Last updated: January 24, 2013

Instructor: Rob Scharpf

Office:	Suite 1101 550 Building
Email:	rscharpf@jhsph.edu (use [bayes] in subject line)
Course dropbox:	email invitation
Office hours:	by appointment

Description Illustrates current approaches to Bayesian modeling and computation in statistics. Describes simple familiar models, such as those based on normal and binomial distributions, to illustrate concepts such as conjugate and noninformative prior distributions. Discusses aspects of modern Bayesian computational methods, including Markov Chain Monte Carlo methods (Gibbs' sampler) and their implementation and monitoring. Bayesian Methods I is the first term of a two term sequence. The second term offering, Bayesian Methods II (140.763), develops models of increasing complexity, including linear regression, generalized linear mixed effects, and hierarchical models.

Topics Probability, random variables, distributions, descriptions of distributions, Bayes' rule, and exchangeability. The Binomial model: priors, posterior calculation, and posterior predictive distribution. The normal model: priors, posterior calculation, Markov Chain Monte Carlo methods (Gibbs' sampler). Implementation of Bayesian data analysis in the computing environment R.

Learning Objective Upon successfully completing this course, students will be able to: (1) explain the difference between the Bayesian approach to statistical inference and other approaches, (2) develop Bayesian models for combining information across data sources, (3) write and implement programs to run analyses, (4) evaluate the influence of alternative prior models on posterior inference, (5) plot and interpret posterior distributions for parameters of scientific interest.

Text Peter D. Hoff (PDH), A First Course in Bayesian Statistical Methods.
Bradley P. Carlin and Thomas A Louis (C&L), Bayesian Methods for Data Analysis, 3rd edition.

Prerequisites Biostatistics 140.651 and 140.652, or instructor consent

Grading Four homework assignments (60%) and a data analysis project (40%).

Course outline

Date	Content	Reading	Assignments
Week 1	As a motivating example, we consider the problem of estimating the prevalance of a rare disease. We begin with non-Bayesian approaches to the problem, namely hypothesis tests and p-values and 95% Wald confidence interval for the prevalance. Potential problems with this approach are discussed. We present an alternative solution that is Bayesian, wherein we introduce Bayes' theorem and the concept of a prior and posterior distributions. We choose a convenient prior for this example, sidestepping the difficulty of eliciting priors and the mathematical details of deriving a posterior distribution.	PDH, 1	

Date	Content	Reading	Assignments
Th	We discuss properties of *rational* beliefs that we refer to as axioms. We demonstrate that our belief axioms are upheld by the axioms of probability · Events, partitions, and Bayes' rule · Random variables, describing distributions, probability densities, cumulative distrubitions, joint distributions, independent RVs, exchangeability, and de Finetti's theorem.	PDH, 2	Homework 1
Week 2 Tu	The binomial sampling distribution. Derivation of posterior using beta priors. Concepts include sufficient statistics, conjugacy, and informative priors. We will illustrate how the same prior can be *informative* in some applications, but *uninformative* in others. Prediction and Bayesian coverage (credible intervals / highest posterior density) are discussed.	PDH, 3	
Th	We discuss the Poisson sampling distribution and its conjugate Gamma prior. We use the birth rate data from PDH for illustration. Following PDH, we conclude this section with a discussion of conjugate priors for one-parameter exponential family models.		Homework 1 due
Week 3, Tu	In many applications, there are other aspects of the posterior distribution that may be of interest. For example, functions of the mean such as risk differences, relative risks, log odds, etc. We show how arbitrary functions can be approximated from random samples of the posterior distribution, a procedure known as Monte Carlo approximation.	PDH, 4	
Th	Monte Carlo approximation for predictive distributions.		
Week 4, Tu	We discuss the normal distribution and Bayesian estimates for the unkown mean with known variance. As in the case of the discrete distributions, the Bayesian estimator for the mean is a weighted average of the prior and the sample mean. We illustrate these concepts using the Midge wing length data in PDH [find a better example].	PDH, 5	
Th	We discuss joint inference in the normal model when both the mean and variance are unknown. Here, we have a joint prior distribution for the mean and variance. Conjugate priors for the mean and variance are discussed. We revisit the midge wing length example and touch on the use of the normal model for non-normal data.		
Week 5, Tu	The joint posterior distribution of multiparameter models can be difficult to derive analytically or may not have a standard form. However, the full conditional distribution of the parameters is often much more tractable. We demonstrate how iteratively sampling from the full conditionals can be used to sample from the joint posterior. The procedure, known as the Gibbs sampler, is illustrated in the context of joint inference for the mean and variance of the normal distribution when the marginal distribution of the inverse variance does not have a standard form. Here, a "semiconjugate" prior distribution is adopted.	PDH, 6	Code a Gibbs sampler
Th	Gibb's sampler continued		
Week 6, Tu	Hierarchical models · data analysis project		
Week 7, Tu	Special topics: TBA		
Week 8, Tu	Student presentations		

Date	Content	Reading	Assignments
Th	Student presentations		