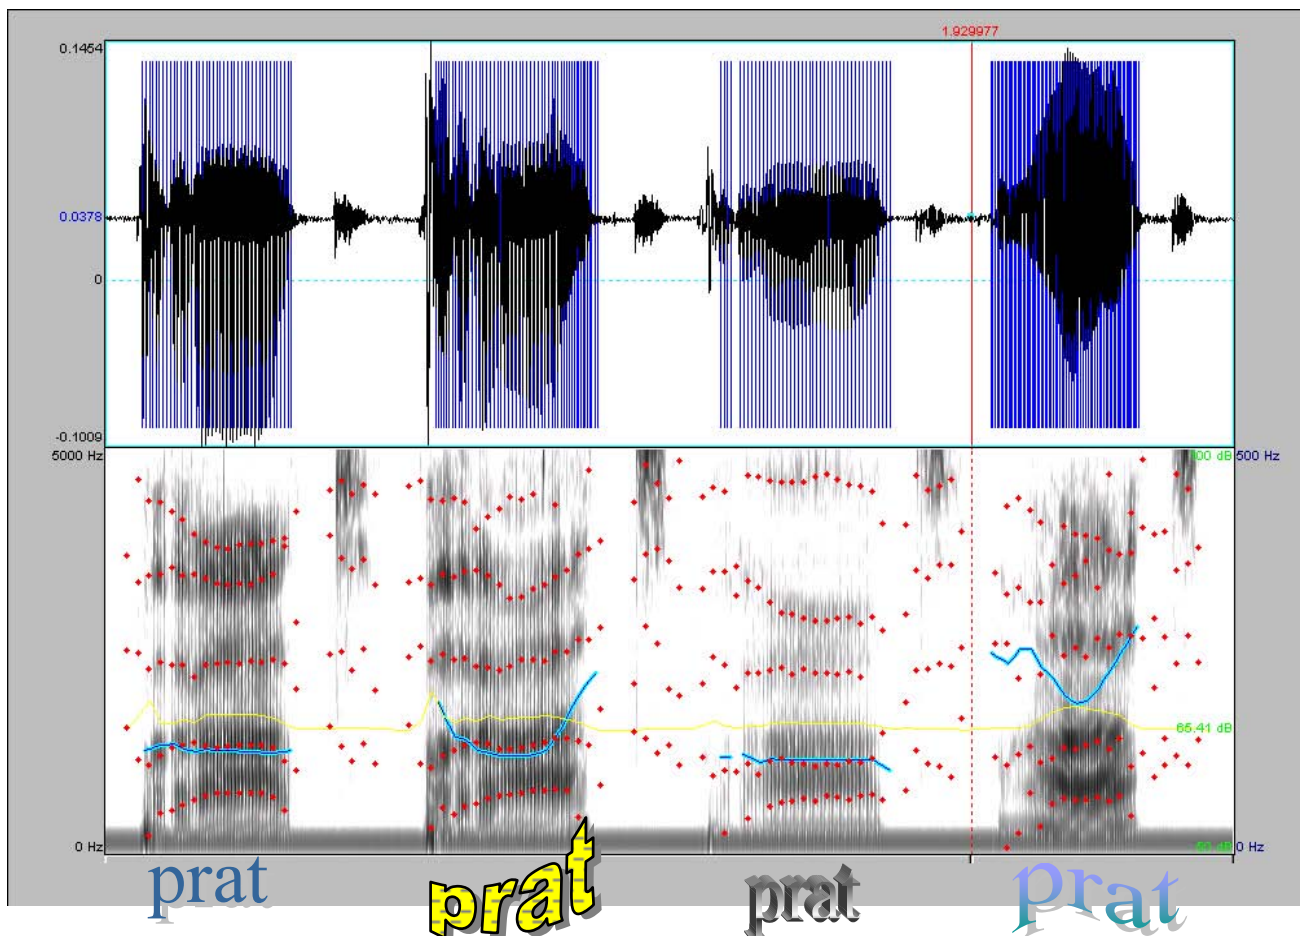# PRAAT
Short Tutorial

A basic introduction

Pascal van Lieshout, Ph.D.
University of Toronto, Graduate Department of Speech-Language Pathology, Faculty of Medicine, Oral Dynamics Lab

V. 4.2.1, October 7, 2003 (PRAAT 4.1.x)

PRAAT VARIATIONS

PRAAT[1]
Short Tutorial

Pascal van Lieshout, Ph.D.
University of Toronto, Graduate Department of Speech-Language Pathology, Faculty of
Medicine, Oral Dynamics Lab (ODL)

A. <u>Introduction</u>

This tutorial provides an introduction to some of the basic procedures in the program
PRAAT. This is a freeware program for the analysis and reconstruction of acoustic
speech signals. The software can be downloaded from the following website:
http://www.fon.hum.uva.nl/praat/. This page also includes a PRAAT beginner's manual
written by Sidney Wood. PRAAT can be used on different operating systems (see
PRAAT website for more information), but this tutorial is based on Windows 2000 OS.

PRAAT is a very flexible tool to do speech analysis. It offers a wide range of standard
and non-standard procedures, including spectrographic analysis, articulatory synthesis,
and neural networks. This tutorial specifically targets clinicians in the field of
communication disorders who want to learn more about the use of PRAAT as part of an
acoustic evaluation of speech and voice samples. The following topics will be covered:

1. Finding information in the Manual
2. Create a speech object
3. Process a signal
4. Label a waveform
5. General analysis (waveform, intensity, sonogram, pitch, duration)
6. Spectrographic analysis
7. Intensity analysis
8. Pitch analysis
9. Using Long Sound files

In this tutorial, readers are assumed to be already somewhat familiar with standard
Windows 9x/NT/2000 tools, like opening and closing windows, making windows smaller
or bigger etc. If not, please ask your colleagues who are more familiar with PC Windows
OS for help or check the following website (http://search.support.microsoft.com) for
more information on Microsoft products. If the instruction mentions the word 'click
(ing)', it simply means that you have to position your mouse cursor on top of the
indicated location and press the *Left-mouse button*. If another mouse-button is required,
this is explicitly mentioned in the text. As usual in Windows, once you made your choice
for a particular window, you confirm by clicking the *'OK' button*. If you want to go back
or cancel, just click the *'Cancel' button*. The main menu options are located on the right-

---

[1] PRAAT (a system for doing phonetics) was developed by Paul Boersma & David Weenink at the
Phonetic Sciences department at the University of Amsterdam.

hand side of the 'Praat Objects' window, but the contents of the menu may change depending on the type of (sound) object you have selected.

Please notice that earlier versions of PRAAT may differ in layout and function from the version described here (PRAAT v. 4.1). Comments and suggestions on this tutorial are welcome; just send an e-mail to: p.vanlieshout@utoronto.ca
PRAAT users can also join a discussion list, which provides a useful forum for asking questions as well as a database for sample scripts. More information can be found here: http://uk.groups.yahoo.com/group/praat-users/

Toronto, May 29, 2003

PvL©

PRAAT
doing phonetics by computer

version 4.1

www.praat.org

B. <u>Working with PRAAT</u>

1. **Finding information in the Manual**

If you open the program[2], the following two windows will appear:
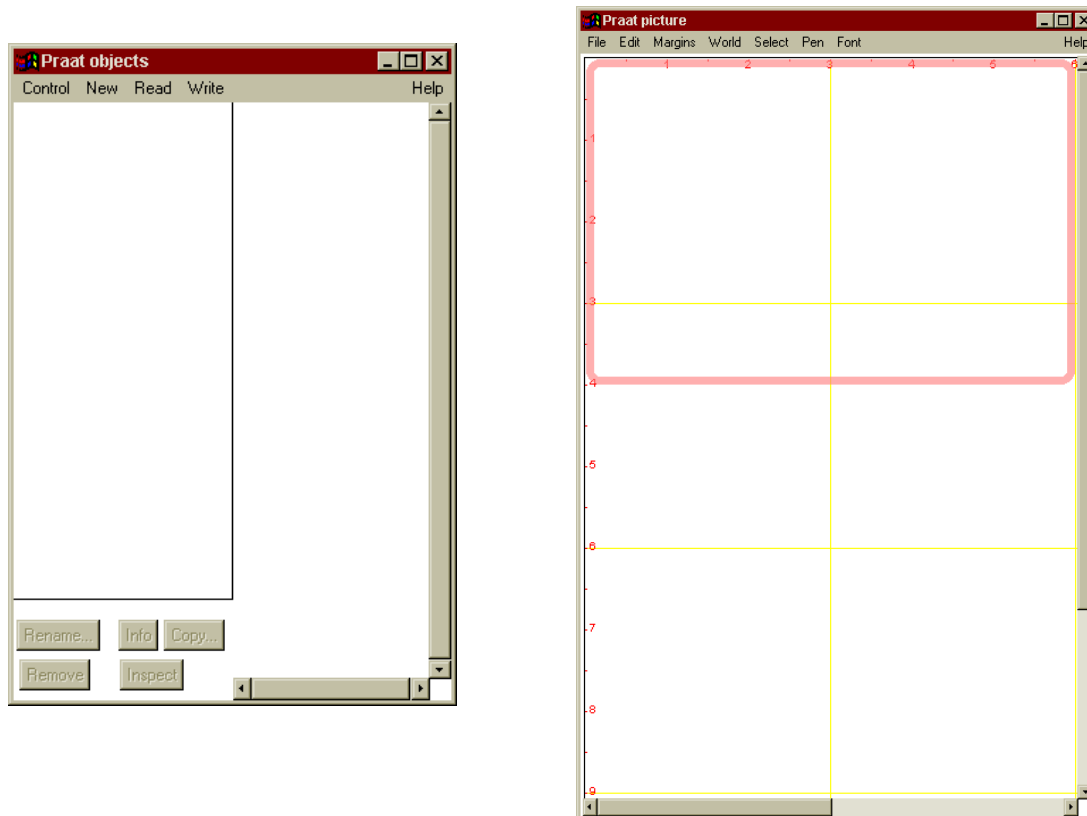
The window to the left is the '**Praat objects**' window. On the left-hand side you will
normally see a listing of your speech files ('objects' in PRAAT language) which can
either be created from scratch (see section 2, #1-15) or read from a file (section 2, #17).
The window on the right is the '**Praat picture**' window and is used for plotting graphs.
These can be saved in various formats, including an EPS postscript[3] or a Windows
Metafile for later word processing purposes or they can be printed directly using "print"
(CTRL-P) in the file menu.

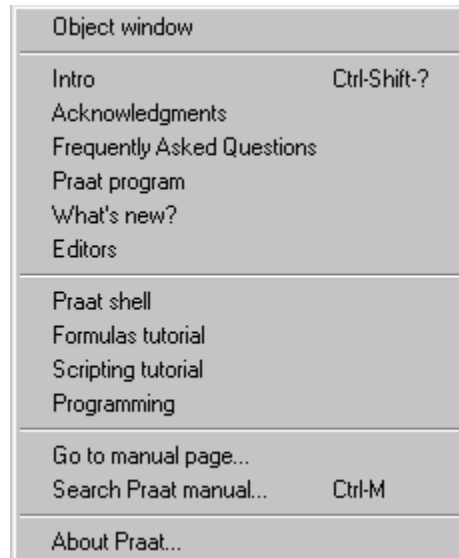Information about the program and all its procedures can be found in the PRAAT manual
by simply clicking on the Help button in the main menu of the PRAAT objects window.
If you do that (try it now), you will find the following options available to you:

---

[2] If not already present, create a shortcut to the "praat.exe" file on your desktop for easy access.

[3] Postscript files can be read and printed using the program Ghostview® (gsview32.exe; version 2.5 or
higher). This is a free downloadable program (see homepage http://www.cs.wisc.edu/~ghost/)

```
Object window

Intro                          Ctrl-Shift-?
Acknowledgments
Frequently Asked Questions
Praat program
What's new?
Editors

Praat shell
Formulas tutorial
Scripting tutorial
Programming

Go to manual page...
Search Praat manual...          Ctrl-M

About Praat...
```

Most options speak for themselves and you can try them out for yourself. The tutorials are useful in that they provide more information about how to deal with specific topics in PRAAT. For those who want to use scripts in PRAAT to automate certain procedures, the 'Scripting tutorial' is highly recommended. More information about the use of formulas, operators, functions etc. can be found in the 'Formulas tutorial'. Check the 'Frequently Asked Questions' section with answers to common issues raised by users and make sure you are aware of recent changes to the program listed in the 'What's new?' section. The option that will be used most often by the majority of users is the 'Search Praat manual' (also notice that some functions have *short-cut keys*, in this case Ctrl-M). Click on the option (try it now), and the following window will appear

```
Search for strings (separate with spaces):

[                                                    ]

[ Help ]   [ Revert to defaults ]     [ Cancel ]   [ OK ]
```
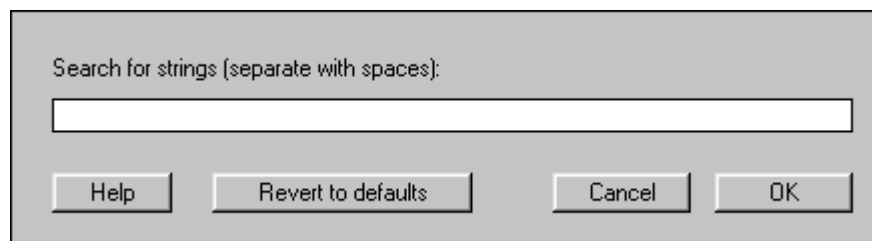
Simply type a search string in the empty space of the window, and you will find the information that is available on your search topic. For example, find information on the following topics:
-   formant
-   pitch
-   intensity
-   spectrogram
-   printing

As you will notice, some queries will give you a lot of options, others are more restricted.

Remember that you can always invoke this Help function from anywhere in the program, and that most procedures allow you to invoke specific help information directly from their window menus (see for example the Search Manual window above).

2. **Create a speech object**

Before analyzing speech samples, it is important to adjust your sound card options properly. To access these options, you have to open the "Volume control" window. For Windows 2000 the following steps apply[4]:

1. Go to 'Start' of the Windows Task bar (left lower corner)
2. Go to 'programs' -> 'accessories' -> 'entertainment' -> select 'volume control'
3. This will bring up the following (or a similar) window (depending on the sound card that is installed on your PC)
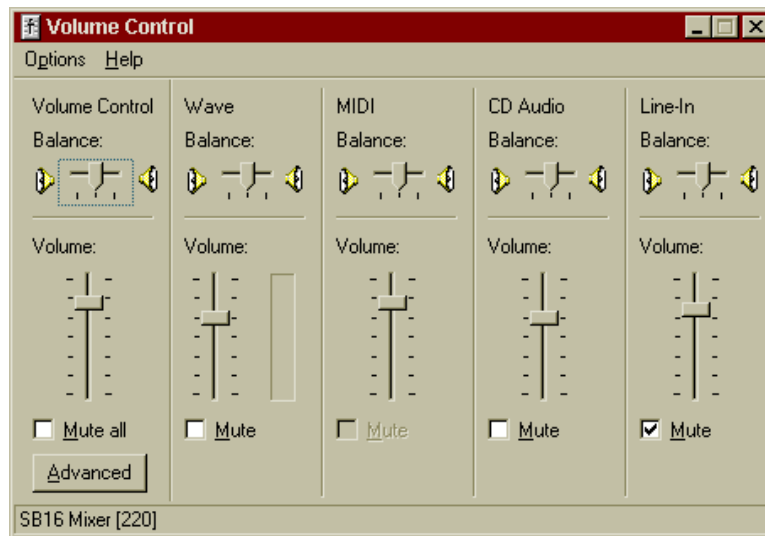
4. Go to 'Options' -> 'properties' -> select 'recording'
5. Now you will see a number of options (including Line-In & Microphone)
6. Select 'microphone' by clicking on 'select' button (White Square below volume meter) and deselect all other options. Adjust 'Volume meter' if necessary to about halfway the scale. You can leave this window open and put it on the Task bar by clicking 'minimize' button [right upper corner, first button to the left { _ }]. This will allow you to adjust settings later on.
7. From the main menu in the 'PRAAT objects window' select 'NEW'.
8. This will open the following window:

Record mono Sound...
Record stereo Sound...
Sound

Matrix
Tables
Tiers

Create TextGrid...
Create Strings as file list...

Optimality Theory
Articulatory synthesis
Polynomial
Multidimensional scaling
Neural nets

9. In most cases, you will record a single speech or voice sample and for that purpose you can select 'Record mono Sound..'. If you want to make stereo recordings, you obviously have to use "Record stereo Sound'. The latter option, for example, can be used to digitize the stereo output signal of the EG-2 PC Electroglottograph from Glottal Enterprises (http://www.glottal.com/electroglottograph.html), thus giving you access to a simultaneous recording of a speech and EGG signal.

10. Next, the SoundRecorder window will appear (shown here for mono recording)
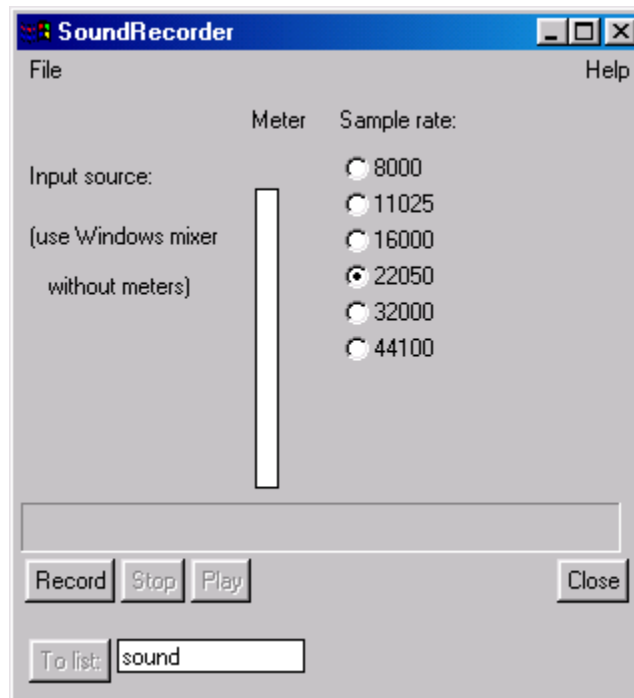
11. First set the sampling rate. In most cases the default (22 kHz) will be more than sufficient. If your computer has less disk space, you may want to use a lower sampling rate (11 kHz). If you want to record at CD quality, select the highest

sampling rate (44 kHz). This means you will have to store 44100 samples per second per channel (= about 176400 Bytes with a 16 bit Sound card!).

12. To record a signal, use a (preferably) high-quality microphone connected to the MIC input (do not use Line Input!) from the sound card, and click the 'Record' button. Some standard (cheap) computer microphones will not pick up frequencies below 100 Hz (check specifications).

13. Take a deep breath and speak the sentence <we stop doing the right thing> three times. Watch how the meter shows input level by green bars. If you are finished, click the 'Stop' button. Now the signal is stored in RAM (Random Access Memory), but not yet available for further processing (except that you can listen to the recording by clicking 'Play').

14. If the recording is to your satisfaction (check with 'Play'), you can add a name for the recording in the 'To list' box (in stereo recording mode you will see two of these boxes, one for the 'left' and one for the 'right' channel) and click on the 'To list' button. This will put your object in the 'Objects window'.

15. If you now go to the 'Objects window' you will find your sound object under the name 'Sound {name}'. You can always change this into any other name if you like. Just click on 'Rename' (lower part of window), and write down a new name (e.g., we_stop). It is a good strategy to give objects easy identifiable names.

16. This is just one example of creating a speech object. You can also digitize a speech sample from tape (DAT or cassette) using the line-input from your sound card. But, make sure you select 'Line-In' from the Volume Control window and deselect the microphone input. Also, you may wish to set the Line-In Balance option (select "Playback" mode under "options" -> properties) to mute, otherwise you get a continuous auditory feedback while you are recording (this however, may be useful to check the contents of a tape).

17. Finally, you can read files from disk (PRAAT supports various formats), including so-called 'long sound files'. Basically, these are pre-recorded sound files that are stored on disk and the program will allow you to select small portions of the total signal for analysis. This way, you can have files of up to several hours (if your computer has enough disk space), which can be handled in a piece-wise manner. In this tutorial, I will deal with the recording and handling of long sound files later (section 9).

3. Processing a signal (optional)

1. There are many things you can do with a speech object in terms of processing. You can filter the signal, enhance specific frequency regions etc. In this section, I will only describe the option of filtering the signal. In general, this is not necessary in PRAAT, but if you want to focus on a specific frequency region, the filter option comes in handy.

2. The first step is to select the original sound object (click on its name in the list).

3. To filter the signal do the following:
   ▪ Select 'Filter' (right-hand menu of Object window) -> Filter (formula)
   ▪ Change the formula to a low and high pass value (in this case I create a high pass filter at 10 Hz and a low pass filter at 5000 Hz):

- o if x<10 or x>5000 then 0 else self fi; rectangular band filter (N.B. the 'x<10' is set to a arbitrary low value; if your microphone does not pick up frequencies below 100 Hz, set this value to 100) and click 'OK'.
  - ▪ This will create a new (filtered) object in the list ({name} + _filt)
1. Play both the original and processed signal (try it now). Can you hear a difference?

4. Label a waveform

1. Sometimes it can be useful to segment a speech waveform and attach labels to each segment.
2. Select the original (= non-processed) sound object by clicking on its name
3. Go to 'label & segment' and select 'To Text grid'. This will bring up the following window
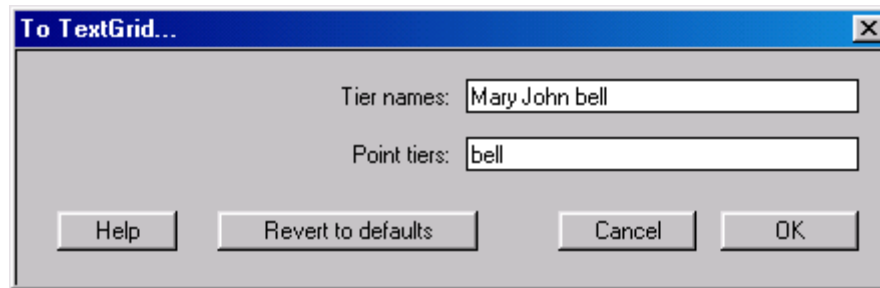
4. Change the names under the 'Tier names' option to identify segmentation *categories*, e.g., words syllables sounds (use space to separate names). So, the labels you input here are meant to indicate a level of segmentation, not the individual items. Make sure you erase the default names (Mary John bell), because they make no sense to you later.
5. The 'Tier names' are used to provide a label for intervals or specific discrete time points. The labels that appear in the 'Point tiers' box are automatically assigned to points, whereas the labels provided in the input window for the 'Tier names' are assigned to temporal intervals for e.g., the durations for the words in a given utterance. I will focus on intervals only in this tutorial, so you can leave the input window for Point tiers blank.
6. Select both the speech object and the Text grid (they share the same name) using the CTRL-key (click on speech object, depress CTRL-key and then click on Text grid)
7. On the right hand side of the window a new menu will appear. Select 'Edit' and the following window will appear (obviously, the speech signals will look different for your samples):
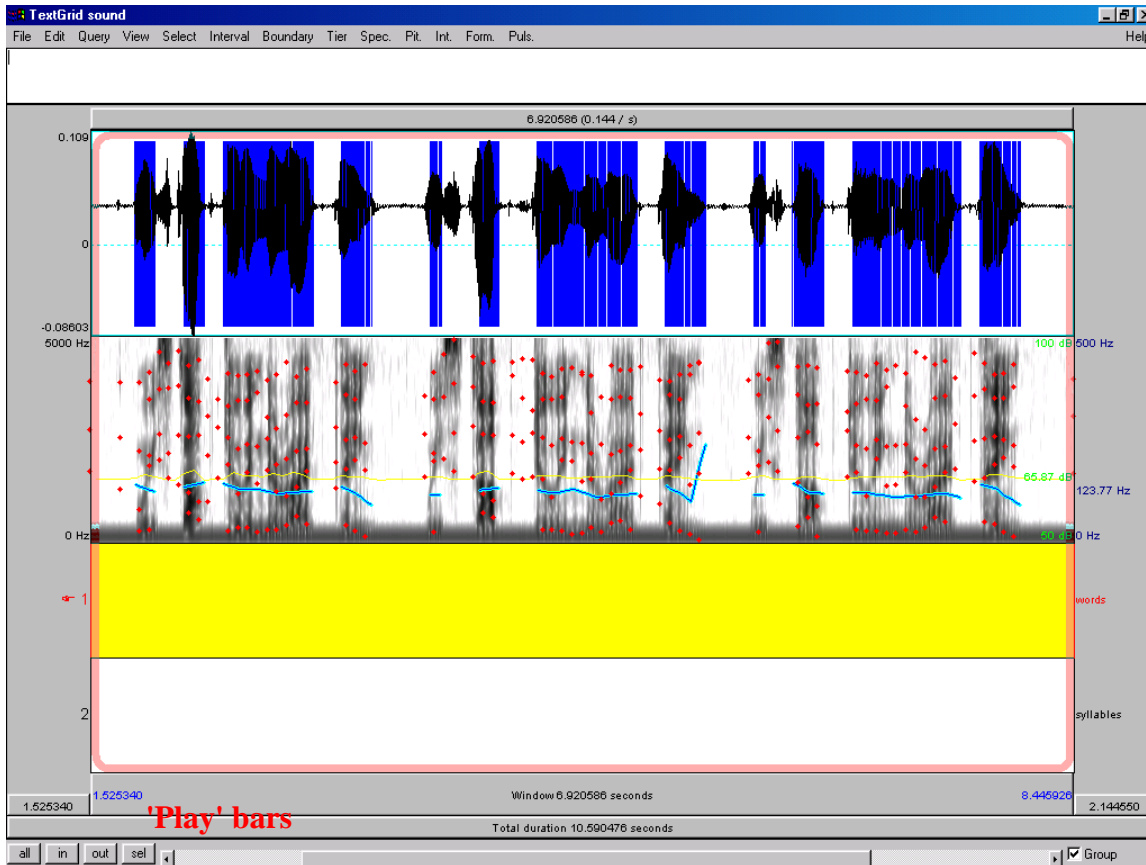
8. Maximize the window using the appropriate Windows button (⊡). You can listen to the entire speech sample by clicking on the lowest horizontal 'Play' bar (see picture). The upper 'Play' bars are divided in columns, as determined by the cursor position and/or selection (see next subsection).
9. Now you can segment words and syllables in the following way:
    - First select a portion of the total signal (remember, you recorded three separate sentences originally), e.g., the middle sentence. You do this by clicking left to the start of the middle sentence and then while keeping the left-mouse button depressed move the selection window to the right, i.e., the end of the middle sentence. Release the mouse button, and the sentence will be selected (a rectangular pink colored shape surrounding it). Then click 'sel' (lower left corner of the window). This will create a new window, zooming in on your selection. You can play certain segments of the signal shown in the window by positioning the cursor anywhere in the displayed signal (just click once at the preferred location). This will bring up a vertical red line in the signal, demarcating the position on the time axis you have selected (see time value in seconds marked in red font at the top of the window above the cursor position). The vertical line will divide the upper 'play bar' in separate segments, which can be played separately by clicking in the appropriate part of the upper bar (alternatively, you can press the TAB key, which will play the

10

segment to the right of the cursor or the part that is selected). The bar below that plays the whole selection in the window and the lowest bar plays the whole signal (3 utterances in this case). Just click on each of them to find out how this works (try it now). You can make further more detailed selections from the original signal to make your segmentations more accurate, but for now let us work with the current zoom-level.

- Position the cursor (using Left-mouse button) at the onset of the first word ("we"). Use the upper 'play bar' to check your selection. Then go to the first ('word') tier and click with left-mouse button on the circle-cursor. This should create a blue[5] vertical line, demarcating the onset of the first word.
- Then put the cursor (in the same way) at the end of the first word, thereby using the upper play-bar to listen carefully where the /e/ stops and the /s/ for the next word is beginning. Leave the cursor at the (for you) correct position and click on the circle-cursor at the same position as the vertical cursor. This again will create a blue line, demarcating this time the end of the first word. Now if you click in the segment demarcated by the two blue lines, push the TAB button, and you will hear the word <we>. The interval will also turn yellow. If you now simply type the word "we", this will appear in the indicated yellow segment.
- Continue this process for all the words (onset + offset) in the selected sentence. Notice that you can change the position of the blue lines, by simply clicking on the line, keeping the left-mouse button depressed and moving the line to a new position.
- If you are finished segmenting the words, it should look similar to this (I left the syllable tier out for display purposes. Do not worry yet about the different aspects of the window contents. I will get to that later):

---
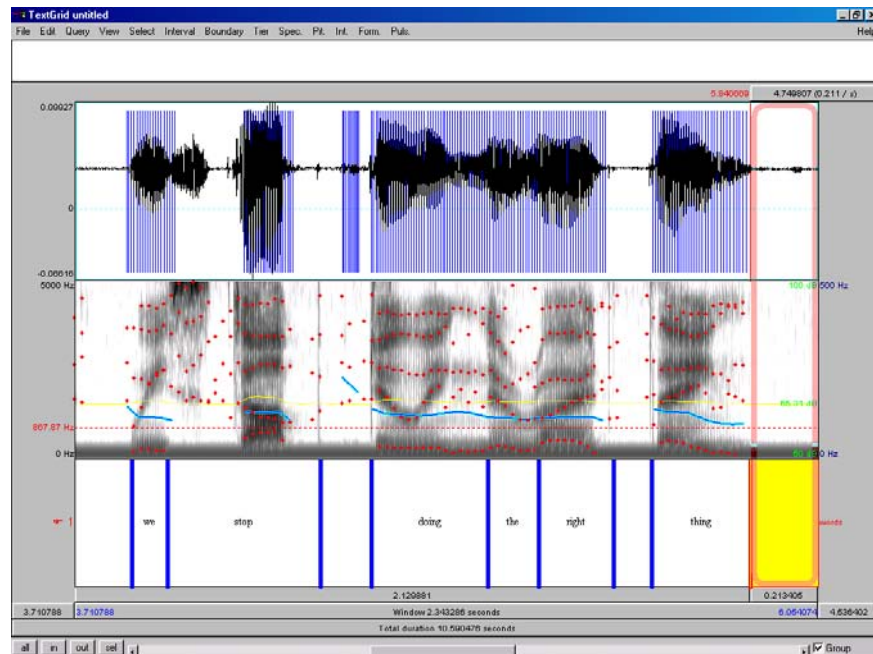
[5] The line initially is red, but if you click anywhere in the signal field, it will turn blue

- If a sentence contains multi-syllabic words, you could do the same thing for the syllables in the sentence, this time using the second ('syllable') tier and similarly, you could repeat this for a third ('sound') tier etc.
- You can also make a plot of this figure. Just close the active window, make sure both objects (sound + TextGrid) are selected, and choose 'draw' from the menu on the right side of the 'object window'. This will create a plot in the 'picture window' of the acoustic signal with the labels below it (in this case only for the middle sentence). If you want to show only the part you have labeled, specify begin and end point (in seconds) of this part, or repeat the process for only a selection of the original signal[6]. You can actually determine the physical size of the plot by changing the selection in the 'Praat picture' window (pink rectangular shape) *before* you draw the graph. Click in an area of the 'Praat picture' window (e.g., left upper corner) and while keeping the left-mouse button depressed draw the new shape. This picture can be saved as a post-script file and printed using Ghostview® (see footnote 2) or printed directly using CTRL-P (or 'print' from the file menu) if a Windows printer is connected to your PC.
- Once you have segmented and labeled the speech signal, you can easily extract each interval separately or all intervals simultaneously. Click on 'Extract intervals' (while both signal and text grid are selected), and in the window that appears, select your tier number (in our case: 1 = words, 2 = syllables in our case) and the label text, e.g., "stop". This will extract the segment labeled "stop" from the speech signal and put it in the 'Praat objects' window as a new object. If you select 'Extract all intervals..', all intervals will be put in separate objects (try it now). Please note that in this case, the empty intervals will also be extracted!
- You can select the extracted signal and by using 'Edit' from the main menu on the right-hand side of the 'object window' you can see and listen to the signal (try it now).
- In the 'Edit' window, you can make selections similar to the way it was described earlier (using the mouse, while keeping the left-mouse button depressed). If the window seems to "disappear" below your screen, simply expand the window to its full size (using the ☐ button).

5. **General analysis (waveform, intensity, spectrogram, pitch, duration)**

1) PRAAT offers a general extremely flexible tool in the 'Edit..' function to visualize, play and extract information from a sound object. For clinicians this is probably the most useful option in the program and it is compatible with the Kay CSL software interface, which is probably the best-known software for speech acoustic analysis in

---

[6] Select sound object -> edit -> select relevant part using mouse -> under 'file', select 'extract selection' option; this will create a new sound object of the size you selected. Repeat the previous steps for labeling the sound signal.

the world of speech-language pathology (at least in North-America). Each of the options will be discussed next.

2) First, create a new speech object of a sustained /ɑ/ vowel production (see section 2).

3) Select the speech object and then choose 'Edit..' from the main menu on the right-hand side of the 'object window'. A new window will appear. If your sound signal only occupies a small part of the entire window you can select the relevant part and extract this selection to the sound object list (see footnote 6). Then close the 'edit' window, select the extracted sound object and choose 'edit' again.

4) In the top menu you will see the following options:
   - ➢ File (this allows you to extract selections in different ways, to open a script file etc.)
   - ➢ Edit (this allows you to copy or paste parts of a signal etc.)
   - ➢ Query (this allows you to get information on the cursor position, selection boundaries, define settings for logs and reports etc.)
   - ➢ View (this allows you to select the contents of the window (spectrogram, intensity etc.) and control zoom settings)
   - ➢ Select (this allows you to control cursor positions)
   - ➢ Spec. (This allows you to control the spectrogram settings and extract information; the frequency value at the cursor position is indicated on the left hand side of the panel in a red font)
   - ➢ Pit. (This allows you to control the voice [pitch] signal settings and extract information; by default the pitch signal is shown in a bright blue solid line and the value at the cursor position is indicated on the right hand side of the panel in a dark blue font)
   - ➢ Int. (this allows you to control the intensity signal settings and extract information; by default the intensity signal is shown in a yellow solid line and the value at the cursor position is indicated on the right hand side of the panel in a bright green font)
   - ➢ Form. (This allows you to control the formant settings and extract information; by default the formants are shown in red dotted lines). The size of the window displaying formant values can be set in the 'Formant settings..' option (set 'maximum duration (s)' to a convenient value).
   - ➢ Puls (this allow you to set pulses (necessary for e.g., pitch analysis) and to extract specific information on voice parameters like jitter and shimmer; pulses are indicated in the top panel with vertical dark blue solid lines)

The following window shows the default view for (in this case) an /a/ sound:
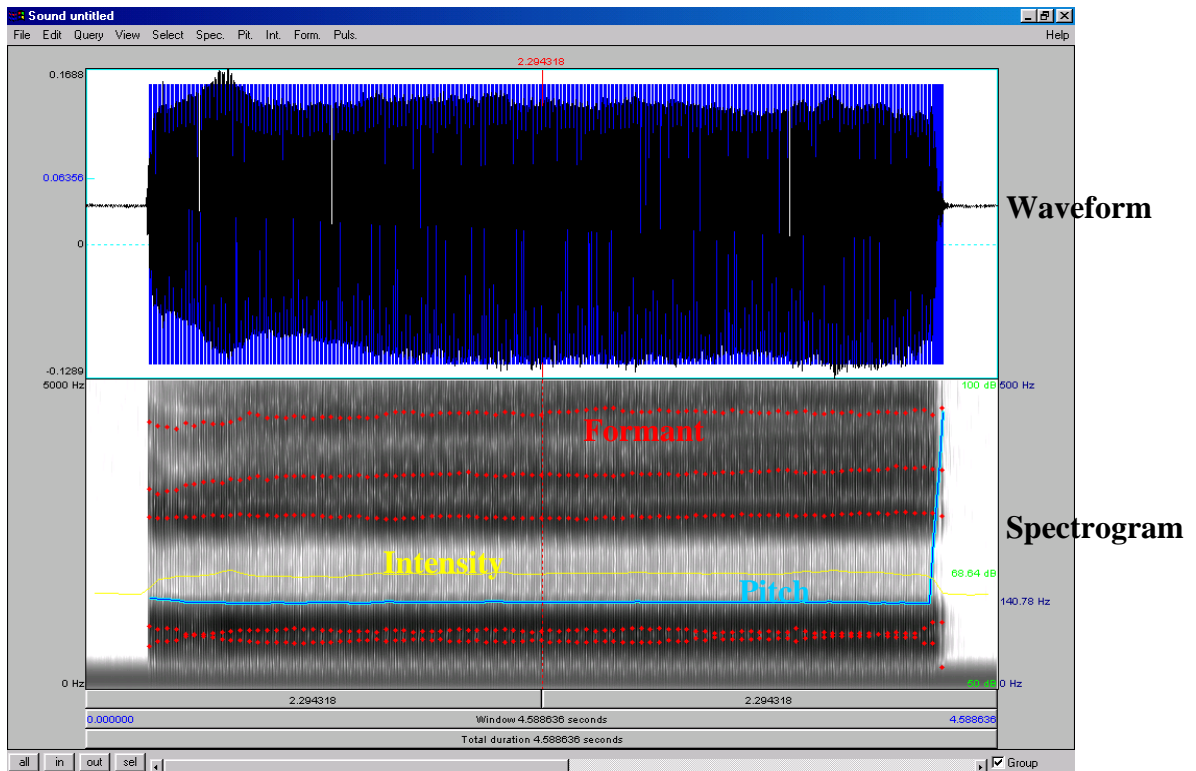
5) If you make a small selection of the signal (in the usual way), and zoom in (click on 'sel'), you will see more detail (try it now). 'Out' obviously zooms out and 'all' gives you the original full display (try it now).

6) If you click on a particular position of any of the displayed signals, it will show you the time at the cursor position, but it will also allow you to extract information about local pitch, intensity, formant and jitter/shimmer values. For example, position the cursor in a stable middle part of the vowel and do the following:

a) Go to 'Pit.' and select 'Get pitch' (F10 will also do the trick). A local pitch value will be displayed in a separate window. Go to 'Int.' and select 'Get intensity' (F11). A local intensity value will be displayed in a separate window.

b) Go to 'Form.' and select 'Get first formant' (F1). The local first formant value will be displayed in a separate window. Do the same for the second formant (F2), third formant (F3), and fourth formant (F4). It is also possible to get a complete report on all formant values at the selected location of the signal by selection the 'formant report' option in the 'Form.' menu. If you make a selection in the signal (instead of simply choosing a single point with the cursor), the formant report will list several temporal indices and their corresponding formant values within the selection boundaries. The number of indices depends on the value specified for the number of time steps ('Form.' -> Formant settings.. -> time steps).

c) F5-F8 are reserved for extracting information on cursor/selection positions; see 'Query' menu.

d) All values can also be stored in a file using the log settings under "Query". However, this option will not be further discussed as this information is extensively covered in the PRAAT manual (search for 'log').

(Try all these options now at different locations in the signal and compare).

e) The settings for each of the signals (spectrogram, formants, pitch, & intensity) can be changed in their respective menu options as indicated above. In general, if there is no specific reason to change the default settings, stick to them! However, if you do want to change specific options, like displaying a narrow band spectrogram, this is where to do it. For example, a narrow band spectrogram can be created as follows:
   ➢ Select in the 'Spec.' menu the 'Spectrogram settings' option
   ➢ Make sure the window shape is set to 'Gaussian'
   ➢ For a wide band setting change the value of the Window length box to 0.0043
   ➢ For a narrow band setting change this value to 0.029
   ➢ Try both versions and see the results (try it now); compare this to the standard (default) setting
   ➢ More information on these options can be found in the help file "Sound: To Spectrogram.." (See also 6.4)[7]

7) In addition to the above information, the 'Edit' window also allows you to do accurate temporal measurements, e.g., determining Voice Onset Time or VOT for a word like /pɛt/. You can do this as follows:
   • Create a speech object for the word /pɛt/ (repeat 3 times, and select best sample)
   • Select the speech object and choose 'Edit..' from main menu
   • If necessary, zoom in to target word by using the selection option (select using mouse and than choose 'sel' from lower left window menu)
   • Now zoom in to the onset of the word /pɛt/ (just before onset /p/ and initial portion of the vowel. Then select the interval between the onset of the release burst for /p/ and the onset of voicing for the vowel /ɛ/. This interval is the VOT interval (try it now).
   • The duration of the interval (in sec) is indicated above the selected portion of the signal (pink rectangle). The value between brackets is a conversion to Hz, which for this analysis is not meaningful. Please note that for accurate positioning, the left and right outer rim of the selection rectangle are placed on the exact location where you clicked with the mouse. If you like, you can go to the 'Select' menu and move begin

---

[7] In PRAAT v. 4.47 a new feature has been added to the 'spec.' menu, viz. 'view spectral slice'. This opens a new window with a (narrowband) spectrum of the signal at the cursor position (or for the selected interval). At the same time, this spectrum slice is placed in the main sound objects window. This object can be used for further processing (see also section 7.4).

and/or end of selection to nearest zero crossing. This will provide a standard onset and offset point close to your original selection.

6. Voice analysis

An option that is of particular interest to clinicians who work with voice patients can be found under the 'Puls.' menu. This menu contains a number of measures for quantifying irregularities in the duration (jitter) and amplitude (shimmer) of individual cycles (demarcated by the dark blue 'puls' lines in the waveform signal) in voicing. However, you have to be aware that the default settings for pitch analysis (see Pitch settings... from the Pit. menu in the Sound Editor window) have been optimized for intonation research. The current version of Praat (v. 4.1) suggest not to change the default values for 'silence threshold' and 'octave jump cost' but to adjust the Pitch range. The default range in Pitch analysis is 75-600 Hertz, but for pathological voices you have to increase the range to include lower values (e.g., 50 Hz) in male voices. To illustrate the use of the voice analysis options do the following:

- Record a sustained vowel /a/ production (6 seconds or longer)
- Select the speech object for this sound and choose 'Edit..' from main menu
- In the Sound Editor, select a stable mid part of the sound (± 4 seconds) and extract this portion as a separate sound object (File -> Extract selection).
- Close Sound Editor window and select extracted sound object in main Praat objects list and choose 'Edit..' again from main menu.
- You should see something similar as shown in the figure of section 5. If not, make sure you have activated the pitch, energy, formant and (most importantly) puls displays.
- Before you continue, you first set the Pitch analysis settings as specified above.
- A first thing you can do is to go to the 'Pit.' menu and select the 'extract the visible pitch contour' option. This will create a Pitch object in the main objects list. Select this object in the list (it is labeled 'Pitch untitled', unless you gave your sound object a name before)
- Click on the info button in the main menu (bottom of Praat Objects window) and a separate window will open that contains a lot of information about this particular vowel production in terms of average values and variations (the most important ones are indicated with a red arrow):
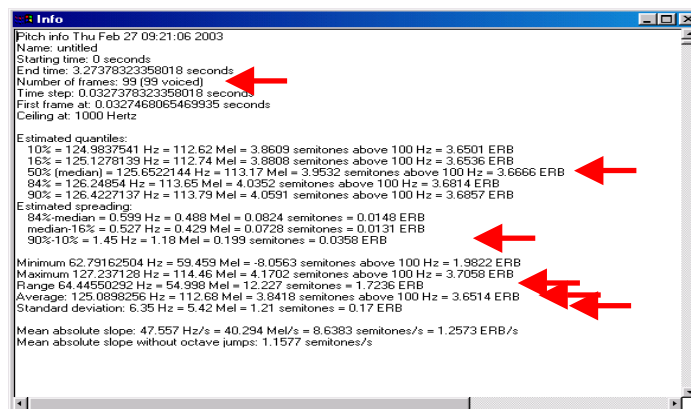
- The median value, the 10-90% spread around the median, the range, average and Standard deviation measures provide information about central moments of the distribution (see your statistics handbook for more information). The values are provided in different units (Hertz [Hz], Mel, Semitones, and ERB). For general users, Hertz will do, but if you are interested in knowing more about the other units, a good introduction on the web (+ links) can be found here: http://www.ling.su.se/staff/hartmut/bark.htm. For readers of Dutch, see also Rietveld and Van Heuven[8] (1997). The measured pointed at in the top of the window (number of frames) specifies two values: the number of frames and the number of frames voiced. As you can see in this case, they are equal. However, in pathological voices there may be voice interruptions (patients unable to continue voicing, e.g., in spasmodic dysphonia) and these two numbers will differ. The contents of this window can easily be selected (use mouse) and copied into a word file for future reference (e.g., as part of treatment evaluation).
- If you go back to the Sound editor window, you will find other measures of voice quality under the 'Puls.' menu. In fact, this menu lists several jitter (frequency perturbation) and shimmer (amplitude perturbation) measures. A very useful overall analysis of several voice features can be gathered by selecting the 'voice report' option in the 'pulses' menu of the 'Edit' window. An example of the output is shown in the next figure.

```
Analysis start: 0.10554608504361296 seconds
Analysis end: 0.18300278970701375 seconds
Analysis duration: 0.0774567046634008 seconds
Fraction of locally unvoiced frames: 0 (0/8)
Number of voice breaks: 0
Degree of voice breaks: 0 (0 s / 0.077457 s)
Jitter (local): 1.091%
Jitter (local, absolute): 100.115E-6 seconds
Jitter (rap): 0.013%
Jitter (ppq5): 0.008%
Jitter (ddp): 0.040%
Shimmer (local): 0.704%
Shimmer (local, dB): 0.061 dB
Shimmer (apq3): 0.266%
Shimmer (apq5): 0.525%
Shimmer (apq11): --undefined--
Shimmer (dda): 0.797%
```

- A detailed description of the various parameters and the way it is calculated can be found in the 'Voice' section of the PRAAT main manual. The following information list specific jitter and shimmer measures and a reference to the pages in Baken, R.J., & Orlikoff, R.F. (2000, Clinical Measurement of Speech

---

[8] Rietveld, A.C.M., & Van Heuven, V.J. (1997). Algemene Fonetiek. Coutinho: Bussum, The Netherlands

and Voice. San Diego: Singular Publishing Group, Inc.) in which
similar/equivalent measures are described:

a) Jitter (local) => see jitter ratio (but without the multiplication by 1000) on p. 201-202
b) Jitter (local, absolute) => as jitter ratio but without division by average period duration
c) Jitter (rap) => see 'relative average perturbation' on p.203-205 and table 6-34 on p. 208
d) Jitter (ppq5) => similar to jitter (rap) but with 5-point (as opposed to 3-point) estimate
e) Jitter (ddp) => original PRAAT jitter measure, which equals 3 times jitter (rap)
f) Shimmer (local) => see p. 133 but this is the non-dB version!
g) Shimmer (local, dB) => see p. 133-134 and table 5-22.
h) Shimmer (apq3) => see p. 134-135 (APQ) and table 5-23 (but read carefully the comments in Baken & Orlikoff as the values cannot be directly compared to the data generated in Praat)
i) Shimmer (apq5) => see shimmer (apq3); the 5-point window size is preferred by some (see B & O, p. 135).
j) Shimmer (apq11) => see shimmer (apq3); this is the original APQ measure suggested by Takahashi & Koike (see p. 134, B & O).
k) Shimmer (ddp) => original Praat shimmer measures; 3 times the value of shimmer (apq3)

All these measures tell you something about frequency or amplitude
perturbation. See table 6-37 in Baken & Orlikoff for a comparison of the
different jitter measures (including ones not used in Praat).

1) The 'pulses' menu also lists many different jitter and shimmer parameters that are available in PRAAT. A detailed description of these measures can be found in the 'Voice' section of the main manual (where these measures are compared to the Multi-dimensional voice program (MDVP) parameters of the Kay company). Other information about these measures can be found in Baken, R.J., & Orlikoff, R.F. (2000)
2) In section 9.1, I will discuss another measure for voice quality (H/R ratio).

**7. Spectrographic analysis (optional)**

In addition to the options available in the 'Edit' menu, which will suffice for most
common users of PRAAT, you can also use more specific spectrographic displays in
PRAAT. This chapter details these options.

1. Close the analysis window and select again the speech object of the vowel /ɑ/.
2. Choose 'Spectrum -' from the main menu (right hand side of Object window).

3. Choose 'To Spectrogram..' from the 'Spectrum' menu, and the following window appears:
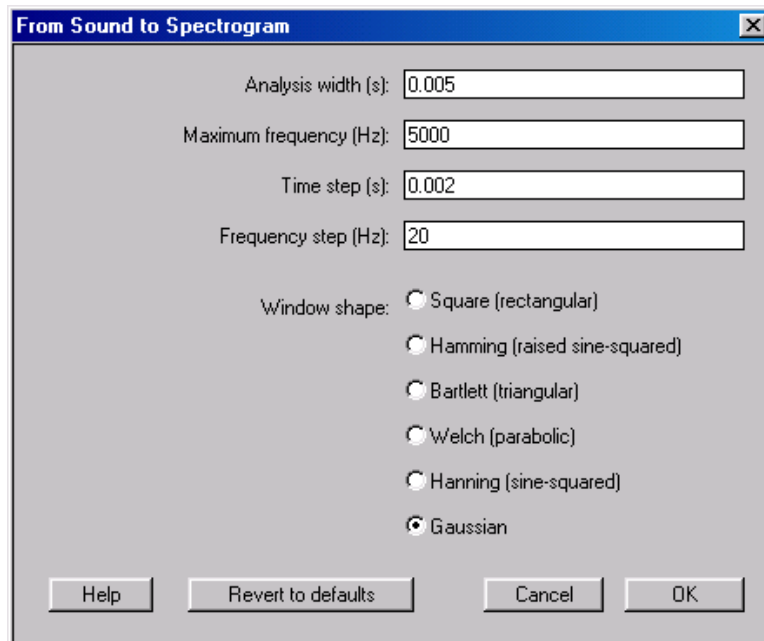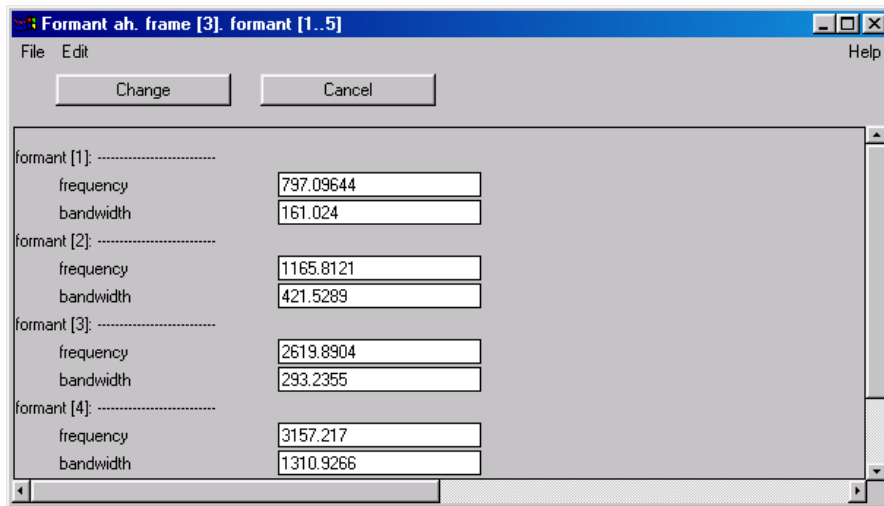

4. Do not worry about the 'Time step (s)' or 'Frequency step (Hz)' parameters (this is for more advanced users) and keep their default values. Also, leave the Window shape at 'Gaussian'. However, change the 'Maximum frequency (Hz)' parameter to the Nyquist-frequency of your speech object or lower (e.g., if you had sampled at 10 kHz, NF = 5000 Hz). Furthermore, you can change the 'Analysis width (s)' parameter value, for this determines the bandwidth of your analysis (as already explained in 5.11). Just to repeat, as a rule-of-thumb the following values apply:
   - Wide-band = 300 Hz = **0.0043 seconds**
   - Narrow band = 45 Hz = **0.029 seconds**
   - The applied formula is: 1.2982804/ analysis width
   - The default value is **0.005 seconds** (thus around 260 Hz bandwidth)
   - Try all three settings and have a look at the resulting spectrograms (select spectrogram signal and choose 'View' from main menu) (try it now)
   - In order to get specific formant values, determine a specific point in time from which you want to retrieve formant values. Close the spectrogram view and choose from the menu on the right the option 'To Spectrum (slice)..'. Specify the selected time value (in sec) and click OK. Now a new object appears in the object list 'Spectrum_{name}'. Select this object and choose the 'Edit' option in the menu to the right. This brings up a window similar to this:
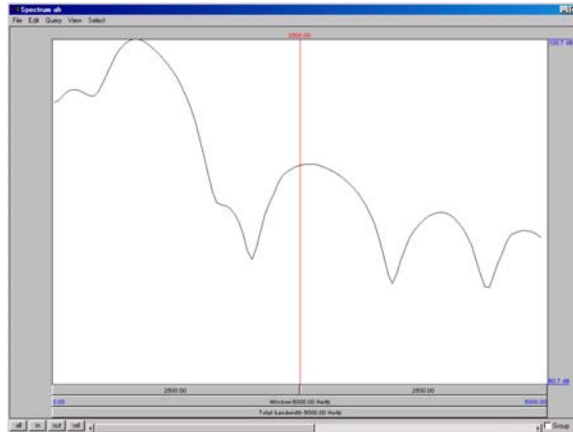
- The cursor can be moved to the desired location in the spectrum (in this case to the second formant peak) and the corresponding formant value can be read at the top (the horizontal cursor also shows the corresponding power value in dB for that particular frequency, but this option makes most sense if your sound recordings are calibrated first). If you compare these results with the formant traces in the Edit window (see 7c) you may notice that for /a/ F1 and F2 are not clearly differentiated in the spectrum. So, in a sense, this spectrum is less accurate than the formant tracing in the 'Edit' window.

5. There is another way to get more details on formant values and bandwidths
   - Close the spectrum window
   - Select the same original /a/ sound object you created earlier
   - From the main menu choose 'Formants & LPC'
   - Select 'To Formant (burg)..'
   - The following window will appear
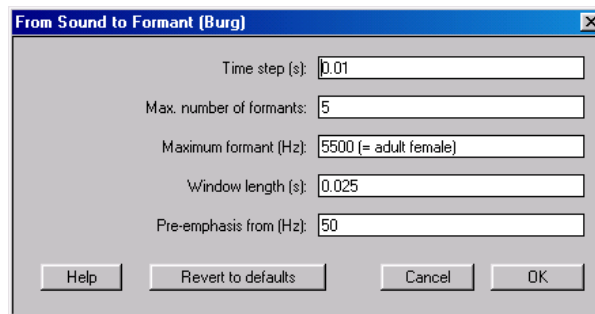
- The 'Window length (s)' option is similar to the 'Analysis width' option in the 'To Spectrogram..' menu (see 6.4). Make sure you change "Maximum formant (Hz)..' to the same value as used for the 'Edit' spectrogram settings (for males 5000 Hz).
- Select 'Formant_{name}' and choose the 'Inspect' from options on the bottom row of the 'object window'
- Click on 'open' two times (two separate windows) and the following menu will appear (F5 not shown):

20

Formant ah. frame [3]. formant [1..5]

File   Edit   Help

Change   Cancel

formant [1]: --------------------------
    frequency   797.09644
    bandwidth   161.024
formant [2]: --------------------------
    frequency   1165.8121
    bandwidth   421.5289
formant [3]: --------------------------
    frequency   2619.8904
    bandwidth   293.2355
formant [4]: --------------------------
    frequency   3157.217
    bandwidth   1310.9266

- These values should be close to the values we see in the formant report in the 'Edit' window. (compare)
- If you like you can plot a spectrogram in the 'Picture window' and in addition you can plot the formants tracks on top of it. Do as follows:
  - Extract spectrogram from 'Edit' window (select sound object -> 'Edit' -> 'Spec.' -> 'Extract visible spectrogram').
  - Select spectrogram object (first select drawing area in 'Picture window' as described in 4.9)
  - Choose 'Paint..' from the menu, and set 'To Frequency' option to Nyquist frequency (e.g., 5000 Hz for 10 kHz sampling rate) and click 'OK' (as you always do in confirming your choices)
  - Next, go to 'Pen' option in the main menu of the 'Praat picture' window and select a different pen color (e.g., red)
  - Then, select the formant object from the same signal (be careful not to click in 'Picture window' !) and choose 'Draw -' from the menu
  - Choose 'Speckle..' from the options window (check that the 'Maximum frequency' is set to 5000 Hz). The formants are overlaid in a red (or whatever color you selected) speckled line across the spectrogram in the 'Picture window'

- For your information: If you select the 'Formant_{name}' object and select 'To LPC[9]..' you create a new object ('LPC_{name}'). If you select the original sound object together with this LPC object, the menu will show the option 'Filter (inverse)'. If you choose this option, the formant structure is used to generate the underlying source signal of the speech object. Simply select the new sound object created this way, and use 'Edit' to view it (do not worry about the strange looking spectrogram). If you listen to the sound, it is more of a buzz, which is similar to what you

---

[9] LPC stands for Linear Predictive Coding

would hear if you would listen to the actual voicing signal without vocal tract modifications (try it now).

8. Intensity analysis (optional)

1. Similar to what is said about the use of spectrographic techniques outside the 'Edit' window, one can do intensity analyses separately. Again, the 'Edit' options are by far the easiest way to go, but the procedures explained below can be used as an alternative if so desired.
2. Select the original speech object of the vowel /ɑ/.
3. From the main menu, choose the option 'To Intensity..'
4. Keep the default values for the selection window (unless your minimal pitch is expected and measured below 100 Hz) and click 'OK'
5. Select again the speech object, and choose 'Edit' from the main menu
6. Determine a selection of your signal, that is, decide on a onset and offset point for an interval for which you want to calculate a mean (and standard deviation) level of intensity. Then, close the window.
7. Select the Intensity object ('Intensity_{name}')
8. Choose from the menu the option 'Query-'.
9. Choose the option 'Get mean..' and define the interval according to your previous selection (see 7.5) (try it now)
10. The mean value will appear in a separate ('Info') window. You can save this information by writing it down or copy and paste it into a document
11. Do the same for 'Get standard deviation..'. (try it now)
12. FYI, if you select the speech object and choose 'Query' from the main menu, you get a lot more options for calculating levels of energy, as shown in the next picture. However, discussing these options would go beyond the purpose of this tutorial.

9. Pitch analysis

The 'Edit' window offers a number of jitter and shimmer measures under the 'Puls.' menu that were explained above (section 5.9). However there is another measure for voice quality that is offered in PRAAT.

1. Select the sound object for the vowel /ɑ/
   - Go to 'Periodicity-' in the main menu
   - Choose 'To Harmonicity (cc)' from the menu and stick to defaults (unless you have good reasons to change it)
   - Select 'Harmonicity_{name}' object and choose 'Query-' from menu
   - Choose 'Get mean..' from menu and the mean Harmonic to Noise ratio (H/N) is written in the 'Info' window. If you choose 'Get standard deviation..' you get the corresponding standard deviation. This measure is explained in more detail on p. 281-282 in Baken & Orlikoff (2000). Search the PRAAT-manual for 'Harmonicity' and you will find some information on this measure as well, including provisional norm values for /a/ and /i/ vowels.

2. Select again the sound object for the vowel /ɑ/
   - Go to 'Periodicity-' in the main menu
   - Choose 'To Pitch (cc)' from the menu and stick to defaults (unless you have good reasons to change it)
   - If you like you can select the 'Pitch_{name}' object and choose 'Edit' to view the signal and to select a suitable interval for analysis
   - Select Pitch object and choose 'Query-' from menu
   - Choose 'Get mean..' from menu and the mean fundamental frequency for the specified interval will be displayed in the 'Info window'. Do the same for 'Get standard deviation..'. A description of the difference between phonation (as in single vowel phonations) and speaking (as in normal speech or in reading aloud fragments of e.g., the Rainbow passage) fundamental frequency can be found in Baken & Orlikoff (2000), including some reference data. Normally, a more robust indicator would be the median using 'Get quantile' in the 'Query' menu, selecting the 0.5 (= median) quantile setting (try it now).
   - The 'Query' option also provides information on the mean absolute slope (and you can choose Hertz, Semitones, or Mel as unit) for the whole object (with or without octave jumps). Simply select 'Get mean absolute slope..' or 'Get slope without octave jumps..' (try it now)

10. Long sound files

1. Depending on your PC memory capacity and the sampling rate you choose, you can record speech samples for a varied amount of time using the normal 'Record Sound' option (in the 'New' menu). You can change the standard buffer size as used in PRAAT in the 'Control' menu, choosing 'Preferences' -> 'Sound input prefs'. Obviously, the amount of RAM you can allocate depends on the memory capacity of your computer.

2. If you want to record really long speech objects, do as follows:
   - Go to 'Start' (normally in lower left corner) from the Windows 2000 task bar[10], go to 'Programs' -> 'Accessories' -> Entertainment' -> 'Sound Recorder'.
   - This will bring up the following window:

   - Before recording check the quality of recording. Select 'File' from the main menu and choose 'Properties'. This will bring up the following window (depending on your sound card):
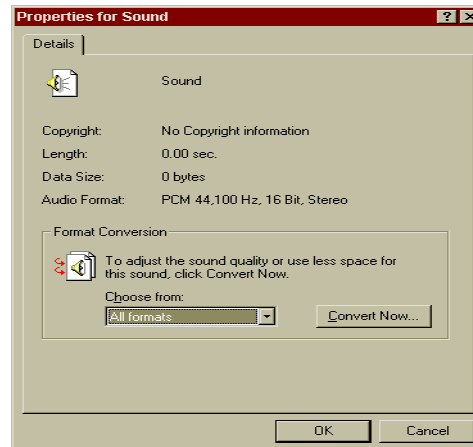
   - By clicking on 'Convert Now' you can set the desired format for recording the speech object (normally, choose CD-quality with 16 bit stereo at 44 kHz sampling rate, unless you expect problems in disk and memory capacity)
   - By clicking the button with the red dot, you can start recording your speech object (the input device is selected in the volume control window, see section 2. Create a speech object). The window will indicate how much recording time you have (TIP! If you stop the recording and then continue, the recording time will increase. However, at some point the program will start to write information to disk, which could slow down the start of continued recording).
   - After recording and checking the speech object (the buttons are similar to that of a ordinary cassette player, and I assume you are familiar with stop,

---

[10] Check help information for higher versions of Windows.

rewind, play and other functions) go to 'File' and choose 'Save'. The sound object can then be saved as a standard Windows *.wav file.

- If you have saved the Long Sound file, go to the 'Object' window of PRAAT, and choose 'Open long sound file..' from the 'Read' option in the main menu and select the file you just saved
- This will create a 'Long sound' object in your list
- In order to work with the file, choose 'view' from the menu after selecting the 'Long sound' object. Within the "view" window you can specify an interval for your selection, use 'Extract selection' from the 'File' menu and a sound object will be added to your list, containing the selected part of your long sound object. In the view window you will see two copies of the same recording, unless the original recording was in stereo. The selected sound object can be handled like any other sound object in PRAAT.

C. Finally

This tutorial is only a basic introduction to some of the functions in PRAAT. You are only going to learn and appreciate more of its possibilities by using the program extensively, using its manual and other sources (e.g., Sidney Wood's beginner's tutorial on the PRAAT homepage). In order to get more familiar with the functions that were introduced in this tutorial, continue with the following assignments:

a. Make a recording of the following words and calculate VOT and vowel durations: "pet", "ket", "tet", and "dead". What can you tell about the differences you find?
b. Make a recording of the following vowels and determine their formant patterns (F1, F2, & F3): /i,æ,ɑ,o,u/. What can you tell about the patterns?
c. Make a recording of sustained long vowels /ɑ/ and /i/ and determine their duration, jitter, and H/N ratio (including standard deviation)
d. Make two recordings of the following sentence, one in the form of a declarative and the other as a question.
   "The Canadian and Dutch players met on the ice" (you can also make up your own sentence, of course)
e. Determine average intensity level and its variation (SD), average fundamental frequency and its variation (SD), and the mean absolute pitch slope for both sentences. What can you tell about the differences?

If you work in groups, make sure you get a chance to work with the program yourself. *Watching is useful, but in the end it is the practice that counts!!*

Questionnaire for Evaluating the Tutorial

I would appreciate feedback on the contents of this tutorial and the form of this manual. Your answers to the following questions will be helpful to me in trying to improve both aspects (if deemed necessary). Some questions ask for a rating, using a five-point scale (1 = poor; 5 = excellent). If you have downloaded this tutorial from my home web page, please complete the form and send it to me by e-mail (p.vanlieshout@utoronto.ca) or fax (416-978-1596).

YOUR FEEDBACK IS VERY MUCH APPRECIATED!

_____

Questions	1	2	3	4	5
1. What is your overall opinion of this tutorial?					
2. What is your overall opinion of this manual?					
3. What is your overall opinion of PRAAT?					
4. How do you judge its applicability in clinical use in the field of speech-language pathology?					
5. How do you judge its applicability in research in the field of speech-language pathology?					

Open questions:

1. What part of this tutorial did you like most? And why?

    _____
    _____
    _____
    _____
    _____

2. What part of this tutorial did you like the least? And why?

    _____
    _____
    _____
    _____
    _____

3. Which part(s) of the manual did you find to be unclear or in need of revision?

_____
_____
_____
_____
_____
_____

4. What in your opinion should be the nature of the revision? (e.g., more examples, more clarification, less details)

_____
_____
_____
_____
_____

5. What parts of the manual did you like? And why?

_____
_____
_____
_____
_____

6. Do you feel that this introduction to PRAAT will encourage you to use the program in the future for clinical and/or other purposes? If yes, why? If not, why?

_____
_____
_____
_____
_____

7. Should there be another tutorial in this course for more advanced applications of PRAAT? Why?

_____
_____
_____
_____
_____

8. Finally, do you have any comments or suggestions you think are important to mention regarding this tutorial and/or manual?

_____
_____