Aligning text to audio and video using ELAN

Institute on Field Linguistics and Language Documentation Workshop
June 22 – July 2, 2010
University of Oregon

Andrea Berez
aberez@umail.ucsb.edu

Christopher Cox
christopher.cox@ualberta.ca

Workshop information:

1   *Workshop description:*

This workshop is intended to provide a basic, hands-on introduction to developing multimedia language resources using ELAN, concentrating upon the application of this software to language documentation and description. Participants will gain experience in time-aligned transcription in ELAN, as well as in transforming ELAN-annotated materials into a variety of output formats using a suite of readily available open source software tools.

2   *Workshop goals:*

The primary goal of this workshop is to provide participants with an opportunity to:

1. Gain a better understanding of the fundamental concepts behind ELAN, as well as practical experience applying this knowledge to the annotation of audiovisual language materials;

2. Consider how ELAN might be of use in their own language-related projects – where ELAN might be used to produce standards-compliant, 'future-friendly' language resources, or assist in facilitating reuse and dissemination of language materials – and weigh these potential benefits against the advantages of other comparable tools and techniques.

Introduction: Introducing ELAN

1   What is ELAN?

ELAN ('EUDICO Linguistic Annotator') is a software tool commonly used in documentary linguistics to produce time-aligned annotations of digital audio and video recordings. That is, using ELAN, it is possible to match up arbitrary pieces of text (*annotations*) with sections of audio or video (*media*), producing documents (*transcripts*) which permit fluid navigation between text and media.

Actively developed since 2002 and made freely available as open-source software, ELAN is a comparatively mature software tool, having grown out of the European Distributed Corpus Project (EUDICO) and ongoing work in the analysis gesture and sign language at the Max Planck Institute for Psycholinguistics (Nijmegen, Netherlands). Besides this institutional commitment to continued development, ELAN has several notable strengths:

1. *ELAN is free, open-source, multi-platform software.* Pre-compiled versions of ELAN are available for Windows, Mac, and Linux platforms, and the entire ELAN source code is freely available for anyone to make changes to.

2. *ELAN produces non-proprietary, standards-compliant output.* ELAN transcripts are stored as Unicode-based XML, conforming to a well-documented technical standards – thus encouraging long-term reuse and archival of these resources.

3. *ELAN is flexible.* ELAN sets no hard limit on the amount of information which may be included in your transcripts. There are no restrictions on how many speakers can appear in a transcript, for instance, or on what languages can be represented. ELAN is capable of annotating a variety of common audio and video formats, works happily with multiple languages and multiple speakers in a single transcript, and can import and export its transcripts in a variety of common formats.

2   What is ELAN *not*?

For all of its strengths, it should be noted that ELAN does not attempt to serve every function to which time-aligned transcripts could be put. There are two particular limitations that are worth bearing in mind:

1. *ELAN does annotation, not editing.* ELAN concentrates upon adding textual annotations to multimedia resources, rather than making any changes to those resources themselves. **ELAN will never change your media in any way during transcription**: it can play recordings, but it can't edit out silences, false starts, or any other features you may want to remove before transcription.

    This is arguably a good thing: since ELAN has no media editing functions, there is much less of a chance of accidentally deleting part of an irreplaceable recording while doing transcription in ELAN!

2. *ELAN's annotations are plain, Unicode text.* ELAN doesn't provide a built-in way of using underlining or boldface selectively within annotations. (Again, this can be a good thing, since it forces you to think about what you mean by using typographical effects. You can certainly use other textual conventions to represent changes in pitch, volume, or voice quality, but ELAN doesn't attempt to stipulate what those conventions should be.)

In general, however, ELAN is capable of integrating with other software tools for media editing and redistribution, on the whole making these restrictions much less serious. The remainder of this document describes methods of using ELAN to help with various common tasks, from segmenting words and phrases out of longer recordings (Section 1) to distributing recordings on subtitled DVDs (Section 2), via YouTube (Section 3), or as web-based, clickable stories (Section 4).

Section One: Segmenting audio recordings with Audacity

In many language documentation projects, it is common to be faced with the task of dividing a single, longer recording into many smaller, independent pieces. This segmentation is often undertaken in order to include audio or video clips in standard Toolbox dictionaries: rather than refer to an entire recording,[1] you produce many smaller, independent audio or video files representing individuals words or phrases that need to be included with particular dictionary entries.

It goes without saying that this kind of segmentation can be a time-consuming process: reviewing an entire recording from within an audio or video editor, selecting individual clips, extracting them to separate files, and assigning these separate files consisting names can take a considerable amount of time (and is usually less than wholly engrossing). Even once the desired clips have been extracted, there are still a number of issues that remain:

1. Performing segmentation like this risks losing the association between the original context in which the word or phrase was produced and any later citation of that same material.

2. It can be more difficult to assign consistent metadata to a host of small audio or video clips than to one source recording. If, at some point, your project decides to start keeping detailed metadata about the recordings it is working with, having to label several hundred small files with the same metadata could end up being quite the inconvenience!

---

1 This is perhaps not the most flexible way of including audiovisual materials in a Toolbox dictionary. Margetts (2009) describes an alternative method, whereby ELAN transcripts of longer recordings are imported into Toolbox as Toolbox texts, and the Toolbox lexicon is then made to refer to individual annotations within those texts with a reference like

```
\rf      Buluhagalagala_01DU_0008
\sound  d:\toolbox_texts\Buluhagalagala_01DU.wav 12.066 14.835
```

This technique avoids many of the problems inherent in segmentation (e.g. losing the association between source context and dictionary citation, producing a host of easily disassociated audiovisual clips for the sole purpose of inclusion in a dictionary, etc.), and may be worth considering for Toolbox-based dictionary projects. See Margetts (2009; http://hdl.handle.net/10125/4426) for details.

3. Manual segmentation tends not to add much reusable information to your project. The process of manually extracting individual audio or video clips from a recording would not only be difficult for anyone else to repeat exactly, but also doesn't generally contribute any information about what's actually being *said* in those clips, which is usually what we're interested in.

This isn't to say that segmentation isn't necessary at times, or shouldn't ever be undertaken: there are certainly cases where segmentation may be the most practical option, even given the above issues. Rather than invest time in *manual* segmentation, however, it's also possible to develop reusable transcripts in ELAN, and to use the annotations within those transcripts to automatically segment a given recording. This section presents instructions on how to automatically extract clips from an ELAN transcript of an audio recording, using Audacity (http://audacity.sf.net) to perform the necessary segmentation.

1   Producing an Audacity label file in ELAN

Since ELAN itself doesn't ever make changes to the media files it's given, we will have to rely on another software package to extract clips from our audio recording for us. In this case, we will use Audacity, a freely-available digital audio editor. Audacity has a number of useful features, including the ability to automatically extract a series of sound clips from a longer recording on the basis of user-defined 'labels' – arbitrary bits of text associated with particular start and stop times, comparable to ELAN's annotations.

In order to have Audacity read in the annotations that we want to have extracted, we need to convert our ELAN transcript into a format that Audacity can understand. As of version 3.9.0, ELAN is able to produce the kinds of labels that Audacity expects. To get started, we need to open an existing ELAN transcript containing annotations that we would like to have extracted. This could be anything: a transcript of an unscripted conversation, or a careful recording of individual words intended for inclusion in a dictionary. Once the transcript is loaded, we will start by exporting the annotations as tab-delimited text (`File > Export As > Tab-delimited Text...`).
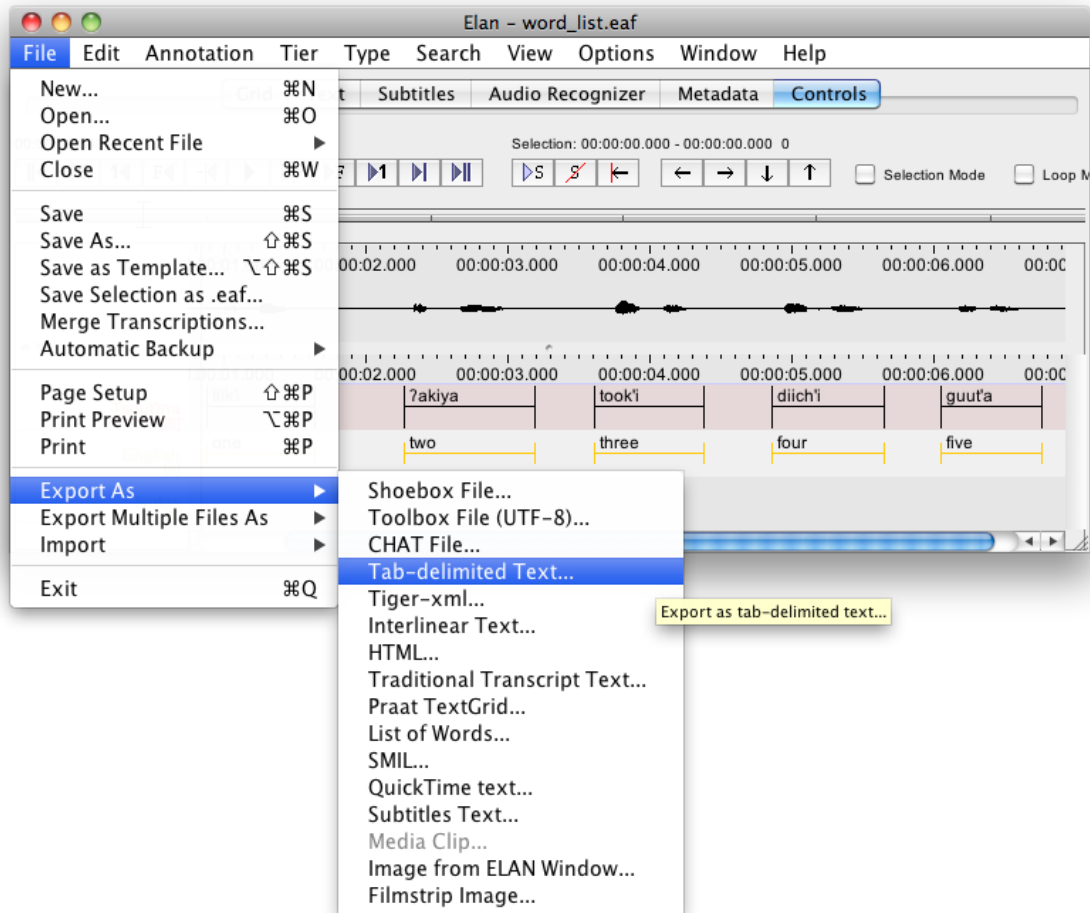
Illustration 1. Exporting tab-delimited text from within ELAN.

The tab-delimited text format that ELAN produces is almost identical to the format that Audacity uses for its labels, containing the start time, end time, and text of each annotation on some number of tiers. To make sure that the two formats are compatible, we need to:

1. Select the tiers whose annotation we would like to have extracted.

2. Select Exclude tier names from output.

3. Under Include time column for, select only Begin Time and End Time.

4. Under Include time format, select only ss.msec.
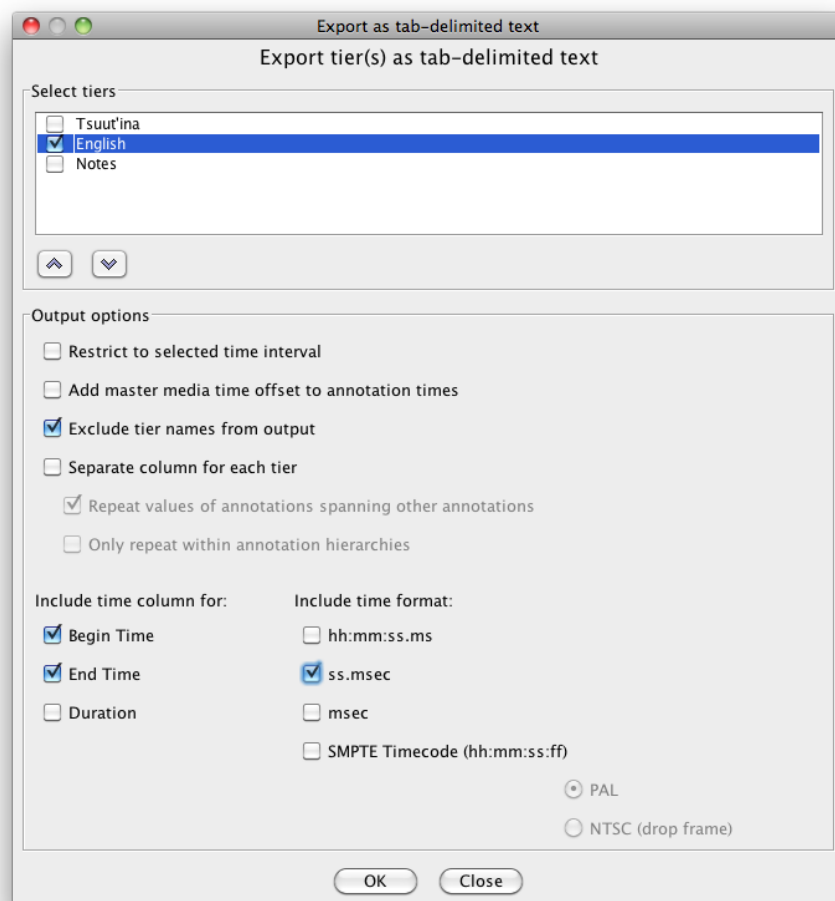
Illustration 2. Exporting Audacity labels in ELAN.

Once this is done, click on OK and save the resulting text file. This file now contains the text of all of the annotations on the tiers that you selected, presented on separate lines and preceded by the start and end time of each annotation (in milliseconds). This is all the information that Audacity needs to extract audio clips automatically; if you want to make changes to this file (e.g. doing a search-and-replace to get rid of particular characters), you can do so now, and Audacity will pick up these changes.
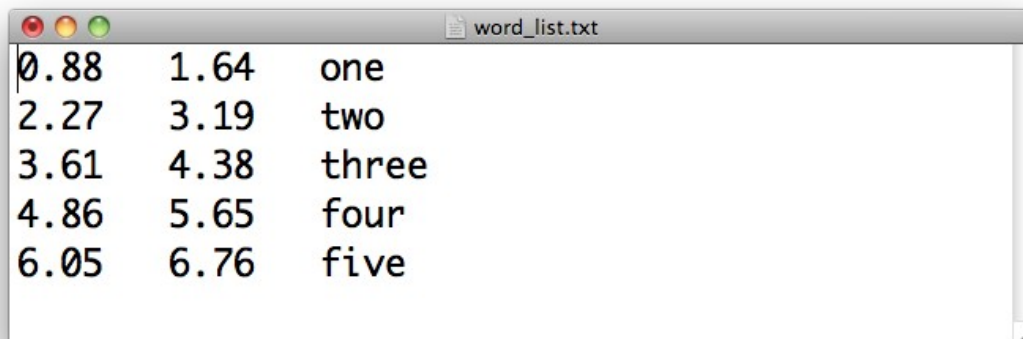
```
● ● ●                           word_list.txt
0.88    1.64    one
2.27    3.19    two
3.61    4.38    three
4.86    5.65    four
6.05    6.76    five
```

Illustration 3. The contents of an Audacity label file produced by ELAN.

2   Automatically segmenting audio in Audacity

Now that we have produced a label file containing the times and text contents of the annotations we want to have made into clips, we can use Audacity to segment the original recording for us. Open the original audio recording in Audacity (File > Open…), and, once it has loaded, import the text file containing the labels produced by ELAN (File > Import > Labels…).

If all goes well, you should see something similar to Illustration 4, which shows not only the waveform for your audio recording, but also a series of text boxes on a separate label track which represents the contents of the tiers you chose to export from ELAN. To make sure that these labels imported properly, you can click on the text of any label, then click on the green 'play' button in the top left-hand toolbar to play the audio associated with that label.
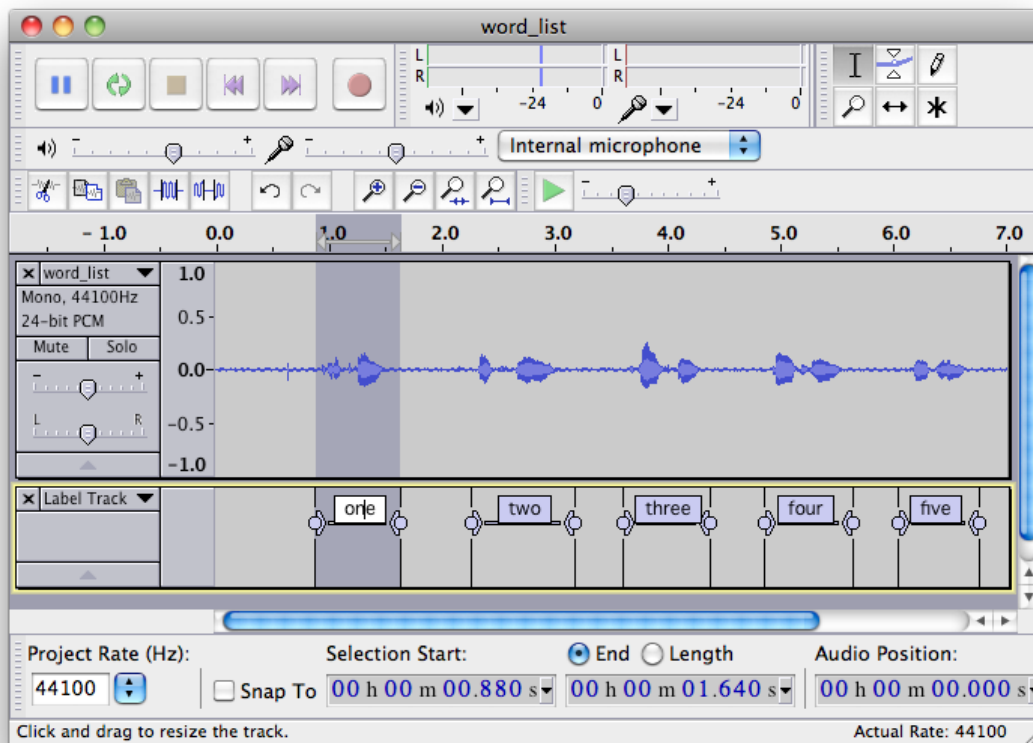
Illustration 4. An audio recording and ELAN-produced labels in Audacity.

Now that the original recording and the labels have been loaded, we are ready to have Audacity segment our audio for us. From the File menu, select Export Multiple... This will bring up a small dialogue box which allows us to specify where the audio clips should be saved (Export location), what format they should appear in (Export format), and how the resulting clips should be named.

Although many other options are available to us, we will choose to export our audio clips here as regular WAV files (WAV (Microsoft) signed 16 bit PCM) into a separate folder from the one where our original recording is stored, using the labels to define the boundaries of our audio clips (Split files based on: Labels), and using label names as the names of our clips (Name files: Using Label/Track Names), as in Illustration 5. Once these settings have been made, click on the Export button to have Audacity begin extracting audio clips.
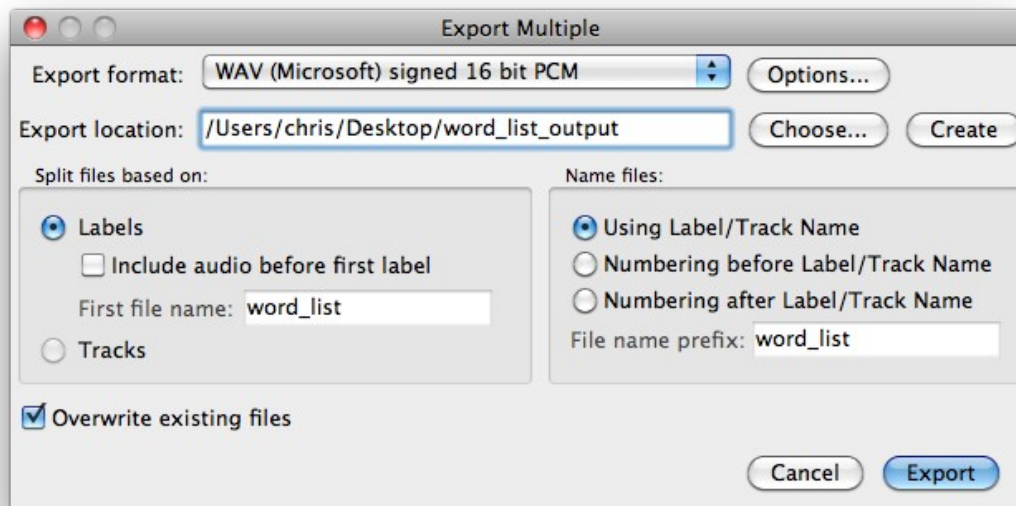
Illustration 5. Exporting multiple audio files in Audacity.

After a short delay, Audacity should inform you once it has completed extraction. The result of this export process should be a series of audio files, each named according to the contents of the label to which it corresponds.

Section Two: Producing DVDs with subtitles

Another way of distributing video materials that have been annotated in ELAN is by producing video DVDs with multilingual subtitles. Much as with the YouTube captions described in the preceding section, viewers of these videos can choose which language (if any) will be displayed as subtitles during playback. Unlike YouTube (cf. Section 3) or CuPED (cf. Section 4), however, the resulting discs can be played by anyone who has access to a DVD player, without requiring a computer, an internet connection, or access to any particular on-line video hosting service.

This section provides information about how to produce video DVDs with multiple subtitles tracks taken from an ELAN transcript. Although this section concentrates on creating DVDs with multilingual subtitles under Windows, it may be possible to adapt these procedures to DVD authoring software for Mac OS X and Linux, as well.

1   Installing DVD authoring software

In order to create a video DVD that can be played on a stand-alone DVD player, you will need to install DVD authoring software. Authoring software helps to automate the task of importing, captioning, and converting videos into the MPEG-2 format expected on standard video DVDs, and may offer facilities for creating interactive menus, as well.

For Windows users, there are several freely-available DVD authoring packages which are capable of producing subtitled discs. In the instructions that follow, we concentrate on using DVD Flick (www.dvdflick.net) for all DVD authoring, since it provides a fairly straightforward interface to basic DVD functions and excellent support for customizing how different sets of subtitles will be displayed. Another possibility would be to use DVD Styler (www.dvdstyler.de), which currently provides many more options than DVD Flick for customizing menus and graphics, although its support for Unicode subtitles seems to be somewhat less robust.
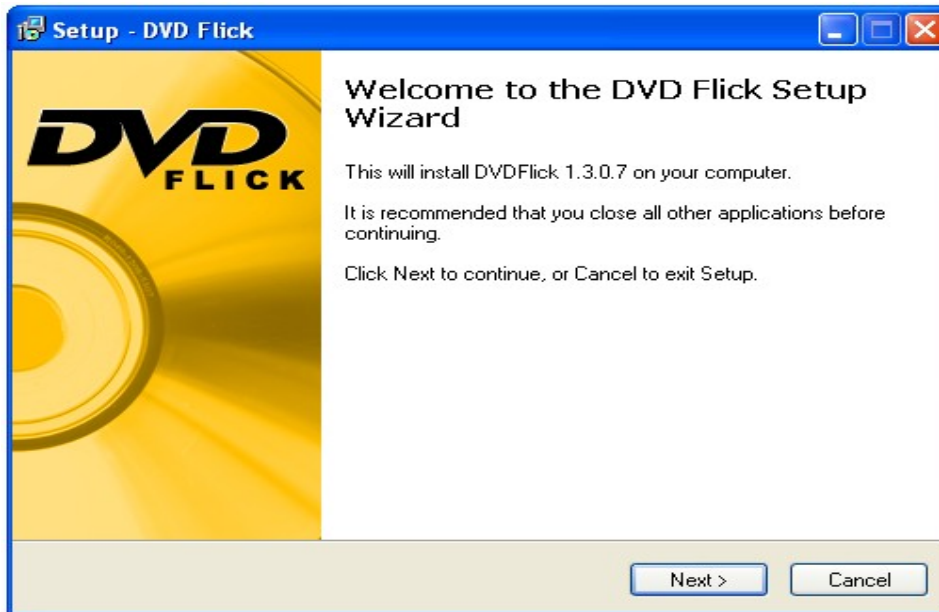
12

Illustration 6. Installing DVD Flick.

To install DVD Flick, download and run the installer from the DVD Flick website (www.dvdflick.net/download.php). Follow the prompts which the installer provides, and DVD Flick should install itself on your computer, adding an icon to your desktop and a folder in your Start Menu. Once installation is complete, you can start DVD Flick with the icons in either of these locations.

2   Adding a video to DVD Flick

To add a new video to a DVD, simply drag and drop the video file that you have annotated in your ELAN transcript onto the DVD Flick window, or click on the Add Title... button on the right-hand side of the DVD Flick window. DVD Flick should be able to load most, if not all of the same formats that ELAN is able to annotate under Windows; if all goes well, DVD Flick should display the video you have added in its main window, along with information about its file name, duration, and associated audio and subtitles tracks.
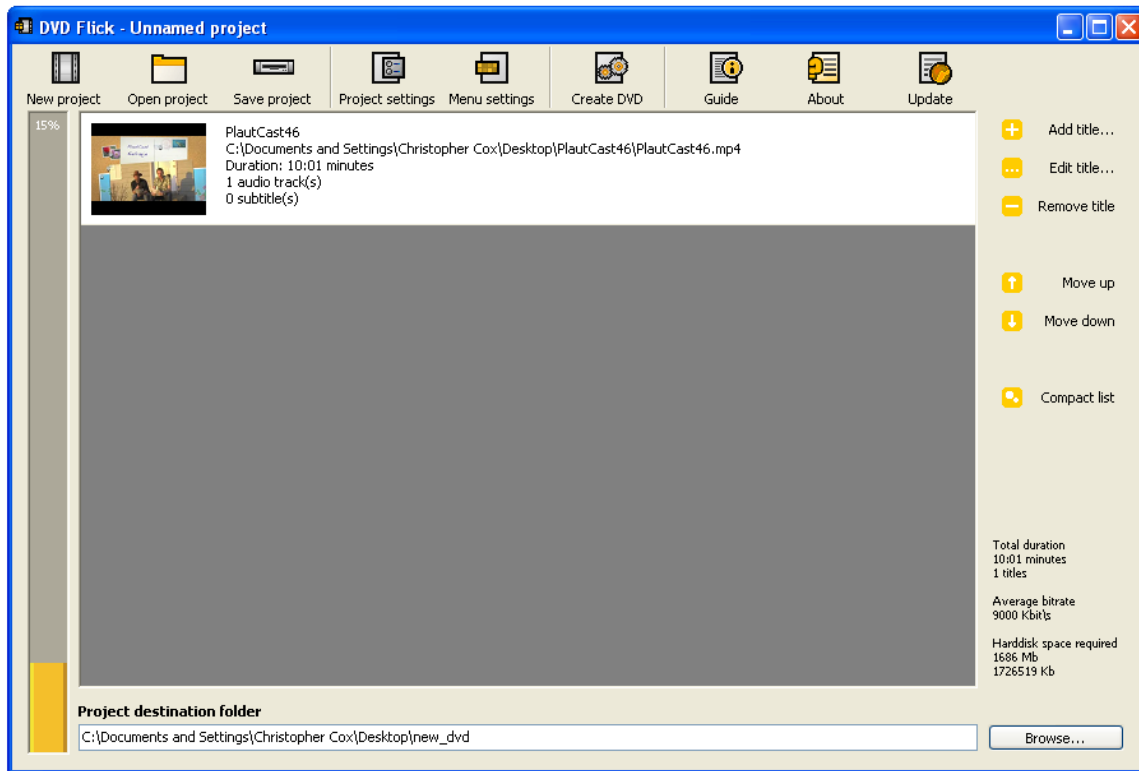
Illustration 7. Adding a new video to DVD Flick.

3 Adding and customizing subtitles in DVD Flick

To add subtitles to this new video, we will first need to export each of the tiers we wish to include as subtitles text (via File > Export As > Subtitles Text...), just as we did to produce captions for YouTube. If, for example, we wanted to have subtitles available both in the language of the recording *and* in the language of any translations we have provided in ELAN, we could export two separate subtitles text files, one with just the tiers in the original language (e.g. PlautCast46-Original.srt), and another with just the tiers of the translation (e.g. PlautCast46-Translation.srt).

DVD Flick is able to associate multiple subtitles tracks with a single video, and permits us to customize how and where these subtitles are displayed. To add subtitles to a video, select the video in the DVD Flick window, then click on the Edit Title... button on the right-hand side. In the Properties window which appears, select the Subtitle Tracks option on the left-hand side, then click on the Add button to locate the

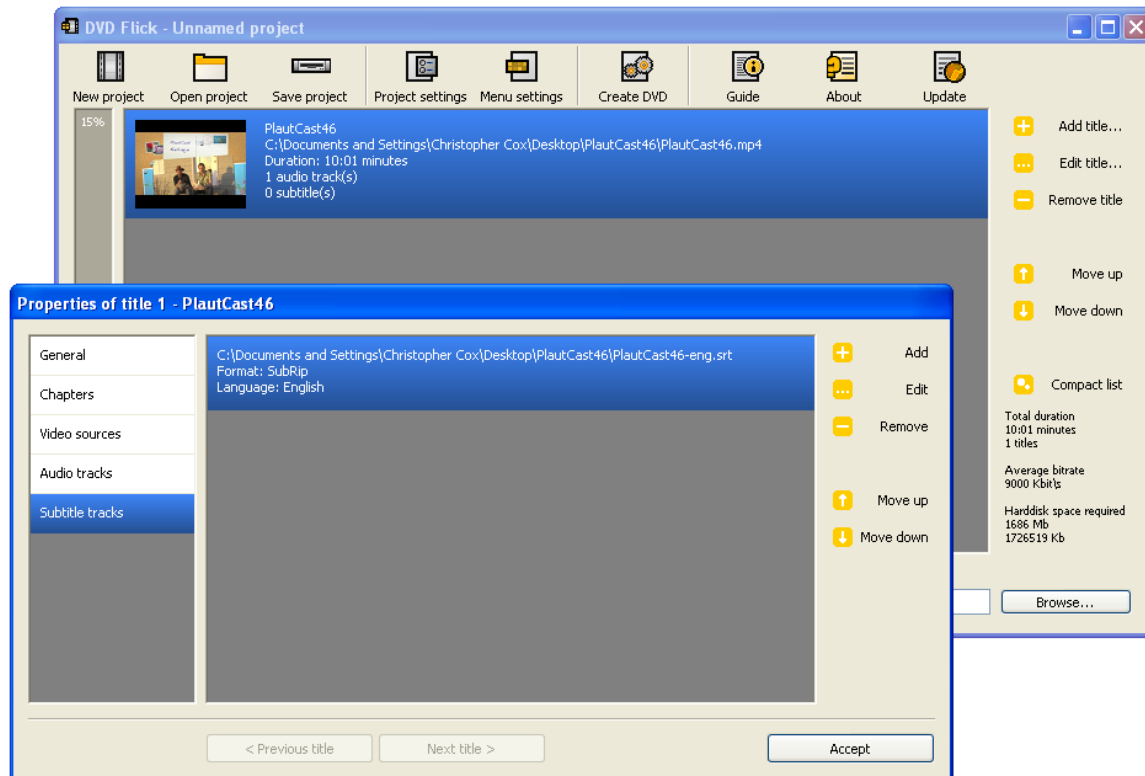subtitles text file that should be included.

Illustration 8. Adding subtitles to a video in DVD Flick.

We can repeat this procedure to add multiple sets of subtitles to the video: each of the sets of subtitles that we add here will be accessible to the viewer on the final DVD as separate subtitle 'tracks'. Again, if we had a set of subtitles in the original language (e.g. PlautCast46-Original.srt) and a set of subtitles that represent translations (e.g. PlautCast46-Translation.srt), we could add them both at this point, and they would both be available to viewers later on.

Not only does DVD Flick permit us to have multiple subtitles tracks, but it also allows us to change how and where each of these sets of subtitles will be displayed. It's possible to customize the font, colour, and positioning of any text that DVD Flick will show as subtitles. To customize a subtitle track, click on that track's name as listed in the Subtitle tracks portion of the Properties window for your video, then click on the Edit button on the right-hand side.
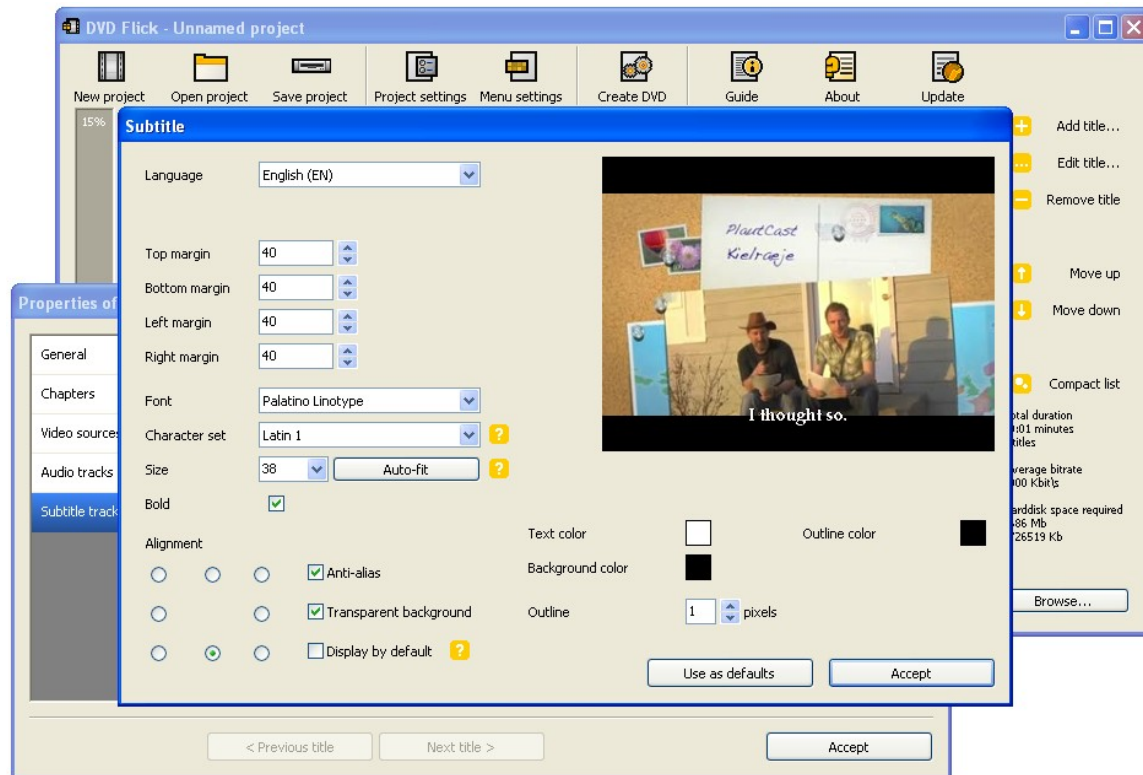
Illustration 9. Customizing subtitle display options in DVD Flick.

From within the Subtitle window which should appear, it is possible to tweak various text formatting and positioning options. Although the default values are usually more than acceptable, this screen can be quite useful when subtitle text needs to be displayed somewhere else on the screen to avoid occluding important parts of the original video. Once you have made the changes you want to this screen, click on the Accept button to return to the Properties window for your video.

4   Authoring and burning a DVD in DVD Flick

Once you have added all the videos and subtitle tracks that you want, you can choose where you would like to save your project by clicking on the Save Project button in the main DVD Flick window. (You can also customize the menus which viewers will see when they play the DVD by clicking on the Menu Settings button at this point, but that is outside of the scope of this description. By default, DVD Flick will produce a DVD that launches directly into playing the video when started, with no menu screen.)

16

We are now ready to compile (or 'author') the final DVD from our video sources and subtitle tracks into a format that can be burned to disc. This authoring process can take some time, especially when DVD Flick needs to convert much of the video into the MPEG-2 format used on video DVDs.

If you only intend to make a single copy of this video, you can insert a blank DVD±R disc at this point and burn away. If you intend to make multiple copies of your DVD, however, it may be better to save the results of the authoring process as an ISO disc image (i.e., an .iso file), which preserves the contents of an entire DVD in a standardized file format. That way, if you decide to make more copies of this DVD, you don't have to wait for all of the video to convert again: the entire DVD will be ready to go, and can be burned to disc with any DVD burning software that can read an .iso file, even on a computer or operating system where DVD Flick is not available.
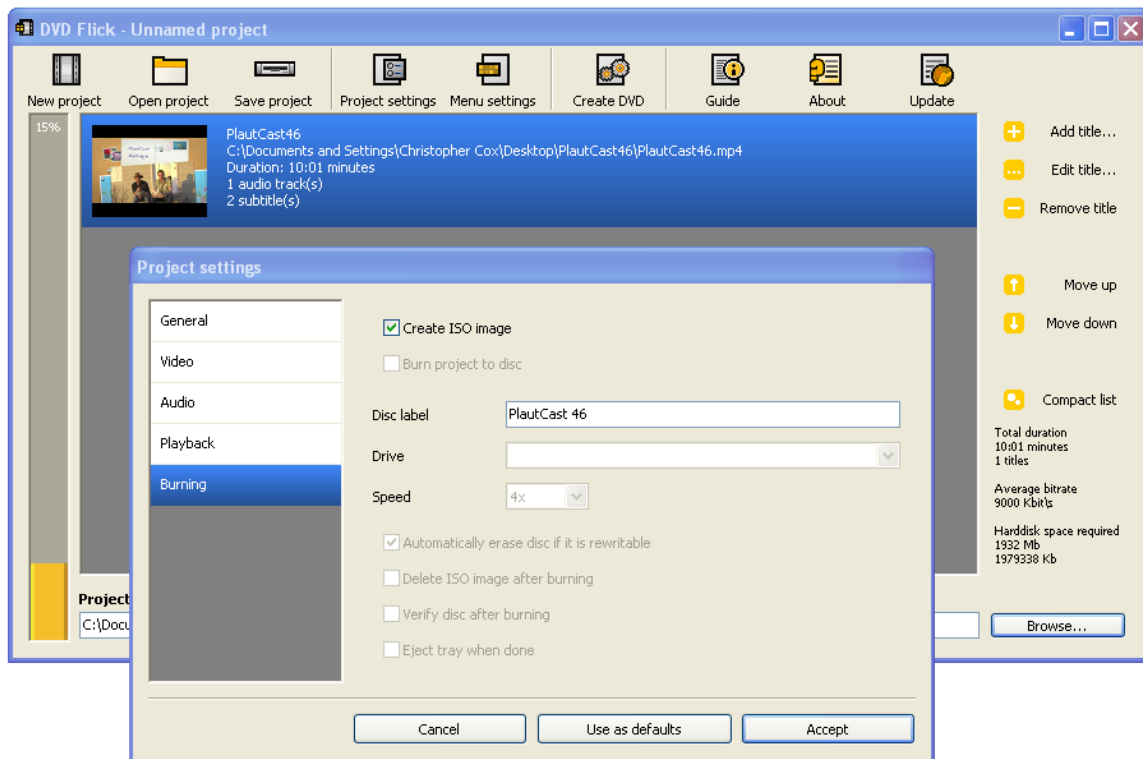

Illustration 10. Changing project settings to create an ISO image in DVD Flick.

To create an ISO image for your DVD Flick project, click on the Project Settings button in the main DVD Flick window, then select the Burning option from the left-

hand list. By selecting the Create ISO image checkbox and clicking on the Accept button, DVD Flick will produce a standard .iso file for this DVD, rather than burning it to disc immediately.[2]
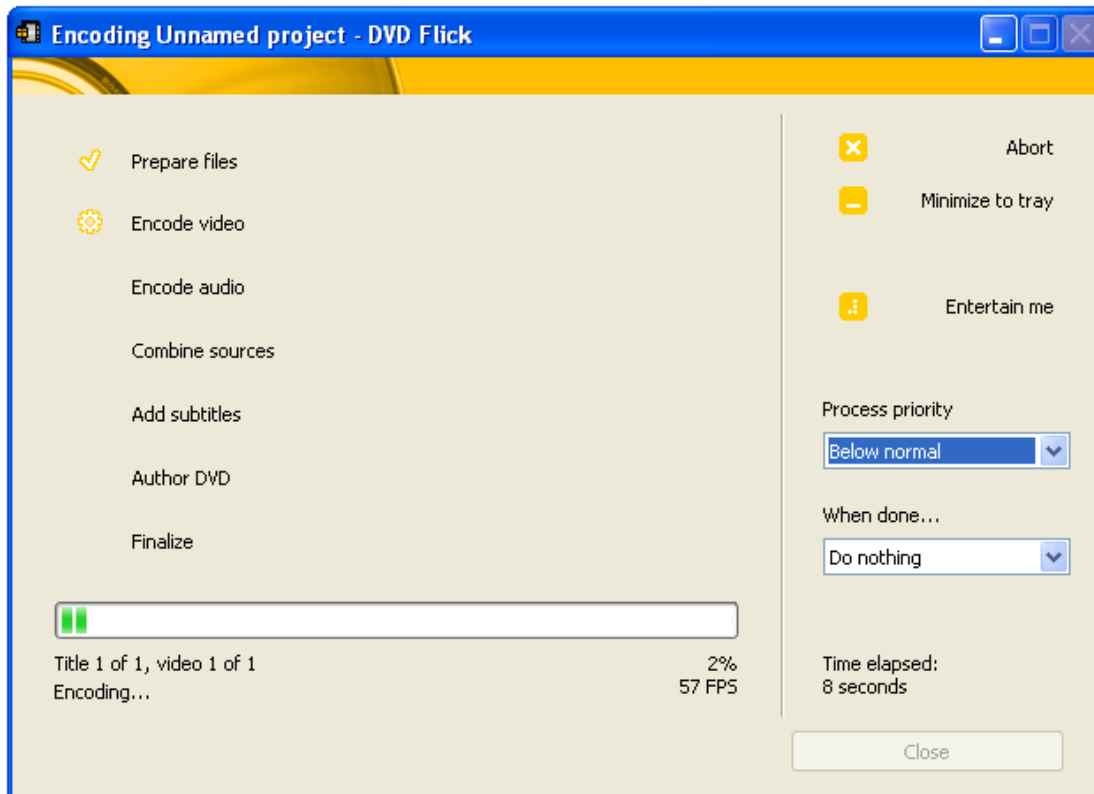

Illustration 11. Encoding a DVD in DVD Flick.

Once you are ready to create your DVD, click on the Create DVD button in the main DVD Flick window. This will begin the DVD authoring process, with DVD Flick automatically converting your video(s) and subtitles into formats needed and either burning or saving the final compilation to an .iso file for you. (Again, this may take a while to finish up; if you want to go get a coffee, this would probably be a good time to do so. Alternatively, DVD Flick also includes a built-in Tetris game for those desperately in need of a diversion!)

---

2   Within the Project settings window, it is also possible to change the video system (i.e. NTSC for most of North and South America, Burma, Korea, Japan, and several other countries; and PAL for most of Europe, Africa, Australia and Oceania, southern Asia, and Brazil) with which this DVD will be compatible.

5   Reviewing the final DVD

If you decided to burn a copy of the DVD to disc, then reviewing the final copy is as simple as inserting the disc in a stand-alone DVD player or in computer with a DVD drive and appropriate playback software. To view the subtitle tracks added earlier, select them from within the playback program (on a computer), or using the 'Subtitles' button on the DVD player remote (on a stand-alone DVD player).

Since DVD Flick does not currently provide a preview option which includes the added subtitles, you may want to consider creating an ISO image of your DVD prior to burning, then loading (or 'mounting') this file as the virtual contents of a DVD drive. Under Windows, Daemon Tools (www.filehippo.com/download_daemon_tools/) is able to load an ISO image as if it were an actual DVD, allowing you to test out a DVD image before burning it to disc. Likewise, Mac users should be able to double-click on the ISO image itself, and it should load just as a normal video DVD would. (Some video players, like VLC (www.videolan.org/vlc), are also able to play videos contained in a ISO image directly, without the need for services like Daemon Tools.)

6   Alternatives

Although the preceding instructions have concentrated on using DVD Flick under Windows to produce subtitled DVDs, there are certainly other DVD authoring packages, both Free and commercial, available for other platforms that are capable of accomplishing much the same task. As mentioned in the introduction, DVD Styler (http://www.dvdstyler.org) can be used under Windows and Linux to produce subtitled DVDs, although issues have been noted in recent versions with its support for Unicode characters outside of the extended Latin alphabet.[3] If your language's orthography only uses letters found in Western European languages, you may still be able to use DVD Styler to process your subtitles – with the following caveats:

1.   When exporting your subtitles from ELAN (File > Export As > Subtitles Text...),

---

3   More specifically, it does not appear that DVD Styler 1.8.1 is able to use .srt subtitles exported from ELAN in anything other than ISO-8859-1 encoding; informal tests suggest that neither UTF-8 or UTF-16 subtitles are fully supported in this version of DVD Styler.

be sure to select the encoding ISO-8859-1, rather than the default UTF-8, for your .srt subtitles.

2. Once ELAN has produced your subtitles, they can be loaded into DVD Styler by dragging them onto the thumbnail image of the corresponding video track in the bottom toolbar in DVD Styler. Alternatively, you can right-click on the video track in DVD Styler, select Properties... from the menu, and proceed to add each subtitle track using the plus and minus buttons in the Properties window which appears, as in Illustration 12 below.
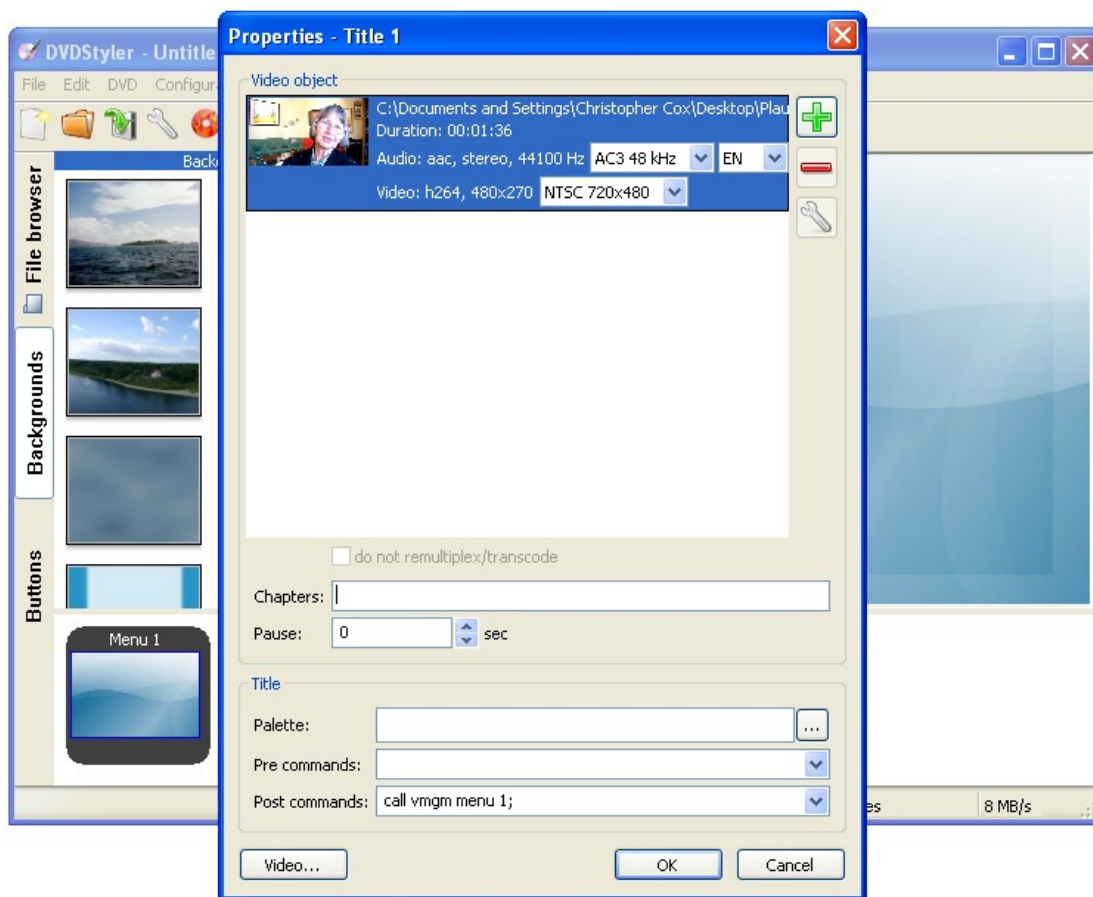

Illustration 12. Adding subtitles to a video in DVD Styler.

For Mac users, commercial software, such as Apple's own DVD Studio Pro (which is bundled with the Final Cut Studio; http://www.apple.com/finalcutstudio/dvdstudiopro/)

is also capable of producing video DVDs with multiple subtitles tracks, and boasts excellent Unicode character support. (Unfortunately, neither iDVD or iMovie are able to add multiple, selectable subtitle tracks to a video being put on DVD)

For users of DVD Studio Pro, a few changes are necessary to the procedure outlined above. Since DVD Studio Pro is unable to load .srt subtitles, we have to have ELAN produce subtitles in another format, .stl (Spruce Subtitles), that this DVD authoring software understands. As of ELAN 3.9.0, it is possible to produce .stl files directly from ELAN through the following procedure:

1. With the transcript loaded for which you want to produce subtitles, select `File > Export As > Subtitles Text...` to bring up the subtitles text export window.

2. In the `Export as Subtitles Text` window, you can select those tiers you wish to export as subtitles, just as before. At this point, however, it is also important to select what kind of *timecode* your subtitles should use: unlike .srt subtitles, which rely on minutes, seconds, and milliseconds for their timings, the .stl subtitles that DVD Studio Pro uses refer to actual *frames* within the video, instead. Since different countries use different video systems, each of which may have a different number of frames being recorded each second, we need to tell ELAN what system our video is intended for so that our subtitles stay perfectly in sync.

    ELAN lists three different video systems as a kind of short-hand for the number of frames per second in your video: `PAL`, `NTSC (DF)`, and `NTSC (NDF)`. If your DVD is being made for use in the Americas (outside of Brasil, Paraguay, Uruguay, Argentina, and French Guiana), Japan, South Korea, Taiwan, the Phillippines, or Burma, your video likely follows NTSC timing, and you can choose either `NTSC (NDF)` (which, in informal tests, has had the best rate of success with staying in sync with NTSC video in DVD Studio Pro) or `NTSC (DF)`. If your DVD is intended for use anywhere else on earth, you're likely using PAL or SECAM-encoded sources, and you can safely choose `PAL` in this window. Illustration 13 shows this export window, with a single tier selected for export and `NTSC (NDF)` as the timecode.
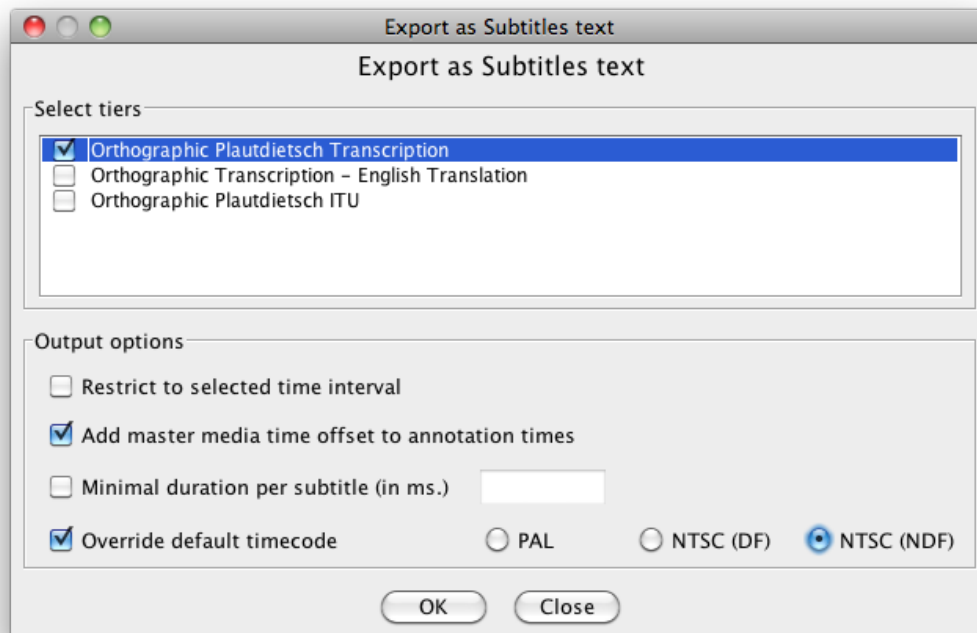
Illustration 13. Selecting tiers to export as .stl subtitles in ELAN.

3. Once you have selected your tiers and a timecode for your subtitles, click OK. You will then be prompted to select a file name and encoding for the subtitles you've chosen to export. You can save the file wherever you want; to make sure that ELAN produces .stl subtitles that DVD Studio Pro can read, be sure to save your file with the extension .stl, rather than the default .srt (e.g. MySubtitles.**stl**); and choose UTF-16 encoding in the drop-down selection at the bottom of the window, rather than the default UTF-8. Illustration 14 shows one such set of subtitles being saved under the name Aupelkoose.stl, and with the encoding UTF-16 selected.
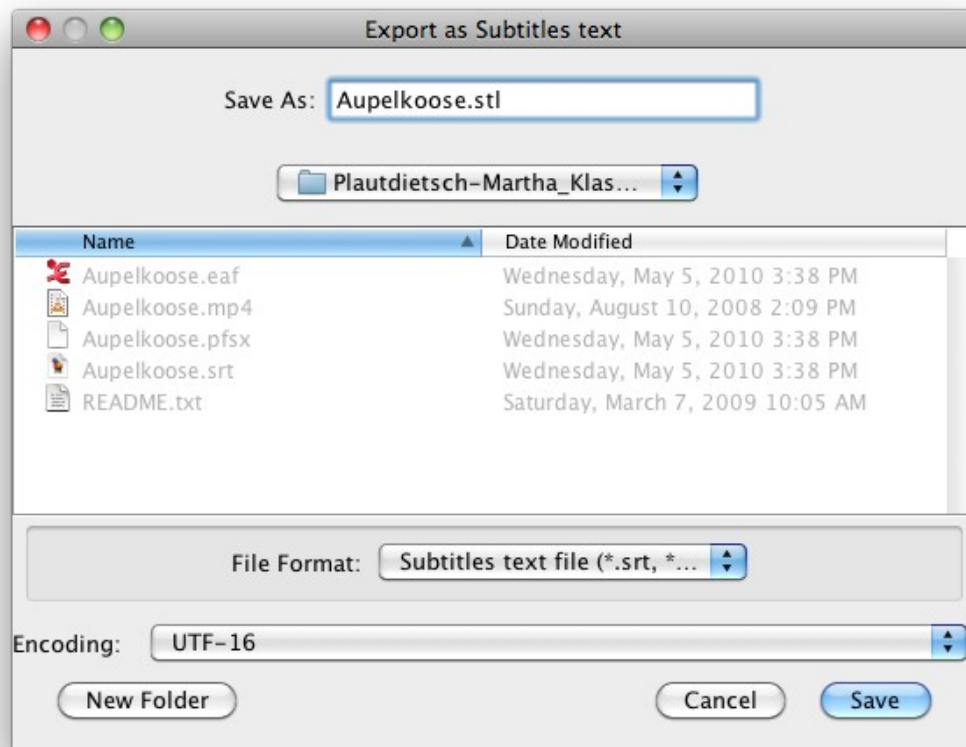
Illustration 14. Exporting subtitles in UTF-16 .stl format for DVD Studio Pro.

The resulting .stl subtitles can then be imported into DVD Studio Pro by right-clicking on one of the subtitles tracks (S1 to S32) and choosing Import Subtitle File... from the context menu which appears. You can add up to 32 separate subtitles in this way. By default, each of the subtitles tracks added can be accessed on the final DVD by pressing the 'Subtitles' button on the DVD player's remote (for a stand-alone DVD player) or by selecting a particular subtitle track (for DVD playback software on a computer).

Section Three: Producing YouTube videos with subtitles

One of the advantages of investing effort in time-aligned transcription using software like ELAN is the variety of output formats which can be produced automatically from the resulting transcripts. Among other export options, ELAN is capable of producing so-called subtitles text files, which present the contents of a selected set of ELAN tiers as captions. These files can be used to provide user-selectable subtitles for both video DVDs and popular on-line video-sharing sites such as YouTube.

This section describes how to produce captioned YouTube videos, using an ELAN transcript of a short video clip as our starting point.

1   Creating a YouTube account

In order to add your own videos to YouTube, you first need to set up a YouTube account. If you don't already have a YouTube account, you can create one by following the 'Sign Up' link, found in the set of links located in the top right-hand corner of the YouTube homepage (www.youtube.com).
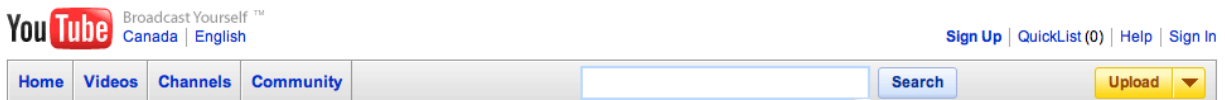

Illustration 15. The header of the YouTube homepage.

YouTube will then guide you through the process of choosing a user name and password, as well as configuring your preferences for use of the site. You will only need to sign up for an account once: when you have your new YouTube account set up, you will be able to log into it by following the 'Sign In' link, which is also found in the top right-hand corner of the YouTube homepage.

2   Adding a video to YouTube

Once you have created a YouTube account, you can begin adding your own videos to the site. To add a new video, select the gold 'Upload' button, again found in the upper

24

right-hand corner of the YouTube homepage. This will present you with two options: one to add a video from a file on your computer, and another to create a video using a webcam. Since we will be working with a video clip that we have already saved using MPEG Streamclip and transcribed using ELAN, we will choose the first option, 'Upload Video File'.
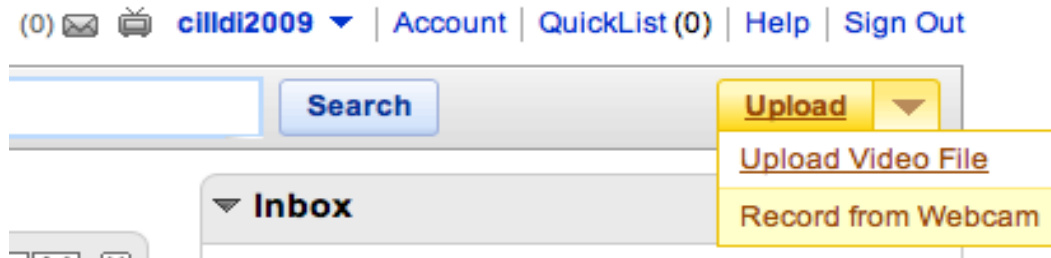

Illustration 16. Uploading a video file to YouTube.

Selecting this option will take you to a screen where you can choose the video file that you want to upload to YouTube. YouTube places several restrictions upon the videos that it accepts, perhaps the most important of which is that clips are limited to a maximum length of ten minutes. If your video clip is longer than ten minutes, then you may want to consider software such as CuPED for helping you share your transcribed language resources over the web. If your video clip is less than ten minutes long, you can add it to YouTube on this page by selecting the 'Upload Video' button.
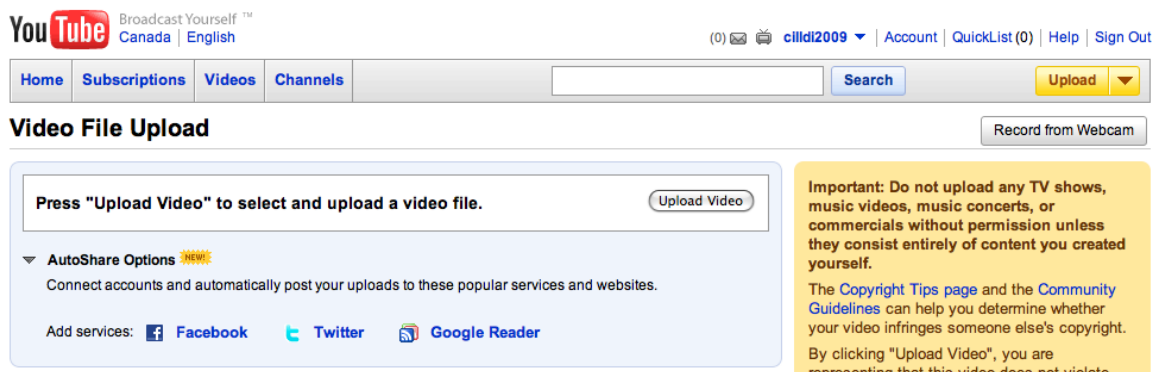

Illustration 17. Selecting a video file to add to YouTube.

The video you choose at this point will be sent to YouTube. This stage can be the

most time-consuming, especially if you are working on a slower internet connection: it will take a while to send your video to YouTube. To make the process take less time, you should consider sending a compressed version of your video clip (e.g. in `.mp4` format, as can be produced by MPEG Streamclip) to YouTube, rather than the full, original Digital Video (`.DV`) file which you captured from your video camera. If you edited your video in MPEG Streamclip, you likely already have your video clip in MPEG-4 format, which should be many times smaller than your original video and take much less time to send to YouTube.

Once you have selected your video file, YouTube will begin to receive it, and you will be presented with a number of fields in which to describe your video and customize how it will be seen by others on the internet. These include the title of your video (e.g. '*Cree Conversations – CILLDI 2009*'), a brief description (e.g. '*Conversations in Plains Cree* (*y-dialect*) *between participants in the Elders and Crafters Forum at the Canadian Indigenous Languages and Literacy Development Institute* (*CILLDI*) *2009, University of Alberta.*'), and, optionally, a set of tags, which are keywords which you think might help people find your video (e.g. '*Cree*', '*Nehiyawewin*' '*CILLDI*'). If you intend to make your video public, for anyone on the internet to see, then this information is very important, since it will be what most people use to find your video!
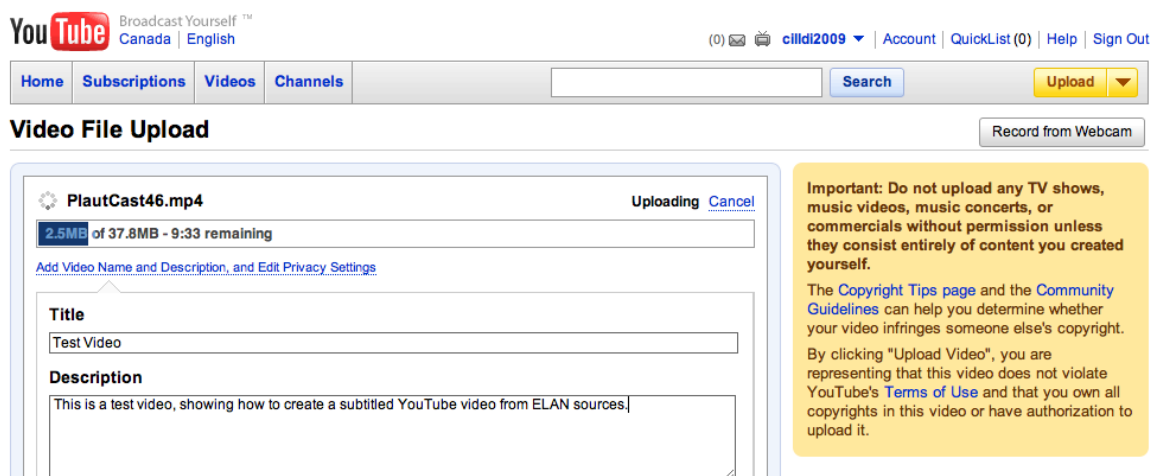

Illustration 18. Uploading a new video to YouTube.

Besides allowing you to describe your video clip, this page also allows you some basic control over who is allowed access to it. In the 'Privacy' section towards the bottom of the page, YouTube allows you to specify whether your video should be public, and available for anyone in the world to watch; or private, and visible only to you and a short list of people you choose. Depending on the conditions which exist around the video clip you transcribed, you may want to let everyone see your video, or you may need to restrict access to only a short list of people; the choice is up to you.

Once you have described your video and decided who should be allowed to access it, you can select the 'Save Changes' button towards the bottom of the screen. YouTube should then tell you that your video settings have been successfully saved. Your video clip is now on YouTube – it may take a couple minutes to show up in the YouTube search listings (if you decided to make your video visible to the general public), but it's now on the site.

3   Adding subtitles

What we want to do now is add subtitles to this video. The subtitles text files that YouTube uses for its subtitles are the same ones that ELAN produces when we wanted to add captions to our DVDs. You can select which tiers you want to include in your subtitles from within ELAN (`File > Export As > Subtitles Text...`); you may want one version which just shows translations in English, or another which only shows your video in the original language. You can export as many different sets of subtitles as you like; we can add each set of subtitles to our new YouTube video, and then allow users to choose between them.

Once you have created subtitles text in ELAN for your video clip, we can add the subtitles to the new YouTube video. In YouTube, you can view all of the video clips which you have added to the site by selecting the 'My Videos' link, found under your user name in the set of links in the top right-hand corner of the screen.
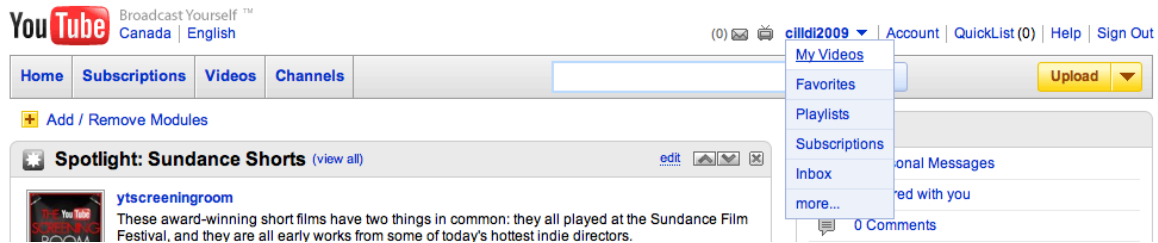
Illustration 19. Accessing your videos in YouTube.

Following this link will bring you to a page which lists all of the videos which you have added to YouTube. What we want to do here is edit the new video and add the subtitles to it. To make changes to a video, click on the 'Edit' button, found just beneath the little preview picture of the video itself.
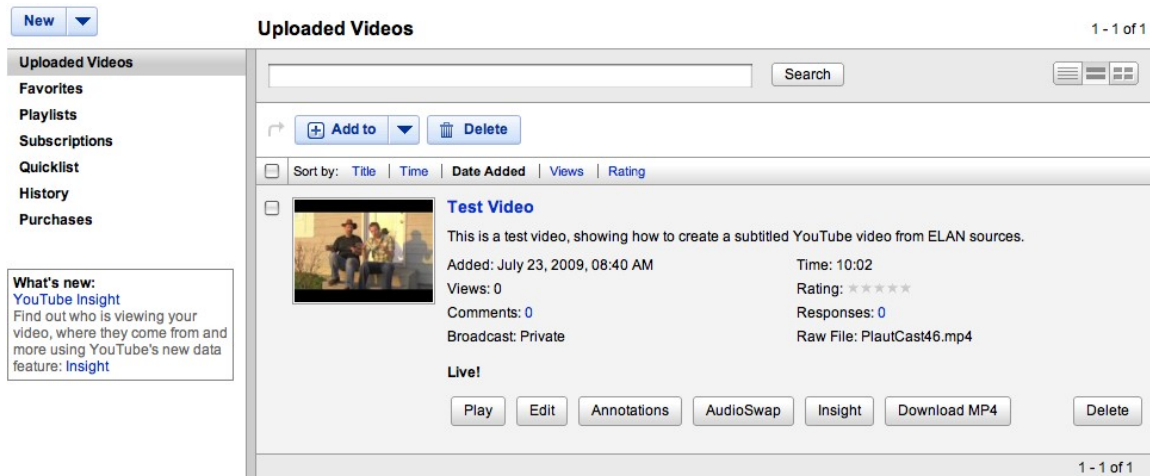

Illustration 20. Reviewing all of the videos which you have sent to YouTube.

Editing your video presents a number of options for customizing how your video is presented and distributed to YouTube viewers: you can select a location for your video on a world map, or choose the thumbnail which will be used as a preview for your video in YouTube search listings, or control how your video can be used in other websites. For our purposes, though, we're most interested in the 'Captions and Subtitles' section, where we can add our subtitles to the video.
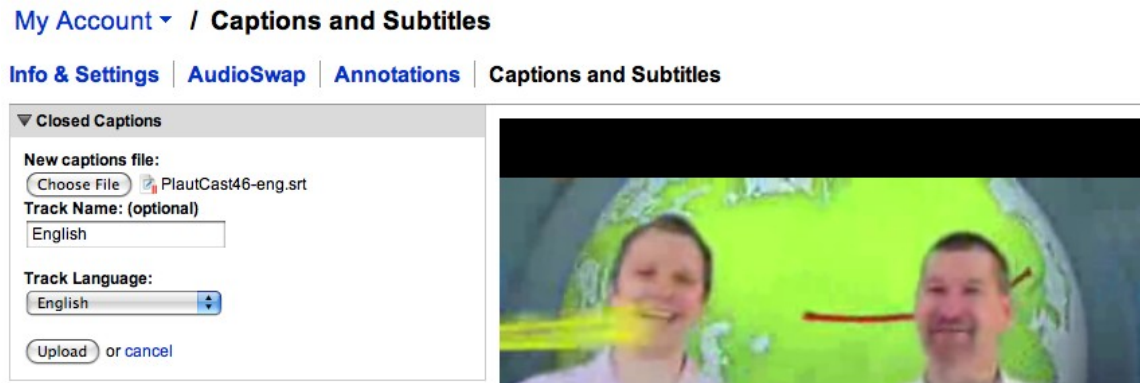
Illustration 21. Adding subtitles to a YouTube video.

Within the 'Captions and Subtitles' section, we can now add subtitles to our new video. Select the 'Choose File' button, and find one of the subtitles text (.srt) files that you made in ELAN. YouTube will then give you the option to give a descriptive name to this set of subtitles (e.g. '*English translation*', '*Cree syllabics*', etc.), and let you choose what language these subtitles are in. (Don't worry if your language isn't listed – YouTube currently presents a very limited list of languages, but we can get around that for now by putting the name of the language of the subtitles in the description, instead.)

Once you have these fields entered, you can select the 'Upload' button, and your subtitles will be sent to YouTube and added to your video. If you want to add more than one set of subtitles – one for what is being spoken in your language, for instance, and one for a translation – you can repeat the steps above, selecting another subtitles text file on your computer and adding it to your new YouTube video, as well.

4   Viewing your YouTube video with subtitles

When you are done adding subtitles to your YouTube video, you can review your work by watching your video in the normal YouTube video display. To access your new video, select the 'My Videos' link within the menu under your user name, found in the set of links in the top right-hand corner of each YouTube page, and then select your new video.

Illustration . Turning on captions for your new YouTube video.

By default, no captions are turned on for your video. To enable captions, click on the arrow in the bottom right-hand corner of your video, and select the 'CC' option from the menu which appears there. This will bring up the list of available subtitle tracks which you have added to your video. You can turn individual sets of subtitles on and off from this menu. (Sometimes it takes a few minutes for newly-added subtitles to appear in this menu. If your new subtitles aren't showing up right away, just wait a couple of minutes, and then reload your video, and the subtitles should appear) You should now see the subtitles appearing underneath your video during playback.

5   Limitations

Although YouTube can be a handy tool for distributing language materials, there are a number of limitations to the site that should be kept in mind:

1. YouTube places a limit on the length and file size of the materials being uploaded. At the time of writing, most YouTube videos are not permitted to exceed ten minutes in length or 2GB in size. While many language materials in compressed video formats (e.g. MPEG-4) may not have to worry about the file size limit, the length limit can prove more troublesome: if your recording is longer than ten minutes, YouTube will

most likely reject it out of hand.

There are certainly ways of working around this limit. Perhaps the simplest method is to segment your language material into multiple 'chunks' of ten minutes or less using video editing software (e.g. MPEG Streamclip), and then producing multiple sets of subtitles from inside ELAN (i.e. by highlighting the sections of video in your original transcript corresponding to each ten-minute chunk, and, for each one, selecting Restrict to selected time interval and unselecting Add master media time offset to annotation times in the Export as Subtitles text dialogue. This way, a separate .srt file is created for each chunk, and can be uploaded to YouTube as described above). This is a rather time-consuming procedure; if you have many longer video files to process, it may be simpler to use a tool like CuPED, instead (cf. Section 4).

2. YouTube concentrates on distributing video materials, rather than audio materials. If you want to share language materials via YouTube that only include audio, you will need to convert your audio recording into a YouTube-compatible video format first. Thankfully, there are a number of tools that can help you to accomplish this task:

   ◦ *Under Windows and Mac OS X:* It's possible to load still images and audio tracks into Windows Movie Maker (under Windows) or iMovie (under Mac OS X), and to save the resulting image-audio combination as a .wmv or .mp4 video that can then be uploaded to YouTube. Although an introduction to either piece of software is outside of the scope of this discussion, instructions are available online (e.g. http://www.youtube.com/watch?v=wFQUWa6cN14).

   ◦ *Under Linux / Unix:* The common multimedia processing tool ffmpeg (http://ffmpeg.org) can be used to turn an audio file and a still image into a video that can be uploaded to YouTube. Assuming your audio file is audio_file.wav, your still image is picture.jpg, and you want the final video to be called output.flv, the corresponding command to invoke ffmpeg would be as follows (all on one line):

      ```
      ffmpeg -f image2 -loop_input -i picture.jpg -aspect 4:3 -i audio_file.wav
        -ar 44100 -ac 2 -ab 192k -shortest output.flv
      ```

Section Four: Producing 'clickable stories' with CuPED

Another way of making materials that have been transcribed in ELAN accessible to a broader audience is by producing so-called 'clickable' stories using a tool called CuPED. Whereas YouTube concentrates on delivering video on-line, and imposes the restriction that all video clips must be ten minutes in length or less, CuPED allows for both audio *and* video materials to be made available in a web browser, whether on a local computer (e.g. on CD, on a laptop, on a library workstation, etc.) or over the internet, and places no limit on the length of the materials being presented. Although the resulting presentation requires a web browser to view, it provides a more interactive interface than either YouTube or a standard video DVD, and may thus be more interesting for some audiences.

This section presents instructions on how to apply CuPED to an ELAN transcript and its associated recording(s) to produce a clickable story which is ready for use on the web or on a local computer. For those who would like further information on how to use CuPED, a user manual is also available from the CuPED homepage (see below).

1   Installing CuPED

CuPED is a free, open-source software tool which provides a user-friendly interface to the media and transcript conversion facilities commonly needed to prepare ELAN-annotated audio or video for distribution over the web. Although CuPED is still in early development, it is capable at the moment of turning most simple, single-speaker transcripts into 'clickable stories', with plans to handle multiple speakers and other output formats in the near future.
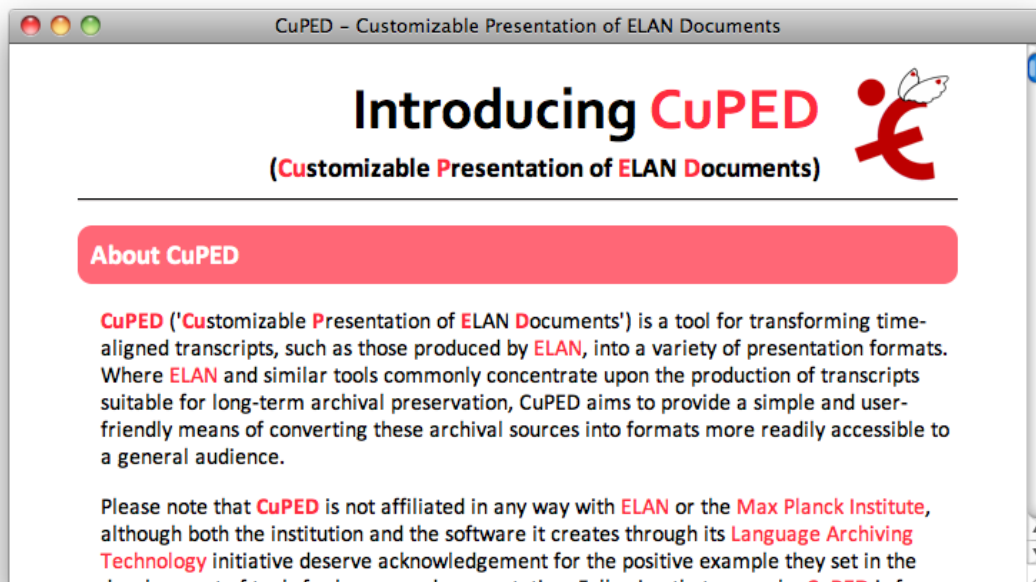
Illustration 22. CuPED homepage. (sweet.artsrn.ualberta.ca/cdcox/cuped)

While the CuPED source code is available for use on any operating system where Python is available, precompiled installers have been prepared for Windows and Mac OS X (Intel/PPC) users. Assuming that you are using either one of these operating systems, you may visit the CuPED homepage (sweet.artsrn.ualberta.ca/cdcox/cuped/) and download and run the appropriate installer, which will guide you through the installation process (or, on a Mac, let you drag CuPED to your Applications folder).

2   Applying CuPED to ELAN transcripts

Once you have downloaded and installed CuPED, you can begin applying CuPED to your ELAN transcripts. The CuPED user interface has been made to resemble a software 'wizard', a set of dialogue windows which gently guide users step by step through the process of converting their transcripts.

Double-click on the CuPED icon to start the application. If all goes well, you will be presented with the CuPED 'welcome' screen. As elsewhere in the CuPED interface, you can advance to the next dialogue by clicking on the Continue button at the bottom

of the CuPED window. By advancing past the welcome screen, CuPED will present you with a list of output templates that you can apply to your transcript.
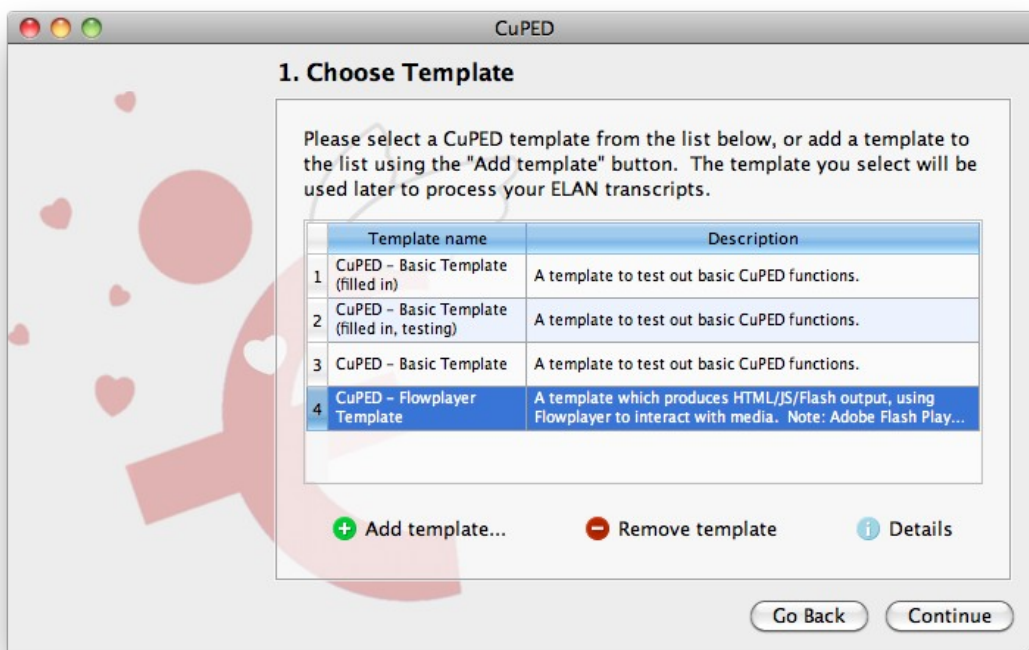

Illustration 23. Selecting a template in CuPED.

You can think of templates as providing different ways of presenting your transcript: each template sets out one way of transforming the transcripts it is given into a particular output format. Even though CuPED currently only produces clickable stories at the moment, it's possible to imagine additional templates that would transform an ELAN transcript and its associated media into, say, a printable Word document, or a LaTeX manuscript, or any number of other output formats. To access the clickable story functions available in the current beta version of CuPED, select the Flowplayer template, then click on Continue.
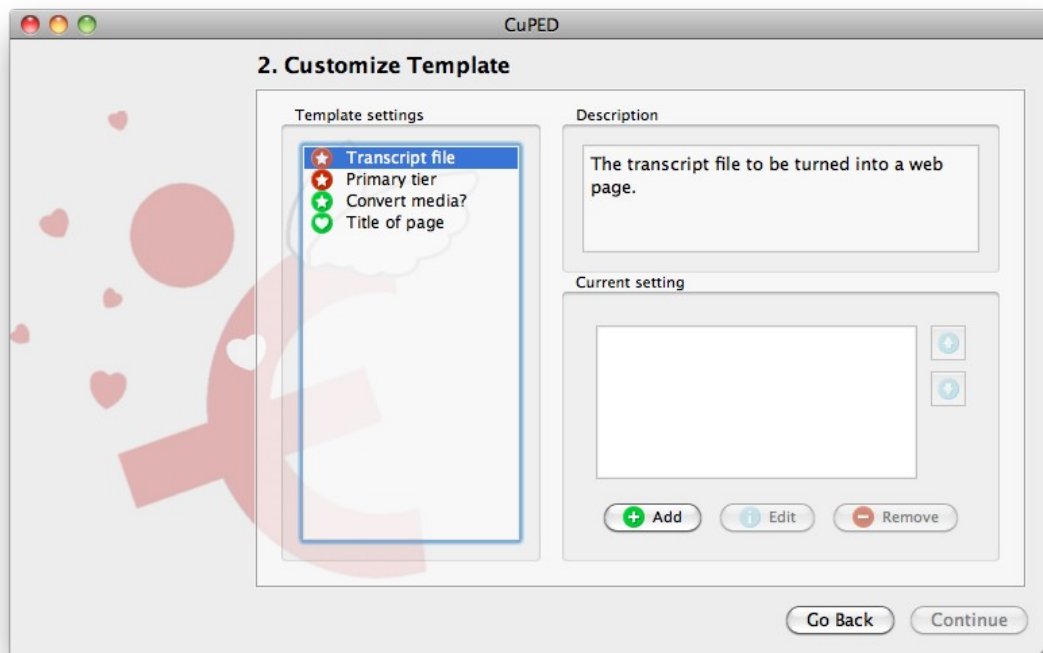
Illustration 24: Customizing the Flowplayer template in CuPED.

CuPED allows users to customize certain settings for the templates they select. In the case of the Flowplayer template, these settings concern the name of the ELAN transcript that should be processed and the primary tier that should be presented in the output. (Remember the limitation in the current pre-release version of CuPED that clickable stories may only have one speaker; this will hopefully be improved in later versions to make full use of ELAN's more robust tier system!). Since these settings are required to produce output – you need to provide a transcript and a tier name before the Flowplayer template can do anything else! – they are marked by default with a star. Optional settings – things which *can* be customized, but can also be skipped over if you're in a hurry – are marked with a heart, instead.

When a setting is currently unavailable or invalid, CuPED indicates it with a red icon. In Illustration 24, since no transcript file and no primary tier have been selected yet, both of their associated icons are red. In both cases, we need to make a selection in order to proceed. Click on the Transcript file option in the left-hand list of template

settings, then click the Add button towards the bottom of the window to find an ELAN transcript to process. If CuPED is able to load this transcript, the icon next to the Transcript file option will turn green.
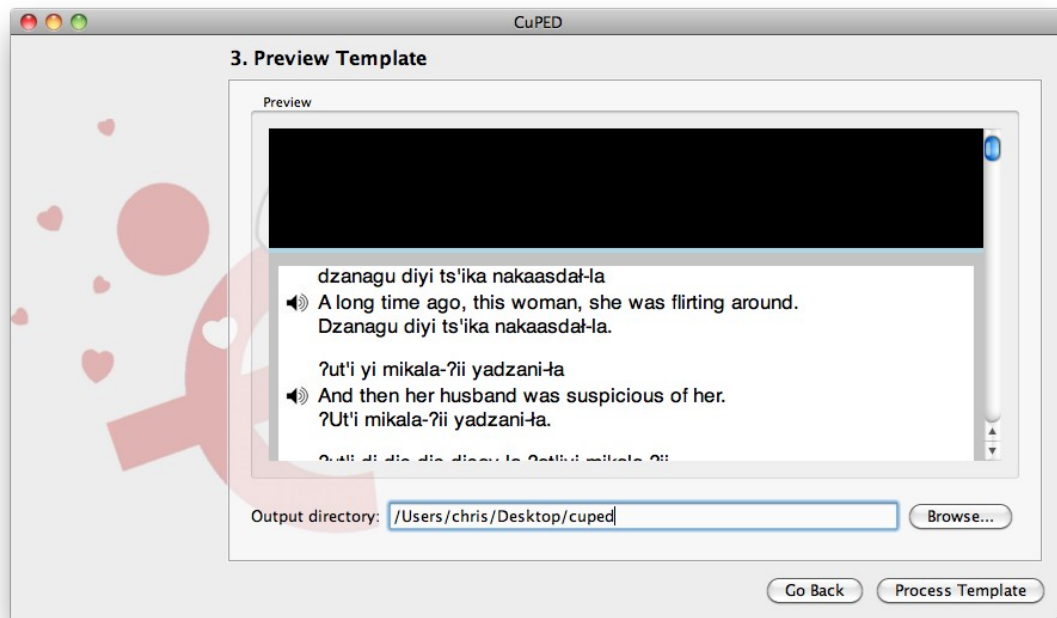

Illustration 25. Previewing a template in CuPED.

Once you have added your transcript, you should be able to select a tier within it to use as the primary tier in the clickable story. The primary tier will be the one displayed most prominently in the final story, along with all of its child tiers. Click on the Primary tier option in the left-hand list of template settings, then check one of the checkboxes that appears next to the names of tiers in the Current setting box. (You can also customize how that tier will be displayed by clicking on the Options button, but that step can also be omitted).

Once all of the required options have been provided, you may click on Continue. This will take you to a preview of the final template. If you want to make changes to any of the display settings – selecting better fonts or hiding particular child tiers, all of which can be done by clicking on the Options button when a particular tier is selected – you can click on the Back button to return to the previous screen, where the template can

36

be customized further. If everything looks alright, click on the `Browse…` button to select a directory into which CuPED should save the final clickable story, then click on `Process Template` to have CuPED begin the conversion.
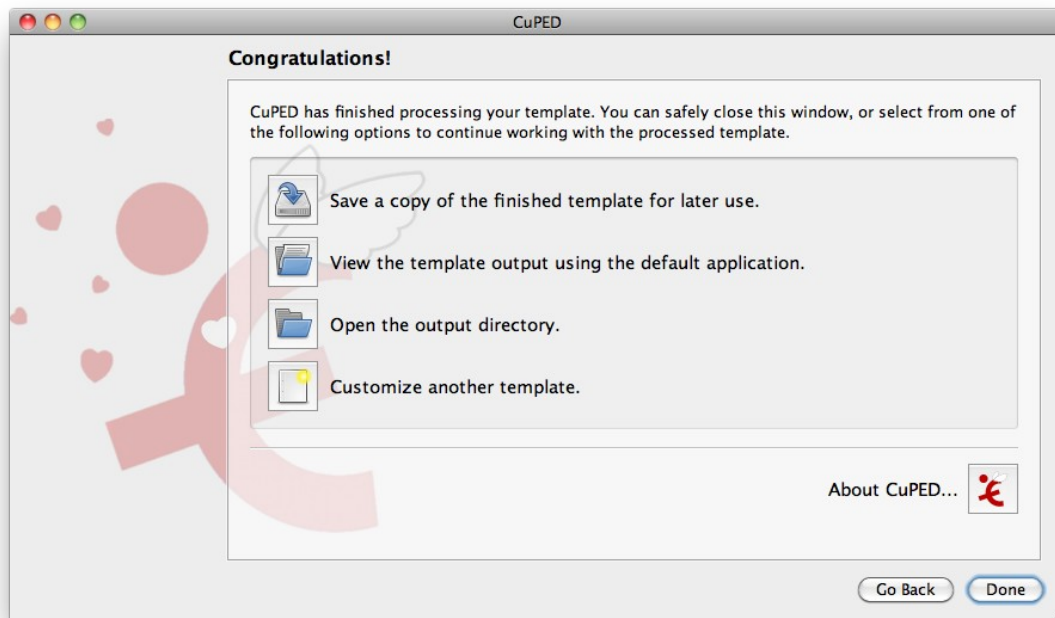

Illustration 26. Reviewing the finished clickable story in CuPED.

(If you are working with a transcript of video material, you will note that the preview screen does not show the video itself. Although the video *will* be present in the final clickable story, most media require some conversion before they can be displayed in the preview, and that processing is usually much too computationally intensive to perform for a quick-and-dirty preview. In other words, don't panic!) Reviewing and sharing a clickable story

Once CuPED has finished processing your template, you may now access the clickable story it has produced. In the final 'congratulations' screen, you may click on the `View the template output` button, and CuPED will automatically open a web browser and load your new clickable story, or you may simply find the folder where you asked CuPED to save the story and open the `index.html` file there yourself.

The clickable stories that CuPED produces use much the same technology as YouTube to provide audio and video through a web-based interface. Unlike YouTube, however, the CuPED output can be placed on individual computers, or on a private web server, or burned onto a data CD or data DVD for off-line distribution. To make the story available to others, simply copy the entire output folder's contents to the desired location – if placed on a web server, the story will immediately be available over the web, or if on a memory stick, it will be accessible to anyone able to open the index.html in the story folder.

Illustration 27. A clickable story produced by CuPED.

In some cases, the technology that CuPED uses to play audio and video in a web browser places restrictions on some of the interactive features of the clickable story: the individual annotations may not be highlighted automatically as viewers play through the recording, and the speaker icons next to each annotation may not be able

to play back sections of the recording.

To fix this, you can run a small work-around program from within the folder where the clickable story is found. Versions of this work-around program are freely available for Windows, Mac OS X, and Linux, and can be downloaded from sweet.artsrn.ualberta.ca/cdcox/cuped/DVD-Workaround.zip. Simply copy this `zip` file into the clickable story's folder, uncompress it, and run the version of the work-around program appropriate to your operating system (e.g. on Windows, double-click on `Windows Setup.bat`; on Mac, double-click on `Mac Setup`, etc.).