

Linear Regression Models

W4315

Instructor: Dr. Frank Wood

Required Text: Applied Linear Regression

Authors: Kutner, Nachtsheim, Neter

Not Registered Yet?

Fill out the form at

<http://tinyurl.com/3jv6mdl>

Additional books we will draw material from in this course:

- ▶ Pattern Recognition and Machine Learning, by Christopher M. Bishop. Springer, 2006.
- ▶ Bayesian Data Analysis, Second Edition, by Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin, Chapman & Hall/CRC Texts in Statistical Science

Course Description

Theory and practice of regression analysis, Simple and multiple regression, including testing, estimation, and confidence procedures, modeling, regression diagnostics and plots, polynomial regression, colinearity and confounding, model selection, geometry of least squares. Extensive use of the computer to analyze data.

Course website or <http://www.stat.columbia.edu/~fwood/w4315/>

Philosophy and Style

- ▶ Easy first half.
- ▶ Very hard second half.
- ▶ Frequent, long digressions from the required book.
- ▶ Understanding \implies proof (derivation) *plus* implementation.
- ▶ Practice makes perfect.
- ▶ Frequentist *and* Bayesian perspectives taught.

If you are looking for a pure applied, pure frequentist treatment of regression as a diagnostic tool and/or you've never programmed before, seek another section.

Available sections

- ▶ MW 10:35am-11:50am, (Zheng)
- ▶ MW 6:10pm-7:25pm (Stodden)
- ▶ F 10:00am-12:30pm (Lindquist)

Goals

- ▶ Deep theoretical understanding
 - ▶ Book provides only recipes
 - ▶ Much detail missing
 - ▶ Can always look up recipes in future
- ▶ Ability to implement/code *all* regression functionality
 - ▶ Different levels of understanding
 - ▶ Not enough to simply be able to apply formula and use pre-built regression software

About me

- ▶ Computer Science PhD, 2007, Brown University
- ▶ Postdoc in Machine Learning, Gatsby Unit, University College London
- ▶ Sports gambling consulting.
- ▶ Former entrepreneur.

My research

- ▶ Inference for nonparametric Bayesian models.
- ▶ Compression.
- ▶ Natural language data modeling.

My website: <http://www.stat.columbia.edu/~fwood>

Course Outline

The first half of the first half is a formal, theoretical review of single variable regression and its classical, frequentist treatment.

- ▶ Roughly 1 chapter per week
- ▶ 3-5 weeks, linear regression
 - ▶ Least squares
 - ▶ Maximum likelihood, normal model
 - ▶ Tests / inferences
 - ▶ ANOVA
 - ▶ Diagnostics
 - ▶ Remedial Measures
 - ▶ Linear algebra review
 - ▶ Matrix approach to linear regression

Course Outline Continued

The second half of the first half covers multiple regression and the various topics that arise from including multiple predictor variables into models.

- ▶ 3-4 weeks multiple regression
 - ▶ Multiple predictor variables
 - ▶ Diagnostics
 - ▶ Tests

Midterm

Course Outline Continued

The remainder of the course will deviate from the book and may be ordered differently than what is shown here. In general we will retain a focus on models that are linear in the parameters, but will look at nonlinear models and Bayesian treatments of linear models.

- ▶ 3-4 weeks on Bayesian regression
 - ▶ MCMC
 - ▶ Bayesian linear regression
 - ▶ Gaussian process regression
 - ▶ Projects

- ▶ 3-4 weeks on generalized regression
 - ▶ Polynomial regression
 - ▶ Logistic regression
 - ▶ Neural networks
 - ▶ Generalized linear models

Requirements

- ▶ Calculus
 - ▶ Derivatives, gradients, convexity
- ▶ Linear algebra
 - ▶ Matrix notation, inversion, eigenvectors, eigenvalues, rank, quadratic forms
- ▶ Probability
 - ▶ Random variables
 - ▶ Bayes Rule
- ▶ Statistics
 - ▶ Expectation, variance
 - ▶ Estimation
 - ▶ Bias/Variance
 - ▶ Basic probability distributions
- ▶ Programming

Projects (homework and final)

- ▶ Software
 - ▶ For homework – Matlab.

Grading

- ▶ Bi-weekly homework (25%)
 - ▶ Due every other week
 - ▶ no late homework accepted
 - ▶ None skipped
- ▶ Participation (up a half grade if I know you by the end, down a half grade if not)
- ▶ Midterm examination (25%)
- ▶ Final project (25%)
- ▶ Quizzes (25%)
- ▶ Curve

Office Hours / Website

- ▶ *http://www.stat.columbia.edu/~fwood*
- ▶ Office hours : TBA
- ▶ Office : Room 1017
- ▶ TA : Ran He
 - ▶ TA office hours TBD
 - ▶ Email: ran@stat.columbia.edu

Why regression?


- ▶ Want to model a functional relationship between an “predictor” (input, independent variable, single or multiple, etc.) and a “response variable” (output, dependent variable, potentially many simultaneous, etc.)
 - ▶ Examples?
- ▶ But real world is noisy, no $f = ma$
 - ▶ Observation noise
 - ▶ Process noise
- ▶ Two distinct goals
 - ▶ Tests about natura of relationship between predictor variables and response variables.
 - ▶ Positive
 - ▶ Negative
 - ▶ No effect
 - ▶ Prediction

History

- ▶ Sir Francis Galton, 19th century
 - ▶ Studied the relation between heights of parents and children and noted that the children “regressed” to the population mean
- ▶ “Regression” stuck as the term to describe statistical relations between variables

Example Applications

- ▶ Epidemiology
 - ▶ Relating lifespan to obesity, smoking habits, and/or other patient *features*.
- ▶ Science and engineering
 - ▶ Relating physical inputs to physical outputs in complex systems
- ▶ Grander



Aims for the course

- ▶ Given something you would like to predict and some number of covariates
 - ▶ What kind of model should you use?
 - ▶ Which variables should you include?
 - ▶ Which transformations of variables and interaction terms should you use?
- ▶ Given a model and some data
 - ▶ How do you fit the model to the data?
 - ▶ How do you express confidence in the values of the model parameters?
 - ▶ How do you regularize the model to avoid over-fitting and other related issues?
- ▶ Not be boring