STAT 380: Spring 2016

Final Examination

Solutions

Richard Lockhart 16 April 2016

Instructions: This is a closed book exam. You are permitted to use 6 sheets of notes, machine-written or hand-written. You may use both sides of the sheets and I place no limits on font size. Calculators are not permitted nor are any other electronic aids. The exam is out of 50. You should have a total of 10 pages; the first page is a grade sheet and the last is extra space. I will be marking for clarity of explanation as well as correctness. Without a clear explanation you should not expect to get more than half marks. You have 3 hours.

1a		3	1b		3	1c		4
2		10				3		10
4a		3	4b		4	4c		3
5a		3	5b		3	5c		4

Total		**50**

1. Each week I buy 0 or 1 or 2 lottery tickets. To decide how many tickets to buy this week I look at how many I bought last week and toss a fair coin. If I bought 0 tickets last week and I get a Heads then I buy 1 ticket next week otherwise I buy 0. If I bought 1 ticket last week and I get a Head I buy 2 tickets this week otherwise I don't buy any. Finally if I bought 2 tickets last week and get a Head I buy two tickets again this week. Otherwise I buy one ticket. Let X_n be the number of tickets I buy in week n. Assume that $X_0 = 0$.

   (a) Write out the transition matrix of the resulting Markov Chain. [3 marks]

$$P = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

   (b) In the long run how many tickets do I buy per week on average? [3 marks]

   We solve

$$\pi = \pi \mathbf{P}$$

   The matrix $\mathbf{P}$ is doubly stochastic so the answer is the uniform distribution: $\pi_i = 1/3$ for $i = 0, 1, 2$. Or you can write out the equations

$$\begin{aligned} \pi_0 &= \frac{1}{2}\pi_0 + \frac{1}{2}\pi_1 \\ \pi_1 &= \frac{1}{2}\pi_0 + \frac{1}{2}\pi_2 \\ 1 &= \pi_0 + \pi_1 + \pi_2 \end{aligned}$$

   The first equation shows $\pi_1 = \pi_0$. Put this in the second to see $\pi_2 = \pi_0$. This shows they are all equal leading to the answer just given.

   (c) If each ticket costs \$1 and wins \$10 with probability 1/20 how much money do I lose per week (losses are cost of tickets minus winnings) in the long run? [4 marks]

   For each ticket you buy you expect to lose

$$1 - \frac{1}{20} \cdot 10 = \frac{10}{20}$$

   dollars. The average number of tickets bought in a week is

$$0\pi_0 + 1\pi_1 + 2\pi_2 = 1$$

   so you expect to lose $10/20 = \$0.50$ per week in the long run. I may get some people who imagine that when you win you get your \$ back so that the cost is either 1, with probability 19/20 or -10, with probability 1/20. The loss per week would be \$0.45 or 9/20. Both are ok though the first answer matches how lotteries actually work.

2. Suppose the $X(t)$ is a continuous time Markov Chain with infinitesimal generator $\mathbf{R}$. Assume that the entries in each column of $\mathbf{R}$ add up to 0. Prove that the transition matrix $\mathbf{P}(t)$ is doubly stochastic for each t; that is, prove that each column of $\mathbf{P}(t)$ adds up to 1. I suggest you use Kolmogorov's Backwards Equations. [10 marks]

Suppose there are N states. Suppose $\mathbf{1}$ is a row vector whose entries are all equal to 1. Then we are given that

$$\mathbf{1R} = \mathbf{0}.$$

Kolmogorov's Backwards Equations are

$$\mathbf{P}'(t) = \mathbf{RP}(t).$$

Multiply this on the left by $\mathbf{1}$ and get

$$\mathbf{1P}'(t) = \frac{d}{dt}\mathbf{1P}(t) = \mathbf{1RP}(t) = \mathbf{0}.$$

That is, the vector $\mathbf{1P}(t)$ is constant as a function of t. At $t = 0$ we have

$$\mathbf{1P}(0) = \mathbf{1I} = \mathbf{1}$$

so that the column sums are all equal to 1 at $t = 0$ and therefore for all t.

3. Particles are detected in a detector according to a Poisson process with rate λ per hour. In one 24 hour period a total of 8 particles are detected. What is the conditional probability that none of these 8 were detected in the last 8 hours of the 24? [10 marks]

Method 1: Given that there are 8 particles in a fixed time period the actual detection times behave like a sample of size 8 from the uniform distribution over the time period. We want the probability that all of a set of $n = 8$ uniforms on the interval [0,1] are less than 2/3. This is

$$\left(\frac{2}{3}\right)^8.$$

Method 2: let N_1 be the number of particles in the first 16 hours and N_2 be the number in the last 8 hours. Let $N = N_1 + N_2$. We are asked for

$$P(N_2 = 0|N = 8) = \frac{P(N_2 = 0, N = 8)}{P(N = 8)} = \frac{P(N_1 = 8, N_2 = 0)}{P(N = 8)}.$$

The top factors because N_1 and N_2 are independent so we get

$$P(N_2 = 0|N = 8) = \frac{\left((16\lambda)^8 e^{-16\lambda}/8!\right)\left((8\lambda)^0 e^{-8\lambda}/0!\right)}{(24\lambda)^8 e^{-24\lambda}/8!} = \left(\frac{16}{24}\right)^8 = \left(\frac{2}{3}\right)^8.$$

4. On the midterm I asked you about an experimental design strategy called play-the-winner and asked you to analyze it with a discrete time Markov Chain. I now want you to consider a continuous time version. Two players, A and B, play a game. On each turn of the game one player 'serves' and can either score a point on that turn or not. If the player who served scores a point she serves again. If not, no point is scored and the other player begins to serve. Now suppose that playing each point takes a random amount of time and that this time has an exponential distribution. Suppose that when A serves the time taken for a turn has rate v_A and that she scores a point with probability p_A. Suppose similarly that when B serves she scores a point with probability p_B and that the time taken for a turn has rate v_B.

   This system can be analyzed using a suitable Markov Chain. Use 4 states, say $\{0, 1, 2, 3\}$. In states 0 or 1 player A serves while in states 2 or 3 player B serves. When A serves and wins the state moves to state 1 if the chain is in state 0 and to state 0 if it is in state 1. Similarly when B wins while serving the state changes from 2 to 3 or 3 to 2. When A loses her serve the chain moves to state 2 and when B loses her serve the chain moves to 0.

   (a) What is the matrix $\mathbf{R}$, the infinitesimal generator, for this chain? [3 marks]

$$\mathbf{R} = \begin{bmatrix} -v_A & p_A v_A & (1-p_A)v_A & 0 \\ p_A v_A & -v_A & (1-p_A)v_A & 0 \\ (1-p_B)v_B & 0 & -v_B & p_B v_B \\ (1-p_B)v_B & 0 & p_B v_B & -v_B \end{bmatrix}$$

   (b) Indicate clearly how you would calculate the long run fraction of time during which the players are playing turns where A served. I want to see you convince me you know exactly what equations to solve without solving them. More explicit answers will get more marks. [4 marks]

   I am looking for the following steps. First write down the equation $\pi\mathbf{R} = 0$ in detailed form such as

$$\begin{aligned} \pi_0 v_A &= \pi_1 p_A v_A + \pi_2(1-p_B)v_B + \pi_3(1-p_B)v_B \\ \pi_1 v_A &= \pi_0 p_A v_A \\ \pi_2 v_B &= \pi_4 p_A v_A + \pi_0(1-p_A)v_A + \pi_1(1-p_A)v_A \\ \pi_3 v_B &= \pi_2 p_B v_B \end{aligned}$$

   I did not ask for the stationary distribution but you see that

$$\pi_1 = p_A \pi_0$$

   and

$$\pi_3 = p_B \pi_2$$

In the first equation

$$\pi_0 v_A(1-p_A^2) = \pi_2 v_B(1-p_B)(1+p_B) = \pi_2 v_B(1-p_B^2)$$

so

$$\pi_0 \left[1 + p_A + (1+p_B)\frac{v_B(1-p_B^2)}{v_A(1-p_A^2)}\right] = 1$$

giving

$$\pi_0 = \frac{v_A(1-p_A^2)}{(1+p_A)(1-p_A^2)v_A + (1+p_B)(1-p_B^2)v_B}.$$

(c) On the midterm I asked you to find the fraction of turns where A is serving. If v_A is larger than v_B will the answer to part (b) of this question be larger than the answer on the midterm or smaller? Why? [3 marks]

When v_A is large you leave states 0 and 1 more quickly than you leave states 2 and 3 so you spend relatively less time in state A. The answer is smaller than the midterm answer.

5. A continuous time Markov Chain has states $\{0, 1, 2\}$. If the chain is started in chain i it stays there for a time T_i and then moves to a new state with transition matrix $\mathbf{P}$ given by

$$\begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{5} & 0 & \frac{4}{5} \\ \frac{1}{2} & \frac{1}{2} & 0 \end{bmatrix}$$

The departure rates from state i are 4, 5, and 6 per time unit for states 0, 1, and 2. That is, the rate for holding time T_0 is 4, for T_1 the rate is 5, and the rate for T_2 is 6.

(a) Find the instantaneous generator, $\mathbf{R}$, for this chain. [3 marks]

$$\mathbf{R} = \begin{bmatrix} -4 & 2 & 2 \\ 1 & -5 & 4 \\ 3 & 3 & -6 \end{bmatrix}$$

(b) Find the stationary initial distribution for this chain. [3 marks]

Solve the equation

$$\pi \mathbf{R} = \pi$$

This becomes

$$\begin{aligned} 0 &= -4\pi_0 + \pi_1 + 3\pi_2 \\ 0 &= 2\pi_0 - 5\pi_1 + 3\pi_2 \\ 0 &= 2\pi_0 + 4\pi_1 - 6\pi_2 \end{aligned}$$

The second equation minus the third gives

$$0 = -9\pi_1 + 9\pi_2$$

so $\pi_2 = \pi_1$. Double the first plus the third shows $\pi_0 = \pi_2$. So all 3 are equal and

$$\pi_0 = \pi_1 = \pi_2 = \frac{1}{3}.$$

(c) Define μ_{ij} to be the mean time for a chain started in state i to reach state j for the first time. Find a formula of the form

$$\mu_{0,2} = a + b\mu_{1,2}$$

and give the values of a and b in terms of the rates and transition probabilities given above. [4 marks]

Starting from state 0 you remain in state 0 for an exponentially distributed amount of time T_0 with mean $1/v_0 = 1/4$. Then you move either to state 1 or state 2. If you move to state 2 you are done and the time involved is just T_0. If you move to state 1 you must now wait a time whose expected value is $\mu_{1,2}$. This gives

$$\mu_{0,2} = \mathrm{E}(T_0) + P_{12}\mu_{1,2} = \frac{1}{4} + \frac{1}{2}\mu_{1,2}.$$