

Computational Statistics. Chapter 5: MCMC. Solution of exercises

Thierry Denoeux

2024-03-11

```
set.seed(2021)
```

Exercise 1

As the density of ϵ is symmetric, the MH ratio is the ratio of the densities at x^* and $x^{(t-1)}$, i.e., we have

$$R(x^{(t-1)}, x^*) = \frac{f(x^*)}{f(x^{(t-1)})} = \exp(|x^{(t-1)}| - |x^*|).$$

The following function `MH_Laplace` implements the random walk MH algorithm for this problem:


```
MH_Laplace <- function(N, sig){
  x<-vector(N,mode="numeric")
  x[1]<-rnorm(1,mean=0,sd=sig)
  for(t in (2:N)){
    epsilon<-rnorm(1,mean=0,sd=sig)
    xstar<-x[t-1]+ epsilon
    U<-runif(1)
    R<-exp(abs(x[t-1]) - abs(xstar))
    if(U <= R) x[t]<-xstar else x[t]<-x[t-1]
  }
  return(x)
}
```

Let us generate a sample of size 10^5 with $\sigma = 10$:

```
x <- MH_Laplace(100000,10)
```


The sample path and correlation plots show good mixing (the chain quickly moves away from its starting value, and the autocorrelation decreases quickly as the lag between iterations increases):

```
plot(x,type="l",xlab='t',ylab=expression(x[(t)]))
```



```
acf(x,lag.max=100)
```


Series x



Plot of the histogram with the Laplace density:

```
u<-seq(-10,10,0.01)
fu<-0.5*exp(-abs(u))
hist(x,freq=FALSE,ylim=range(fu))
lines(u,0.5*exp(-abs(u)))
```

Histogram of x




Let us now generate another sample of the same size, this time with $\sigma = 0.1$:

```
x<-MH_Laplace(100000,0.1)
```


This time, the sample path and correlation plots show poor mixing (the chain remains at or near the same value for many iterations, and the autocorrelation decays very slowly):

```
plot(x,type="l",xlab='t',ylab=expression(x[(t)]))
```



```
acf(x, lag.max=1000)
```

Series x



Exercise 2

Question a

The likelihood function is

$$L(\beta; y_1, \dots, y_n) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(y_i - \beta x_i)^2\right) = (2\pi)^{-n/2} \exp\left(-\frac{1}{2} \sum_{i=1}^n (y_i - \beta x_i)^2\right).$$

The density of the Gamma distribution with shape parameter a and rate b is $f(\beta) \propto \beta^{a-1} \exp(-b\beta) I(\beta > 0)$. Here $a = 2$ and $b = 1$, so $f(\beta) \propto \beta \exp(-\beta) I(\beta > 0)$. Consequently, the posterior density is

$$f(\beta | y_1, \dots, y_n) \propto \beta \exp(-\beta) \exp\left(-\frac{1}{2} \sum_{i=1}^n (y_i - \beta x_i)^2\right) I(\beta > 0).$$

Question b

We first write a function that computes the log-likelihood:

```
loglik <- function(beta,x,y){  
  n<- length(x)  
  return(-0.5 * sum((y-beta*x)^2) - n/2*log(2*pi))  
}
```

We then write a function that generates a MC of size N for a given data set:

```

gen_MH<-function(x,y,N){
  beta <- vector(N,mode="numeric")
  beta[1] <- rgamma(1,shape=2,rate=1)
  for(t in (2:N)){
    beta_star <- rgamma(1,shape=2,rate=1)
    u <- runif(1)
    logR <- loglik(beta_star,x,y)-loglik(beta[t-1],x,y)
    if( log(u) <= logR ) beta[t] <- beta_star   else beta[t]<- beta[t-1]
  }
  return(beta)
}

```

Question c

Data generation:

```

beta0<- rgamma(1,shape=2,rate=1) # Generation of beta
cat(beta0)

```

```
## 1.085444
```

```

n<-50
x<-rnorm(n)
y<-x*beta0+rnorm(n)


```

Plot of the data:

```

plot(x,y)
abline(0,beta0)

```



Running the MH algorithm:

```


N <- 100000
beta <- gen_MH(x,y,N)

```

Question d


Sample path:

```
plot(beta,type="l",xlab="iterations")
```



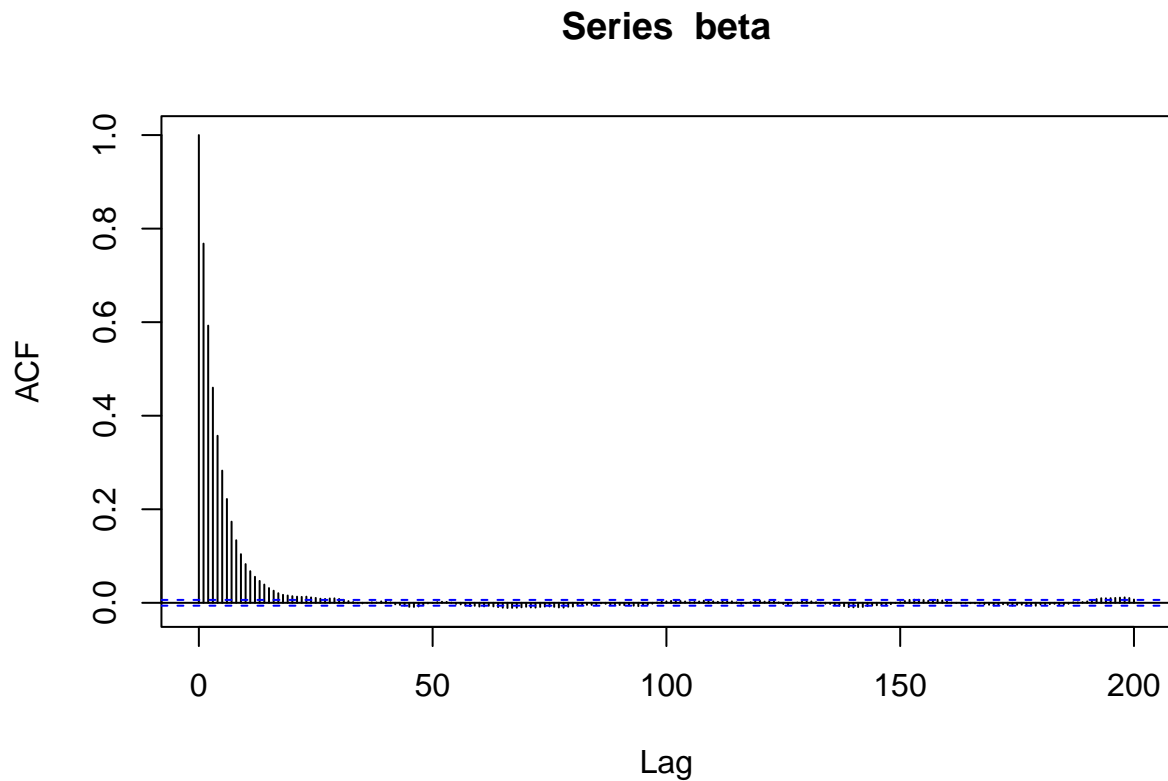
Histogram (leaving out the first 500 values):

```
hist(beta[500:N],xlab=expression(beta))
```



Autocorrelation plot:

```
acf(beta, lag.max=200)
```



Question e

We use the batch means method. We first determine the lag k_0 such that the autocorrelation is small enough to be neglected:

```
ACF<-acf(beta, lag.max=200, plot=FALSE)
k0<-ACF$lag[min(which(abs(ACF$acf)<0.01))]
cat(k0)
```

```
## 26
```

We fix the burn-in period and we compute the number of batches:

```
D<-1000 # burn in
B<-floor((N-D)/k0)
```

We compute the means within each block:

```
Z<-vector(B, mode="numeric")
for(b in (1:B)) Z[b]<-mean(beta[(D+(b-1)*k0+1):(D+b*k0)])
```

The estimated simulation standard error is the standard deviation of the batch means divided by the square root of the number of batches:

```
se <- sd(Z)/sqrt(B)
```

Estimated posterior expectation of β and simulation standard error: