithm. The combo box shows both the number of samples and the resulting **minimum detectable frequency** (a 1024-sample window at 44100 Hz reaches down to about 43.1 Hz).

Longer time windows extend the algorithm's reach to lower frequencies but reduce how quickly it can track changes. The default is appropriate for most situations.

Frequency Range to Detect

===== Lowest Frequency

Frequencies below this value are ignored, which helps avoid false detections when the recording contains low-frequency rumble or mains-hum.

===== Highest Frequency

Frequencies above this value are ignored, which helps avoid false detections from high-frequency noise or from a strong harmonic that might otherwise be mistaken for the fundamental.

Minimum Clarity

The minimum confidence the algorithm must reach to accept a pitch result as valid, expressed as a percentage. Higher values reject ambiguous detections; lower values accept more — at the cost of more false positives.

Minimum Intensity

The minimum signal level (in dB) required for a pitch reading to count. Signals quieter than this threshold are ignored, even when their clarity passes the test above. This keeps background noise from producing spurious pitch tracks during silent passages.

Prefer Harmonic Fundamental

When checked, the pitch detector tries to find the **fundamental** frequency of harmonic sounds rather than reporting whichever harmonic is loudest. This is the right choice for analyzing musical instruments and the singing voice, where the fundamental is often weaker than its upper harmonics.

Profile

A **Profile** bundles every setting on this page into a named configuration that you can switch between in one click.

Default Profiles

VoceVista ships with four built-in profiles, tuned for different kinds of material:

Singing — Narrowband

The default. Narrowband spectrogram (8192-point FFT) optimised to show the harmonics of the singing voice. Pitch detection runs at 50 Hz across a 65 Hz – 1400 Hz range with prefer-harmonic-fundamental enabled.

Voice — Broadband

Same frequency settings as Narrowband, but with a much higher spectrogram update rate (300 / s). The shorter analysis window emphasises individual glottal cycles and formant transitions, at the cost of frequency resolution.

Speech

Same window as Narrowband, but with stricter pitch detection: the clarity threshold is raised to 85% and the frequency range narrowed to 65 Hz – 400 Hz, so vowel pitch is detected reliably and unvoiced segments are rejected.

Gongs / Singing Bowls

Ultra-narrowband: the FFT size is forced to its maximum (131072 points) for the highest possible frequency accuracy on long, inharmonic decays. Prefer-harmonic-fundamental is turned off so the pitch tracker follows the strongest partial rather than guessing at a missing fundamental.

Manage Profile

The **Manage Profile** button next to the profile dropdown opens a menu with the available actions:

Reload

Discard changes you have made on this page and reload the saved values of the currently selected profile.

Save

Save the current values back into the selected profile.

Rename / Copy

Rename the current profile, or save it under a new name to create a user-defined copy.

Delete

Remove the currently selected user-defined profile. Built-in profiles cannot be deleted.

Switching profiles applies all of that profile's settings at once, so profiles are the fastest way to retune the analyzer for a different kind of material.

3.2.4. Acoustic Analyzer View




Figure 54. Acoustic Analyzer View

This settings page controls the appearance and behaviour of the [Analyzer View](#): which displays are visible, how they are laid out, and what additional information is overlaid on them. Settings here do not affect **what** the analyzer measures (that is on the [Analyzer Settings](#) page) — only how the results are shown.

Long-Term view

Toggles the Long-Term view, which scrolls past as audio plays and can contain the spectrogram, the fundamental pitch, and several overlays. The checkboxes inside this group decide which of those layers are drawn.

Spectrogram

Show how frequency intensities change over time. Colors are mapped from intensity using the current colormap (see [Colormap Editor](#)).

Pitch

Mark the fundamental pitch on top of the spectrogram as a thin line.

Mean Pitch

Overlay the mean pitch line used by the vibrato analysis.

Peaks

Mark spectral peaks detected on the spectrogram.

Generator Resonances

Show the resonances of the active filter generator on the spectrogram.

Piano

Show a piano keyboard alongside the spectrogram's frequency axis. The current note is highlighted while audio plays.

Staff View

Show a musical staff alongside the spectrogram, with the current note rendered as a notehead.

Frequency Axis Orientation

vertical

Spectrogram has a vertical frequency axis and a horizontal time axis — the most common arrangement.

horizontal

Spectrogram has a horizontal frequency axis and a vertical time axis, i.e. time runs from top to bottom of the screen.

NOTE

The spectrogram orientation radio buttons are hidden by default, since most users keep the spectrogram in its standard vertical-frequency arrangement. Turn them on with **Show Spectrogram Orientation Buttons** on the [Advanced Settings](#) page.

Short-Term view

Toggles the Short-Term view, which shows the spectrum at the current cursor position. When both the Long-Term and Short-Term views are active, moving the cursor on the Long-Term view updates the Short-Term view.

Spectrum

Show the intensity of each frequency at the current time position.

Pitch

Mark the current fundamental pitch on top of the spectrum.

Average Spectra

Display the Long-term Average Spectrum (LTAS) of the current selection, and of any markers checked for comparison.

Unfiltered Spectrum

When a frequency filter is active, also draw the unfiltered spectrum for reference.

Generator Resonances

Show the resonances of the active filter generator on the spectrum.

Piano

Show a piano keyboard alongside the spectrum's frequency axis.

Staff View

Show a musical staff alongside the spectrum.

Frequency Axis Orientation

vertical

Spectrum has a vertical frequency axis and a horizontal intensity axis.

horizontal

Spectrum has a horizontal frequency axis and a vertical intensity axis — the more familiar "amplitude over frequency" plot.

below Spectrogram

When the spectrogram has a vertical frequency axis and the spectrum a horizontal frequency axis, this option places the spectrum stacked **under** the spectrogram instead of beside it. Useful when the time axis is the most important dimension and you want to use the screen's full width for it.

Smoothed Curves

A smoothed curve is a low-pass filtered version of a spectrum that emphasizes the spectral envelope rather than individual peaks. This group enables smoothing on the short-term spectrum and on average spectra, and configures how the smoothing radius is interpreted.

Short-term Spectrum

Display a smoothed version of the short-term spectrum.

Average Spectra

Display smoothed versions of the current LTAS and of marker average spectra checked for comparison.

Radius

The radius over which the source curve is averaged to produce the smoothed curve. The radio buttons next to the radius value set how the number is interpreted:

Relative (% of Pitch)

The radius is expressed as a percentage of the current fundamental frequency. The smoothing scales with pitch, which is the right choice for following formant shapes across a sung phrase.

Absolute (Hz)

The radius is a fixed number of Hertz, regardless of pitch. Useful when you want consistent smoothing in a frequency-band-of-interest rather than a pitch-relative one.

Spectrum Type

Four icon buttons choose the visual style of the spectrum:

- **Outline only** — single line tracing the spectrum.
- **Outline and Monochrome Fill** — outline with a single-color fill below.
- **Stacked Fill** — outline with a colored fill matching the colormap.
- **Peak Fill** — solid fill emphasizing the peaks.

Use gradient fill

When checked, the spectrum's filled area uses a transparency gradient that fades according to intensity. This is mostly cosmetic but can make overlapping spectra easier to read.

Normalize Spectrum to loudest intensity

When checked, the loudest spectral peak is shown at 0 dB and every other point is scaled relative to it. Unchecked, the spectrum shows absolute intensity values, which is what you usually want when comparing recordings or judging signal level.

Pitch and Cursor Text Display

These overlays put the current pitch and cursor information directly on the secondary (short-term) spectrum, so you can read the values without moving your eye away from the analysis.

Note name

Print the note name of the pitch under the current cursor position on the secondary spectrum.

Frequency

Display the frequency of the current pitch on the secondary spectrum.

Intensity

Display the intensity of the current pitch on the secondary spectrum.

Show current note on piano and staff view

Highlight the note at the cursor position as a pressed key on the piano view, and as a notehead on the staff view.

Label harmonics

Label each detected harmonic with its number (H1, H2, ...) on the spectrum.

3.2.5. Spectral Peaks Settings




Figure 55. Spectral Peaks Settings

This page configures the detection and display of spectral peaks on the spectrogram. A **spectral peak** is a local maximum in the spectrum — a frequency that stands out above its neighbours — and is typically a harmonic, a formant, or a resonance.

Display Peaks on Spectrogram

When checked, the strongest spectral peaks at each time position are drawn on top of the spectrogram. Uncheck to compute peaks without drawing them on the spectrogram (the [Statistics Pane](#) still lists them).

Color

The color used to draw peak markers on the spectrogram. Click the color button to change it.

Number of peaks

How many peaks to detect per analysis frame, from 0 to 12. Higher values reveal more spectral structure but clutter the display. The default is 2.

Min Intensity

Spectral peaks quieter than this threshold (in dB) are ignored. Raise it to suppress noise-floor peaks; lower it to find weaker partials. The default is -60 dB.

Frequency Range

The lowest and highest frequencies considered when detecting peaks. Peaks outside this band are skipped, which is useful for ignoring mains hum at the low end or noise above the signal's useful range at the high end.

3.2.6. Vibrato Display Settings




Figure 56. Vibrato Display Settings

This page controls the appearance of the [Vibrato View](#) pane and the parameters of the underlying vibrato analysis. The page is divided into three groups: what the Vibrato View shows, what is drawn on top of the spectrogram, and the detection thresholds used by the analyzer.

Vibrato view

The boxed group at the top of the page configures the Vibrato View itself. The top checkbox toggles the view entirely; when it is off the other controls in this group are disabled.

Rate

When checked, the Vibrato View plots vibrato **rate** — how many cycles per second the pitch oscillates. The color button to the right sets the curve's color, and the two numeric fields set the lower and upper bound of the rate axis (in Hz).

Extent

When checked, the Vibrato View plots vibrato **extent** — the depth of the pitch oscillation in cents. Same controls as Rate: color and from/to range for the axis (in cents).

Overlays on the spectrogram

These checkboxes draw vibrato-related lines on top of the spectrogram rather than inside the Vibrato View.

Show mean pitch

Draw the **mean pitch** line — the slowly-varying centre around which vibrato oscillates — on top of

the spectrogram. Useful for separating the melodic line from the vibrato modulation. The color button to the right sets the line color.

Show individual cycles

Mark the boundaries of each detected vibrato cycle on the spectrogram. Useful when tuning the detection thresholds or verifying that the analyzer is locking on to the intended cycles.

Highlight current cycle

When **Show individual cycles** is on and the [Statistics Pane](#) is active, highlight the boundaries of the cycle that the Statistics Pane is currently reporting on. Lets you see which exact cycle the displayed numbers refer to.

Detection thresholds

These numeric fields tune the underlying cycle-detection algorithm. The defaults are appropriate for the singing voice; you usually only need to touch them if the analyzer is splitting one vibrato into several, or merging cycles that should be separate.

Mean interval

The sliding time window used to estimate the mean pitch around each vibrato cycle. Wider windows smooth the mean line; narrower ones follow melodic motion more closely.

Min. period

The shortest cycle duration the analyzer accepts as a vibrato cycle. Oscillations faster than this are ignored — useful for rejecting tremor or jitter at the edges of a sustained note.

Min. extent

The smallest pitch swing (in cents) required for a cycle to count as vibrato. Raise to ignore narrow pitch wobble; lower to detect shallow vibrato.

Max. mean step

The largest jump in mean pitch (in cents) allowed between consecutive cycles. A larger jump terminates the current vibrato segment and starts a new one — which is what you want when a singer moves between sustained notes.

Time Resolution

The number of vibrato analysis steps per second. Higher values give finer time resolution at the cost of more processing.

Reset to defaults

Restore every setting on this page to its factory default.

3.2.7. EGG Display Settings




Figure 57. EGG Display Settings

This page controls the appearance and behaviour of the Electroglottograph (EGG) display, including the long-term EGG history, the short-term audio/EGG view, the polarity and timing of the signal, the Contact Quotient calculation, and the Wavegram colormaps.

NOTE

Available in VoceVista Video Pro

The EGG features are only available in *VoceVista Video Pro*. See [our products page](#) for a comparison of the different editions.

EGG History

The EGG History is the long-term EGG view, equivalent to the spectrogram for audio: it shows how EGG-derived measurements change over time. The checkboxes inside this group choose which curves and diagrams are drawn, and a color button to the right of each line choice sets that curve's color.

Contact Quotient (CQ)

Plot the Contact Quotient over time. The exact CQ method used is configured in the **Contact Quotient Determination** group below.

Close / dEGG Max

Plot the position of the dEGG maximum (the moment of fastest contact) for each glottal cycle.

Open / dEGG Min

Plot the position of the dEGG minimum (the moment of fastest de-contact) for each glottal cycle.

EGG Wavegram

Show the EGG Wavegram — a stacked plot of one EGG cycle per row, colored by amplitude with the colormap chosen below.

dEGG Wavegram

Show the same kind of stacked plot using the derivative of the EGG signal, which highlights the contact and de-contact events.

Mini Short-term EGG

Show a small short-term EGG view inside the EGG History area in addition to the main Short-Term view. This is convenient when you want to compare a few cycles in detail without giving up the wider EGG History.

Short-term Audio / EGG

The short-term EGG view shows a few periods of the EGG and the audio signal at the current cursor position. It is the EGG counterpart of the short-term spectrum.

Pick which signals to display:

Audio

Display a few periods of the audio channel.

Audio Derivative

Display the derivative (rate of change) of the audio signal.

Electroglottograph (EGG)

Display the EGG signal.

EGG Derivative (dEGG)

Display the dEGG, the derivative of the EGG signal. The dEGG emphasizes the moments of fastest contact and de-contact, which are often the cleanest landmarks in a glottal cycle.

Audio Delay

A time offset (in milliseconds) applied to the audio signal to compensate for the delay between EGG and microphone recordings. EGG electrodes pick up vocal-fold contact almost instantaneously, while microphone audio arrives slightly later because sound has to travel through the vocal tract and air to the microphone. Adjust this until the EGG and audio peaks line up the way you expect them to.

EGG on top

When checked, the EGG signal is drawn above the audio waveform on the Short-Term view. When unchecked, the audio is drawn on top. Pick whichever order makes the relationship you are studying easiest to see.

Signal

Invert displayed polarity

Invert the displayed polarity of either signal on the Short-Term view.

Audio

Flip the audio waveform vertically.

EGG

Flip the EGG signal vertically. EGG hardware sometimes outputs the signal with the polarity reversed (closing-as-downward instead of closing-as-upward); use this checkbox if the EGG curves point the wrong way for your hardware.

Smoothing radius

Smoothing applied to the EGG signal before its derivative is computed, expressed as a percentage of the cycle length. A small amount of smoothing helps suppress noise that would otherwise produce spurious peaks in the dEGG. Too much smoothing blurs out the very events the dEGG is meant to highlight.

Contact Quotient Determination

The *Contact Quotient (CQ)* is a measure of how much of each glottal cycle the vocal folds spend in contact. A CQ of 50% means contact for half the period; 25% means contact for a quarter of the period. Note that EGG cannot give a complete picture of vocal-fold closure — it can only indicate that contact has increased. The term *Contact Quotient* is preferred over *Closed Quotient* for that reason.

The two controls in this group choose which method is used.

Use dEGG peaks

When checked, the contact and de-contact instants are taken from the peaks in the dEGG. The contact instant is the dEGG maximum (steepest rise); the de-contact instant is the dEGG minimum (steepest fall).

When unchecked, the contact instants are determined by the **Criterion Level** (described below) instead.




Figure 58. Using dEGG maximum for glottal closure, and dEGG minimum for opening

Criterion Level

When **Use dEGG peaks** is unchecked, the Criterion Level draws an imaginary line across the EGG signal: parts above the line are counted as closed, parts below as open. The CL value is the threshold's height on the EGG signal, expressed as a percentage between the cycle's minimum and maximum.

A CL of 50% places the threshold halfway up; a CL of 35% places it lower so a larger portion of each cycle is counted as closed. The default is somewhere in between; the optimal value depends on the recording and the speaker.

You can also drag the criterion-level line directly on the EGG display.




Figure 59. Using Criterion Level at 50% for estimating glottal contact period




Figure 60. Using Criterion Level at 35% for estimating glottal contact period

The Criterion Level is generally less accurate than the dEGG-peaks method on a clean signal, but it is more **robust** to noise: the dEGG method requires sharp, identifiable peaks, which a noisy signal may not provide.

NOTE

A hybrid mode — close from dEGG max, open from Criterion Level — is useful when only the rising flank of the dEGG is clean. When **Use dEGG peaks** is unchecked but the criterion level is high enough that above-line is short, the result approximates this hybrid behaviour.

Wavegrams

EGG Wavegram Colormap

Choose the colormap used to color the EGG Wavegram. The dropdown lists the available colormaps; the same options are managed on the [Colormap Editor](#) page.

dEGG Wavegram Colormap

Choose the colormap used to color the dEGG Wavegram. Different colormaps suit the dEGG well because its values are signed (positive and negative); a diverging colormap can make the contact / de-contact direction more obvious.

Reset to defaults

Restore every setting on this page to its factory default. This affects all groups above (history toggles, short-term toggles, signal polarity and smoothing, CQ method and value, and wavegram colormaps).

3.2.8. Colormap Editor




Figure 61. Colormap Editor

The Colormap Editor configures how the analyzer maps signal intensity to color. You can pick a stock colormap, modify its control points, generate a new colormap from a mathematical function,

or override the linked UI colors (background, axes, ...) just for this colormap.

Selected Colormap

The combo at the top of the page selects the active colormap. Three buttons next to it manage the colormap library:

Rename / Copy

Save the current colormap under a new name. Useful for creating a user-defined copy of a built-in colormap before customizing it.

Delete

Delete the currently selected colormap. Built-in colormaps cannot be deleted.

Read-only

Mark the current colormap as read-only to prevent accidental changes or deletion.

Control Point Editor

The colored field on the left of the dialog shows which colors represent which intensity. Colors are interpolated between the control points marked along the gradient. You can move and edit the control points with the mouse, or create new ones by double-clicking into empty space:

Table 5. Colormap Editor Mouse Commands

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	Open context menu (set color or delete point)
Double-click empty space	Create new control point

NOTE	The lowest control point is always rendered as the background color of the Analyzer View. To change it, use the Fonts and Colors page or the Linked UI Colors tab below.
------	---

Space Evenly

Arrange all control points so they have the same distance from each other. Useful after you have changed the number of control points and want them spread cleanly across the gradient.

Reverse Order

Reverse the order of the control points so the high-intensity colors are at the bottom and vice versa. This inverts the appearance of the spectrogram and can be a quick way to audition a new color scheme.

Interpolation

How colors transition between control points:

None

Hard color boundaries between control points. Each segment is a single solid color.

Linear

Straight gradients between control points (the most common choice).

Smooth

Cubic interpolation. Slightly softer transitions than linear, useful on colormaps with widely spaced control points.

A related but separate setting, **Smooth Spectrogram Colors**, controls whether the spectrogram itself smooths between adjacent samples. The Interpolation combo here only governs the colormap gradient.

Displayed Dynamic Range

The range of intensities mapped to the colormap, in decibels (dB). 0 dB is the loudest possible sine wave at the recording's sample format (a 16-bit recording can theoretically reach about 96 dB of dynamic range). On the spectrogram it is usually better to **narrow** the displayed range to suppress noise and bring out the signal.

The same two values can also be set with the dynamic-range sliders on the toolbar — useful for tuning contrast on the fly.

Brightness (top level)

The saturation point of the colormap. Every intensity above this value is drawn at the colormap's maximum-intensity color.

Contrast (range)

The width of the displayed dynamic range. A wider range shows more detail at the cost of clarity; a narrower range makes loud features stand out and pushes quiet detail to the floor color.

Height Shading

use shaded relief

When checked, the spectrogram is rendered with a shaded-relief effect that treats intensity as elevation, similar to a terrain map. This emphasizes peaks and ridges in the spectrum at the cost of some absolute-intensity readability.

Generate

Instead of placing control points by hand, you can generate the entire colormap from a mathematical function. The controls in this group are disabled while **Use Function** is unchecked, and the **Control Points** group above is disabled while it is checked.

Use Function

When checked, generate the colormap from the function below instead of the manual control points. Switching this off again preserves the control points you previously placed.

Generator Function

The function used to produce the colormap. **Sinebow** combines sine waves on each color channel to produce a continuous rainbow.

Gamma

Gamma exponent applied to the generator function. Values below 1 brighten the colormap; values above 1 darken it.

Offset

Phase offset of the generator function. Shifts the colormap along the hue cycle, which lets you pick the "starting color" of the rainbow.

Range

The portion of the function's output range used by the colormap. Negative values reverse the color order — equivalent to clicking **Reverse Order** on a control-point colormap.

Linked UI Colors

This tab opens an editor similar to the [Fonts and Colors](#) page, but every change is local to the **current colormap**. For example, the spectrum background color is normally the global Analyzer View background; with Linked UI Colors you can override it (e.g. white) just for this colormap, and similarly tune the cursor lines, axis labels, and other UI colors so they read well against the chosen gradient.

Linked UI Colors are saved with the colormap, so a user-defined "high-contrast" colormap can carry its own background and axis colors without disturbing the global theme.

3.2.9. Scale Settings




Figure 62. Scale Settings

The Scale Settings page configures the frequency axis used by every analyzer view: its range, whether it is drawn linearly or logarithmically, how note names are written, and what guide lines are drawn behind the spectrogram. Most of these settings can also be changed from the toolbar.

Scale Range

The lowest and highest note (or lowest and highest frequency) shown on the piano and analyzer views. The combos pick a note; the adjacent spin boxes show the corresponding frequency in Hertz, which you can edit directly when you want a non-note boundary.

The scale range can also be changed without opening this dialog by clicking and dragging the frequency scale in the main view, or by scrolling the mouse wheel over it.

Frequency Display

linear

Frequencies are spaced evenly along the axis. Consecutive harmonics end up the same distance apart, but piano keys grow wider toward the high end.

logarithmic

Frequencies are spaced logarithmically. Piano keys are all the same size, harmonics grow more closely spaced higher up, and the scale matches how the ear perceives pitch — making it the preferred setting for analyzing music.

The example below illustrates the difference between linear (left) and logarithmic (right) over the same range, ~80 Hz to 1100 Hz (E2 to C6). The logarithmic view shows much more detail in the lower octaves.

Musical Notes

This group controls how note names are written on the piano keyboard, the staff view, and the labels of note sliders.

show note name

Display the nearest tempered tone next to each slider label, along with its difference from the slider's pitch in cents.

show frequency

Display the frequency of each slider label in Hertz.

Accidental

Sharp

Black keys on the piano are written using the sharp symbol ($\sharp$).

Flat

Black keys are written using the flat symbol ($\flat$).

You can also right-click the staff view to switch between sharp and flat.

Concert Pitch A

The frequency used as the reference for the note A4. Every other note frequency is calculated relative to this pitch, so changing it transposes the entire keyboard. The default is 440 Hz; baroque ensembles often use 415 Hz, and some classical orchestras tune to 442 or 443 Hz.

The **Reset to 440 Hz** button next to the spin box restores the standard concert pitch.

Notation System

Helmholtz Notation

Common in German-speaking countries. Notes range from **A2** to **C**, then **c–c5**. The case of the letter conveys the octave.

Scientific Notation

Common in English-speaking countries. Notes range from **A0** to **C8**, with the octave number after the letter.

Piano Keyboard Notation

Keys are numbered 1–88 as on a real piano (white **and** black keys). Mainly used by piano tuners.

Solfège

Notes are written **Do, Re, Mi, Fa, Sol, La, Si**. Common in Romance-language traditions.

Background Lines

Background lines are thin guide lines drawn behind the spectrogram and spectrum to help you

read frequencies at a glance.

draw background lines

When checked, background lines are drawn. When unchecked, the background of the analyzer view is left empty.

The combo to the right chooses what the lines represent:

Piano Notes

One line per tempered note, with octave lines emphasized.

Frequency Scale Ticks

Lines at the same positions as the tick marks on the frequency axis.

Fixed increments

Lines at a fixed Hz interval. The interval is set by the **line spacing** spinbox that appears below this combo, and is independent of the current notation or scale type.

The color of the background lines is set by the color button to the right of the combo. This color is part of the current colormap if the **Linked UI Colors** tab on the [Colormap Editor](#) is configured to override it, otherwise it follows the global setting on [Fonts and Colors](#).

Color by Frequency

When this option is enabled, the pitch line is drawn in a color that varies with frequency, using the colormap chosen below.

Pitch Line

When checked, the displayed pitch line is colored according to its frequency, using the **Frequency Colormap** below. This makes vibrato and pitch transitions more visually obvious than a single-color line.

Frequency Colormap

The colormap used to color the pitch line when **Pitch Line** above is checked. The same colormaps available on the [Colormap Editor](#) are listed here.

3.2.10. Note Slider Settings




Figure 63. Note Slider Settings

This page configures the appearance and behaviour of the [Note Sliders](#) (formerly called Overtone Sliders). Many of these settings are also available on the [sliders toolbar](#).

show sliders

When checked, sliders are drawn on top of the spectrogram and spectrum. When unchecked, the sliders continue to exist (and continue to play back) — they just are not drawn. Useful when you want to hear a slider arrangement without obscuring the visual analysis.

fixate sliders on screen

When checked, sliders stay in the same screen position as the recording scrolls past them — like a measuring overlay. When unchecked, each slider is bound to a specific time range of the recording, like the notes in a piece of music. Time-bound sliders can also be exported as a MIDI file.

Number of note sliders

The number of distinct sliders shown. Each slider can have a separate fundamental frequency.

The spin box can only set the slider count when sliders are **fixated on screen**. When sliders are time-bound, add and delete them via **Tools > Insert note slider at selection** and **Tools > Delete selected note sliders**.

Note Slider Labels

Each visible slider position can be labeled with a note name, a frequency in Hz, or both.

note name

Show the nearest tempered note next to each slider label, plus its distance from the slider's pitch in cents (e.g. **A3+12ct**).

frequency

Show the frequency of each slider label in Hertz (e.g. **220 Hz**).

hide middle labels

When checked, only the fundamental and the highest visible harmonic (and lowest visible undertone) of each slider show full note name and frequency. Middle harmonics show just their harmonic number, which keeps the display readable when many harmonics are visible.

Transparency

When checked, sliders are drawn semi-transparently so the spectrogram underneath shows through, which makes it easier to position them accurately. The spin box to the right sets the transparency percentage. The same value can also be adjusted from a slider on the status bar.

undertone octaves only

When checked, only undertones that are octaves of the fundamental are drawn. Useful when constructing scales that span many octaves.

Snapping

When you drag a slider along the frequency axis, it can **snap** to nearby landmarks. Hold **Alt** while dragging to invert the snapping decision temporarily; for the full set of keyboard modifiers (including slider-to-slider snapping) see [Note Slider Mouse Commands](#).

no snapping

Sliders move freely along the frequency axis.

snap to nearest musical note

Sliders snap to the frequency of the nearest **tempered tone**.

snap to nearest peak in spectrum

Sliders snap to the nearest spectral peak in the current recording at the current cursor position. Useful for accurately measuring prominent frequencies.

Editing

highlight note borders

When checked, the borders between time-bound notes are drawn with high contrast. This makes it much easier to grab a slider's start or end edge while editing.

Playback

These options control which slider tones are played, and when.

play sliders during playback

When checked, time-bound sliders are played as audio playback progresses past them — turning the slider arrangement into an audible melody on top of the recording. Only applies to time-bound sliders; fixated sliders are not affected.

play top labels

When checked, clicking the fundamental of a slider also plays its highest visible harmonic, so you hear the overtone melody together with the fundamental.

play bottom labels

When checked, the slider of the first harmonic (the fundamental) is played during slider playback alongside the upper harmonics.

play multiple sliders in stereo

When checked, multiple simultaneous sliders are split between the left and right output channels. With four sliders the left two are sent to the left channel and the right two to the right. When unchecked, all sliders mix to mono.

Practice

Two convenience options for practising single notes from a longer arrangement.

set single slider when playing tones in other view

When two views are active and one of them has only a single slider, clicking on a slider (or piano key) in the other view sets the fundamental of that single slider. This lets you "jump" the practice slider around without picking it up directly.

Piano sets active slider

When checked, pressing a piano key sets the active slider's fundamental to the corresponding frequency. Useful for quickly auditioning a known pitch on the active slider.

3.2.11. Fundamental Ranges

The Fundamental Ranges page configures the colored bands drawn behind the Staff View, used to mark the frequency ranges of choral voices and of singable overtones. You can pick from several built-in profiles, edit the included ranges, or add new ranges of your own.

Voice Range Profiles

The dropdown at the top of the page selects the active voice-range profile. Built-in profiles are read-only; user-defined copies can be edited, renamed, and deleted freely.

SATB

Shows the four classical choir voice ranges: **Soprano, Alto, Tenor, Bass**.

Standard Voice Range Profile

The four choir voices plus the range of overtones that can be sung in a controlled manner by consciously amplifying specific harmonics.

Overtones NG / L

The singable overtones split by technique used to produce them:

NG technique

Uses the entire oral cavity without the tongue. Suited to lower overtones.

L technique

Uses the tongue to divide the oral cavity into two chambers. Used to produce the higher overtones.

Managing profiles

Three buttons to the right of the profile dropdown manage the profile library:

Rename / Copy

Save the current profile under a new name. Built-in profiles must be copied before they can be edited.

Delete

Delete the currently selected profile. Built-in profiles cannot be deleted.

Read-only

Mark the current profile as read-only to prevent accidental edits.

A **visible** checkbox lets you hide the entire current profile from the Staff View without deleting it. Hidden profiles still appear in the profile dropdown so they can be restored later.

Editing the ranges

The property browser in the middle of the page lists every range in the current profile. Each range has its own group of editable properties below its name.

Inner and Outer values

Each range fades in from the background color over a short frequency band, holds at full color across the main band, and fades out at the upper edge.

Outer value (low / high)

The frequency at which the range **starts to fade in** from the background, and the frequency at which it has fully faded out at the top.

Inner value (low / high)

The frequency at which the range **reaches its full color**, and the frequency at which it begins to fade out at the top.

The Outer and Inner values together produce a soft-edged colored band that does not visually shout the boundaries — useful when the boundary is itself fuzzy (different sources disagree about the exact range of an alto, for example).

Color Type

Each range is drawn in one of two semantic colors:

Voice

Used for normal singing-voice ranges (Soprano, Alto, etc.).

Overtone

Used for ranges of singable overtones.

The actual RGB values for **Voice** and **Overtone** are configured on [Fonts and Colors](#), under **Colors / Staff View / Range Voice** and **Range Overtone**. Editing them there changes every range that uses that color in every profile.

Adding, removing, and reordering ranges

Four buttons next to the property browser manage the list of ranges within the current profile (only enabled for editable, user-defined profiles):

Add

Append a new range with default values. Edit it in the property browser to set the frequency boundaries and color type.

Remove

Delete the selected range from the current profile.

Move Up / Move Down

Reorder the ranges. Order matters when ranges overlap because later ranges are drawn over earlier ones.

3.2.12. Selection and Pointer




Figure 64. Selection and Pointer Settings

This page configures how the mouse pointer interacts with the spectrogram and time scales: how the **Selection** is set and protected, what cursor lines are drawn, and which keyboard modifiers reshape clicking and dragging.

Pointer Mode

Lock Selection

Selection Locked: Ctrl/Cmd-Click sets Selection

When the selection is locked, a single mouse click cannot change or clear it. To set the selection, hold the **Control** key (Windows / Linux) or **Command** key (macOS) and drag on the spectrogram or a time scale. Use this when you have set up a careful selection and do not want a stray click to wipe it out.

Selection Unlocked: Click sets Selection

When unlocked, a single click resets the selection. Drag on the spectrogram or a time scale to set a new one.

Always Show Crosshair

When checked, the crosshair follows the mouse on every view (the spectrogram, the LTAS, the staff view, ...). When unchecked, the crosshair is shown only on the Short-Term view (spectrum) and on the scales.

The four cursor-line toggles below this option choose which lines make up the crosshair.

===== Time Cursor Line

Show a vertical line through the cursor's time position when the crosshair is visible.

===== Frequency Cursor Line

Show a horizontal line through the cursor's frequency when the crosshair is visible.

===== Intensity Cursor Line

Show a horizontal line through the cursor's intensity value on intensity axes such as the spectrum.

===== time range on crosshair

When checked, small helper lines around the time cursor mark the **effective time range** of the spectrum and pitch window functions — i.e. how much audio on either side of the cursor actually contributes to the current measurement. Useful when tuning the analyzer's time window: you can see how much context is being folded into each update.

NOTE

Enabling this option also turns on the Time Cursor Line and Frequency Cursor Line, since the time-range markers are drawn on top of them.

Enable Hover Spectrum

By default, moving the mouse across the spectrogram has no effect on the short-term spectrum.

When **Enable Hover Spectrum** is checked, moving the pointer across the spectrogram, waveform, or timeline immediately updates the short-term view to that time position — so you get a continuously refreshing spectrum slice as you scrub. Enabling hover also turns on the crosshair and time cursor line; disabling them again turns hover off.

Long Term Average Spectrum

Sync Average Spectrum with Selection

When checked, the Long-term Average Spectrum (LTAS) automatically recomputes whenever the selection changes, so the LTAS always represents the current selection. When unchecked, the LTAS holds whatever range it was last computed for, until you ask for a new one.

Keyboard Modifiers when Clicking

These keyboard modifiers change the pointer's behaviour while clicking or dragging.

Command / Control

Holding **Command** (macOS) or **Control** (Windows / Linux) while clicking always sets the selection, regardless of whether the selection is locked. + Holding the same key while clicking on sliders or piano keys selects multiple elements at once.

Shift

Holding **Shift** while dragging the pointer slows the motion down for more precision. Useful for fine-grained moves of selection edges, sliders, or scales.

Alt

Holding **Alt** while dragging a [Note Slider](#) **inverts** its snapping. When the slider is set to snap to

musical notes or spectral peaks, it does not snap while **Alt** is held — and vice versa.

3.2.13. Keyboard Shortcuts

This page lets you view and modify keyboard shortcuts for every VoceVista command. To change a shortcut, click on it in the list and press the key combination you want to use. Duplicate assignments are automatically removed.

Profile management

The **Profile** bar on the right of the page saves and restores entire keyboard layouts. This is convenient when you switch between, say, a recording-and-editing layout and a teaching layout, or when you want to share your shortcut customizations with someone else.

Rename / Copy

Save the current profile under a new name. Built-in profiles must be copied before they can be edited.

Delete

Delete the currently selected profile. Built-in profiles cannot be deleted.

Read-only

Mark the current profile as read-only to prevent accidental edits.

Add Key / Clear Key

A single command can have more than one shortcut bound to it.

Add Key

With a command selected in the list, click **Add Key** to attach a new shortcut. The next key combination you press becomes the new binding. The original shortcut is preserved.

Clear Key

Remove the keyboard binding from the currently selected command.

Default shortcuts

Below are the commands that ship with a default keyboard shortcut, organized by the same categories that appear in the dialog. Commands that have no default binding are still listed in the dialog — open it to find them and assign your own.

Command	macOS	Windows
File		
New	Cmd + N	Ctrl + N
Open	Cmd + O	Ctrl + O
Close	Cmd + W	Ctrl + W
Loaded File Info	Cmd + Shift + I	Ctrl + Shift + I

Command	macOS	Windows
Save	Cmd + S	Ctrl + S
Save As	Cmd + Alt + S	Ctrl + Alt + S
Save All	Cmd + Shift + S	Ctrl + Shift + S
Save Slider Layout	Cmd + Alt + L	Ctrl + Alt + L
Export Image	Cmd + I	Ctrl + I
Export Data	Cmd + D	Ctrl + D
Screenshot Hotkey	F12	F12
Export Data Hotkey	F11	F11
Exit	Cmd + Q	Alt + F4
Edit		
Undo	Cmd + Z	Ctrl + Z
Redo	Cmd + Y	Ctrl + Y
Cut	Cmd + X	Ctrl + X
Copy	Cmd + C	Ctrl + C
Paste	Cmd + V	Ctrl + V
Delete selection	Del	Del
Delete All	Cmd + Del	Ctrl + Del
Crop to selection	Cmd + R	Ctrl + R
Select All	Cmd + A	Ctrl + A
Select Visible	V	V
Set selection	S	S
Reset selection	Esc	Esc
Lock selection	Shift + L	Shift + L
Normalize	N	N
Fade In	Cmd + Shift + [	Ctrl + Shift + [
Fade Out	Cmd + Shift +]	Ctrl + Shift +]
Reverse	Cmd + Shift + R	Ctrl + Shift + R
Replace with silence	Cmd + Shift + Space	Ctrl + Shift + Space
Insert silence	Cmd + Alt + Space	Ctrl + Alt + Space
Reinsert silence	Cmd + Alt + Shift + Space	Ctrl + Alt + Shift + Space
Combine mono to stereo	Cmd + Shift + C	Ctrl + Shift + C
Split stereo to mono	Cmd + Alt + I	Ctrl + Alt + I
Generate sound	G	G

Command	macOS	Windows
View		
Zoom in	1	1
Zoom out	2	2
Zoom to fit selection	3	3
Zoom out fully	4	4
Reset Amplitude Zoom	Cmd + 4	Ctrl + 4
Center on reticule	5	5
Timeline	Shift + F2	Shift + F2
Waveform	Shift + F3	Shift + F3
Staff View	Shift + F4	Shift + F4
Piano View	Shift + F5	Shift + F5
Status Bar	Shift + F6	Shift + F6
Marker List	Shift + F7	Shift + F7
Vowel Chart	Shift + F8	Shift + F8
Statistics Pane	Shift + F9	Shift + F9
Secondary Time Scales	Cmd + Shift + F2	Ctrl + Shift + F2
Frequency Scale	Cmd + Shift + F3	Ctrl + Shift + F3
Spectrum Intensity Scale	Cmd + Shift + F4	Ctrl + Shift + F4
Spectrogram Time Scale	Cmd + Shift + F5	Ctrl + Shift + F5
Reset Window Layout	Cmd + Shift + F1	Ctrl + Shift + F1
Toggle Language	Cmd + Shift + L	Ctrl + Shift + L
Reload Language	Cmd + Alt + Shift + L	Ctrl + Alt + Shift + L
Full Screen	Alt + Return	Alt + Return
Rulers		
Add Ruler	Shift + R	Shift + R
Delete All Rulers	D	D
Filters		
Add Filter	F	F
Delete All Filters	Shift + D	Shift + D
Enable All Filters	Cmd + Shift + F	Ctrl + Shift + F
Disable All Filters	Cmd + Alt + F	Ctrl + Alt + F
Enable Realtime Preview	Alt + F	Alt + F
Apply Filters Permanently	Cmd + Alt + Shift + F	Ctrl + Alt + Shift + F

Command	macOS	Windows
Filter Additively	Alt + =	Alt + =
Filter Subtractively	Alt + -	Alt + -
Create Spectrum Filters	Shift + F	Shift + F
Average Spectrum		
Set Average Spectrum	L	L
Reset Average Spectrum	Cmd + L	Ctrl + L
Sync Average Spectrum with Selection	Alt + L	Alt + L
Markers		
Add and edit Marker	Cmd + M	Ctrl + M
Add Marker	M	M
Add Range Marker	Shift + M	Shift + M
Add Thumb up Marker	Cmd + G	Ctrl + G
Add Thumb down Marker	Cmd + B	Ctrl + B
Delete Marker	Cmd + Shift + Del	Ctrl + Shift + Del
Go To Previous Marker	Q	Q
Go To Next Marker	W	W
Zoom To Previous Marker	Shift + Q	Shift + Q
Zoom To Next Marker	Shift + W	Shift + W
Merge Markers	Cmd + Shift + G	Ctrl + Shift + G
Search Markers	Cmd + F	Ctrl + F
Edit Current Marker	Cmd + E	Ctrl + E
Edit File Description	Cmd + Shift + E	Ctrl + Shift + E
Zoom to Current Marker	E	E
Tools		
Rewind to start of file	Cmd + Home	Ctrl + Home
Rewind to start	Home	Home
Rewind page	PgUp	PgUp
Rewind	Cmd + Left	Ctrl + Left
Play and Stop	Space	Space
Stop	Shift + Space	Shift + Space
Record	Cmd + Space, R	Ctrl + Space, R
Forward	Cmd + Right	Ctrl + Right
Forward page	PgDown	PgDown

Command	macOS	Windows
Forward to end	End	End
Forward to end of file	Cmd + End	Ctrl + End
Replay last take	Alt + R	Alt + R
Use Video	Cmd + Shift + U	Ctrl + Shift + U
Loop playback	Cmd + Shift + P	Ctrl + Shift + P
Play Selected Tones	Return	Return
Stop Playing Tones	Backspace	Backspace
Deselect all Tones	Cmd + Backspace	Ctrl + Backspace
Profile Manager	Cmd + P	Ctrl + P
Options	F10	F10
Volume		
Increase input volume	Cmd + Up	Ctrl + Up
Decrease input volume	Cmd + Down	Ctrl + Down
Increase output volume	Cmd + Shift + Up	Ctrl + Shift + Up
Decrease output volume	Cmd + Shift + Down	Ctrl + Shift + Down
Window		
New Window on Current Selection	Cmd + Shift + N	Ctrl + Shift + N
Tile horizontally	Cmd + Shift + H	Ctrl + Shift + H
Tile vertically	Cmd + Shift + V	Ctrl + Shift + V
Maximize current window	Cmd + Shift + M	Ctrl + Shift + M
Close all	Cmd + Shift + W	Ctrl + Shift + W
Close all but active	Cmd + 0	Ctrl + 0
Activate Next Window	Cmd + Tab	Ctrl + Tab
Activate Previous Window	Cmd + Shift + Tab	Ctrl + Shift + Tab
Note Sliders		
Show Sliders	Alt + O	Alt + O
Select All Sliders	Cmd + Shift + A	Ctrl + Shift + A
Insert Note Slider	Alt + S	Alt + S
Use Note Transcription Tool	T	T
Fix Sliders on Screen	Alt + X	Alt + X
Display		
Toggle Spectrum	Alt + 1	Alt + 1
Toggle Spectrogram	Alt + 2	Alt + 2

Command	macOS	Windows
Toggle Average Spectrum	Alt + 3	Alt + 3
Toggle Normalized Spectrum	Alt + 4	Alt + 4
Disable Analyzer Display	7	7
Show Spectrogram	8	8
Show Spectrum	9	9
Show Spectrogram and Spectrum	0	0
Linear Scale	Cmd + 1	Ctrl + 1
Log Scale	Cmd + 2	Ctrl + 2
Show Notename	Cmd + 5	Ctrl + 5
Show Frequency	Cmd + 6	Ctrl + 6
Use sharp accidental	A	A
Transparent Sliders	Cmd + 7	Ctrl + 7
Snap to tempered	Cmd + 8	Ctrl + 8
Draw tempered lines	Cmd + 9	Ctrl + 9
Track current note	Cmd + T	Ctrl + T
Mini EGG	Cmd + Alt + E	Ctrl + Alt + E
Animate Mini EGG	Cmd + Alt + Shift + E	Ctrl + Alt + Shift + E
Settings		
Recording Settings	F2	F2
Playback Settings	F3	F3
Analyzer Settings	F4	F4
Analyzer View	F5	F5
Colormap Editor	F6	F6
Scale	F7	F7
Sliders	F8	F8
Profiles		
Select Profile 0–9	Cmd + Shift + 0 ... 9	Ctrl + Shift + 0 ... 9
Select Profile 10–12	Cmd + Alt + 0 ... 2	Ctrl + Alt + 0 ... 2

3.2.14. Fonts and Colors

This page customizes the global fonts, colors, and line widths used across the VoceVista interface — the analyzer view, the spectrum, the spectrogram, the staff view, and so on. Use it to tune contrast against a particular spectrum colormap, or to make the display friendlier for projection or screen-sharing.

Editing properties

Properties are organized in a tree at the center of the page, with two top-level groups:

Fonts

The fonts used by the various views (axis labels, slider labels, staff view notes, ...). Click a font row to change its family, size, and bold / italic style.

Colors

Every recolorable element (background colors, axis colors, cursor lines, voice-range bands, ...). Click a color row to open the color picker.

Some color rows also expose a **line width** sub-property that controls how thick the corresponding line is drawn.

Expand a category to see its members; the tree remembers which categories were last expanded.

Profile management

The **Profile** bar on the right of the page stores complete fonts-and- colors layouts. Profiles let you switch between different visual themes (e.g. a high-contrast theme for projection vs. a softer theme for daily use) without having to re-edit every property.

Rename / Copy

Save the current profile under a new name. Built-in profiles must be copied before they can be edited.

Delete

Delete the currently selected profile. Built-in profiles cannot be deleted.

Read-only

Mark the current profile as read-only to prevent accidental edits.

Global colors vs. colormap-linked colors

Settings on this page define the **global** fonts-and-colors profile — they apply unless something else overrides them.

A colormap can override specific UI colors using the **Linked UI Colors** tab on the [Colormap Editor](#). For example, a "White background" colormap can pin the spectrum background to white no matter which fonts-and-colors profile is active. When such an override exists, the colormap's value wins for as long as that colormap is selected.

NOTE

Fonts can only be set globally — the Linked UI Colors tab does not override font choices, only colors and line widths.

3.2.15. Advanced Settings

The Advanced Settings page collects fine-grained controls intended for power users. Settings are

organized into a tree of categories you can expand or collapse on the left of the dialog. Some settings — language, audio backend choice — only take effect after restarting the program.

NOTE

A few categories are conditional. The **Statistics Panel** category appears only in *VoceVista Video Pro*. The **Audio Backend** row appears only on macOS, and the **Enable DirectSound input**, **Enable Loopback recording**, **Enable ASIO driver**, and **Graphics Performance** entries appear only on Windows.

System

Reopen last session at startup

When checked, VoceVista remembers the files that were open at shutdown and reopens them on the next launch.

Number of uncompressed files to cache

When loading a compressed file (Mp3, Ogg, ...), VoceVista first decompresses it into a temporary uncompressed file for fast scrolling and editing. This setting limits how many decompressed files are kept in the temp folder. Reopening a file that is still cached skips decompression entirely, which is much faster.

Cache Folder

A drop-down with two commands. The label also shows the current size of the cache folder in parentheses (or *empty* if nothing is cached).

Clear Cache Folder

Empty the decompressed-file cache. Useful if disk space is tight.

Show in Explorer / Show in Finder

Open the cache folder in the OS file manager.

Logging Level

The amount of detail written to the log file: **None**, **Errors**, **Warnings**, **Messages**, **Events**, or **Debug** (most verbose). Raise it when troubleshooting; lower it again to keep the log file small.

Language

Set the UI language. Changing the language requires an application restart for every string to update.

Encoders

Default File Type

The file format that **File > New** uses, and the default format suggested in save dialogs.

MP3 Compression Quality

The MP3 LAME-encoder preset used when exporting MP3 files:

Insane (320 kbps)

Highest quality, largest files.

Extreme (200–240 kbps)

Near-transparent quality.

Standard (170–210 kbps)

Default. Good quality / size balance.

Medium (150–180 kbps)

Smaller files, audible compression artefacts on demanding material.

Display

Show Timeline on top

Show the timeline above the spectrogram instead of below it.

Waveform

===== Waveform uses Decibel

When checked, the waveform and timeline use a logarithmic (dB) intensity scale instead of a linear one.


===== Use symmetric waveform

When checked, the waveform and timeline draw intensity symmetrically above and below the zero line. Unchecked draws the actual signed sample values.

Scales

===== Use vertical scale labels

When checked, scale labels along the Y axis are drawn vertically. Vertical labels save horizontal space; horizontal labels are easier to read. The setting is also accessible by right-clicking the actual scale.

	
Y-Scale with horizontal text	Scale with vertical text

===== Show only scale units

By default the scales include the name **and** unit (e.g. "Time (s)", "Frequency (Hz)"). When checked, only the unit is shown ("s", "Hz").

===== Show all secondary time scales

Show the time scale on every secondary view (Waveform, EGG, Vibrato), not only the main spectrogram.

===== Show Timeline Amplitude Scale

Show the amplitude axis on the Timeline view.

Markers

These visibility toggles refine which markers are drawn where, on top of the per-marker-type visibility set in the [Markers](#) menu.

===== Draw marker heads on timeline

When checked, marker labels appear on the timeline. When unchecked, only the marker boundaries are drawn there.

===== Show markers above spectrogram

When checked, markers are drawn above the spectrogram (subject to the per-type visibility settings in the **Markers** menu).

Initial Time Range (in seconds)

The time range used for new empty documents. Default: 10 seconds.

Fixed Timeline Range (in minutes)

By default, the timeline range starts at 10 seconds and grows as the recording grows. Setting a value here pins the timeline to that length at all times.

Initial number of note sliders

The number of Note Sliders used for new empty documents and for existing documents loaded for the first time.

Spectrogram Scroll Mode

What the spectrogram does when the time cursor reaches the end of the visible range during recording or playback.

Scrolling

Keep the time cursor at a fixed position and slide the spectrogram to the left.

Paging

Switch to a fresh empty page and let the time cursor advance across it.

Show Spectrogram Orientation Buttons

Show the small icons that switch the spectrogram between vertical and horizontal frequency axes (the same toggle is on the [Acoustic Analyzer View](#) page).

Decibel distance for background lines on spectrum

Draws horizontal background lines on the spectrum at this dB interval. The default of 10 dB is a useful coarse grid.

Decibel range of normalized spectrum

The dynamic range shown on the spectrum when **Normalize Spectrum** is enabled (see [Acoustic Analyzer View](#)).

Toolbars are floatable

When checked, toolbars can be undocked and floated as separate windows. When unchecked, they are pinned to their docking positions.

Toolbar Icon Size

The size of toolbar icons: **Extra Small**, **Small**, **Normal**, **Large**, or **Extra Large**. Useful on high-DPI displays where the default size feels too small.

Show context menu titles

Show category headers inside right-click context menus, which clarifies what each section is for at the cost of menu height.

Statistics Panel

NOTE

Available in VoceVista Video Pro

These settings only have an effect in *VoceVista Video Pro*, which is the edition that ships the [Statistics Pane](#).

Statistics Panel updates/second

How often the Statistics Pane recomputes its measurements during playback or recording. Lower this on slow computers.

Show full spectrum time range

When checked, the Statistics Pane computes its values across the full spectrum-analysis time window rather than just the cursor instant.

Recording

Insert silence after each take (in seconds)

The number of seconds of silence appended to the file each time recording is stopped. Useful for keeping takes visually separated on the timeline.

Automatically start recording when program starts

When checked, the program starts recording into a new empty file at launch. The same effect can be triggered by passing `/record` on the command line.

Clear recording after this many seconds

Allows the program to run continuously without filling the disk: when the recording reaches this length, it is discarded and a new recording starts. Useful for long unattended monitoring sessions.

Close unsaved files without confirmation

When checked, closing an unsaved file does not prompt for confirmation. Saves time when working with throwaway recordings; risky when working with takes you care about.

Open Web Links

Auto paste weblinks when opening web link dialog

When checked, opening the **Open Web Link** dialog automatically pastes a URL from the clipboard if one is found. Saves a click when you copy a link before opening the dialog.

File Export

Show exported files in Explorer / Finder

When checked, the OS file manager opens to show the exported file after each export. Unchecked keeps focus on VoceVista.

Signal Generator

The Signal Generator plays a test signal over a sound output, used mainly to measure the resonance of the oral cavity by playing a known signal into it and analyzing the echo.

Play signal while recording

When checked, the test signal plays continuously during recording.

Wave Output Device

The output device used for the test signal. Pick a separate device from your main playback device when you want the test signal in one place (e.g. a small speaker pointed at the singer) and the recording playback elsewhere.

Frequency (Hz)

The frequency of the test signal in Hertz. The current implementation plays a square wave at maximum amplitude.

Signal Phase Ratio

The ratio of the up phase to the total period of the square wave. 0.5 produces a symmetric signal;

smaller values produce a short up spike followed by a longer down period — useful for varying the harmonic content of the test tone.

Cursor Lines

Show Playback Cursor Line

Show the green vertical line that indicates the current playback position. Turning it off is mostly useful when taking screenshots.

NOTE

The visibility of the time, frequency, intensity, and time-range cursor lines moved to the [Selection and Pointer](#) page in a recent release. Set them there.

Vowel Chart

These settings control the appearance and behaviour of the [Vowel Chart](#).

Show Vowel Chart Toolbar

Show the small toolbar at the top of the Vowel Chart with quick toggles for the most common chart options.

Show Harmonic Matrix

Overlay a grid of harmonics on the vowel chart for reference.

Show Harmonic Vowels

Annotate the chart with the vowel positions corresponding to each harmonic.

Show note names for f1/f2

Annotate the chart axes with note names for the first two vocal-tract resonances (fR1 and fR2).

Scale Mode

Whether the fR1 / fR2 axes use a **linear** or **logarithmic** scale.

Lock aspect ratio

When checked, the fR1 / fR2 axes are forced to the same scale, which keeps vowel-chart shapes recognizable when the window is resized.

Play Vowels

When checked, hovering over a vowel position in the chart synthesizes the corresponding sound through the [Wave Output Device](#).

fR1 / fR2 Min and Max (Hz)

Four separate fields — **fR1 Min (Hz)**, **fR1 Max (Hz)**, **fR2 Min (Hz)**, **fR2 Max (Hz)** — set the lower and upper bounds of the two resonance axes on the vowel chart.

Audio Performance

Audio Backend (*macOS only*)

Choose which low-level audio API VoceVista uses for capture and playback. The default is the cross-platform PortAudio backend; RtAudio and Qt Audio are available as fallback options when a device misbehaves under PortAudio. Changing the backend takes effect after restarting the program.

Enable DirectSound input (*Windows only*)

When checked, the DirectSound API is available for input. DirectSound performs automatic sample-rate conversion in software, which is convenient when the device's hardware rate doesn't match the recording rate.

Enable Loopback recording (*Windows only*)

When checked (and WASAPI support is available), you can record "what you hear" by selecting your output device as the input source. Disabling Loopback hides such devices from the input list.

Enable ASIO driver (*Windows only*)

When checked, the ASIO API is available for low-latency input and output. The same toggle exists on the [Recording Settings](#) page.

Recording Latency (ms)

The approximate time between a sound arriving at the microphone and its appearance on the screen. Lower values mean the analyzer feels more responsive; values that are too low can cause clicks or gaps in the recording when the computer cannot keep up.

File Playback Latency (ms)

Same as the recording latency, but for playback. Lower values reduce the lag between actions (changing playback position, toggling a filter) and the audible response. Values that are too low can cause audible clicks during playback.

Fade-In/Out Duration (ms)

The fade-in / fade-out time applied at the boundary of selections and takes during playback. A small fade prevents clicks at edit points; too much fade audibly softens the transitions.

Disable sleep for recording and playback

When checked, the operating system is prevented from sleeping while VoceVista is recording or playing. Long unattended recordings won't be cut short by power-management policies.

Frequency Filters

Filter Size

The FFT size used by the frequency-filter engine. Larger sizes give sharper filter cut-offs at the cost of higher latency. The values are powers of two from 1024 up to 65536.

Tone Generator

These settings affect only the Sine Wave Generator that plays overtones and piano keys.

Sampling Rate

The sampling rate at which the tone generator synthesizes samples.

Playback Latency (ms)

The latency of the tone generator's output. A lower value makes the program feel more responsive when you slide a slider's pitch in real time.

Input Level Meter

The Input Level Meter is the slider on the toolbar that shows input volume and the current input signal strength.

Enable monitoring

When the Input Level Meter is enabled, VoceVista continuously records from the input device so the meter can move. Turn this off to suspend that recording.

The same toggle is on the right-click context menu of the level meter on the toolbar.

Displayed range (in dB)

The dynamic range displayed on the meter.

Max. updates per second

How often the Input Level Meter is redrawn. Lower this if the meter takes too much CPU on a slow computer.

Peak History (seconds)

The duration over which the peak indicator is held. With a 3-second peak history, the meter shows the maximum value seen in the last 3 seconds before decaying.

Graphics Performance (*Windows only*)

MSAA Samples

The number of multisample-antialiasing samples used when rendering the spectrogram and other visualizations. Higher values smooth edges at the cost of GPU work.

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 is already running, click on **File › Open** to open an existing file. You can also drag files into the VoceVista window. In Windows Explorer, you can right-click on a file and select "Open with" and then chose VoceVista. You may have to click on "Choose program..." if VoceVista 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 can open most audio file formats. It can also open some video formats and load their audio stream.

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

3.3.2. Export Screenshots

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

Click **File** > **Export Screenshot** to open the **Export Screenshot** dialog box. The following screen appears.




Figure 65. Export Screenshot Dialog Box

This window has two main sections. On the left side in the **Export Area** section, you specify whether to take a screenshot of the whole application window or just some of the views in the user interface, and on the right side in the **Output Images** section, you specify the location where you want to save the graphic and the filename for the resulting image.

Export Area

In this section, select **entire window** to take a screenshot of the VoceVista window, or select **selected views** and then check any of the views that appear below this option. Only those views you select appear in the screenshot. VoceVista exports the smallest rectangle that includes all selected views, so in some cases other views may be included. As you select views, a transparent colored box appears in the application over the views that will appear in the resulting screenshot.

NOTE

Only selected views that are currently visible in the application window will appear in the screenshot.

Image Size

As you select different views to include in the screenshot, the values in the **Image Size** reflect the size in pixels of the graphic that will be created. Click **Set size** to change this value.

Set size

Click **Set size** to open the **Set Window Size** dialog box where you can select from preset dimensions, or enter custom values for the exported screenshot.

VoceVista will attempt to resize the main window such that the exported views fit your requested size.

Output Image File

This section enables you to specify where to save the exported screenshot and its filename.

Output Folder

The path to the location where the screenshot will be saved appears in this field. Click the three dots to the right of the field to browse to a different location. Check **use location of current recording** to save the screenshot in the same folder as the recording.

File Name

The file name for the exported screenshot appears in this field. You can edit the file name directly, or click the three dots that appear to the right of the field to enter the file name and path for the exported screenshot. Check **use name of current recording** to save the screenshot in the same folder as the recording.

TIP

If you save multiple screenshots, the application appends a number to the file name so that each is unique. VoceVista will not overwrite existing files.

Export Now

After you have made all of your selections, press **Export Now** to create the screenshot.

Export with Hotkey

If you want to save your current settings and bypass the **Export Screenshot** dialog box every time you want to save a screenshot, click **Export with Hotkey**. Now, every time you press **[F12]** (configurable through **Keyboard Shortcuts** on the **Settings** dialog), the application creates a screenshot using the saved settings.

TIP

After you click **Export with Hotkey**, you have the option to check the **include mouse cursor option** so that the mouse cursor appears in screenshots that you create using the hotkey. This can be helpful when you are highlighting information in the application.

3.3.3. Export Data

VoceVista not only enables you to view detailed information about your recordings, you can also export the following data to a spreadsheet for further analysis:

Long-term data

- pitch
- mean pitch
- vibrato

Short-term data

- statistics result (pitch)
- spectrum
- long-term average spectrum
- checked marker average spectra
- smoothed spectra
- marker list

Select **File** > **Export Data** to open the **Export Data to Spreadsheet** dialog box. The following screen appears.




Figure 66. Export Data to Spreadsheet Dialog Box

This window has two main sections. On the left side in the **Export Data** section, you select:

- long-term and/or short-term data you want to export
- the time range for the data

On the right side in the **Output File** section, you specify:

- location to save the file
- file name for the resulting image
- format for the name of the resulting spreadsheet
- whether to add charts and size
- whether or not to export the audio at the same time

Export Data

The controls in this section of the dialog box enable you to select the types of data you want to export and the time range for the data.

There are two tabs near the top that reveal different potential data sources.

- Select **Long-term** to select/de-select
 - pitch
 - mean pitch
 - vibrato
- Select **Short-term/Aggregate** to select/de-select
 - statistics result (pitch)
 - spectrum
 - long-term average spectrum
 - checked marker average spectra
 - smoothed spectra
 - marker list

In the **Exported time range** section, you can choose one of the following time ranges for the data you want to export:

- **selected range** (if you have made a selection in the recording)
- **visible range** (the time range that appears in the spectrogram)
- **entire recording**

Output File

The controls the **Output File** section enable you to specify different information about the file(s) you are exporting including location, name, name format, chart selections, and whether to export the audio as well.

Output Folder

The path to the location where the spreadsheet will be saved appears in this field. Click the three dots to the right of the field to browse to a different location. Check **use location of current recording** to save the spreadsheet in the same folder as the recording.

Filename

The file name for the exported spreadsheet appears in this field. You can edit the file name directly, or click the three dots that appear to the right of the field to enter the file name and path for the exported file. Check **use name of current recording** to save the spreadsheet in the same folder as the recording.

Name Format

The **Name Format** drop-down enables you to select the information that the application will append to the exported spreadsheet so that the file names are unique.

- **Running number** - the application adds the next number in a series, e.g. 1, 2, 3 ...
- **Seconds** - the application uses the time range information for the data from the original recording
- **Frames** - the application uses the frames information for the data from the original recording.

NOTE

The **Seconds** and **Frames** options are only available if you choose **selected range** or **visible range** as the exported time range.

Charts

VoceVista enables you to create charts from the data that appear in the exported spreadsheet.

- **add charts** - Check this box to include charts for the data you selected to export.
- **use current scale ranges** - Check this box to use the scale ranges for the data your are exporting from your recording. If you leave the box empty, the application will attempt to determine the best scale to use for the charts based on the data you are exporting. This may result in a chart that is easier to analyze since it doesn't include ranges that are outside your data results. You must check **add charts** in order to check this box.
- **Chart size** - If you are adding charts to the spreadsheet, select the chart size (small, medium, or large) from the drop-down list.

Exported Files

The name of the file that will be exported appears in this section of the **Export Data to Spreadsheet** dialog box next to **Spreadsheet**:

Check the **Audio** box to export the audio source for the data at the same time as you export the data. When you export the data, VoceVista creates a .wav file with the audio in the same location as the spreadsheet. The spreadsheet and audio files will have the same name and different extensions (.wav for the audio file, .xlsx for the spreadsheet).

Export Now

After you have made all of your selections, press **Export Now** to create the spreadsheet (and audio file, if selected).

Export with Hotkey

If you want to save your current settings and bypass the **Export Data to Spreadsheet** dialog box every time you want to export the data, click **Export with Hotkey**. Now, every time you press **[F11]** (configurable through **Keyboard Shortcuts** on the **Settings** dialog), the application creates a spreadsheet using the saved settings.

3.3.4. Export Sound Generator Profiles

Exporting and migrating Sound Generator Profiles is explained at [Section 3.5.4, “Manage Sound Generator Profiles”](#).

3.3.5. Loading and Saving Note Slider Layouts

If you have created a complex configuration of Note 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 **Notes** › **Save note slider layout** and specify where you want to save the layout file. The file type ending is .oa (for Overtone Analyzer Files).

Loading a slider layout

Click on **Notes** › **Load note slider layout** to open an existing .oa file. This will replace your current slider layout.

Export note sliders as MIDI

Click on **Notes** › **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 Note 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 Note 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 Note Sliders.

You can also save a slider layout as DefaultSliderLayout.oa into the folder where VoceVista 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 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 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. 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.

3.4.7. 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.

3.4.8. 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.

3.4.9. 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 67. 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.

3.4.10. Change File Format

Change sampling rate and bit depth

This is done by opening the file format dialog with **File › Loaded File Format**.

Convert between mono and stereo

This can also be done on the file format dialog, and with the edit commands **Edit › Combine mono to stereo** and **Edit › Split stereo to mono**.

3.5. Sound Generator

3.5.1. Fundamental and Overtones




Figure 68. Sound Generator - Fundamental and Overtones

Fundamental

The fundamental is the basic waveform from which the generated sound is composed. You can select from the following types:

- **Sine Wave** - The resulting sound is composed of one or more sine waves added together.
- **White Noise** - The resulting sound is pure white noise. This fundamental sound type has no additional parameters.
- **Coherent Noise** - The resulting sound is somewhere between sine waves and white noise. Coherent noise is a waveform that is more irregular than a sine wave, but still has a frequency.

White Noise (Breathiness)

This setting enables to add a custom amount 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

The controls in this section determine if the pitch stays constant over the duration of the sound, or if it moves, for example by creating a triad or a scale.

When the selected **Type** is not **Constant Pitch**, the pitch moves from the first value in the **Note** or **Frequency** fields (**from:**) to the second value in these fields (**to:**). Check **return to start** to have the pitch follow the same pattern from the second note back to the starting note.

Type

The pitch type options are:

- **Constant Pitch** - The generated pitch is constant over the entire duration of the sound.
- **Single Step** - The generated pitch moves smoothly from the start to the finish note.
- **Major / minor triad** - The generated pitch forms a major or minor triad.
- **Major / minor scale** - The generated pitch forms a major or minor scale.
- **Fifth and Fourth** - The generated pitch moves from the first note to the fourth and then another fifth to the octave.
- **Intervals** - The generated pitch moves from the first note to second and then back in intervals of the major scale.
- **Sweep (linear)** - The generated pitch forms a straight interpolation along the linear scale.
- **Semitone steps** - The generated pitch moves in semitones from the **from** note to the **to** note

Smoothness

Use this control to specify whether the pitch changes gradually from one tone to the next (high smoothness value), or with a more sharply defined step (low smoothness value).

Harmonics

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 (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.

The options for the **Number** drop-down are:

- **Specify Max amount** - Use the control on the right to specify how many overtones to use. 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** - Use the control on the right to specify the frequency below which all overtones will be included (or half the sampling rate, whichever is lower).

Spacing

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

The options for the **Spacing** drop-down are:

- **Harmonic Series** - Each overtone is 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 has 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 are all octaves of the fundamental.
- **Multiply by** - Each overtone has 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: $\text{frequency}(n) = \text{factor} * \text{frequency}(n-1)$
- **Fundamental factor** - Each overtone has the frequency of the fundamental, multiplied by a factor and by the number of the overtone. If the factor is 1.0, this generates the harmonic series. In other words, the factor sets the distance between consecutive overtones. + In mathematical terms: $\text{frequency}(n) = n * \text{factor} * \text{fundamental}$

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).

The options for the **Amplitude** drop-down are:

- **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: $\text{intensity}(n) = 1 / (\text{frequency}(n) / \text{fundamental})$
- **Flat** - Select this option to set all of the harmonics at the same amplitude.
- **Custom** - This option enables you to click and drag your mouse across the black screen to specify custom amplitude values.

3.5.2. Modulation




Figure 69. 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.

3.5.3. Resonances

Available in VoceVista Video and higher

NOTE

The Resonances section of the Sound Generator is only available in *VoceVista Video* and higher.




Figure 70. 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.

3.5.4. Manage Sound Generator Profiles

How to use Sound Generator Profiles

3.6. Profile Manager




Figure 71. Profile Manager

NOTE

Available in VoceVista Video Pro

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.

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 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, 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.

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 Note 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.

Predefined Profiles

VoceVista 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 Note 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.

3.7. Customizing the User Interface

This section is about customizing the user interface of VoceVista, in particular its appearance and its behavior.

Use View Menu

something about using the view menu

Change colors and line thickness

use the fonts and colors dialog

Change Keyboard shortcuts

use the keyboard dialog

Customize Toolbar

use the customize toolbar dialog

Use Profiles

use the profile manager

Change the vowel chart voice

go to the Sound Generator Settings

Chapter 4. Theory and Concepts

What does VoceVista measure? Unlike the [user interface overview](#), which shows how the various displays fit into the overall presentation of the program, this chapter aims to explain in more detail what the program actually measures, and what this means from the perspective of (Psycho-)Acoustics and Voice Science.

VoceVista provides multiple ways to analyze and visualize sound, each offering unique insights into different aspects of vocal production and acoustics. Starting with the fundamental representation of sound as a waveform, we progress through increasingly sophisticated analysis tools that reveal frequency content, intensity patterns, and spectral characteristics.

VoceVista attempts to visualize certain aspects of sound, such as the frequencies that are in the sound, and their intensity. Sound is a vibration, and that is a motion that happens over *time*. If the motion is periodic, we can describe it by its *frequency* (how many vibrations per second), and its *intensity* (how large is the amplitude of the vibrations). Consequently, most displays in VoceVista represent a relationship of something against an axis of time, frequency, or amplitude/intensity, or a combination of those.

Each visualization method measures specific aspects of sound:

- The **waveform** shows amplitude variations over time - the basic structure of sound
- **Frequency** analysis reveals pitch and harmonic content
- **Intensity** measurements quantify loudness and energy distribution
- The **spectrum** breaks down sound into its frequency components
- The **spectrogram** shows how these components change over time

Together, these tools provide complementary views that help us understand and analyze vocal sound production from both physical and perceptual perspectives.

4.1. Waveform

The waveform is the most fundamental way to visualize sound. It plots the [amplitude](#) of the signal against time, giving a direct view of the recorded audio in the [time domain](#).

Basic Principles

A waveform represents sound as a sequence of [samples](#) — measurements of air pressure (more precisely, of the displacement from atmospheric pressure) taken at regular intervals. Each sample is stored as a number scaled into the range -1 to +1, where the extremes correspond to the loudest signal the file format can represent.

The number of samples per second is the [sampling rate](#). CD-quality audio uses 44 100 samples per second; studio recordings often use 48 000, 96 000, or 192 000.

Two Views of the Same Signal

What the waveform shows depends on the zoom level.

Oscilloscope view

At extreme zoom, each sample becomes visible as a point, and the connecting line traces the actual shape of the wave. This view reveals the periodicity of a held tone, lets you spot discontinuities or artifacts in the recording, and is useful for studying the very first cycles of a note's attack.

Aggregated view

When zoomed out, a single screen pixel covers many samples. Rather than picking one sample arbitrarily, the waveform aggregates the samples within that pixel's time range and draws two things:

- the **maximum and minimum** values within the range — the outer envelope of the signal;
- the **RMS** (root mean square) value, drawn closer to the centerline, which tracks perceived loudness more closely than the peaks alone.

Aggregation makes it easy to see overall dynamics, identify louder and quieter sections, and locate specific events in long recordings.

Reading the Waveform

The horizontal axis is time, from earlier on the left to later on the right. The vertical axis is amplitude, with 0 at the centerline (silence) and ± 1 at the extremes — the maximum the file format can encode. When the signal would exceed those limits, samples are cut off at ± 1 , producing distortion known as [clipping](#).

Relationship to Other Views

The waveform shows the overall intensity of the signal but not its frequency content; the [Spectrum](#) and [Spectrogram](#) reverse that emphasis. Sudden amplitude changes — note onsets, plosive

consonants, glottal closures on the EGG — are typically more obvious on the waveform than on the spectrogram.

In VoceVista

See [Section 3.1.4, “Timeline and Waveform”](#) for how to display and navigate the Waveform and the Timeline, including the available display modes and the linear/logarithmic amplitude scales.

4.2. Frequency

Frequency is one of the most fundamental properties of sound and is central to musical analysis. It is measured in **Hertz (Hz)**, the number of vibrations per second.

Basic Principles

When air molecules vibrate in a regular pattern, we perceive the frequency of the vibration as a musical pitch — the faster the vibration, the higher the perceived pitch. In the displays of VoceVista, faster vibrations appear as more closely spaced peaks in the Waveform, and as higher positions on the frequency axis of the Spectrum and Spectrogram.

Musical Context

Frequencies map to musical notes:

- The musical note A4 (concert pitch) corresponds to 440 Hz by convention. VoceVista can be reconfigured for non-standard tunings.
- Each **octave** represents a doubling of frequency.
- Each semitone of **equal temperament** multiplies the frequency by the twelfth root of 2 (about 1.0595).
- One **cent** is one hundredth of a semitone — a useful unit for measuring fine deviations in **intonation**.

Linear and Logarithmic Frequency Scales

The Spectrum and Spectrogram can show frequency on a linear or a logarithmic scale, and the choice has a strong effect on what stands out in the display.

A **linear** scale spaces frequencies equally: every Hz takes the same amount of space. The harmonics of a held tone are evenly spaced, which makes harmonic structure easy to see and to count, and the upper part of the spectrum gets a generous amount of room.

A **logarithmic** scale gives each octave the same amount of space. This matches the way we hear pitch: equal musical intervals appear as equal distances on the screen, and piano keys, drawn against the scale, are evenly spaced. A log scale is usually preferred for musical and pedagogical work.

Frequency Resolution

The **frequency resolution** of the analysis is the smallest difference in Hz that the analyzer can resolve as two separate frequencies. It depends on three interacting choices:

FFT size

The number of samples in each **FFT** calculation. Larger FFTs resolve finer differences in frequency.

Time window

A longer analysis window lets the FFT see more cycles of a slow vibration, and so resolves low frequencies more precisely. A shorter window, conversely, responds faster to changes over time.

Window function

The shape applied to the samples before the FFT (see [FFT Window Function](#)). It controls how cleanly two close frequencies are separated and how much each peak smears into its neighbors.

There is no single best setting: better frequency resolution always costs [time resolution](#), and vice versa.

Limits

The lowest measurable frequency is set by the analysis window: if the window is shorter than one cycle of the signal, there is no full cycle for the FFT to measure. A 50 Hz tone has a 20 ms cycle, so a window shorter than that simply cannot represent it as a frequency.

The highest measurable frequency is set by the [sampling rate](#) of the recording. The **Nyquist** limit — half the sampling rate — is the highest frequency a digital recording can encode. A 44.1 kHz recording reaches up to 22.05 kHz; a 48 kHz recording reaches 24 kHz. The microphone and playback chain further restrict what can be captured and reproduced in practice.

For comparison, healthy young human hearing covers roughly 20 Hz to 20 kHz, so the analyzable range usually exceeds what you can hear.

Applications

Understanding frequency is at the core of many uses of VoceVista: precise pitch measurement for intonation work, harmonic analysis of complex tones, identification of [formants](#) in vocal pedagogy, assessment of tuning, characterization of [vibrato](#), study of vocal registration, and comparison of [timbre](#).

In VoceVista

The frequency scale of the Spectrum and Spectrogram is configured in [Scale Settings](#); FFT size, window function, and resolution are in [Analyzer Settings](#).

4.3. Intensity

Intensity is a measure of how loud or strong a signal is. VoceVista shows it in two complementary ways: the **overall** intensity of the signal at each moment, and the intensity of **each frequency component** in the signal.

Overall and Per-Frequency Intensity

The **Waveform** shows the combined strength of all frequencies at each instant in time. Values in the waveform are **amplitudes**, scaled into the range -1 to +1 for digital audio. When the signal exceeds those limits, samples are cut off at ± 1 , producing distortion known as **clipping**.

The **Spectrum** and **Spectrogram** show how the energy is distributed across frequencies. The Spectrum shows intensity as height on its vertical axis; the Spectrogram shows intensity as colour. Reading these views together — overall intensity in the waveform, frequency-by-frequency intensity in the spectrum — gives a complete picture of the loudness structure of a recording.

Decibels

Spectrum and Spectrogram measure intensity in **decibels** (dB). The dB scale is *logarithmic* and *relative to a reference level*, and is built around three rules that take some getting used to:

- **0 dB is the maximum** — the loudest signal the file format can represent. There is no positive dB value; anything stronger would clip.
- All other values are **negative**: -6 dB, -12 dB, -24 dB, and so on.
- Each **-6 dB** step is roughly a halving of perceived loudness.

So -6 dB is half as loud as 0 dB, -12 dB is a quarter, -18 dB is an eighth, and so on. The scale extends downward indefinitely; the practical floor is set by the noise of the recording and by the bit depth of the file.

Why the Scale Is Inverted

This arrangement seems backwards at first, but it reflects a real asymmetry of digital audio:

- There is **no natural lower limit** to how quiet a sound can be. You can always make a vibration smaller; the bottom of the scale runs to negative infinity.
- There **is** a natural upper limit: the largest number the file format can encode. That hard ceiling becomes the obvious reference, and so it is set to 0 dB.

A useful analogy is depth below the surface of a body of water. The surface is the natural reference (0 m); every depth below it is negative (-1 m, -2 m, ...); and there is no theoretical limit to how deep one can go. Saying "I'm at -10 m" gives a clear, absolute position. Saying "I'm at 50 % of the depth" doesn't, because the answer depends on how deep the water is at that point. Sound intensity in dB works the same way: 0 dB is the surface, and every quieter signal sits somewhere below it.

Dynamic Range

The actual range of intensities you can see depends on the recording format and on the analysis:

Recording format

16-bit audio supports about 96 dB of [dynamic range](#); 24-bit, about 144 dB; 32-bit float, an unrealistically wide theoretical range — useful primarily as headroom against clipping during recording and processing.

Spectrum analysis

The spectrum display typically shows usable detail across about 100 dB. The bottom of that range is set by the FFT window function, by background noise in the recording, and by the microphone.

For most analytical purposes, focusing on the range from 0 dB down to -60 dB shows the relevant content. Sounds quieter than -60 dB are usually indistinguishable from noise.

In VoceVista

The visible dynamic range of the Spectrum and Spectrogram is set with the brightness and contrast controls on the toolbar; the Waveform's amplitude axis can be switched between linear and logarithmic display modes, see [Section 3.1.4, “Timeline and Waveform”](#). Colour mapping for the Spectrogram is configured in the [Colormap Editor](#).

4.4. Spectrum

The fundamental question that the [Spectrum](#) answers is: at a given point in time, which frequencies are present in a sound, and how loud is each one? More precisely, for a given short period of time, what is the average intensity of each frequency over that period?

Basic Principles

Where the [Waveform](#) shows a sound in the [time domain](#) — amplitude against time — the Spectrum shows it in the [frequency domain](#): intensity against frequency. The two views describe the same signal but answer different questions.

The transformation between the two views is performed by the [Fast Fourier Transform](#) (FFT). Given a short window of audio samples, the FFT computes how much energy is present at each frequency. The result is most revealing on stable, sustained sounds, where the spectrum settles into a clear pattern; on quickly changing material, each spectrum captures only a fleeting snapshot.

Reading the Spectrum

The horizontal axis shows frequency, typically up to about 5000 Hz for voice analysis; it can be set to a linear or logarithmic scale. The vertical axis shows intensity in [decibels](#), with 0 dB at the top.

Peaks on the spectrum correspond to strong frequency components:

- Sharp, narrow peaks at evenly spaced frequencies are [harmonics](#) — frequency components at integer multiples of a [fundamental](#).
- Broad peaks that envelop several harmonics are [formants](#) — resonances of the vocal tract that shape vowel sounds.

The relative heights of these peaks are what give a sound its [timbre](#): the same fundamental sung as different vowels produces the same harmonic positions but different formant peaks.

Spectrum Types

VoceVista can compute the spectrum from windows of different lengths, and the choice has a strong effect on what the spectrum reveals.

Instantaneous Spectrum

Computed from a short window centred on a single point in time. Responds quickly to changes in the signal — the most detailed view, but also the most jittery.

Averaged Spectrum

Averaged over a longer window. Smooths out short-lived fluctuations and produces a more stable view of sustained sounds.

Long-Term Average Spectrum (LTAS)

Averaged over a longer section of the recording — often the whole file — to reveal the overall spectral fingerprint of a voice or instrument. See [Long-Term Average Spectrum](#) for more.

Applications in Voice Science

For singing teachers and voice scientists, the Spectrum makes objective measurements available for phenomena that are otherwise described only subjectively. From a single sustained tone you can read:

- the **fundamental frequency** (the pitch of the note);
- the **distribution of energy** among the harmonics, which determines the brightness or warmth of the timbre;
- the **formants**, in particular the strength and position of the [singer's formant](#) in trained classical voices;
- the effect of different vowels and different vocal techniques on the resonance pattern.

In VoceVista

The Spectrum is part of the [Analyzer View](#). FFT size, window function, and the choice between instantaneous and averaged spectra are configured in [Analyzer Settings](#); for the dedicated long-term display, see [Long-Term Average Spectrum](#).

4.5. Spectrogram

The [Spectrogram](#) extends the [Spectrum](#) into the time dimension. Where a spectrum shows the frequency content of a single moment, a spectrogram shows how that frequency content evolves over time. You can think of a spectrogram as many spectra placed side by side and viewed from above.

From Spectrum to Spectrogram

The two views differ in which axes they use:

	Spectrum	Spectrogram
Horizontal axis	Frequency	Time
Vertical axis	Intensity	Frequency
Intensity shown as	Height of the curve	Colour (or brightness)
Time shown as	(none — single instant)	Horizontal position

The spectrogram trades the spectrum's intensity axis for a time axis, and represents intensity through colour instead. Brighter or more saturated colours indicate stronger frequency components; darker colours, weaker ones.

Reading the Display

What you see on a spectrogram lines up with a small vocabulary of patterns:

- a **bright horizontal line** — a frequency that is sustained over time (a held tone);

- **several parallel horizontal lines** — the [harmonics](#) of a musical tone, evenly spaced on a linear scale and progressively closer on a logarithmic one;
- **vertical bright stripes** — moments where many frequencies sound at once, such as a percussive attack or an unvoiced consonant;
- **gently wavy harmonics** — [vibrato](#), the periodic modulation of pitch;
- **moving formant bands** — bright zones that drift across several harmonics, typical of a vowel transition;
- **dark areas** — frequencies with little or no energy at that moment.

Why the Spectrogram Is Powerful for Singing

Consider a singer holding one note and then changing to another:

- On the **spectrum**, the peaks suddenly jump to new positions, and the previous note is gone.
- On the **spectrogram**, the harmonics of the first note appear as horizontal lines, then move to the new positions of the second note's harmonics. The transition between the two is preserved.

This is what makes the spectrogram the central display for singing analysis: it shows not only the notes themselves, but the way the singer moves between them — register transitions, vibrato, vowel changes, and shifts in vocal timbre.

In VoceVista

The Spectrogram is part of the [Analyzer View](#). The frequency scale, FFT size, time resolution, and colour mapping all affect what shows up in the display; see [Analyzer Settings](#) and the [Colormap Editor](#).

4.6. Pitch / Fundamental Frequency

[Pitch](#) is a perceptual property of a sound that lets us classify tones as *higher* or *lower*. For a sustained tone from a single voice or instrument, pitch corresponds closely — but not identically — to the [fundamental frequency](#): the lowest frequency in a periodic signal, of which all the [harmonics](#) are integer multiples.

The Distinction Between Pitch and Fundamental Frequency

The fundamental frequency is a measurable physical property of the signal. The pitch is what a human listener actually hears, and the two can drift apart for several reasons:

- Two listeners can disagree about the pitch of the same sound, especially if it is short, noisy, or in an unusual register.
- The ear can perceive a pitch even when the corresponding frequency is missing from the signal — the **missing fundamental** effect. If only the 2nd, 3rd, and 4th harmonics of 100 Hz reach the ear, most listeners still hear the pitch as 100 Hz.
- Sounds whose components are not on a [harmonic series](#) — singing bowls, bells, chimes — may have no single defensible pitch. Different listeners, and different analyzers, will give different

answers.

For most singing and most pitched instruments, the perceived pitch and the harmonic fundamental coincide closely enough that the two terms can be used interchangeably.

How VoceVista Estimates Pitch

To estimate pitch, the analyzer looks for periodicity in the signal — in effect, asking "what is the fundamental frequency of which the prominent peaks are integer multiples?" When the signal is clean and harmonic, this is reliable. The estimate becomes less stable on:

- unvoiced consonants and other noise-like sounds, where there is no periodicity;
- very low or very high pitches near the limits of the analyzer's range;
- polyphonic material with two or more independent voices.

In these cases the displayed pitch line may flicker, jump, or disappear; the spectrogram remains an honest record of what the signal contains, even where the pitch estimate fails.

In VoceVista

The estimated pitch is overlaid on the [Analyzer View](#) and can be turned on or off there. The expected pitch range can be configured in [Fundamental Ranges](#), which helps the analyzer avoid octave errors when tracking very low or very high voices.

4.7. Harmonics

A [harmonic](#) is a frequency component of a sound that is an integer multiple of a [fundamental frequency](#). If the fundamental is 100 Hz, its harmonics are at 200 Hz, 300 Hz, 400 Hz, and so on. The full set of these multiples is the [harmonic series](#).

Why Harmonics Exist

Most pitched sound sources — vocal folds, vibrating strings, columns of air in a wind instrument — do not vibrate at one frequency only. They vibrate at all integer multiples of a base frequency simultaneously. Each of those frequencies adds its own component to the radiated sound; the relative strengths of those components are what give an instrument or voice its characteristic [timbre](#). Two trumpets playing the same A4 sound similar because both produce harmonics at 440 Hz, 880 Hz, 1320 Hz, ... but they sound recognisably *like trumpets* because the relative loudness of those harmonics follows a pattern that is typical of trumpets.

In the human voice, the source of harmonics is the regular opening and closing of the [vocal folds](#). The shape of the resulting waveform — sharp closure, smoother opening — determines how much energy ends up in the higher harmonics. The [vocal tract](#) then filters this source by amplifying frequencies near its [resonances](#) and attenuating others. The harmonics carry the source's energy through the vocal tract; the resonances shape which harmonics come out loudest.

Harmonics and Overtones

The two terms refer to almost the same set of frequencies, with a small difference in counting:

- The first **harmonic** is the fundamental itself. The second harmonic is one octave above the fundamental, the third is an octave plus a fifth above it, and so on.
- The first **overtone** is whatever sits above the fundamental — i.e., the second harmonic. The "first overtone" and the "second harmonic" are the same frequency.

In VoceVista the harmonic numbering is used.

Reading Harmonics in the Display

In the [Spectrum](#), harmonics appear as a series of peaks at evenly spaced frequencies. In the [Spectrogram](#), they appear as parallel horizontal lines: evenly spaced when the frequency axis is linear, progressively closer together when the axis is logarithmic.

In VoceVista

A [Note Slider](#) can be drawn out to show the position of each harmonic of a chosen fundamental, which makes it easy to verify whether a peak in the Spectrogram is a harmonic of a given note and to read off harmonic numbers. See [Note Sliders](#) for how to place and configure sliders.

4.8. Subharmonics

A [subharmonic](#) is a frequency that is a whole-number fraction of a fundamental frequency. If the fundamental is 200 Hz, its subharmonics lie at 100 Hz, 66.7 Hz, 50 Hz, 40 Hz, ... — the same set of frequencies as the [undertones](#) of the fundamental.

The terms **subharmonic** and **undertone** are largely interchangeable; **subharmonic** emphasises the mathematical relationship to a reference frequency, while **undertone** emphasises the audible tone.

As a Mathematical Construct

Mathematically, the subharmonics of a fundamental are the inverse of its harmonics: harmonics multiply the fundamental, subharmonics divide it. Drawn on a logarithmic frequency scale, the harmonic series spreads upward and the subharmonic series mirrors it downward. VoceVista can show both at once on a [Note Slider](#).

Unlike harmonics, subharmonics do not arise spontaneously from most ordinary vibrating systems. A plucked string, a column of air, or a freely vibrating vocal fold radiates harmonics, not subharmonics. Subharmonics show up only when something more complicated is going on.

As a Vocal Phenomenon

In the human voice, a subharmonic at half the fundamental frequency typically signals **period doubling**: the [vocal folds](#) are vibrating in alternating cycles instead of in a single repeating cycle. Two consecutive cycles differ slightly — for example in amplitude, in closing pattern, or in duration

— and the result is an extra component at half the fundamental frequency, an octave below the apparent pitch.

The most common occurrence is **vocal fry** (also called **creaky voice**), at the lower end of the modal range. Period doubling can also appear as a transient feature in stressed singing, in some forms of overtone singing, or in pathological voices. Subharmonics at one third of the fundamental, from period tripling, are rarer but follow the same mechanism.

On the [Spectrogram](#) a subharmonic appears as a horizontal line *between* the harmonics of the apparent fundamental — at the half-way position for period doubling, the third-way position for period tripling.

In VoceVista

A [Note Slider](#) can be drawn to show subharmonics (undertones) below its fundamental, mirroring the harmonics above it. See [Note Sliders](#) for how to set this up.

4.9. Resonances and Formants

A [resonance](#) is a frequency at which a system naturally vibrates with high amplitude. The [vocal tract](#) has multiple resonances simultaneously, each of which amplifies sound near its own frequency. The vocal tract is the air-filled space above the vocal folds — throat, mouth, and (when open to it) nasal cavity — and its shape determines where these resonance frequencies fall.

A [formant](#) is the peak that appears in the spectrum when one or more [harmonics](#) happen to fall near a resonance. The literature on the voice does not always sharply distinguish formants from resonances, and the words **formant**, **resonance**, and **overtone** are sometimes used interchangeably. The useful distinction:

- A **resonance** is a property of the vocal tract — it exists whether or not anyone is making a sound.
- A **harmonic** is a frequency component of the source signal — what the [vocal folds](#) (or another source) actually produce.
- A **formant** is what you see in the spectrum: a peak where harmonic energy and a resonance line up.

Source and Filter

A common framing for voice acoustics is the **source–filter** model:

- The **source** is the vocal-fold vibration, which produces a buzz rich in harmonics. The fundamental frequency of that buzz determines the perceived pitch.
- The **filter** is the vocal tract, whose resonances amplify some frequencies and attenuate others.

The same source can produce different vowels by reshaping the filter; the same vowel can be sung at different pitches by changing the source while keeping the filter shape. This separation is what makes it possible to recognise a vowel independently of its pitch, and to recognise a singer’s voice across the range.

Formant Numbering and Vowels

Vocal-tract resonances are numbered upward from the lowest: F1, F2, F3, F4, ... The first two are the most important for distinguishing vowels:

- F1 reflects the openness of the mouth — roughly, how far the jaw is dropped.
- F2 reflects the front-back position of the tongue.

Each vowel corresponds to a characteristic (F1, F2) pair. The [Vowel Chart](#) in VoceVista is exactly this: a two-dimensional plot of F2 against F1, with the standard [IPA](#) vowel symbols positioned at their typical formant pairs.

F3 contributes to the perception of "rhotacised" sounds (such as the English /r/) and to brightness. F4 and F5 are involved in voice timbre and individual voice identity.

The Singer's Formant

In trained classical voices, F3, F4, and F5 cluster together into a strong peak in the 2.5–3.5 kHz region, known as the [singer's formant](#). Because orchestral instruments produce relatively little energy in that band, the singer's formant lets a voice project over an orchestra without amplification. It is one of the most studied features of vocal acoustics and shows up clearly on the [LTAS](#) of an operatic voice.

Reading Resonances in the Display

In a Spectrum, resonances appear as broad humps in the curve, with one or more harmonics riding on top of each hump. In a Spectrogram, they appear as bright horizontal bands that are wider than the harmonic lines and slowly drift up or down as the singer changes vowel.

In VoceVista

The [Vowel Chart](#) shows the F1/F2 pair for the current cursor position, mapped against IPA vowel symbols. A [Note Slider](#) can be drawn to mark a specific frequency in the Spectrogram, useful for tracking a single formant. For visualising the formant envelope as a smoothed curve, see [Smoothed Curves and the Spectral Envelope](#).

4.10. Smoothed Curves and the Spectral Envelope

A spectrum is naturally jagged: every harmonic of the sung note shows up as a sharp peak, with deep valleys between them. Most of the time, what we want to see is not the individual harmonics but the **envelope** underneath them — the gentle hills and valleys that reflect the resonances of the vocal tract. The smoothed curve makes that envelope visible.

What the Smoothed Curve Shows

The vocal-tract [filter](#) shapes the source by amplifying frequencies near each [resonance](#) and attenuating frequencies away from them. The harmonics riding on this filter inherit its shape: a harmonic that lands close to a resonance frequency comes out loud; one that lands between

resonances comes out quiet.

The unsmoothed Spectrum shows both the harmonics and the filter response, mixed together. The smoothed curve hides the harmonics and leaves the filter response — the **spectral envelope**. Reading a smoothed curve is therefore much closer to reading the vocal tract directly:

- A bump in the smoothed curve is a **formant**, regardless of which harmonic happens to fall on it.
- The shape of the curve between formants reflects the overall balance of the voice — bright, warm, breathy, ringing.
- Two singers producing the same vowel produce smoothed curves whose major bumps are in similar places, even if the two are at different pitches.

How the Smoothing Works

VoceVista produces the smoothed curve by **low-pass filtering** the spectrum: at each frequency, the displayed value is averaged over a small window of nearby frequencies. The width of that window — the **radius** — determines how much detail is preserved:

- A **narrow** radius leaves much of the harmonic detail intact and only smooths over very local jaggedness.
- A **wide** radius smooths through the harmonics entirely, leaving only the slow envelope.

The right radius depends on what you want to see. To follow formants reliably across a sung phrase, the smoothing has to adapt to the pitch: at higher pitches harmonics are spaced further apart, so the smoothing window must be correspondingly wider. VoceVista therefore lets you express the radius as a **percentage of the current pitch**; this is the right setting for most pedagogical and analytical work. A radius expressed as a fixed number of Hertz is also available, useful when you want consistent smoothing within a fixed frequency band.

When Smoothing Helps and When It Hides

Smoothing is most useful when you are looking for the **vocal-tract** part of the signal — formants, vowel identity, timbre. It is less useful when you want to see individual harmonics — pitch, intonation, register transitions in detail. A common workflow is to keep both views available and switch between them, or display the smoothed curve as an overlay on the unsmoothed Spectrum.

In VoceVista

The smoothed curve is enabled in [Display Settings](#), separately for the short-term Spectrum and for the LTAS or marker-averaged spectra. The smoothing radius and its interpretation (relative to pitch, or absolute Hz) are configured in the same place; see [Radius](#).

4.11. Long-Term Average Spectrum

A single Spectrum captures one moment of a recording; an averaged Spectrum captures a short window. The **Long-Term Average Spectrum (LTAS)** goes further: it averages spectra over an extended section of the recording — often the whole file or a long marked region — to reveal the

overall spectral fingerprint of a voice or instrument.

Why Average Over a Long Range?

When the analysis window is short, the spectrum changes constantly: each new pitch, vowel, or breath rearranges the picture. A long-term average smooths through these short-term changes and leaves what stays roughly constant across the recording — the typical balance between low and high frequencies, the prominence (or absence) of the singer’s formant, the overall noisiness of the voice.

This makes the LTAS the right tool whenever the question is about **style**, **identity**, or **training**, rather than about a specific note:

- a stable, settled spectral picture, less affected by the particular notes being sung;
- a clearer view of [resonance](#) features that persist throughout the recording — for example, the [singer’s formant](#) cluster around 2.5–3.5 kHz in trained classical voices;
- a useful basis for **comparison**: between two sections of the same recording, between two recordings of the same singer, or between two different singers.

What the LTAS Does Not Show

The LTAS averages intensities across time, so it discards everything about **when** something happened. A passage with strong vibrato, a sudden register change, or a pitched ornament all dissolve into the same long-term picture. For those phenomena, the Spectrogram remains the right tool. The LTAS and the Spectrogram are complements: the spectrogram shows *what changes*, the LTAS shows *what stays*.

Choosing the Section to Average

The averaging only makes sense over material that is reasonably uniform. Averaging across silence inflates the noise floor; averaging across very different material — loud and quiet, sung and spoken, two different vowels — blurs together styles that are worth keeping separate. A typical workflow is to mark a passage of comparable material with a Range Marker and compute the LTAS only over that marker.

In VoceVista

The Long-Term Average Spectrum is part of VoceVista Pro. It is computed over the current selection or over a marker; see [Long-Term Average Spectrum](#) for the operational details. To compare LTAS curves between markers, see also [Markers](#). The smoothed-curve overlay (see [Smoothed Curves and the Spectral Envelope](#)) is often used together with the LTAS to highlight the formant envelope.

4.12. Vibrato

[Vibrato](#) is a periodic modulation of the pitch of a sustained tone, typically a few cycles per second. It gives a sung or played tone its characteristic warmth and is one of the things singers and instrumentalists train deliberately. Vibrato is described along two largely independent axes: its **rate** (how fast the modulation cycles) and its **extent** (how wide the pitch swing is).

Rate and Extent

Rate

The number of vibrato cycles per second, measured in Hz. Most classical singing vibrato falls between 5 and 7 Hz. Faster than about 8 Hz is usually heard as nervous or unstable; slower than about 4 Hz tends to be heard as a **wobble**.

Extent

The peak-to-peak depth of the pitch swing, usually expressed in [cents](#) or in semitones. Typical operatic vibrato has a peak-to-peak extent in the range of about 50 to 150 cents — between a half and one and a half semitones of total swing. Smaller extents read as a "straight tone with shimmer"; larger extents become hard to identify as a single pitch at all.

A given singer's vibrato is not a single number: rate and extent vary with pitch, dynamic, and vowel, and they often become more consistent as a singer's training matures.

Vibrato vs. Tremolo and Wobble

Vibrato is a **pitch** modulation. A periodic **intensity** modulation, with the pitch held constant, is called a **tremolo**. Many singers produce a small amount of tremolo as a side effect of vibrato; some instruments (organ tremulant, Leslie speaker, the violin technique called *tremolo*) produce one but not the other. **Wobble** is the colloquial term for vibrato that has become too slow, too wide, or too irregular to read as a single pitch — the line between expressive vibrato and unwanted wobble is largely a matter of context and taste.

Reading Vibrato in the Display

Vibrato is one of the easiest features to spot on a Spectrogram. The harmonics of the held tone, normally drawn as horizontal lines, instead trace gentle wavy lines that all move up and down in lockstep. Higher harmonics show a larger vertical excursion than lower ones — not because the vibrato is wider there, but because a given percentage swing in frequency translates to a larger absolute number of Hz at higher frequencies. On a logarithmic frequency axis the wavy lines run parallel.

The pitch line of the analyzer follows the same modulation, which makes the rate and extent measurable directly: the rate is the inverse of the period of the wave, and the extent is its peak-to-peak amplitude.

In VoceVista

The dedicated [Vibrato View](#) in VoceVista Pro shows vibrato rate and extent as continuous traces over time, separated from the underlying pitch contour. This is the easiest way to compare vibrato characteristics across phrases, sections, or singers.

4.13. EGG (Electroglottograph)

The [electroglottograph](#) (EGG) is a small device that lets us follow the vibrations of the [vocal folds](#) from outside the body. When paired with VoceVista, it gives a second, **physiological** signal

alongside the audio — one that comes directly from the source of the voice and is unaffected by the vocal tract that shapes everything downstream of the folds.

How the Signal Is Acquired

A small high-frequency current flows between two electrodes held in place on either side of the larynx. The conductance between the electrodes increases slightly when the vocal folds make contact, because the contacting tissue forms a continuous conductive path; it falls again when the folds part. The modulations in conductance, recorded over time, give the EGG signal — a measure of vocal-fold **contact** over the cycle of vibration.

The procedure is non-invasive: nothing enters the mouth or throat, and the small current is harmless. The signal is acquired through a separate hardware module and arrives at VoceVista as a second channel synchronised with the audio.

What the EGG Reveals

Because the EGG is a near-direct view of the glottal source, it complements what the audio shows in two important ways:

- It is unaffected by the resonances of the vocal tract. Two singers producing the same vowel may produce very similar audio spectra, but their EGG traces can be quite different — and conversely, the same EGG pattern can underlie very different vowels.
- It makes voice phenomena that are hard to hear in the audio easy to see. Register transitions, fold approximation, breathiness, pressed phonation, and pathology all leave characteristic signatures in the cycle-by-cycle shape of the EGG.

The most commonly read features of the EGG are the **closed quotient** (the fraction of each cycle during which the folds are in contact) and the **closing** and **opening** events — the moments when contact begins and ends. These events are sharper and more reliable in the **derivative** of the EGG (often called **dEGG**), where they appear as positive and negative peaks rather than as transitions in slope.

Applications

The EGG is most widely used by:

- **singers and singing teachers**, to study register transitions, onset coordination, and fold approximation in ways the audio alone does not show;
- **speech-language pathologists**, to characterise voice disorders and monitor therapy through measurable contact behaviour rather than perceptual judgement alone;
- **voice scientists**, who use the EGG as a near-source reference signal in studies of the source-filter division.

In VoceVista

The EGG signal is shown alongside the audio in the [EGG \(Electroglottograph\)](#) view, available in VoceVista Pro. Display options for the EGG, the dEGG, and the Wavegram are configured in [EGG](#)

Display Settings.

Chapter 5. Troubleshooting

Sometimes the program does not work as expected. This section explains what you can do to diagnose the problem, and how to get support.

5.1. Fix Recording

5.1.1. Program is not recording anything

- check correct input source
- check microphone permissions
- check input format
- switch to direct sound

5.1.2. Recorded sound is distorted

- reduce input volume
- increase distance to microphone
- use better microphone

5.2. Contact Us for Support

From within the program, click on Help / Send support request.

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 and contains one or two sub-windows that can show the Spectrogram, the Spectrum, or both. [More about the Analyzer View](#)

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.

dB

Short for [Decibel](#).

Decibel

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 \(Wikipedia\)](#), the opening between the vocal folds.

Equal Temperament

A tuning system that divides the octave into twelve equal semitones, each 100 [cents](#) apart. It is the default tuning of the modern piano and most Western music. Tones in this system are referred to as [tempered tones](#); in VoceVista, the snap-to-note feature snaps slider lines to the nearest tempered tone.

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. Read more about the [FFT Window function on Wikipedia](#).

Fast Fourier Transform

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.

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 peak in the spectrum resulting from harmonics meeting a [resonance](#). The literature on the voice does not always clearly distinguish between formants, resonances, and [overtones](#). Overtones are frequency components of a sound that may be amplified by the vocal tract if they match the frequency of a resonance, and the resulting peak could be called 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 Domain

A representation of a signal as a function of frequency rather than time. The [Spectrum](#) is a frequency-domain view: it shows which frequencies are present in the signal and how strongly, but not when each event occurred. Compare to [Time Domain](#).

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 can show is therefore not the loudness experienced by a human listener, but the sound pressure level recorded by the microphone.

Intonation

How accurately a singer or instrumentalist produces the intended pitches in a piece of music. Good intonation means the notes match the chosen [tuning](#) system closely enough that the music sounds in tune. VoceVista makes intonation visible by showing the pitch of each note against musical notation and by reporting deviation from the nearest tempered tone in [cents](#).

IPA

Short for **International Phonetic Alphabet**, a standardized set of symbols for representing the sounds of spoken language. VoceVista uses IPA symbols on the [Vowel Chart](#) to label vowels independently of any specific language's spelling.

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.

LTAS

Short for **Long-Term Average Spectrum**. An LTAS is computed by averaging many spectra over an extended period of a recording — often the entire recording or a long marked section. This averaging smooths out the detail of individual notes and reveals the overall spectral fingerprint of a voice or instrument, such as the singer's formant cluster around 3 kHz, or the typical balance of low and high harmonics. [more about LTAS](#)

MIDI

Short for Musical Instrument Digital Interface, a standard protocol to encode messages to electronic instruments. In VoceVista, MIDI output is used to play the keys of the piano keyboard and of Note 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.

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. [More about Markers](#)

Mono

A mono recording has one channel, for example the input of a single microphone.

Note

In music, a note is a single pitched sound of a defined duration, and the written symbol that represents it. Within VoceVista, "note" may refer to (1) the musical note as written on the [Staff View](#), (2) a [Note Slider](#) overlaying the spectrogram, or (3) an entry in a [MIDI](#) sequence — all of which share the same underlying idea of a pitch with a time extent.

Note Slider

A Note Slider is a visual tool that is laid over the Spectrogram. Each slider has a fundamental frequency and, optionally, a start and end time, so it can be understood as a musical note — much like a note in a MIDI sequence. A slider can be played as a sound. It can also be drawn out to show the overtones and undertones of its fundamental frequency, which makes it useful for highlighting a specific note, illustrating principles of music theory and acoustics, or transcribing a piece of music. Note Sliders are sometimes referred to simply as *Notes*, and were called *Overtone Sliders* in earlier versions.

Octave

An octave is the interval between two pitches whose frequencies are in a 2:1 ratio. A note one octave above another sounds like the "same note" at a higher pitch and shares its letter name (e.g., A3 and A4). Because the ear perceives pitch logarithmically, equal octaves appear as equal distances on a logarithmic [frequency scale](#) and on the piano keyboard.

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, 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

Earlier name for [Note Slider](#). The two terms refer to the same visual tool; **Note Slider** is the current name.

Partial tone

Other word for [overtone](#).

Pitch

Pitch is a perceptual property of a sound that corresponds to the fundamental 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. [more about Pitch](#)

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.

Resonance

A resonance is a property of the [vocal tract](#) with a specific frequency. The vocal tract has multiple resonances that will amplify sound with the frequency of that tone. The sound can come from the [vocal folds](#), but it may also come from other sources.

Ruler

A ruler is a visual aid that marks a specific frequency or amplitude. Over the Spectrogram, rulers are similar to [Note 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. Also referred to as [Bit Depth](#).

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.

Singer's Formant

A cluster of [formants](#) in the range of approximately 2.5 - 3.5 kHz that is characteristic of trained classical singing voices. The clustering of formants in this region produces a strong peak in the spectrum that allows the voice to project over an orchestra without amplification, since orchestral sound has relatively little energy in that band. The singer's formant is typically visible as a bump on the high end of the [LTAS](#) of an operatic voice.

Snapping

A behavior in VoceVista where a moved [Note Slider](#) line aligns to a nearby reference — most commonly the nearest [tempered tone](#) (snap to musical note) or the nearest peak in the spectrum (snap to spectral peak). Snapping can be enabled, disabled, or temporarily inverted with the **Alt** key while dragging. Useful both for setting a precise reference pitch and for measuring the actual frequency of a sung or played note.

Sound Generator

A built-in tool in VoceVista for synthesizing reference tones, glides, scales, and harmonic stacks directly from the program. Useful for ear training, illustrating acoustic phenomena, generating playback material to practice along with, and providing a controlled signal for testing analyzer settings. [more about the Sound Generator](#)

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 Note 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.

Subharmonic

A frequency that is a whole-number fraction of a fundamental tone — the same set of frequencies as [undertones](#). The two terms are largely interchangeable; **subharmonic** emphasizes the mathematical relationship to the fundamental, while **undertone** emphasizes the audible tone. Subharmonics also appear as a vocal phenomenon (often called **vocal fry** or **creaky voice**) when the [vocal folds](#) vibrate in alternating cycles, producing pitches at 1/2 or 1/3 of the perceived fundamental.

Tempered Tone

A tone whose frequency matches one of the twelve notes of the [equal-tempered](#) scale.

Timbre

Often called **tone color**, timbre is the perceptual quality of a sound that lets us tell apart two notes of the same pitch and loudness — for example, the same note sung as an "ah" and an "ee", or played on a violin versus a flute. Timbre is determined by the relative strengths of the [harmonics](#) (and other components) in the sound, and by how those strengths change over time. In VoceVista, the [Spectrum](#) and [Spectrogram](#) are the main tools for visualizing timbre.

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 Domain

A representation of a signal as a function of time. The [Waveform](#) is the classic time-domain view: it shows the changing amplitude of the signal sample by sample, but does not directly reveal which frequencies are involved. Compare to [Frequency Domain](#).

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.

Tuning

The process of adjusting an instrument or voice so that its notes match a chosen reference pitch system (most commonly [equal temperament](#) with A4 = 440 Hz). Also: the system itself, e.g. **equal temperament tuning**, **just intonation tuning**. VoceVista can be used as a precise tuning aid by

showing the deviation of a played note from the nearest tempered tone in [cents](#).

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.

Vibrato

A periodic modulation of the pitch of a sustained tone, typically a few cycles per second. Vibrato gives a sung or played tone its characteristic warmth and is one of the things singers and instrumentalists train deliberately. VoceVista can visualize the rate (cycles per second), extent (depth in cents or semitones), and regularity of vibrato in the [Analyzer View](#) and the dedicated [Vibrato View](#). [more about Vibrato](#)

Vocal Folds

Two folds of muscle and mucosa in the larynx that open and close to produce the buzzing source sound of the human voice. Often also called **vocal cords**. The frequency at which the vocal folds vibrate determines the [pitch](#) of the voice; the resulting source sound is then shaped by the resonances of the [vocal tract](#) into the recognizable timbre of speech and singing.

Vocal Tract

The air-filled space above the [vocal folds](#) — comprising the throat (pharynx), the mouth, and (when open to it) the nasal cavity — through which sound from the vocal folds passes before leaving the body. The shape of the vocal tract creates [resonances](#) that amplify some frequencies and attenuate others, producing the [formants](#) that define [vowels](#) and a singer's individual sound.

Vowel

A speech sound produced with an open vocal tract, where its quality (a, e, i, o, u, ä, etc.) is determined by the shape of the tract — particularly the position of the tongue and lips. Each vowel corresponds to a characteristic pattern of [formants](#); the first two formants alone are usually enough to distinguish most vowels, which is why the [Vowel Chart](#) is a two-dimensional plot of F1 against F2.

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 & 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.