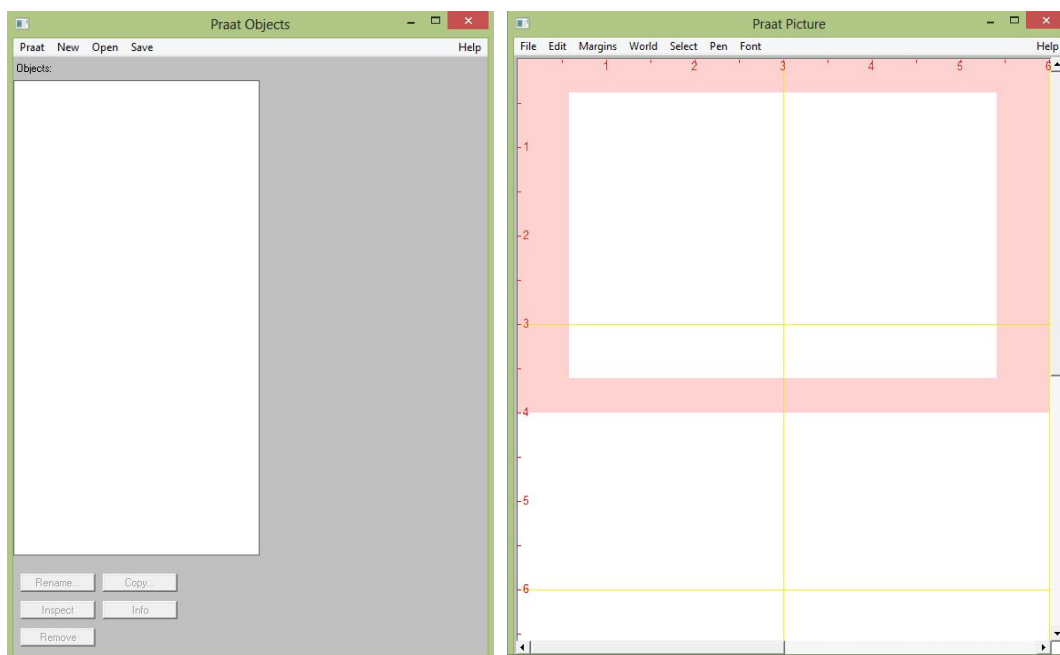


Praat — the basics

1 interface

When you open Praat, two windows will pop up. One is called 'Praat picture'. It is pretty unlikely that you will ever need this window in the course of this class and if indeed you do, it will pop up automatically so you might as well close it for now.

The other one is called 'Praat Objects'. This is your main workspace. You load and save files from here, start all sorts of analyses and generally manage all the files you're working with. You will find further explanation of the different things you can do here below.



(a) Praat Objects window

(b) Praat Picture window

2 loading data

2.1 recording with Praat

Praat can, of course, work with sound files that you recorded using some other software (see the website for the list of formats). However, it is also possible to record directly in Praat. To do this:

1. Click on 'New' and 'Record mono Sound...' (shortcut: CTRL + R).
2. Choose a sampling frequency.
(44100 Hz is enough for everything we are interested in; if you are only looking at vowels even 22500 Hz should be fine)
3. Click on 'Record' to start the recording and on 'Stop' to stop it.
4. Name the recording (bottom right corner) and click on 'Save to list & Close'.

Your recording should now be listed in the objects window.

2.2 loading files

1. To load a sound file that is shorter than a few minutes just
 - a) Click on 'Open' and 'Read from file...' (shortcut: CTRL + O).
 - b) Choose your file and click 'Öffnen'.

NOTE: If you try to load very large files this way, Praat will freeze!
2. If you are working with longer sound files (anything up to 2 GB) then
 - a) Click on 'Open' and 'Open long sound file...' (shortcut: CTRL + L).
 - b) Choose your file and click 'Öffnen'.

3 spectrograms

3.1 viewing

To view a spectrogram, click on a sound in your objects window and then on 'View & Edit' (second button from the top). A new window will open, the spectrogram of your sound file.

- time is marked on the x-axis
- frequency is on the y-axis
- different shades of black and grey indicate intensity




Figure 2: spectrogram of /i/

The darker an area in the spectrogram is the more energy there is in the sound around this frequency and at this given point in time. If the lower part of your screen is largely white you will probably have to zoom in (see below) to see the spectrogram.

You can play (parts of) the sound by either clicking on the bars at the bottom of the window or by clicking anywhere in the spectrogram and hitting `TAB`. The file will then start playing from the position you clicked. Stop the playback with `ESC`.

3.2 zooming, selecting, editing

In the bottom left corner you will find a number of buttons for zooming. 'in' and 'out' should be self-explanatory. 'sel' zooms in so that your selection fills the entire window. 'all' shows the whole recording and 'bak' returns to the previous zoom window.

Selecting is easy. Just click (in the spectrogram) where you want your selection to start, and hold the key while dragging to where you want the selection to end. See the 'Select' menu for options like moving the cursor to a precisely specified time. One special command might be of interest, especially if you want to edit your sounds. 'Move start (or end) of selection to nearest zero crossing' can be used to make sure there is no unwanted noise when you copy-paste bits of sound.

Editing is straightforward as well. Once you have made a selection click on 'Edit' and 'Cut' (shortcut: `CTRL + X`) to cut it. Placing the cursor somewhere and clicking 'Edit', 'Paste after selection' inserts a sound from the clipboard at the current location. Finally, selected bits can be saved as separate sounds via 'File' and 'Save selected sound as...'

3.3 formant measurements (vowels)

Praat can do a lot more than just select, edit, and play sounds. It is possible to actually *measure* them. This is particularly easy for vowels. Vowels have a spectral makeup that consists of a number of prominent frequency ranges. These frequencies are called formants, and they are somewhat similar to overtones in music. Essentially, the first two formants are enough to define a vowel. These formants can be measured in Praat.

1. Make sure formant analysis is actually enabled by clicking on 'Formant' and 'Show formants'. The formants should now be visible as dotted red lines.
2. Adjust the formant settings ('Formant', 'Formant settings...')
At the very least, you have to use different settings for men and women. The relevant variables are 'Maximum formant (Hz)' and 'Number of formants'. Standard values are 5000-4 for men and 5500-4 for women.
3. Click where you want to take your measurement (e.g. in the middle of the vowel).
4. Click 'Formant' and 'Get first formant' (shortcut: F1) to get the measurement for the first formant. The same works for the second, third, and fourth formant.
Doing this while you have an active selection will give you the mean formant value for this interval.




Figure 3: formants of /i/

If you want to save the position of the formants as an image file click on 'Formant' and 'Draw visible formant contour...'. The Praat picture window will open again and you can save the picture as a file.

3.4 pitch

Measuring pitch is very similar to measuring formants.

1. Make sure pitch analysis is enabled by clicking on 'Pitch' and 'Show pitch'. Pitch is drawn as a blue line.
2. Again, you have to adjust the settings ('Pitch', 'Pitch settings...')
Suggested pitch range: 75 to 300 Hz for men, 150 to 500 Hz for women
3. Click where you want to take your measurement.
4. Click 'Pitch' and 'Get pitch' (shortcut: f5).

Again, doing this while you have an active selection will give you the mean pitch for this interval. The commands 'Get minimum pitch' and 'Get maximum pitch' do just what their name implies.

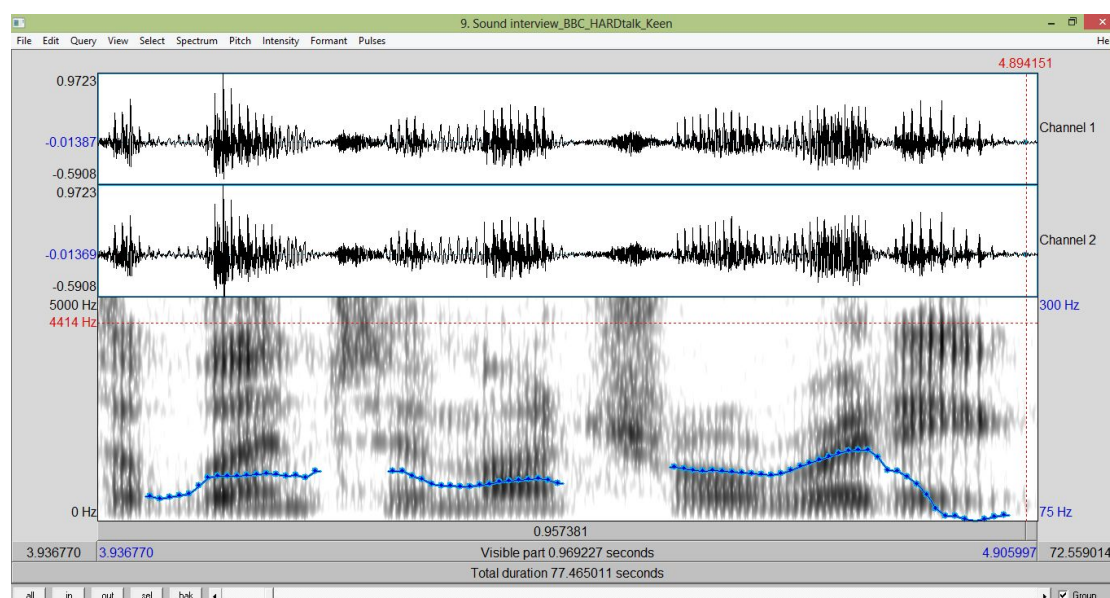


Figure 4: pitch of *they've lent too much already*

Just as with formants you can display the pitch contour in the picture window ('Draw visible pitch contour...').

Finally, pitch can also be modified.

1. Go to the objects window and select the sound you want to alter.
2. Click on 'Manipulate -' (fourth button from the bottom) and 'To Manipulation...'
3. A new object has appeared in your objects list. Mark it and click 'View & Edit'.

4. You can now alter the pitch contour (hold `SHIFT` to move several anchor points together) and save the altered sound (to the objects list) by clicking on 'File' and 'Publish resynthesis'.

3.5 intensity

Intensity (\approx loudness) works exactly like pitch, except that you usually shouldn't need to change any presets.

4 text grids

TextGrids are a way of annotating text and are great for stuff like conversation analysis or just better orientation within the sound file. Also, many of the scripts (see below) that you can use for automatic data processing need a TextGrid as a basis.

1. Mark a sound in your objects list and click 'Annotate -', 'To TextGrid...'.
2. Name the tiers, e.g with 'speaker A' and 'speaker B'. You probably won't need any point tiers, so make sure this box is empty.
3. The TextGrid appears as a new object in your objects list. Mark both the sound and the corresponding TextGrid and click on 'View & Edit'.
4. The active tier is yellow. Now click somewhere in the spectrogram and start the playback. Pressing `ENTER` will add a boundary on the selected tier at the corresponding time.
5. These boundaries mark off intervals on the tiers. You can label those intervals by simply marking them and typing in any text you like, e.g. the transcription of the sound passage.

This can be done individually on every tier, so if you had a tier for every speaker in your recording you could, for instance, easily see who interrupts whom, how often and how much they talk at the same time etc.

Boundaries can be moved around by simple drag-and-drop. If you want to delete a boundary, mark it and press `ALT + BACKSPACE`. Don't forget to save your TextGrid by clicking 'File' and 'Save TextGrid as text file...' (shortcut: `CTRL + s`)!

5 scripts

Scripts are a great way of (semi-)automatizing all the things that Praat can do. They are essentially little programs that take care of all the tedious, repetitive bits of data processing. A good script can easily save you hours and hours of work. Load a script by clicking on "Praat", "Open Praat script". Run it by clicking "Run" and "Run" (shortcut: `CTRL + R`).

You can either write a script yourself, or – and that's by far the easier version – you can try to find a ready-made one that you can use or that you only have to adapt a bit. Dozens of scripts can be downloaded for free and there's a good chance there's already one around that does just what you want. Examples of what you can do automatically with scripts:

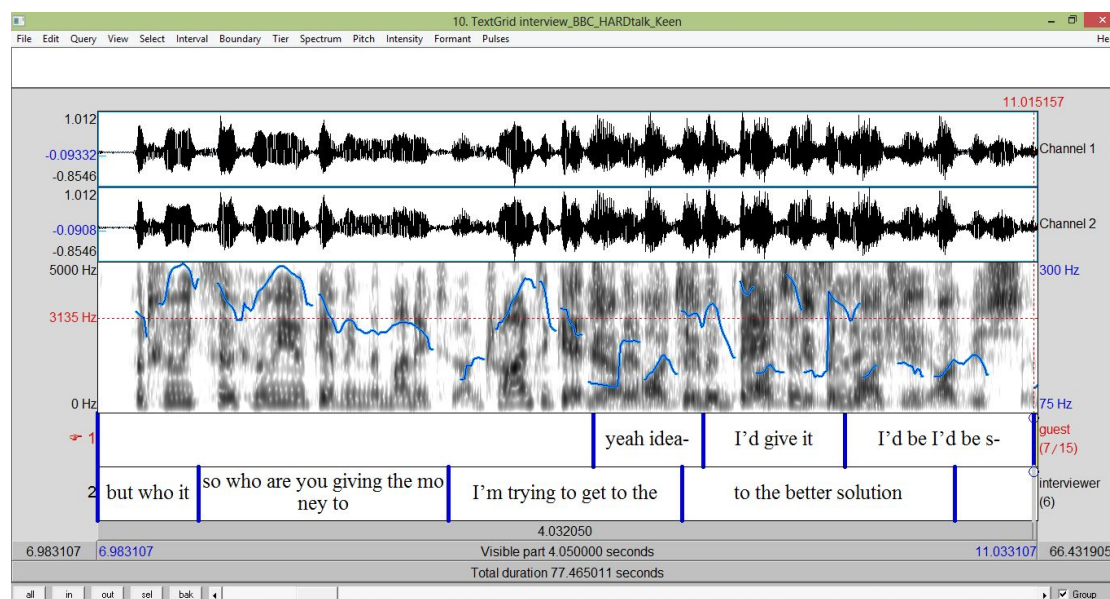


Figure 5: example of a TextGrid, aligned with sound file

- split sounds or TextGrids into separate files
- measure pitch and formants
- measure the duration of sounds or of whole intervals (→ TextGrid)
- write all of this into a nice table that you can directly feed into your statistics software.
- modify pitch (easy!), duration (easy!), or formants (more complicated...)
- ...

An interesting website regarding this topic can be found here. Don't hesitate to get in touch with me. I'll gladly help you to find a suitable script if applicable.