BSTA 771: Applied Bayesian Analysis
Spring 2023

Instructor:	Jeffrey Morris, PhD Qi Long, PhD
Class time:	Tuesday/Thursday at 1:45pm – 3:15pm (January 12-April 25, 2023) Lectures on March 7 and March 9 cancelled due to spring break
Location:	Blockley 418
Office Hour:	To be determined
TA:	To be determined
TA Office Hour:	To be determined

Overview: This course will focus on modern Bayesian methods for analyzing biomedical data. Once students have learned the basics, we will cover more advanced topics including computational algorithms (Gibbs; Metropolis Hastings; rejection sampling; slice sampling); model checking using posterior predictive distributions; model averaging; shrinkage priors; linear models; hierarchical and linear mixed models; generalized linear models and advanced topics including semiparametric/additive models, functional data models, spatial models and factor models, as time permits. These ideas will be illustrated using data that are relevant for biostatistics graduate students.

By the end of the course students should be able to:

- Determine what type of prior distributions to use (non-informative, weakly informative, or informative)
- Be able to fit complex models by either writing their own Gibbs sampler, using available software, or some combination of the two
- Be able to implement parametric Bayesian methods
- Know how to check modeling assumptions
- Understand the strengths and limitations of Bayesian inference
- Understand how Bayesian methods can be used to deal with common problems in biomedical data such as: missing data; censoring; multiple comparisons; sparse cells

Text books:

Gelman, A., Carlin, J.B., Stern, H.S. Dunson, D.B., Vehtari A., and Rubin, D.B., 2020. Bayesian data analysis. Chapman and Hall/CRC

Hoff, P.D., 2009. A first course in Bayesian statistical methods (Vol. 580). New York: Springer.

Software:
R code will be provided for data examples in lectures: R (http://www.r-project.org/).
In addition to R, students can also use OpenBUGS or Stan for homework assignments: OpenBUGS (https://www.mrc-bsu.cam.ac.uk/software/bugs/openbugs/) and Stan (https://mc-stan.org/).

Schedule: Dr. Morris will teach Lectures 1-2 and Lectures 17-28; Dr. Long will teach Lectures 3-16

Lectures 1-2 (01/12-01/17): Introducing the Bayesian paradigm

Lectures 3-16 (01/19-03/14): Cover basic topics in Bayesian modeling and computation

- Single parameter models
- Prior specifications (conjugate; proper vs improper; informative vs noninformative)
- Multiparameter models (Gaussian/Multinomial)
- Bayes Factors and model criticism/checking
- Bayesian computation (Monte Carlo/Adaptive-Rejection SamplingMCMC; Gibbs; Metropolis Hastings; rejection sampling; slice sampling; Hamiltonian MC)

Lectures 17-28 (03/16-04/25): Cover regression analysis and related topics, with some advanced topics

- Linear Models
- Hierarchical Models and Linear Mixed Models
- Generalized Linear Models
- Bayesian Variable Selection and Shrinkage
- Finite Mixture Models
- Advanced Topics: Semiparametric Models, Functional Data Models, Factor Models

Final Project: Students are expected to choose one of the following two options: a) conduct Bayesian analysis of a data set of your choice; and b) review two or more papers related to Bayesian analysis. For each option, students are expected to give a presentation during the class session on April 27 (tentative) and submit a written report by May 5 (tentative). Each student will have 15 minutes for your presentation followed by 2 minutes for Q&A. Students are expected to send the course instructors an outline of the final project by April 13.

Grading Policy:

Attendance and Participation @ 20%
Homework @ 40%
Final project presentation @ 40%

Grades:

A: (85; 100]
B: (75; 85)
C: [60; 75]
+/- grades will be given accordingly.

While students are encouraged to discuss homework problems together, the actual document that is turned in (including computer code) must be each student's own work.