MAS275 Probability Modelling Chapter 5: Hitting times and probabilities

Dimitrios Kiagias

School of Mathematics and Statistics, University of Sheffield

Spring Semester, 2020

Introduction

The main idea of this section is that we are interested in a particular state (or sometimes group of states) and we are interested in answering questions about, for example,

- whether the chain will ever reach ("hit") the state(s);
- if we know that it does so, how long the expected time until this happens is.

Absorption probabilities

We first consider Markov chains with finite state space in which there are some transient states.

As we have noted before, there must be at least one closed class.

We also know that if there is only one closed class then the process will eventually end up in it, regardless of which state it started in.

More than one closed class

When there are two or more closed classes the situation is more interesting, and a process starting in one of the transient states may end up in any one of the closed classes.

It is therefore meaningful to refer to the probability of eventually entering any closed class C starting in any transient state i.

This will be denoted by q_{iC} and is called an **absorption probability**.

Note that in previous notation it coincides with f_{ij} for any $j \in C$, because the process visits j eventually if and only if it enters C.

Notation

For completeness, we define q_{iC} for all states $i \in S$,...

...noting that $q_{iC} = 1$ if $i \in C$ and $q_{iC} = 0$ if i is in any closed class other than C.

Sometimes we write q_{iC} simply as q_i if we are only considering entry into one closed class.

Often C consists of a single state, known as an **absorbing state**.

Finding absorption probabilities

To find absorption probabilities, we condition on the first step and use the law of total probability and the Markov property.

Letting D_C mean the event that the Markov chain ends up in C, we get

$$\begin{aligned} q_{iC} &= P(D_C | X_0 = i) \\ &= \sum_{j \in S} P(X_1 = j | X_0 = i) P(D_C | X_1 = j) \\ &= \sum_{j \in S} p_{ij} q_{jC} \end{aligned}$$

for all transient states i.

Recurrent states

This equation is also true, but tells us nothing new, if i is in a closed class:

Both sides are 1 if $i \in C$ and both sides are zero if i is in any other closed class.

This means that the column vector whose components are the q_{iC}'s is a right eigenvector of P corresponding to the eigenvalue 1, and each equation corresponds to a row of the transition matrix.

Because the Markov chain is not irreducible, we are in a situation where the space of such eigenvectors is more than one-dimensional.

Equations

In practice we have a collection of simultaneous linear equations, one for each transient state.

The conditions that $q_{iC} = 1$ for $i \in C$ and $q_{iC} = 0$ for i in any other closed class are sometimes known as **boundary conditions**, and may be exploited wherever appropriate.

Examples

Example

University course model

Example

Gambler's ruin: probability of ruin

Irreducible chains

Sometimes, in an irreducible chain, we may be interested in which of two sets of states, C_1 and C_2 say, the chain reaches first.

We can treat such cases by using essentially the same method as above, but now setting $q_{iC_1} = 1$ if $i \in C_1$ and $q_{iC_1} = 0$ if $i \in C_2$, and discarding equations which come from rows of the transition matrix corresponding to states in either C_1 or C_2.

(One way to think about this is that we modify the chain so that the classes C_1 and C_2 are absorbing, and then apply exactly the same method as before to the modified chain.)

Example

Example

Random walk on cube

Expected times to absorption

Suppose that a Markov chain has some transient states but that, starting in any one of these states, there is probability 1 of eventually being absorbed in some closed class or other.

This will necessarily be the case if the state space is finite.

Then the length of time, or number of steps, until this absorption takes place will be a proper random variable.

Notation

For starting state i, denote this random variable by T_i and its mean or expected value by $e_i = E(T_i)$.

For completeness, we also define $T_i = 0$ and $e_i = 0$ if i is a state which is already in a closed class, in which case there is zero time to absorption.

Finding expected times to absorption

To get some equations which can be solved for the numbers $\{e_i\}$, we again condition on the first step, and this time use conditional expectation and the Markov property.

We note that, starting in transient state i, if the first step is to state j, say, then the total time to absorption is 1 (corresponding to the step from i to j itself) plus a time to absorption starting in state j, say T_j.

Equations

So

$$\begin{aligned} e_i &= E(T_i) \\ &= \sum_{j\in S} p_{ij} E(T_i | X_1 = j) \\ &= \sum_{j\in S} p_{ij}(1 + E(T_j)) \\ &= 1 + \sum_{j\in S} p_{ij} e_j . \end{aligned}$$

Equations II

Since $e_j = 0$ if j is in a closed class, the above equations may be simplified to

$$e_i = 1 + \sum_{j \in T} p_{ij} e_j$$

for $i \in T$ where T denotes the set of all transient states.

So we have a collection of simultaneous linear equations, the number of unknowns and the number of equations both being equal to the number of transient states.

If the state space is finite, they will always have a unique solution.

Example

Example

University course model

Generalised method

As with the absorption probability method, we can extend this method above to find expected times until the chain arrives in a state (or group of states) which is not absorbing.

For example, if we have a state j in an irreducible chain and we are interested in the expected time until the chain reaches j, we use the same method as above, setting $e_j = 0$, and discard the equation corresponding to row j of the transition matrix.

Example

Example

Patterns in coin tossing