

Statistics 208

An introduction to the bootstrap

Spring quarter, 2004

The bootstrap is a computer-based method for assigning measures of accuracy to statistical estimates. By substituting computation in place of mathematical formulas, it permits the statistical analysis of complicated estimators.

We will start this course by doing a few easy examples, and build up both the computational and theoretical machinery necessary to the use of the bootstrap in complex situations.

Topics: nonparametric assessment of standard errors, biases, and confidence intervals; related resampling methods including the jackknife, cross-validation, and permutation tests. Theory and applications (multivariate statistics, trees, dependent data).

This course will have a hands-on approach with lab sessions and homework projects involving computations on actual data.

Students may choose to use either matlab or R/Splus, both languages will be used in the lab sessions and as examples during the lectures.

The final grade will include a project in 2 parts, half graded at midterm and half at the final exam date.

Prerequisite: at least one course in statistics or probability.
(No previous computer experience is necessary.)

Text: Efron B., and Tibshirani R. (1993), An Introduction to the Bootstrap, Chapman and Hall.

Class meeting: Mondays, Wednesdays and Fridays, 11:00-12, Sequoia 200.

Instructor: Susan Holmes, TAs: Brit Katzen and Jie-Hua Chen.

Web page: <http://www-stat.stanford.edu/~susan/courses/s208/>