STAT J535, Introduction to Bayesian Data Analysis

Instructor:
David Hitchcock, associate professor of statistics
209A LeConte College
Phone: 777-5346
Email: hitchcock@stat.sc.edu
Course Web Page: http://www.stat.sc.edu/~hitchcock/stat535.html
(Also will be accessible via Blackboard)

Classes:
Meeting Times: MW 11:00 a.m.-12:15 p.m., Wardlaw College 116 (also available online)

Office Hours:
MWF 10:05-10:40 a.m. and Tues-Thurs 11:00-11:45 a.m. or **please feel free** to make an appointment to see me at other times.

Prerequisites:
STAT/MATH 511 and STAT 515 or equivalent, or CSCE 582 (= STAT 582).

Textbook :
Gill, Jeff. *Bayesian Methods: A Social and Behavioral Sciences Approach*, Second Edition. Chapman & Hall/CRC Press, 2007.

Course Outline: Topics covered include: Principles of Bayesian statistics; one- and two-sample Bayesian models; Bayesian linear and generalized linear models; Monte Carlo approaches to model fitting; Prior elicitation; Hypothesis testing and model selection; Complex error structures, hierarchical models; Statistical packages such as BUGS/WinBUGS, R, or SAS.

Learning Outcomes: Upon completion of the course, successful students will:
- Understand the philosophy of Bayesian statistical modeling
- Understand Bayesian models for numerous common data analysis situations, including prior elicitation
- Be able to use software such as R, WinBUGS, or SAS to implement Bayesian analyses
- Understand basic principles of both conjugate analyses and MCMC-based Bayesian analyses

Homework: Homework will be assigned on the course web page. Due dates will be posted given on the course web page. Late homework will be penalized. You must do each homework problem independently. You may not look at another student's work while doing the homework. You may ask me for help on the homework problems. If homework is found to have been copied, all students involved will receive a 0.

Exams: There will be two midterm exams (February 15, March 21) and a take-home final exam due April 27. Each **midterm** exam will be given in the classroom during the regularly scheduled class time. If you are not able to come to campus for the exams, you must contact the distance education office to set up a proctor. If you are on campus and not able to attend class live, you may either contact the distance education office to set up a proctor, or may arrange with me to take it at some other time.

Graduate Student Project: Since 500-level courses are required to contain more rigor for graduate students than for undergraduates, there will be an extra short project required for graduate students. *Undergraduate students may do this project for extra credit*. The project will be due near the end of the semester and will involve collecting or obtaining a real data set and analyzing it using the methods discussed in this class. More information will be given out later in class.

Grading:
The course grade will be based on homework (20%), 2 midterm exams (25% each), and a final exam (30%). The overall course average will result in the following grades: 90-100 = A, 87-89 = B+, 80-86 = B, 77-79 = C+, 70-76 = C, 67-69 = D+, 60-66 = D, 59 and below = F.
The grading scale will be slightly more rigorous for graduate students, as required by university policy. For graduate students, the mandatory project will represent 10% of their overall grade, with the other grade components scaled proportionally.
For graduate students only: 91-100 = A, 88-90 = B+, 81-87 = B, 78-80 = C+, 71-77 = C, 68-70 = D+, 61-67 = D, 60 and below = F.

Computing:
Some problems in this course involve significant computations, and for these, we will learn to use the statistical software packages R and WinBUGS. It is required to download R (for free) and recommended to download WinBUGS (also free); instructions are given on the course web page.

Time Allocation Framework:

Topics Covered	Time
Review of Probability Concepts	1 week
Bayes' Law and the Basic Bayesian Framework	1 week
Bayesian Analyses for Basic One-Sample Models	1.5 weeks
Bayesian Linear Models	1.5 weeks
General Classes of Prior Distributions and Prior Elicitation	1 week
Some Useful Monte Carlo Methods (along with use of R and BUGS)	1.5 weeks
Assessing Model Quality	1.5 weeks
Bayesian Hypothesis Testing	1 week
Bayesian Analyses for Two- and *k*-Sample Models	1 week
Hierarchical Bayesian Models	1 week
Advanced Bayesian Models: Count Regression, Mixed Models, Models for Clustered/Longitudinal Data (Time permitting)	2 weeks