Lecture 4 (8 November 2012)

Assignment 7: read Hayward chapter 4.

7a. What are the F1, F2 and F3 for a child with a vocal tract length of 8.5 centimetres who pronounces a [ə]?

7b. Record a small song (no longer than 10 seconds) whose melody we all may know, and turn this into a thing that sounds as the same song sung only on the vowel [a]. In class, the others can then guess what you sung.

Here is how it goes. You first record the song in the normal way in Praat. You then click **View & Edit**, and then switch on **Show pulses** in the **Pulses** menu. Zoom in on the blue pulses and look at where they are: this is Praat's thoughts on when your vocal folds collide. With **Extract visible pulses**, you can then copy the whole pulse train to the Objects window. There you can choose **To Sound (phonation)...**. In the window that appears, you just click OK, without minding about all the settings. As a result, you will see a Sound, and if you **View & Edit** that, you can see the glottal waveform (the 'phonation', the 'source'). If you play it, it will sound like you're singing that song with your head screwed off.

The glottal waveform does not look much like the ones you find in Hayward's figures 2.16a or 4.6a. This is because the radiation (conversion of airflow in liters per second to sound pressure in Pascal) has been included in this sound (you know, the –6 dB per octave).

It also does not much sound like an [a] yet. To make it do so, you *filter* it with the formants of [a] that you can find in figure 4.7b. From that figure, find the values of F1, F2, F3, and F4 (please write them up) and apply all four of them to the source sound in sequence. That is, select the source sound, then choose **Filter (one formant)...** from the **Filter** menu, typing your F1 in the **Frequency** field and leaving the **Bandwidth** at 100 Hz (the bandwidth is the width of the spectral peak); this will give a new Sound that has one resonance (listen to it). Then filter this resulting sound again with F2, then the result with F3, and so on. The result should sound like somebody singing your tune on the vowel [a] only. Please send me that sound.

From the end product, create a Spectrum (with **To Spectrum...**) and draw that for the range from 0 to 4000 Hz (and send it in). Does it look like figure 4.7b or more like 4.7c? And what are the remaining slight differences?

For class, please think a bit about this 'quantal theory' of Kenneth Stevens.

Assignment 8: factoring out the variation between speakers. Confidence interval of the effect size.

An (unnormalized) *effect size* is the difference in a certain quantity (e.g. the duration of a vowel) between two conditions (e.g. whether the vowel was intended as /a/ or as /i/).

In this assignment you are going to compute how much /a/ is longer than /i/, and how much longer /i/ is than /ɪ/. The answer is not just a single value ("/a/ is 55.6 ms longer than /i/"), but a range of values: "with 95 percent confidence I can say that /a/ is between 41.8 and 71.8 ms longer than /i/".

First you note again how good it was that /a/ and /i/ were spoken by the same set of speakers. The /i/ of the slowest speaker might be longer than the /a/ of the fastest speaker, so in themselves the duration distributions for /i/ and /a/ may overlap. But all of you found that every speaker's /a/ was longer than his /i/, and that is what matters.

8a. In the table that you created before, you create a column for the difference between the durations of /a/ and /i/. You could call that column "a–i" or so. Fill the column with 8 values, one for each speaker. Also make an analogous column "i–I". Include the table in your Word document.

8b. Compute the mean μ_{ai} of the eight difference values in the column "a–i". Show how you did it (choose any of the four known methods). Please also compute (from the mean durations that you computed before) the difference between the mean duration of /a/ and the mean duration of /i/. What do you notice?

8c. Now that you know the *mean* "a–i" difference, you can compute the *standard deviation* σ_{ai} of the eight "a–i" differences. Show how you did it.

So now you know the three important statistics of the observed distribution of a-i differences: the number of measurements (N=8), the mean (μ_{ai}), and the standard deviation (σ_{ai}). These are sufficient to compute the confidence interval of the a-i difference.

8d. First you compute the *standard error* of the a–i difference. The standard error is a measure for how accurately the true mean of the distribution is known if we know the observed mean of the eight values. Of course it depends on the standard deviation: if the standard deviation is large, there is a lot of spreading in the data and the mean must be relatively poorly known. But it also depends on the number of measurements: the more participants you have tested, the closer your observed mean will be to the true mean of the distribution. The formula for the standard error is: the standard deviation divided by the square root of the number of measurements, i.e. $\sigma_{ai}/\sqrt{N}$. Compute this value (don't forget to give the right units).

So now we know the standard error, i.e. the accuracy (in milliseconds) with which the true a–i mean is known. An estimate of the true mean (as you can expect) is the mean that you computed from the eight values. The observed mean is obviously greater than zero (all /a/ measurements are longer than the corresponding /i/ measurements), but by how much? We now compute the *t-value*.

8e. You divide the observed mean μ_{ai} by the standard error. What is the result?

The result of this computation (the *t*-value) tells us how many standard errors fit in the space between zero and μ_{ai}. The more standard errors fit in this space (i.e. the higher the *t*-value is), the more likely it is that the true mean a-i difference is different from zero. As a rule of thumb one can say that if the *t*-value is greater than approximately 2.0, then the true mean is indeed reliably different from zero (with 95% confidence). The precise reliability, however, depends a bit on *N*, so in practice we have to do the computation.

8f. In Praat's calculator you can compute `invStudentQ (0.025, N-1)`. This should give a value near 2.0, but not exactly 2.0. How much is it for *N*=8?

That was easy. This number (around 2.0) is the number of standard errors that make up the 95% confidence interval around the observed mean. The "0.025" comes from what is left outside the confidence interval: 5 percent in total, so 2.5 percent on each side.

8g. Finally the computation of the confidence interval. Multiply the standard error with the value of 2.0 or so that you just computed. The result (what is it?) is the maximum distance that the true mean will be away from the observed mean (with 95% confidence). If you subtract this distance from the observed mean, you get the lower limit of the true mean. What is it? If you add the distance to the observed mean, you get the upper limit of the true mean. What is it?

8h. Now summarize the whole procedure in the following phrase: "with 95% confidence I can say that the true mean difference between the durations of /a/ and /i/ lies between [fill in the lower limit here] and [fill in the upper limit here] ms."

8i. Does the confidence interval include zero? If not, what does that mean?

8j. Similarly, give the confidence interval of the difference between /i/ and /ɪ/. Is one of these two vowels reliably longer than the other according to the *t*-test?

8k. Do the significance results of the *t*-tests match those of the sign tests you did before?