T-61.5040 Oppivat mallit ja menetelmät
T-61.5040 Learning Models and Methods
Pajunen, Viitaniemi

Exercises 8, 16.3.2007

Problem 1.

Assume that the true posterior is $p(\theta|y) = N(\theta|0, \sigma^2 I)$, where $\theta \in \mathbb{R}^{1000}$. You try to perform rejection sampling using a proposal density $g(\theta) = N(\theta|0, \sigma_g^2 I)$. Suppose that your proposal density is close to the true posterior so that $\sigma_g = 1.1\sigma$. Compute approximately how many samples are rejected.

Problem 2.(demonstration)

The Metropolis algorithm simulates a posterior by starting from a value θ^0. Then the algorithm repeats a step n which produces value θ^n give θ^{n-1}, $n = 1, 2, 3, ...$ In the nth step of the algorithm a new value θ^* is picked from a jumping distribution $J(\theta^*|\theta^{n-1})$. The new value is accepted ($\theta^n = \theta^*$) with a probability $p_r = \min\{1, r\}$ where $r = \frac{p(\theta^*|y)}{p(\theta^{n-1}|y)}$. If it is not accepted, the next sample is $\theta^n = \theta^{n-1}$.

Assume that the Markov chain defined by the Metropolis algorithm has a unique stationary distribution. Show that this distribution is the posterior $p(\theta|y)$ (assume that the jumping distribution is symmetric).

What if the posterior is $1/2 p_1(\theta|y) + 1/2 p_2(\theta|y)$, where p_1 and p_2 are uniform distributions over $[0, 1]$ and $[2, 3]$, respectively? Can you think of a jumping distribution that prevents a unique stationary distribution?

Problem 3.

Observe data $y_1, \ldots, y_n$ from a Normal distribution $N(\mu, \sigma^2)$. Assume both μ and σ^2 are unknown. Choose the priors as $p(\mu|\sigma) = N(\mu|\mu_0, \sigma_0^2)$ and $p(\sigma^2) = IG(\sigma^2|a, b)$.

i) Formulate the Gibbs Sampler for the unknowns μ and σ^2.

ii) Describe how you estimate the posterior mean of μ using the simulated posterior. Recall that Monte Carlo approximation is $\mathrm{E}(h(\mu, \sigma^2)|y) \approx \frac{1}{N}\sum_i h(\mu_i, \sigma_i^2)$ where (μ_i, σ_i^2) is the i:th simulated posterior sample.

Hint: inverse-gamma distribution is $IG(z|a, b) \propto z^{-(a+1)} \exp(-b/z)$. It is a conjugate prior for σ^2 when the model is $N(\mu, \sigma^2)$ where μ is known. Writing $y = \frac{1}{n}\sum_i (y_i - \mu)^2$, the posterior for σ^2 is

$$p(\sigma^2|D) = IG(\sigma^2|\frac{n}{2} + a, \frac{1}{2}(2b + ny)).$$

Problem 4.

You have observed independent samples $y_1, \ldots, y_n$ from a distribution

$$p(y|a,b) = b\exp(-b(y-a)), \text{ when } y \geq a \text{ and } p(y|a,b) = 0, \text{ when } y < a.$$

The parameters a and b are nonnegative.

i) Choose an uninformative prior $p(a,b) \propto b^{-1}$ and compute the unnormalized posterior of a and b. Which two scalar functions of data $y_1, \ldots, y_n$ determine the posterior?

ii) Write a Gibbs Sampler for the posterior.
Hint: Gamma distribution for x is $Gamma(\theta|c,d) \propto x^{c-1}\exp(-dx)$.
Another hint: you don't have to simulate a if it seems difficult. You might want to simulate $exp(abn)$, since you can solve a from these values when b is known.