

Sociolinguistics Laboratory

revision history: 10/2016 ABW, 1/2012 ABW (v. 4.0), 5/2006 JCMiller

An overview of ELAN – Linguistic Annotation Software

ELAN stands for EUDICO Linguistic Annotator, and EUDICO for European Distributed Corpora Project. ELAN is an annotation tool that allows you to create, edit, visualize and search annotations for video and audio data. Developed at the Max Planck Institute for Psycholinguistics, Nijmegen, The Netherlands, ELAN provides a technological basis for the annotation and exploitation of multi-media recordings in service of a range of language-related research projects, such as study of:

- spoken language
- signed language
- gestures
- any media corpora with video and/or audio data, for purposes of annotation, analysis, and documentation (i.e., subtitling)

Sociolinguists will mainly use ELAN for transcription of conversational data, search, and descriptive statistics about speaker productions. It is more powerful than simple dictation software in its interoperability with signal analysis software, like Praat. However, it has many functions sociolinguists may not be interested in. Therefore, analysts will want to balance the costs of learning ELAN against the values of ELAN for a particular project:

- search (particularly in long transcription files)
- linking of annotations to other annotations
- interoperability with Praat
- consistent application of metadata ontologies (e.g., Dublin Metadata) (See Sec. 5)
- manipulate the contents of a transcription

In some cases, e.g., where you have a short file, or reading list, ELAN is unnecessary: use Praat!

In this document:

1. Supported Functions..... p. 2
2. Getting Started p. 2
3. The Anatomy of ELAN p. 3
4. Your Work Session p. 4
5. Audio Player Control Reference Chart.....p.17
6. Tier (Stereo)Types: not what you think! p. 20

1. Supported Functions in ELAN:

- display audio and/or video signals, together with their annotations
- link annotations to other annotations
- input different character sets (foreign language support)
- import and export between ELAN, Praat, Shoebox and other software
- time link annotations to media streams
- unlimited number of user-defined annotation tiers
- export as tab-delimited text file or Praat Text Grid
- search

2. Getting started

Obtain ELAN (free from <http://www.lat-mpi.eu/tools/elan/>)

NOTE: ELAN must reside in a folder with its extensions. You can create an alias (shortcut) that will sit elsewhere on your desktop, but if you separate the main software from its extensions, it will not launch.

NOTE: It appears that ELAN does not support foot pedal peripheral devices. Your only option is to use the audioplayer controls onscreen to control audio playback.

One or more media files:

- .wav
- .mpg
- .mov
- .avi

An annotation file:

- .eaf (EUDICO default annotation format – created after you load your media and set up your tier template, etc.)

.txt
 .trs (from Transcriber AG)
 .cha (from CHILDES)

Note: all information generated in ELAN work sessions is saved to annotation files, never to your media file.

3. The anatomy of ELAN

ELAN main screen shows the Menu bar, the Media Player options, and up to 7 Viewers while you are working (indicated by numbers in Figure at right):

1. video viewer (not pictured)
2. waveform viewer
3. timeline ("Annotation Density Viewer"*)
4. interlinear gloss viewer (not pictured)
5. grid viewer tab
6. text viewer tab
7. subtitle viewer tab

* A couple of things about the Annotation Density viewer:

- This is a "timeline" bar that allows users to navigate through the whole media file. The length of this viewer always corresponds to the whole media file, so e.g. by clicking in the middle you will always go to the middle of the media file. The selection is represented as a small grey bar.
- It will, as you work, show small vertical lines indicating places in the media file where annotations have been added. When you start a new ELAN session, it will have no "tickmarks" along its length. See how many annotations are concentrated at a particular moment of the time (the Annotation Density). The more annotations available for a particular moment, the more the Annotation Density bar is filled. This can be useful to track places in the media file that still have to be annotated.

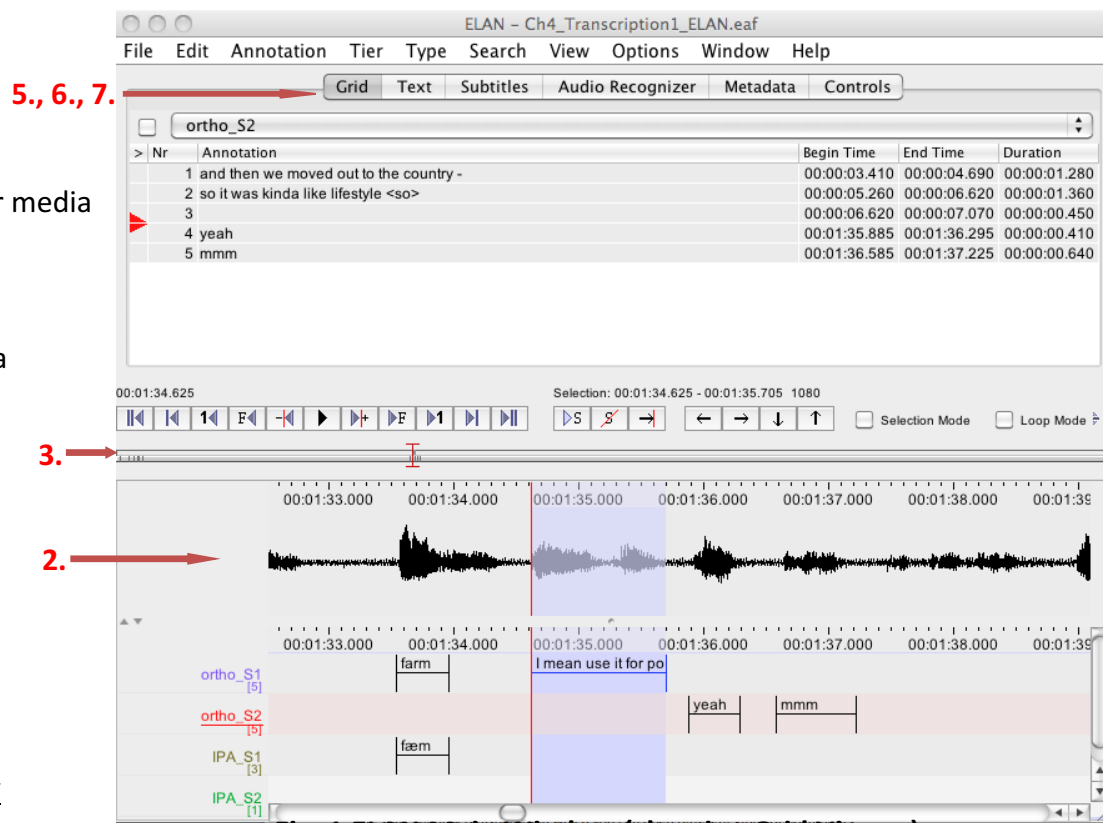


Fig. 1 ELAN Main Window (showing Grid Viewer)

4. Your Work Session

To start ELAN:

1. Double-click on the ELAN icon (on your desktop, or contained within the Java Web Start icon on your desktop). The start window is displayed in which you can open different kinds of documents.

2. Click on the File menu (see Figure 4.1¹)

The main options in the File menu are:

- New (Section 4.2.1)
- Open (Section 4.2.3)
- Import (Section 4.2.15, Section 4.2.16, Section 4.2.17, Section 4.2.18, Section 4.2.19, Section 4.2.20 and Section 4.2.21)
- Exit (Section 4.2.42)

3. To begin your work, Click on:

- Open... if you have a pre-existing ELAN file (*.eaf) (Section 4.2.3)
- New... if you have a media file (e.g. *.mpg, *.wav), but no annotation file (*.eaf, *.txt) (Section 4.2.1).

4. Browse to the directory that contains your media files (e.g., audiofile you desire to transcribe)

5. If you want to use a predefined set of tiers (a template), select the Template radio button and choose the template (i.e. *.etf) to be used:

6. Click OK to open the new annotation document; otherwise click Cancel to exit the dialog window without creating a new file. An ELAN window containing the new document appears.

What can I do to an audiofile in ELAN?

- playback
- tagging (e.g., for acoustic events, prosodic events, etc.)
- annotation (e.g., associate an orthographic transcription of a conversation)
- tokenization/segmentation

¹ Section reference numbers in this guide refer to the numbered sections of the online ELAN 3.8.1 manual

- zoom in and out
- open a portion of the soundfile in Praat (use the SendPraat function)


NOTE: Tapping the space bar will execute the last movement you initiated (unlike Praat, which restricts spacebar functions to play or pause).

What can't I do to an audiofile in ELAN?

Signal analysis (formant measures, etc.) BUT: YOU CAN SEND AN ELAN TRANSCRIPTION TO PRAAT AS A TEXT TIER FOR EASY ANALYSIS.

File > Export Multiple Files as... > Praat Text Grid...

Particularly useful Audiofile Navigation Buttons:



ELAN's play selection button: You'll typically want to use this one! Not the unmarked arrow in the left bank of controls, which plays out the entire file till you stop it.



Scroll 1sec right (ahead) in audiofile. Corresponding control for scrolling 1sec left (back).



Jump ahead to next scroll view (window) in audiofile.

Zooming into and out of the waveform:

The default zoom for the Waveform and the Timeline Viewer is 100%, corresponding to 10 milliseconds per pixel. The zoom can be changed simultaneously for both Viewers. Do either of the following:

1. Click with the **right mouse button** on either the Waveform Viewer or the Timeline Viewer. A dropdown menu appears.
 - a. Select desired percentage, or...
 - b. If you have a selection highlighted, choose Zoom to selection.
2. Use the **Scroll wheel** of the mouse:

- a. Place mouse cursor on the Waveform Viewer or the Timeline Viewer.
- b. Press the **CTRL** key and keep it pressed.
- c. Move the scroll wheel of your mouse. Moving down is zooming out and moving up is zooming in.

Creating Tiers

Your first step of transcription is to determine the number of layers or tiers of annotation you need for your project. This may include orthographic tiers (one for each of your speakers), gloss tiers, morphology tiers, phonetic tiers (with IPA symbols), word list numbers, etc.

Steps:

Click on **Tier > Add New Tier....** The Add Tier dialog window appears.

Define the tier attributes:

1. Enter the following information:
2. Go to Tier name. Enter the tier name.
3. The tier name is the name that is displayed in the Timeline, Interlinear and Subtitle Viewers.
4. Go to Participant. Enter the name of the participant whose utterance is being transcribed.
5. Go to Parent tier. Select a parent tier from the pull-down menu. (e.g., the IPA tier may be a child tier associated with the orthographic tier)

Your tiers will be visible in the area below the waveform, and each one can be independently selected, colors may be selected, etc. **You can also drag and move tiers to reorder them (or right-click to determine their ordering)!**

Creating a transcription

Now, the tier structures you need are in place, and you are ready to create annotations.

Annotation Mode: DEFAULT

Options > Annotation Mode

In this mode, there are two ways to add annotations: either from the pulldown menus, or from the waveform itself.

Note: Each annotation field you create must be saved after you enter it! (CMD+RETURN)!!!!

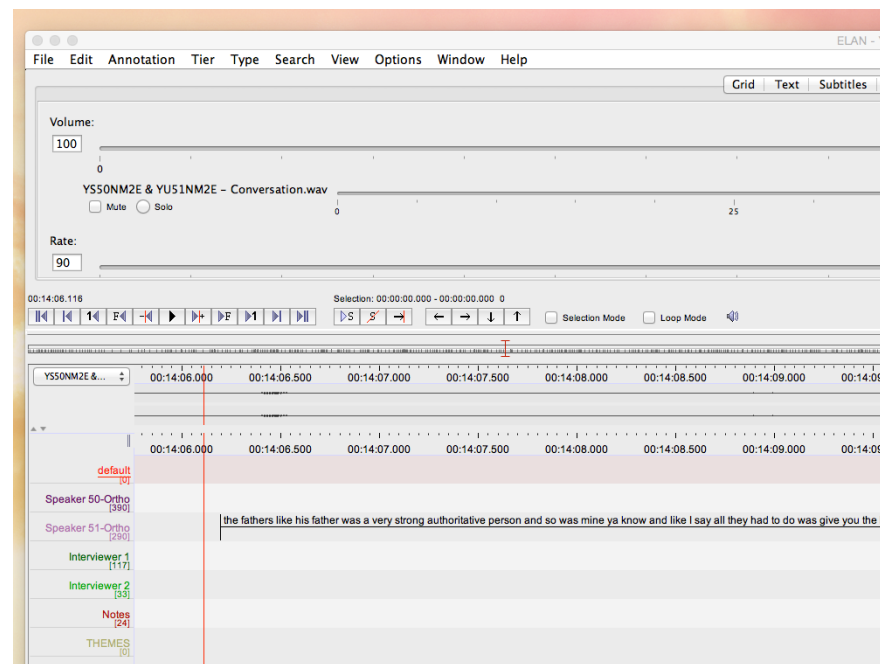


Fig. 2 ELAN's Annotation Mode

Pulldown menu method:

1. **Highlight a portion of the waveform.**
2. Select from menu **Annotation > New Annotation here** to create a time-aligned window **in the active tier**, in which you may type your annotation. A selection-length window, **called the Inline annotation editor (Fig. 2)** should now appear. If it appears in color (not white), click it once again to get a blank white box containing a text-entry crosshair cursor.
3. Tip: Click the boxes “selection mode” and “Loop Mode” to enable you to replay audio while you type. To stop looping, press “spacebar” or just click the “play selection button” (Sec. 3).
4. **Type annotation.**
Tip: You can enter text, but you can also enter IPA using any IPA character inputting system that works for your computer’s system.
Tip: If you are making annotations in a metadata tier, you may find it useful to refer to “Linguistic Stereotypes” (Sec. 5)
CMD + RETURN to save annotation

Waveform Method:

1. **Highlight a portion of the waveform.**
2. **Hover your mouse over the tier to which you desire to add an annotation.**
3. **Double-click on the selected region.** A selection-length window should now appear (Fig. 2). If it appears in color (not white), click it once again to get a blank white box containing a text-entry crosshair cursor.
4. **Type your annotation** into this box.
5. **CMD + RETURN** to save annotation.

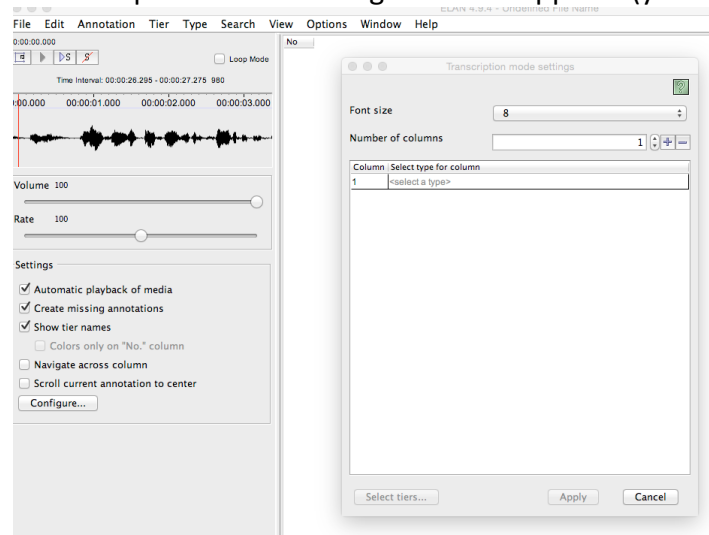
Annotations may be deleted by double-clicking to select the desired annotation, then right-click and select “**delete annotation**”.

Speed Transcription Mode: HIGHLY RECOMMENDED!


Options > Transcription Mode

In this mode (Fig. 3), ELAN shows you your pre-generated segments of audio/video, and you transcribe while it plays chunks for you. You won't see the waveform in this mode, so it's best to do segmentation first (into turns, by breath groups, or at pauses).

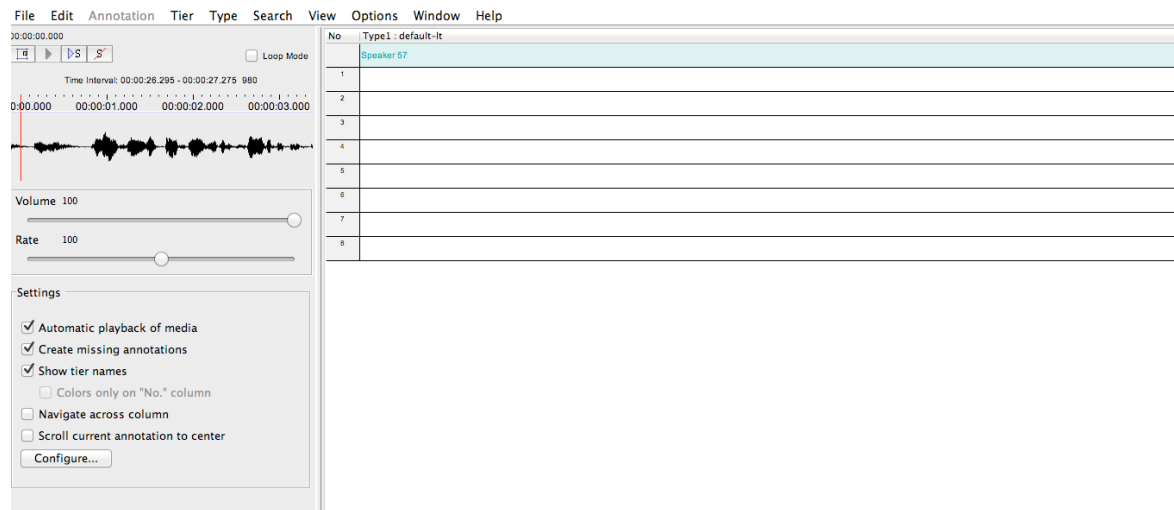
1. First, working in annotation mode (see above) or segmentation mode, segment your audio/video (creating EMPTY annotation fields).
2. Navigate to Transcription mode.
3. First, choose your settings in the Transcription Mode Settings box that appears (you will do this only once for each .eaf file):



- Choose a size for your screen font (8 is the default). 12 is recommended for easier viewing.
- Select a type for your column (ELAN gives you the types you selected when you created your tiers earlier. If you are only creating orthographic text tiers of type default-It, this is the only choice you will have, but you must select something).
- Select tiers by clicking on the **Select tiers...** box at the bottom left of this pane.
- The below window appears. You might choose to hide default-It at this time.



4. Click **Apply** to return to the previous pane. Click **Apply** again to return to the main Transcription mode editor.
5. Your Tier names (differentiated by color) now appear in the Transcription Mode editor, and a blank, numbered field appears corresponding to each annotation field you created in the Annotation editor in step 1. A mini waveform editor displays the stretch of audio corresponding to the numbered annotation fields.
6. **Click inside a blank annotation field** to begin playback. **Type** to transcribe. Press **TAB** to replay/pause the audio segment. Press **SHIFT+TAB** to play back the current annotation from the start. **Press ENTER** to save annotation and advance to the next field.



7. **Customize your experience by adjusting settings, including:** Activate/deactivate Loop mode, Adjust speech rate (faster/slower), Add missing annotations (ELAN creates fields for unselected segments of audio), Show/hide tier names, Scroll to keep the current field at the center of your window (ELAN will automatically advance the screen upward as you work).
8. Adding an annotation on another field. At this time, it appears that if you need to add an annotation field on another tier (e.g., if you find an instance of overtalking and you created an annotation field for just one talker's utterance), you must leave Transcription Mode, return to segmentation mode, and add an annotation field on the desired Tier.

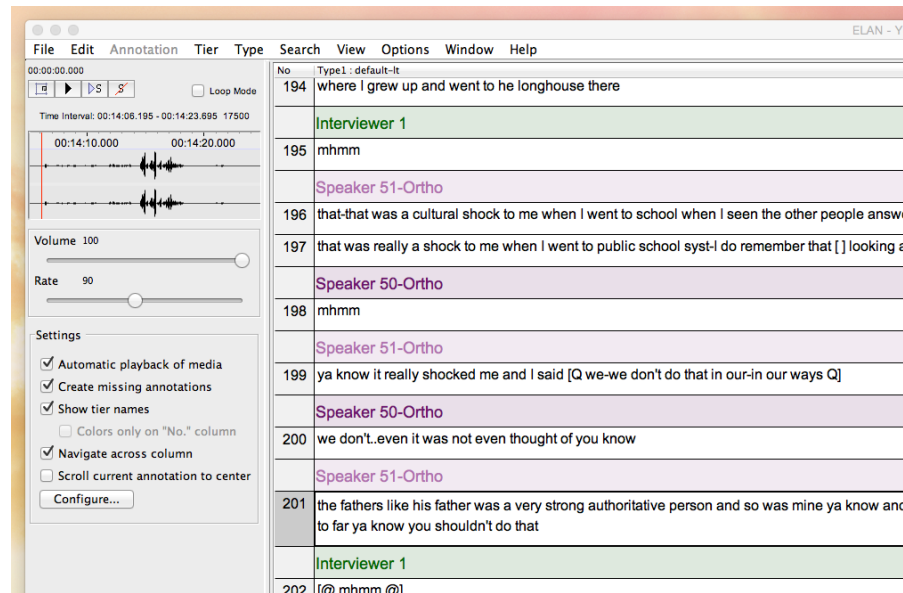


Fig. 3 ELAN's Transcription Mode

Other functions:

Note: You can Right-click or ctrl+click in either Transcription mode or Annotation mode!

Merging annotations

Sometimes, you will want to combine adjacent annotation fields. Right-click or ctrl+click on the waveform and select **Merge annotations**. (Options in the drop down menu will be **Merge with Annotation Before/After**.)

Deleting annotations

To delete an annotation, make sure you have selected the annotation field you desire to delete. Right-click on the annotation and select **Delete annotation**.

Splitting annotations

Sometimes, you will want to split annotation fields. There are two ways to do this:

- To split a annotation exactly in the center, select an annotation field and click on **Annotation > Split Annotation**. This will split your annotation exactly in the center and both the annotations will have the same values.
- To split an annotation at a specific point, select an annotation and right-click on the point where you want to split the annotation and select **Split Annotation**. This will split the annotation at the point where the right click is made.

The Grid Viewer

The Grid Viewer (below, Fig. 4) displays the content as well as the begin time, end time and duration of all annotations from a single tier. You can activate this Viewer through selecting the Grid tab next to the video window. You can use the Grid Viewer window **to navigate to a particular annotation** in the ELAN window. You have the following two options:

Working between waveform viewer and grid:

- a. In the waveform viewer, move the crosshair to the beginning of an annotation.

- b. The grid viewer window will change to show the selected annotation.

From Grid:

- a. In the Grid Viewer window, select the desired tier from the Grid Viewer pulldown menu.
- b. Click on an annotation.
- c. A red triangle appears next to the annotation in the Grid Viewer window, and the crosshair moves to the beginning of that annotation in the Waveform Viewer.

You can also use the Grid Viewer window to **Select a time interval** for playback.

- a. In the Grid Viewer window, click on the first annotation that you want to select.
- b. Shift-click to select a range of annotations. Or, keep the mouse button down and drag the mouse to another annotation. In the Grid Viewer window, all selected annotations are highlighted in light blue color. In all other windows, the corresponding time interval is selected and highlighted in light blue color (starting with the beginning of the first annotation and ending with the endpoint of the last).
- c. **SHIFT+SPACE** to play selection.

Here's what the grid viewer looks like:

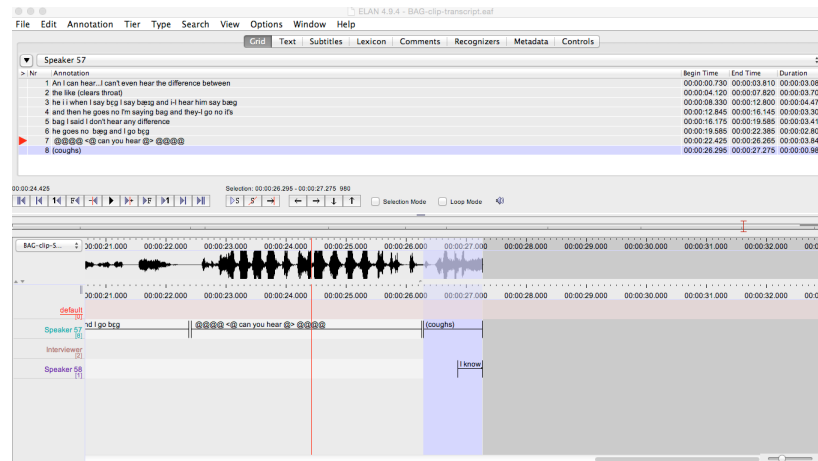


Fig. 4 The Grid Viewer

The Text Viewer

This works similarly to the Grid viewer. The text viewer shows all values of all annotations on a selected tier as ongoing text. Within this viewer, there are 3 kinds of highlighting shown:

- Text inside a red box: the value of the annotation on the selected annotation tier that matches the current point of time.
- Text inside a dark blue box: the active annotation (see also Section 4.3.9)
- Text with a light blue background: the selected time interval.

Here's what the Text Viewer looks like:

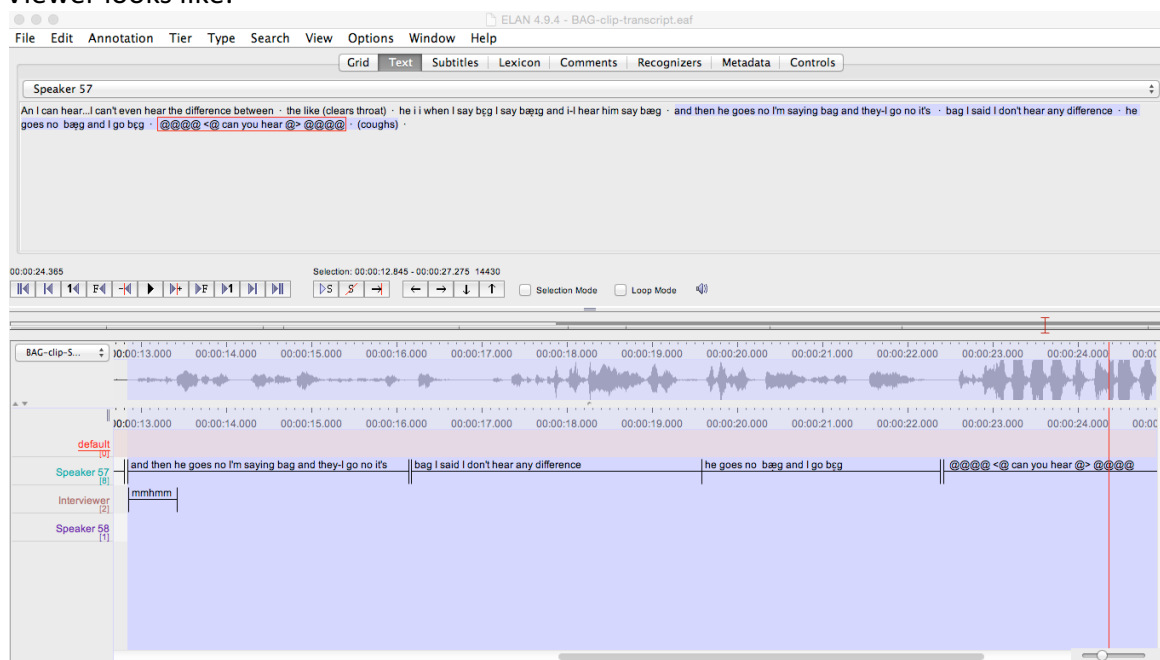


Fig. 5 The Text Viewer

Here's what the Text Viewer looks like with a video clip:

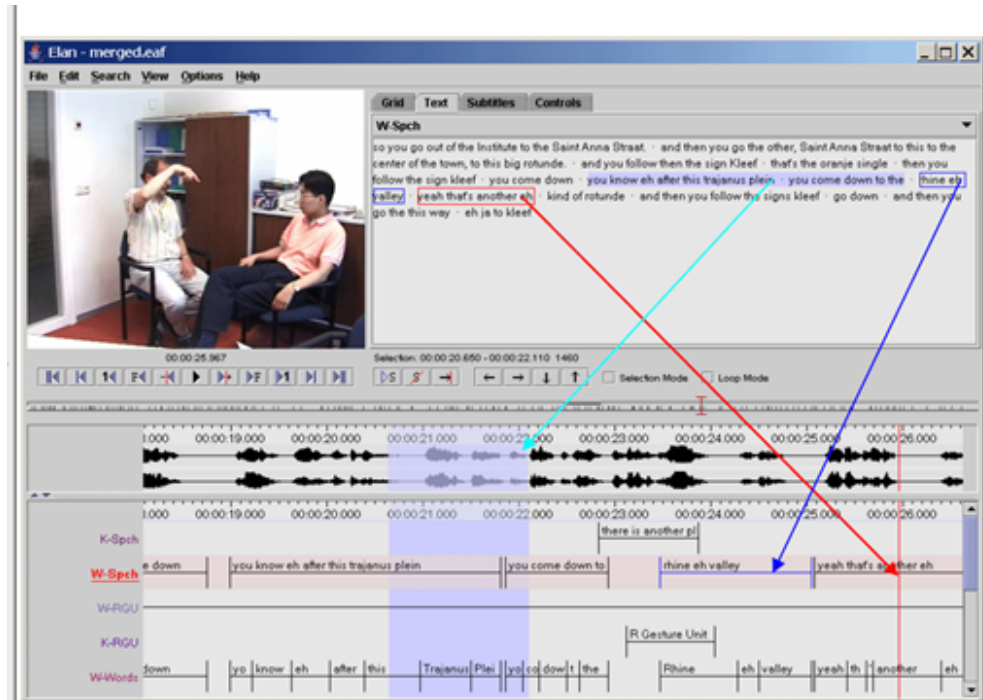


Fig. 6. The Text Viewer with a video clip

Ending an ELAN session

Saving a File

All documents can be saved as ELAN files (*.eaf, “EUDICO Annotation Format”). This includes documents that were created by ELAN itself as well as documents that were imported into ELAN from Shoebox/Toolbox, CHAT or Transcriber. **No changes are made to your media file.** To save a document as an ELAN file:

1. Either use the menu options:
 - a. Click on **File** menu.
 - b. Click on **Save** or **Save as**.
2. Or use the shortcut key **CTRL+S**.

Closing a file and Exiting ELAN

1. To close a file that is being viewed/edited without exiting ELAN use the **File > Close** menu option. Alternatively, you can also close an open file by clicking on the red box in the upper corner of the screen window or by pressing **CTRL+W** (PC) or **CMD+W** (Mac).
2. **File > Exit** to Exit ELAN.

5. Fig. 7: Reference Table of Audio Player Controls

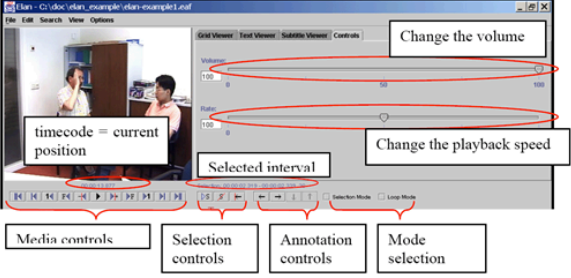
ELAN – Linguistic Annotator

Find: zoom

- 17 4.3.20. Changing the
- 4 4.3.21. Changing the
- 3 4.3.10. The Timeserie
- 2 ELAN – Linguistic Anno
- 2 Part II. USER'S GUIDE
- 1 4.3.12. The Media Pla

4.3.12. The Media Player options












With the Media Player options, you can control the playback of the file. The following options are available at the bottom and at the left side of the ELAN window:



The screenshot shows a video player window with several callouts: 'Change the volume' pointing to the volume slider, 'Change the playback speed' pointing to the speed control, 'Selected interval' pointing to the timeline selection area, and 'timecode = current position' pointing to the timecode display. Below the player, four boxes categorize the controls: 'Media controls', 'Selection controls', 'Annotation controls', and 'Mode selection'.

Figure 4.74. The media player options

Table 4.1. Media Controls

Icon	Meaning	Shortcut
	Go to the beginning of the video/audio fragment	CTRL+B
	Go to the previous scroll view (make the beginning point of the current timeline view the end point)	CTRL+PAGE UP
	Go back one second	SHIFT+LEFT
	Go back one frame (= 40 ms for PAL, 33.4 ms for NTSC)	CTRL+LEFT
	Go back one "pixel" on the timeline viewer (smallest unit, depends on the zoom factor of the timeline viewer, default value 10 ms)	CTRL+SHIFT+LEFT
	Start / Pause the playback	CTRL+SPACE
	Go to the next "pixel" on the timeline viewer (smallest unit)	CTRL+SHIFT+RIGHT
	Go to the next frame	CTRL+RIGHT
	Go to the next second	SHIFT+RIGHT
	Go to the next scroll view	CTRL+PAGE DOWN
	Go to the end of the media fragment	CTRL+B

ELAN – Linguistic Annotator

Find: zoom

- 17 4.3.20. Changing the
- 4 4.3.21. Changing the V
- 3 4.3.10. The Timeserie
- 2 ELAN – Linguistic Anno
- 2 Part II. USER'S GUIDE
- 1 4.3.12. The Media Pla

Table 4.2. Selection Controls

Icon	Meaning	Shortcut
	Play the selected interval.	SHIFT+SPACE
	Clear the selection.	CTRL+C
	Move the crosshair to the begin / end of selection	CTRL+/ or CTRL+SHIFT+K

Table 4.3. Annotation Controls

Icon	Meaning	Shortcut
	Go to the previous annotation on the active annotation tier	ALT+LEFT
	Go to the next annotation on the active annotation tier	ALT+RIGHT
	Go to the annotation above.	ALT+UP
	Go to the annotation below.	ALT+DOWN

Table 4.4. Selection Mode

Icon	Meaning	Shortcut
<input type="checkbox"/> Selection Mode	While playing, select an interval automatically	CTRL+K
<input checked="" type="checkbox"/> Loop Mode	Keep playing the selected interval (if used together with play selection)	CTRL+L

- The timecode can be displayed in the following formats:hh:mm:ss.ms (hours:minutes:seconds.milliseconds), PAL (hours:minutes:seconds.frames), NTSC (drop frame) (hours:minutes:seconds.frames) and msec (milliseconds). It can be used to access points in time (by clicking on it and entering a value in "hh:mm:ss.ms" format). To change the display format, right click on it and select a format from the menu. The SMPTE timecode formats PAL and NTSC only indicate the way time values are converted to frame numbers; this is independent of the actual video standard of the associated video(s).

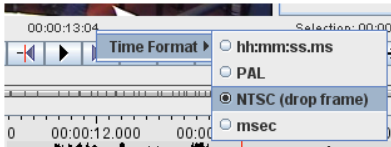


Figure 4.75. Timecode formats

- The sliders available when the Controls tab is selected allow you to control the playback rate and the volume.

Fig. 7 Audio Player Controls Reference Panel

6. Linguistic Stereotypes

It is possible to build up nested hierarchies, i.e., tier A can be the parent tier of tier B, and tier B can be the parent tier of tier C, etc. Each tier is assigned to a linguistic type (see also Section 5.3). A linguistic type denotes the linguistic data that is contained in the referring tier. Examples of names for linguistic types are utterances, words, orthography, phonetic transcription, PoS (part of speech), but any name can be used. Each linguistic type specifies a number of constraints that hold for all tiers assigned to that type. Such constraints are bundled into so-called ‘stereotypes’. The following stereotypes are currently available:

1. None: The annotation on the tier is linked directly to the time axis. The annotation is on an independent tier. Two annotations cannot overlap.
2. Time subdivision: The annotation on the parent tier can be subdivided into smaller units, which in turn, can be linked to time intervals. Note: there are no time gaps allowed: i.e., the smaller units have to immediately follow each other. An utterance transcribed on a parent tier can be subdivided into words, each of which is then linked to its corresponding time interval.
3. Symbolic subdivision: Similar to Time Subdivision, except that the smaller units cannot be linked to a time interval. E.g., a word on a parent tier can be subdivided into individual morphemes (which are not linked to a time interval).
4. Included in: All annotations fall within the borders of the parent tier. However, there can be gaps between the child annotations. E.g., a sentence with a silence can be split up into words, while the silence corresponds to a gap in the child annotations (i.e., the separate words).
5. Symbolic association: The annotation on the parent tier cannot be sub-divided further, i.e., there is a one-to-one correspondence between the parent annotation and its referring annotation. E.g., one sentence on a parent tier has exactly one free translation. Or one word has exactly one gloss.

Metadata Viewer

If you have a metadata protocol (IMDI, Dublin, OLAC, etc.), you may associate a metadata file with ELAN, which can then be completed as you work. Click the Metadata Tab, then **Select Metadata Source**. The file in which you maintain metadata for the audiofile and transcription will be listed in the box. (The sociolab file is called MetadataTemplate.rtf).

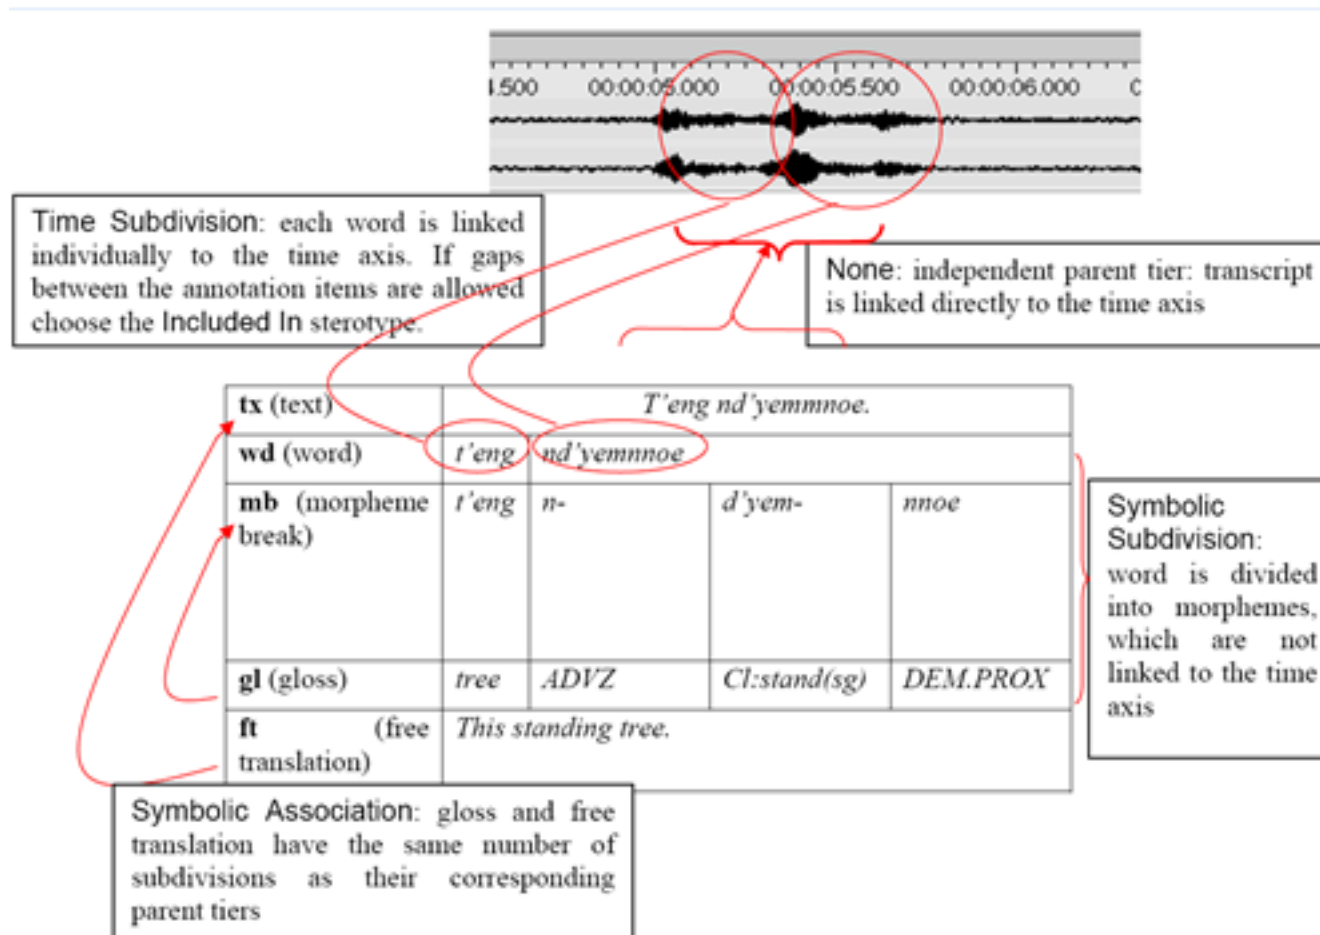


Fig. 8 Linguistic Types and dependencies for five tiers: text, word, morpheme break, gloss and free translation.