SCHOOL OF MEDICINE & DENTISTRY
UNIVERSITY of ROCHESTER MEDICAL CENTER

Syllabus for SMD Graduate Course BST413: Bayesian Inference (83122)

Spring 2017
Location: SRB 1.404. Class meetings: Tuesdays and Thursdays 11 am - 12:30 pm

Instructor: Sally W. Thurston, PhD
Office Hours: See below.
Course website: N/A. Course email list: may be set up at a later date.
Prerequisites: *BST411 (Statistical Inference)*

Instructional Staff

Instructor: Sally W. Thurston, PhD, Sally_Thurston@urmc.rochester.edu.
Phone: 275-2406. Office: SRB 4.145. Office hours: by appointment
TAs: Joe Ciminelli, Joseph_Ciminelli@URMC.Rochester.edu. Office hour: Wed 1-2 (& appt.)
Valeriia Sherina, Valeriia_Sherina@URMC.Rochester.edu. Office hour: Th 12:45-1:45 (& appt.)

Course Description: Major topics

1 - Introduction to Bayesian statistics including Bayes theorem, the likelihood principle, comparison to the frequentist approach
2 – Bayesian inference for single parameter models with conjugate priors
3 – Bayesian inference for multiple parameter models with conjugate, semi-conjugate, and non-informative priors
4 – Large sample inference and frequency properties of Bayesian inference (very briefly covered)
5 – Hierarchical models
6 – Model checking including posterior predictive distribution and graphical checks
7 – Introduction to Markov Chain Monte Carlo, including Gibbs sampling and the Metropolis-Hastings algorithm, and assessing convergence (Note: this topic is covered in more depth in BST512)
8 – Regression from a Bayesian perspective
9 – Decision analysis: brief overview
10 – Student projects

The course material includes extensive programming in R, and an introduction to knitr. Previous knowledge of R is helpful but is not a course prerequisite.

Course Aims and Objectives

By the end of the course students should be very familiar with the framework for Bayesian inference and how this differs from frequentist inference. Students will be able to derive posterior distributions in closed form for parameters in many models, when using conjugate priors. Students will be familiar with hierarchical models and their advantages, and with methods for model checking. Due to the extensive programming required in this course, students will be able to write MCMC code in R for a variety of models, and will be familiar with some of the ways of checking for convergence and adequate run length of the MCMC chain.

Course Policies and Expectations

Students are expected to attend class, unless ill, out of town, or under extenuating circumstances. Class participation and questions are encouraged.

Students are encouraged to discuss the homework with other students. However all computer code should be written individually (your code may be based on my programs), and all written answers to the homework should be your own.

While I don't have assigned office hours, I am happy to set up appointments to answer questions and/or discuss anything related to the class. I also find that many short questions can be answered quickly and easily through email. I use email extensively to communicate with students, including sending homework assignments, any corrections to the homework, a reminder of course announcements, etc.

Materials and Access

The required text is: Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB "Bayesian Data Analysis", third edition, CRC Press, 2014. Copies of other required reading, e.g. journal articles, will be passed out in class. The text will be available on reserve in Karin Gasaway's office in Biostatistics, and at the Carlson library (loan period: 1 day).

Assignments and Grading Procedures

Homework will generally be assigned every two weeks and will be due two weeks after being assigned. I anticipate six homework assignments. Homework assignments will be graded by the TAs, and generally will be ready to hand back to students about a week after being due. The course also includes a final project, which will include a project proposal and class presentation which I will grade. New this year will be a single quiz given about 1/2 to 2/3 of the way through the semester. Grades will be based on class participation (5%), the final project (15%) and the quiz and homework (total 80% of the grade; quiz weighted as the average homework weight). Grades are not curved.

Academic Integrity

Academic integrity is a core value of the University of Rochester. Students who violate the University of Rochester University Policy on Academic Honesty are subject to disciplinary penalties, including the possibility of failure in the course and/or dismissal from the University. Since academic dishonesty harms the individual, other students, and the integrity of the University, policies on academic dishonesty are strictly enforced. For further information on the University of Rochester Policy on Academic Honesty, please visit the following website:
http://www.rochester.edu/college/honesty/docs/Academic_Honesty.pdf

Accommodations for Students with Disabilities

Students needing academic adjustments or accommodations because of a documented disability must contact the Disability Resource Coordinator for the school in which they are enrolled:
http://www.rochester.edu/eoc/DisabilityCoordinators.html

Course Schedule

1. Background: Why Bayesian statistics? (Gelman et al, ch 1): Bayes theorem, the likelihood principle, the Bayesian approach and relationship to the frequentist approach, example calculations for posterior odds.

2. Single parameter models (Gelman et al, ch 2): review of exponential families, single-parameter models with conjugate priors, posterior predictive distributions, noninformative priors and Jeffrey's invariance principle.

3. Multiple parameter models (Gelman et al, ch 3): The normal model with noninformative priors for the mean and variance, marginal distribution for the mean, posterior predictive distribution for a new observation. Normal model with conjugate priors. Normal model with semi-conjugate prior. Multinomial model, and sampling from the Dirichlet distribution. Multivariate normal model with unknown mean and unknown variance. Non-conjugate example: a logistic regression of animal death as a function of dose, with intercept, slope and LD50 the parameters of interest.

4. Large sample inference and frequency properties of Bayesian inference (Gelman et al., ch 4): The delta method for two parameters and how this relates to a normal distribution, why most posterior distributions can be approximated by a normal distribution, and counterexamples when this fails.

5. Hierarchical models (Gelman et al, ch 5 and extra material): Exchangeability including when this does and does not make sense. Sampling from the posterior for a hierarchical model in the context of several examples including a binomial-beta model and a hierarchical normal model. Example of when the prior gives an improper posterior. Comparison of a hierarchical model to other models. Hierarchical models in R; possible brief overview of hierarchical models in Stan.

6. Model checking (Gelman et al, ch 6): the posterior predictive distribution and graphical checks, posterior predictive p-value, model expansion, model comparisons, Bayes factors.

7. Posterior simulation and introduction to MCMC (Gelman et al, ch 10-11, plus lots of extra material): overview of methods to sample from the posterior. Introduction to MCMC including the Metropolis-Hastings algorithm and Gibbs sampling. Examples of R code for Gibbs sampling and the Metropolis-Hastings algorithm. Inference and assessing convergence of MCMC (this is an introduction – it covered in more depth in BST512), including traceplots, the Gelman-Rubin diagnostic, and the Raftery-Lewis diagnostic. Possibly: comparison of R and Stan.

8. Regression (Gelman et al., ch 14): brief introduction to regression with noninformative priors including deriving posteriors for the model parameters. Using the Gibbs sampler to obtain draws from the posterior. Posterior predictive distribution for new data.

9. Decision analysis (Gelman et al., ch 22): presentation from both a frequentist and a Bayesian perspective, using several examples.

10. Student projects: the last few classes are devoted to student presentations of projects (3-4 students present per class). Typically, projects for this class consist of the student presenting a paper to the class on some aspect of Bayesian statistics, but other project types are possible.

There will be no final exam. An absence by the instructor will likely be handled by having a TA give the lecture, but might be handled by scheduling a make-up class at a different time, and/or by scheduling a guest lectures during an instructor absence. There are no classes or homework assigned during Spring Break. The work for this class will end by the last lecture, unless it is necessary to add an additional lecture for some unexpected reason.

Other potentially useful information for students

Guidelines for homework and the project will be handed out separately, at the appropriate time. If you need help during the semester, please email me or speak to me after class or at a different time.

(last edited: 1/23/17)