STA 7934 Markov Chain Monte Carlo Methods in Statistical Inference
Spring 2013

Instructor Hani Doss, 222 Griffin-Floyd Hall; Office Hours: 3rd period, i.e. 9:35–10:25am. (If you come see me, don't come later than 10:10am).

Course Web Page `http://www.stat.ufl.edu/~doss/Courses/mcmc`

Course Description Monte Carlo methods are increasingly used in many scientific areas, including statistical physics (where they originated), Bayesian and frequentist statistical inference, and image reconstruction. The basic idea is to carry out a simulation to estimate quantities of interest that cannot be computed analytically. This course will begin with a brief discussion of some standard Monte Carlo schemes, before moving to Monte Carlo methods based on Markov chains.

Consider the situation where there is a distribution π on some space, and we are interested in estimating π or $\int f\, d\pi$ where f is some function, but π is analytically intractable. Markov chain Monte Carlo proceeds as follows. We set up a Markov chain with the property that its transition function has π as its stationary distribution. Then we run a chain $X_1, X_2, \ldots$ with this transition function. If the Markov chain converges to its stationary distribution (i.e. for large n, the distribution of X_n is approximately π), then by running the chain long enough, we can obtain a sample from π. This sample can be used to estimate π or some feature of it such as $\int f\, d\pi$.

In this course I will explain the method in detail, describe the main implementations, and discuss some classes of problems in statistics, primarily in Bayesian inference, where it has had success. The method is not fool-proof. I will talk about some of the mathematical results pertaining to convergence issues, and also discuss some practical convergence diagnostics.

Prerequisites STA 6466-7 (Probability Theory I and II) and STA 7346 (Statistical Inference I). You also need to be familiar with the statistical computing language R. This is a course intended for Ph.D. students in the Statistics Department. Students who are not in the Statistics Department and who do not have the prerequisites should not take this course. I will not assume you know anything about Markov chains. If you are not sure you have the prerequisites, see me.

Grading There will be about eight homework assignments, some of a theoretical nature and some involving computer implementation of the methods we discuss on specific data sets, counting for 20% of the grade; two midterms, each counting for 25% of the grade; and a final, counting for 30% of the grade.

The two midterms will be given in the evening. Tentative dates are Monday February 18 and Monday March 25, both at 6:30 pm.

Course Policies Homework must be turned in at the beginning of the lecture on the due date. Late homework will not be accepted. The solutions to the homework assignments must be entirely your own (this applies also to R code).

Topics

- Issues in practical implementation of Bayesian statistics
- Censored data
- Basic Monte Carlo methods
- General idea of Markov chain Monte Carlo
- The Gibbs sampler (general properties; application to latent variable models, including hierarchical Bayesian models and censored data models; application to high dimensional problems)
- Rao-Blackwellization
- BUGS
- Convergence diagnostics
- Application of the Gibbs sampler to nonparametric Bayes problems
- The Metropolis-Hastings algorithm (general properties; application to Ising model, Mallows model; random walk chains and independence chains; adaptive rejection Metropolis sampling)
- Theory of convergence (ergodic theorems and central limit theorems)
- Estimating ratios of normalizing constants (bridge sampling; estimation using output from several distributions)
- There will be additional topics, such as application of MCMC to estimation of Bayes factors, and to Bayesian model selection (if time permits)