

Module 9: Exercises

Exercise 1: Minimise auto-correlation

Write a function in R that uses a Metropolis-Hastings algorithm to simulate from a standard normal distribution, using normally distributed proposals centred at the current value (in other words, this is a random walk Metropolis algorithm). Make sure your code allows you to adjust the standard deviation of the proposal distribution and monitor the acceptance probability. (**Hint:** You already implemented this in the exercises for module 7 (see the solution if you didn't finish the exercise) and you may use this as the main part of your function). The skeleton of such a function is given below:

```
myMH <- function(N, sigma){
  x <- rep(0, N) # Empty Markov chain
  a <- 0 # Number of accepted proposals
  # Below implement the M-H alg. and add 1 to `a` every time you accept
  #####
  # YOUR CODE HERE #
  #####
  # Finally return a list of the results
  return(list(x=x, a=a/(N-1)))
}
# A short run with proposal std. dev. 1
rslt <- myMH(10, 1)
```

To plot the auto-correlation function use the function `acf()` in R. To approximately calculate $V = 1 + 2 \sum_{m=1}^{\infty} \rho_m$ use the following command in R:

```
V <- 2 * sum(acf(rslt$x, plot = FALSE)$acf) - 1
```

Make some experiments and find the standard deviation for the proposal distribution that gives the smallest value of V . What is the acceptance probability in this case? (Make sure that you have run the MH-algorithm long enough (really long), so that the auto-correlation function is well estimated.)

Exercise 2: Beetles

This example is concerned with different beetles exposed in 5 hours to different concentrations of carbon disulphide. The data contains for 481 beetles, in the first column, the concentration, and in the second column, the status of the beetles after the exposure. The status is coded as zeros (dead beetles) and ones (survived beetles). Read in the data set `beetles.dat` from the course website:

```
data_url <- "https://asta.math.aau.dk/course/bayes/2018/?file=beetles.dat"
beetles <- read.table(data_url)
```

Let $x_i > 0$ and $y_i \in \{0, 1\}$ denote the concentration and status for the i th beetle, respectively. We use a logistic regression model for the data of survived and dead beetles, i.e., we condition on the concentrations and assume that the probability that the i th beetle survives its given dose x_i is

$$\frac{\exp(a + bx_i)}{1 + \exp(a + bx_i)}$$

where a and b are real parameters. In R you can define the data vectors and log-likelihood function as:

```

beet_x <- beetles[,1]
beet_y <- beetles[,2]
llik <- function(a, b){
  ## BEWARE: We refer to global variables beet_x and beet_y here!
  return(sum(beet_y*(a+b*beet_x)) - sum(log(1+exp(a+b*beet_x))))
}

```

Now, let $\text{Cauchy}(\gamma)$ denote the symmetric Cauchy distribution with scale parameter $\gamma > 0$; this distribution has density

$$\pi(c|\gamma) = \frac{1}{\pi} \frac{\gamma}{\gamma^2 + c^2}, \quad c \in \mathbf{R}.$$

Assume a priori that $a \sim \text{Cauchy}(10)$ and $b \sim \text{Cauchy}(2.5)$ are independent (the argument for using this prior is given in a 2008 paper by Andrew Gelman and coauthors in *Annals of Applied Statistics*). Then construct a Metropolis-Hastings algorithm which generates a sample from the posterior distribution of (a, b) .

Hint: You have to work with the **logarithm** of the Hastings ratio for numerical stability. The logarithm of the prior density can be calculated as

```

lprior <- function(a, b){
  dcauchy(a, scale = 10, log = TRUE) + dcauchy(b, scale = 2.5, log = TRUE)
}

```

Exercise 3: The banana

Construct a Metropolis-Hastings algorithm which has the following (unnormalised) invariant two-dimensional density

$$\pi(x, y) \propto \exp(-x^2/200 - (y + 0.1 * x^2 - 10)^2) \quad x, y \in \mathbf{R}.$$

Notice that $\pi(x, y)$ is symmetric in x . Do your simulations reproduce this symmetry?