STA 6276 Statistical Computing II: Monte Carlo Methods in Statistical Inference
Spring 2022

Instructor Hani Doss—222 Griffin-Floyd; email: doss@stat.ufl.edu (email is primarily for administrative purposes, not for questions regarding the course material; for such questions, talk to me during office hours). Zoom Office Hours: MWF period 6, i.e. 12:50pm–1:40pm (if you want to talk to me, please join zoom before 1:20pm). If this time doesn't work for you then let me know (either in class, or by email, or by phone) and I will arrange a meeting at a different time. In-person office hours are at the same time; however, note that social distancing makes in-person office hours in my office somewhat awkward, and I allow only one student in my office at a time. For short questions, you may talk to me in person right after class in the hall. I will email the class the following information, which you should not give out to anyone who is not in the class: my zoom personal ID, and the username and password for the parts of the course webpage that are password protected.

Course Description Monte Carlo methods are now used in virtually every scientific area, including statistical physics (where they originated), Bayesian and frequentist statistical inference, image reconstruction, and various parts of machine learning. 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.

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

Prerequisites/Corequisites STA 6326 (Introduction to Theoretical Statistics I) is a prerequisite and STA 6327 (Introduction to Theoretical Statistics II) is a corequisite. You also need to know some probability theory beyond what is covered in the Master's program, but I will go over the facts you need to know. Additionally, you need to be familiar with the statistical computing language R. I will not assume you know anything about Markov chains.

Grading Your course grade will be based on the four components below, with the stated weights.

Exam 1:	Friday February 11, 8:20pm.	25%
Exam 2:	Friday March 18, 8:20pm. Covers the material after Exam 1.	25%
Exam 3:	Date and time to be determined. Covers the material after Exam 2.	25%
HW:	There will be about 6 or 7 homeworks assigned during the semester.	25%

Some of the homework assignments will be of a theoretical nature, and some will involve computer implementation of the methods we discuss on specific data sets. 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
- Illustrative example: 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 and variants thereof
- Convergence diagnostics
- Application of the Gibbs sampler to nonparametric Bayes problems
- The Metropolis-Hastings algorithm (general properties; application to Ising model; random walk chains and independence chains; adaptive rejection Metropolis sampling)
- Hamiltonian Monte Carlo
- Theory of convergence (ergodic theorems and central limit theorems)
- There will be one additional in-depth application of MCMC to some nontrivial problem in statistics, which is likely to be Bayesian variable selection in regression.

If you need special arrangements because of a disability please contact me.