

# Math 365 Rejection Sampling and MCMC

Submit your R script to [tleise@amherst.edu](mailto:tleise@amherst.edu) by next Tuesday.

In this lab we will look at rejection sampling and the Metropolis-Hastings algorithm for Markov chain Monte Carlo simulations.

## 1 Monte Carlo Integration

We start with the basic idea behind Monte Carlo simulations. Suppose we want to approximate  $\pi$ . Imagine throwing lots of darts at a 1 meter by 1 meter square board. Draw a quarter circle of radius 1 on the board. If your aim is uniformly random so that you are equally likely to hit any point on the square, then after throwing many many darts, you would expect the proportion of darts landing inside the quarter circle to be equal to its area,  $\pi/4$ . The more darts you throw, the better the approximation of  $\pi$  you will tend to get.

**Exercise 1** Simulate this experiment in R by generating many points on the unit square, then counting how many landed in the unit circle. Multiply by 4 to obtain an estimate of  $\pi$ . How good is your approximation if you throw 100 darts? 1,000 darts? 10,000 darts? 100,000 darts?

```
Ndarts <- 1000
x <- runif(Ndarts) # uniform on [0,1]
y <- runif(Ndarts)
theta <- seq(0, pi/2, .01) # to draw circle
plot(cos(theta), sin(theta), col="red", type="l",
      xlim=c(0,1), ylim=c(0,1))
points(x,y,pch=".") # darts on board
r2 <- x^2+y^2
sum(r2<=1)/Ndarts*4 # estimate of pi
```

## 2 Rejection Sampling

Now suppose you want to generate random numbers from an unusual probability density function, for example, something strange like  $f(x) = \frac{x(1-x)e^x}{3-e}$  on  $[0, 1]$ . We want a method that uses random numbers from a density  $h(x)$  we know how to calculate, e.g., uniformly distributed random numbers with  $h(x) = 1$  on  $[0, 1]$ , and outputs random numbers according to the desired density  $f(x)$ . Rejection sampling is one way to accomplish this. First find a constant  $c$  such that  $f(x) \leq ch(x)$  for all  $x \in [0, 1]$ , then follow these steps:

1. Generate a random number  $x$  with the density  $h(x)$ .
2. Generate a uniformly distributed random number  $u$  from the interval  $[0, 1]$ . If  $u \leq \frac{f(x)}{ch(x)}$ , then output  $x$ ; otherwise reject  $x$  and return to step 1.

This procedure uses the same underlying idea as the darts example above: Throw many darts at a board that is  $c$  high and with base  $[0, 1]$  on which we have drawn our desired density  $f(x)$ . Reject the darts that land above the curve, and accept those that land under it. There will be more darts near values of  $x$  for which  $f(x)$  is large than where  $f(x)$  is small, so with enough darts you will approximate the density function.

**Exercise 2** Use rejection sampling to generate a set of random numbers according to the density function  $f(x) = \frac{x(1-x)e^x}{3-e}$  on  $[0, 1]$ .

```
f <- function(x) x*(1-x)*exp(x)/(3-exp(1))
t <- seq(0,1,0.01)
plot(t,f(t),type="l") # graph of the desired density function

rdist <- function(c) { # choose c with f(x)<=c for all x
for (i in 1:1000) {
x <- runif(1,0,1) # step 1
u <- runif(1,0,1) # step 2
ifelse(c*u<=f(x),return(x),-1)
}
return(-1) # if tried 1000 times and rejected all
}
```

The command `rdist(c)`, where you put an appropriate value in for  $c$  based on the graph of  $f(x)$ , will generate a single random number. Generate 10,000 such random numbers and plot a histogram using `freq=FALSE` (so it plots the density). Add a plot of  $f(x)$  using the `points` command in R to check that the histogram indeed closely matches the desired density function.

### 3 Metropolis-Hastings Algorithm

Suppose you want to generate a discrete-time Markov chain whose invariant distribution is a given vector  $\bar{\pi}$ . The Metropolis-Hastings algorithm takes any Markov chain and uses it to generate the desired Markov chain. We'll look at the algorithm in this lab, and go over the theory of why it works in lecture during a later class.

To illustrate the algorithm, let  $S = \{0, 1, 2, 3\}$ . Let the desired invariant probability vector be  $\bar{\pi} = [0.1 \ 0.2 \ 0.3 \ 0.4]$ . That is, we want the chain to be at state 0 10% of the time, at state 1 20% of the time, etc. Take  $P$  to be the transition matrix of any convenient Markov chain, for example, random walk on the circle with  $p = \frac{1}{2}$ . To generate a new Markov chain, we need to specify the transitions. Suppose  $X_m = i$ . The next state  $X_{m+1}$  is determined by a two-step procedure:

1. Choose a proposal state  $j$  using the transition probabilities  $P(i, j)$ . For the random walk on a circle, the proposal state  $j$  will be either  $i - 1$  or  $i + 1 \pmod{4}$ .
2. Decide whether or not to accept the proposal state, according to the acceptance function

$$a(i, j) = \frac{\pi(j)P(j, i)}{\pi(i)P(i, j)}.$$

Generate a uniformly distributed random number  $u$  from the interval  $[0, 1]$ . The next state of the chain is

$$X_{m+1} = \begin{cases} j, & \text{if } u \leq a(i, j) \\ i, & \text{if } u > a(i, j). \end{cases}$$

In this simple case, since  $p = \frac{1}{2}$ ,  $P(j, i) = P(i, j)$ , so the acceptance function simplifies to  $a(i, j) = \frac{\pi(j)}{\pi(i)}$ . Hence the next step will be  $j$  if  $u \leq \frac{\pi(j)}{\pi(i)}$ ; otherwise it is again  $i$ .

```
pi <- c(0.1, 0.2, 0.3, 0.4)
```

```
mcmc <- function(n) { # n=number of steps to take
  currentstate <- 0 # initial state X0=0
  for (i in 1:n) {
    proposal <- (currentstate + sample(c(-1,1),1)) %% 4 # mod 4
    accept <- pi[proposal+1]/pi[currentstate+1]
    u <- runif(1,0,1)
    if (u < accept) currentstate <- proposal }
  currentstate }
```

The `mcmc(n)` function outputs  $X_n$ . We can choose, for example,  $n = 100$  and examine the distribution of  $X_{100}$  from 10,000 simulations:

```
trials <- 10000
simlist <- replicate(trials,mcmc(100))
table(simlist)/trials
```

**Exercise 3** Run the MCMC simulation and compare your observed distribution to the desired distribution  $\bar{\pi}$ .