Group Lab 5 – Metropolis-Hastings

Here is a Metropolis-Hastings sampler for a (non-normalized) probability distribution function f:

```
theta <- rep(0,1000)          # space to store results
t <- 0                        # must pick a starting t with f(t) > 0
for (step in 1:1000) {
  t_proposed <- t + rnorm(1)
  acceptance <- f(t_proposed)/f(t)
  if (runif(1) < acceptance) { t <- t_proposed }
  theta[step] <- t
}
```

1. Let $f(x) = \begin{cases} 1 - x^2 & -1 \leq x \leq 1 \\ 0 & \text{otherwise} \end{cases}$.
   (a) What normalization constant C would make $Cf(x)$ into a PDF?
   (b) Set up and run M-H to sample θ distributed as f.
   (c) Make a density plot of θ and check that it matches the function $f(x)$.
   (d) Make a trace plot of your simulated θ values.
   (e) Estimate the standard deviation of θ.

2. The M-H sampler you used for question 1 has a proposal distribution given by adding `rnorm(1)`, a standard normal value with mean 0 and sd 1.
   (a) Use a proposal SD of 10 (add `rnorm(1,0,10)`), so the proposed values are taking much larger jumps. Repeat 1b,c,d. Explain what happened.
   (b) Use a proposal SD of 0.1, so the proposed values are taking much smaller jumps. Repeat 1b,c,d. Explain what happened.

3. Let $f(x) = \begin{cases} 1 & 0 \leq x \leq 1 \\ 3 & 10 \leq x \leq 11 \\ 0 & \text{otherwise} \end{cases}$.
   (a) In R, you can use `dunif(x,0,1)+3*dunif(x,10,11)` to implement f (why?)
   Run the M-H sampler for f.
   (b) How do the starting value of `t` and the proposal sd affect the results?