



1. This figure shows the unnormalized density f^* and its envelope Mg from Section 3.5.3. I have also generated a sample of (x, y) pairs uniformly under the graph of the envelope Mg but plotted only those pairs which also lie under the graph of f^* . Produce a similar plot. (R hints: I calculated the envelope on a dense grid of points `t <- seq(-5, 5, length = 501)`, then did the initial plot with a call of the form `plot(t, env, 'l')` and added a line plot of f^* with a call of the form `lines(t, f.ast)` and finally added the simulated (x, y) pairs with a call of the form `points(x, y, pch = '.')`).

If you take the code in Section 3.5.3 as your starting point, you should figure out how the entities in the verbal description relate to the quantities calculated in the code.

2. Suppose that one uses the accept–reject algorithm for the unnormalized target f^* using the proposal density g and the constant M but the majorizing condition

$$f^*(x) \leq Mg(x)$$

does not hold in some region of the space. Then the algorithm does not simulate the distribution corresponding to f^* but a certain other distribution. Which distribution?

Hint: once you have digested the geometric reasoning behind the accept–reject method, you can solve this problem by sketching a figure, without making any calculations.

3. We want to simulate the standard normal $N(0, 1)$ truncated to values greater than $\alpha > 0$. The target distribution has the unnormalized density

$$f^*(x) = \exp\left(-\frac{1}{2}x^2\right) 1_{(\alpha, \infty)}(x).$$

When the truncation point α is large enough (try $\alpha = 9$ in R), the inverse transform method does not work anymore because of numerical problems. Nevertheless, simulating the truncated normal turns out to be quite easy, when we use the accept–reject method and utilize log-concavity. Proceed as follows.

The function $-x^2/2$ is concave, and so any tangent line gives an upper bound for it. Use the tangent line at $x = \alpha$, and exponentiate to get an upper bound for $\exp(-x^2/2)$. Multiply your upper bound with the indicator of the set (α, ∞) , and you get a nice upper bound for the function f^* . Your upper bound is of the form $Mg(x)$, where g is the density of a shifted exponential distribution. (I.e., the distribution of $k + Z$ where k is a constant and $Z \sim \text{Exp}(\lambda)$. What is the pdf of that?)

You end up with an accept–reject algorithm, which can be implemented by using generators for the $\text{Exp}(\lambda)$ and $\text{Uni}(0, 1)$ distributions. Write such an implementation (at least in pseudo code). Calculate the acceptance probability of your algorithm for some values of $\alpha > 0$, in particular for $\alpha = 5$ (for which the acceptance probability of the naive method is about $2.9 \cdot 10^{-7}$).

4. We consider a continuous joint distribution corresponding to the unnormalized density

$$h(x, y) = \begin{cases} (1-x)y(x-y)^2, & \text{when } 0 < y < x < 1, \\ 0, & \text{otherwise.} \end{cases}$$

Show how this joint distribution can be simulated using generators for standard univariate distributions (listed in the Appendix) and the multiplication rule.

Hint: the conditional distribution $Y \mid (X = x)$ is “obviously” that of a scaled beta variate $x \text{Be}(c, d)$ for a certain choice of constants c and d . (Here $x \text{Be}(c, d)$ means the distribution one obtains by multiplying a $\text{Be}(c, d)$ random variable with the constant x .)

5. (Multivariate t distribution as a scale mixture of multivariate normals.) Multivariate t distribution in \mathbb{R}^d with $k > 0$ degrees of freedom (a positive real number, not necessarily an integer), location parameter $\mu \in \mathbb{R}^d$ and shape parameter $\Sigma \in \mathbb{R}^{d \times d}$ (a symmetric and positive definite matrix) can be defined as the marginal distribution of the random vector X in the hierarchical model

$$X \mid Y \sim N_d\left(\mu, \frac{1}{Y} \Sigma\right) \quad \text{where} \quad Y \sim \text{Gam}(k/2, k/2).$$

(This means that we are given the joint pdf in the form $f_{X,Y}(x, y) = f_{X|Y}(x \mid y) f_Y(y)$.)

Show that the marginal density of X is the density $t_d(x \mid k, \mu, \Sigma)$ given in the Appendix.

Hint: integrate y out from $f_{X,Y}(x, y)$; the integral equals the normalizing constant of a certain gamma density. Note that

$$\det(s\Sigma) = s^d \det(\Sigma) \quad \text{and} \quad (s\Sigma)^{-1} = \frac{1}{s} \Sigma^{-1},$$

when s is a scalar.