

13 Tugg: A transcription tool for language documentation

Magnus Ahltop
Nakajima Koen Research
Institute

Harald Berthelsen
Institute for Language
and Folklore

Modern language documentation is less about field linguists treating the speakers of a language as rather passive informants, and more about actively involving them in the work as language consultants. This includes handling tools like recording equipment and being able to transcribe language data on their own. Catering to non-expert users of transcription software presents new challenges for a category of tools that traditionally was the exclusive domain of expert users.

We will describe a new transcription tool that aims to be directly usable without any special knowledge of the tool, and only requiring modest amounts of previous computer knowledge for the end user. The tool is geared towards producing fully Leipzig glossing rules compliant transcriptions with a minimum of effort.

This chapter will describe the tool, how to set up the server part of the tool, and practical instructions on how to use it, including suggested workflows. The chapter ends with a brief description of a practical use case and future directions.

1 Introduction

Traditionally, the field linguist has treated the speakers of the described language as fairly passive informants. The research has often consisted of a linguist travelling to an area, collecting data, travelling home, and writing a book. The speakers of the language have often not received anything back and have not really learned anything new about their language or gained skills. To use the terminology of [Cameron et al.](#), this was *ethical research*, “research on” ([Cameron et al. 1992](#): 14) subjects.

The field of language documentation instead aims to be more reciprocal ([Dwyer 2006](#): 31) and involve the speakers, normally called consultants ([Dwyer 2006](#): 52). Involving the consultants can mean that they should be

able to handle technical equipment like recording tools and, after some basic linguistic training, also be able to transcribe the material on their own (Mosel 2006: 74). To once again use the terminology of Cameron et al., language documentation aims to be *empowering research*, “research on, for and with” (Cameron et al. 1992: 22) subjects.

Many tools for transcribing speech data are, however, very complicated to work with and export data from. This chapter describes a language-independent tool for manual transcription of sound files that is directly modelled on the Leipzig glossing rules (Max Planck Institute for Evolutionary Anthropology 2015) format in conjunction with a phonetic/phonematic transcription.

Two popular tools commonly used for transcribing speech are Praat (Boersma & Weenink 2022) and Elan (Max Planck Institute for Psycholinguistics, The Language Archive 2024). Praat has many features for analysing speech, making it suitable for users with a need for deep phonetic analysis. It can be configured to work in many different ways, but this also makes it less suited to users that only need to transcribe speech.

Elan is oriented more towards video transcription, and is often used with many annotation layers. A common use case is creating many overlapping types of annotations, such as what is said, eye gaze, and gestures. Both these tools support arbitrary layers.

Tugg is instead specialised only on transcribing speech, with no advanced acoustic analysis and a fixed layer structure, aiming to get the user up to speed as fast as possible with almost no training in the tool itself.

This chapter is not meant as a substitute for an education in either field linguistics or phonetics, so any advice given should only be seen as giving context to the use of Tugg.

2 Overview of Tugg

Tugg presents one sound file at a time on a web page that is scrolled in the vertical up-and-down manner that most users are familiar with (Nielsen 2005, Sherwin 2014). This avoids the unfamiliar and difficult to navigate horizontal timeline that is common in speech manipulation software, and also makes it natural to use the Page up/Page down keys. A horizontal timeline can both make even a relatively short recording seem unmanageably long and make the user often resort to zooming in and out of the material, switching between overview and detail. There are few spatial cues that facilitate navigation, and zooming changes what the timeline looks like, making it harder to associate a position with its content.

To be able to present the sound data in a vertically scrollable canvas, the sound file is automatically divided into chunks, approximately 5 seconds long, that are placed under each other. Tugg tries to divide the sound file at pauses in the speech, but if the user finds more suitable line breaks they can easily be inserted.

Each part of the sound file is represented by a spectrogram. A spectrogram is a visual representation of the sounds in the file, with white representing low intensities and black representing high intensities. The spectrogram shows not only the total intensity at a certain point in time but also how the intensity is distributed in the frequency spectrum. Intensities at the bottom of the spectrogram represent low frequencies, and intensities at the top represent high frequencies. Short periods of low intensity followed by a vertical line often do not indicate pauses in the speech, but are plosives, such as *p t k*. For examples of this, see Section 4.1, Figure 6.

The spectrograms used in Tugg are too small to be used for a detailed phonetic analysis, but help the user navigate and determine the start and end of utterances. Unless the user inserts a line break, the spectrograms will never change, which makes them a reliable spatial cue. The goal here is that even users that are not familiar with spectrograms will find the visual appearance useful as a navigation tool.

Tugg provides three transcription layers: phonetic/phonematic, gloss, and free translation. The purpose of the phonetic/phonematic layer is primarily as support for the glossing, but can be used to research intra- or inter-speaker variation if the researcher so chooses.

The gloss layer is for Leipzig glossing rules style morphosyntactic transcription and morphosyntactic glosses. Finally, the free translation layer is meant to run across the whole utterance.

The intended use of Tugg is that the user annotates one layer at a time by:

- marking a region (see definition in Section 4.1) in the spectrogram with the mouse,
- listening to the region with the apostrophe (') key, and
- entering the transcription.

Each layer behaves differently, according to the needs of that specific layer. Characters in the phonetic/phonematic layer are, for example, spread out over the whole region to maximise the probability that the phone/phoneme will appear close to the sound in the spectrogram. Text in the gloss layer is formatted according to Leipzig glossing rules, with each transcription and gloss token together.

Tugg can produce directly publishable L^AT_EX code formatted according to Leipzig glossing rules with automatic timestamps for each transcribed utterance.

3 *Setting up Tugg*

3.1 *Downloading*

Start by downloading the latest version of Tugg from <https://codeberg.org/nakajimakoen/tugg> and unpack it on your computer.

3.2 *Installing Python and necessary packages*

Install Python from your preferred package/software manager or download it from <https://www.python.org/downloads/>. The oldest Python version supported by Tugg is 3.9.

Tugg requires some Python packages:

- flask
- scipy
- pypng

Install them for instance using the package manager in your Python installation. If you use a unix system, you can probably install the packages you need by running the command:

```
pip install flask scipy pypng
```

3.3 *Configuring*

Place your audio files in the “audio” subdirectory of Tugg. All audio files need to be in the WAV audio format and end in “.wav”. You can also later place additional audio files in this directory at any time.

3.4 *Running the server*

Now start the server “server.py”. This starts a web server on your computer that only you can access.

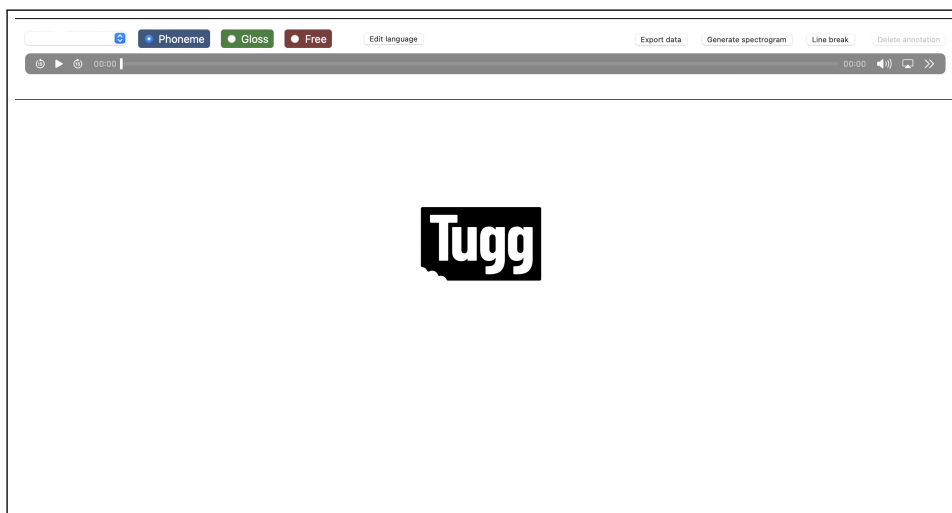


Figure 1: The interface when accessing the page the first time.

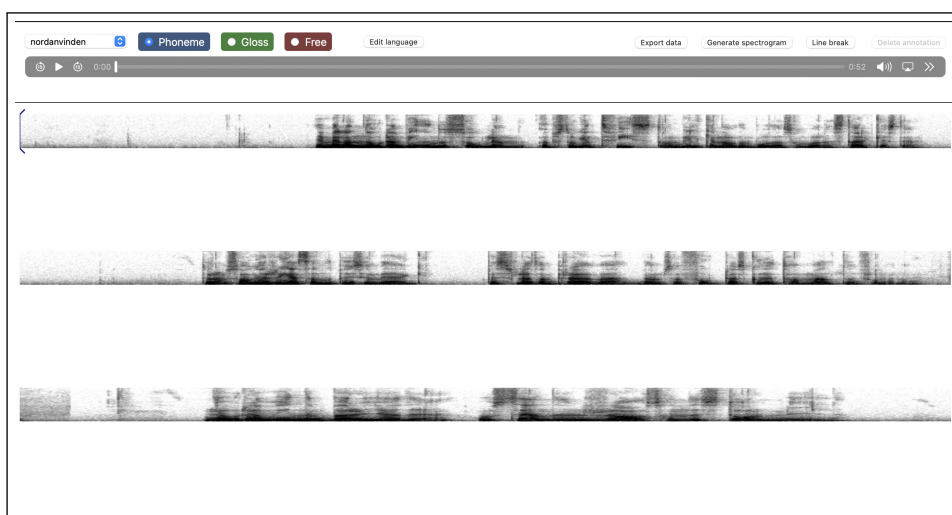


Figure 2: The interface after having selected a file to work with.

4 Using Tugg

Access the server using a web browser, using the address <http://127.0.0.1:5000>.

You should now see a screen similar to Figure 1. Select the audio file you want to work on in the upper left corner. You should now see the spectrogram of the speech data (Figure 2) with a sound player showing the whole sound

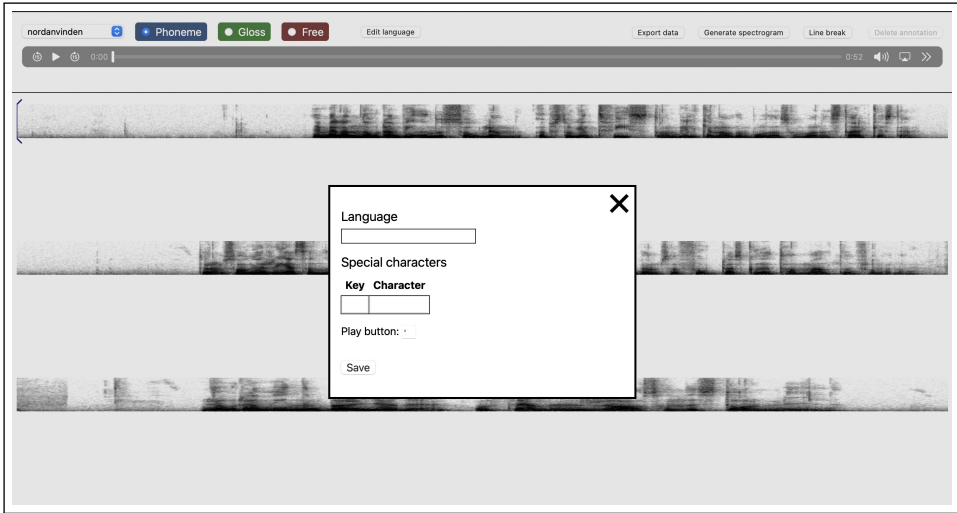


Figure 3: Settings dialog box for language and special characters.

file at the top of the screen. You can navigate either by using the sound player or clicking directly in the spectrogram. The recommended way of playing and pausing is by using the apostrophe (') key on your keyboard. This is to make it as easy as possible to play the contents of the region, without clicking any buttons with the mouse and at the same time leaving letter keys and the space bar available for typing.

If input of the apostrophe (') character is needed frequently enough in the annotation work that this becomes untenable, or the key is not placed conveniently on the keyboard, the key can be changed in the settings dialog box. Playing can also be started and stopped by using the sound player controls or the play/pause button on your keyboard (if you have one).

Tell Tugg what language you are working on by clicking the “Edit language” button at the top of the screen (Figure 3). Enter the language in the way you want it to look in your exported output. You can also assign special characters to keys in this dialog box. In the example in Figure 4, upper-case “D” has been assigned the IPA character for a voiced retroflex plosive, a sound needed to make the phonetic transcription we want. Saving the preferences makes the language name and the special characters appear at the top of the screen (Figure 5).

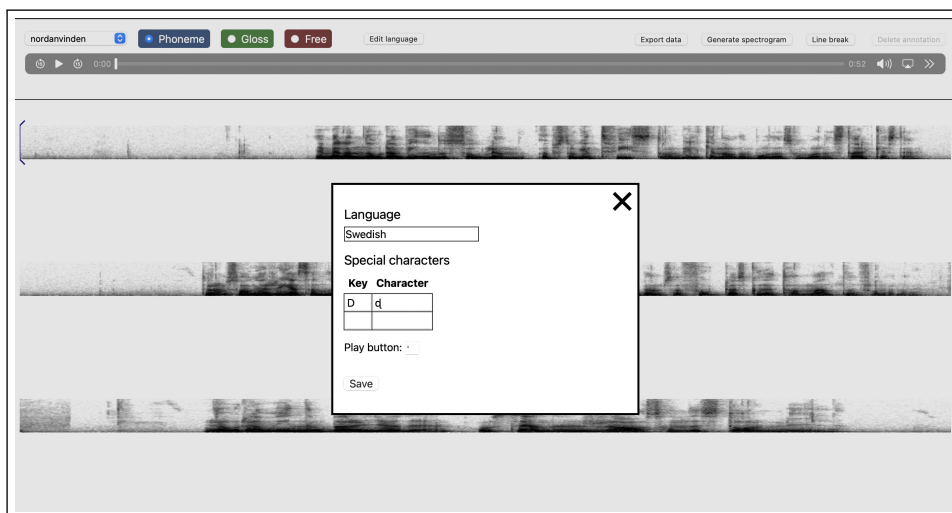


Figure 4: Entering language name and special characters.

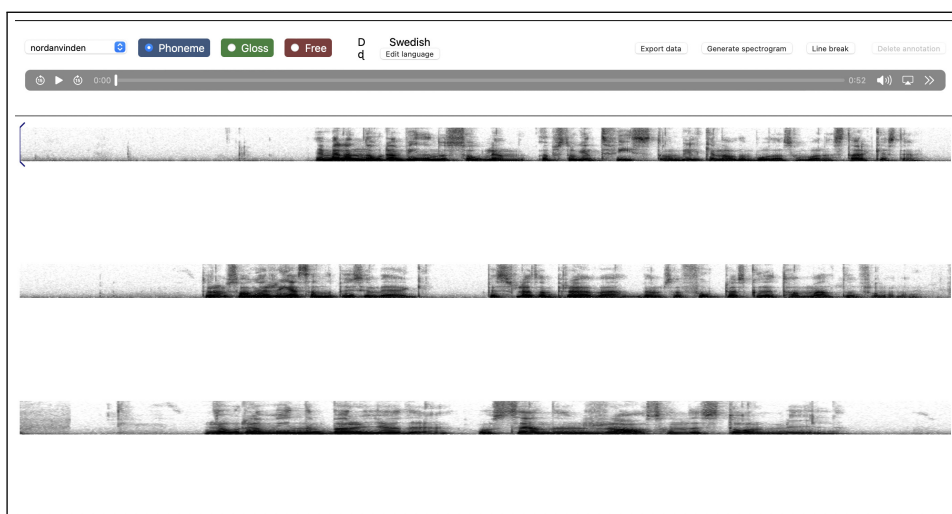


Figure 5: Interface after having saved language name and special characters.




Figure 6: The main interface elements of Tugg.

4.1 *The interface elements*

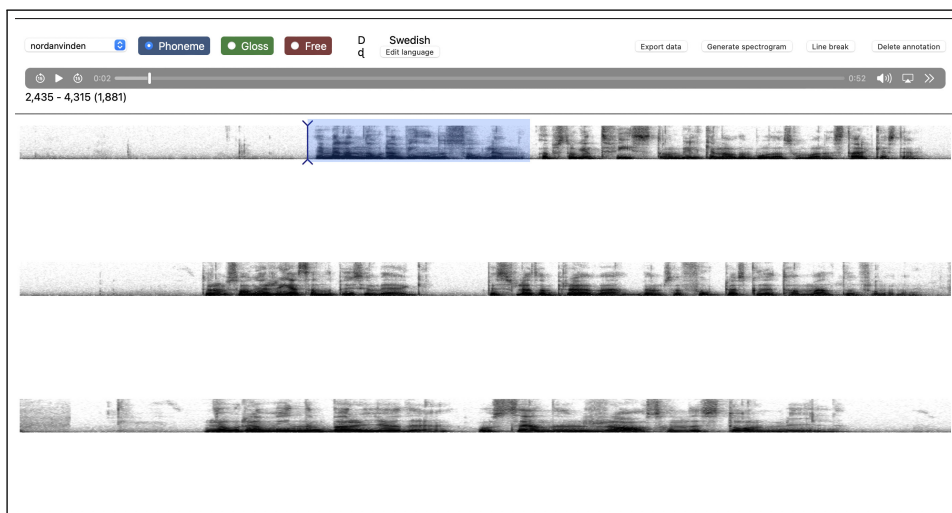
The main interface elements of Tugg are shown in Figure 6:

- the *cursor*: indicating where playback will start or (if playback is in progress) the current position of playback
- the *region*: a part of the recording
- the *spectrogram*: the time/frequency representation of the sound in the recording
- the *annotation*: the text corresponding to a region for a specific layer

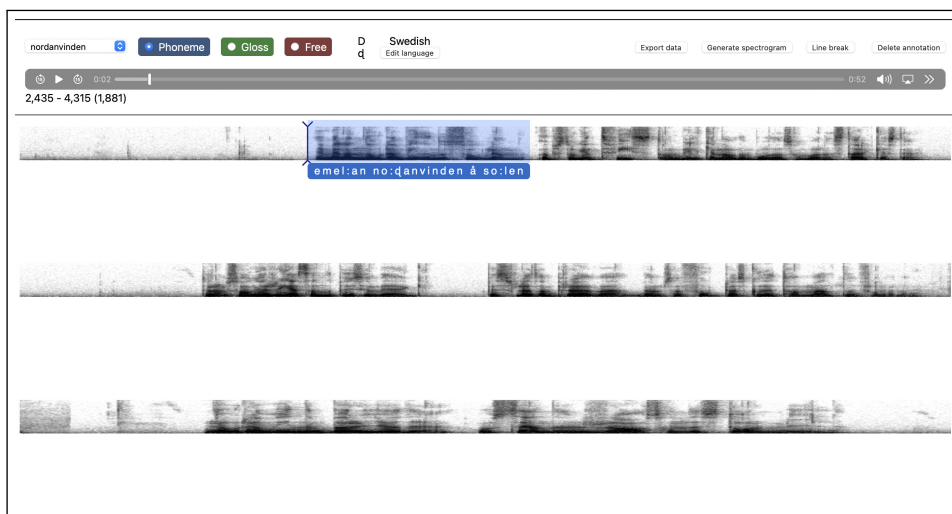
4.2 *Marking a region and creating an annotation*

Now we are ready to start annotating. An annotation is created by holding down the mouse button while dragging the mouse cursor from the position in the spectrogram where we want the annotation to start to where we want the annotation to end (see Figure 7). As soon as a region is marked, we can start entering text by typing normally. The appearance of the text will be different depending on the layer type. In the lower part of Figure 7, we can see an example of text that has been entered in the phonetic/phonematic layer. If you want to select a region that already has an annotation, click the annotation.

In the current version of Tugg, a region cannot span two spectrogram rows.



a.



b.

Figure 7: Selecting a region of the spectrogram. In (a), the selected region is marked in light blue and the spectrogram cursor is placed at the start of the region. Screenshot (b) shows the region after entering text.

4.3 *The different layers*

There are three different layers that can be used for annotations. The layer types are shown at the top of the screen in three coloured boxes (Phoneme, Gloss and Free), see for example Figure 7. The layer types are:

- *The phonetic/phonematic layer.* This layer is suited for both broad and narrow transcriptions. There is, however, no specific support for easily adding IPA diacritics, such as voiced/voiceless, advanced/retracted, and raised/lowered which might be needed in certain narrow transcription situations.

Annotations are limited to one row, and the entered characters will automatically fill the whole region. Text in this layer uses the “special characters” setting mentioned in Section 4.

This layer is presented in blue, and is always the top-most layer.

- *The glossing layer.* This layer is suited for Leipzig glossing rules annotations. Annotations use two rows, and pressing the return key will change to the second row. Tokens are separated by space, and have to contain the same amount of hyphens (-) in the top and bottom rows.

This layer is presented in green, and is always the middle layer.

- *The free translation layer.* This layer is suited to the idiomatic translation part of the Leipzig glossing rules. No special formatting is used in this layer.

This layer is presented in red, and is always the bottom layer.

An example of entered text in all layers can be seen in Figure 8.

4.4 *Playing a selected region*

You play a selected region by pressing the apostrophe (') key on your keyboard. Tugg plays the region and then stops. If you press the key before the region is fully played, playing stops and the cursor is placed at the beginning of the region again. This differs from playing when a region is not selected, where pressing the apostrophe (') key pauses, and pressing it again resumes from that position.

4.5 *Editing annotations*

Annotations can be edited in two ways: changing the boundaries of the region and changing the text. First select the annotation by clicking on it.

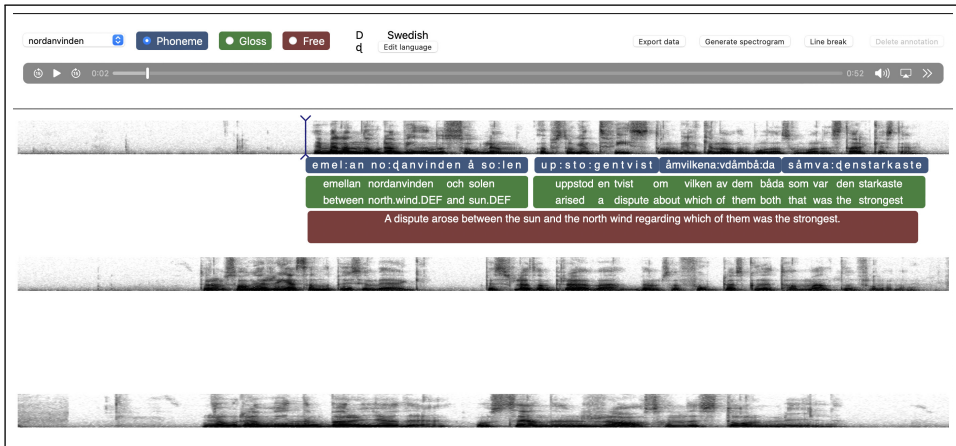


Figure 8: Example of entered text in all layers.

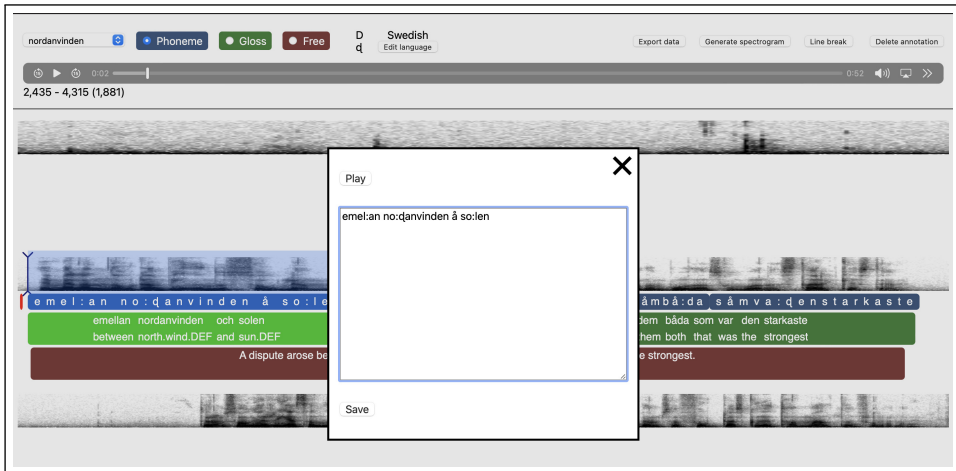


Figure 9: Editing an annotation in more detail. The user has here double-clicked an annotation, and a dialog box containing the text appears, allowing the user to edit the annotation as normal text.

Changing the boundaries is done by dragging the start or end of the region to the desired place.

The text is most commonly changed by just typing on the keyboard. This way, text can easily be entered without any extra steps. However, since changes can only be made to the end of the annotation, either by adding or removing characters, a more complicated but flexible way is also available.

By double-clicking an annotation, an edit dialog box appears (see Figure 9). In this dialog box, all standard edit tools can be used, such as editing in the middle of the annotation, selecting characters, and copy/paste.

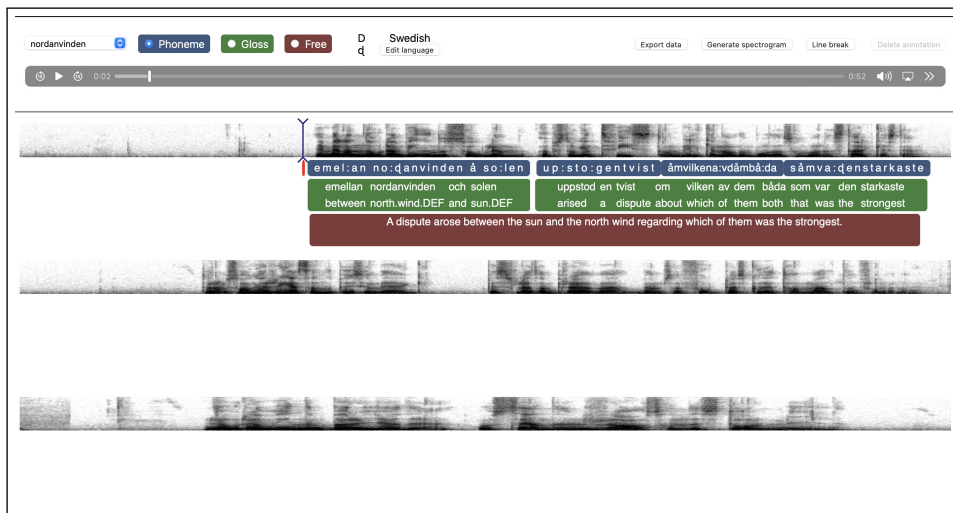


Figure 10: Creating a manual line break. The little red line under the spectrogram cursor represents the new line break.

4.6 Making line breaks

Tugg tries to place automatic line breaks where there are pauses in the speech data. When this automatic method is not enough, the user can insert manual line breaks. This is done by placing the cursor at the desired position and clicking the button “Line break” at the top of the screen. The line breaks are indicated by a little red line under the spectrogram cursor (see Figure 10).

When the manual line breaks have been placed, click the “Generate spectrogram” button at the top of the screen. All line breaks are now recalculated (see Figure 11).

4.7 Exporting data

Data can be exported into different formats. At the moment, three formats are supported directly from the interface: the Tugg format, \LaTeX gb4e or covington, and Praat TextGrid.

The Tugg format contains all the annotation and configuration data, and can be imported again into Tugg. Exporting to this format can be used for backup purposes, sending data to another person, or further processing.

Data exported as \LaTeX gb4e/covington can be used directly in \LaTeX documents using the common glossing packages gb4e and covington. The only necessary addition to the document is to include the correct package by using the `\usepackage{gb4e}` command for gb4e or `\usepackage{covington}`

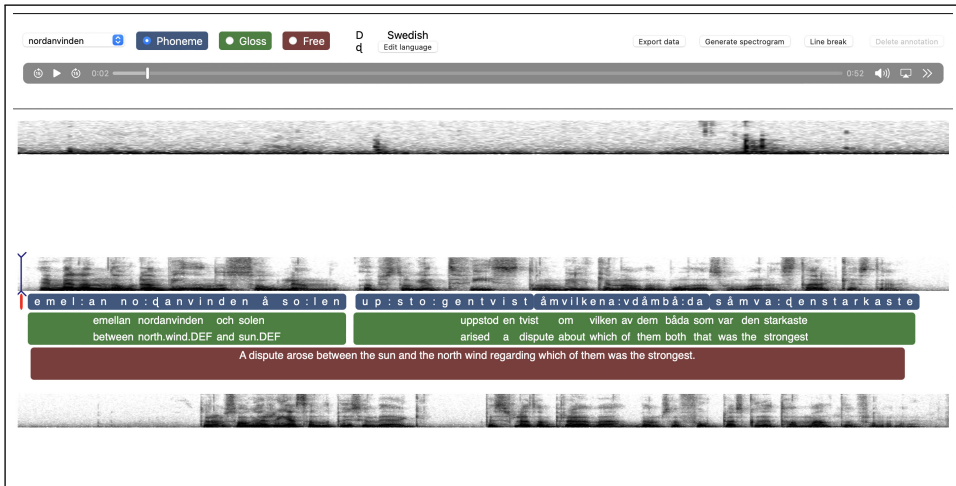


Figure 11: Recalculated spectrogram after creating a manual line break. The red line break symbol is now near the beginning of the second row. The row that was previously the first row is now split between the first and second row.

for covington in the beginning of the document. An example of this output and the formatted result can be seen in Figure 12.

Exporting is done by clicking the “Export data” button at the top of the screen. This brings up the export dialog box, where a format can be chosen. After a format has been chosen, the exported data appears in a text field. This can either be copied directly, or downloaded by using the “Download as file” link below the text field (see Figure 13).

5 Suggested workflows

Tugg can be used in several different ways. Here we present two different approaches, but there are of course many other workflows possible. The workflow suggestions are meant to be possible to be read independently of each other, so some text will be repeated.

5.1 Workflow 1: phonematic first

One possible workflow is the “phonematic first” workflow. This is suitable when the annotator has little knowledge about the language and first has to decode the phonological structure of the language. When using this workflow, start by playing the recording from the beginning and simultaneously watch the spectrogram.

```
\digloss[ex,preamble={Swedish (nordanvinden.wav 2.4--7.9s)}]
{emellan nordanvinden och solen uppstod en tvist om
vilken av dem båda som var den starkaste}
{between north.wind.\textsc{def} and sun.\textsc{def}
arised a dispute about which of them both that was
the strongest}
{A dispute arose between the sun and the north wind
regarding which of them was the strongest.}
```

a.

(1) Swedish (nordanvinden.wav 2.4–7.9s)
emellan nordanvinden och solen uppstod en tvist om
between north.wind.DEF and sun.DEF arised a dispute about
vilken av dem båda som var den starkaste
which of them both that was the strongest
‘A dispute arose between the sun and the north wind regarding which of
them was the strongest.’

b.

Figure 12: Example output from Tugg (a) and the same output formatted with the \LaTeX covington package (b).

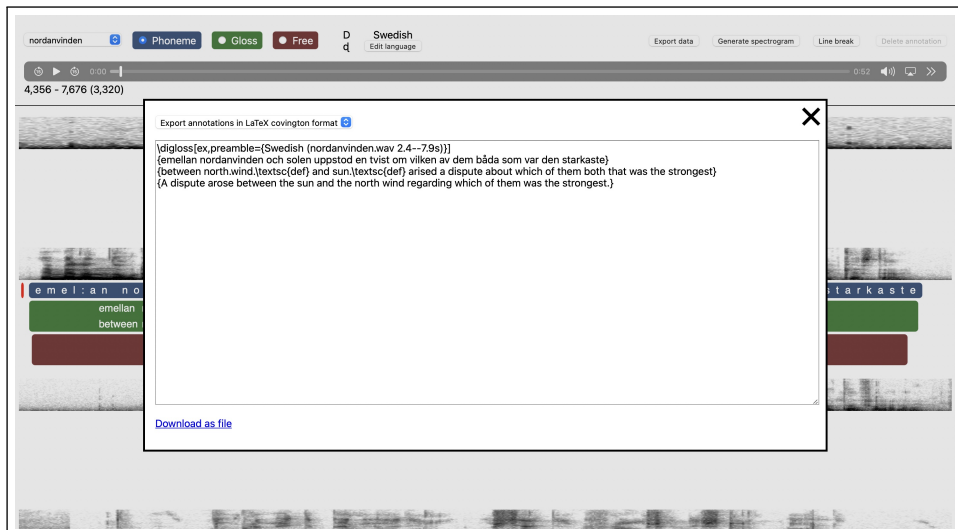


Figure 13: Exporting data from Tugg, in this case in the \LaTeX covington format.

When you have located a place that you are interested in transcribing, select the approximate region by clicking and dragging the mouse from the beginning to the end. If the region would cross two lines, first place a line break before the relevant part and press “Generate spectrogram” to move it up to the same row. Try using quite short regions for the phonematic transcription, a couple of seconds.

Play the region by pressing the apostrophe (') key. The sound of the region will now play, and stop at the end of the region. Adjust the boundaries of the region by clicking and dragging the left or right edge and play again until the region matches what you want to transcribe.

Now start typing on the keyboard, alternating between typing and playing. If you make a mistake, you can use the backspace key to erase one character at a time.

If you later discover an error in the middle of the annotation, you can double click the annotation to bring up an edit dialog box where you can make full edits to your annotation.

Everything is saved immediately (except in the edit dialog box), so no additional action is necessary.

If you need to go back and edit a previous annotation, just click the annotation and change either the text or the boundaries.

Depending on how you want to approach the task, you can either annotate everything in the phonetic/phonematic layer first, before you proceed to the glossing stage, or gloss when you have completed one or a few rows of phonetic/phonematic annotations. If you choose the second approach you will have to switch between different layers often, so remember to switch layers before you select a region, since data entered into one layer cannot be moved to another layer without using the edit dialog box.

Annotations in the the glossing layer are meant to be longer than the annotations in the phonetic/phonematic layer, but in the end it depends on the nature of the task and what feels natural for the annotator. Since the annotator has the support of the phonetic/phonematic layer, it is often convenient to use a bit longer annotations. The regions can be played in the same way as before, by pressing the apostrophe (') key. One reason to not have very long annotations in the glossing layer is that edits are made on the whole annotation, starting with the transcription row and only afterwards on the gloss row.

Finally, the free translation layer (called the idiomatic translation in the Leipzig glossing rules) is meant to span a whole utterance. This is what is used as the basis for what counts as an utterance when exporting the glosses to a text document.

5.2 *Workflow 2: translation first*

If you already have easy access to a translation of the utterances, perhaps because you are a native speaker or because you have asked a native speaker for translations as the first step of your analysis, it can be useful to use the “translation first” workflow.

When using this workflow, start by playing the speech from the beginning and simultaneously watch the spectrogram in the same way you would in the “phonematic first” workflow, but listen for utterance boundaries. Before the first utterance, insert a line break by pressing the “Line break” button. Then insert a line break after each utterance, and if there is a long pause, also insert a line break just before the pause is over.

When you have inserted line breaks throughout the recording where needed, regenerate the spectrogram by pressing the “Generate spectrogram” button. Tugg will now take your line breaks into consideration when deciding what is displayed on each row of the spectrogram.

Continue by selecting the “Free” layer type and making a region that encompasses the entire utterance. The length of the free translation annotations is what will be used as the basis for what counts as an utterance when exporting the glosses to a text document.

Play the region by pressing the apostrophe (‘) key. The sound of the region will now play, and stop at the end of the region. Adjust the boundaries of the region by clicking and dragging the left or right edge and play again until the region matches what you want to transcribe.

Now start typing on the keyboard, filling in what the translation of the whole utterance is. If you want access to full editing tools for your text, simply double click the annotation to bring up an edit dialog box, otherwise the only editing available is pressing the backspace key to erase one character at a time. Everything is saved immediately (except in the edit dialog box, where you have to press the “Save” button), so no additional action is necessary.

If you need to go back and edit a previous annotation, just click the annotation and change either the text or the boundaries.

5.2.1 *Subworkflow 2a: glossing directly from speech*

If the annotator has direct access to the words of the language, for example by being a native speaker, or otherwise knows it well, glossing can start directly from the speech data. Beware that if the orthography you use is not phonological, or you want to investigate phonological processes, a better approach might be to use subworkflow 2b instead.

First switch to the glossing layer type by clicking the “Gloss” button at

the top of the screen. It is often practical to use smaller gloss regions than the whole utterance. This ultimately depends on the nature of the task and what feels natural for the annotator. One reason to avoid very long annotations in the glossing layer is that edits are made on the whole annotation, starting with the transcription row and only thereafter on the gloss row.

In this subworkflow, the phonetic/phonemetic layer is only necessary if you need this data for other purposes than Leipzig style glossing. The data will not be used in formatted glossing output.

5.2.2 *Subworkflow 2b: glossing from phonetic/phonemetic data*

If the annotator has little knowledge about the language and first has to decode the phonological structure of the language, creating a phonetic/-phonemetic annotation before glossing might be a good idea.

Follow the instructions in Section 5.1 except the paragraph about the translation layer.

6 *Practical use case*

One of the authors initially created Tugg as part of a course in field linguistics, and this section will briefly walk through the process of using Tugg in this course.

6.1 *Background*

The language studied was the Indo-Aryan language of Khowar, spoken primarily in northern Pakistan. The language consultant involved in this project grew up in the district of upper Chitral in Pakistan and is a native Khowar speaker. The consultation sessions took place over a teleconferencing platform, with the consultant in Islamabad (Pakistan), recording their own voice and sending the audio files after each session, which ensured good audio quality that did not depend on the quality of the teleconferencing platform or any network problems.

The aim of the study was to investigate the construction of polar (yes/no) and content questions in Khowar. An additional goal was to provide enough clearly sourced transcriptions showing examples of the language data.

6.2 *Session structure*

Interactions during the sessions consisted mainly of elicitations, but also explanation parts, where the documenter asked the consultant for clarifica-

tion regarding certain previous utterances, sometimes by playing back these utterances using Tugg and sometimes by asking follow-up questions.

The elicitations ranged from how to greet people in different situations and common daily interactions to more specific questions on how to reformulate a certain type of utterance that had already been encountered.

6.3 *Transcription*

Some of the transcriptions were made offline, between sessions. This was primarily the phonematic transcriptions, which required a lot of analysis to construct a hypothesis of the phoneme inventory of Khowar. This has been described before, but the described study was purposefully made without any prior knowledge of Khowar.

The glossing, and especially the free translations, were in large parts entered directly during the explanation parts of the sessions. Since all sessions were recorded, there was always the possibility to go back and fill in translations afterwards, but this was made somewhat more difficult by Tugg not having any system for linking audio files. When working on the translation of one file, the spoken translation from the consultant was recorded in another session and thus stored in another file. This necessitated the use of auxiliary text files and manually keeping track of what position in which file contained the explanations/translations.

A mix of workflow 1 and 2 was used. When the language consultant had provided translations before or during the annotation work, workflow 2 was used, otherwise workflow 1 was used. As the amount of recorded speech data was quite large, only a small portion of the data was phonematically annotated. Data that was neither phonematically annotated nor translated was left unannotated. In total, 228 annotations were made over 65 utterances (about 3.5 annotations per utterance, or about 1.2 annotations per utterance and layer). Most utterances were selected for their relevance to the study aim of polar/content question construction. Especially very early in the study, some were instead randomly selected because of the need to investigate the phoneme inventory.

The speed of making annotations with Tugg was the most obvious benefit. In Praat, annotations are made by splitting up the timeline, while in Tugg regions are created with one dragging action with the mouse. Also, a lot of zooming in and out is needed in Praat to switch between overview and detail, something that is not needed in Tugg thanks to the vertical organisation of the annotations.

Using Praat, the examples would have had to be formatted manually to present in the final report. Since all the formatting work was done inside

Tugg, going back and making changes to the analysis did not feel as much of a burden as if every modified example would have had to be reformatted manually.

The vertical organisation also feels more like a traditional document, making it easier to show to the language consultant.

7 *Conclusions and future improvements*

Tugg has now reached the point where it can begin to be useful in real-world contexts. Audio files can be transcribed using three specialised layers, a phonetic/phonematic layer, a glossing layer, and a free translation layer. It uses the more familiar vertical scrolling method for navigation, combining the automatic splitting of lines with the possibility of manual partition. Tugg can also export the transcriptions to a few widely used formats: code for the two L^AT_EX glossing packages `gb4e` and `covington`, making the transcriptions directly publishable, as well as the Praat TextGrid format for further analysis and processing.

Even though Tugg has been tested in a small user study and used for a few transcription tasks, leading to several improvements, it has not yet been tested on non-linguists. Going forward, testing Tugg on non-linguists should be one of the main priorities, since non-expert usability is one of the goals of the tool.

At the moment, regions and annotations cannot span multiple spectrogram rows. Since this limits the length of utterances that can be annotated, implementing multirow annotations would enable the annotation of longer utterances. It is important, however, to do this while keeping much of the simplicity of the current interface, since that is a key advantage of Tugg.

More export and import options would be useful for users wanting to annotate the same speech data in different applications. Tugg already supports exporting to the Praat TextGrid format, but handling other formats and being able to import the TextGrid format would improve interoperability.

The current version of the edit dialog box shows the raw text without any formatting. A future version could align the words in the upper and lower row, in the same way the normal editing does. Combined with features for controlled vocabulary of abbreviations and lexicon support would greatly improve the glossing process.

For the phonetic/phonematic annotations, a way to align the annotations to the sounds in the spectrogram could be useful for certain applications. A more detailed spectrogram view, in for example the edit dialog box, could enable certain simple analyses without resorting to a more complicated application.

Adding a linking feature, where timestamps in one file can be linked to timestamps in another would make recordings mentioning other recordings much easier to keep track of. Also, adding a notes feature, where the annotator can scribble non-structured metadata about a certain section of the recording, would be very helpful.

Finally, all annotation types would benefit from the ability to merge and split annotations.

We plan to use and evaluate Tugg in other transcription projects. This will lead to more features being incrementally added, but Tugg is already usable in its current state, supporting the basic needs of many transcription tasks.

Acknowledgements

The authors would like to thank Natalia Kalyva and Henrik Liljegen for their valuable input.

The development of Tugg was partially funded by the Swedish Research Council: Swe-CLARIN/The National Language Bank of Sweden (2017-00626).

References

- Boersma, Paul & David Weenink. 2022. *Praat: doing phonetics by computer*. <http://www.praat.org/>.
- Cameron, Deborah, Elizabeth Frazer, Penelope Harvey, M. B. H. Rampton & Kay Richardson. 1992. *Researching language*. Routledge.
- Dwyer, Arienne M. 2006. Ethics and practicalities of cooperative fieldwork and analysis. In Jost Gippert, Nikolaus P. Himmelmann & Ulrike Mosel (eds.), *Essentials of language documentation*, chap. 2, 31–66. De Gruyter Mouton.
- Max Planck Institute for Evolutionary Anthropology. 2015. *Leipzig glossing rules*. <https://www.eva.mpg.de/lingua/pdf/Glossing-Rules.pdf>.
- Max Planck Institute for Psycholinguistics, The Language Archive. 2024. *Elan (version 6.8)*. Nijmegen. <https://archive.mpi.nl/tla/elan>.
- Mosel, Ulrike. 2006. Fieldwork and community language work. In Jost Gippert, Nikolaus P. Himmelmann & Ulrike Mosel (eds.), *Essentials of language documentation*, chap. 3, 67–86. De Gruyter Mouton.
- Nielsen, Jakob. 2005. *Scrolling and scrollbars*. <https://www.nngroup.com/articles/scrolling-and-scrollbars/>.

Sherwin, Katie. 2014. *Beware horizontal scrolling and mimicking swipe on desktop*.
<https://www.nngroup.com/articles/horizontal-scrolling/>.

Corresponding author

Magnus Ahltop
Nakajima Koen Research Institute
magnus@nakajimakoen.org