APPM 4/5720: Computational Bayesian Statistics, Spring 2018

Problem Set Three (Due Wednesday, February 14th)

1. Let $\vec{X} = (X_1, X_2, \ldots, X_n)^t$ be a vector of random variables. Let $V = Var(\vec{X})$ be the $n \times n$ positive definite variance-covariance matrix. This means that, for any $n \times 1$ non-zero vector $\vec{x}$, V satisfies $\vec{x}^t V \vec{x} > 0$. Positive definiteness is an important to us because it implies that V is invertible.

   Let $1 \leq k \leq n$. Show that the variance covariance matrix for $\vec{Y} = (X_1, X_2, \ldots, X_k)^t$ is also positive definite.

2. (Gelman et. al Ch 3, Exercise 5) Rounded data: It is a common problem for measurements to be observed or recorded in rounded form. For a simple example, suppose we weigh an object five times and measure weights, rounded to the nearest pound, of 10, 10, 12, 11, and 9. Assume that the unrounded measurements are nonmally distributed with independent noninformative priors for μ and σ^2.

   (a) Give the posterior distribution for (μ, σ^2) obtained by pretending that the observations are exact unrounded measurements.

   (b) Give the correct posterior distribution for (μ, σ^2) treating the measurements as rounded.

   (c) Correct the correct and incorrect posterior distributions. (Ex: means, variances, any plots you might find interesting, etc...)

   (d) Let $z = (z_1, z_2, \ldots, z_5)$ be the original unrounded measurements. (This is not equal to $(10, 10, 12, 11, 9)$, rather $(10, 01, 12, 11, 9)$ is one observation of z.) Simulate values from the posterior distribution of z. Estimate the posterior mean of $(z_1 - z_2)^2$.

3. In this exercise, we will generate iid values from the $Beta(3, 2)$ distribution using the accept-reject algorithm. Recall that the first step to generating values from a distribution with pdf f is to find a function g such that $g(x) \geq f(x)$ for all x in the space of interest.

   (a) Take g to be a flat line. Code up and run the accept-reject algorithm to generate $100,000$ values from the $Beta(3, 2)$ distribution. Submit code and a histogram of results with the superimposed target density.

   (b) Take g to be a piecewise linear function, with two pieces, that more closely resembles the target density.

      i. Carefully describe how you would simulate values from the normalized version of g.

      ii. Code up and run the accept-reject algorithm to generate $100,000$ values from the $Beta(3, 2)$ distribution. Submit code and a histogram of results with the superimposed target density.

   (c) Which of (a) and (b) do you expect to be faster and why? Report average times to acceptance from both of your simulations.

4. (Gelman et. al Ch 3, Exercise 1) Suppose that the vector $(X_1, X_2, \ldots, X_k)$ has a multinomial distribution with parameters n and $(\theta_1, \theta_2, \ldots, \theta_k)$. Assume a Dirichlet prior for $(\theta_1, \theta_2, \ldots, \theta_k)$ with known hyperparameters $(\alpha_1, \alpha_2, \ldots, \alpha_k)$. Assume that n is fixed and known. Define

$$\theta := \frac{\theta_1}{\theta_1 + \theta_2}.$$

   (a) Find the marginal posterior distribution for θ given observations $(x_1, x_2, \ldots, x_n)$.

   (b) Show that this distribution is identical to the posterior distribution for θ by treating x_1 as an observation from the binomial distribution with probability θ and sample size $x_1 + x_2$, ignoring the data $x_3, \ldots, x_n$.