

A brief introduction to Praat

Shigeto Kawahara

1 Before we begin

This is a handout that I use for my introductory phonetics classes (graduate and undergraduate). It will be revised as I get feedback from my students. Your feedback will of course be appreciated. (This handout was fairly substantially revised in June 2014, following some major changes in Praat.)

1.1 Why use Praat?

- It is free (downloadable from <http://www.fon.hum.uva.nl/praat/>).
- It allows us to perform multiple operations (e.g. acoustic editing, acoustic measurements, creating pictures, creating sounds, etc).
- It comes with a useful scripting function.
- Many people use it—so it's easy to get help, and many scripts are available online; I am happy to share mine (some are available from my website).

1.2 To begin

- Go to <http://www.fon.hum.uva.nl/praat/> and download Praat, a version that is appropriate for your OS.
- Get some speech sample (available on Sakai).

2 Very basics

2.1 Opening sound files

To open sound files:

1. Open Praat.
2. Open → Read from file...
3. Find and choose the sound file you want to use.

4. You can also click on a file and open the file by right-clicking (this would be easier if you want to open multiple files).
5. If you like, you can associate .wav files with Praat, so that doubling clicking .wav files will automatically open the files in Praat.

If you want to open all sound files in a folder, there is a nice script. Download the script at:

http://www.helsinki.fi/~lennes/praat-scripts/public/open_all_files_in_folder.praat.

I usually save my scripts with .praat extension, and associate .praat files with Praat. That way, you can open a script by double-clicking.

2.2 Saving sound files

Once you open a sound, it will appear in **the object window**. You can edit the sound file in various ways, but **new sounds files are not saved automatically**. To save sound files in the object window, `Save` → `Save as WAV file...`

2.3 Closing

To close a sound file, click the remove button on the bottom left. **REMEMBER TO SAVE IT** before you close if you created a new sound. Always remember, just because a sound is in the object window, it does not mean that the sound is automatically saved. It isn't.

2.4 Inspecting sound files

To see the phonetic properties of a sound in the object window, highlight the sound you're interested in and click `View & Edit`. You see a waveform on top and a spectrogram in the bottom panel. Praat has many functions in this editing mode.

2.4.1 What's displayed

First, we realize that Praat plots various acoustic properties. Red dotted lines on the spectrogram display represent formants, the blue line the pitch contour, the yellow line the intensity contour, and the blue vertical lines on the waveform display pulses. You do not always need to show all of these—Go to the `view` pull-down menu, and select `show analysis` to select the properties that you want to be displayed.

2.4.2 Selection

- If you want to select a part of the sound file, click the beginning and click the ending. You see the duration of that interval in the bar at the bottom (in seconds).
- You can play the selection by clicking a short bar at the bottom.
- If you want to zoom to selection, `Command+N`, or click on the rightmost button in the bottom left corner.

- You can zoom in and zoom out by using `Command+I` and `Command+O`.
- You can zoom to all with `Command+A`.
- If you want to extract the selected portion, go to the `Select` pull-down menu and `Move start of selection to nearest zero crossing`¹ and also `Move end of selection to nearest crossing`. Otherwise, the extract sounds will have clip-pings!
- Then from the `File` pull-down menu, you can `Extract selected sound...` This will put the selected sound in the object window (you should remember to save it).
- Or alternatively, you can directly do `File` and `Save selected sound to WAV file...`
- By going to `Select` and `Select...` you can specify the beginning time and end time. This is very useful if you want to extract some stretch of portions with some exact duration (for example, if you select from 0 to 1 second, then the resulting sound is going to be exactly 1 second long).²
- You can also copy and paste the selected portion as you would in other softwares.

2.4.3 Spectrum

Pull down the `Spectrum` tab, and go to `Spectrogram settings`.

- `View range` stands for the lower and upper bounds of the spectrogram analyses. The default is 5,000, which is fine unless you're interested in fricatives. If you're analyzing fricatives, change the upper bound to 8,000.
- `Window length controls...` well the window length for spectral analysis (in terms of frequency).
- Confusingly, if you set the window length high, say 0.05 (from the default setting of 0.005), then you get so called “narrow-band” spectrograms. This is because a “narrow-wide” distinction refers to a frequency band for the spectral analysis. The longer the window length, the narrower the frequency band.
- A “wide-band spectrogram”, which is the default for Praat, has nice time resolution, but it does not show harmonics well. A “narrow-band spectrogram” on the other hand shows harmonics really well, but it does not track changes in formant frequencies well.
- A narrow-band spectrogram does not capture spectral changes very well, so it is not good at catching short acoustic events, like bursts. I thus myself use a wide-band spectrogram unless I want to see individual harmonics.

¹Zero-crossing is the point where the sound pressure becomes zero.

²Remember that if you want to compare two or more spectrograms in your paper, you should always use the same time scale. This technique is very useful for this purpose.

- `Dynamic range` determines energy thresholds that Praat picks out for its spectral analysis. It represents a range within which a sound is considered to be present. Therefore, lower values would make everything "thinner", but you may miss important acoustic signals. If you set dynamic range too high, however, the spectrogram may get too dark to see.

Another function that is helpful is `View spectral slice`.

- Spectral slices show spectral properties of a sound at a certain point (it is timeless).
- Spectral slices are useful when you want to see the formant structure of a vowel at some specific point or when you want analyze fricatives.
- Spectral slices are used when we calculate spectral tilt as well.
- The default range is the Nyquist frequency of your sound files; e.g. if the sampling rate is 44.1k, the range is about 22k.
- Therefore you may `Zoom` from the `View` pull-down menu.

2.4.4 Pitch

You can show pitch and measure it. Depending on whether you are analyzing male or female speakers, you may want to set the pitch range in `Pitch settings` in the `Pitch` menu. You can also get a pitch at a particular point by using the `getting pitch` function.

2.4.5 Formant

In analyzing formants, go to `Formant settings`.

- `Maximum formant`: The standard recommendation is to set it to 5000Hz for male speech and 5500Hz for female speech.
- `Keep Number of formants to 5`.
- These are, of course, rules of thumbs, and you should be flexible (and explicit about it in your write-up).
- `Window length`: If you're analyzing formants, use a wide band spectrogram.
- `Dymanic range`: See above.
- You can also get formant measurements by using `Get F1, F2`

3 Annotation

Go back to the object window, select the sound file, and go `Annotate` in the menu on your right.

1. `Annotate` → `To TextGrid`....
2. Decide how many tiers you need: if you're just putting segmental boundaries, one should suffice. If you want to include segmental and tonal information, you may want to have two. So decide how many kinds of info you want to put in.
3. Delete "Mary John Bell" unless you like those names. I usually put in "segment" for the segmental tier and "tone" for the tonal tier.
4. Point tiers annotate points. Non-point tiers (or interval tiers) annotate intervals. I usually use non-point tiers for segmental transcriptions, and point-tiers for tonal transcriptions.
5. Press `Apply` and go back to the object window. Select BOTH the sound file and the TextGrid file. Hit `Edit`
6. Now you have a tier (or two) for transcription at the end. Use circles at the top of a gray line to put a boundary.
7. Add and edit boundaries as you like. To delete a boundary, use `Option+Delete` (in windows, it is `Alt+Delete`).
8. DO NOT FORGET TO SAVE THE TEXTGRID FILE!!

4 Making (nice) pictures

4.1 Basics

- Find the picture window. It is possible that you accidentally closed it. Reopen Praat if that's the case.
- Decide how big the picture should be. The white space inside the blue bars will determine the size of the picture. The default selection can be very small.
- Let's first choose a sound file, and click `Draw` and hit `OK`. You should see the picture of the waveform of the sound in the selected portion. So, Praat draws a waveform based on a sound. What if you want to make pictures of other properties?

4.2 Drawing pitch & spectrograms

- If you want pitch or spectrograms, you first have to create them.
- Select a sound in the object window and go to `Analyze` → `Periodicity` → `To pitch`... to generate pitch.
- To generate a spectrogram, `Analyze` → `Spectrum` → `To spectrogram`....

- If you want to analyze or illustrate fricatives, you want to set `Maximum frequency` to a higher value, say, 8,000 Hz.
- Now you can select the generated pitch or spectrogram and `Draw` (and `Paint` for spectrograms).
- When drawing a pitch line, the line may be too thin. Go to the `Pen` pull-down menu, and choose appropriate `Line width`.
- When drawing a spectrogram, you may readjust `dynamic range` to determine the “thickness” of the picture.
- To draw a pitch track or a spectrogram, you can select a new range or superimpose on something else.
- You can add other info in the picture. One thing that I often do is to put dotted lines with a 1000Hz interval. To do this, go to `Margins` → `Marks` → `Marks left` and select $(F/1000) + 1$ (where F is your frequency range) for `Number of marks`.
- `Command/Control+z` if you want to undo.

4.3 Drawing transcription

1. Once you have a picture of a waveform, pitch and spectrogram, you may want to superimpose the segmental boundaries.
2. For this, you need a `TextGrid` file. See above for how to make it.
3. Select the `TextGrid` file and `Draw`. I usually superimpose segmental boundaries on the pictures of a spectrogram and a waveform.

4.4 Saving

I usually save the picture files as `eps`, because I use `LATEX`. If you’re using a Microsoft word processor, you cannot copy and paste (well, you can, but the picture wouldn’t look nice). Save the picture file as `eps` first. You can later convert it to `jpeg` or other formats you like, or you can insert an `eps` file to your document from the `insert` pull-down menu. These days, at least on mac, it seems that you can save the files directly as `PDF` files.

4.5 Screen shooting?

Screen shooting should be the last resort, since Praat has such a nice picture-generating function. But I have to admit that my first published paper contains Praat pictures that I took with screen shooting. But try to avoid screenshooting if possible,

5 Recording and creating sounds

New → Record mono sound and press record. When you're done, press Stop and Save to list. Remember again that saving to list does not mean saving to a file. To save it, click the sound in the object menu, and from the Write pull-down menu, save it as a wave file.

5.1 Creating sine waves

New → Sound → Create sound as pure tone. You can specify various parameters; starting time, ending time (duration=ending time-starting time), sampling frequency, its f0, amplitude (Pa, not in dB!), rise time, etc.

5.2 Creating sine wave harmonics complex

New → Sound → Create sound as tone complex. You can play with various harmonics complexes. Isn't this fun?

5.3 Creating silence

New → Sound → Create sound from formula, change the name to "silence", specify the start time and end time (i.e. the duration), put "0" in the formula.

6 Other useful stuff

6.1 Modifying intensity

You can modify the intensity of sound files. Select a sound file(s) and click Modify → Scale peak if you want to adjust peak intensity (I usually avoid using the default .99 because it may result in clipping, try .95). If you want to adjust average intensity in dB, click Modify → Scale intensity.

6.2 Manipulating duration

Select a sound, click Convert and Lengthen.

7 Resources

- An online intro material by Sidney Wood:
<http://person.sol.lu.se/SidneyWood/praat/frames.html>
- Praat script resources:
<http://www.linguistics.ucla.edu/faciliti/facilities/acoustic/praat.html>

8 Exercise

- (1) Exercise
 - a. Record a few words of your choice.
 - b. Annotate them
 - c. Draw pictures consisting of their waveforms, spectrograms and your annotation.
- (2) Exercise II
 - a. Record all the cardinal vowels in your native language in some frame, say [h_t].
 - b. Measure F1 and F2.
 - c. Create a vowel chart (a chart template will be provided for you).
- (3) Exercise III
 - a. Record morphologically related pairs like *atom* and *atomic*, *write* and *writer*, and *pho-
netic* and *phonetician* a few times.
 - b. Compare the durations of [t] and [r].