Syllabus
LING 219 R: Advanced Phonology
Computational Modeling in Phonology and Beyond
Spring 2017
Tuesday 11–1, Boylston 303

Instructor

Aleksei Nazarov
Email: anazarov@fas.harvard.edu
Office: 314 Boylston Hall
Office hours: Monday 11-1 (subject to change!)
Course site: https://canvas.harvard.edu/courses/23912

Course Overview

This course will offer an overview of topics in computational learning as applied to both phonology and syntax. We will focus on approaches to inducing hidden variables (which include, e.g., syntactic category labels; syntactic and phonological parameters; underlying forms; and foot structure).

Previous computational background is not necessary: an introduction to R, a simple scripting environment that is also highly useful for statistics and data analysis, will be offered.

In keeping with the official name of the course ("Advanced Phonology"), we will look at models that are useful for phonology. However, each model is also equally applicable and useful in the domain of syntax, and I hope to provide a balance of phonological and syntactic examples.

Course Setup

Each 2-hour session of this course will be divided into two halves. In the first half, we will focus more on theory and motivation behind models, while the second half will focus will be a workshop session, where will put the ideas discussed in the first half into practice.

As the course unfolds, there will be more emphasis on student presentations of papers, and at the very end, we will do brief presentations of your final projects.

Assessment and Final Project

The main graded element of this course will be a final paper, which will contribute 50% of your grade. The final project should be a computational modeling project, although various options can be discussed, as long as they employ the skills you learn in this course.

Other graded elements will include in-class presentations of papers, which will contribute 35%, as well as minor programming assignments that will help you gain proficiency in the models that we will be learning about, which will contribute 15%. Papers for presentation in class will be chosen in consultation with students, and may or may not coincide with the background readings listed in the schedule below.

To summarize:

Programming assignments:	15%
Presentations:	35%
Final paper:	50%

Course Schedule (subject to change)

Week	Day	Topic	Background readings	Assignments
1	Jan 24	Introduction, Programming in R	Baayen 2003: ch 1,2,4	
2	Jan 31	Programming in R (cont'd)		
3	Feb 7	Markov Models (n-gram models)	Manning and Schutze 1999: ch 5-6	Assignment 1 due (programming/statistics basics)
4	Feb 14	Hidden Markov Models	Manning and Schutze 1999: ch 9-10	
5	Feb 21	Hidden Markov Models (cont'd)	Goldwater and Griffiths 2007	
6	Feb 28	Parameter models	Yang 2002: ch 2	Assignment 2 due (Markov)
7	Mar 7	Parameter models (cont'd)	Nazarov and Jarosz to appear	
8	Mar 14	**Spring Recess, no class**		
9	Mar 21	Maximum Entropy Models	Goldwater and Johnson 2003; Ratnaparkhi 1998: ch 5(-6)	Assignment 3 due (Parameters)
10	Mar 28	Learning serial derivations	Staubs and Pater 2016; Nazarov and Pater to appear	
11	Apr 4	Finding hidden structure through clustering	Shih 2014: ch 5; Nazarov 2014	Assignment 4 due (MaxEnt)
12	Apr 11	Iterative Learning/Historical Predictions	Staubs 2014: ch 2	
13	Apr 18	TBD		
14	Apr 25	Final paper presentations		