Ma 322: Biostatistics
Homework Assignment 5

Prof. Wickerhauser

Read Chapter 10, "Stochastic Processes and Markov Chains," pages 160–184 of our text.

Note: Although our text has no index or table of contents, it is easy to locate words in the electronic version using the Find function of your favorite PDF reader.

1. Suppose that the list of variables $p = (p_1, \ldots, p_K)$ has a Dirichlet prior density

$$f_\alpha(p) \propto p_1^{\alpha_1 - 1} \cdots p_K^{\alpha_K - 1},$$

where $\alpha = (\alpha_1, \ldots, \alpha_K)$ is a list of shape parameters. We perform an experiment that yields counts $n = (n_1, \ldots, n_K)$ having the multinomial likelihood

$$L_n(p) \propto p_1^{n_1} \cdots p_K^{n_K}.$$

(a) For what values of α does one get a non-informative Dirichlet prior pdf?

(b) Determine the shape parameters for the posterior pdf $L_n(p) f_\alpha(p)$.

Solution: (a) Shape parameters $\alpha_1 = \cdots = \alpha_K = 1$, namely $\alpha = (1, \ldots, 1)$, yield the uniform pdf $f_\alpha(p) = 1$.

(b) The exponents combine to give the posterior pdf

$$L_n(p) f_\alpha(p) \propto p_1^{n_1 + \alpha_1 - 1} \cdots p_K^{n_K + \alpha_K - 1},$$

so the posterior Dirichlet pdf has shape parameters $\alpha + n = (\alpha_1 + n_1, \ldots, \alpha_K + n_k)$. $\square$

2. Suppose that 100 individuals selected randomly from a population are blood-typed and the results are 55 type O, 25 type A, 15 type B, and 5 type AB.

(a) Using a non-informative prior, and assuming Hardy-Weinberg equilibrium, generate a contour plot of the posterior pdf on the proportions p_A and p_B of the A and B blood-type alleles, respectively, in the population. HINT: see `r-eg-35.txt` on the class website.

(b) Find, at least approximately, the maximum-likelihood estimator of the proportion of A,B, and O alleles in the population.

Solution: (a) Use the following code modified from `r-eg-35.txt`

```
pA <- seq(0,1, by=0.01); pB <- seq(0,1, by=0.01)
z <- matrix(0, nrow=length(pA), ncol=length(pB))
nA <- 25; nB<-15; nO<-55; nAB<-5; # experimental data
```

```
    for(i in 1:length(pA)) {
        for(j in 1:length(pB) ) {
           a <- pA[i]; b <- pB[j];
           if( a+b < 1) {   # Otherwise not in the domain.
             c <- 1-a-b      # Shorthand for pO=1-pA[i]-pB[j].
             sA <- a**2 + 2*a*c; sB <- b**2 + 2*b*c; sO <- c**2; sAB <- 2*a*b;
             z[i,j] <- sA** nA * sB**nB * sO**nO * sAB**nAB
             } else { z[i,j] <- 0 }
    }
    }
```

(Note: "**" is the same as exponentiation with a carat ˆ.) Then the requested plot is produced by

```
pdf("hw5ex2a.pdf"); contour(pA,pB,z,xlab="pA",ylab="pB"); dev.off()
```

(b) Judging the location of the peak by eye gives $p_A \approx 0.15$, $p_B \approx 0.10$, and so $p_O = 1 - p_A - p_B \approx 0.75$.
□

3. Implement the function `Walk1d()` on p.167 of our text and graph three 100-step simulations starting from three random seeds: your student ID, the year of your birth, and the last four digits of your favorite telephone number.

   Solution: Modify the function to make specifying the seeds and printing the plots a bit easier:

```
Walk1d<-function(n=100, seed=NULL) {
 if( !is.null(seed) ) set.seed(seed);
   y<-vector(length=n); y[1]<-0;
   for(i in 2:n) y[i]<-y[(i-1)]+sample(c(-1,1),1);
   plot(1:n,y,type='l',ylim=c(-20,20));  }
```

   For this model solution the three seeds will be 123, 4567, and 89120. Then the three requested plots are produced by

```
 pdf('w1d123.pdf'); Walk1d(seed=123); dev.off();
 pdf('w1d4567.pdf'); Walk1d(seed=4567); dev.off();
 pdf('w1d89012.pdf'); Walk1d(seed=89012); dev.off();
```

   □

4. Implement the function `Walk2d()` on p.168 of our text and graph three 500-step simulations starting from the same random seeds you used in the previous problem.

   Solution: Modify the function to make specifying the seeds and printing the plots a bit easier:

```
        Walk2d<-function(n=500, seed=NULL) {
          if( !is.null(seed) ) set.seed(seed);
          x0<-0; x<-x0+cumsum(sample(c(-1,1),n,replace=TRUE));
          y0<-0; y<-y0+cumsum(sample(c(-1,1),n,replace=TRUE));
          plot(x,y,xlim=c(-40,40),xlab='x',ylim=c(-40,40),ylab='y',type='l');}
```

Then the three requested plots are produced by

```
pdf('w2d123.pdf'); Walk2d(seed=123); dev.off();
pdf('w2d4567.pdf'); Walk2d(seed=4567); dev.off();
pdf('w2d89012.pdf'); Walk2d(seed=89012); dev.off();
```

□

5. A restless koala moves among three eucalyptus trees labeled 1, 2, and 3. A patient park ranger watches and makes notes every morning and evening on the koala's position, producing the following table:

Koala Tree-Change Counts

Morning Tree	Evening Tree	Count
1	1	4
1	2	10
1	3	7
2	1	11
2	2	14
2	3	12
3	1	6
3	2	13
3	3	12

(a) Treat the koala's movements as a Markov process and determine the transition matrix M from this table of counts.

(b) Starting with a uniform prior distribution on the three trees and assuming the koala's tree-change preferences remain the same, compute the posterior koala distribution, namely the stationary distribution determined by M.

Solution: (a) Input the table into a matrix of counts as follows:

```
M<-matrix(0,nrow=3,ncol=3);
M[1,1]<-4; M[1,2]<-10; M[1,3]<-7;
M[2,1]<-11; M[2,2]<-14; M[2,3]<-12;
M[3,1]<-6; M[3,2]<-13; M[3,3]<-12;
```

The row is the starting location and the column is the ending location for each move, just as for the frog situation on p.172 of our text.

To get a matrix of transition probabilities, the matrix elements must be divided by their respective row sums:

```
for(i in 1:3) M[i,] <- M[i,]/sum(M[i,])
```

This produces the transition matrix

$$M = \begin{pmatrix} 3/21 & 12/21 & 8/21 \\ 10/37 & 5/37 & 20/37 \\ 13/31 & 4/31 & 7/31 \end{pmatrix} = \begin{pmatrix} 0.1904762 & 0.4761905 & 0.3333333 \\ 0.2972973 & 0.3783784 & 0.3243243 \\ 0.1935484 & 0.4193548 & 0.3870968 \end{pmatrix}$$

(b) Take powers of M applied (on the right) to the uniform prior $(1/3, 1/3, 1/3)$ until the result stops changing in the first few decimal places (10,000 iterations is more than we need, but the computer will not complain):

```
t0<-c(1,1,1)/3; t<-t0; for(i in 1:10000) t<-t%*%M; t
```

The result is `0.2359551, 0.4157303, 0.3483146`, so the koala may be expected to spend 23.6% of the time in tree 1, 41.6% of the time in tree 2, and 34.8% of the time in tree 3.

Note: Taking powers of just M until its coefficients stop changing gives a matrix whose rows are identical copies of this stationary distrubution. This may be done by the following R commands:

```
MM<-M; for(i in 1:10000) MM<-M %*% MM;   MM
```

□

6. Consider the following transition matrix for a 4-state Markov chain:

$$M = \begin{pmatrix} 0.2 & 0.1 & 0.4 & 0.3 \\ 0.1 & 0.2 & 0.3 & 0.4 \\ 0.3 & 0.4 & 0.1 & 0.2 \\ 0.4 & 0.2 & 0.3 & 0.1 \end{pmatrix}.$$

(a) Is M periodic or aperiodic?

(b) Is M irreducible?

(c) Is M ergodic?

(d) Does M have a stationary distribution?

(e) Is M reversible?

Solution: First use the following R code to test if $\lim_{n\to\infty} M^n$ exists:

```
data<-c(2,1,3,4,1,2,4,2,4,3,1,3,3,4,2,1)/10;
M<-matrix(data,4,4); MM<-M; M
for(i in 1:100) MM <- M %*% MM;
MM
```

The result is

$$M^\infty \overset{\text{def}}{=} \lim_{n\to\infty} M^n = \begin{pmatrix} 0.2535211 & 0.2288732 & 0.2711268 & 0.2464789 \\ 0.2535211 & 0.2288732 & 0.2711268 & 0.2464789 \\ 0.2535211 & 0.2288732 & 0.2711268 & 0.2464789 \\ 0.2535211 & 0.2288732 & 0.2711268 & 0.2464789 \end{pmatrix}.$$

(a) The existence of the limit of M^n as $n \to \infty$ implies that M is aperiodic.

(b) M has all nonzero elements, so M is irreducible.

(c) Since M is aperiodic and irreducible, it is ergodic.

(d) M is ergodic by part c, and from M^∞ we conclude that M has a stationary state

$$\pi = (0.253521, 0.2288732, 0.2711268, 0.2464789).$$

(e) Check the 16 detailed balance equations with this code:

```
pp <- MM[1,]; pp  # stationary distribution: any row of MM
dif <- matrix(0,4,4);  # initialize to all zeros
for ( i in 1:4 ){
 for(j in 1:4){
  dif[i,j] <- pp[i]*M[i,j]-pp[j]*M[j,i]  # reversible iff all 0s
}};  dif
```

The final output is the matrix `dif` of differences:

```
             [,1]          [,2]         [,3]         [,4]
[1,]  0.000000000  0.002464789  0.02007042 -0.02253521
[2,] -0.002464789  0.000000000 -0.03978873  0.04225352
[3,] -0.020070423  0.039788732  0.00000000 -0.01971831
[4,]  0.022535211 -0.042253521  0.01971831  0.00000000
```

This shows in particular that $\mathtt{dif[1,3]} = \pi[1]M[1,3] - \pi[3]M[3,1] \approx 0.02 \neq 0$, so M is not reversible. But any choice of $i \neq j$ would yield the same conclusion. □

7. Consider the following transition matrix for a 5-state Markov chain:

$$F = \begin{pmatrix} 0.4 & 0.3 & 0.2 & 0.1 & 0.0 \\ 0.0 & 0.4 & 0.3 & 0.2 & 0.1 \\ 0.0 & 0.0 & 0.5 & 0.3 & 0.2 \\ 0.0 & 0.0 & 0.0 & 0.6 & 0.4 \\ 0.0 & 0.0 & 0.0 & 0.0 & 1.0 \end{pmatrix}.$$

(a) Is F periodic or aperiodic?

(b) Is F irreducible?

(c) Is F ergodic?

(d) Does F have a stationary distribution?

(e) Is F reversible?

Solution: First use the following R code to test if $\lim_{n \to \infty} M^n$ exists:

```
fdat<-c(c(4,3,2,1,0)/10,c(0,4,3,2,1)/10,c(0,0,5,3,2)/10,
c(0,0,0,6,4)/10,c(0,0,0,0,1)); F<-matrix(fdat,5,5,byrow=T); F
rowSums(F)  # check that F is row stochastic
# Iterate F to see if a limit exists:
FF <- F; for(i in 1:100) FF <- FF %*% FF;  FF
```

The result is

$$F^\infty \overset{\text{def}}{=} \lim_{n \to \infty} F^n = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

(a) The existence of the limit of F^n as $n \to \infty$ implies that F is aperiodic.

(b) Every pair of states i, j with $i > j$ has $F(i, j) = 0$, so state i can never transition to any state j with $j < i$. Thus F is reducible.

(c) Since F is not irreducible, it is not ergodic.

(d) Though F is not ergodic, the limit F^∞ nonetheless exists, and we conclude that F has a stationary state $\pi = (0, 0, 0, 0, 1)$. This means that state 5 is *absorbing*, or equivalently that there are no transitions possible out of state 5.

(e) Check the 25 detailed balance equations with the stationary distribution from part d:

```
pp <- FF[1,]; pp  # stationary distribution: any row of FF
dif <- matrix(0,5,5);  # initialize to all zeros
for ( i in 1:5 ){
 for(j in 1:5){
  dif[i,j] <- pp[i]*F[i,j]-pp[j]*F[j,i]  # reversible iff all 0s
}};  dif
```

The final output is the 5×5 matrix `dif` of differences:

```
     [,1] [,2] [,3] [,4] [,5]
[1,]   0    0    0    0    0
[2,]   0    0    0    0    0
[3,]   0    0    0    0    0
[4,]   0    0    0    0    0
[5,]   0    0    0    0    0
```

This shows that $\pi[i]F[i,j] - \pi[j]F[j,i] = 0$ for all i, j. Conclude that F is reversible. □

Figure 1: HW 5, Ex.2a: Contour plot of posterior pdf in (p_A, p_B)

Figure 2: HW 5, Ex.2i: 1-D random walk, seed=123

Figure 3: HW 5, Ex.2ii: 1-D random walk, seed=4567

Figure 4: HW 5, Ex.2iii: 1-D random walk, seed=89012

Figure 5: HW 5, Ex.3i: 2-D random walk, seed=123

11

Figure 6: HW 5, Ex.3ii: 2-D random walk, seed=4567

Figure 7: HW 5, Ex.3iii: 2-D random walk, seed=89012