BSTA 771: Applied Bayesian Analysis
Fall 2024

Instructor: Jeffrey S. Morris, PhD (Blockley 600, jeffrey.morris@pennmedicine.upenn.edu)

Class time: Tuesday/Thurs at 1:45pm – 3:15pm (August 27-December 5, 2024)
No class and office hour on November 28 (Thanksgiving Break)
No class/office hour on Nov 14, Nov 26 (due to JSM's travels)

Location: Blockley 701

Office Hour: Jeffrey S. Morris (Blockley 600), Thursday at 3:15-4:15pm

Overview: This course will focus on an introduction to 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

Textbooks:

- 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:

Lectures 1-2 (August 27-29): Introducing the Bayesian paradigm

Lectures 3-8 (September 3-19): Cover basic topics in Bayesian modeling

- Single Parameter Models
- Prior specifications (conjugate; proper vs improper; informative vs noninformative)
- Multiparameter models and Hierarchical Models
- Bayes Factors and Model Criticism/Checking

Lectures 9-14 (September 24-October 15): Covers topics in Bayesian computing

- Monte Carlo/Adaptive-Rejection sampling
- MCMC
- Gibbs sampling

- Metropolis Hastings
- Rejection sampling
- Slice sampling
- Hamiltonian MC
- Other advanced computing topics (time allowing)

Lectures 15-20 (October 17-November 5): Bayesian regression analysis

- Linear Models
- Hierarchical Linear Models and Linear Mixed Models
- Generalized Linear Models
- Survival Analysis
- Nonparametric Additive Models

Lectures 21-24 (November 7-November 21): Advanced Topics

- Bayesian Variable Selection and Shrinkage
- Bayesian Functional Data Analysis
- Other Advanced Topics (time allowing): Factor models, Finite Mixture Models, Bayesian nonparametrics

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 December 3 or December 5 at 1:30pm-4:30pm (tentative) and submit a written report by noon on December 6 (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 November 8.

Grading Policy:

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