Rejection sampling

Ege Rubak and Jesper Møller

Rejection sampling (or von Neumann sampling or acceptance-rejection method)

Rejection sampling is a general method for simulation of a random variable (or random vector) X with a pdf $f \propto f_0$, assuming there is another random variable (or random vector of the same dimension) Y with pdf $g \propto g_0$ so that $f_0(x) \leq M g_0(x)$ for all possible realizations x of X, where M is a positive number.

Rejection sampling requires only that we need to know g_0 and M, and it works as follows.

Repeat generating independently

- realizations $Y = y$ and $U = u$ until $u \leq f_0(y)/[M g_0(y)]$,
- and then return $X = y$ as a simulation from f.

The idea is that if it is hard to simulate directly from f, it should be simpler to simulate from g.

The algorithm works best if f and g_0 are approximately proportional.

If $f = c_1 f_0$ and $g = c_2 g_0$, where c_1 and c_2 are normalizing constants, then the probability of getting acceptance is

$$P(U \leq f(Y)/[M g_0(Y)]) = \cdots = c_2/(c_1 M),$$

so the best choice of M is $M = c_2/c_1$ (whose value may be unknown).

Target density

Assume that the target density f has support on $I = [-4, 4]$ where it is proportional to

$$f_0(x; a, b) = \exp(a(x - a)^2 - bx^4),$$

where a, b are known parameters (we will use $a = 0.4$ and $b = 0.08$).

```
f0 <- function(x, a=.4, b=.08){exp(a * (x - a)^2 - b * x^4)}
curve(f0(x), from = -4.5, to = 4.5)
```

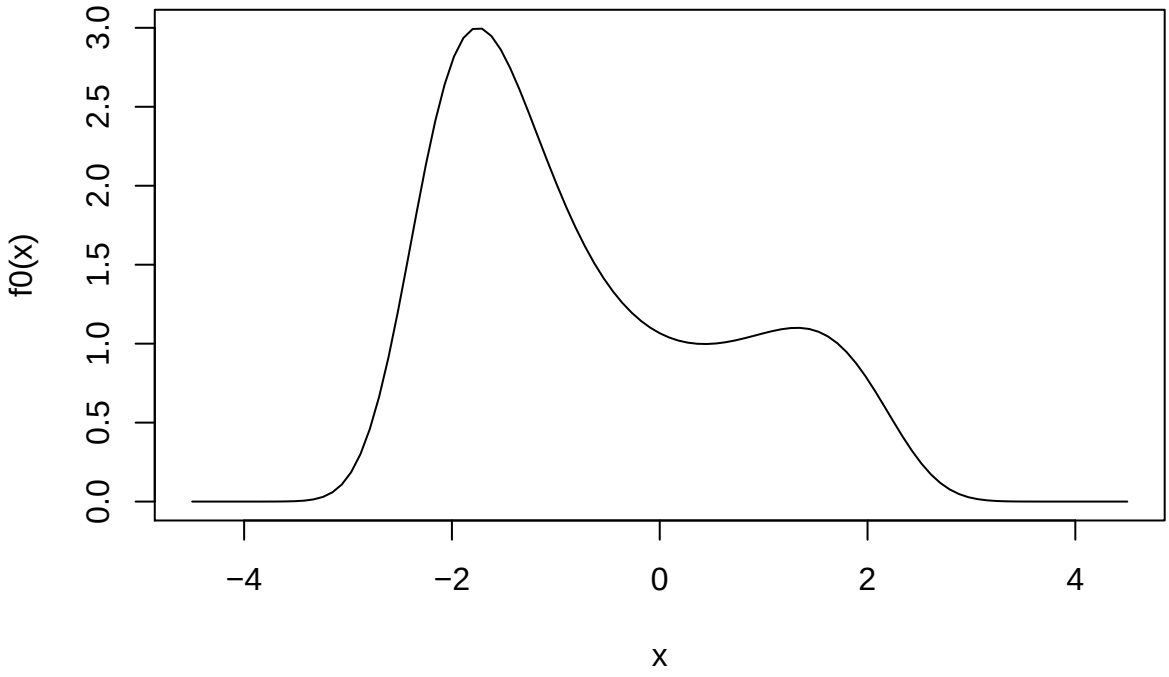

Proposal distribution

Assume we make uniform proposals on $I = [-4, 4]$, i.e. $g(x) = 1/8$ on I and zero outside I. Note that $f_0(x) \leq 3.1$ on I, so we can use $M = 3.1/(1/8) = 24.8$:

```
curve(f0(x), from = -4.5, to = 4.5)
M <- 24.8
g <- function(x){dunif(x, -4, 4)}
curve(M*g(x), add = TRUE, col = "red")
```

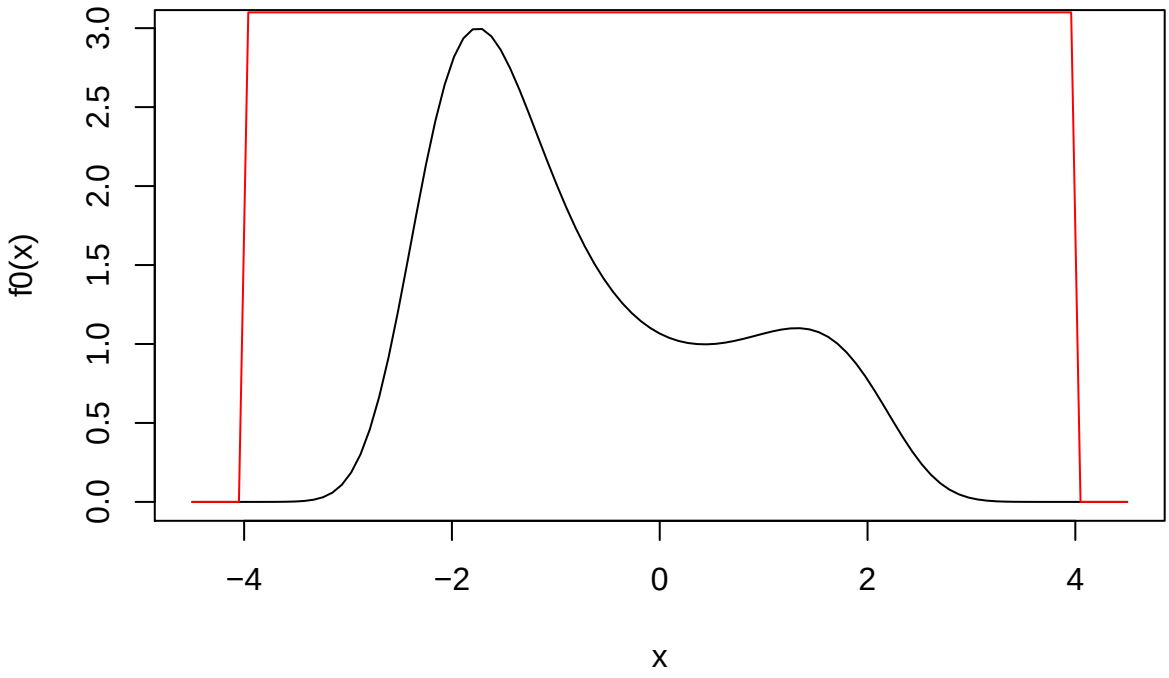

Running the algorithm

Now we can generate proposals from g (`runif()`) and note whether they are accepted as realizations from the target distribution or not:

```
N <- 10000
y <- runif(N, -4, 4)
p_accept <- f0(y)/(M*g(y))
u <- runif(N, 0, 1)
keep <- u<p_accept
head(round(cbind(y, percent = 100*p_accept, keep), 3))
```

```
##           y percent keep
## [1,] -3.729   0.006    0
## [2,]  0.106  33.395    0
## [3,] -0.326  39.795    1
## [4,] -1.224  77.404    1
## [5,]  2.119  20.955    0
## [6,]  1.902  27.922    0
```

Results

The resulting acceptance rate is

```
mean(keep)
```

```
## [1] 0.3142
```

A histogram of the samples:

```
hist(y[keep], prob = TRUE, col = "gray")
norm_const <- integrate(f0, -4, 4)$value
curve(f0(x)/norm_const, col = "red", add = TRUE)
```

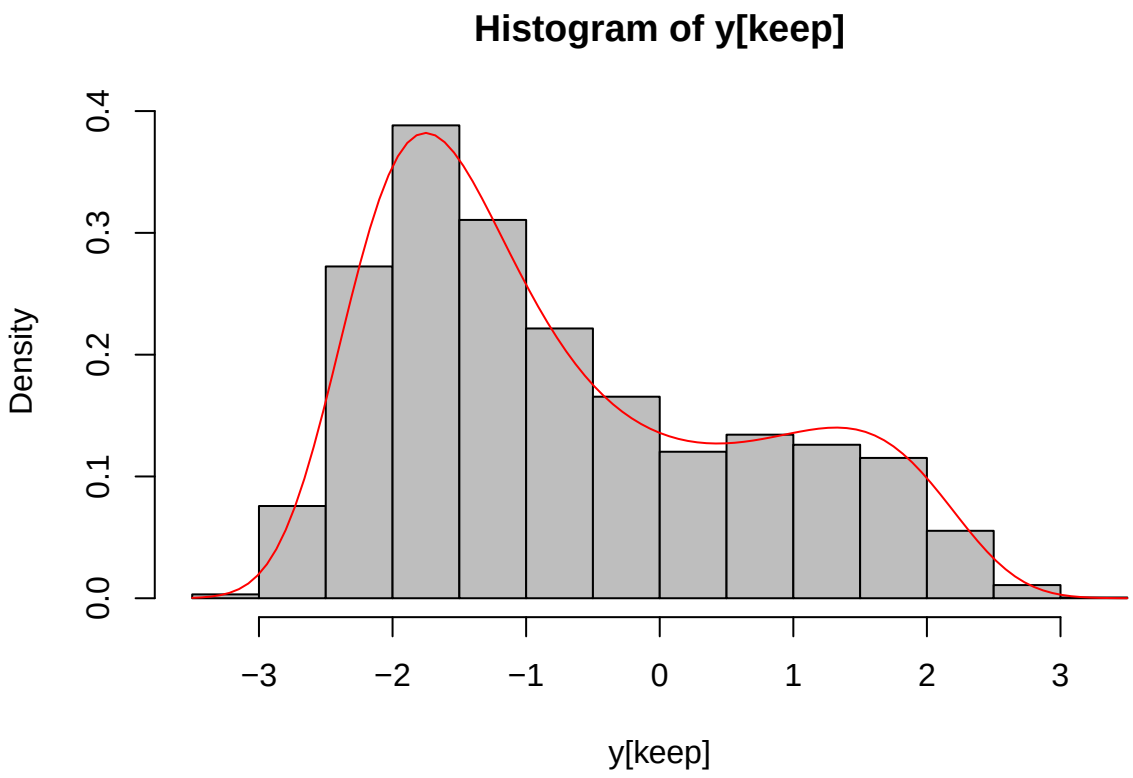
Histogram of y[keep]
Density
0.0
0.1
0.2
0.3
0.4
−3
−2
−1
0
1
2
3
y[keep]