

Praat Tutorial and Resources

Jean-Philippe Goldman

February 2004

The purpose of this document is to give an overview of Praat, its features, its on-line resources and some quick steps to start using it.

What is Praat?

Praat is an open-software tool for speech signal edition and labeling, as well as for various acoustic analyses (spectral, formant, pitch, intensity) and prosodic manipulation. It has also some abilities to plot figures, build Optimality-Theory grammars, make articulatory synthesis, simulate feed-forward neural networks and make discriminative statistics. Paul Boersma and David Weenink from the Institute of Phonetic Sciences, University of Amsterdam (www.fon.hum.uva.nl), have been developing Praat for almost 10 years and are still very active (frequent software updates and quick answers on praat-users group). It's free and available for most platforms. Many tutorials exist and the discussion group is very active (see resources below).

It has been designed for both beginners (easy interface, many default options to learn by trying, searchable manual, various possibilities of analysis, manipulation and labeling) and expert users (logging of every action for easy batch scripting, or resource for advanced application development).

Quick Tutorial

Praat's layout is based on 'objects'. Any object (*sound*, *pitch* curve, *matrix*, group of *strings*,...) can be loaded, created, saved, queried, modified or used to create another object (of the same type or not).

General information about this tutorial:

This tutorial will teach you how to create or load sounds, to make first analyses and to label them. Praat needs no right-click button nor double-click and has no contextual menus. Follow the instructions in red.



Figure 1 Praat icon

Start Praat

Praat starts with two windows: *Praat objects*, and *Praat picture* (you can close the second one as we will not draw pictures here). The "Praat objects" window is the main one. It lists the objects in memory. Menus and buttons are dynamic, i.e.

they may change (appear, disappear or be disabled) according to the selected objects. When it's just started, Praat has no objects in its list, therefore the list is empty, the buttons (bottom) are yet disabled, and the 'Write' menu (4th) is empty.

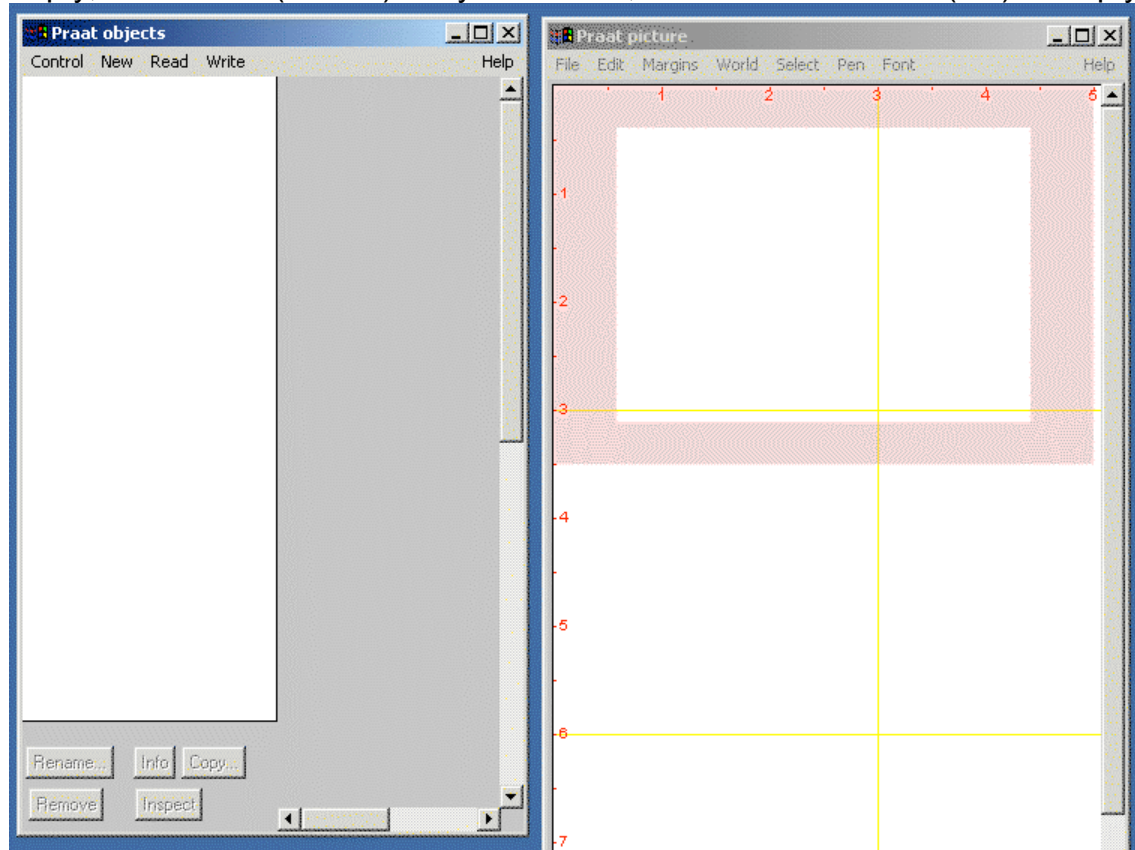


Figure 2 Praat windows at start

Create a Sound

Execute **Menu 'New' > Sound > Create Sound from Shepard tone...> OK**

The sound you've just created is now in the list of objects (type + name). As it is selected (grey), many buttons appear on the right, the buttons at bottom are now enabled, and the **Write** menu displays options.

In fact the 5 buttons at the bottom are called the *fixed buttons* (they are always present though they might be disabled if nothing is selected or present in the list). These functions are common too all objects whatever the type. With them, you may:

Rename an object (**Rename...**)

Duplicate an object (**Copy...**)

Delete an object (**Remove**)

Get the results from various queries on an object (**Info**)

Browse the internal data of an object (**Inspect**)

The buttons on the right of the list are the *dynamic buttons*. We'll go through most of them. The first one is always a Help button referring to the type of the selected

object (it won't appear if several objects of different types are selected at the same time).

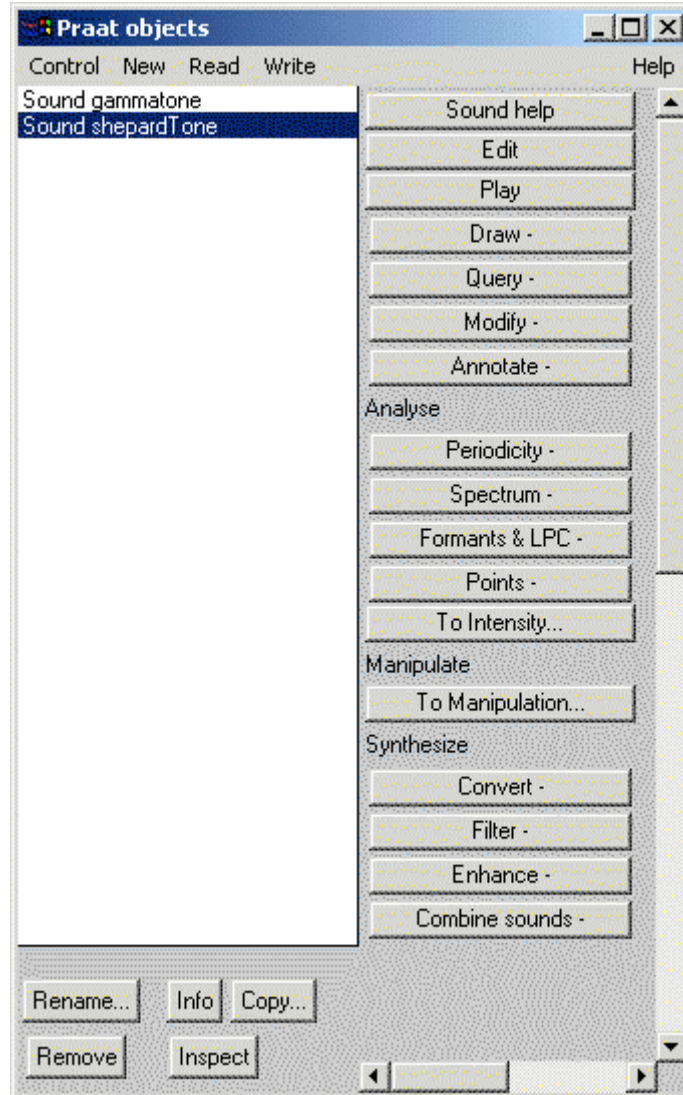


Figure 3 "Praat objects" window with "Sound" objects

Record a Sound

With a microphone, you can record a sound or some speech with **Menu > New > Record Mono Sound...** A window 'SoundRecorder' opens (see below). Plug your microphone. Press **Record** to start recording. You should see a real time meter in green/yellow/red while you speak. If not, check your general recording settings. Press **Stop** to stop recording. **Play** to listen to your recording. You can start over if you are not satisfied. When finished, press **Save to list**: (change the name 'sound' to something else if you like) then **Close**. The sound appears now in the object list.

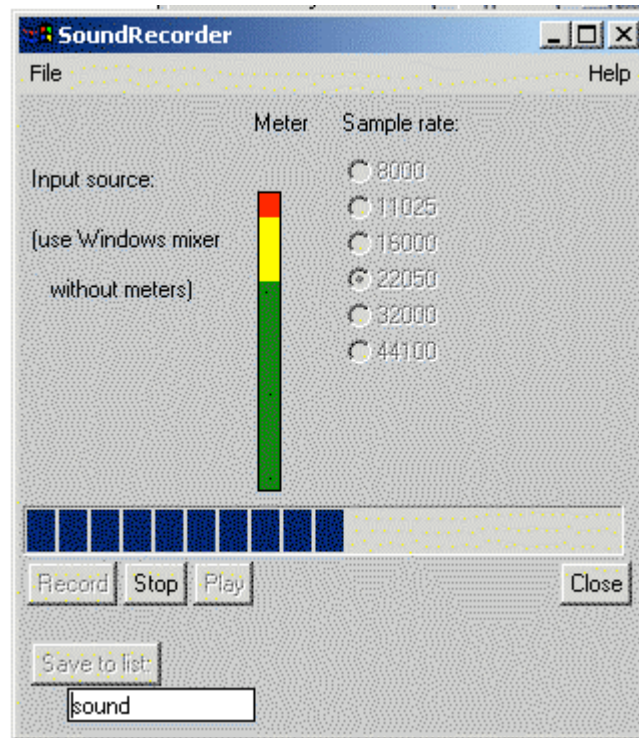


Figure 4 SoundRecorder window

Load a Sound

From **Read menu > Read from file...**, Praat reads whatever it can recognize. For sound files, WAV, AIFF, non-compressed AIFC, NeXT/Sun (.au) and NIST files are supported.

See www.praat.org/manual/Sound_files_3__Files_that_Praat_can_read.html

Raw (headerless) files can also be read with other options in this **Read** menu.

Sound Object

Once you have a Sound object in your list (created, recorded or loaded), select it by clicking on it if it's not selected (see figure 3). Here is a quick description of the dynamic buttons:

Sound Help opens the Help window with many useful information

Edit opens an Editor window with signal and various acoustic analyses. We'll describe it in the next section

Play plays the sound. Press the **Esc** key to stop playing.

Draw- opens a local menu with options. It is useful to plot picture in the window we've closed. You can give a try.

Query- opens a local menu with many options. Try some of them, like **Get duration...** An Info window will open to display results. Try also **Get nearest zero crossing...** A form window asks for a time value. In this form, notice the **Help** button referring to the right page in the Manual. Try **Info** fixed button (at bottom); the info window will give you all the possible queries on the whole file.

Modify- allows to **Reverse** the sound file (try it and **Play**). Notice that

modifications are done with the object. Reverse it once more to get back to original sound.

Modify- > Formula. See Formulas tutorial for more help (from Help menu, or Help button once you clicked **Formula...** or see www.praat.org/manual/Formulas.html).

Annotate- > To TextGrid... create a labeling object (see [TextGrid object](#) section)

The *Analyze* section or submenu is a group of dynamic buttons that create new objects of different types from the sound object after one acoustic analysis among various explicit possibilities:

Periodicity- deals with pitch contour: Pitch analysis, Harmonicity, Glottal pulses

Spectrum- deals with all kind of spectral analyses (spectrum, spectrogram, wavelets)

Formant&LPC- creates formant contours and (linear predictive) filter coefficients

To Intensity calculates to intensity contour

When you try these tools, notice that the dynamic buttons change. Click the **Edit** dynamic button to see results. Reselect the original Sound object to try another one.

The *Manipulation* section contains a unique button:

To Manipulation: creates a Manipulation object, which contains the original sound, a pitch tier, a duration tier and a PointProcess (i.e. a sequence of points in time) representing the glottal pulse.

Most of the buttons in the *Synthesize* section are self-explicit and create Sound objects from the original one(s). You'll find here tools to extract a part, to resample, to filter (see Filter tutorial in Help Menu), or to concatenate sounds.

Edit Window

As a Sound object is selected, press **Edit**, a new window opens. In the top half, see the speech signal in black as well as the glottal *pulses*. The bottom half shows the analyses: *pitch* in blue, *intensity* in yellow (or green), *formants* in red and *spectrogram* in background. Each of the five analyses (these four and the pulses) has its own menu, all starting with a 'Show/Hide' switch and a 'Settings...' option.

The three bars below the analysis pane are 'play' buttons. You can click them to play the sound and see a real time cursor (vertical bar moving from left to right) on the signal. The top *cursor bar* is divided in two equal sections as the *main cursor* or *marker* is set by default right in the middle of the file when Edit window opens. Figures indicate durations in seconds. The second bar (*window bar*) represents the part of the file that is shown in the window. The third is the entire file (*full bar*). The last two are now the same as Edit window shows all the file when it opens.

Select a part in the middle of the signal by a click-and-drag on signal or analysis pane. The selection appears in transparent pink. Now the top bar should be

divided in three sections: before the selection, the selection itself, after the selection. Play them individually. You can expand or shrink this selection by pressing-and-holding the shift button before you click the signal pane or make another selection.

Now look at the bottom left corner, you see four handy buttons to help you modify the window settings

all shows the entire file

in let you zoom in by a factor of two (i.e. you'll see half of the current window)

out zooms out by a factor of two

sel adjusts the display window to the selection

Try them. As the window displays a portion of the signal, notice how the *window bar* and the *full bar* change. You can still play what is before the window, after the window and the complete sound.

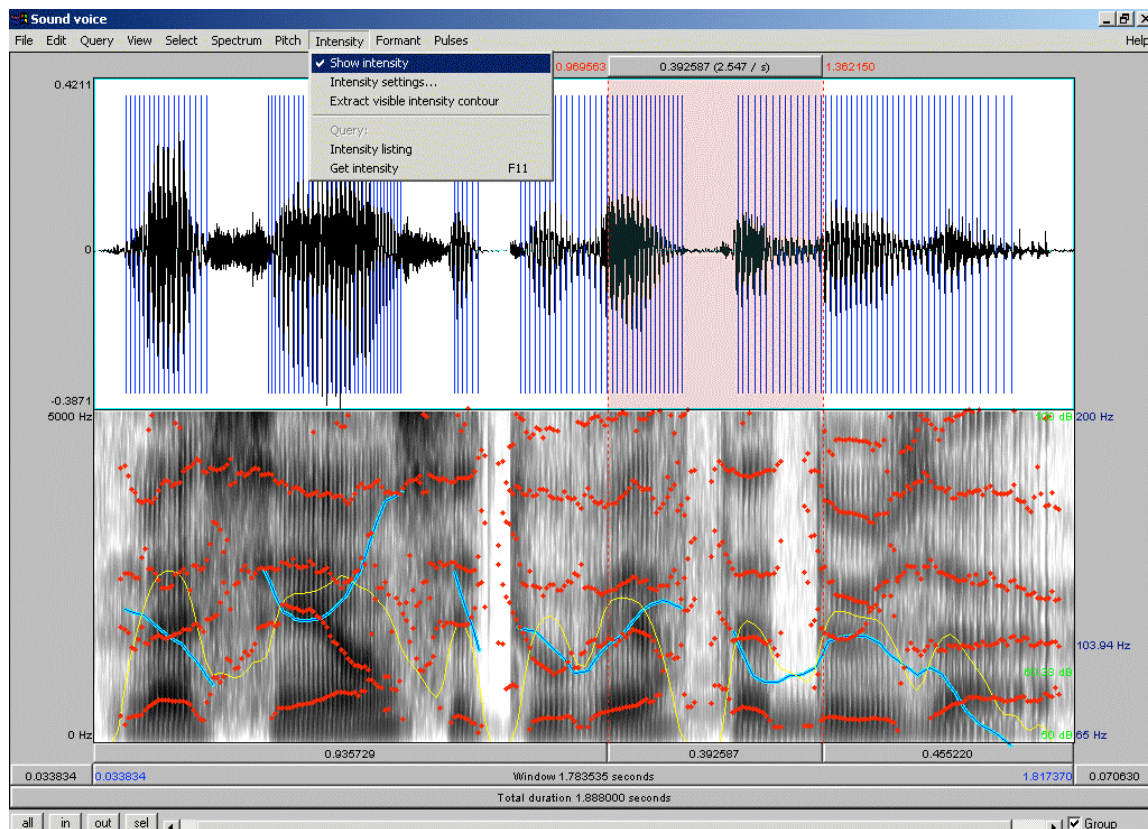


Figure 5 Sound Edit window

As you click on the signal, see in blue, on the left of the sound, the value of the waveform. If you click on the analysis pane, you may have up to three vertical scales depending on the number of shown analyses:

frequency on the left. It shows the value in Hz of the marker in the frequency domain (as long as spectrum or formant analysis is shown)

pitch in blue on the right. marker value as well as minimum and maximum. Note that the marker value is the pitch estimation of the sound at that moment, i.e. on the blue curve, not at the marker position which can be anywhere in the pane.
intensity in green. It might be either side, or even inside the analysis pane.

Here is briefly what you can do with the other menus:

File: extracts parts of the signal to the object list or to a file

Edit: the usual Undo/Cut/Copy/Paste menu for signal edition. Reverses and sets to silence options too.

Query: is used to get cursor position or log your measures into text files

View: zooming and playing

Select: moves the 3 cursors (which are the main cursor, the beginning and the end of the selection)

TextGrid Object (Labeling)

A **TextGrid** is an annotation/labeling object with one or more layers or *tiers*. Each tier can be either an *interval tier* (contiguous labeled intervals) or a *point tier* (labeled points).

Create a TextGrid object with the **New menu > Create TextGrid...** or, better, select the sound you want to annotate, then use the dynamic buttons: **Annotate-> To TextGrid...** In the latter case, the TextGrid time range will be adjusted from the Sound.

A form window opens and asks for the tiers names (be sure to write all the tiers names) and, among these, which are the point tiers (the others are interval tiers by default)

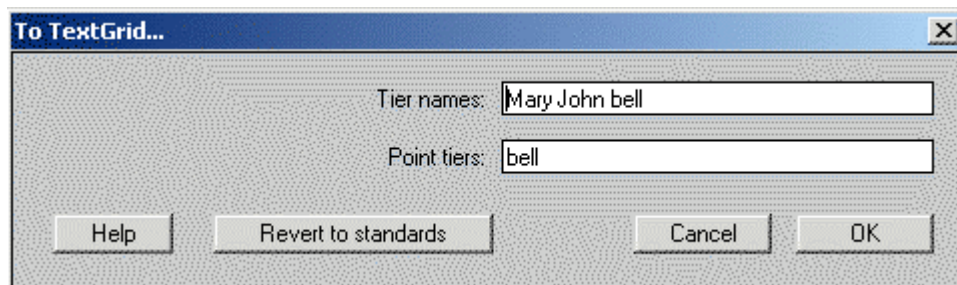


Figure 6 Defining the tier names

Then, select BOTH Sound and TextGrid objects (with the help of shift button), and press **Edit** (dynamic button)

The TextGrid window looks like the previous Edit window with a third pane (below analysis) and three additional menus (Interval, Boundary, and Tier). The new pane shows the tier(s) for annotation. If several tiers are shown, a red hand

on the left tells which one is selected. You can zoom and select the sound as in the editor window but not edit the sound (cut/copy/paste).

To add a boundary, click on the signal (or analysis) pane to set the main cursor wherever you want. Notice the vertical red dashed line is doubled with a grey thick line in each tier with a *circle* at the top of it. Then, three possibilities exist:

- **Enter**, and a new boundary will be created in the selected tier
- **Ctrl+F1** to add on the first (top) tier, **Ctrl+F2** to add on the second, and so on... **Ctrl+F9** to all on all tiers
- Click on the circle of each tier you want to add a boundary to

Then the boundary becomes red. It is selected. Unselected boundaries appear in blue.

To insert text, simply select the boundary or interval to annotate and type in some text. You can do it just after the boundary creation. The label will appear at the boundary (if point tier) or after the boundary (if interval tier) in the text editing pane, just below the menus. Note that the Cut/Copy/Paste options in the Edit menu are meant to edit this text (not the signal). You can edit the text as long as the boundary/interval is active.

To remove a boundary, make sure you select the right boundary, then Menu **Menu Boundary>Remove** or **Alt+Backspace**.

To move a boundary, select it and drag it.

You may want to hide the analysis (by unchecking 'Show' in each analysis menu, or with the **View > Show Analyses...** option) for a larger view of the signal and the tiers, although the analysis may be helpful for an accurate annotation.

In the following figure, see the waveform and two interval tiers (phones, words) and one point tier.

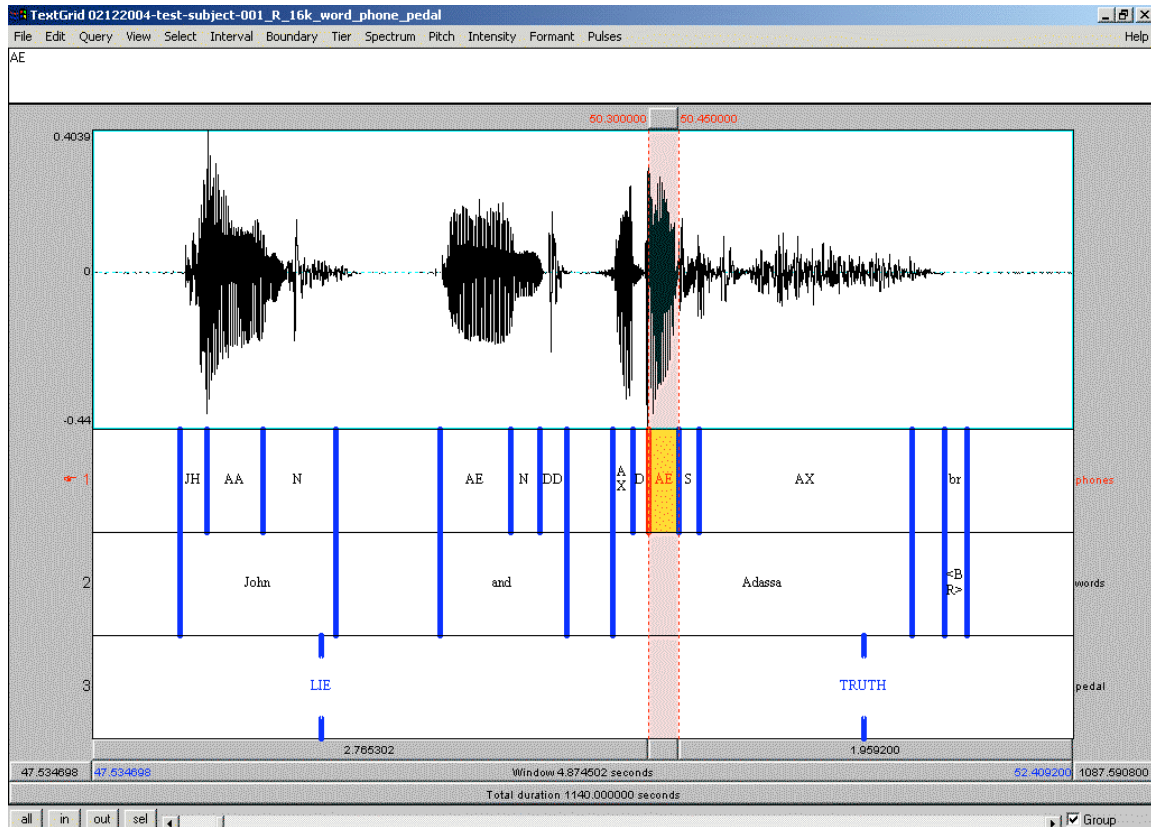


Figure 7 Annotation window

Resources on the Web

Download: from www.praat.org for Macintosh, Windows, Linux, FreeBSD, SGI, Solaris, and HP-UX. These archives contain a unique executable file named 'praat'. No installation or compilation is required. You can launch it immediately. The source code is also available for expert developers.

Discussion group: groups.yahoo.com/group/praat-users

Help: www.praat.org/manual or Help menus and buttons within Praat

Tutorials in order of (my) preferences (best first):

- www.ling.lu.se/persons/Sidney/praat/frames.html
Basic good tutorial by Sidney Wood, Lund University. Website structure with frames- 10 main sections - glossary - quick guides - nice images - more to come
- ots.utoronto.ca/users/vanlieshout/praat-manual.pdf
Also good Praat tutorial by Pascal van Lieshout, U. of Toronto (Jan 2002) 26 pages - basic descriptions of windows, objects, spectrographic analysis, good tips for recordings
- ling.ohio-state.edu/~welby/TUTORIALS/welby-ito-praat-tutorial.pdf
9 pages - jan 2002 - basic

- www.ling.ed.ac.uk/facilities/help/praat/guide.shtml
1page - basic button-click summary
- www.unc.edu/~jlsmith/ling120/pdf/hdt2-4.pdf Handout 2.4, "Using Praat for speech analysis," from the Ling 120 course pack
9 pages- Aug 2003 - no image but a lot of info

Advanced tutorials

- www.ling.rochester.edu/people/cross/intro-to-praat.pdf
Introduction to advanced functions, by Katherine Crosswhite, U. of Rochester - 8 pages - December 2002
- www.praat.org/manual/Scripting.html
- www.praat.org/manual/Formulas.html

Scripts collections:

- www.ling.rochester.edu/people/cross/scripts.html
- uk.groups.yahoo.com/group/praat-users/files/
- www.ling.lu.se/persons/JohanF/PRAATSCRIPTS/
- www.helsinki.fi/~lennes/praat-scripts/
- www.cphling.dk/pers/johtnd/praat/my_praat.htm