Computational Statistics
First course exam, 18 Mar 2011.

1. We have an observation $y \in \mathbb{R}$ which we think is the observed value of a normally distributed random variable Y with mean zero and precision parameter $\theta > 0$, i.e.,

$$Y \mid \theta \sim N(0, \frac{1}{\theta}).$$

We have only vague prior information on the precision and therefore we use an improper prior, which is proportional to $1/\theta$ on the positive real axis.

a) Why do we call this an *improper* prior? (2 points)

b) Identify the posterior distribution. Is the posterior distribution proper for all possible observations y? (For the familiar distributions listed in the problem sheet, any parameter value outside the stated domain results in an imporper distribution.) (4 points)

2. We have n conditionally independent observations $y_1, \ldots, y_n$ from the Exp(θ) distribution, and we have carefully formulated a prior density $\pi(\theta), \theta > 0$ based on our prior knowledge. Now we want to reparametrize the model using

$$\phi = 1/\theta$$

as the new parameter.

a) Write the likelihood for the original parameter θ.

b) How can we calculate the likelihood as a function of the new parameter ϕ?

c) Find the prior density for ϕ implied by the original prior $\pi(\theta)$.

3. Consider the distribution with the unnormalized density

$$f^*(x) = \begin{cases} \dfrac{e^{-x}}{2 + (\sin(x))^2}, & \text{if } x > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Present some concrete method for simulating independently from this distribution. (Your fellow student who has not taken this course should be able to program your method based on your description.) You may use draws from any of the familiar distributions listed on this problem sheet.

4. We want to estimate the value of the integral

$$I = \int x\, f^*(x)\, \mathrm{d}x = \int_0^\infty \frac{x\, e^{-x}}{2 + (\sin(x))^2}\, \mathrm{d}x,$$

where f^* is defined as in the previous problem.

a) Explain how you can calculate an estimate $\hat{I}$ for I by Monte Carlo, where you use values drawn from the Exp(1) distribution.

b) Explain how you can calculate the Monte Carlo standard error of your estimate $\hat{I}$.

c) Explain how you can estimate the mean EX of a random variable X whose distribution has f^* as its unnormalized density, by using importance sampling.

Familiar distributions

Beta distribution $\mathrm{Be}(a, b)$ with parameters $a > 0, b > 0$ has pdf

$$\mathrm{Be}(x \mid a, b) = \frac{1}{B(a, b)}\, x^{a-1}(1-x)^{b-1}, \quad 0 < x < 1.$$

$B(a, b)$ is the beta function with arguments a and b,

$$B(a, b) = \int_0^1 u^{a-1}(1-u)^{b-1}\, \mathrm{d}u = \frac{\Gamma(a)\,\Gamma(b)}{\Gamma(a+b)}.$$

Exponential distribution $\mathrm{Exp}(\lambda)$ with rate $\lambda > 0$ has pdf

$$\mathrm{Exp}(x \mid \lambda) = \lambda\, \mathrm{e}^{-\lambda x}, \quad x > 0.$$

Gamma distribution $\mathrm{Gam}(a, b)$ with parameters $a > 0, b > 0$ has pdf

$$\mathrm{Gam}(x \mid a, b) = \frac{b^a}{\Gamma(a)}\, x^{a-1}\mathrm{e}^{-bx}, \quad x > 0.$$

$\Gamma(a)$ is the gamma function,

$$\Gamma(a) = \int_0^\infty x^{a-1}\mathrm{e}^{-x}\, \mathrm{d}x, \qquad a > 0.$$

It satisfies $\Gamma(a+1) = a\Gamma(a)$ for all $a > 0$, and $\Gamma(1) = 1$, from which it follows that $\Gamma(n) = (n-1)!$, when $n = 1, 2, 3, \ldots$

Normal distribution $N(\mu, \sigma^2)$ with mean μ and variance $\sigma^2 > 0$ has pdf

$$N(x \mid \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}}\, \exp\left(-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right).$$

Uniform distribution $\mathrm{Uni}(a, b)$ on the interval (a, b), where $a < b$, has pdf

$$\mathrm{Uni}(x \mid a, b) = \frac{1}{b-a}, \qquad a < x < b.$$

Binomial distribution $\mathrm{Bin}(n, p)$, n positive integer, $0 \le p \le 1$, has pmf

$$\mathrm{Bin}(x \mid n, p) = \binom{n}{x} p^x(1-p)^{n-x}, \quad x = 0, 1, \ldots, n.$$

Poisson distribution $\mathrm{Poi}(\theta)$ with parameter $\theta > 0$ has pmf

$$\mathrm{Poi}(x \mid \theta) = \mathrm{e}^{-\theta}\, \frac{\theta^x}{x!}, \quad x = 0, 1, 2, \ldots$$

Multivariate normal distribution (in d dimensions), $N_d(\mu, \Sigma)$ with mean $\mu \in \mathbb{R}^d$, covariance matrix Σ (a symmetric, positive definite $d \times d$ matrix) and precision matrix $Q = \Sigma^{-1}$ has pdf

$$\begin{aligned} N_d(x \mid \mu, \Sigma) &= N_d(x \mid \mu, Q^{-1}) \\ &= (2\pi)^{-d/2}(\det\Sigma)^{-1/2}\exp\left(-\frac{1}{2}(x-\mu)^T\Sigma^{-1}(x-\mu)\right) \\ &= (2\pi)^{-d/2}(\det Q)^{1/2}\exp\left(-\frac{1}{2}(x-\mu)^T Q(x-\mu)\right). \end{aligned}$$