SoundWriter 2.0 Manual

John W. Du Bois

Overview

SoundWriter (available free from **http://www.linguistics.ucsb.edu/projects/transcription**, for Windows only) is software designed for researchers who work with spoken discourse. SoundWriter helps you link a transcription of spoken discourse to the corresponding digitized audio. The result is a transcription you can hear: when you click on a line of transcription, you hear the words as originally spoken. Beyond this, SoundWriter provides a number of further aids to the transcriber, such as tools to help in accurately identifying the precise beginning and end time for intonation units, and for working with intonation units in overlapping speech, a common feature of naturally occurring spoken interaction. SoundWriter also displays a unique TimeLine that shows the temporal relation of adjacent intonation units, including those involved in complex patterns of overlapping turns at talk.

The objective of this tutorial is to help you learn how to use SoundWriter to link your discourse transcriptions to the corresponding digitized audio. You will learn how to import a transcription file into SoundWriter, and will begin to learn the basics of using SoundWriter to align each line of your transcription with the corresponding segment of sound in your audio file.

Caveats

SoundWriter should be considered beta software. It has some known bugs; it is not as user-friendly as commercial software; it lacks an undo function; it crashes from time to time; and it is possible to lose your hard-earned work of aligning a transcription if you don't save your data properly (using SoundWriter's quirky approach to saving data—see below).

That said, SoundWriter does introduce some cutting-edge features that are designed to serve the linguistic researcher in unique ways. SoundWriter has innovative features that were introduced a decade ago that still have not been implemented in any other software we have seen. Ideally, SoundWriter would be rewritten from the ground up in a new programming language (or re-built as an extension to an existing free and open-source program such as Elan, EXMARALDA, PRAAT, or Audacity), but we do not have the resources for this. What SoundWriter can offer is a prototype which presents a certain vision of a discourse research tool, designed to help the researcher hear and visualize relationships between utterances in conversational interaction.

Install SoundWriter

1. Download SoundWriter from the website of the UC Santa Barbara Linguistics Department. Save the downloaded file to your computer's hard drive—for example, to the desktop.
2. You should have a file named SW.zip on your computer. Open it and start the installation by double-clicking on SETUP.EXE.
3. The SoundWriter installation process will begin. By default, SoundWriter will install files in three directories on the C: drive of your computer: C:\SW, C:\RUNTIME, and C:\CORPWAV. It is best to stick with these default locations for SoundWriter files, since SoundWriter is rather inflexible about where it expects to find the files it needs.

Start SoundWriter

1. Start SoundWriter: From the Windows **Start** menu, choose **All Programs** > **CSAE SoundWriter 2.0**.
2. SoundWriter will open with a sample transcription already loaded in the box at the top of the SoundWriter screen.
3. You should be able to click on any line of the transcription and hear the corresponding recording.
4. If you don't hear sound when you click on a transcription line, check to make sure your sound card and sound system are working, and that the necessary files are installed in the correct location (see above).

Experimenting with SoundWriter

1. The best way to learn about SoundWriter is to click on the various buttons and see what happens. The safest way to do this is to use the transcription and audio files that came with SoundWriter for your initial experiments, so you don't have to worry about losing your own hard-earned alignments of your transcription data.
2. To begin with, focus on the text box at the top, which contains a transcription in which each line represents a separate intonation unit. We will also learn about the brightly colored bars and colored buttons that appear beneath the text box. (For now we can ignore the gray area on the right side of the screen, since these are mostly just default settings that don't need to be modified at this point.)
3. Click on a line of transcript. Immediately you will hear the sound corresponding to that line. Also, the TimeLine changes when you click on a new line.
4. You can use the **down arrow** to step through the lines, playing them one at a time.
5. The **TimeLine** presents a graphical display with colored bars representing the currently selected line of text (i.e. the current intonation unit) plus the two lines before it and the two lines after it. The bar representing the currently selected line in the transcript is always in yellow, and the other four bars each have their own distinctive color, as you can see.
6. Each time you click on a new line, all five colored bars in the TimeLine are automatically refreshed and resized to represent the relative temporal duration of the spoken utterances they correspond to.
7. Note that the iconicity of the TimeLine bars is both visual and auditory: each bar is a live button that plays the corresponding sound, in addition to representing the temporal duration. Click on a TimeLine bar to hear the sound.
8. Next, try out the buttons in the **Playback** section. In addition to allowing you to play the currently selected text line, these buttons let you play the previous or next line, without changing which text line is currently selected. You can also play two adjacent lines with a single click—useful for checking whether it is really two intonation units, or should be retranscribed as just one intonation unit.
9. Other functions of interest in the Playback section include **Turn**, which plays the entire turn of the current speaker, and **Screen**, which plays the material currently visible on the screen (in the text box).
10. If you want to hear a line repeated several times, set the **Loop** number to the number you are interested in, and click Loop.
11. The big red **Stop** button is useful when you want to stop audio playback.
12. If you scroll down to where there is some overlap between speakers, you can see that the TimeLine is especially useful in displaying temporal relations between intonation units

involved in overlapping speech. Click on a few different transcript (text) lines that contain overlap brackets [ ], and watch what happens to the TimeLine display. Then click on different bars in the TimeLine display to get an auditory picture of what the overlap is like.

13. Look at the buttons in the **Overlap** section. These buttons automatically reconfigure themselves to graphically display the iconic relations between overlapping intonation units (the current unit in yellow, and the previous unit in blue).
14. Use the buttons in the **Boundary** section (just below the TimeLine) to change the start and end of an intonation unit. (These buttons automatically modify the timecodes corresponding to the start and end of the intonation unit.) Note that these buttons do modify the data, so before you start working for real on your own data, you will want to understand well what they do. Try them out now with this test file.

Edit functions

1. Remember, SoundWriter has no **Undo** function, so think before you act. (Some functions ask for confirmation first, others do not.)
2. The **Join** function joins the current line to the previous line. It merges the text of the two lines, and also adjusts the start and end timestamps accordingly. (It asks for confirmation first.)
3. The **Transpose** function moves the current line, placing it before the previous line. It works only when the lines are attributed to different speakers. (It asks for confirmation first.)
4. The **Time** function lets you type in the timestamp information directly. Because SoundWriter handles time codes automatically, you shouldn't need to use this.
5. The **Go to Line** function lets you jump to a text line by typing in the line number (not used very often).
6. The **Speaker** function lets you reassign the current intonation unit to a different speaker.
7. **Split Line** lets you split the current intonation unit into two new intonation units. (It asks for confirmation first.) This function simply makes two copies of the current line, and then lets you adjust the text and the timestamps for each line. First, use the **Text** button to edit the text of each line. Then use the various **Boundary** buttons to adjust the start and end times for the audio.
8. **Delete Line** deletes the current line.
9. **Add Line** inserts a new line.
10. **Delete Comment Line** deletes the *following* line (whether or not it's actually a comment line).
11. **Save Text** exports the text of your transcription (including timestamps) to an external text file, but this is *not the best way* to save your work in SoundWriter. For the right way to save your work, see the next section.

Preparing your transcription for importing

1. Although it is theoretically possible to do your transcription from scratch within SoundWriter, SoundWriter is better used for light editing of a previously prepared transcription (and of course for time-aligning audio, its main strength). It is far more convenient to prepare at least the first draft of your transcription using other audio software, such as VoiceWalker (available from the UC Santa Barbara Linguistics Department website), which you may want to use in combination with your favorite word

processor. So, assuming you have a pre-existing transcription prepared already (at least a first draft), you will need to import it into SoundWriter.

2. If you are just using SoundWriter primarily to explore the research concepts and methods it implements, or to consider the design possibilities for this kind of software tool, you may decide you don't want to get involved in the process of importing transcriptions. If so, you can still conveniently get the idea of what SoundWriter does by working with the transcription and audio demo files that came with the software.
3. If you decide you do want to import a transcription file, you will need to prepare your transcription for importing by doing the following:
4. Make a copy of your file first. You will be making a few changes to this file in preparation for importing, so it is wise to preserve your original file.
5. Remove ALL blank lines. There should be no blank lines (lines containing no characters) at all in your transcription file. Don't use double spacing. Check the beginning and end of your file to eliminate blank lines there, too.
6. On the very last line of your transcription, there should be NO CARRIAGE RETURN. (This would create an invisible "blank line" at the end of your file.)
7. Avoid comment lines if possible. All lines in the transcription should correspond to real-world, real-time audible/visible phenomena (transcriptions of words, vocal noises, actions, etc., but not your commentary on the transcription, for example).
8. If you have lines containing interlinear glosses, or other analytical coding, you will want to remove them.
9. When you have finished preparing your file, save it in "Text Format" or "Text Only" or "Generic ANSI" or some such simple text format. If you have special characters (e.g. non-English characters) in your transcription, you may need to experiment a bit to find which format best preserves the work you have done.
10. Now that you have prepared a text-only version of your transcription, you may want to use WordPad (a Windows accessory) for the next few steps.
11. Using WordPad or other similar text editor, select (highlight) all of the text you wish to import. This will normally be your entire transcription file. (You may use the Edit function **Select All** or **Ctrl-A** for this.)
12. Select **Copy** from the Edit menu of your word processor or text editor (or use **Ctrl-C**).

Importing your transcription into SoundWriter

1. Start your copy of SoundWriter (e.g. by double-clicking on it). The file as originally installed should have the name **SW.exe**. Each person who uses SoundWriter should have their own copy of it, so you may find it useful to make an additional copy of this file and give it a new name, like **SW-*name*.exe** (where *name* corresponds to your name). (SoundWriter is unusual software in that you will save your data in the .exe file itself, so this is why you want to have your own copy for your own files.)
2. In SoundWriter, click on "New Transcript". A green screen will appear. There are instructions on the bottom of the screen about what to do next. You should follow these instructions, except for a few updates that appear on this handout.
3. You can ignore the section called "Preparing to Import Text", since that's what you just did by following the instructions above.
4. Follow Steps 2-5 under "Importing and Preparing Text".
5. The instructions refer to "Sample numbers", which appear initially as rows of zeros in the format 00000000. (Eventually, as you work on your transcription using SoundWriter,

these zeros will be replaced by time codes in the format 00:00:00:00, corresponding to HOURS:MINUTES:SECONDS:HUNDREDTHS.)

6. Be prepared to tell SoundWriter the location of your audio file (i.e. filename and directory). Your file name should contain no more than 8 characters, and no spaces, plus the extension **.WAV**. (The same restriction to 8 characters may apply to the directory names, but you shouldn't have to worry about this.)
7. Also, be prepared to tell SoundWriter where you want it to save the SoundWriter text file that it will create. Again, your file name should contain no more than 8 characters, and no spaces. It should end with the extension **.SW**. (SoundWriter is rather inflexible regarding the location of its files, and may insist on placing the SoundWriter text file in the C:\SW file on your local (C:) drive.)
8. Once you have imported your transcription into SoundWriter in this way, you should not move or rename your audio file, nor your SoundWriter text file. This is because SoundWriter is rather inflexible about where it expects to find your files. (If for some reason you have to move your files, you will need to reimport the data—sorry!)

Aligning Audio with a Newly Imported Transcription

1. Assuming you have learned about how to use the various SoundWriter functions by practicing with the test file, as described above, you can now use the Boundary, Overlap, and other groups of functions to align your transcription with the corresponding sound.
2. It is important to *start by aligning the first line of your transcription* with the beginning of your audio file.
3. You should then proceed to work through your transcription *in sequence*, one line at at time, progressing from the first line to the second, second to third, and so on to the end of the transcript.
4. Note that when you click on a new transcription line, SoundWriter will try to guess the start and end boundaries of that line, based in part on the boundaries you have assigned to the previous line. This is why it's useful to go through the transcription from beginning to end, one line at a time: SoundWriter can make better guesses that way.
5. Press **Guess** to make SoundWriter guess the boundaries of the currently selected unit. (Notice that SoundWriter will automatically guess whenever it moves to a line that has no previous timecodes for the boundaries, i.e. a line that has only zeros instead of timecodes. So when you move through the transcript in sequence, you will normally not have to press the Guess button; it's automatic.)
6. When you have your approximate start and end boundaries as guessed by SoundWriter, you can then refine them. Using the Boundaries section, click on the left and right **arrowheads < >** labeled **Beginning** and **Endpoint**, depending on which boundary of the current intonation unit you want to change.
7. If you click on **Extend**, this will begin playing at the beginning of the current intonation unit and keep on playing until you press the button again, to signal that's where you want the intonation unit to end. SoundWriter will then play back the intonation unit you have just defined, so you can hear if it ends in the right place.
8. The **Snap Left** button is useful for when you already have specified the correct endpoint for the previous unit, and want to make the beginning of the current intonation unit match the previous endpoint.
9. The **Snap Right** button performs a similar function for marking the endpoint of the current intonation unit, making it match the start of the next intonation unit.

10. The **Snap Skip** check box modifies the behavior of the Snap Left and Snap Right buttons slightly, making them skip one intonation unit to match up with the intonation unit two lines back, or two lines after.
11. Remember to **save your work often**. *Be sure to always use the special procedure described below* for saving your data.

Saving Your Work

1. SoundWriter has an unusual procedure for saving. To securely save all of your work, you must fully exit the program—and give the correct answer as you exit. You can do this be clicking on **Quit** (in the lower right corner of the screen), or by exiting the program in the usual Windows fashion. You will then see a dialog box that says "**Save current changes to: SW.EXE?**" (or whatever name you have assigned to the copy of SoundWriter you are using, e.g. **SW-yourname.exe**). *You should answer **Yes** to this*, in order to save the work you have done in the current session using SoundWriter. (Yes, you will be changing the .exe file itself—most unusual—but that's just the way SoundWriter saves its data.)
2. *Save early and save often!* Aligning a transcription with audio is fairly labor-intensive, even with the assistance of SoundWriter, so you will want to save your work often. (Also, SoundWriter is beta software and does have some bugs; it may crash at unexpected times.)
3. The bottom line is: You should fully exit from SoundWriter *frequently*—say every fifteen minutes—in order to avoid losing your work. And remember to say *yes* each time you do this, in order to save your changes as explained above. To continue working, you then need to restart the program (in the standard way, e.g. by double-clicking its icon).

The Grey Area

1. The gray area on the right half of the SoundWriter screen contains a variety of items, most of which are settings that can be modified by the user. Most of these items are self-explanatory, and the default values don't really need to be changed much, so users don't need to be use this area as much as they do the more colorful parts of the interface.
2. **Resolution** is the one item that users may wish to adjust fairly often as they work on time-aligning an intonation unit. Choosing a setting from among the five options in the Resolution box allows the user to determine how large a change will be made by the various Boundary buttons when they are pressed (and also the Overlap buttons). The Resolution can be set so that all buttons that modify the start and end times for the current intonation unit adjust the timestamps by as little as a hundredth of a second, or as much as half a second (shifting the relevant timestamp to the left or right, as the case may be, by the amount specified in the Resolution).
3. The **Walk** button, and the **Step**, **Context**, and **Continuous** buttons next to it, perform functions similar to those of VoiceWalker. VoiceWalker is a program that is simpler, easier to use, and more stable than SoundWriter. And it is specifically dedicated to just these functions (along with a handful of closely related functions), so all in all, VoiceWalker implements these features better. So if you want the Walk function—well worth having—in a more user-friendly package, don't use it here, just download VoiceWalker (available free from the UC Santa Barbara Linguistics Department website).

[Revised 30 July 2006]