Lab Assignment #02 Measuring sound waves

Due M Aug 29 at 11:15am on Sakai (Tests & Quizzes)
20 points total

Lab session: Part of class on F Aug 26 is a lab session for this assignment. You will have an opportunity to collaborate with classmates, get help with Praat, and ask questions.

Purpose

This assignment provides an opportunity to:

- Explore **real-world examples** of one way in which languages use **fundamental frequency**, abbreviated ***f*0** or **F0**: namely, **tone**
- Practice **calculating *f*0** directly, using the **waveform** of a sound file in Praat
- Learn how to use the Praat **spectral slice** (=**power spectrum** in *AAP*) function to analyze the **components** of a complex wave and **measure *f*0**
- Learn how to turn on and use the **"pitch" tracker** (strictly speaking, an *f*0 tracker!) in Praat, and how to check for **pitch-tracker errors**

Task

A. Prepare to complete the lab assignment

- **Download** the following files from the "Lab assignments" page and save them on your computer, and then open them with Praat. You will probably need to right-click on the link and choose "Download as...", "Save as...", "Save link as...", or something similar.
  `[lab02_thai.wav]`
  `[lab02_thai.TextGrid]`

B. Answer questions (1)–(11) directly in Sakai

- Go to Tests & Quizzes, "Lab 02 | Measuring sound waves" (may be saved; no time limit)
- Reminder: Use Praat handouts #2–4 to help you navigate in Praat, open sound files and TextGrids, play sounds, and take screenshots

Part 1: Calculating *f*0 from a waveform

Highlight the `[lab02_thai.wav]` and `[lab02_thai.TextGrid]` files in the Praat Objects window, and then click on the `View & Edit` button in the right-side menu to open them together in a TextGrid window.

Thai is a **tone language**: a language in which two words can be distinguished by their tones (pitch levels or contours) alone. The five words in this sound file have the same consonant and vowel, and differ only in their *f*0 contours. Play the whole sound file to hear the five tone categories of Thai. The traditional labels for these tone categories are *low, mid, high, falling,* and *rising*.

Click on ‘field’ in the text tier to select that word, and then click the `sel` button at the bottom of the TextGrid window to zoom in so that your selection fills the view. Now try selecting small pieces of the word and listening, to determine where the [a] sound is in this waveform.

(1) Select pieces of the sound file and zoom in (and out) on the [a] until you get the TextGrid window to show about 5 cycles of the [a] sound wave. **Take a screenshot, as you did for Lab #1, and upload this to the Lab 2 T&Q on Sakai.** (Please remember to include your name in the filename.) *PDF is preferred, but image files in .jpg, .png, or .gif formats are also acceptable; please do not upload a .doc(x) file.*

(2) Using information from **near the midpoint of the waveform** of [a] that is visible in the TextGrid window, **determine the period (T) of this [a] sound wave**.

- Hints: This is a *waveform* (graph of a sound wave as produced by a microphone). What physical property is shown on the x-axis? What information does Praat give you on the x-axis if you *click at one spot* in the waveform? If you *click and drag to highlight a selection* of the waveform?

(3) Using your answer to (2), **calculate the (fundamental) frequency (f_0) of this [a] sound wave**.

Part II: Praat spectral slices and components of complex waves

As we have seen, a complex wave can be described by describing the amplitudes and frequencies of the simple waves that are its **components**. A **power spectrum**, sometimes just called **spectrum**, is a graph that displays each component as a line (or a “wide line”, somewhat like an upside-down icicle). Note the values represented on the **axes** for a spectrum: **amplitude** on the y-axis (the same as for a waveform), but **frequency**, not time, on the x-axis.

In Praat, a power spectrum is called a **spectral slice**. Here is how (one way) to make one:

- In the TextGrid window, zoom out until the whole word ‘field’ is visible again.
- In the top menu of the TextGrid window, click on `Spectrogram > Spectrogram settings`. A box with various preset values should appear. In the field labeled `Window length (s)`, it should say 0.005. Change this value to **0.05** and click `OK`.
- Look at the spectrogram—it should have changed noticeably, so that now you see many thin horizontal stripes. This is a **narrow-band spectrogram**. (The default view with ‘window length’ at 0.005 s is a wide-band spectrogram. We will talk more about the difference between the two kind of spectrograms a little later in the course.)
- Place the cursor approximately in the center of the [a] part of the word ‘field’. Then, on the top menu in the TextGrid window, click `Spectrogram > View spectral slice`. Go back to the Praat Objects window and you should now see a **Spectrum object** with the name `lab02_thai_(some value)`. Click on this Spectrum object to highlight it, and then click on `View & Edit`.

(4) You now have a Spectrum window. (You can make it bigger or smaller as needed by dragging a corner of the window.) **Take a screenshot of this Spectrum window and upload it to the Lab 2 T&Q on Sakai.** (Please remember to include your name in your document and in the filename.) *PDF is preferred, but image files in .jpg, .png, or .gif formats are also acceptable.*

Try clicking on the spectral slice you created for question (4). You can click on the spectrum in the Spectrum window to get amplitude (in dB, decibels) and frequency (in Hz) values for any point on the graph. Try it.

Your spectral slice consists of a series of "wide lines" or "upside-down icicles". Each of these corresponds to the lines on the power spectrum in Fig 1.7 of *AAP* (p 16)—an infinitely long and unchanging complex wave would have actual lines in its spectrum, but the shorter the duration of the wave, the wider the "lines" get. Therefore, each of these "wide lines" represents one **component** of the complex sound wave for this [a]. Click at the tip of an "icicle" (at its *maximum* amplitude point) to get amplitude and frequency information for that component. Try it.

(5) There is probably a "half icicle" very close to zero on the frequency axis (x-axis) in your [a] spectrum—ignore this; it is not actually a component. Disregarding this near-zero line, **what is the frequency of the lowest component in this spectrum**? (Click on the spectrum to get the value. You will probably want to select an interval on the x-axis and zoom in with `sel` first, so that you can click more accurately.)

(6) Now click to find the **frequencies of the next four components**, up to component #5.

(7) Allowing for a little measurement error, based on the first five components, **what would you estimate the f_0 of this [a] to be, and why? How does your answer compare to your answer to question (3) above?**

Part III: Using, and checking, the pitch tracker

Praat has a tool that estimates f_0 from a sound wave, typically known as a **pitch tracker**. You can turn it on in the top menu of the TextGrid window by clicking on `Pitch` and selecting `Show pitch`; a blue line should appear superimposed on the spectrogram. Note that **the frequency scale for the pitch tracker is shown on the *right* side of the window**—the frequency scale on the left side is for the spectrogram.

Pitch trackers are convenient, but they can sometimes be disrupted by particular characteristics of a sound file. This is why it is a good idea to check pitch-tracker results against other ways of measuring f_0, especially when the pitch-tracker results look strange.

Make sure the pitch tracker is turned on. Now change the settings to expand the relevant pitch range (this speaker has a low voice): From the `Pitch` menu, select `Pitch settings`. The default settings for `Pitch range (Hz)` are 75.0 to 500.0. Change the minimum value to 55.0, and the maximum value to 200, and click `OK`. Changes in f_0 should now be easier to see on the pitch track.

(8) Click on the sound at a point somewhere in the sound file where, based on what you hear, the pitch tracker is making an error. (*Hint:* Look for large jumps in the pitch track that do not correspond to large pitch changes in the audio.) What is the **timepoint** you have clicked on, and **what (wrong) f_0 value** is the pitch tracker reporting? (Remember to read pitch-tracker frequency values from the ***right*** side of the TextGrid window, not the left!)

(9) Find **another method of determining the f_0** at or close to the point you have clicked, other than using the pitch tracker. **Explain** what method you chose, and **report the f_0 value** you have found. Now **discuss your results:** Did you find the same f_0 value as the pitch tracker? If not, does your hand-measured value seem more plausible? Did you encounter any

problems or questions in trying to determine the f_0 by hand here?

Criteria for success

This lab assignment is worth a total of 20 points. Each question is worth 2 points except for (9), which is worth 4 points. Points will be awarded for accuracy, and partial credit will be given where appropriate. Some of the parts will be automatically graded on Sakai, but I will double-check everything by hand in case of Sakai errors.