

Linguist/SphHrng 286: Analyzing the Sounds of Language Spring 2006

Call number: XXXXXXXX; 5 Credit Hours

Class meetings: Tuesdays, Thursdays 1:30-3:18, Room XX, Yyyyy Hall

Instructor: Mary Beckman

Office hours: Tuesdays at 4:00, Thursdays at 11:00, and by appointment;

Contact information: office in room 07 Oxley, tel. 292-9752, e-mail mbeckman@ling.osu.edu

Course web page: <http://ling.osu.edu/~mbeckman/LxSHS286>

Course description: Speech is a very complex behavior. Saying even a simple sentence such as *Luce hit another home run*. takes more motor coordination than actually hitting the home run. Yet we are not surprised when we talk with a child as young as three or four who can say a sentence like this very naturally and intelligibly. Understanding speech also is a very complex skill. We take it for granted that we can listen to this sentence and correctly identify the first word as *Luce* rather than as *Ruth* and the second word as *hit* rather than as *heat*. Yet speakers of some other languages besides English have a great deal of difficulty distinguishing these two pairs of words. How can we describe the similarities between the consonants in *Luce* versus *Ruth* or between the vowels in *hit* versus *heat* that make these two word pairs difficult for speakers of some other languages? Can we analyze speakers' behavior in saying words and sentences, and listeners' behavior in listening to them, well enough to be able to construct models that predict which sounds will be difficult for children acquiring their first language or for adults acquiring another language? In this course, we will introduce pertinent ideas and results from research in the various disciplines that have contributed to our understanding of the sounds of language. We will introduce some of the quantitative analytical tools that are used in the phonetic sciences, and do several experiments in class, to give a flavor of the diverse research methods that speech scientists have developed to try to determine how speech is produced and perceived by humans. We will also take a brief look at how speech engineers have applied this knowledge to develop computer speech synthesis systems.

Course objectives:

Participation in this course should lead to:

1. an understanding of the general character of speech sounds, and of how they are produced and perceived,
2. a familiarity with some of the measurements and models that are used to study speech sounds,
3. an understanding of some basic concepts in probability theory and statistics, and
4. an appreciation of how probability theory and statistics can be applied to understanding speech sounds.

Textbooks:

- Peter Ladefoged. (2004). *Vowels and Consonants: An Introduction to the Sounds of Language*. 2nd edition. Blackwell. ISBN 1-4051-2459-8 [The sound files on the CD that accompanies this book are also available at <http://hctv.humnet.ucla.edu/departments/linguistics/VowelsAndConsonants>.]
- *Notes on Probability and Statistics in Speech Science*. [COPEZ course packet.]

Coursework and grading: Letter grades will be assigned, based on class preparation and participation, as described in the Weekly Schedule below. The grade will be based on three components, as follows:

30%	reading and other preparation for primary data analysis (7 homework assignments)
30%	primary data analysis and review quiz in some weeks (6 quizzes)
24%	primary data analysis and written report in other weeks (4 reports)
16%	final report due at the time of the final examination, which will be as determined in the schedule of classes.

A more detailed description of each of these three components is given below.

Reading and preparation for data analysis (7 homework assignments): Almost every week, you will need to read a chapter or two from one or both of the textbooks, and think about some issue related to the reading. In most weeks, you will be given the opportunity to demonstrate to me that you have read and thought about the issue by doing a short on-line homework. The homework will consist either of four or six questions with multiple-choice answers, or of a two or three questions asking for short one-sentence answers. In either case, it will be possible to add a one- or two-sentence explanation if you wish to justify or elaborate on your answer. In order to get credit for doing the homework, you must complete it **by 8:00 a.m. on Monday of the week that we cover the material in the reading.** (30% of grade)

Primary data analysis: Each of the issues that we will be covering will require you to analyze data together in-class or individually outside of class. The data will be of two types: (1) data that have already been gathered in research by established speech scientists, and (2) data that you will collect together in class using yourselves as speakers and listeners. Relating the data to the issue will require numerical reasoning and visualization techniques that are useful in many other areas of science (and of life!) and I will introduce these techniques in the first class session of the week. After we have covered the material, you will demonstrate that you understood the data and the data analysis methods in one of two ways.

1) **Post hoc review quiz (6 quizzes):** In six weeks of the quarter, you will demonstrate that you understood the data analysis that you did by doing a short online review quiz, consisting of questions similar to those on the homework. To get credit for the quiz, you must complete it **by 8:00 a.m. on Monday of the following week.** (30% of grade)

2) **Written report (4 reports):** In each of the other four weeks of the quarter, you will turn in a short written report. Instructions for doing the report will be made available online. Typically you will be instructed to summarize the data analysis by making a specific type of graph and by calculating some relevant numbers (such as the means of two samples, or a set of regression coefficients for a regression model related two sets of numbers), and then to write a short paragraph relating the graph and numbers to some aspect of our discussion in class. In order to get credit for a report, you must turn it in **by the beginning of the first class session of the week after we have covered the topic.** (24% of grade)

Final report: During the sixth week of classes, you will meet individually with me to discuss and develop a small experiment extending one of the four topics examined in the reports. You will execute and write up this experiment in the format of a research paper (Introduction, Methods, Results, Discussion, References, Figures, Tables). This report will be due at the time of the final examination as listed in the schedule of classes.

Instructor's philosophy: By my offering this course and your enrolling in it, we have instigated the following agreement. I am responsible for designing and implementing a course that engages your intellect and for executing assessment procedures that fairly test your knowledge of the course material. You are responsible for coming to class, doing the readings on time, participating in class activities and discussions, independently completing the assigned homework exercises, quizzes, and reports. That is, your part of the agreement is not just to memorize or to do the analyses by rote, but to actively think about the questions that we tackle through the weeks. I pledge to do my part to make this class interesting, challenging, and thought-provoking. By your presence, you are promising to do your part, too.

Students with disabilities: If you need some accommodation for the impact of a disability, please make an appointment to meet with me as soon as possible to discuss the course format in terms of anticipated needs and to explore potential adaptations of the format to those needs. I rely on the Office of Disability Services to verify the need for accommodation and to help develop accommodation strategies. Students with disabilities who have not previously contacted the Office of Disability Services are encouraged to do so, by looking at their web site (<http://www.ods.ohio-state.edu>) and calling them for an appointment (tel. 614 292-3307).

Academic misconduct: I am required by my contract with the university to report suspected cases of academic misconduct to the Committee on Academic Misconduct. The most common form of misconduct

is plagiarism. Remember that any time you use the ideas or the statements of someone else, you must acknowledge the source in a citation. This includes material that you found on the Web. The University provides guidelines for research on the Web at <http://gateway.lib.ohio-state.edu/tutor>.

Collaboration: I encourage collaborative learning such as study groups and discussion of assignments out of class. However, all homework assignments, quizzes, and reports must be your own individual effort.

Weekly schedule

Week 1 (March 28 & 30)

Puns, rhymes, and pig Latin: things we already know about some speech sounds.

In every culture that has been studied to any extent, there are word games that require a great deal of very sophisticated knowledge about the sounds of the language of that culture. In the first class session, we will use some of these ways of playing with words as a first analytical tool, to see how much we already know about the sounds of English. We will also begin to learn how to use Praat, a free speech analysis tool that you can download off the web from <http://www.praat.org>.

How many sounds can a word contain? (Or should we say “how much sound”?)

The word games that we play on the first day should give us a sense of the kinds of word-internal units that we can manipulate and count. In the second class session, we will think about the relationship between word length measured in terms of those countable units and word length measured another way. We will measure some words both ways and then look at the two different sets of measurements using histograms and various summary statistics. We will compare our measurements using the second method for different words that we counted as being the same length the first way, to see if we can guess at things that will affect our measurements. We will also compare different people’s measurements of the same word, to see how we can sort out measurement “error” from other sources of variability.

Class preparation. There is no reading this week, and no homework. We will be gathering some data from ourselves in the first class session, and during the second class session I will introduce some visualization techniques that will help us analyze those data. Prepare for the second class session by downloading Praat onto your home computer, or onto some portable medium that you can take with you to use in a computer lab. Begin to familiarize yourself with this program so that we can begin doing the data analysis assignment during class, where you can ask me any technical questions about the program and its use.

Quiz 1. The quiz will involve some questions about the data that we gathered in the first class session and about the measurements that we gathered in the second class session. You must complete it by 8:00 a.m. on Monday of week 2 to get credit for it.

Week 2 (April 4 & 6)

Can we know how many sounds there are in a word from a language we don’t know?

This week we will broaden our understanding of the two ways of measuring word length in three ways. First, we will learn to use histograms and averages to summarize our measurements visually and numerically. Second, we will listen to and look at words in another language, to see whether it is possible to measure word length in a language that we don’t know using either of the two kinds of measures of word length that we introduced in the week before. Third, we will learn to use two more data analysis tools (scatterplots and regression) for visualizing and evaluating the **relationship** between the lengths of words measured in the two different ways.

Class preparation (homework 1). Read chapter 1 of *Notes on Probability and Statistics in Speech Science* (on “Data types and data distributions”) and answer the online homework questions about averages and the relationship between sample means and sets of real-world measurements. Do this by 8:00 a.m. on Monday of week 2 to get credit for doing the homework. Also read chapter 2 of *Notes on Probability and Statistics in Speech Science* (on “Correlation and regression”).

Report 1. In this report, you will make two histograms to show the distribution of measured values of word lengths that we gathered last week, and then you will calculate averages across various subsets of the measurements and mark these averages on the histograms, to get an appreciation of the relationship between the graphical and the numerical ways of summarizing the measurements. You will then make a scatterplot and use a regression model to see whether words that are longer by the first way of counting are generally also longer by the second way of counting. Turn in your report at (or before) the beginning of the first class meeting in the Week 3.

Week 3 (April 11 & 13)

Does the number of vowels and consonants constrain word lengths?

In preparing for class this week, you will be asked to think about these questions: What is the typical length of a word in English and how many vowels and consonants does English have? Do you know any other languages well enough to estimate these two numbers for that language? Is the relative word length and relative number of distinct vowels and consonants similar to English, or different? If it is different, is the relationship what you would expect from the cognitive pressures that Ladefoged identifies on pp. 2-3? In the class sessions and data analysis assignment this week, we will explore this relationship across some of the languages for which we have reliable estimates. We will extend the tools that we used last week, by showing how to use a log transform of one of the relevant estimates to be able to fit a meaningful regression.

Class preparation (homework 2). Read chapter 3 of *Notes on Probability and Statistics in Speech Science* (on “Ratios, logarithms, and other transforms”) and chapter 1 of Ladefoged (on “Sounds and languages”). Think about what Ladefoged says (on pp. 2-3) about the relationship between the possible lengths of words in a language and the number of vowels and consonants. What is the typical length of a word in English and how large is the set of distinct vowels and consonants in the language? Can you think of any good way to estimate these two numbers? The homework will involve answering these questions, and you must do it by 8:00 a.m. on Monday to get credit for it.

Quiz 2. In this quiz, you will be answer some multiple-choice questions about (other-than-straight-line) relationships between two variables. Complete the quiz by 8:00 a.m. on Monday of Week 5.

Week 4 (April 18 & 20)

Pitch and rhythm.

What are the physical cues that signal the difference between pairs of sentences such as:

(Figure 2.8 in Ladefoged) When danger threatens your children, call the police

(Figure 2.9 in Ladefoged) When danger threatens, your children call the police

That is, what exactly are we doing when we “say the comma” in these two sentences? I will bring recording equipment to class, so that we can look together at your own productions of these sentence and also of some other pairs that you think of. We will review notions relevant for analyzing time-series data, and show how the statistical techniques that you mastered in the previous weeks help us derive data such as those plotted in Ladefoged’s figures. We will also play with some of the tools that speech engineers have used to model the differences between these two sentences, and try to change one sentence in the pair into the other.

Class preparation (homework 3). Read chapter 4 of *Notes on Probability and Statistics in Speech Science* (on “Time series data”) and chapter 2 of Ladefoged (on “Pitch and loudness”). Listen to the recordings on the CD that go with the figures in chapter 2, particularly to the files that go with the pair of sentences in Figures 2.8 and 2.9. The homework will ask you to come up with another pair, where the two sentences contain exactly the same string of words, but differ in meaning in a comparable way. If the pair that you think of does not differ in punctuation (as the two in Ladefoged’s figure did), you should try to describe the meaning difference in a short paraphrase. As always, the homework is due by 8:00 a.m. on Monday of the current week

Quiz 3. The quiz will include multiple-choice questions about such concepts as windowing, step size, and the auto-correlation technique that we covered in class.

Week 5 (April 25 & 27)

Introduction to vowels.

How many vowels are there in English, and how can we be sure of our answer to this question? How can we count the number of vowels in any language? This week, we will cover some of the analytic tools used to answer these questions. For example, you will learn how to find all of the words in a dictionary that might contain a particular vowel sound and then to use rhyming judgments and other types of word play in a way that will let you evaluate whether they do in fact all contain the same vowel. You will also learn about the “sound spectrogram” — a type of three-dimensional scatterplot that lets us visualize the relationships among three physical parameters that can be used together to examine and compare a speaker’s productions of syllables that may or may not contain the same vowel.

Class preparation. Read chapters 3 and 4 of Ladefoged (on “Vowel contrasts” and “The sounds of vowels”). The homework for this week will ask you to listen to the recordings on the CD that go with the tables, particularly the files that go with Tables 3.2 and 3.3. Go through the lists of words in the tables and try to figure out which of these two speakers has a vowel system more like your English vowel system. Once you have decided this, decide how many different vowel sounds you have. You should enter your two decisions and a sentence or two describing how you made the decisions by 8:00 a.m. on Monday of Week 6.

Quiz 4. The quiz will be multiple-choice questions about spectrograms and vowel inventories.

Week 6 (May 2 & 4)

The vowel space.

How can we measure the degree of similarity or difference between two different vowel sounds? For example, how can we evaluate the claim that “For speakers of BBC English, the vowel in *sing* is very like that in *sin*; but for many speakers of American English, ... it is almost identical with the vowel in *seen*”? This week, we will begin to learn about acoustic measures of vowel quality, and see how they correspond to perceptual distance among vowels. We will also see how we can represent each vowel as a distribution of instances in a two-dimensional space, and learn about some statistical tools to measure and evaluate distance in that space. These tools include tests related to the normal distribution, which help in evaluating central tendencies (i.e., average values) and variability around the centers of the distributions.

Class preparation (homework 4). Read chapter 5 of *Notes on Probability and Statistics in Speech Science* (on “Clustering algorithms and related techniques”) and chapter 5 of Ladefoged (on “Charting vowels”). As you read the assigned chapter of Ladefoged, think back to the claim that Ladefoged makes on p. 30 in chapter 3 that speakers of British English and American English differ in whether their pronunciation of the vowel in *sing* is more like the vowel in *sin* or in *seen*. How could you test this claim? The homework assignment is to write a short paragraph (three or four sentences) describing any test that you can think of. As usual, the paragraph is due by 8:00 a.m. on Monday.

Report 2. In this report, you will be asked to plot vowel formant data from two different speakers of the same dialect of American English and to apply one of the techniques that we learned about to assess similarities and differences between the plotted measurements for different vowels within a single speaker and for the same vowel across the two different speakers.

Week 7 (May 9 & 11)

Analyzing our personal vowel spaces.

This session continues the previous week’s topic, using the results of an online experiment that we will do together during the first class session of the week. We will compare our results, to see if we can figure out how our mental map of the vowel space works, and how these maps might differ from one language to another. We will also discuss results of an experiment comparing the mental vowel spaces of speakers of Ohio English with speakers of Athenian Greek. In discussing that experiment, we will learn about the analysis techniques (e.g., analysis of variance) that the experimenters used in analyzing these data.

Class preparation (homework 5). Read chapter 6 of *Notes on Probability and Statistics in Speech Science* (on “Calculating and analyzing variance”) and chapter 7 of Ladefoged (on “Acoustic components of speech”). The homework assignment is to answer some questions about variance.
Quiz 5: This will be a multiple choice questions on the design of multi-factor experiments like the one that we are doing together in class.

Week 8 (May 16 & 18)

How many consonants does English have and where can they occur?

This week we will apply the methods that we used to analyze the vowel inventory to the consonant inventory. We will explore the ways in which the auditory properties of different consonants are affected by their position. We will also explore the constraints on where different sounds can occur, which (in English) affect the consonants even more than the vowels. We will extend the tools that we use in Week 5, by learning about the Chi-square distribution and a related test that will let us evaluate whether some of the seemingly odd patterns that we find are due to chance.

Class preparation (homework 6). Read chapter 7 of *Notes on Probability and Statistics in Speech Science* (on “Frequencies, probabilities, and independence”) and chapter 6 of Ladefoged (on “The sounds of consonants”). Think about the examples that he gives in Table 6.1 to illustrate the three different places of articulation for nasals. Why does he use *ram*, *ran*, and *rang* instead of word-initial consonants as he did for the analogous contrast among *p*, *t*, and *k* in *pie*, *tie*, and *kite*? Can *p*, *t*, and *k* occur in word-final position? Try to think of a set of words in which the stops contrast after the same vowel as in *ram*, *ran*, and *rang*. Which of these six consonants can occur in word-medial position? Try to find a set of words that contrasts all of the consonants there. The homework assignment will involve answering some questions about where different consonants can occur.

Report 3: In this report you will apply the concepts in probability theory that we covered in class to some tables of phoneme frequencies in different positions in order to evaluate whether some of the seeming “holes” could be due to chance.

Week 9 (May 23 & 25)

The consonant chart.

As you do the reading for this week, think about what Ladefoged said (on p. 3 in chapter 1) about “gestural economy”. Some sets of English consonants illustrate this principle very well, whereas others do not. For example, there are voiced stops, voiceless stops, and nasals at each of the labial, alveolar, and velar places of articulation, but there is no velar fricative. How can we evaluate the strength of this principle relative to the principles of “articulatory ease” and “auditory distinctiveness” that Ladefoged also identifies as factors constraining consonant inventories? In order to try to answer this question, we will extend the tools that we use in Week 8, by learning about the Chi-square distribution and a related test that will let us evaluate whether some of the seemingly odd patterns that we find are due to chance.

Class preparation. Read chapter 8 of *Notes on Probability and Statistics in Speech Science* (on “Contingency tables and the χ^2 distribution”) and chapters 10 and 12 of Ladefoged (on “Making English consonants” and “Actions of the larynx”). In reading chapter 12, focus especially on sections 12.1 and 12.2. There will be no homework assignment, so that you have time to do the reading.

Report 4: In this report you will use contingency tables of consonants across languages and the χ^2 -test to evaluate the relative strength of “gestural economy” in different parts of the consonant chart.

Week 10 (May 30 & June 1)

Stringing together consonants and vowels.

So far we have focused on what vowel and consonant categories there are without thinking very much about how vowels and consonants affect each other when they are strung together. If there were no drastic interactions between adjacent sounds, we should be able to just cut out the chunk of the waveform corresponding to the *l* sound in an audio recording of a talker producing the word *black* to create a new

artificial “audio recording” of the word *back*. As Ladefoged points out, this will not work. Our question this week is: Why doesn’t speech production work that way?

Class preparation (homework 7). Read the first three sections of chapter 15 of Ladefoged (on “Putting vowels and consonants together”) and also chapter 8 (on “Talking computers”). Think about the example that he gives (in Figure 8.1) about possible sources of the three diphones in *black*. How many diphones do you need to synthesis the *mighty* in this way and what are the diphones? The homework will consist of multiple-choice questions that this topic.

Quiz 6: This quiz will be short answers to questions relating the material from this week to the material we covered in week 9.