STAT 27400/37400: NONPARAMETRIC INFERENCE

Syllabus, Fall 2015

Nonparametric Inference is an advanced undergraduate and beginning graduate level course in modern nonparametric methods for statistical estimation and inference.

Nonparametric inference is about statistical methods and models that make weak assumptions. A typical nonparametric approach estimates a nonlinear function from an infinite dimensional space, rather than a linear model from a finite dimensional space. This course gives an introduction to nonparametric statistics, with a focus on density estimation, regression, confidence sets, orthogonal functions, random processes, and kernels. The course treats nonparametric methodology and its use, together with theory that explains the statistical properties of the methods.

Schedule

LECTURES	Tuesdays and Thursdays 10:30-11:50 am	SS 122

Contact Information

Instructor:

John Lafferty	Ry 257	lafferty@galton.uchicago.edu
Office hours:	Monday, 5:00 pm, Ry 257	

Course Assistants:

Yuancheng Zhu	Eck 8	yuancheng@galton.uchicago.edu
Office hours:	Wednesday, 6:00 pm, Eck 117	
Zhe Liu	Eck 8	zheliu@galton.uchicago.edu
Office hours:	Monday, 6:00 pm, Eck 117	

Prerequisites

Stat 24400 is required; alternatively Stat 22400 and exposure to multivariate calculus and linear algebra.

Course Structure and Grading

The course will have a standard lecture format. During a portion of some classes students will work problems in small groups. Assignments will be handed out and due every other week. Two short 15-minute quizzes will be given in class. These quizzes are not intended to be too challenging, and are a way reinforcing some of the basic concepts in the class. An in-class midquarter exam will be given.

There are two sections of the course. The sections will be graded using separate scales. The undergraduate section, Stat 27400, will have a final exam (schedule to be determined). The graduate section, Stat 37400, will require a final project in lieu of an exam. This project will require analysis of a data set, to be chosen by the student. The analysis should include a comparison of parametric against nonparametric methods. A 4-8 page report should be submitted, describing the data, the statistical problem, the methods used, and the findings.

Each of these components will be weighted as follows to determine a final grade:

- Assignments: 40%
- In-class midquarter exam: 25%
- Final exam/project: 30%
- In-class quizzes: 5%

Policy on Assignments

Assignments will be posted on Thursdays, and are due fortnightly in class. Assignments may be handed in up to three days late, with a 20%-per-day penalty. That is, an assignment passed in on Friday will have a 20% penalty, on Saturday a 40% penalty, and on Sunday a 60% penalty. Students are encouraged, but not required, to use LaTeX to prepare solutions to the assignments.

Assignments will have at least one problem that involves computation. Students are required to use the R programming language for solutions. R is the language of choice for statistical computing; it is easy to learn, and the course does not require previous exposure to the language. Tutorial introductions can be organized as needed. Free downloads of the language, together with an extensive set of resources, can be found at `http://www.r-project.org`.

Collaboration on homework assignments with fellow students is encouraged. However, such collaboration should be clearly acknowledged, by listing the names of the students with whom you have had discussions concerning your solution. You may *not* share written work or code—after discussing a problem with others, the final solution must be prepared and written by yourself.

Course Calendar

The course calendar and other materials will be posted on the course website, through the chalk system at `chalk.uchicago.edu`, and will be updated throughout the semester.

The schedule of topics, exams, and assignments follows.

Week	Date	Topic	Exams and Assignments
1	September 29	what is nonparametric inference?	
	October 1	nonparametric regression	assn 1 out
2	October 6	nonparametric regression	
	October 8	density estimation	
3	October 13	density estimation	
	October 15	normal means	assn 1 due, assn 2 out
4	October 20	SURE and minimax risk	
	October 22	statistical functionals	quiz 1
5	October 27	the bootstrap	
	October 29	minimax theory	assn 2 due, assn 3 out
6	November 3	shape-constrained estimation	
	November 5	—	midquarter exam
7	November 10	nonparametric Bayes	
	November 12	nonparametric Bayes	assn 3 due, assn 4 out
8	November 17	nonparametric Bayes	quiz 2
	November 19	wavelets	
9	November 24	nonparametric sparsity	
	November 26	Thanksgiving (no class)	—
10	December 1	review	assn 4 due
	December 3	reading period	—
	TBD		final exam

Textbook

The book used in the class is *All of Nonparametric Statistics*, by Larry Wasserman (Springer, 2006). The book will be complemented by extra notes. The book is succinct and technical in places. While we will cover much of the material in this text, it will be presented at a more elementary level.

Other books that may be of complementary interest are:

- Alexandre B. Tysbakov, "Introduction to Nonparametric Estimation," Springer Series in Statistics, 2009.
- L. Györfi, M. Kohler, A. Krzyżak, H. Walk, "A Distribution-Free Theory of Nonparametric Regression," Springer, 2002.
- T. Hastie, R. Tibshirani, J. Friedman, "The Elements of Statistical Learning: Data Mining, Inference, and Prediction," second edition, Springer, 2009.

The first two of these are books are technical and theoretical; the third is strong on intuitive motivation and methods.