Petter Mostad
Applied Mathematics and Statistics
Chalmers and GU

MSA100 / MVE186 Computer Intensive Statistical Methods

Exam 22 October 2016, 14:00 - 18:00
Examiner: Petter Mostad, phone 031-772-3579
visits the exam at 15:00 annd 17:00
Allowed aids: None.
The appendix of this exam contains informaion about probability distributions.
Total number of points: 30. To pass, at least 12 points are needed

1. (4 points) Assume $x \mid \theta \sim \text{Uniform}[0, \theta]$, so that we have the density

$$\pi(x \mid \theta) = \frac{1}{\theta} I[0 \leq x \leq \theta]$$

   (a) Assume θ has a Pareto$(1, 1)$ distribution, i.e., the density is

$$\pi(\theta) = \frac{1}{\theta^2} I[1 \leq \theta]$$

   find the posterior distribution $\pi(\theta \mid x)$.

   (b) Prove that the Pareto family is a conjugate family to the distribution of $x \mid \theta$.

   (c) Find the prior predictive distribution for x when θ has the prior indicated in (a).

2. (4 points) In each of the following cases, assume that you want to generate a sample of size one billion from the distribution in an efficient way. Write down an algorithm where generation of random numbers is done only from the Uniform$[0, 1]$ distribution.

   (a) A Cauchy$(3, 1)$ distribution.

   (b) A Poisson(2.9) distribution.

   (c) An Exponential(3) distribution conditional on values being above 10.

3. (8 points) We assume that, for $i = 1, \ldots, n$, counts c_i have been generated from Poisson distributions with intensities λ_i. We also assume that the λ_i are independently drawn from a Gamma(α,β) distribution, where $\alpha > 0$ and $\beta > 0$, and that we use flat priors on α and β, so that $\pi(\alpha) \propto 1$ and $\pi(\beta) \propto 1$. Finally, we assume that for each i, instead of the counts c_i we have only observed censored counts y_i with possible values $0, 1, 2$, and "many". In other words, if the count is less than 3, y_i gives the count, otherwise it is equal to "many".

   (a) Write down a function proportional to the joint posterior density for all the variables above.

(b) For each of the variables above, write down a function proportional to its conditional density given all the other variables, and, when possible, identify the name and parameters of this conditional distribution.

(c) For each of the variables above where the conditional distribution is not equal to a standard named probability distribution, describe how you would perform the conditional simulation of that variable in an implementation of Gibbs sampling for this problem.

(d) Assume that, given the data $y_1, \ldots, y_n$, you want to find the maximum likelihood estimates for the parameters α and β. Outline how the EM algorithm could be used make such an estimate (I do not expect you to produce all details, as this may take you too much time).

4. (4 points)

(a) Write down, using precise notation, the Metropolis Hastings algorithm.

(b) Define what it means that a Markov chain satisfies the detailed balance condition relative to a density f.

(c) Prove if a Markov chain satisfies the detailed balance condition relative to a density f then f is a stationary distribution for the Markov chain.

(d) Prove that the Markov chain defined by the Metropolis Hastings algorithm satisfies the global balance condition relative to the target density.

5. (4 points) Assume $x > 0$ has the probability distribution specified by

$$\pi(x) \propto \exp(-x^{2.3}) \cdot \exp(-x^{1.3}) \cdot \frac{1}{1 + x^3}$$

Our goal is to simulate from this distribution.

(a) Consider the ideas of slice sampling. Define the three extra variables y_1, y_2, and y_3 suggested by slice sampling in this case, and write down a function proportional to the joint distribution on x, y_1, y_2, y_3.

(b) Write down the details of the slice sampler algorithm in this case.

6. (4 points) Assume a variable x with $0 < x < 1$ has density

$$\pi(x) = C\,(\sin(14x) + \cos(19x))^2$$

Our goal is to simulate from this distribution.

(a) Suggest a specific function to use in a rejection sampling algorithm for this problem.

(b) Describe the rejection sampling algorithm in this case.

(c) Extend the algorithm so that it also produces an estimate for the number C.

7. (2 points) Below is a trace plot from a Metropolis Hastings simulation using independent proposals. Do you think the output indicates problems with the simulation, and if so, why? Do you have a general suggestion for how the proposal function should be changed?

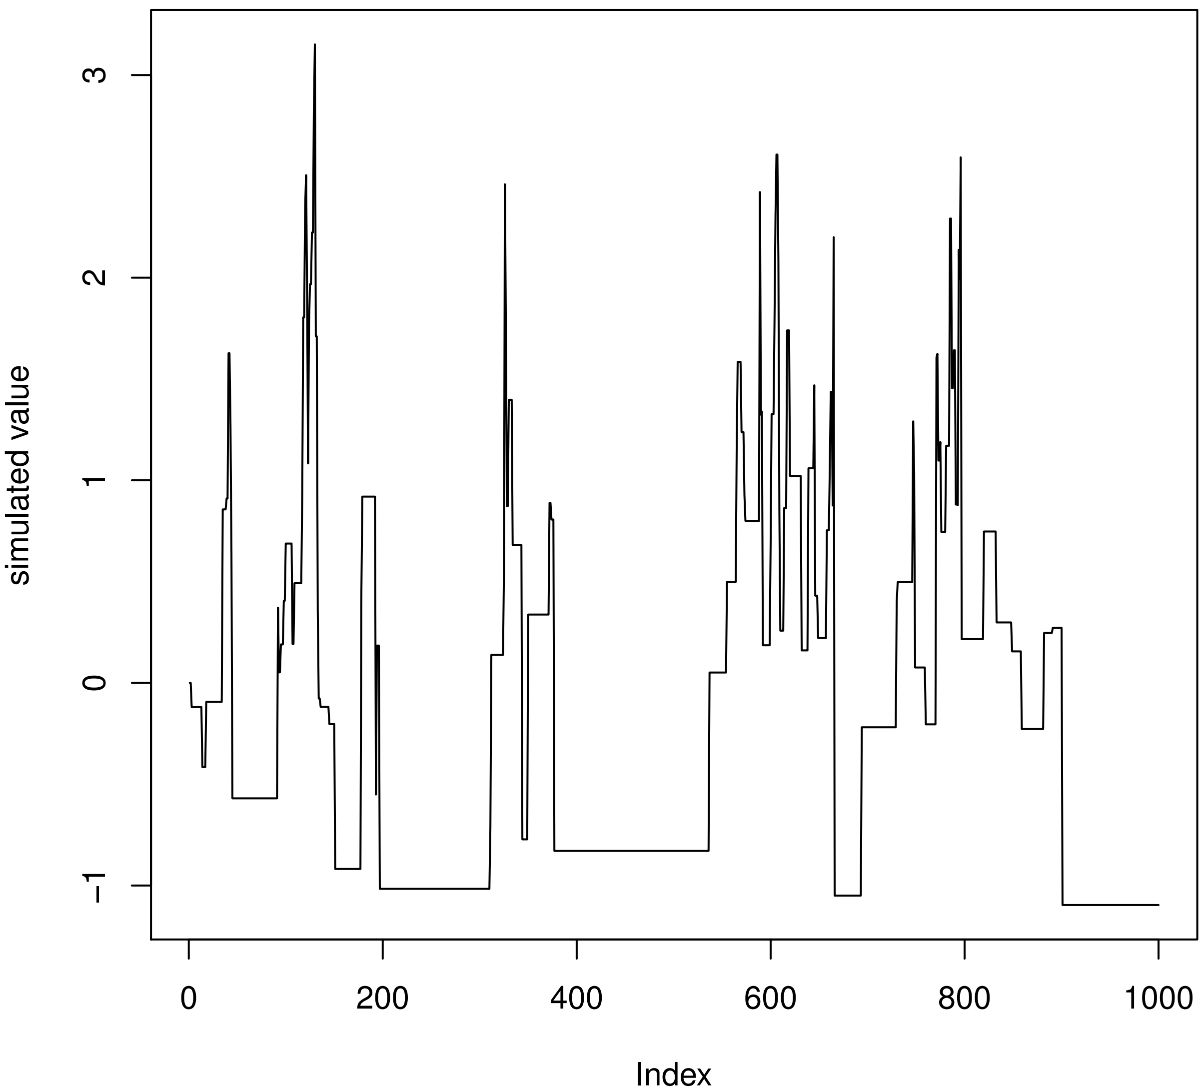

Appendix: Some probability distributions

The Beta distribution

$x \geq 0$ has a Beta distribution with parameters $\alpha > 0$ and $\beta > 0$ if it has density

$$\pi(x \mid \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}$$

The Cauchy distribution

If x has a Cauchy(μ, γ) distribution ($\gamma > 0$), then the probability density is

$$\pi(x \mid \mu, \gamma) = \frac{1}{\pi\gamma\left[1 + \left(\frac{x-\mu}{\gamma}\right)^2\right]}$$

and the cumulative distribution is

$$F(x) = \frac{1}{\pi}\arctan\left(\frac{x-\mu}{\gamma}\right) + \frac{1}{2}$$

The Exponential distribution

If $x > 0$ has an Exponential(λ) distribution, with $\lambda > 0$, the density is given by

$$\pi(x \mid \lambda) = \lambda \exp(-\lambda x)$$

and the cumulative density function is given by

$$F(x) = 1 - \exp(-\lambda x)$$

The Gamma distribution

If x has a Gamma(α, β) distribution, with $\alpha > 0$ and $\beta > 0$, then the density is

$$\pi(x \mid \alpha\beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1}\exp(-\beta x)$$

The Negative Binomial distribution

Given a number r of failures and a probability p of success, the Negative Binomial specifies the number k of successes observed until r failures are observed. The probability mass function:

$$\pi(k \mid r, p) = \binom{k+r-1}{k}(1-p)^r p^k$$

The Pareto distribution

If θ has the Pareto distribution with parameters M and α,

$$\theta \mid M, \alpha \sim \text{Pareto}(M, \alpha)$$

then the density is

$$\pi(\theta \mid M, \alpha) = \alpha M^\alpha \frac{1}{\theta^{\alpha+1}} I(M \leq \theta)$$

and the cumulative distribution is

$$F(\theta) = \left(1 - \left(\frac{M}{\theta}\right)^\alpha\right) I[M \leq \theta]$$

The Poisson distribution

If k has a Poisson(λ) distribution with intensity $\lambda > 0$, then the probabilities are given by

$$\pi(k \mid \lambda) = \exp(-\lambda)\frac{\lambda^k}{k!}$$