Absorption probabilities and related topics

Thursday, October 02, 2008
12:03 PM

Homework 2 is now posted, due Tuesday, October 14 at 12 PM.

We now prepare to develop deterministic formulas to answer the questions posed at the end of last lecture.

As one further preparatory step, we'll write the probability transition matrix in canonical form by reordering the states so that we can write the matrix as:

$$P = \begin{array}{c} \\ T^c \\ T \end{array} \begin{array}{c} \begin{array}{cc} T^c & T \end{array} \\ \begin{bmatrix} \bar{P} & 0 \\ S & Q \end{bmatrix} \end{array}$$

$$T^c = \bigcup C_k$$

$$\bar{P} = \begin{array}{c} \\ C_1 \\ C_2 \\ C_3 \\ \vdots \\ C_r \end{array} \begin{array}{c} \begin{array}{ccc} C_1 & & C_r \end{array} \\ \begin{bmatrix} \bar{P}^{(1)} & & 0 \\ & \bar{P}^{(2)} & \\ & \ddots & \\ 0 & & \bar{P}^{(r)} \end{bmatrix} \end{array}$$

block matrix

This just means relabelling the states so that you first run through closed recurrent classes $C_1, C_2, \ldots, C_r$ and then transient states T.

It's useful to represent the results in this canonical form, but in practice, one does not usually have to go through the actual reordering process.

Absorption probabilities

For $i \in T$, what is the first recurrent state $j \in T^c$ that will be visited?

$$U_{ij} = P(\underbrace{X_\tau}_{=\text{first recurrent state visited}} = j \mid X_0 = i)$$

where $\underbrace{\tau}_{=\text{epoch at which a recurrent state is first visited}} = \min\{n \geq 0 : X_n \notin T\}$

is the probability that recurrent state j will be the first recurrent state visited if the Markov chain starts in transient state i. Knowing these probabilities also tell us the probabilities for the Markov chain starting in transient state i to end up in each of the closed recurrent classes, because

$$P(X_\tau \in C_k \mid X_0 = i) = \sum_{j \in C_k} U_{ij}$$

(sum over disjoint events $\{X_\tau = j\}$)

is the probability to end up in class C_k.

The technique we will use for deriving formulas for these absorption probabilities is called first step analysis:

- idea is to break down the paths corresponding to the desired event by looking at what happens at the first step, and then using this as the basis of a recursion argument. Markov property is key for this to work.

$$U_{ij} = P(X_\tau = j \mid X_0 = i)$$

$$= \sum_{k \in S} P(\overset{A}{X_\tau = j}, \overset{B}{\underbrace{X_1 = k}} \mid \overset{C}{X_0 = i})$$

(sum over partition)

$$= \sum_{k \in S} P(\overset{A}{X_\tau = j} \mid \overset{B}{X_1 = k}, \overset{C}{X_0 = i}) P(\overset{B}{X_1 = k} \mid \overset{C}{X_0 = i})$$

$P(A \text{ and } B \mid C) = P(A \mid B, C)\, P(B \mid C)$

This trick using the definition of conditional probability often helps computations by shifting the calculation from simultaneous events to looking at the probability of one event at a time, but maybe with multiple conditions.

$$= \sum_{k \in S} P(X_\tau = j \mid X_1 = k, X_0 = i) P_{ik}$$

We now consider how to express the first factor for various possible states k.

1) $k = j \in T^c$.

$$P(X_\tau = j \mid X_1 = k, X_0 = i) = \underbrace{P(X_\tau = j \mid X_1 = j, X_0 = i)}_{=1} \quad (\tau = 1)$$

2) $k \in T^c$, $k \neq j$

$$P(X_\tau = j \mid X_1 = k, X_0 = i) = 0$$

3) $k \in T$: ... past $n < 1$

3) $k \in T$:

$$P(X_t = j \mid \underset{\text{past } n<1}{X_1 = k}, X_0 = i) = \quad (\text{future } n>1)$$

$$P(X_t = j \mid X_1 = k) \quad \text{(Markov property)}$$

$$= U_{kj}$$

(the actual label of the initial epoch doesn't matter since the conditional probability being computed doesn't refer to absolute time; more rigorously, one can show a one-to-one correspondence between the paths and their probabilities when you condition on the value of X_0 or X_1. (see Resnick Sec. 2.11)

Assembling these results:

$$U_{ij} = 1 \cdot P_{ij} + \sum_{\substack{k \in T^c \\ k \neq j}} 0 \cdot P_{ik} + \sum_{k \in T} U_{kj} P_{ik}$$

$$U_{ij} = P_{ij} + \sum_{k \in T} P_{ik} U_{kj} \quad \text{for } i \in T, \; j \in T^c$$

Using the canonical form, we rewrite this as:

$$U_{ij} = S_{ij} + \sum_{k \in T} Q_{ik} U_{kj} \quad \text{for } i \in T, \; j \in T^c$$

$$U = S + QU \quad \text{as matrices,}$$

$$(I - Q) U = S$$

identity matrix $I = \begin{pmatrix} 1 & & 0 \\ & \ddots & \\ 0 & & 1 \end{pmatrix}$

$U = (I - Q)^{-1} S$ provided that the matrix

$I - Q$ is not singular.

You can prove this by using the Perron-Frobenius theorem to argue that all eigenvalues of Q satisfy $|\lambda| \leq 1$.

The formula we derived in fact can be given an intuitive meaning by expanding the inversion as a geometric series, using the fact that the eigenvalues of Q have modulus strictly less than 1:

$$\boxed{U = (I-Q)^{-1} S = \left(\sum_{n=0}^{\infty} Q^n\right) S}$$

$$U_{ij} = \sum_{n=0}^{\infty} (Q^n S)_{ij}$$

$$\| \quad P(X_\tau = j, \tau = n+1 \mid X_0 = i)$$

(partition over $\{\tau = n\}_{n=1}^{\infty}$)

Accumulated cost/reward during transient period

$$\sum_{n=0}^{\tau-1} f(X_n)$$

$$\tau = \min\{n \geq 0 : X_n \notin T\}$$

f: cost/reward function per epoch

Important special cases:

1) $f = 1$: number of epochs spent in transient states

2) $f(j) = \delta_{ij} = \begin{cases} 1 & j = i \\ 0 & j \neq i \end{cases}$

then the accumulated cost/reward sum would return the number of epochs at which (transient) state i is visited.

$$w_i = \mathbb{E}\left(\sum_{n=0}^{\tau-1} f(X_n) \,\middle|\, X_0 = i\right) \quad \text{for } i \in T$$

$$\boxed{\mathbb{E}(Y \mid A) = \sum_{j \in S} j \, P(Y = j \mid A)}$$

First-step analysis again, using now the law of total expectation:

For a partition $\{A_j\}_{j \in I}$ of probability space,

$$\mathbb{E}(Y) = \sum_{j \in I} \mathbb{E}(Y \mid A_j) \, P(A_j)$$

$$\mathbb{E}(Y \mid B) = \sum_{j \in I} \mathbb{E}(Y \mid A_j, B) \, P(A_j \mid B)$$

$$\mathbb{E}(Y|B) = \sum_{j\in I} \mathbb{E}(Y|A_j, B)\, P(A_j|B)$$

$$W_i = \mathbb{E}\left(f(X_0) + \sum_{n=1}^{T-1} f(X_n) \,\middle|\, X_0 = i\right)$$

$$= f(i) + \mathbb{E}\left(\sum_{n=1}^{T-1} f(X_n) \,\middle|\, X_0 = i\right)$$

$$= f(i) + \sum_{k\in S} \mathbb{E}\left(\sum_{n=1}^{T-1} f(X_n) \,\middle|\, X_1 = k, X_0 = i\right) P(X_1 = k | X_0 = i)$$

(law of total expectation w/ partition $\{X_1 = k\}_{k\in S}$)

$$= f(i) + \sum_{k\in S} \mathbb{E}\left(\sum_{n=1}^{T-1} f(X_n) \,\middle|\, X_1 = k, X_0 = i\right) P_{ik}$$

Consider the first factor for various possible states k:

$$k \notin T \Rightarrow T > 1 \Rightarrow \sum_{n=1}^{T-1} f(X_n) = 0$$

$$k \in T: \mathbb{E}\left(\sum_{n=1}^{T-1} f(X_n) \,\middle|\, X_1 = k, X_0 = i\right)$$

$$= f(k) + \mathbb{E}\left(\sum_{n=2}^{T-1} f(X_n) \,\middle|\, X_1 = k, X_0 = i\right)$$

future, $n > 1$ — past, $n < 1$

$$= f(k) + \mathbb{E}\left(\sum_{n=2}^{T-1} f(X_n) \,\middle|\, X_1 = k\right)$$

(Markov property)

$$= \mathbb{E}\left(\sum_{n=1}^{T-1} f(X_n) \,\middle|\, X_1 = k\right)$$

$$= W_k$$

(again by equivalence of paths; the label of the initial epoch does not matter).

$$w_i = f(i) + \sum_{k \in T^c} 0 \, P_{ik} + \sum_{k \in T} w_k P_{ik}$$

$$w_i = f(i) + \sum_{k \in T} w_k \underbrace{P_{ik}}_{= Q_{ik}} \quad \text{for } i \in T$$

$$\vec{w} = \{w_i\}_{i \in T} \qquad \vec{f} = \{f(i)\}_{i \in T}$$

$$\vec{w} = \vec{f} + Q\vec{w}$$

$$(I - Q)\vec{w} = \vec{f}$$

$$\boxed{\vec{w} = (I - Q)^{-1}\vec{f}}$$

$$\vec{w} = \sum_{n=0}^{\infty} Q^n \vec{f} =$$

$$w_i = \sum_{n=0}^{\infty} (Q^n f)_i = \sum_{n=0}^{\infty} \underbrace{(Q^n)_{ij}}_{= P(X_n = j \mid X_0 = i) \text{ for } i,j \in T} f_j$$

expected reward at epoch n

$$\mathbb{E}(f(X_n) \mid X_0 = i)$$

Actually these techniques for reducible Markov chains also allow us to address new questions for irreducible Markov chains that involve indefinite time horizons but are not of the form for which law of large numbers works well.

For example, in an irreducible Markov chain:

- Starting from state i, what is the expected number of epochs until a certain state j is visited?
- Starting from state i, how many times is state k visited before state j is visited?
- What is the probability, starting from state i, that state k is visited before state j?

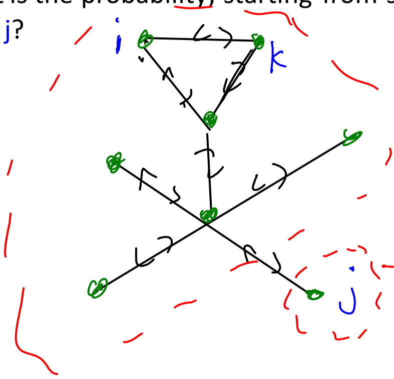

become transient

make absorbing state by setting

$$P_{jk} = \delta_{jk}$$

This works for the first two questions, with appropriate cost/reward functions.

For the third question, we can make both states j and k absorbing, in which they become their own recurrent classes and every other state is transient. Calculate absorption probabilities!