Module 7: Supplementary exercises

Exercise: Potential rejection sampling problems

First try to answer the following questions without using the computer – then reuse the code from the supplementary slides to check your answer:

- Suppose we could not easily determine M and hence had to make a conservative choice; say M = 100 or M = 500 in this context.
    1. Which effect will that have on the number of accepted samples?
    2. How would you have to compensate for a too large value of M if you want a given number of samples from the target distribution?
- What happens if you do not choose M large enough (e.g. M = 10 in our example)?
- What would be the effect of using a uniform proposal distribution on $[-10, 10]$?
- What happens if the proposal distribution is an standard normal distribution (i.e. mean zero and standard deviation 1? Hints:
    1. You can use `dnorm()` for the normal density.
    2. You may have to create a sequence `x <- seq(-4, 4, by = 0.01)` to numerically evaluate the bound M relating `f0(x)` and `dnorm(x)`.

Exercise: Improving the proposal distribution

If $f(x)$, $x \in [0, 1]$ is a pdf on $[0, 1]$ then for $a > 0$, $1/a \cdot f(x/a)$, $x \in [0, a]$ is a pdf on $[0, a]$. Furthermore, a pdf on $[b, a + b]$ can be obtained by simple translation.

- Based on these facts how can a beta distribution Be(α, β) indirectly be used as the proposal distribution for our example? -Implement the rejection sampling algorithm using Be(2.5,3.5) transformed to $[-4.1, 4.1]$ (but with M determined on $[-4, 4]$).
- Check with a histogram that you are sampling the correct distribution.
- Find the acceptance rate.