Ma 322: Biostatistics
Homework Assignment 6

Prof. Wickerhauser

Due Friday, March 2nd, 2012

Read Chapter 11, "Algorithms for MCMC," pages 185–209 of our text.

Note: Although our text has no index or table of contents, it is easy to locate words in the electronic version using the Find function of your favorite PDF reader.

1. Suppose that K is a positive constant and $f(x) = K/(1+x^2)$ is a pdf on $\mathbf{R}$.

   (a) Find K.

   (b) Find the cumulative distribution function (cdf) of f.

   (c) Find the inverse cdf of f.

   (d) Use the formula in part c to simulate sampling 1000 times from the pdf f. Plot the resulting histogram over 30 equal-width bins. Use random number seeds 12345 and 6789.

2. Consider the pdf $f(x) = Kx$ defined on the probability space $X = [0, 1]$.

   (a) Find the value of K.

   (b) Implement a rejection sampler for f and plot the histogram of at least 1000 samples in 20 bins. Use random seeds 12345 and 6789.

3. Suppose that X and Y are discrete random variables that have the following joint pdf:

   Joint pdf of X and Y

$f(X,Y)$	$Y=1$	$Y=2$	$Y=3$	$Y=4$
$X=1$	0.05	0.10	0.15	0.20
$X=2$	0.20	0.15	0.10	0.05

   (a) Compute the marginal pdfs f_X and f_Y.

   (b) Compute the complete conditional pdfs $f(X|Y)$ and $f(Y|X)$.

   (c) Implement a Gibbs sampler for this joint pdf and simulate taking 10 000 samples. Print the resulting matrix of counts as well as the normalized matrix that approximates the joint pdf. (Hint: see the example `R` code on p.193 of our test.)

   (d) Compute the marginal pdfs for X and Y from the simulation in part c, giving both the raw counts and the normalized vectors that approximate f_X and f_Y.

4. Suppose that X and Y are continous random variables on $[0,1] \times [0,1]$ with the joint pdf

$$f(x,y) = cx(1-x)(1-xy)(1+y),$$

where c is a positive constant.

(a) Compute c. (Hint: use Macsyma.)

(b) Compute the marginal densities $f_X(x)$ and $f_Y(y)$. (Hint: use Macsyma.)

(c) Are X and Y independent?

(d) Use the formulas from part b to compute the conditional probabilities $f(X|Y)$ and $f(Y|X)$. (Hint: use Macsyma.)

5. Implement the Gibbs sampler `gibbsBVN()` on p.195 of our text and use it to generate 1000 mean-$(0,0)$ samples with (a) $\rho = 0.33$ and (b) $\rho = -0.88$. Show your results in scatterplots.

   Use `mvrnorm()` to generate 1000 mean-$(0,0)$ samples from bivariate normal densities with covariances corresponding to $\begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$ with (c) $\rho = 0.33$ and (d) $\rho = -0.88$. Show your results in scatterplots.

6. Implement the Metropolis algorithm as on p.203 of our text, only using an experimental outcome of 4 successes out of 9 trials as the likelihood, and a symmetric uninformative prior $f(\theta) \propto 1$, which is a beta density with $\alpha = 1, \beta = 1$ and thus a mean of $1/2$. Start with initial $\theta = 0.04$ as in the text's example, perform 1000 steps, and plot the histogram of the result after a 50-step burn-in against the known posterior beta-density having $\alpha = 5$ and $\beta = 6$.

7. Implement the Metropolis-Hastings algorithm by modifying the code on p.203 of our text, again using an experimental outcome of 4 successes out of 9 trials as the likelihood, but with the nonsymmetric prior $f(\theta) \propto \theta^2$, which is a beta density with $\alpha = 3, \beta = 1$ and thus a mean of $3/4$. Start with initial $\theta = 0.04$ as in the text's example, perform 1000 steps, and plot the histogram of the result after a 50-step burn-in against the known posterior beta-density having $\alpha = 7$ and $\beta = 6$.