Exercise 1 (chapter 3.4): Let the random variable X follow a Normal distribution with mean μ and variance σ^2. The density $g_X(x)$ of the truncated Normal distribution with support on the interval $I = (a, b)$ and $a < b$ is

$$g_X(x) = \frac{\phi(x; \mu, \sigma^2)}{\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)} 1_{(a,b)}(x) \,.$$

Derive the inverse transformation method to simulate from the truncated Normal distribution.

Solution: The cumulative distribution function of the truncated distribution is

$$G_X(x) = \int_{-\infty}^{x} g_X(t)\,\mathrm{d}t = \frac{1}{\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)} \int_{a}^{x} \phi(t; \mu, \sigma^2)\,\mathrm{d}t = \frac{\Phi\left(\frac{x-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)}{\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)} \,.$$

The quantile function is now straightforward to derive

$$G_X^{-1}(u) = \mu + \sigma\Phi^{-1}\left\{\Phi\left(\frac{a-\mu}{\sigma}\right) + u\left[\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)\right]\right\} \qquad 0 < u < 1 \,.$$

The inverse transform method is then

$$\begin{aligned} &\text{Simulate } U \sim \text{Unif}(0,1) \\ &\quad \text{Set } X = \mu + \sigma\Phi^{-1}\left\{\Phi\left(\frac{a-\mu}{\sigma}\right) + U\left[\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)\right]\right\} . \end{aligned}$$

Exercise 2 (chapter 3.5): Let $p^*(\theta \,|\, y)$ be an unnormalized posterior density of a parameter Θ and $g(\theta)$ the density of the proposal distribution in the accept-reject method. The conditional acceptance probability of a proposed value θ' is

$$\Pr(\text{accepted} \,|\, \theta') = \frac{p^*(\theta' \,|\, y)}{Mg(\theta')} \,,$$

where $M > 0$ is a known majorizing constant such that $p^*(\theta \,|\, y) \leq Mg(\theta)$ for all θ. Derive the unconditional acceptance probability. For univariate distributions, this happens to be equal to

$$\frac{\text{Area under } p^*(\theta \,|\, y)}{\text{Area under } Mg(\theta)} \,.$$

Solution: The unconditional acceptance probability is

$$\begin{aligned}
\Pr(\text{accepted}) &= \int \Pr(\text{accepted} \mid \theta) g(\theta)\, \mathrm{d}\theta \\
&= \int \frac{p^*(\theta \mid y)}{M g(\theta)} g(\theta)\, \mathrm{d}\theta \\
&= \frac{1}{M} \int p^*(\theta \mid y)\, \mathrm{d}\theta \\
&= \frac{\text{Area under } p^*(\theta \mid y)}{\text{Area under } M g(\theta)} \,.
\end{aligned}$$

Exercise 3 (chapter 3.5): Let $f_X(x)$ be the density of a random variable X with support on the interval $I = [a, b]$. Consider the special version of the accept-reject method:

1. Generate independently u_1 and u_2 from standard Uniform distributions.
2. Accept $x' = a + (b - a)u_1$ as a sample from the distribution of X if

   $$M u_2 \leq f_X(a + [b - a]u_1)\,,$$

   where $M = \max_{a \leq x \leq b} f_X(x)$.

Derive the unconditional acceptance probability and demonstrate why x' is a sample from the distribution of X with cumulative distribution function

$$F_X(x) = \int_a^x f_X(t)\, \mathrm{d}t\,.$$

Solution: The unconditional acceptance probability is

$$\begin{aligned}
\Pr(\text{accepted}) &= \int_0^1 \Pr(\text{accepted} \mid u_1) g(u_1)\, \mathrm{d}u_1 \\
&= \int_0^1 \frac{f_X(a + [b - a]u_1)}{M}\, \mathrm{d}u_1 \\
&= \frac{1}{M} \int_a^b f_X(x) \frac{1}{b - a}\, \mathrm{d}x \\
&= \frac{1}{M(b - a)}\,,
\end{aligned}$$

where the second line used the change-of-variables $x = a + (b - a)u_1$. The probability $\Pr(X \leq x \mid \text{accepted})$

is equal to $F_X(x)$, because

$$\begin{aligned}
\Pr(X \le x \,|\, \text{accepted}) &= \frac{\Pr(X \le x \cap \text{ accepted})}{\Pr(\text{accepted})} \\
&= M(b-a)\Pr\left(a + (b-a)U_1 \le x \cap U_2 \le \frac{f_X(a+[b-a]u_1)}{M}\right) \\
&= M(b-a)\Pr\left(U_1 \le \frac{x-a}{b-a} \cap U_2 \le \frac{f_X(a+[b-a]u_1)}{M}\right) \\
&= M(b-a)\int_0^{(x-a)/(b-a)} \mathrm{d}u_1 \int_0^{f_X(a+[b-a]u_1)/M} \mathrm{d}u_2 \\
&= (b-a)\int_0^{(x-a)/(b-a)} f_X(a+[b-a]u_1)\,\mathrm{d}u_1 \\
&= (b-a)\int_a^x f_X(t)\frac{1}{b-a}\,\mathrm{d}t \\
&= F_X(x)\,.
\end{aligned}$$

Exercise 4 (chapter 3.5): Let the random variable X and Y be independent and follow Uniform distributions on the interval $I = (-0.5, 0.5)$. The density of the random variable $W = X + Y$ is known as the triangular density

$$f_W(w) = 1 - |w|\,, \quad |w| < 1\,.$$

with support on the interval $I = (-1, 1)$. Use the special accept-reject method to simulate from the distribution of W.

Solution:

```
   target <- function( w ) { 1 - abs( w ) }
   a <- -1; b <- 1 ; nSamples <- 10000 ; nProposed <- 0
   w <- numeric( nSamples ); ii <- 0
   while( ii < nSamples ) {
      nProposed <- nProposed + 1
      x <- a + ( b - a ) * runif( 1 )
      if( runif( 1 ) < target( x ) ) {
         w[ ii <- ii + 1 ] <- x
      }
   }
# Estimated acceptance probability:  0.4987282
# Theoretical acceptance probability:  0.5
```

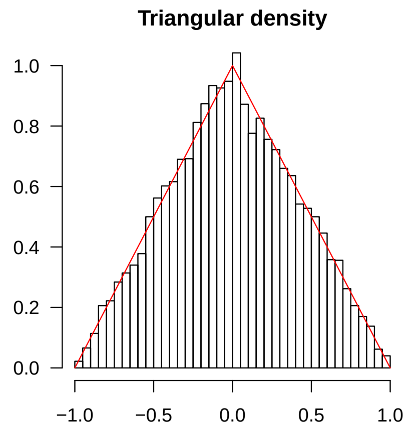

Exercise 5 (chapter 3.5):

1. Let the random variable X follow a standard Normal distribution. The density of the standard Normal distribution is

$$f_X(x) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{x^2}{2}\right\}.$$

Determine the value of the majorizing constant M in the accept-reject method using a standard Cauchy distribution as the proposal distribution. The density of the standard Cauchy distribution is

$$f_X(x) = \frac{1}{\pi(1+x^2)}.$$

2. Demonstrate that it is not possible to simulate from the standard Cauchy distribution using the accept-reject method with a standard Normal proposal distribution.

Solution: The accept-reject method requires

$$M \geq \frac{f_X(x)}{g_X(x)} = \frac{\pi\big(1+x^2\big)}{\sqrt{2\pi}} \exp\{-x^2/2\} = h(x)\,.$$

The least upper bound of $h(x)$ is

$$\frac{\mathrm{d}\log h(x)}{\mathrm{d}x} = -x + \frac{2x}{1+x^2} \stackrel{!}{=} 0 \Rightarrow x = 1$$
$$M \geq h(1) = \sqrt{\frac{2\pi}{e}}\,.$$

Using a standard Normal proposal distribution to simulate from the standard Cauchy distribution, requires

$$h(x) = \frac{\sqrt{2\pi}}{\pi(1+x^2)} \exp\big\{x^2/2\big\} \leq M\,.$$

However, an exponential function grows faster than a polynomial function as $x \to \infty$ so that there is no least upper bound M for any x.

```
  M <- sqrt( 2 * pi / exp( 1 ) )
  nSamples <- 10000 ; nProposed <- 0
  x <- numeric( nSamples ); ii <- 0
  while( ii < nSamples ) {
     nProposed <- nProposed + 1
     xProposed <- rcauchy( 1 )
     if( runif( 1 ) < dnorm( xProposed ) / M / dcauchy( xProposed ) ) {
        x[ ii <- ii + 1 ] <- xProposed
     }
  }
# Estimated acceptance probability:  0.6625588
# Theoretical acceptance probability:  0.6577446
```

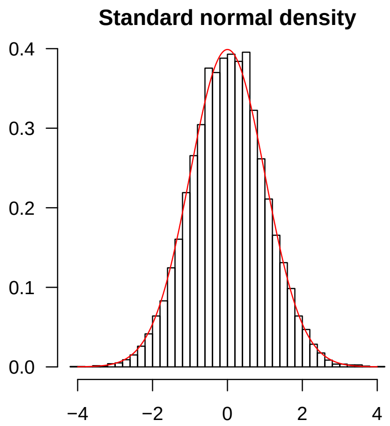

Exercise 6 (chapter 3.5):

1. Suppose that it is possible to compute the maximum likelihood estimate $\hat{\theta}_{\mathrm{MLE}}$ of a parameter Θ, that is,

   $$\hat{\theta}_{\mathrm{MLE}} \subseteq \left\{ \underset{\theta \in \Theta}{\operatorname{argmax}}\, p(y \,|\, \theta) \right\}.$$

   Show that the prior of Θ can be used as the proposal distribution in the accept-reject method to simulate from the posterior with unnormalized density

   $$p^*(\theta \,|\, y) = p(y \,|\, \theta) p(\theta)\,.$$

   Derive the acceptance condition of the accept-reject method.

2. Suppose that the likelihood $p(y \,|\, \theta)$ can be normalized to yield a so-called normalized likelihood. Let the prior density $p(\theta)$ be bounded and assume that direct simulation from the normalized likelihood is feasible. Show that the normalized likelihood of Θ can be used as the proposal distribution in the accept-reject method to simulate from the posterior. Derive again the acceptance condition of the accept-reject method.

Solution: The accept-reject method requires

$$h(\theta) = \frac{p(y \,|\, \theta)p(\theta)}{p(\theta)} \leq p(y \,|\, \hat{\theta}_{\mathrm{MLE}}) \leq M \,.$$

The acceptance condition with $M = p(y \,|\, \hat{\theta}_{\mathrm{MLE}})$ is therefore

$$U \leq \frac{p(y \,|\, \theta)p(\theta)}{Mp(\theta)} = \frac{p(y \,|\, \theta)}{p(y \,|\, \hat{\theta}_{\mathrm{MLE}})}$$

yielding the following algorithm

$$\begin{aligned} \text{Simulate } \theta \sim p(\theta) \text{ and } U \sim \mathrm{Unif}(0,1) \\ \text{Accept } \theta \text{ if } U \leq \frac{p(y \,|\, \theta)}{p(y \,|\, \hat{\theta}_{\mathrm{MLE}})} \,. \end{aligned}$$

If the prior density $p(\theta)$ is bounded, then

$$h(\theta) = \frac{p^*(y \,|\, \theta)p(\theta)}{p^*(y \,|\, \theta)} \leq \max p(\theta) \leq M \,,$$

where $p^*(y \,|\, \theta)$ is the normalized likelihood. The acceptance condition with $M = \max p(\theta)$ is therefore

$$U \leq \frac{p^*(y \,|\, \theta)p(\theta)}{Mp^*(y \,|\, \theta)} = \frac{p(\theta)}{\max p(\theta)}$$

yielding the following algorithm

$$\begin{aligned} \text{Simulate } \theta \sim p^*(y \,|\, \theta) \text{ and } U \sim \mathrm{Unif}(0,1) \\ \text{Accept } \theta \text{ if } U \leq \frac{p(\theta)}{\max p(\theta)} \,. \end{aligned}$$