

Chapter 9 Correlation

About this chapter

This chapter explains how Raven's correlation tool works. Correlations are a way to perform quantitative comparisons between spectrograms or waveforms. In this chapter you'll learn:

- some general information about correlation
- how correlation functions are calculated
- how to use the correlation function and specify options common to both spectrogram and waveform correlations
- about issues and options specific to each type of correlation
- about batch correlation

Overview of Correlation

Correlations are performed by 'sliding' two inputs (either two spectrograms or two waveforms) past each other in time. At each time offset, a *correlation value* between the inputs is calculated. These correlation values are then plotted versus time in a *correlation plot* to show a measure of similarity between the inputs. The time axis of the plot is shown relative to the first input, indicating how far the second input has been offset, so a peak at a positive lag, or time offset, indicates that the second signal occurs at an earlier time than the first (see [Figure 9.2.](#)).

- Correlation types
- Raven provides the capability to run correlations between either two spectrograms or two waveforms.
- **Correlating spectrograms:** Usually when correlating spectrograms, the most important information provided is the peak correlation value which shows the similarity between the spectrogram images.
 - **Correlating waveforms:** Waveform correlation can help determine the *lag* at which two signals most closely resemble each other, as shown in [Figure 9.2.](#)

Calculating Correlation Functions

- Spectrogram correlations
- For each lag Δt , Raven calculates either a normalized or non-normalized correlation value $C_{\Delta t}$ between two spectrograms (For more on the effect of normalization, see ["Normalization"](#) on page 226). If the **Normalize** option is selected, each correlation value is calculated using the following:

$$\text{Formula 9.1} \quad C_{\Delta t} = \frac{\sum_{t=1}^n \sum_{f=1}^{FFT} (X_{t,f} \cdot Y_{t+\Delta t,f})}{\sqrt{\left(\sum_{t=1}^n \sum_{f=1}^{FFT} (X_{t,f})^2 \right) \left(\sum_{t=1}^n \sum_{f=1}^{FFT} (Y_{t,f})^2 \right)}}$$

where n equals $(N_1+N_2) - 1$ and N_1 and N_2 are the numbers of frames in the two spectrograms (for a discussion of these variables, see [“Spectral analysis of time-varying signals: spectrograms and STFT analysis”](#) on page 331). Note that this formula corresponds to a correlation using a biased rather than an unbiased normalization (For more on the distinction between biased and unbiased normalization, see [“Biased/Unbiased”](#) on page 227). FFT equals the number of frequency bins, which must be the same for the two spectrograms being correlated. $X_{t,f}$ and $Y_{t+\Delta t,f}$ are the amplitude values (in this case, the power of the spectrogram) in the two spectrograms at frequency f and times t and $t+\Delta t$, respectively. The normalized correlation value for spectrograms can vary between 0 and 1. A correlation of 0 means that the non-zero values in the two spectrograms do not coincide at all; a correlation of 1 indicates that the two signals are identical (given the lag Δt .) Successive correlation values are calculated by incrementing the value of Δt in steps equal to the time grid resolution (which must be the same for both spectrograms) in effect sliding the two spectrograms past each other in time. If normalization is turned off, only the numerator of [Formula 9.1](#) is used.




Figure 9.1. A spectrogram correlation between two calls of a Black-capped Vireo.

Waveform correlations For each lag, Raven calculates either a normalized or non-normalized correlation value between two waveforms. If the **Normalize** option is selected, each correlation value is calculated using:

$$\text{Formula 9.2} \quad C_{\Delta t} = \frac{\sum_{t=1}^n (x_t \cdot y_{t+\Delta t})}{\sqrt{\left(\sum_{t=1}^n x_t^2 \right) \left(\sum_{t=1}^n y_t^2 \right)}}$$

where n equals $(N_1+N_2) - 1$ and N_1 and N_2 are the numbers of digitized samples in the two waveforms. x_t and $y_{t+\Delta t}$ are the values of sample numbers t and $t+\Delta t$ of the two waveforms, respectively. If the two signals differ in length, the shorter signal is zero-padded at the end to the length of the longer signal. The correlation value for waveforms can vary between -1 and 1 (If you choose to plot the *complex envelope* of a waveform correlation, the values that are plotted vary between 0 and 1, as discussed in “[Waveform correlations](#)” on page 237).

A correlation of 0 means that the signals are orthogonal¹; a correlation of 1 indicates that the two signals are identical; a correlation of -1 indicates that the signals are identical in magnitude, but opposite in phase.

Successive correlation values are calculated by incrementing the value of Δt in steps equal to the inverse of the sampling frequency, in effect sliding the two waveforms past each other in time. If normalization is turned off, only the numerator of [Formula 9.2](#) is used.



Although [Formula 9.1](#) and [Formula 9.2](#) are written in the time domain, Raven actually performs these computations in the frequency domain. This improves the performance of the correlator.

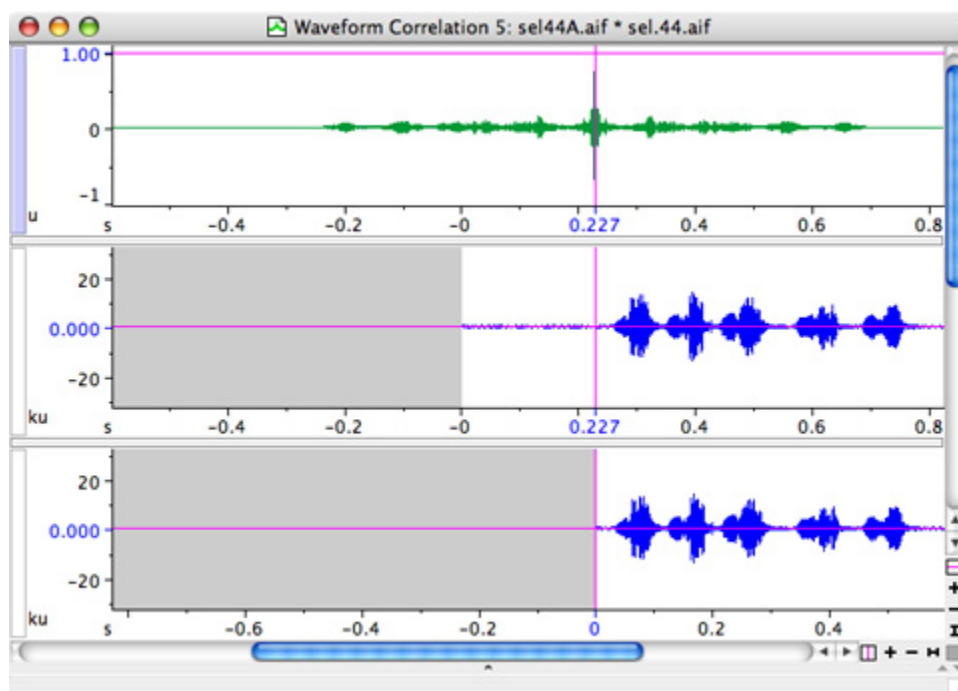


Figure 9.2. A normalized waveform correlation plot. The peak correlation lag is shown on the x-axis by the vertical magenta line. Note that a positive correlation lag indicates that the first waveform is ahead of the second in time.

1. Whether or not two signals are orthogonal depends on their frequency content and on their relative phase. For example, sinusoidal signals of different frequencies are orthogonal, as are signals of the same frequency that are 90° out of phase with each other.

Using the correlation tool

To perform a correlation, choose **Tools > Correlator** and select the two files to be compared. Next, choose whether to correlate the waveform or the spectrogram views and select any appropriate parameters (see below for more information on specific options).



Note that both signal files must be recorded at the same sample rate in order to be correlated.

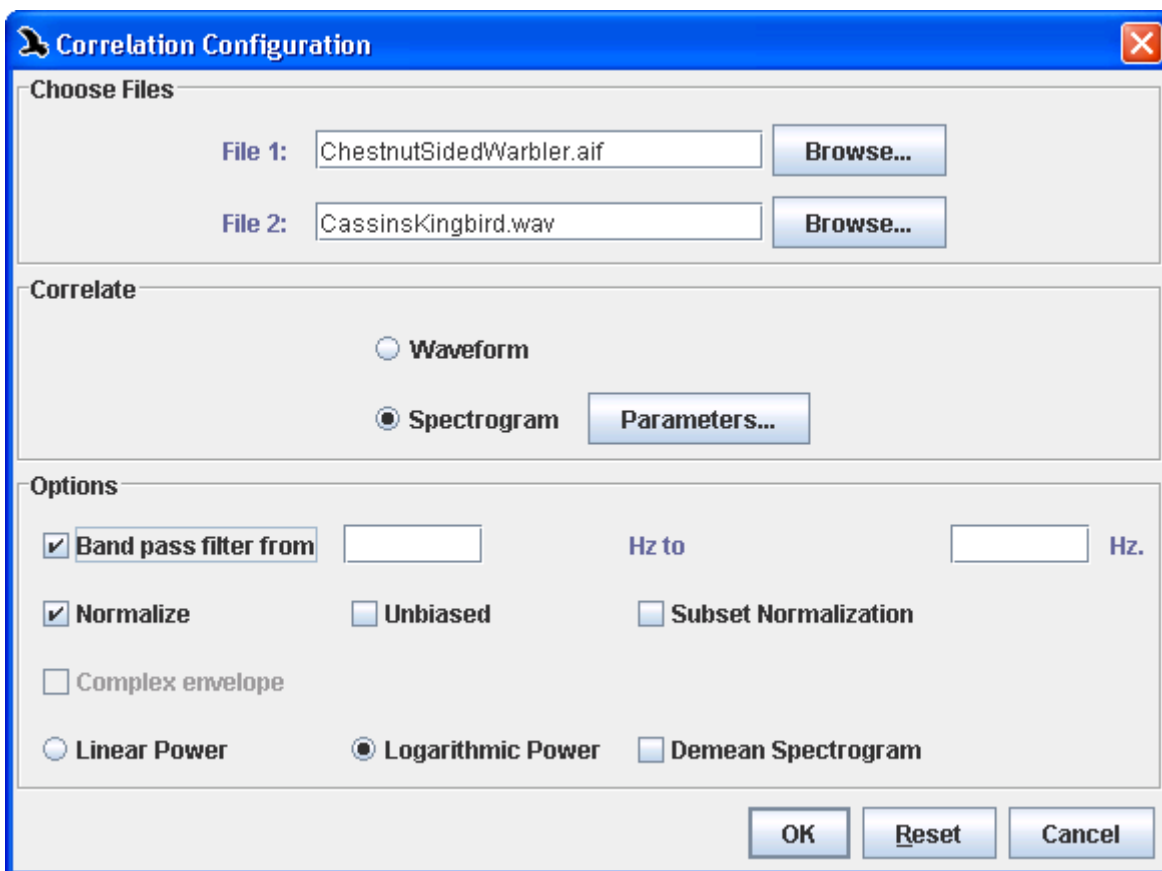


Figure 9.3. The Correlation dialog box. After selecting which two files to correlate, you can choose which type of view to correlate (waveform or spectrogram) along with other parameters and options as well.

Peak correlation values can be important when comparing spectrograms; however, it is important to note that there are serious limitations to correlation comparisons. Spectrogram correlation is not a tool for generalized “pattern recognition” and it is important to remember that the measured similarities are simple and narrowly defined and may not be relevant in a given context. (see [“Spectrogram correlation”](#) on page 229)

The timing of a peak correlation value from a waveform correlation plot is often useful as well. It can determine the lag at which two inputs (possibly from the same source signal) most closely match each other. Usually, waveform correlations are less useful than spectrogram correlations for assessing “similarity” between signals in an intuitive way. This is partly because (unlike spectrogram correlations) waveform correlations are sensitive to phase differences in inputs that our auditory system does not detect.

Canary users will find much of the correlation functionality similar, although there are a few added features.

Band Pass Filter Choosing to add a **Band pass filter** is helpful for signals containing noise. Checking the filter box allows you to enter a lower and upper frequency limit that will filter the data before it is correlated.

In most applications, filtering is advisable for both waveform and spectrogram correlations. By selecting a frequency band corresponding to the relevant signal(s), you reduce the effect of any other noise or signals on the correlation values.

The effect of filtering on the correlation function (for both waveforms and spectrograms) depends very much on the particular signals being correlated. If neither file contains much energy outside the frequency band occupied by the signals, filtered correlations may not differ much from unfiltered correlations. If there is a significant amount of energy outside the frequency band of interest, the difference between the filtered and unfiltered correlations can be much larger. Also, it is important to note that filtering may either raise or lower correlation values, depending on the particular signals being correlated.

Normalization If **Normalize** is checked, the sum of the products of the data values from the two signals is divided by the square root of the product of the sums of values from the two signals, as indicated in [Formula 9.1](#) and [Formula 9.2](#). The units in the numerator and denominator cancel and the correlation value is scaled to a dimensionless value. For spectrograms, which contain only non-negative amplitude values, the normalized correlation value is always between 0 and 1. For waveforms, which can contain positive, negative, and zero values, the normalized correlation varies between -1 and 1.

If **Normalize** is left unchecked, the correlation is calculated as the sum of the products of the data values from the two signals (i.e. just the numerator of [Formula 9.1](#) and [Formula 9.2](#).) A non-normalized correlation is given in arbitrary units.

Normalizing a correlation will compare the overall 2-dimensional shape of signals but will ignore any amplitude differences. Examples of normalized and non-normalized correlations, and how the overall amplitude level affects the correlation values, can be seen in [Figure 9.4](#). Unless your application requires that similarity measurements

incorporate information about the absolute amplitude levels of the signals, you should most likely leave the **Normalize** option checked.

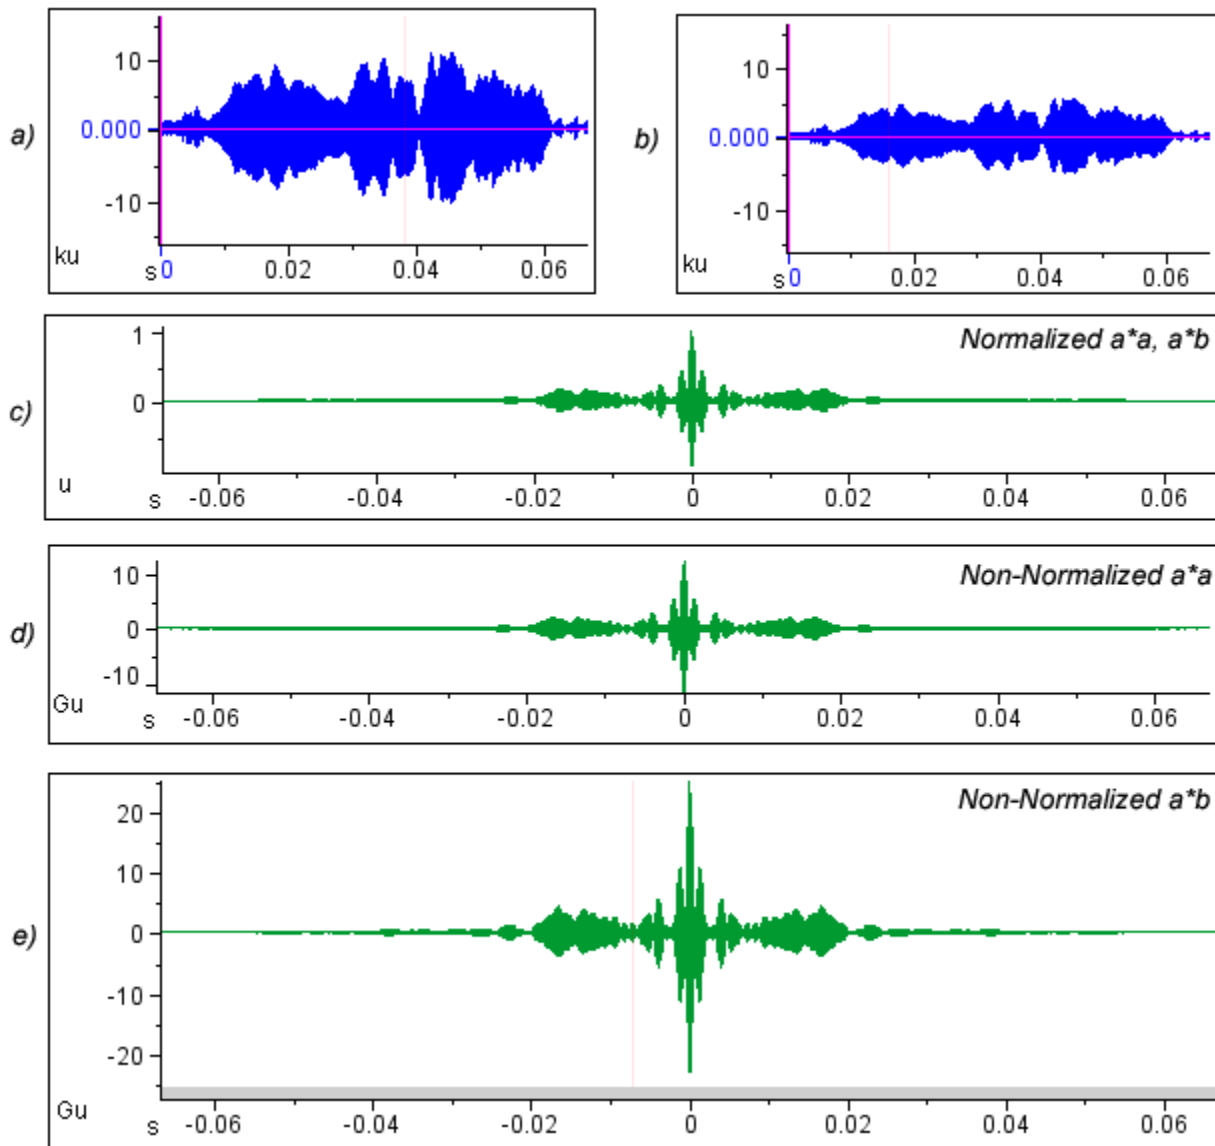


Figure 9.4. Comparison of normalized and non-normalized correlations. **(a)** Waveform of a portion of song from a Chestnut Sided Warbler. **(b)** The same waveform amplified by a factor of 2. **(c)** Normalized correlation between (a) and (b). This correlation is identical to the correlation between (a) and itself. **(d)** Non-normalized correlation between (a) and itself. Note the change in axis scale from units to gigaunits **(e)** Non-normalized correlation between (a) and (b). Notice the differences in correlation value amplitudes among the plots.

Biased/Unbiased By default, a *biased normalization* is performed during correlations. However, you can choose to perform an *unbiased normalization* simply by

checking the **Unbiased** box. The expected value of an unbiased normalization equals the quantity it estimates; however, this sometimes leads to large variance in endpoints because only a few data points are used. To avoid large variations at correlation endpoints, do not check the **Unbiased** box.

Subset Normalization Full normalization uses values from the entire length of both parent signals in calculating the normalization coefficient. However, this can result in high correlation values at locations of greater intensity in a longer sound, even if the features of the sounds are dissimilar. When the Subset Normalization checkbox is checked, only the overlapping portions of each sound are used in normalization. By accounting for variations in noise level within the longer sound, this provides a more accurate correlation between a short segment and a longer sound. [Figure 9.5](#) illustrates this difference when the Black-capped Vireo recording is correlated with a single call using both full and subset normalization.

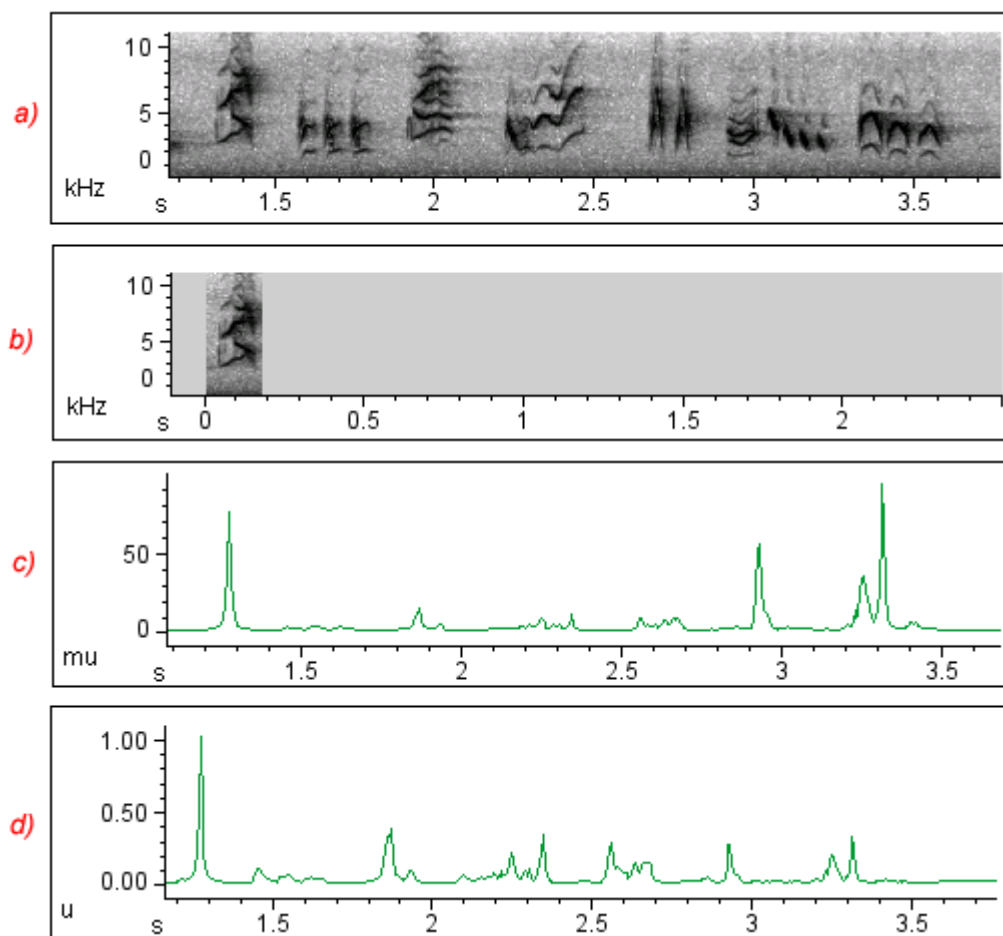


Figure 9.5. Comparison of correlations performed with full and subset normalization. **(a)** a Black-capped Vireo recording. **(b)** A short selection taken from the recording. **(c)** Correlation between (a) and (b) performed using full normalization. Notice that the peak correlation is relatively small (about 93 mu) and occurs at the incorrect location. **(d)** Correlation between (a) and (b) performed using subset normalization. The peak correlation value is now 1, and occurs at the predicted location.

However, when performing a correlation between two short clips, subset normalization can introduce extra peaks near the endpoints of the sounds, when only a small portion of the sounds are used in normalization. For example, in [Figure 9.6](#), a high correlation value is obtained when the end of one selection is correlated with the beginning of the other, despite the dissimilarity of the sounds.




Figure 9.7. A correlation between two Black-capped Vireo calls. Using subset normalization (left image), a peak is introduced when the end of one call is compared with the beginning of the other. However, using full normalization (right image) this peak is not noticeable.

Spectrogram correlation

Often, the only point of interest in a spectrogram correlation plot is the maximum value of the correlation function. The peak correlation value can provide a quantitative measure of one type of similarity between spectrograms.

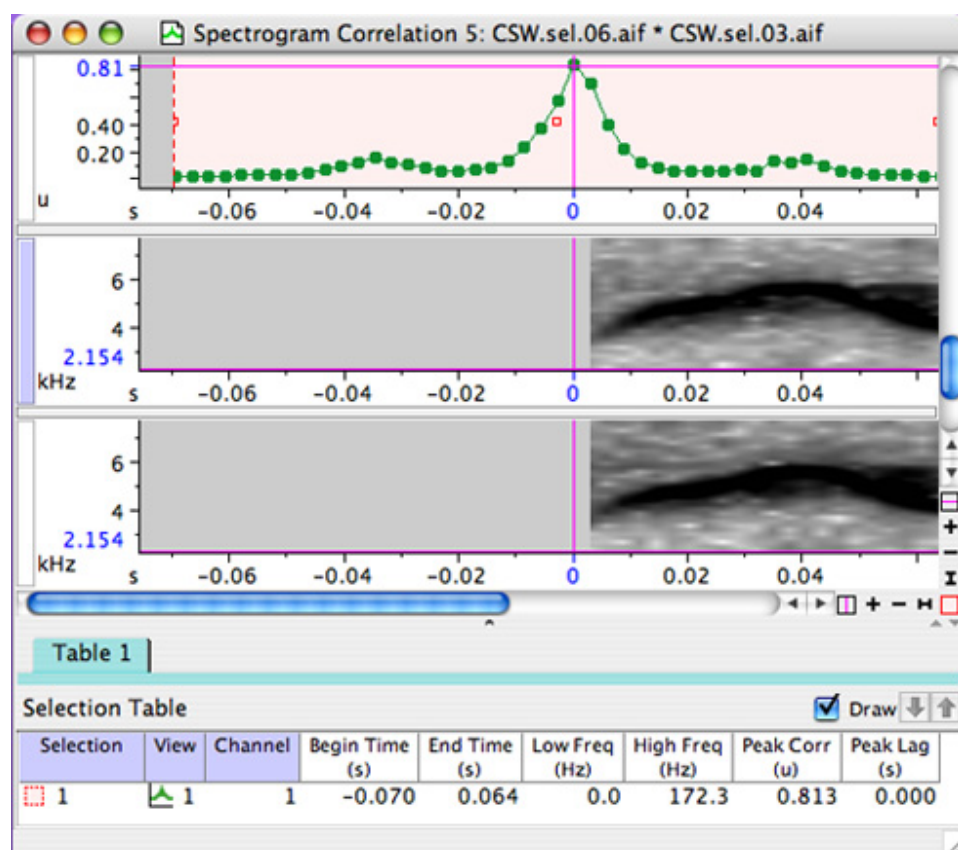


Figure 9.8. The two spectrograms shown were produced in succession during the song of a Chestnut Sided Warbler. Above in green, you can see the correlation between the two spectrograms. The entire correlation graph is selected using Edit > Select All, and in the selection table, the peak correlation measurement of 0.813 is included. Note that the peak correlation and peak lag are shown on the axes and by the magenta lines on the correlation plot, in addition to being shown in the selection table in their measurement columns. The spectrograms are zoomed in frequency to show the detail.

While the peak value of a spectrogram correlation function can provide an objective, well-defined, repeatable, and comparable measure of the similarity of two spectrograms, it is NOT a tool for generalized “pattern recognition”. The “similarity” that is measured is simple and narrowly defined, and may or may not be appropriate to the research question being asked. The usefulness of spectrogram correlations as a measure of similarity thus depends very much on the specific context in which they are being used. The best way to develop a feel for how to interpret spectrogram correlations is to experiment with correlating a variety of spectrograms.

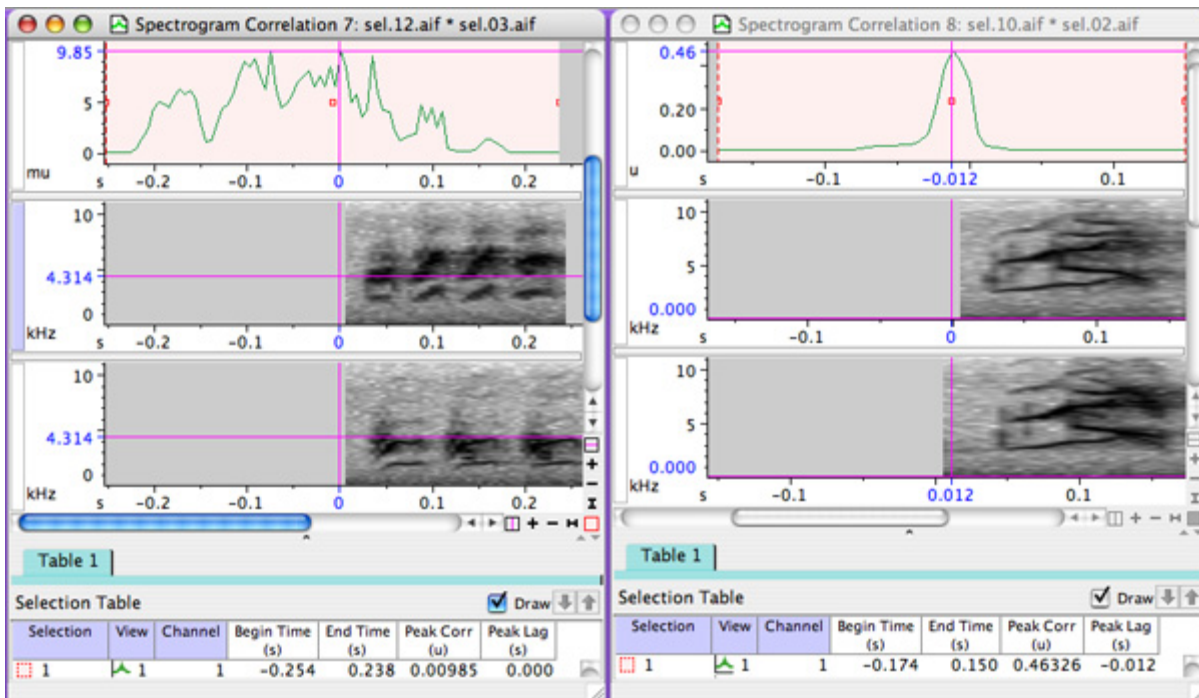


Figure 9.9. Two correlation windows, one showing two sounds with a low peak correlation value (left), and the other showing two sounds with a high peak correlation value (right). Note that although the signals on the left appear to be similar, their high energy components are offset in frequency around the 4314 Hz line, marked by the magenta frequency position marker in the view. Raven does not do any frequency shifting when it performs spectrogram correlations. The peak correlation value on the left is 9.85 milli-units, whereas the peak on the right is 0.463 units, or 463 milli-units. Always be sure to check the units on the y-axis of the correlation view.

When Raven displays a correlation plot, it allows you to scroll the time position of the view to see how the two parent signals line up at different points in the correlation plot. The default view when a correlation plot is initially shown is to show the lag at the peak correlation value, and to have the parent signals lined up at their peak correlation. Moving the time scroll bar in the view allows you to see how the parent signals line up at other correlation values, as shown in [Figure 9.10](#). for spectrogram correlation. The same type of time scrolling can be done with a waveform correlation.

After running a spectrogram or waveform correlation, you can choose to playback either of the parent sounds directly from the correlation window. When a sound view is active, the playback features are available

and apply to the active sound. However, the playback features are disabled in the correlation view.




Figure 9.10. Three correlation windows of the same correlation plot scrolled to three different time positions. The first shows a lesser peak to the left of the highest peak, the second shows the main peak, or peak lag, the time position at which the signals are most highly correlated, and the third shows a lesser peak to the right of the highest peak. Observe how in both the first and the third, two of the three calls in each signal are aligned, but that the correlation value is much higher when all three are aligned in time.

Scale of spectrogram power values

Raven can compute spectrogram correlations using either the logarithmic power values from the spectrogram, in units of decibels, or using the linear power values, in units of squared amplitude units. Linear power values are the traditional method used to compute spectrogram correlations and provide a good spread between spectrograms that look similar and those that look different. Because the logarithmic scale compresses the range in which power values exist, the corresponding correlation plots also tend to be more concentrated with higher peaks, which can lead users to think that signals are similar when they are really dissimilar. However, it is often possible to minimize this effect by adjusting the spectrogram clipping parameters of demeaning spectrogram values before correlation. (For more information on spectrogram clipping and demeaning see [“Demeaning of spectrogram values”](#) on page 233 and [“Spectrogram correlation parameters”](#) on page 234). Also, since logarithmic power values are used to display spectrograms and make measurements in Raven, using logarithmic values in correlations better coincides with these tools.

[Figure 9.11](#) shows two normalized spectrogram correlation plots using linear and logarithmic power values. Notice that using logarithmic as opposed to linear power values changes the maximum correlation value and location, as well as the shape of the correlation plot. The peak correlation values are marked on the y-axis by the horizontal magenta line

and the spectrograms are shifted in time to show the best alignment of the images, as indicated by the peak correlation lag, which is the vertical magenta line.




Figure 9.11. A comparison between a correlation using linear power values (on the left) and the same correlation using logarithmic power values (on the right).

Demeaning of spectrogram values When performing a spectrogram correlation, you can opt to demean the spectrogram values before correlating the files. When **Demean Spectrogram Values** is checked, Raven subtracts the average value from each value in a given spectrogram before applying the correlation formula. This can be useful for comparing sounds with very different average power values or to expand the range of logarithmic power values. Unlike standard correlation, demeaning uses negative correlation values for sounds that are very dissimilar.

For example, [Figure 9.12](#) shows the correlation between two dissimilar Black-capped Vireo selections, performed with both linear and logarithmic power values, with and without demeaning. When using linear power values, demeaning changes the relative heights and shapes of the peaks only slightly, but the effect is much more apparent for logarithmic power values. Without demeaning, using logarithmic power values results in a deceptively high correlation value and shows few of the surrounding features. However, once the spectrograms are demeaned, the peak correlation value is more realistic and the smaller features become apparent.

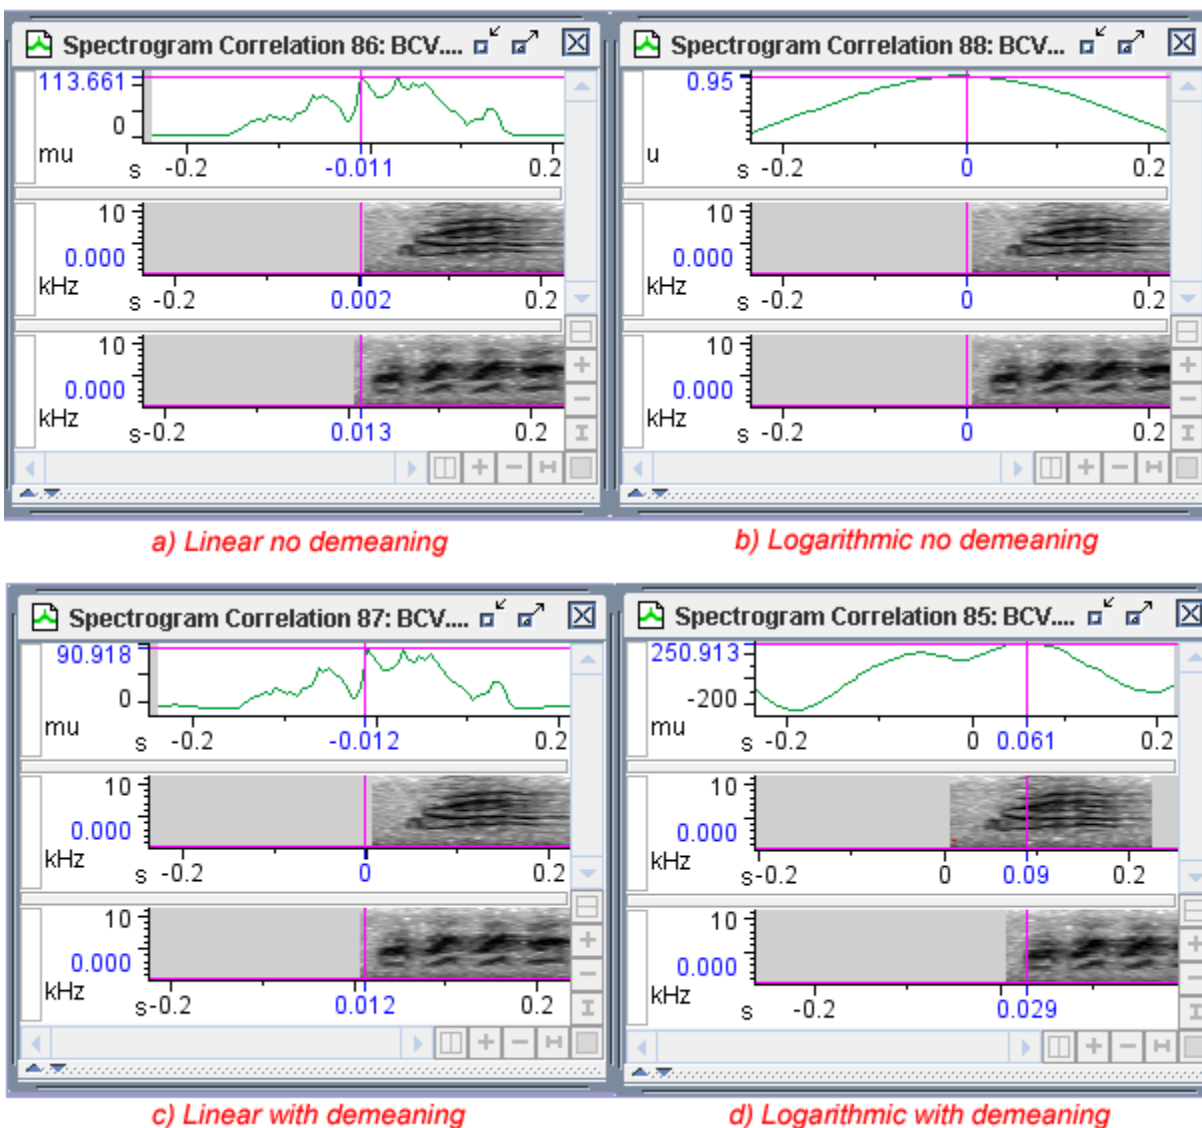


Figure 9.12. Spectrogram correlations for two Black-capped Vireo calls using (a) linear values and no demeaning, (b) logarithmic values and no demeaning, (c) linear values with spectrograms demeaned, and (d) logarithmic values with spectrograms demeaned.

Spectrogram correlation parameters You can configure the view and results of spectrogram correlations (see [Figure 9.15](#)) as you would adjust display parameters for typical spectrograms. In general, increasing the overlap will decrease the hop size and increase the number of frames in the spectrogram for a more detailed correlation. Decreasing the overlap will oppositely affect the hop size and number of frames. Increasing the window size will result in more frequency bins in the DFT and a more detailed analysis of frequency, but at the cost of time resolution. Decreasing the window size will oppositely affect the frequency bins and time resolution.

Adjusting these parameters can noticeably affect the resulting correlations. Since the correlator scans the spectrograms past each other in time, not frequency, an increase in time resolution provides a more detailed correlation, but reduces sensitivity to the frequency distribution of the sounds. For example, [Figure 9.13](#) shows a correlation between two Black-capped Vireo calls. Notice how the top call has a band of high energy around 3800 Hz, while the bottom call has a gap in that range. When the correlation is performed with a large window size, this difference is reflected in a low correlation value. However, as the window size is decreased, the loss in frequency resolution results in a higher correlation value as this difference is smoothed out. For more information on configuring spectrogram display parameters, see [Appendix B, “A Biologist’s Introduction to Spectrum Analysis”](#).




Figure 9.13. Correlations between two Black-capped Vireo calls performed using different window sizes. **(a)** 256 frames **(b)** 64 frames **(c)** 16 frames.

Effect of spectrogram clipping on correlations When performing a correlation, you can choose to clip the individual spectrograms below a certain level before performing the correlation. For more on spectrogram clipping, see [“Clipping level” in Chapter 5](#) (page 123). Because of the way correlations are calculated, the effect of spectrogram clipping is much more significant when using logarithmic, as opposed to linear, power values. Spectrogram clipping can reduce the effect of noise on the correlation, but may also have undesired side effects. For example, clipping spectrogram values to a small finite value can help to accentuate distinct features of a signal that might otherwise be blurred by noise. However, clipping values to no power (-infinity dB) while performing a logarithmic correlation may return an empty correlation view.

Figure 9.14 shows two autocorrelations of a Spotted Hyena call using logarithmic power values. The correlation on the left does not clip any spectrogram values, whereas the one on the right clips all values below 60 dB to 0 dB. By severely reducing the surrounding noise, clipping these spectrogram values exposes the secondary peaks in the correlation that were not previously visible.

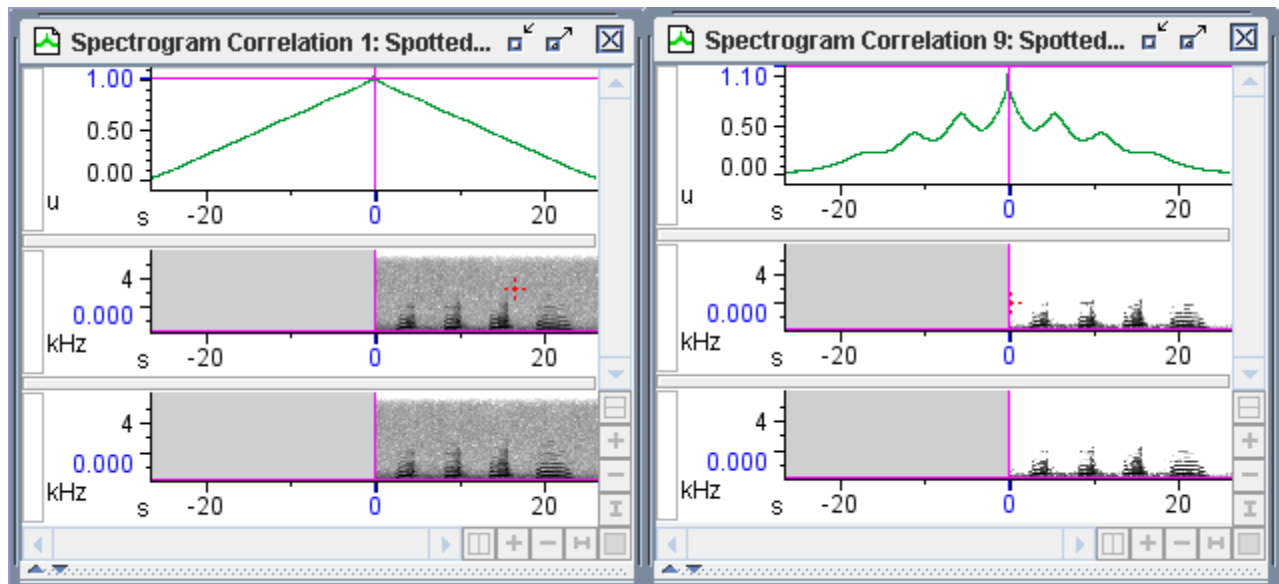


Figure 9.14. Autocorrelation of a Spotted Hyena call performed without spectrogram clipping (left) and with values below 60 dB clipped to 0 dB(right).




Figure 9.15. The spectrogram configuration dialog box. You can adjust the view of spectrogram correlations by changing these parameters.

Waveform correlations

When working with waveforms, one might be interested in the time at which the peak correlation value occurs. To determine the position of a sound in space, for example, an array of microphones in a known geometry can be used to record onto separate but synchronized recording tracks. The lags of the correlation peaks between signals on these synchronized recordings then indicate the delays between the arrival times of the sound at different microphones. These time delays can then be used to calculate the location of the sound source relative to the positions of the microphones, based on known information (including the speed of sound and the microphone array geometry). However, in most applications, waveform correlations are less useful than spectrogram correlations for assessing the degree of “similarity” between signals in a way that is intuitively satisfying (See [“Using the correlation tool”](#) on page 225)

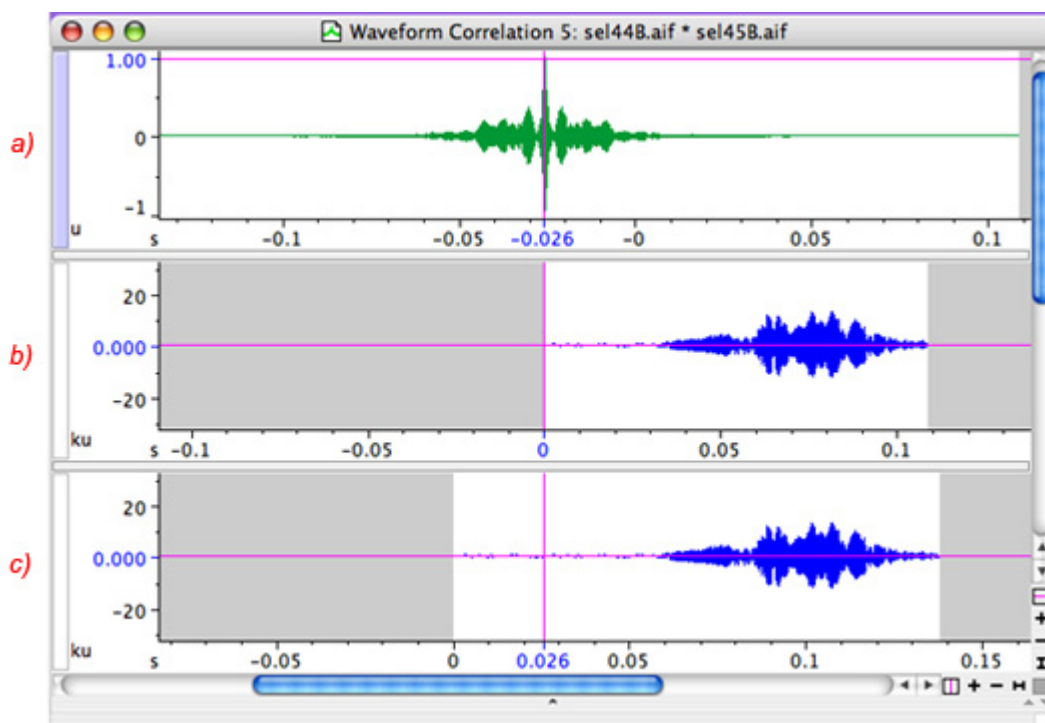


Figure 9.16. Waveform correlation between song syllables simultaneously recorded at two different microphones. **(b)** Waveform from microphone 1. **(c)** Waveform from microphone 2. **(a)** The correlation between (b) and (c). The time delay of -0.026 s indicates that the bird was 8.9 m closer to microphone 1 than to microphone 2 (assuming a speed of sound of 344 m/s).

Complex envelope If you are performing a correlation between two waveforms, you can check the **Complex envelope** box which will display the complex envelope of the correlation function, as opposed to the correlation function itself. A complex envelope varies between 0 and 1. The relationship between the complex envelope and the correlation function itself is illustrated in [Figure 9.17](#).

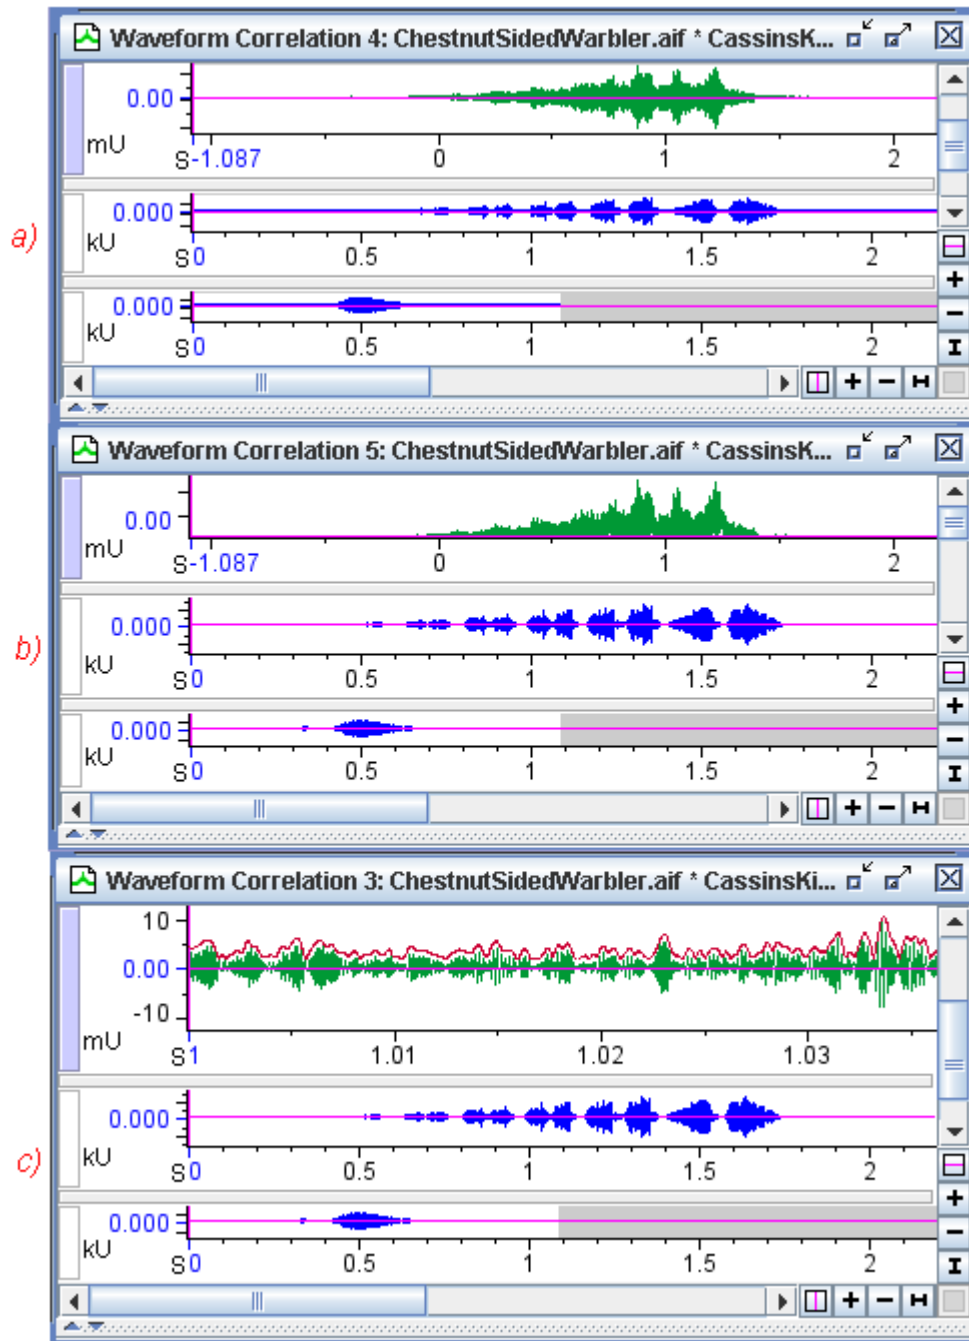


Figure 9.17. (a) The waveform correlation of two sounds. (b) The complex envelope of the correlation function in (a). (c) The complex envelope function superimposed over the correlation function. For clarity of illustration, the complex envelope (in red) is shifted slightly upward in the figure from the raw correlation plot (in green) and the horizontal axis has been zoomed in to show more detail.

Correlation functions always contain high-frequency oscillations, which are related to the frequencies present in the signals being correlated. If the signals are approximately sinusoidal (i.e., at any moment, most of the energy in each signal is concentrated at a single frequency, as in the frequency-modulated whistles common in bird song), their correlation function will itself be close to an amplitude modulated sinusoid. In this case, the complex envelope is roughly equivalent to the amplitude envelope of the absolute value of the correlation function. If the signals being correlated are spectrally complex (with energy distributed over many frequencies, as in human speech), their correlation function contains high-frequency oscillations that are generally not sinusoidal.

Taking the complex envelope removes much of the high-frequency oscillation in a correlation function, which can make it easier to visually identify the peak of a waveform correlation. Since the qualitative relationship between the appearance of the complex envelope and the raw correlation plot depends somewhat on the signals being correlated, you should experiment with the type of signals that you work with in order to get a feel for the relationship between the two types of plot.

Batch correlation

Raven's batch correlator provides a mechanism for automatically performing the same correlation operation on an arbitrarily large number of files and saving the results. For example, a batch spectrogram correlation lets you perform correlations of many files at a time, using a set of spectrogram parameters that you specify only once. The batch correlation process runs in the foreground so you have to wait for it to finish before you can perform any more analysis within Raven, but you can use other applications while the batch correlator is running.

To run the batch correlator, choose **Tools > Batch Correlator...** from the menu bar. This will display the batch correlation configuration dialog [Figure 9.18](#).

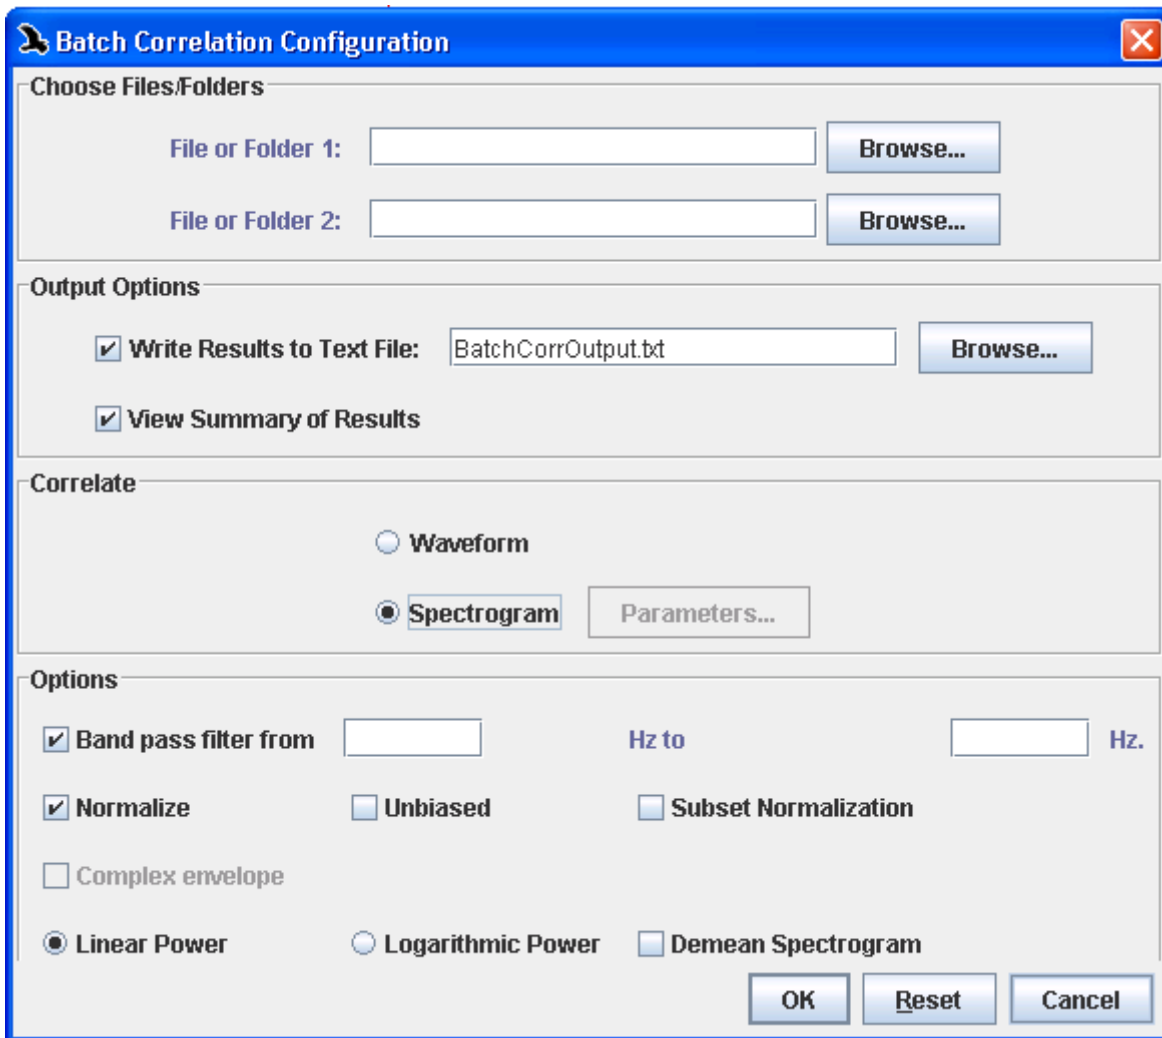


Figure 9.18. The batch correlation configuration dialog.

Input and output For Raven’s batch correlator, you must specify which files are to be used as input and where the output file should be saved. The input files must all be stored in the same folder, and the output file can be saved in a folder, which can be different from the input folder.

For batch correlation, there are two sets of file inputs (see [Figure 9.18](#)). Each set can consist of a file or folder. If you choose a folder, then all sound files within that folder will be used as input to the batch correlation process.

After you select the input files or folders, you must select the output options (see [Figure 9.18](#)). The default settings are to save the output to a text file and to view the results within the Raven window. One of these options must be chosen for the correlator to run. You can change the output folder and name by using the Browse... button, or just enter the name of the output file by typing in the text box.

All other options within the batch correlation configuration dialog are identical to those in the correlation configuration dialog box, as discussed earlier in this chapter.

Running the correlator When you click the OK button, Raven starts calculating the correlations and displays a status window, titled “Spectrogram Correlation” or “Waveform Correlation” depending on your selected option, as shown in [Figure 9.19](#). The progress bar in the status window shows the progress of the entire batch process. Clicking the Cancel button stops the batch correlation.

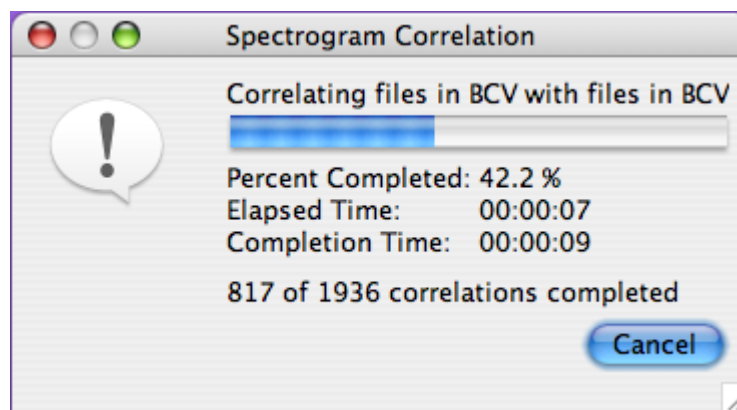


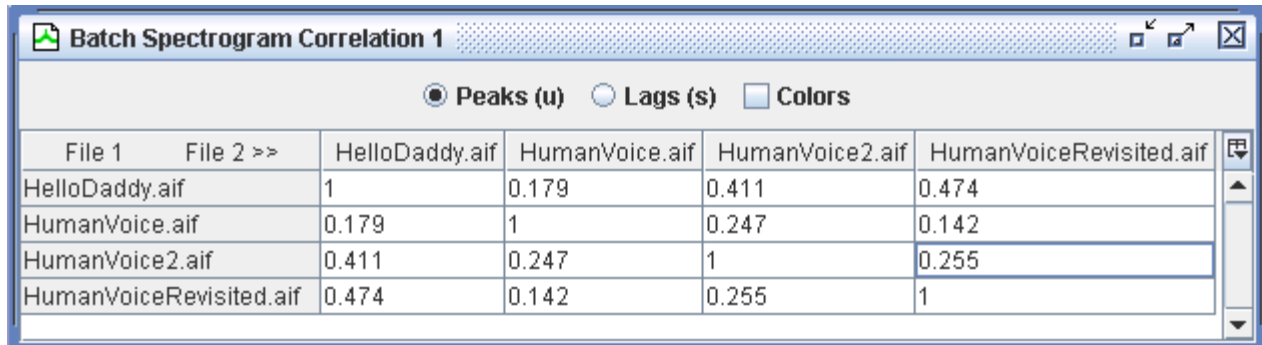
Figure 9.19. The correlation status window (for a batch spectrogram correlation, in this case).

The batch correlator computes correlations for each possible pair of files in the two input sets. When all of the correlations have been calculated, Raven displays a window containing the correlation table ([Figure 9.20](#)). The title of the window specifies whether the correlation was of waveform or spectrogram views. The correlation table contains a peak correlation value (in correlation units) for each pair of files. Clicking on the Lags radio button will change the view to show the lag values (in seconds), or lags at which the peak correlation values occur. If the Colors option at the top of the window is selected, Raven will color code the peak correlation values using the following system: Blue for low scores, then green, yellow, orange, and red for higher scores.



To use a continuous colormap for batch correlation results as opposed to the ten default colors, change the preferences entry: `raven.ui.correlation.colorMap.continuous` from false to true. Once this preference has been changed, you can also change the RGB for low, medium, and high correlation values by editing the corresponding preferences.

By default, the correlation table rows are sorted by the name of File 1. However, clicking on a column heading will sort the table by that column's values. The columns of the correlation table may be resized by positioning the mouse pointer on the line between the column headings, and then dragging the line to a new position. The window itself may be resized in order to show more columns, or the horizontal scroll bar at the bottom of the window may be used to scroll to different columns in the table. If there are enough rows in the table such that they cannot all be seen, then a vertical scrollbar is provided to allow you to scroll to other rows.



File 1	File 2 >>	HelloDaddy.aif	HumanVoice.aif	HumanVoice2.aif	HumanVoiceRevisited.aif
HelloDaddy.aif		1	0.179	0.411	0.474
HumanVoice.aif		0.179	1	0.247	0.142
HumanVoice2.aif		0.411	0.247	1	0.255
HumanVoiceRevisited.aif		0.474	0.142	0.255	1

Figure 9.20. A batch correlation table showing peak correlation values for each file pair using linear power values in the correlations.

To view the correlation function of a single correlation, just double-click on any peak or lag value in the table. Raven will launch a new correlation window showing the correlation function in one view along with the two files that were correlated, each in their own views (see [Figure 9.21](#)). This window is identical to the one that would be shown if a correlation had been run on the individual files. Any of the individual sound files can also be opened by double clicking on the row heading.

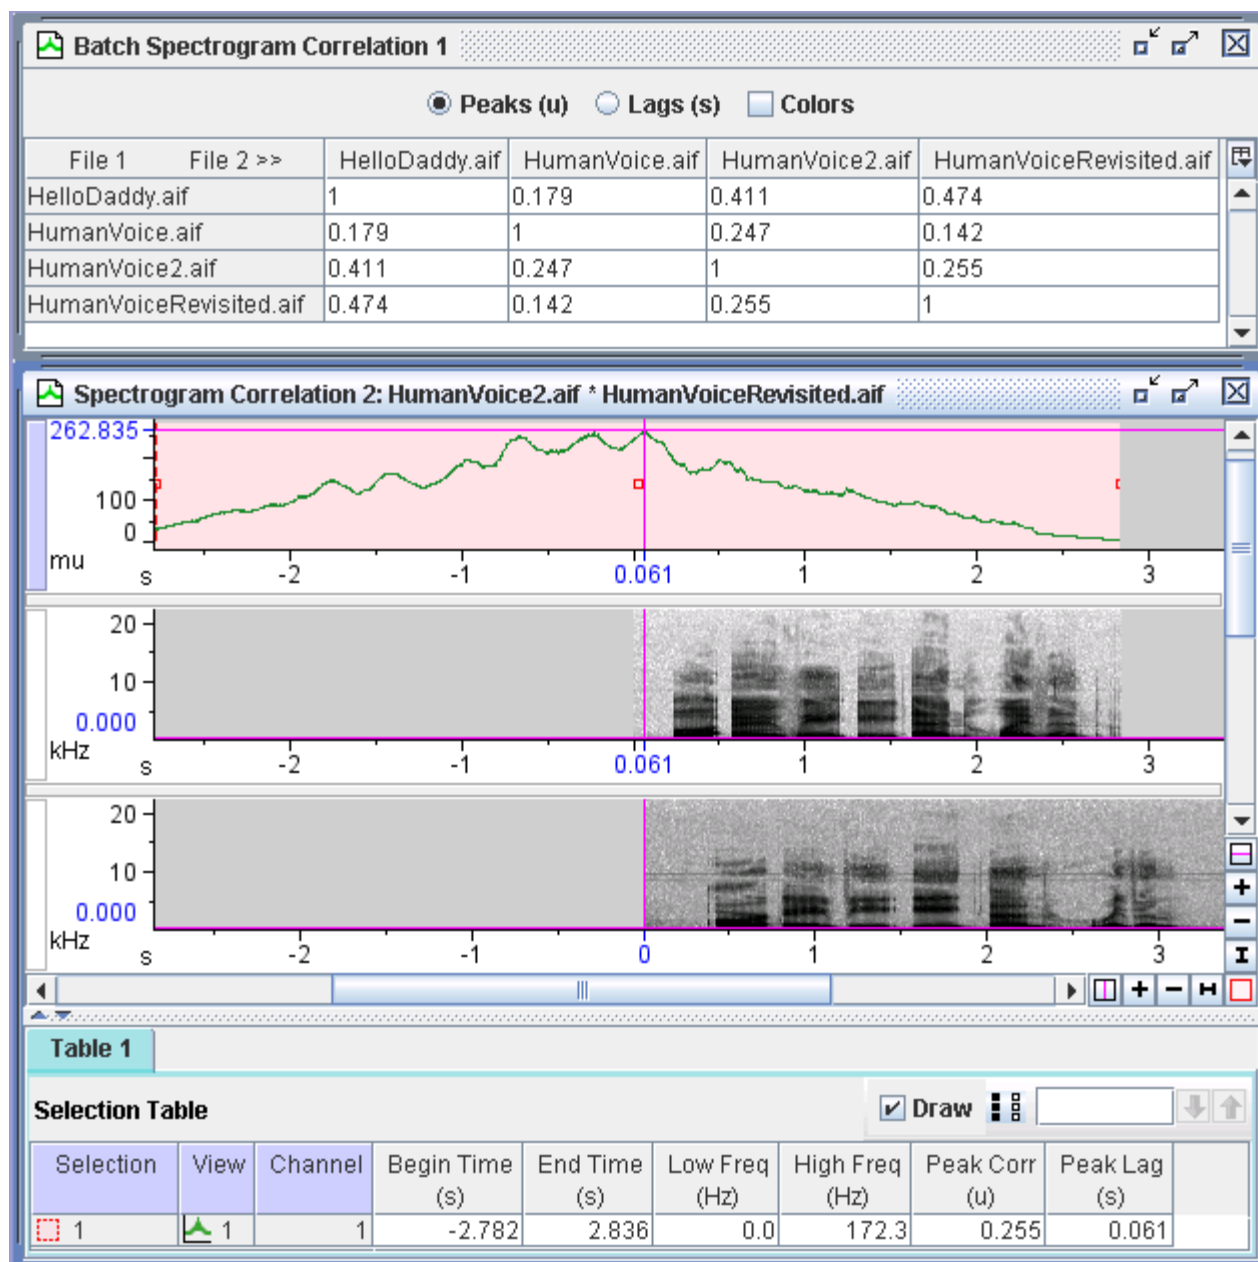


Figure 9.21. By clicking on a single cell in the batch correlation table (top), Raven displays the individual correlation view between the files in a new sound window. Note that the frequency position marker in the correlation view marks the peak correlation value, and the time position marker marks the peak lag value. You can also add these measurements to the selection table, as shown in the figure.

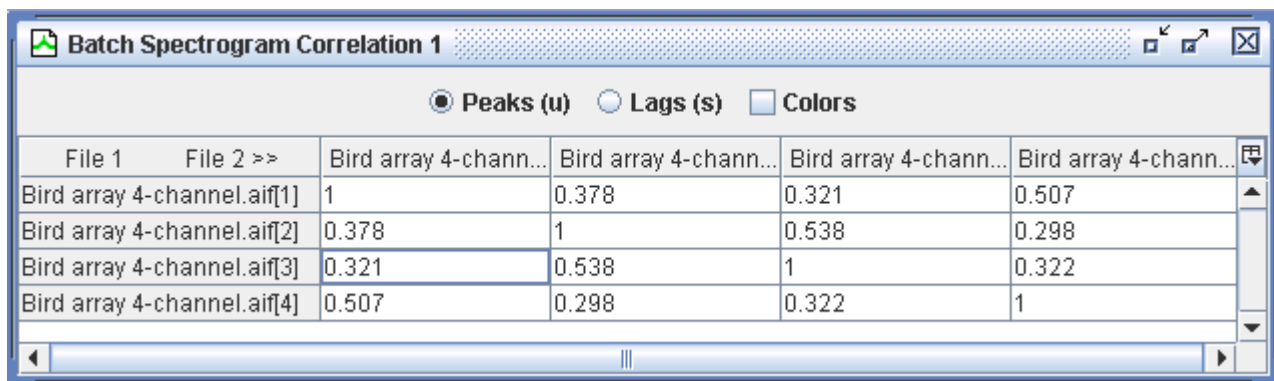
Correlator output can also be saved to a tab-delimited text file, where the correlation peak table and correlation lag table are displayed one after the other (see Figure 9.22). You can specify the output file name on the initial correlator dialog.

	A	B	C	D	E
1	Batch Correlation Peaks (u)				
2					
3	File 1 File 2 >>	HelloDaddy.aif[1]	HumanVoice.aif[1]	HumanVoice2.aif[1]	HumanVoiceRevisited.aif[1]
4	HelloDaddy.aif[1]	1	0.179	0.411	0.474
5	HumanVoice.aif[1]	0.179	1	0.247	0.142
6	HumanVoice2.aif[1]	0.411	0.247	1	0.255
7	HumanVoiceRevisited.aif[1]	0.474	0.142	0.255	1
8					
9	Batch Correlation Lags (s)				
10					
11	File 1 File 2 >>	HelloDaddy.aif[1]	HumanVoice.aif[1]	HumanVoice2.aif[1]	HumanVoiceRevisited.aif[1]
12	HelloDaddy.aif[1]	0	0.058	0.25	-0.299
13	HumanVoice.aif[1]	-0.058	0	0.177	0.717
14	HumanVoice2.aif[1]	-0.25	-0.177	0	0.061
15	HumanVoiceRevisited.aif[1]	0.299	-0.717	-0.061	0
16					

Figure 9.22. A correlation table saved as a file, and subsequently opened using Microsoft Excel.

If you are correlating the contents of a folder with itself, the peak correlation value for $X*Y$ and $Y*X$ will be in the same table, but one lag will have a positive value and the other a negative value, as shown in the lag section of the spreadsheet of [Figure 9.22](#).

For multi-channel files correlated in the batch correlator, each channel is displayed as its own row or column ([Figure 9.23](#)). If the same files were correlated using **Tools > Correlator**, the resulting batch correlation table would not be created, and all channel correlations would be in their own windows. For ease of viewing the results, you should use the batch correlator when dealing with multi-channel files.



File 1	File 2 >>	Bird array 4-chann...	Bird array 4-chann...	Bird array 4-chann...	Bird array 4-chann...
Bird array 4-channel.aif[1]		1	0.378	0.321	0.507
Bird array 4-channel.aif[2]		0.378	1	0.538	0.298
Bird array 4-channel.aif[3]		0.321	0.538	1	0.322
Bird array 4-channel.aif[4]		0.507	0.298	0.322	1

Figure 9.23. A correlation table with results from a multi-channel file correlation. In this correlation, the input files are both the same (Bird array 4-channel.aif), so each channel of the file is correlated with itself and with the other channels.

If a single file is batch correlated with a folder of files, the resulting table will either be a single row (if the file is specified as input 1) or a single column (if the file is specified as input 2). See [Figure 9.24](#) for an example.

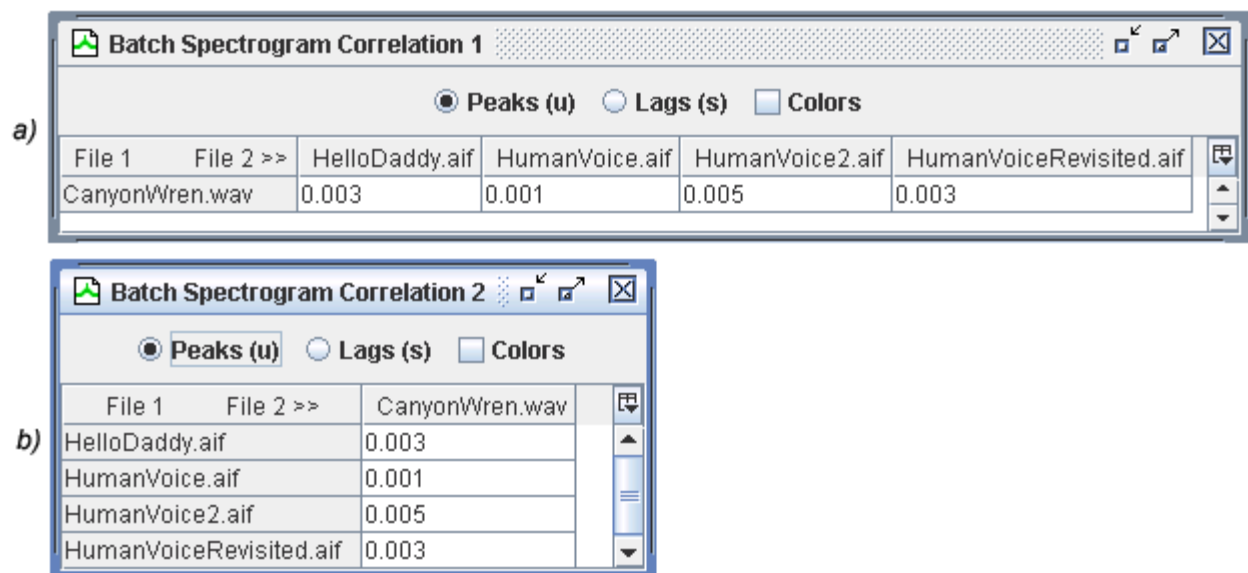


Figure 9.24. Correlations between one file and a folder of files. **(a)** the file was specified as input 1 and the folder as input 2. **(b)** the folder was specified as input 1 and the file as input 2.

Correlator example

This example shows, step by step, how to run a batch correlation on a set of provided files and compares two specific correlations to demonstrate how the correlator works.

Open sound file and selection table

Shipped with Raven, you will find an audio file (located in the Examples folder) named BlackCappedVireo.aif and a text file containing selections (located in the Selections folder) named BlackCappedVireo.selections.txt. Opening the sound file and then opening the selection table should give you the following image (Figure 9.25).




Figure 9.25. Recording of a Black-capped Vireo, showing 44 selections within the recording (achieved by opening the BlackCappedVireo.aif file along with the BlackCappedVireo.selections.txt file).

Save selections into a new folder

Next, choose File > Save All Selections in Current Table As... from the menu and click on the File Names tab. In the Directory field, type `"/BCV"` after the default "Selections" entry, creating a new folder named BCV (see Figure 9.26). Also, adjust the Sound File entry to read `"BCV<ii>.aif"` (this will name the selection files BCV01.aif, BCV02.aif, BCV03.aif and so on.) Skip over the list file name and annotations box, as they are not important for this example. For information on these options, see "Saving all selections" in Chapter 6 (page 161). Finally, clicking the Save button will complete the process, saving the named selections into the newly created folder.




Figure 9.26. The Save All Selections As dialog box with fields completed as described in “Save selections into a new folder” on page 247.

Run batch spectrogram correlation After saving all of the selections to the BCV folder, you can now run a batch correlation of all the files in the folder against each other. (Note that you could also choose to correlate a single file against all the files in the folder. However, this example uses the BCV folder as both input 1 and input 2.)

To perform the batch correlation, choose Tools > Batch Correlator... which will display the appropriate dialog box. For both input files, choose Browse... and select the new folder you created, BCV (it should be in the Selections folder in your Raven directory). Next, adjust the name of the output text file to read “BlackCappedVireo.specCorr.txt” and make sure that Spectrogram (under the correlate section) and Normalize (under the options section) are selected. After verifying your inputs with Figure 9.27, select OK to begin the correlation.


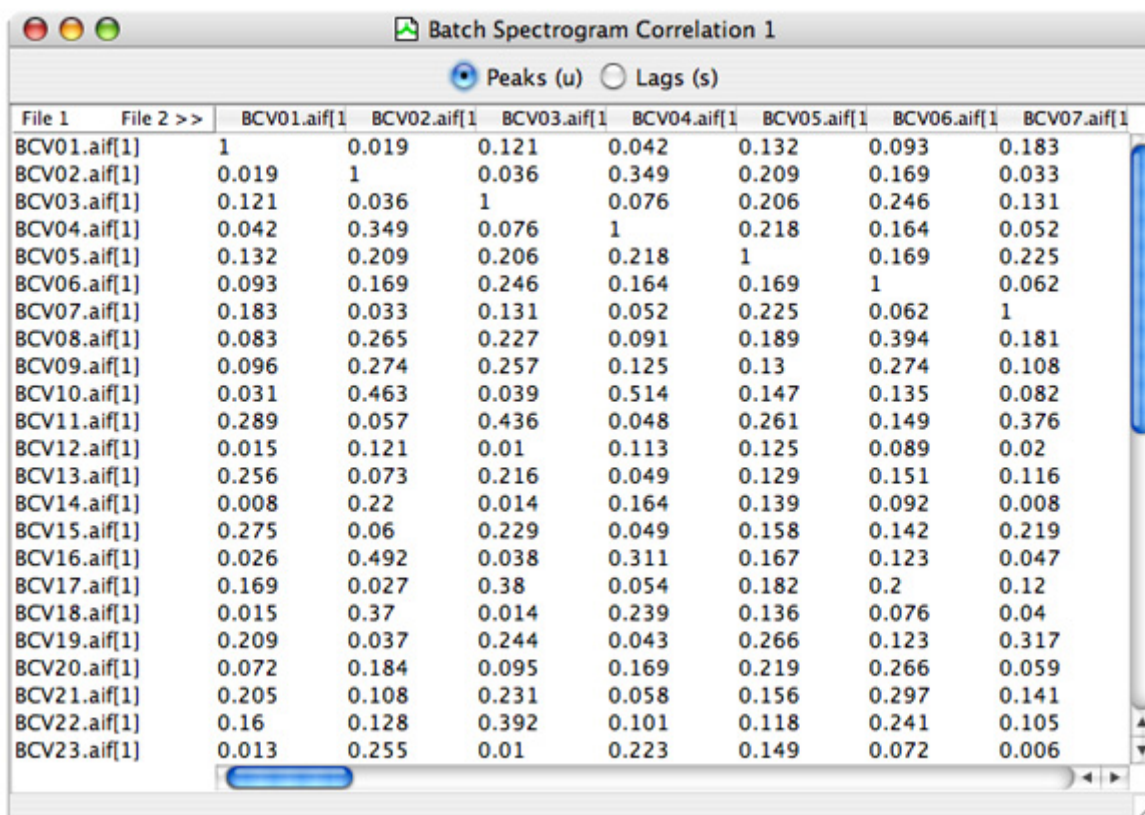


Figure 9.27. The Batch Correlation Configuration dialog with information filled in to correlate the selections contained in the BCV folder against themselves.

The correlation table After running the correlation, you should get a 44 x 44 table of correlation values (Figure 9.28) where each cell in the table represents the correlation of one of the files against another of the files.



File 1	File 2 >>	BCV01.aif[1]	BCV02.aif[1]	BCV03.aif[1]	BCV04.aif[1]	BCV05.aif[1]	BCV06.aif[1]	BCV07.aif[1]
BCV01.aif[1]		1	0.019	0.121	0.042	0.132	0.093	0.183
BCV02.aif[1]		0.019	1	0.036	0.349	0.209	0.169	0.033
BCV03.aif[1]		0.121	0.036	1	0.076	0.206	0.246	0.131
BCV04.aif[1]		0.042	0.349	0.076	1	0.218	0.164	0.052
BCV05.aif[1]		0.132	0.209	0.206	0.218	1	0.169	0.225
BCV06.aif[1]		0.093	0.169	0.246	0.164	0.169	1	0.062
BCV07.aif[1]		0.183	0.033	0.131	0.052	0.225	0.062	1
BCV08.aif[1]		0.083	0.265	0.227	0.091	0.189	0.394	0.181
BCV09.aif[1]		0.096	0.274	0.257	0.125	0.13	0.274	0.108
BCV10.aif[1]		0.031	0.463	0.039	0.514	0.147	0.135	0.082
BCV11.aif[1]		0.289	0.057	0.436	0.048	0.261	0.149	0.376
BCV12.aif[1]		0.015	0.121	0.01	0.113	0.125	0.089	0.02
BCV13.aif[1]		0.256	0.073	0.216	0.049	0.129	0.151	0.116
BCV14.aif[1]		0.008	0.22	0.014	0.164	0.139	0.092	0.008
BCV15.aif[1]		0.275	0.06	0.229	0.049	0.158	0.142	0.219
BCV16.aif[1]		0.026	0.492	0.038	0.311	0.167	0.123	0.047
BCV17.aif[1]		0.169	0.027	0.38	0.054	0.182	0.2	0.12
BCV18.aif[1]		0.015	0.37	0.014	0.239	0.136	0.076	0.04
BCV19.aif[1]		0.209	0.037	0.244	0.043	0.266	0.123	0.317
BCV20.aif[1]		0.072	0.184	0.095	0.169	0.219	0.266	0.059
BCV21.aif[1]		0.205	0.108	0.231	0.058	0.156	0.297	0.141
BCV22.aif[1]		0.16	0.128	0.392	0.101	0.118	0.241	0.105
BCV23.aif[1]		0.013	0.255	0.01	0.223	0.149	0.072	0.006

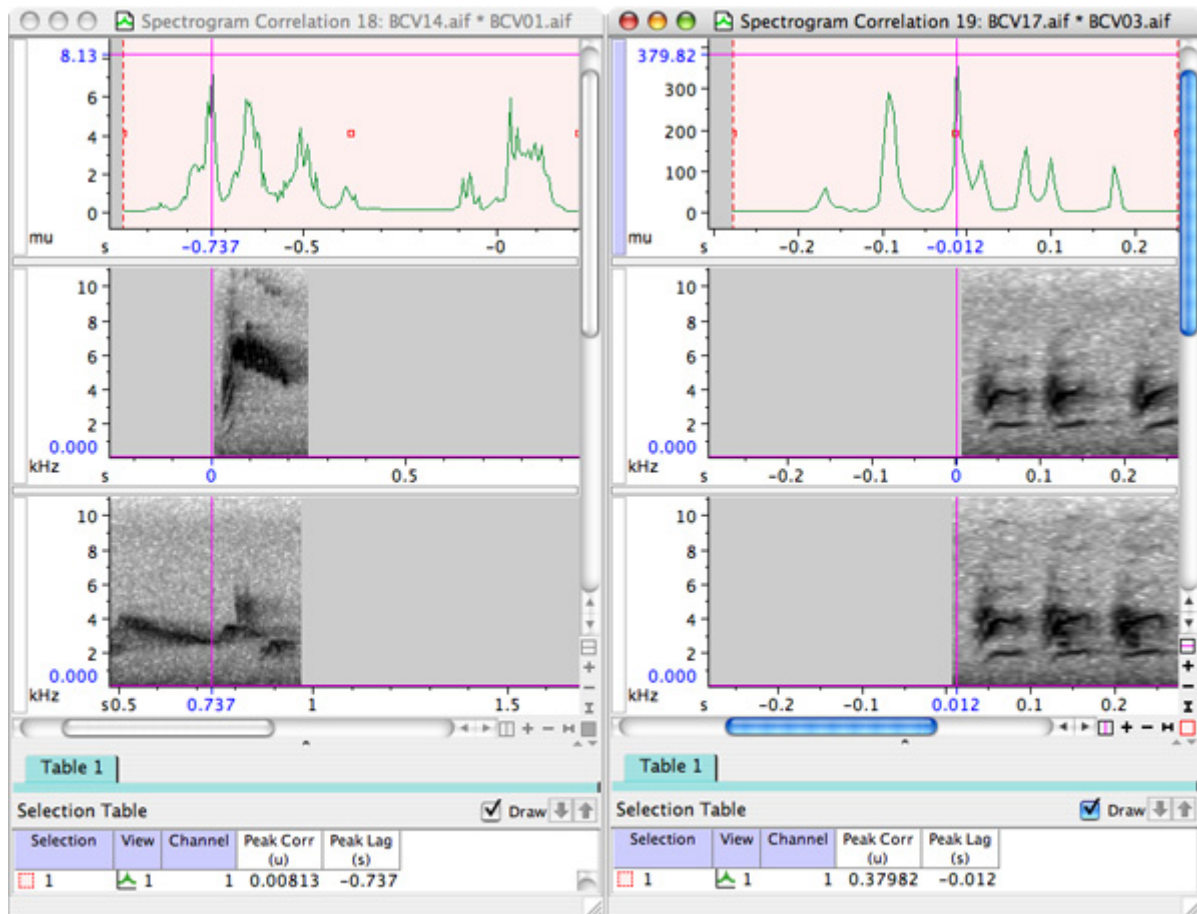
Figure 9.28. A portion of the correlation table from the previous example (correlating all files in BCV against each other).

You can view individual correlation functions by double-clicking in any cell showing a peak value (or lag value). Double-clicking the boxes for BCV14 * BCV01 and for BCV17 * BCV03 displays two sound windows. Each shows the correlation function view (top) as well as the two spectrogram views for the sounds that were correlated.

As shown in [Figure 9.29](#), the correlation peak values and peak lags are marked with the magenta frequency position marker and time position marker, respectively. In each window, select the entire correlation function, and then display the selection table. You can format the table to include only the measurements for Peak Correlation and Peak Lag. For a refresher on working with selection tables, see [“Selection Tables” in Chapter 6](#) (page 150).

After doing this for both sound windows, you should have something similar to [Figure 9.29](#). This example demonstrates the difference between correlation functions with high and low peak correlation values. In this example, the correlation values seem to correspond with the similarity between the calls being compared. While this is an important example, it is essential to remember that not all results will be this evident; not all

correlation values will correspond to visual similarity/dissimilarity of spectrograms.



a)

b)

Figure 9.29. (a) Linear power spectrogram correlation (top view) of BCV14 (middle view) with BCV01 (bottom view) with a low peak correlation value of 0.00813 (see selection table at bottom of window). (b) Linear power spectrogram correlation (top view) of BCV17 (middle view) with BCV03 (bottom view) with a high peak correlation value of 0.37982 (shown in selection table at bottom of window).