

BIOS 779: Bayesian Statistics

Fall 2013, 4 credits

Location and Time:

Class: 9:00-10:45 Tuesday and Thursday, 2308 McGavran-Greenberg Hall

Instructor: Amy H. Herring, Professor of Biostatistics

Office: 3104-D McGavran-Greenberg

e-mail†: aherring@bios.unc.edu

office hours: Tuesdays and Thursdays after class by appointment

Grader	E-mail†
Jing Zhou	jingzhou@live.unc.edu
Xin Zhou	xinzhou@live.unc.edu

†You are welcome to e-mail us at any time.

BIOS 779 examines basic aspects of the Bayesian paradigm in the context of observational studies and clinical trials. Topics include Bayes' theorem, the likelihood principle, prior distributions, posterior distributions, and predictive distributions. General topics include Bayesian modeling (including linear, generalized linear, hierarchical, and survival models), informative prior elicitation, model comparisons, Bayesian diagnostic methods, variable subset selection, and model uncertainty. Markov chain Monte Carlo methods for computation are discussed in detail. Computation in R software is emphasized.

Learning Objectives:

- Gain a basic working knowledge of Bayesian statistical methods.
- Understand how to choose appropriate prior distributions for answering a research question for a given study design.
- Learn to carry out Bayesian statistical analysis in modern statistical software and interpret analysis results.

Prerequisites:

- Basic understanding of statistical inference, matrix algebra, and theory of linear models (BIOS 772)
- Knowledge of (or willingness to learn on your own) R software

Texts:

- Primary Resource
 - Hoff, *A First Course in Bayesian Statistical Methods*, available as free e-book through UNC libraries (hard copies on Amazon for roughly \$55)
 - Ibrahim, course notes, available on Sakai
 - Relevant papers as assigned
 - Laptop with R installed for in-class work (days marked as case studies on syllabus; will divide into groups so fine if you have no laptop)
- Supplemental Resources
 - Gelman, Carlin, Stern, and Rubin, *Bayesian Data Analysis*, 2nd edition (3rd edition coming in November)
 - Albert, *Bayesian Computation with R*, available as free e-book through UNC libraries
 - [Handy R tips from Germán Rodríguez](#), UNC BIOS PhD graduate
 - Christensen, Johnson, Branscum, and Hanson, *Bayesian Ideas and Data Analysis: An Introduction for Scientists and Statisticians*

Grading and Assessments:**Honor Code:**

Students in BIOS779 are expected to abide by the UNC Honor Code. All suspected Honor Code violations will be reported to the UNC Dean of Students, who will investigate the case. These investigations typically involve lengthy hearings of the Honor Court, and as outlined in the Instrument of Student Judicial Governance, “The usual sanction for a first academic violation is definite suspension for at least one academic semester and a grade penalty of an ‘F’ for the course, a portion of the course, or the assignment.”

Homework:

Homework will be assigned approximately every other week and is required. You are welcome to discuss strategies of problem solving with each other and with the graders or myself, but the work you turn in must be your own. Copying answers will be considered a violation of the UNC honor code. There will be no possibility of making up missed homework assignments. However, the lowest homework score will automatically be dropped before grades are averaged. (You will not see this reflected online in Sakai.)

Assignments will be based on numerous important health studies, including the China Health and Nutrition Survey (CHNS). We thank the CHNS, funded by NIH (R01-HD30880,

DK056350, and R01-HD38700), the Carolina Population Center, the Chinese CDC, Dr. Barry Popkin, and Mr. Jim Terry for providing these data.

Exam:

The final exam will be in-class, 8-11am on Thursday, December 12 (location TBD). Absolutely no electronic devices or resource materials are permitted during the exam.

Important note: If you need any special accommodation for the exam, you must be registered with [UNC Accessibility Resources and Services](#), who will then provide me with an official letter.

Assignment of Grades:

Each student's numeric grade will be 25% homework, 25% course project (your group's paper and presentation), 25% course project (your anonymous peer review of other projects), and 25% final examination. The course project and review procedure are described in a separate handout. Note that due to the rolling group project schedule, some project reviews are due the last week of class (these due dates are specified in the handout).

The Graduate School uses the grades H (clear excellence), P (entirely satisfactory), L (low pass), and F (failure). Graduate students are expected to earn "P" grades, with remarkable performances rewarded with other grades from the scale as appropriate. Class participation may modify the association between the numeric average and assigned letter grade in exceptional cases.

Instructor Evaluation:

Each student is expected to complete an online course evaluation during the last week of class. These are strictly anonymous and are very helpful in providing feedback about the course.

Class Etiquette:

- Behavior in and out of the classroom should enhance the learning process.
- At all times we will use common courtesy and respectful behavior.
- Attendance at lectures is expected. You are responsible for any announcements or content delivered in the classroom, whether you are in attendance or not.
- Please make every effort to arrive on time, and I promise to do so as well. If you must arrive late or leave early, please do so quietly.
- Use of electronic devices should be limited to viewing course notes.
- Eating in class is problematic in our fabulous, newly-renovated classroom! Properly contained, spill-proof beverages should not be a problem.

Schedule

Date	Topic
Tuesday, August 20	A Bayesian grace: “All models are wrong, but some are useful”
Thursday, August 22	Thomas Bayes, preacher man, and one parameter models
Tuesday, August 27	Money, sports, and sex: Bayesian basics
Thursday, August 29	<i>Great Expectations</i> and the Gaussian distribution
Tuesday, September 3	(Don’t) show me your posterior: In class case studies
Thursday, September 5	Free samples! Introduction to MCMC
Tuesday, September 10	“I once was lost”: multivariate Gaussian distribution
Thursday, September 12	“But now am found”: missing data
Tuesday, September 17	Project proposal presentations
Thursday, September 19	Project proposal presentations
Tuesday, September 24	Hierarchical models: magical land of shrinkage!
Thursday, September 26	Leave that BLUE cryin’ in the rain: linear regression
Tuesday, October 1	(Really don’t) share samples from your posterior: In class case studies
Thursday, October 3	More free samples: Metropolis-Hastings algorithms
Tuesday, October 8	Tuning your instruments: Metropolis-Hastings in practice
Thursday, October 10	Random walk with the linear mixed effects model
Tuesday, October 15	Random walk with the generalized linear mixed effects model
Thursday, October 17	Fall break
Tuesday, October 22	Have you committed statistical malpractice? Model selection
Thursday, October 24	Bayesian model selection
Tuesday, October 29	Project presentations
Thursday, October 31	Project presentations
Tuesday, November 5	New frontiers: Incorporating space
Thursday, November 7	Project presentations
Tuesday, November 12	Project presentations
Thursday, November 14	“At first I was afraid; I was petrified”: Survival analysis
Tuesday, November 19	Nonparametric Bayes: an oxymoron?
Thursday, November 21	Nonparametric Bayes
Tuesday, November 26	Bayesian clinical trials: when and why?
Thursday, November 29	Thanksgiving
Tuesday, December 3	250 years of Bayes
Thursday, December 12	Final exam, 8:00-11:00 am