STAT 535, Introduction to Bayesian Data Analysis – Spring 2024

Instructor: David Hitchcock, associate professor of statistics
Office: 219B LeConte College
Phone: 777-5346
Email: hitchcock@stat.sc.edu
Course Web Page: http://people.stat.sc.edu/hitchcock/stat535.html
(Also accessible via Blackboard)

Classes:
Meeting Times: MWF 12:00 - 12:50 pm, LeConte College 103

Office Hours:
Mon-Wed-Fri 1:00-2:00 pm, Tuesday 1:30-2:30 pm, 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 : *Bayes Rules! An Introduction to Applied Bayesian Modeling*, by Alicia A. Johnson, Miles Q. Ott, Mine Dogucu. CRC Press, 2022. Online version available at: https://www.bayesrulesbook.com/
Other Resource (not required): Kruschke, John K.. *Doing Bayesian Data Analysis*, Second Edition. Academic Press, 2015.

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 and Stan to implement Bayesian analyses
- Understand basic principles of both conjugate analyses and MCMC-based Bayesian analyses

Class Lectures / Attendance Requirement:
You are urged to attend the lectures live on Mondays, Wednesdays, and Fridays in LeConte 103. If you are forced to miss a class or if you would like to review material from a class, I plan to record and post the lectures on the Blackboard STAT 535 course page (click Blackboard Collaborate Ultra link, then navigate to the recordings using the top left corner menu). This is not a replacement for attendance and it's not guaranteed that the recording will work every class.

Since this is an in-person class, you are expected to attend at least 80% of the class sessions in person. Attendance will be taken each class, and your grade on the attendance component will be 1.25 times the percentage of class sessions that you attend live (with a maximum of 100% for the attendance grade). For example, if you attend 60% of the class sessions in person, your attendance grade (which is 5% of the overall course grade) will be 75%. If you attend 80% or more of the class sessions in person, your attendance grade will be 100%.

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 and will not be accepted once solutions are posted. You may work with other students in this class on these problems, but you should write your answers independently of others (do not copy another student's answers). Test problem(s) will sometimes be similar in nature to assigned homework problems. Therefore you are personally responsible for knowing how to do each homework problem (even if you worked in a group on the homework). So it is important that you understand how to solve the homework problems! Please write your homework answers NEATLY.

Exams: There will be two midterm exams (an in-class exam scheduled for February 16, and a take-home exam around late March or early April). There will also be an in-class final exam on Wednesday, April 24 at 12:30 p.m. The first midterm exam will be given in the classroom during the regularly scheduled class time.

Data Analysis Project: 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. There will be the option of working in teams or individually. More information will be given out later in class.

Graduate Students: Any students enrolling in the course for graduate credit will do some extra homework problems during the semester. Graduate students will also be required to write a paper discussing an article from the scientific literature that uses Bayesian statistical analysis and describe in detail how the Bayesian approach is used in the article.

Grading:
The course grade will be based on attendance (5%), homework (12.5%), the project (12.5%), 2 midterm exams (22.5% each), and a final exam (25%). 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.
For graduate students only: The extra graduate student paper described above will be worth 11% of the overall course grade, with all other grade components adjusted proportionally. Besides that, the same grading scale will apply.

Computing: Some problems in this course involve significant computations, and for these, we will learn to use the statistical software packages R and several packages needed to run Stan code within R. It is required to download R (for free) and the associated packages; instructions are given on the course web page, within the links at the online textbook website, and via a video in Blackboard Collaborate Ultra on the STAT 535 Blackboard course page.

During Class: No cell phones may be on during class. In the classroom, laptops and tablets (e.g., iPads) may be used only for note-taking and other course-related activities. Students may not use laptops and tablets to look at unrelated web pages, play games, etc.

Disabilities: Any student with a documented disability should contact the Student Disability Resource Center at 777-6142 to make arrangements for appropriate accommodations.