

Acoustic Phonetics

A Brief Introduction to Praat

Dafydd Gibbon

Guangzhou, Autumn 2019

Objectives

At the end of the course, participants should be able to

- extract duration and fundamental frequency information from a speech recording, using Praat

Part I: Annotation

- Basic knowledge of signal analysis with Praat.
- Basic knowledge of syllable annotation with Praat.
- Practical application to the recorded data of course participants.
-

Part II: Syllable timing analysis

- Examination information in Praat annotation file.
- Analysis of Praat annotation file with Time Group Analyser.

Praat

- Praat is a phonetic workbench application developed in Amsterdam by Paul Boersma and David Weenink.
- “Praat” means ‘talk’ in Dutch.
- The basic functionality of Praat includes:
 - Input: speech recordings
 - Methods:
 - analysis of properties of speech signals such as spectral analysis, pitch analysis, annotation of signals with transcription labels
 - Outputs:
 - files with information about the speech signal
- The annotation information files which Praat produces
 - can be re-structured, and analysed with other means,
 - with Excel or Calc
 - with the online tool Time Group Analyser, for efficient analysis of timing relations in the speech signal.

Praat Input

Praat Input

Pre-recording phase:

- definition of purposes for which the data will be used
- scenario: domain, activities, speakers
- equipment and technical operator:
 - general: digital audio (recorder / laptop), digital video
 - specialised: laryngograph, etc

Recording phase:

- negotiate scenario with chiefs, elders, speakers
- ensure the recording location is quiet
- if possible ensure the microphones, video tripod etc. can be stably positioned

Post-recording phase:

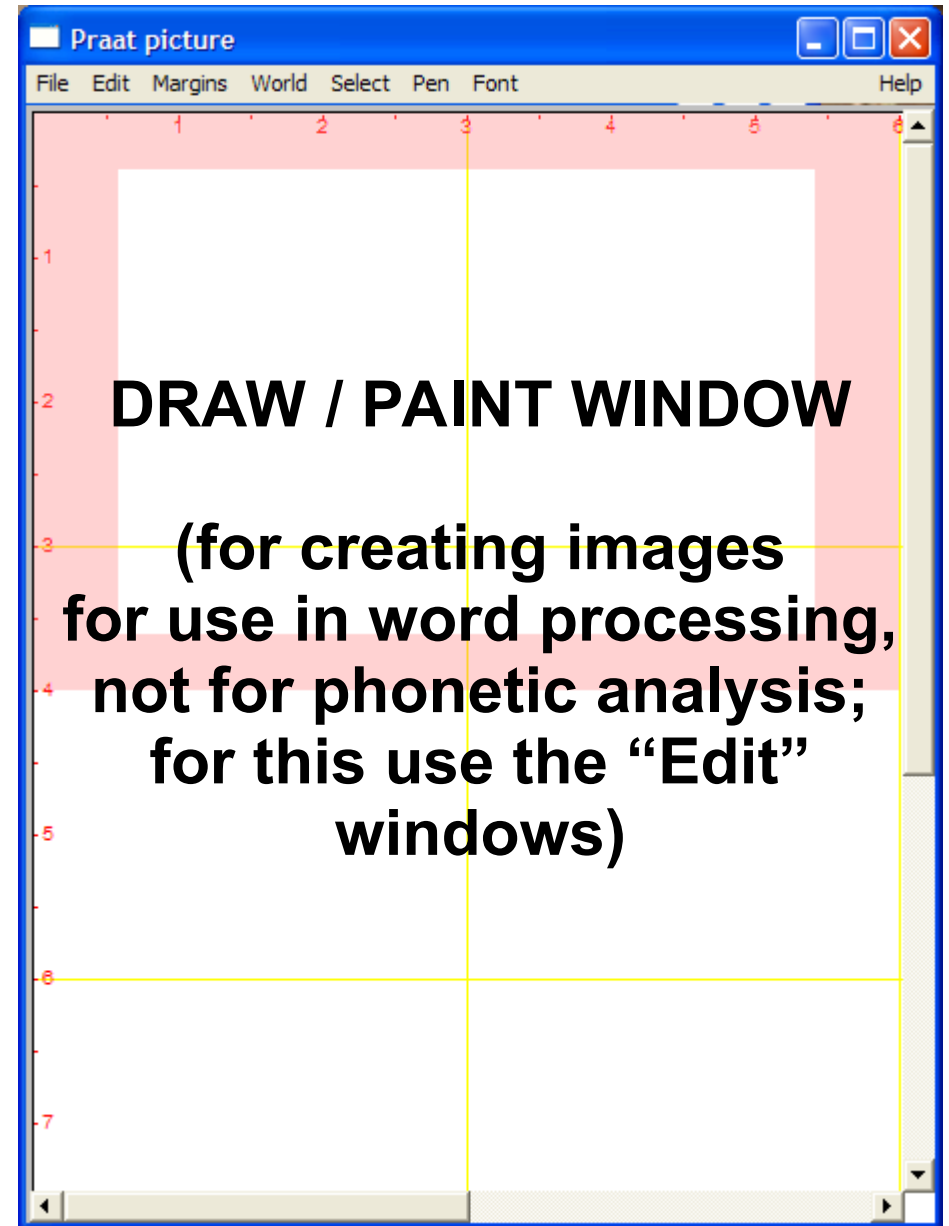
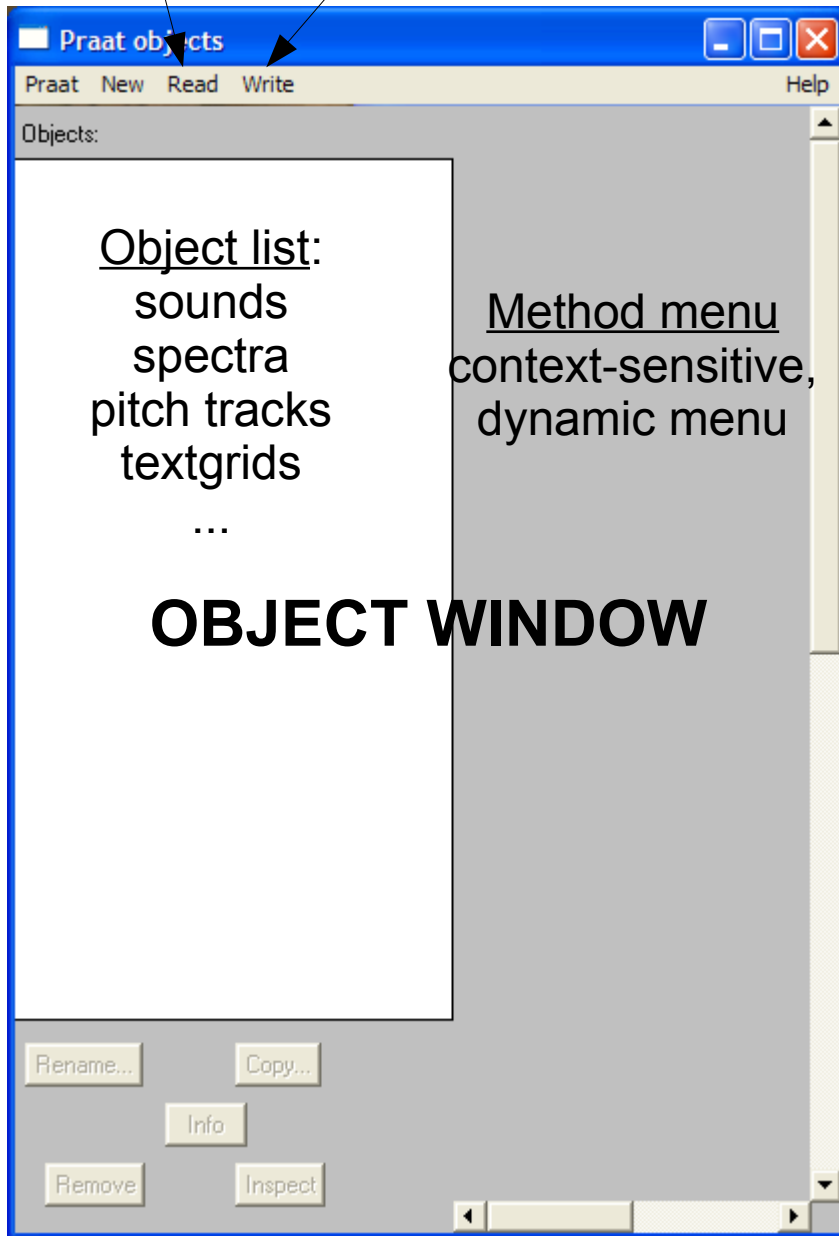
- provide recordings with metadata immediately
- label the data media immediately
- make safety copies immediately

Basic Praat Methods and Operation


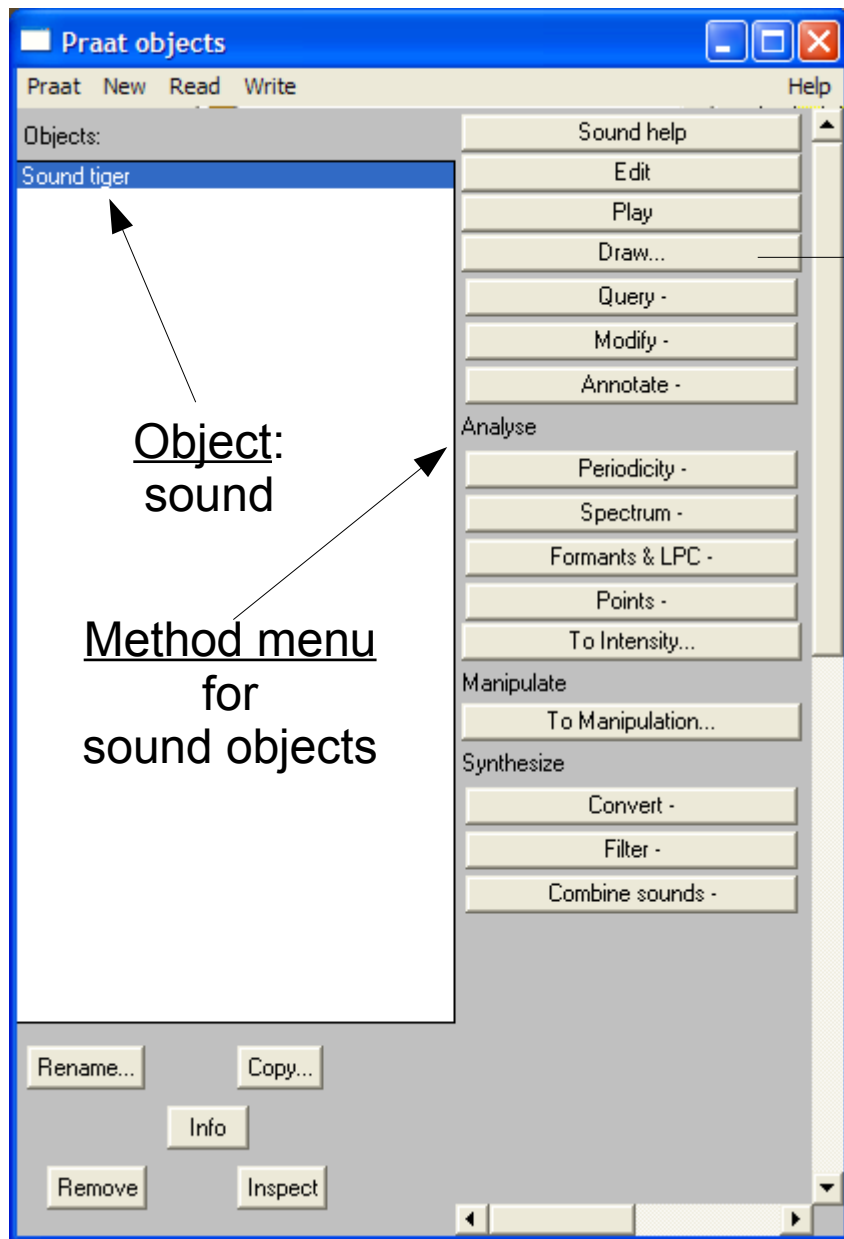
Praat Windows

Load file


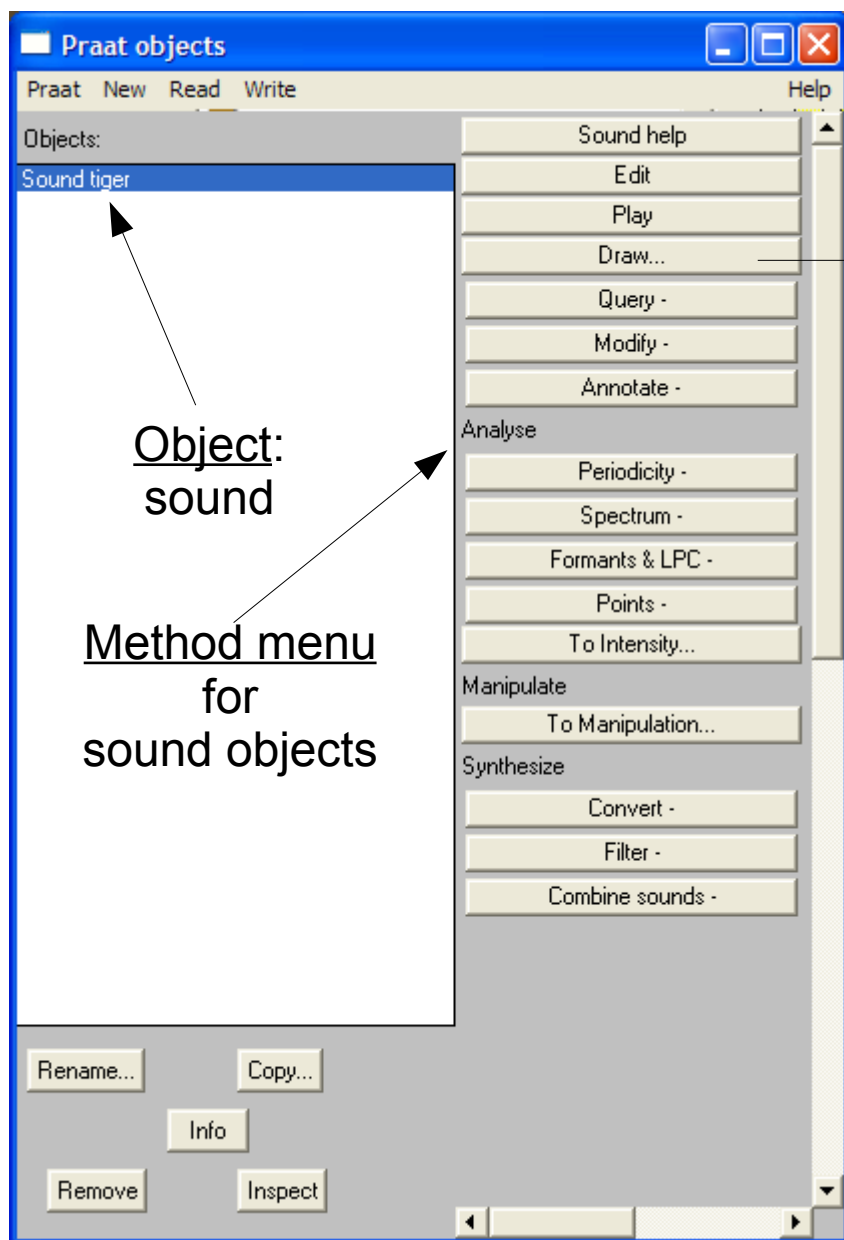
Save file



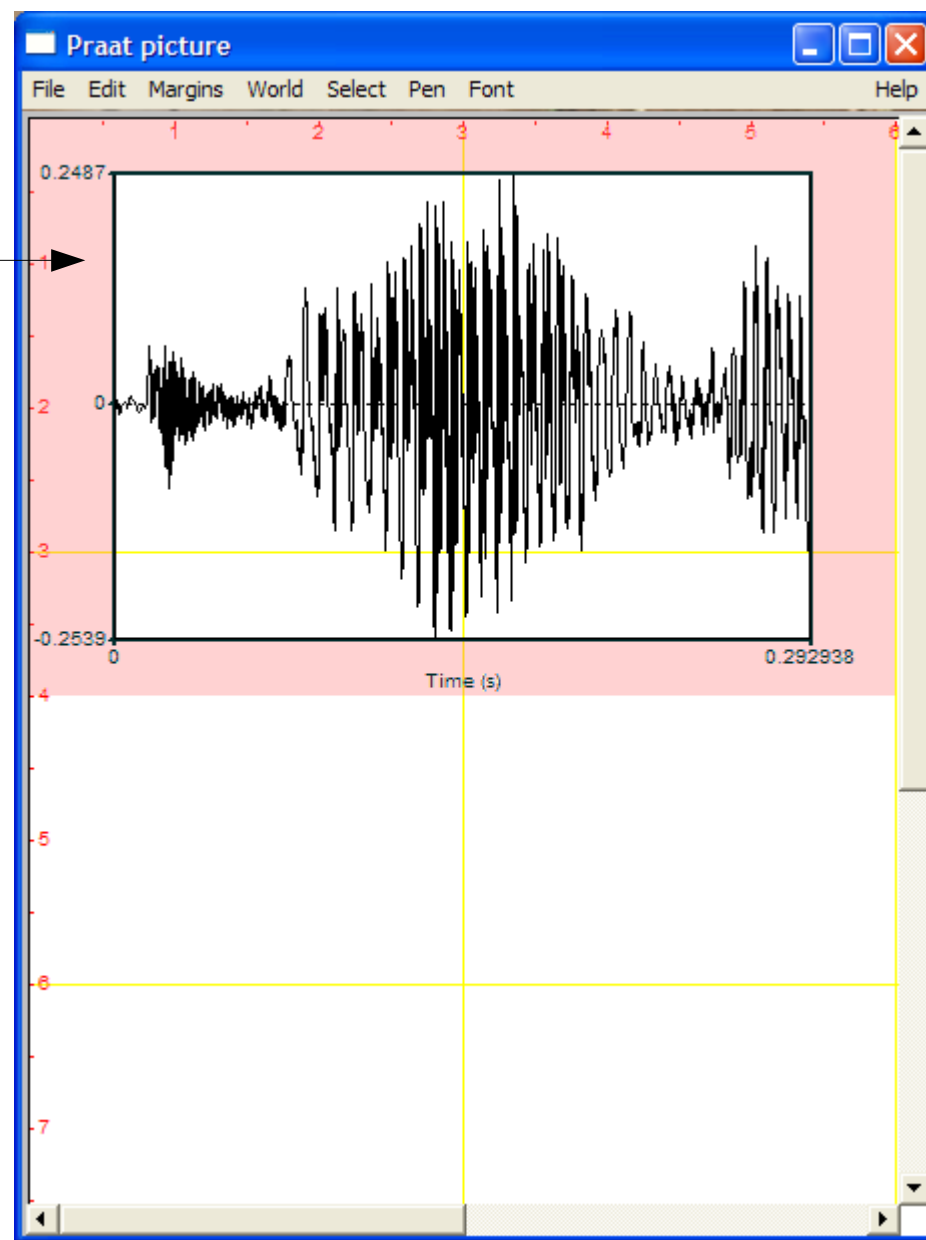
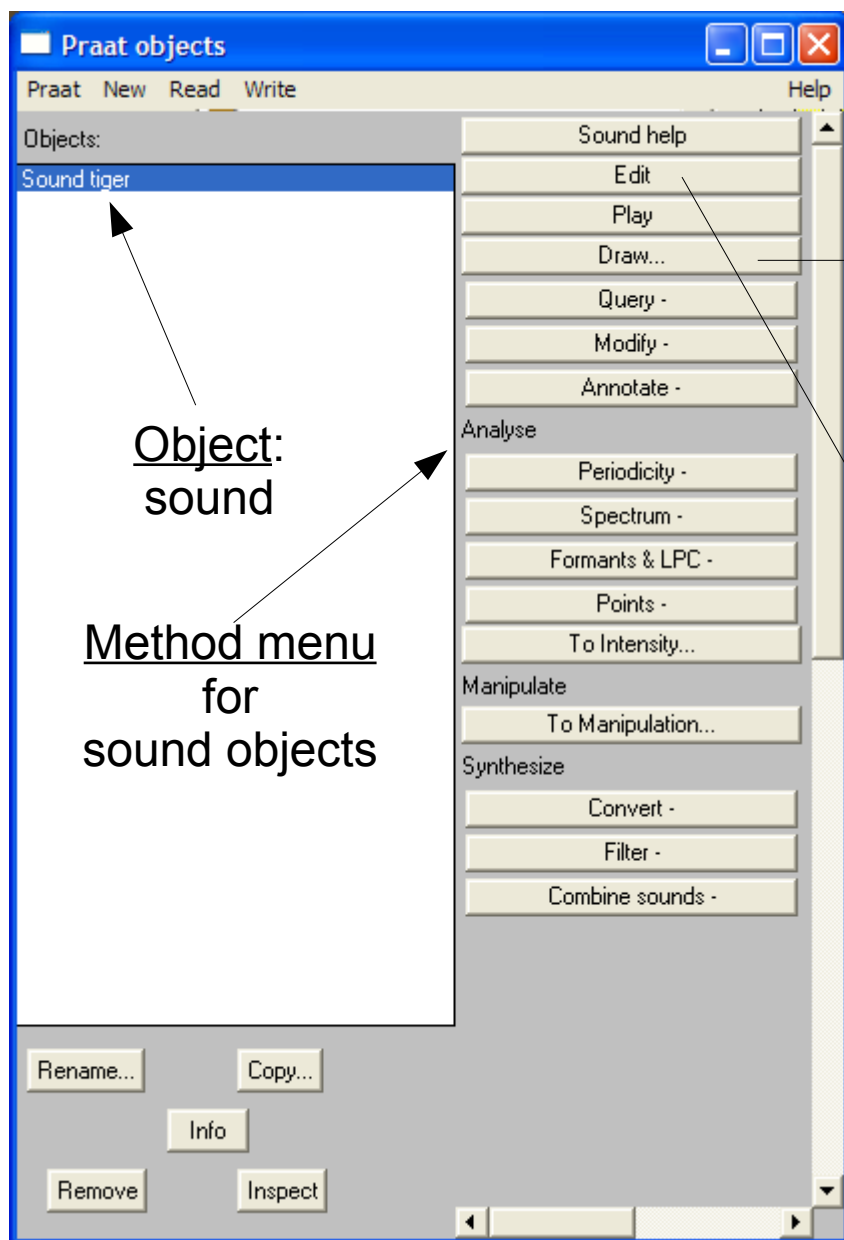
Processing Praat Objects



Processing Praat Objects



Processing Praat Objects



Processing Praat Objects

Praat objects

Praat New Read Write Help

Objects:

Sound tiger

Sound help
Edit
Play
Draw...

Praat picture

File Edit Margins World Select Pen Font Help

1 2 3 4 5 6

0.2487

Sound tiger

File Edit Query View Select Spectrum Pitch Intensity Formant Pulses Help

0.2487

0.146469

0.01224

-0.2539

0.146469 0.146469

0.000000 Visible part 0.292938 seconds 0.292938

Total duration 0.292938 seconds

all in out sel Group

Object: sound

Method menu for sound objects

Processing a Sound Object

The image displays the Praat software interface. On the left, the 'Praat objects' window shows a list of objects with 'Sound tiger' selected. A context menu is open over 'Sound tiger', listing options: 'Sound help', 'Edit', 'Play', and 'Draw...'. An arrow points from the 'Edit' option to the waveform in the 'Sound tiger' window. The 'Sound tiger' window shows a waveform plot with a red vertical dashed line at 0.146469 seconds and a blue horizontal dashed line at 0.01224. The x-axis is labeled '0.000000', '0.146469', and '0.292938'. The y-axis is labeled '0.2487' and '-0.2539'. The status bar at the bottom indicates 'Visible part 0.292938 seconds' and 'Total duration 0.292938 seconds'. Below the waveform window, there are buttons for 'Rename...', 'Copy...', 'Info', 'Remove', and 'Inspect'. At the bottom left, there are buttons for 'all', 'in', 'out', and 'sel'. At the bottom right, there is a 'Group' checkbox.

Object: sound

Method menu for sound objects

Selecting Part of a Sound Object

The image shows two windows from the Praat software. The top window, titled "Praat objects", has a menu bar with "Praat", "New", "Read", "Write", and "Help". Below the menu is a list of objects with "Sound tiger" selected. To the right of the list is a "Method menu" with buttons for "Sound help", "Edit", "Play", and "Draw...". The bottom window, titled "Sound tiger", has a menu bar with "File", "Edit", "Query", "View", "Select", "Spectrum", "Pitch", "Intensity", "Formant", "Pulses", and "Help". The main area of this window displays a waveform of a tiger sound. A red shaded region highlights a portion of the sound, bounded by vertical dashed red lines. The time values for these boundaries are 0.071819 and 0.233795. The duration of the selected part is 0.161977 (6.174 / s). The bottom status bar shows "Visible part 0.292938 seconds" and "Total duration 0.292938 seconds".

Object: sound

Method menu for sound objects

Time (s)	Description
0.071819	Start of selected portion
0.161977	End of selected portion
0.233795	End of visible part
0.292938	Total duration

Displaying More Properties of a Sound Object

The image shows two windows from the Praat software. The top window is titled 'Praat objects' and contains a list of objects with 'Sound tiger' selected. A menu is open for 'Sound tiger' with options: 'Sound help', 'Edit', 'Play', and 'Draw...'. The bottom window is titled 'Sound tiger' and displays a waveform and a spectrogram. The waveform is a black line on a white background, and the spectrogram is a grayscale image below it. The waveform has a red shaded region between 0.071819 and 0.233795 seconds. The spectrogram shows frequency components up to 5000 Hz. The bottom of the 'Sound tiger' window shows time markers: 0.000000, 0.071819, 0.161977, 0.059142, 0.292938, and 0.292938. The total duration is 0.292938 seconds. The bottom of the 'Praat objects' window has buttons: 'Rename...', 'Copy...', 'Info', 'Remove', and 'Inspect'.

Object: sound

Method menu for sound objects

0.071819 0.161977 (6.174 / s) 0.233795

0.2487

0

-0.2539

5000 Hz

0 Hz

0.071819 0.161977 0.059142

0.000000 Visible part 0.292938 seconds 0.292938

Total duration 0.292938 seconds

all in out sel Group

Displaying More Properties of a Sound Object

The image shows two overlapping windows from the Praat software. The top window is titled 'Praat objects' and contains a list of objects with 'Sound tiger' selected. A menu is open for 'Sound tiger', showing options: 'Sound help', 'Edit', 'Play', and 'Draw...'. The bottom window is titled 'Sound tiger' and displays a waveform and a spectrogram. The waveform is a black line on a white background, and the spectrogram is a grayscale image with red diamond markers. The time axis is marked with values: 0.071819, 0.161977 (6.174 / s), and 0.233795. The frequency axis is marked with 0 Hz, 5000 Hz, and 0.2487. The bottom of the window shows 'Visible part 0.292938 seconds' and 'Total duration 0.292938 seconds'. There are also buttons for 'Rename...', 'Copy...', 'Info', 'Remove', and 'Inspect'.

Object: sound

Method menu for sound objects

Displaying More Properties of a Sound Object

The image shows two overlapping windows from the Praat software. The top window is titled 'Praat objects' and contains a list of objects with 'Sound tiger' selected. A menu is open for 'Sound tiger', showing options: 'Sound help', 'Edit', 'Play', and 'Draw...'. The bottom window is titled 'Sound tiger' and displays a waveform and a spectrogram. The waveform is a black line on a white background, with a red shaded region highlighting a segment. The spectrogram is a grayscale plot with red diamonds and a blue line. The x-axis represents time in seconds, with markers at 0.071819, 0.161977 (6.174 / s), and 0.233795. The y-axis represents frequency in Hz, with markers at 0 Hz, 75 Hz, 155.34 Hz, and 500 Hz. The bottom status bar shows 'Visible part 0.292938 seconds' and 'Total duration 0.292938 seconds'. There are also buttons for 'Rename...', 'Copy...', 'Info', 'Remove', and 'Inspect' in the bottom left corner of the 'Sound tiger' window.

Object:
sound

Method menu
for
sound objects

Creating a New Object

The image shows two overlapping windows from the Praat software. The top window is titled 'Praat objects' and contains a list of objects with 'Sound tiger' selected. A menu is open over 'Sound tiger' with options: 'Sound help', 'Edit', 'Play', and 'Draw...'. An arrow points from the text 'Object: sound' to the 'Sound tiger' entry. Another arrow points from the text 'Method menu for sound objects' to the menu options. The bottom window is titled 'Sound tiger' and displays a waveform and a spectrogram. A red shaded region highlights a segment of the audio, with time markers at 0.071819, 0.161977 (6.174 / s), and 0.233795. The spectrogram shows frequency components up to 500 Hz. A blue line tracks the pitch contour. At the bottom, it shows 'Visible part 0.292938 seconds' and 'Total duration 0.292938 seconds'. A label 'Extract sound selection' points to the red shaded region. At the bottom left of the 'Sound tiger' window, there are buttons: 'Rename...', 'Copy...', 'Info', 'Remove', and 'Inspect'.

Creating a New Object

The image shows two overlapping windows from the Praat software. The top window is titled 'Praat objects' and contains a list of objects with 'Sound tiger' selected. A context menu is open over 'Sound tiger', showing options: 'Sound help', 'Edit', 'Play', and 'Draw...'. An arrow points from the text 'Extract sound selection' to the 'Play' option. The bottom window is titled 'Sound tiger' and displays a waveform and a spectrogram. The waveform shows a signal with a red shaded region between time markers 0.071819 and 0.233795. The spectrogram below it shows frequency components up to 500 Hz, with a blue line representing the pitch contour. Annotations include 'Object: sound' pointing to the 'Sound tiger' object in the top window, and 'Method menu for sound objects' pointing to the 'Play' option in the context menu.

Praat objects

Praat New Read Write Help

Objects:

Sound tiger

Sound help
Edit
Play
Draw...

Extract sound selection

Object: sound

Method menu for sound objects

Sound tiger

File Edit Query View Select Spectrum Pitch Intensity Formant Pulses Help

0.071819 0.161977 (6.174 / s) 0.233795

0.2487
0
-0.2539
5000 Hz

500 Hz

155.34 Hz
75 Hz

0 Hz

0.071819 0.161977 0.059142

0.000000 Visible part 0.292938 seconds 0.292938

Total duration 0.292938 seconds

all in out sel Group

Creating a New Object

The image displays two windows from the Praat software interface. The top window, titled "Praat objects", has a menu bar with "Praat", "New", "Read", "Write", and "Help". Below the menu is a list of objects: "Sound tiger" and "Sound untitled". A context menu is open over "Sound untitled", showing options: "Sound help", "Edit", "Play", and "Draw...". An arrow points from the text "Extract sound selection" to the "Play" option. The bottom window, titled "Sound tiger", has a menu bar with "File", "Edit", "Query", "View", "Select", "Spectrum", "Pitch", "Intensity", "Formant", "Pulses", and "Help". It displays a waveform and a spectrogram. The waveform has a red selection box from 0.071819 to 0.233795 seconds. The spectrogram shows a blue pitch contour and red formant markers. The text "Object: sound" and "Method menu for sound objects" are annotated with arrows pointing to the "Sound tiger" window. The text "Extract sound selection" is annotated with an arrow pointing to the "Play" option in the context menu.

Saving a Praat Object

Save file

Praat objects

Praat New Read Write Help

Objects:

Sound tiger

Sound untitled

Sound help

Edit

Play

Draw...

Query -

Modify -

Annotate -

Analyse

Periodicity -

Spectrum -

Formants & LPC -

Points -

To Intensity...

Manipulate

To Manipulation...

Synthesize

Convert -

Filter -

Combine sounds -

Rename...

Copy...

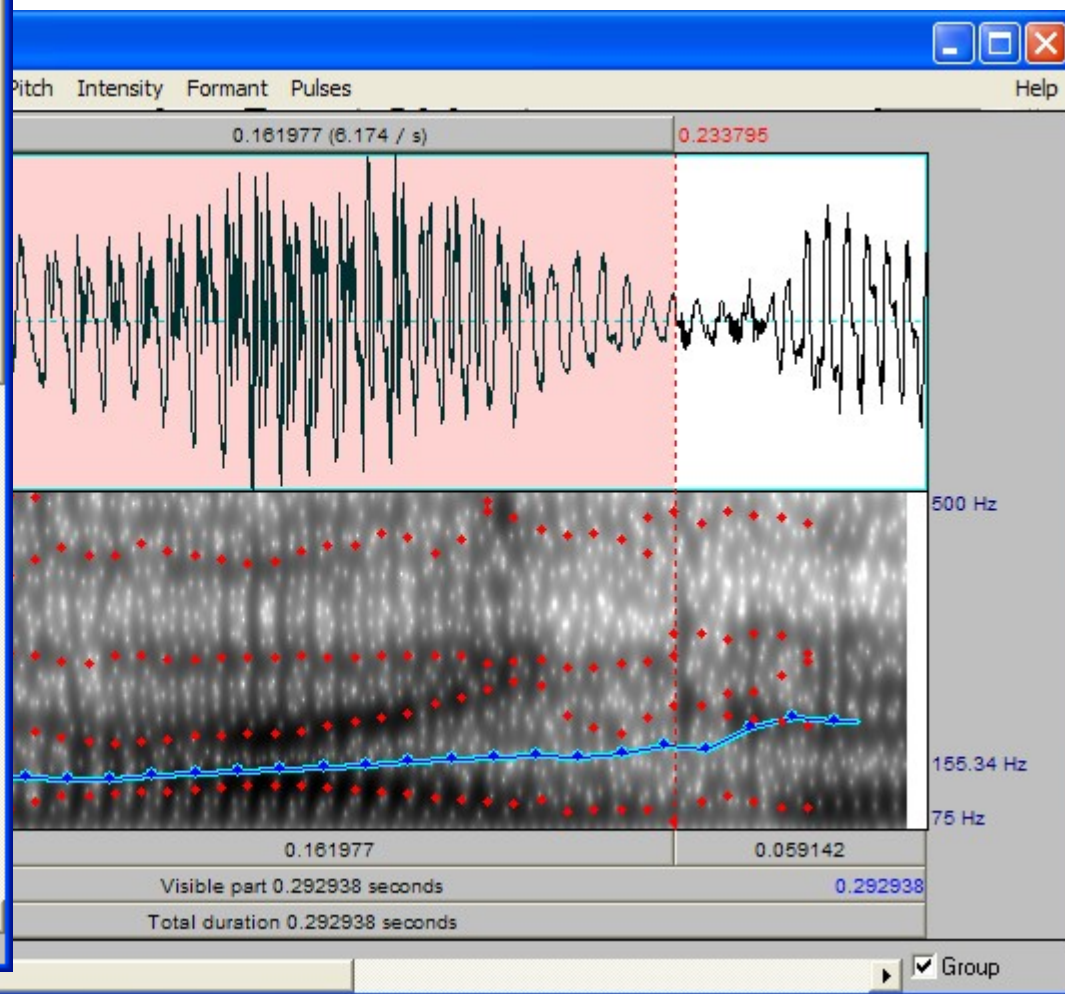
Info

Remove

Inspect

Object:
sound

Method menu
for
sound objects



Saving a Praat Object

Save file

The image shows two windows from the Praat software. The left window, titled 'Praat objects', has a menu bar with 'Praat', 'New', 'Read', 'Write', and 'Help'. Below the menu is a list of objects: 'Sound tiger' and 'Sound untitled'. A menu of actions is displayed to the right of the object list, including 'Sound help', 'Edit', 'Play', 'Draw...', 'Query -', 'Modify -', 'Annotate -', 'Analyse' (with sub-items: 'Periodicity -', 'Spectrum -', 'Formants & LPC -', 'Points -', 'To Intensity...'), 'Manipulate' (with 'To Manipulation...'), and 'Synthesize' (with 'Convert -', 'Filter -', 'Combine sounds -'). At the bottom of this window are buttons for 'Rename...', 'Copy...', 'Info', 'Remove', and 'Inspect'. The right window shows a spectrogram with a waveform above it. The waveform has a red shaded region from 0.161977 to 0.233795 seconds. The spectrogram has a blue line representing a formant track. The y-axis is labeled with frequencies: 500 Hz, 155.34 Hz, and 75 Hz. The x-axis has markers at 0.161977 and 0.059142. At the bottom, it shows 'Visible part 0.292938 seconds' and 'Total duration 0.292938 seconds'. A 'Group' checkbox is checked at the bottom right.

Object: sound

Method menu for sound objects

Annotating a speech signal

1. TextGrid TEM-all-amp-mono-normal-cropped

File Edit Query View Select Interval Boundary Tier Spectrum Pitch Intensity Formant Pulses Help

The spectrogram shows a blue line representing the fundamental frequency (F0) of the speech signal. A red dashed line is drawn at 3889 Hz. The TextGrid below shows the following annotations:

1	-	117	132	155	127	142	144	116	124	137	109	125	-	F0 (209 / 233)											
2	-	l	!h	h	l	!h	h	l	!h	h	l	!h	-	Pitch (233)											
3	-	L	H	L	L	H	L	L	H	L	L	H	-	Tone (233)											
4		k	O	d	O	N	a	r	i	k	e	J	a	z	I	w	u	r	o	t	a	s	i	-	Phoneme (427)
5	-	kO	dO	Na	ri	ke	Ja	zI	wu	ro	ta	si	-	Syllable (233)											
6	-	kOdONarike					JazI			wuro		ta	si	-	Word (139)										
7	-	18c. kOdONarike JazI wuro ta si?											-	Utterance (77)											
8	-	rire comme si le roi n'était pas mort!											-	Gloss Fr (77)											
9	-												-	Comments (77)											

53.688130 53.688130 2.000030 Visible part 2.001972 seconds 55.690102 2.610601 Total duration 58.300703 seconds

all in out sel bak Group

Praat Output

Praat Output

- Praat produces information which can be stored in a file.
- The file contents are not normally seen by the user.
- However, they can be seen and used for further analysis:
 - Excel, OpenOffice Calc
 - Praat scripting
 - Shell, Perl, Python scripting
 - Time Group Analyser online tool

```
intervals [1]:  
  xmin = 0  
  xmax = 0.3559744193778952  
  text = " _"  
intervals [2]:  
  xmin = 0.3559744193778952  
  xmax = 0.500147057910385  
  text = "ta"  
intervals [3]:  
  xmin = 0.500147057910385  
  xmax = 0.614452757446077  
  text = "la"  
intervals [4]:  
  xmin = 0.614452757446077  
  xmax = 0.8853950267508599  
  text = "sin"  
intervals [5]:  
  xmin = 0.8853950267508599  
  xmax = 1.096059981756913  
  text = "Ge"  
intervals [6]:  
  xmin = 1.096059981756913  
  xmax = 1.5079951315848832  
  text = " _"
```


Exercises

- Install Praat: <http://www.fon.hum.uva.nl/praat/>
- Make a folder “Praat” on your Desktop, for your recordings and annotations
- Record a speech signal:
 - menu New>Record mono sound
- Listen to individual words and syllables.
- Examine different displays (spectrogram, pitch track).
- Get highest, lowest and average frequencies
- Annotate the syllables, producing a ‘TextGrid’
- Save the Sound and TextGrid files with useful names
- Examine the TextGrid file using a text editor
 - make an Excel table in Excel, labels and times in columns
 - calculate durations (time differences)
 - calculate average duration

Extracting information about the speech signal

Tasks for next time:

- Plan a speech recording scenario:
 - Which questions are to be answered using this recording?
 - Which speech style or genre: reading / spontaneous?
 - Which speaker parameters: gender / age?
 - Which recording equipment, location?
- Record a short speech utterance with Praat
- Save the speech file with a systematic file name
- Read the speech file into Praat
- Open an “Edit” window
- Select sections of the signal and listen carefully for tones
- Experiment with the menu options to show
 - pitch / spectrum / formants
- Experiment with pitch:
 - modify the analysis range for pitch