

Lesson 2: Exercises

Learning goals

- Multi-tier speech annotation and some use cases for it
- Using various sound visualisation methods (waveform, spectrograms, intensity and pitch curves) as listening aids when transcribing and annotating speech
- Phonetic annotation (segmenting and labeling speech sounds) in Praat
- Measuring the durations of annotated units manually
- First contact with a Praat script (automatic calculation of a number of segmental durations)

Exercises

Tip: Divide the work into several sessions - do not try to finish everything in one go!

Preparing your speech material for research

1. Review the background material about [annotating speech](#) (see lesson 1). How would you need to annotate your speech material? How many annotation tiers would you need?

Can you think of any properties in speech that should be described in a PointTier rather than in a regular IntervalTier?

Using the analysis displays in the annotation editor

The following exercises are intended to give you a first impression of what the different analysis displays mean and how they can be used to support listening. It is very useful if you discover situations where your own experience seems different from what the acoustic visualisations and graphs are telling you.

2. Open a sound file that contains some read-aloud text and not much background noise. For instance, use a sentence from the American English recording of the story *North Wind and the Sun* (look in the folder *Narrative* in the package you downloaded from the [IPA Handbook website](#) during lesson 1). To start from scratch, create a new TextGrid object for the sound object. Create three annotation tiers: one for phones (or speech sounds), one for words, and one for sentences.
3. In the TextGrid editor window, switch on the spectrogram (if required, select **Spectrum: Show spectrogram**), the pitch curve (**Pitch: Show pitch**), and the intensity curve (**Intensity: Show intensity**). Zoom in and out, scroll back and forth, and listen to short portions of the sound, one at a time, while looking at the acoustic displays. Try to find out if and how you could identify the following types of events, at least to some extent, just by looking at the sound waveform and the analysis displays:
 - pauses / silence
 - background noise
 - the highest pitch, or a falling pitch in the speaker's voice

- voiced sounds
 - [s] sounds (sibilants) or other fricatives
 - voiceless stops [k, p, t]
 - different types of vowels (e.g., [a] vs. [i])
 - What other discoveries did you make?
- (Not all of the analysis displays will be useful for all of the above tasks...)

Phonetic annotation (segmentation and transcription)

4. In the sound file you are working with, delineate and annotate one sentence, the individual words in it and all the speech sounds within one word. For each type of unit, use the corresponding annotation tier. Note that there are usually no pauses between consecutive sounds and words in speech! The locations of the boundaries will remain inaccurate and “fuzzy” to some extent and ultimate “correct answers” will not be available. However, try to be as precise as you can!
5. Save the TextGrid you edited.
6. Zoom in on the word whose individual sounds you annotated. Take a screenshot of this view, save it in an image file (e.g., PNG or PDF) and submit it as your answer to the speech annotation assignment in the course area.
7. Can you find out the exact duration of a selected portion of sound, in seconds, with lots of decimals? (Tip: take a look at the menu commands or other things you see in the editor window...)

Think about where you could use duration measurements in your own studies or research.

8. Now, let's take a first look at Praat scripts! Scripts are sort of computer programs: they include sequences of commands that can be run within a particular application program or environment. There is a scripting language built in Praat, too. With Praat scripts, the regular features in Praat can be extended to suit particular purposes, and many repeated tasks can be made automatic and less error-prone. In the rest of the lessons in this course, we will be test-running several scripts, just to give you an idea of what they can do for you. If this sounds completely new to you, *don't worry, you will not be required to write scripts of your own in this course!* All you need to learn is how to run a script written by someone else.
9. Download the script called **calculate_segment_durations.praat**. Save the script in a convenient location on your computer, somewhere you can easily find it. You can create a folder for Praat scripts if you like.
10. The purpose of the script you just downloaded is to measure the durations of all the labeled segments in a TextGrid object that is selected in the Object list in Praat, and to save the results in a plain text file. So, in Praat, open a TextGrid file where you have annotated some sound segments with labels. If you do not have such a TextGrid yet, you can for instance use the TextGrid file with the labeled vowels for a Finnish question, **kysymys7.TextGrid**.

11. Now, still in Praat, open the script file you just downloaded: *calculate_segment_durations.praat* (**Open: Read from file...**, or **Praat:Open Praat script...**). Then, make sure the TextGrid object to be analysed is selected in the Object list.
12. Try to run the Praat script by selecting Run:Run in the script editor window. Note that the script requires the Object window to be active when the script starts running, and the script will fail if you were working in an editor window just before you started the script. (To be sure, you can just close all other editor windows except the Praat script.)

The script starts by asking for the name of the tier where durations are to be measured. The tier name must be typed exactly the way it is called in the selected TextGrid. Next, you need to select the folder and the file name for saving the results. By default, the file will be called *durations.txt*.

(Note that if using university workstations, you may not have writing privileges to all folders on the machine.)

13. After providing the requested information and pressing OK, the script should measure the durations and save them. If you find the result file where it should be, try to open it in a text editor or Word, for example, to see the contents.

Backups

14. When collecting speech material, you must be prepared for the worst: something may happen to the computer's hard disk, or you might lose your USB memory stick. Accidents happen! However, you will probably end up blaming yourself if your entire collection of speech is destroyed.

What is required from good backups? What means do you have at your disposal for making and storing backups, e.g., of speech recordings? Can you use them? How often do you make backups? Note that speech itself can be regarded as personal data and should be handled accordingly.