POLITECNICO DI TORINO

An introduction to stochastic vs. deterministic models of CRNs: 1. Intro to Discrete Time Markov Chains (DTMC)

Enrico Bibbona
enrico.bibbona@polito.it

DISMA Dipartimento di Scienze Matematiche G. L. Lagrange
ECCELLENZA 2018 · 2022

An overview of the course

Markov processes are stochastic processes such that the probability distribution of the future observations is completely determined by the present state (or distribution), regardless of any knowledge of the *past* history.

The interest in MP follows two main lines of thought

- (math) The original idea by Markov, is that the basic theorems (LLN, CLT...) that holds for a sequence of i.i.d. random variables may be extended to some sequence of dependent r.v. The easiest kind of dependence that could be introduced is that of the Markov property
- (physics and other applications) It is natural to assume that the dynamics of a physical system that starts with some initial condition (ic) is only determined by such ic, without any need for further informations. cf. ODE

Markov Property in discrete time

Let $(X_n)_{n=0}^{\infty}$ be a discrete time stochastic process with a discrete state space S. $(X_n)_{n=0}^{\infty}$ is a **DTMC** if for any $j, i, i_{n-1}, \ldots, i_0 \in S$, the Markov property

$$P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \ldots, X_0 = i_0) = P(X_{n+1} = j \mid X_n = i) \stackrel{\text{def}}{=} p(i, j).$$

holds. It says that the probabilities associated with a future state only depends upon the current state, and not on the full history of the process.

$p(i, j)$ are called the (one step) transition probabilities (also called transition matrix)

Note that the transition probabilities are not time (n) dependent. We restrict to this so-called *temporally homogeneous* case.

Transition probabilities

If the state space S is finite (say of cardinality k), the transition probabilities can be organized into a $k \times k$ matrix

$$p = \begin{pmatrix} p(1,1) & \cdots & p(1,k) \\ \vdots & & \vdots \\ p(k,1) & \cdots & p(k,k) \end{pmatrix}$$

with non-negative entries such that the row sums are 1:

$$\sum_{j \in S} p(i,j) = 1.$$

Matrices of this kind are called *stochastic matrices*.
If S is countably infinite you can still think at a matrix with infinite dimensions.

Consequence 1: joint prob of a trajectory

We can apply the Markov property to compute the probability of observing a trajectory $X_0, X_1 \cdots X_m$. We have

$$P(X_n = i_n, X_{n-1} = i_{n-1}, \cdots, X_0 = i_0) =$$
$$= P(X_n = i_n | \, X_{n-1} = i_{n-1}, \cdots, X_0 = i_0) \cdot P(X_{n-1} = i_{n-1}, \cdots, X_0 = i) =$$

by the Markov property

$$= p(i_{n-1}, i_n) \cdot$$
$$\cdot P(X_{n-1} = i_{n-1} | \, X_{n-2} = i_{n-2}, \cdots, X_0 = i_0) \cdot P(X_{n-2} = i_{n-2}, \cdots, X_0 = i)$$
$$= p(i_{n-1}, i_n) \cdot p(i_{n-3}, i_{n-2}) \cdots p(i_0, i_1) P(X_0 = i_0)$$

To compute the join probability of a trajectory it is sufficient to know the transition matrix and the initial distribution $\alpha(i) = P(X_0 = i)$.

Consequence 2: Chapman-Kolmogorov eqs. and n-step transitions

A natural question is how to calculate transition probabilities in two steps, e.g.

$$p^{(2)}(i,j) = P(X_{n+2} = j \mid X_n = i)$$

Again, conditioning on $X_{n+1} = k$ and applying the law of total probabilities, the Markov property, and temporal homogeneity we have

$$\begin{aligned} p^{(2)}(i,j) =& P(X_{n+2} = j \mid X_n = i) = \\ =& \sum_k P(X_{n+2} = j \mid X_{n+1} = k, X_n = i) \cdot P(X_{n+1} = k \mid X_n = i) \\ =& \sum_k p(i,k)p(k,j) \end{aligned}$$

this formula has a nice geometric interpretation, as the sum of the probabilities of all the paths that leads from i to j in two steps.

Consequence 2: Chapman Kolmogorov and n-step transitions

Please note that the formula above can be read by saying that the two steps transition probability matrix $p^{(2)}$ is noting but the matrix product between the transition matrix with itself (the second power of the transition matrix p)

$$p^{(2)} = p \cdot p = p^2$$

More in general, the have that the following Chapman Kolmogorov equations hold

$$p^{(m+n)} = p^{(m)} \cdot p^{(n)} = p^{m+n}$$

that, componentwise, means that

$$p^{(m+n)}(i,j) = \sum_k p^{(m)}(i,k)p^{(n)}(k,j)$$

Summary

A Markov chain is fully specified if we know:

- the state space (somehow implicit in the next requirements)
- the transition probability matrix $p(i,j)$
- the initial distribution α

Another way of specifying the chain is through its transition graph that is a graph whose nodes are the states, and the (weighted) edges are arrows that connect two states if the corresponding transition probability is positive. Weights are the transition probabilities. Cf. Ex. 2 below

Examples

Ex. 1: Weather forecast

Let us define the following Markov model of weather conditions. Let the state be either rainy (state 1) or sunny (state 2). Let the probability that a sunny day is followed by another sunny day be 0.8, and the probability that a rainy day is followed by another rainy day be 0.6.

Calculate the probability of observing 4 sunny days in a raw, and to observe the sequence SRS (212).

Calculate the probability that if today is sunny the day after tomorrow will be sunny again.

Given an initial distribution such that of $P(X_0 = 2) = 0.6$, calculate the distribution of X_1. Please note that both the initial distribution and the distribution of X_1 can be seen as vectors whose lengths are equal to the cardinality of the state space. Write a relation between the two vectors.

Examples

Ex. 2: Gambler model

Consider the discrete time Markov chain associated with the following transition graph:

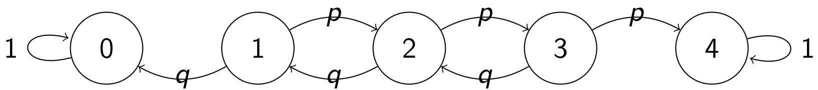

it is a model of gambling. Le's say that your initial fortune is 2 euros. Calculate the probability that after the third bet you still have a fortune of 2 euros.

Examples

Ex. 3: Random walk in 1d

Consider the discrete time Markov chain on $S = \mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ where for some $0 < p < 1$ we have

$$p(i, i+1) = p, \quad p(i, i-1) = q, \quad \text{with} \quad q \stackrel{\text{def}}{=} 1 - p.$$

This chain is very important.It is referred to as *one-dimensional random walk*.

Calculate $p^{(k)}(0,0)$ for any integer k.

Classification of states, communication classes

We say that state j is reachable from i (and write $i \to j$) if there exist an $m \geq 0$ such that

$$p^{(m)}(i,j) > 0.$$

By definition, $p^{(0)}(i,j) = \delta_{ij}$, where δ_{ij} is the Kronecker symbol, i.e. the (i,j) element of the identity matrix.

If both $i \to j$ and $j \to i$ we say that state i communicates with state j , and write $i \leftrightarrow j$.

Communication classes, irreducibility

It is easy to prove that communication is an equivalence relation. Both symmetry and reflexivity are obvious. Transitivity is proved in this way.
Since $i \to j$, there exist an m such that $p^{(m)}(i,j) > 0$.
Since $j \to k$, there exist an n such that $p^{(n)}(j,k) > 0$, therefore

$$p^{(m+n)}(i,k) = \sum_s p^{(m)}(i,s)p^{(n)}(s,k) \geq p^{(m)}(i,j)p^{(n)}(j,k) > 0.$$

As any equivalence relation, communication partitions the state space into equivalence classes, also called communication classes.

A Markov chain is is called irreducible, if all the states communicate with each other. A subset $A \subset S$ of the state space is called absorbing (or a trap), if when

the chain is started in A, there is no way to get out of A.

Example

Consider the following transition graph (where transition probabilities are not explicitly written, but are understood to be numbers strictly between 0 and 1):

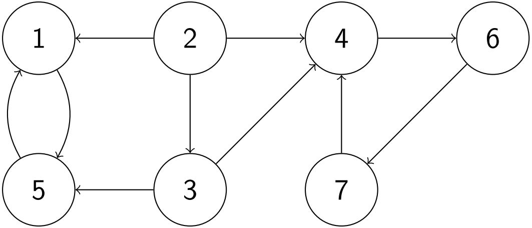

(1)

Describe the communication classes, and find out which one of them are absorbing.

Classification of states, recurrence and transience

Let T_y be the time of the first visit to y, without counting X_0.

$$T_y = \min\{n \geq 1 : X_n = y\}$$

T_y is called **hitting time of** y, and if the chain starts in $X_0 = y$ **return time to** y. T_y is a random variable expressing how many steps are needed to visit y.
Let

$$\rho_{xy} = P_x(T_y < \infty) = P(\text{we will visit } y \text{ again} | X_0 = x),$$

be the probability of returning to y in a finite time if we start at y.

There are two distinct types of states:

- y is recurrent if $\rho_{yy} = 1$;
- y is transient if $\rho_{yy} < 1$.

Classification of states, recurrence and transience

The names *recurrent* and *transient* are better justified by the following properties:

- Recurrent states are visited infinitely many times.
- on the contrary, the number of visits to any transient state is finite. Therefore you will always find a large enough time such that after that time a transient state is never visited any more.

Recurrence and expected number of visits

The random variable $N_n(y) = \sum_{i=1}^{n} \mathbb{1}(\{X_i = y\})$ represents the number of visits to y before time n. We define $N(y) = \lim_{n\to\infty} N_n(y)$. We can easily prove that

$$\mathbb{E}_x[N(y)] = \sum_{n=1}^{\infty} p^{(n)}(x, y)$$

indeed the sequence $N_n(y)$ is a.s. increasing and monotone convergence holds

$$\mathbb{E}_x[N(y)] = \sum_{n=1}^{\infty} P_x(X_n = y) = \sum_{n=1}^{\infty} p^{(n)}(x, y)$$

- it can be proved that y is recurrent if and only if

$$\mathbb{E}_y[N(y)] = \sum_{n=1}^{\infty} p^{(n)}(y, y) = \infty$$

Communication and recurrence

- Recurrence is a class property, i.e., if $i \leftrightarrow j$ and i is recurrent, so is j.

Since i communicates with j, there exists k and m such that $p^{(k)}(i,j)$ and $p^{(m)}(j,i)$ are both positive. Now for any integer n we have

$$p^{(m+n+k)}(j,j) \geq p^{(m)}(j,i)p^{(n)}(i,i)p^{(k)}(i,j)$$

and by summing over all n

$$\sum_{r=1}^{\infty} p^{(r)}(j,j) \geq \sum_{r=m+k+1}^{\infty} p^{(r)}(j,j) \geq \sum_{n=1}^{\infty} p^{(m+n+k)}(j,j)$$

$$\geq p^{(m)}(j,i)p^{(k)}(i,j) \sum_{n=1}^{\infty} p^{(n)}(i,i) = \infty$$

since i is recurrent and $p^{(k)}(i,j)$ and $p^{(m)}(j,i)$ are both positive.

- Transience is also class property, i.e., if $i \leftrightarrow j$ and i is transient, so is j.

Finite state spaces and recurrence

When S is finite, we have four nice properties relating absorbing classes and recurrence.

1. If $A \subset S$ is a finite absorbing communication class, it contains at least a recurrent state
2. all its states are recurrent
3. if S is finite, it can be partitioned as $S = T \cup R_1 \cup \cdots \cup R_k$, where T contains all the transient states and the R_i are absorbing communication classes (and therefore they only contain recurrent states)
4. if a MC on a finite state space is irreducible, it is recurrent (that means that all states are recurrent)

Stationary distributions, limit distributions and asymptotic frequencies

In the next slides we are going to understand the relations between three distinct notions that we will see are strongly related one to each other

- A stationary distribution π is a distribution on S such that if $X_0 \sim \pi$, that means that the chain is initialized in state i with probability $\pi(i)$, then $X_1 \sim \pi$, and for every n, $X_n \sim \pi$
- A limiting distribution is a distribution π su that

$$\lim_{n\to\infty} p^{(n)}(x,y) = \pi(y)$$

  regardless of x.
- The asymptotic frequency of state i is the long run proportion of visits to i

$$\lim_{n\to\infty} \frac{N_n(y)}{n}$$

Stationary distributions

An exercise we have done (cf. slides 9) showed us that we can find the probability distribution ν of X_1 if we know the initial probability distribution α and the transition matrix p, by calculating $\nu = \alpha \cdot p$. Therefore if we want to find a stationary probability distribution we need to solve the following algebraic equation

$$\pi = \pi \cdot p$$

and require that $\sum_i \pi(i)) = 1$ in coordinates

$$\pi(i) = \sum_j \pi(j) \cdot p(j, i)$$

that means π is an eigenvector of p corresponding to the eigenvalue 1.
We want to investigate conditions that may assure that a stationary distribution of a MC exists and is unique.

Example

Consider a discrete time Markov chain with $S = \{1, 2, 3\}$ and transition matrix. You can interpret it as a (very bad) social class mobility model. State 1 is working class, state 2 middle class, state 3 upper class

$$P = \begin{pmatrix} 0.7 & 0.2 & 0.1 \\ 0.3 & 0.5 & 0.2 \\ 0.2 & 0.4 & 0.4 \end{pmatrix}$$

Find a stationary distribution.

Solution:

$$\pi(1) = \frac{22}{47}, \quad \pi(2) = \frac{16}{47}, \quad \pi(3) = \frac{9}{47}.$$

Example: doubly stochastic matrices

A transition matrix p is called doubly stochastic, if besides $\sum_j p(i,j) = 1$ for all i, also for every j, $\sum_i p(i,j) = 1$. That means that both rows an columns sums up to 1.

(Exercise) Prove that for a MC with a $k \times k$ doubly stochastic transition matrix, the uniform distribution

$$\pi(i) = \frac{1}{k} \text{ for all } i$$

is a stationary distribution.

Stationary distributions and stationary measures

Finding a stationary distribution amount at solving $\pi = \pi \cdot p$ and requiring $\sum_i \pi(i)) = 1$. It might be easier to find non negative vectors η that solves

$$\eta = \eta \cdot p$$

without requiring the normalisability, allowing for the case $\sum_i \eta(i) = \infty$.

A vector η that solves the above equation is called a stationary measure. If, moreover, it is normalizable, it is also a stationary distribution.

Theorem

Suppose a Markov chain is irreducible and recurrent. Then there exist a stationary measure η with $\eta(i) > 0$ for all $i \in S$ (not necessarily finite!)

For finite chains, it follows from Perron Frobenius theorem. To guarantee the existence of a stationary distribution a stronger condition is needed.

Positive recurrence

A state i is positive recurrent if

$$\mathbb{E}_i(T_i) = \mathbb{E}(T_i | X_0 = i) < \infty$$

It can be proved that

- Positive recurrence is a class property, i.e., if $i \leftrightarrow j$ and i is positive recurrent, so is j.

A MC is positive recurrent if all states are.

Theorem

Suppose a Markov chain is irreducible and positive recurrent. Then there exist a stationary distribution π with $\pi(i) > 0$ for all $i \in S$

Finite chains

Irreducible finite chains are always positive recurrent, therefore they always allow for a stationary distribution, and this distribution is also unique

Finite chains that are not irreducible The state space is partitioned in $S = T \cup R_1 \cup \cdots \cup R_k$. For each R_i we can find a stationary distribution $\pi_i(\cdot)$ with support equal to R_i. Therefore different stationary distributions exists and it can be proved that a distribution π is stationary if and only if it is a linear combination of the π_i.

I urge you to prove the theorem yourself and to think at the Gambler ruin example.

Limit distributions

Assume that α is a limit distribution, meaning that α is a probability distribution and for any $x \in S$ we have

$$\lim_{n\to\infty} P(X_n = x) = \alpha(x).$$

Then, it holds (if you trust me enough in changing the limit with the sum in the third equality)

$$\begin{aligned}\alpha(x) &= \lim_{n\to\infty} P(X_n = x) = \lim_{n\to\infty} \sum_{y\in S} p(y,x) P(X_{n-1} = y) \\ &= \sum_{y\in S} p(y,x) \lim_{n\to\infty} P(X_{n-1} = y) = \sum_{y\in S} p(y,x)\alpha(.y)\end{aligned}$$

That is, we must have $\alpha = \alpha P$ and therefore a limit distribution must be a stationary distribution!

Limit distributions and periodicity

A limit distribution must be a stationary distribution, but the opposite implication does not hold. This is a counterexample. Consider a discrete time Markov chain with $S = \{1, 2\}$ and transition matrix

$$P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

This is a finite irreducible chain, therefore positive recurrent. We can see that $\pi = (\frac{1}{2}, \frac{1}{2})$ is a stationary distribution.

However a limit distribution cannot exist since if we start at state 1, at all even times we will be in 1 with probability one, while in all odd times we will be in state 2 with probability one.

Limit distributions and periodicity

The problem with the previous chain is that the chain is periodic.

A state y is periodic with period N if $p^{(k)}(y, y)$ can be strictly positive only if k is multiple of N.

Theorem

Suppose a DTMC is irreducible, all states are aperiodic, and there is a stationary distribution π. Then, for all $i, j \in S$,

$$\lim_{n\to\infty} P(X_n = j | X_0 = i) = \lim_{n\to\infty} p^{(n)}(i, j) = \pi(j).$$

Asymptotic frequency and stationary distributions

Theorem (Theorem 1.21 in Durret book. Asymptotic frequency)

Suppose a DTMC is irreducible and recurrent (all states). Then, a.s.

$$\lim_{n\to\infty} \frac{N_n(y)}{n} = \frac{1}{\mathbb{E}_y[T_y]}.$$

This is an interpretation of positive recurrence.

Theorem (Theorem 1.22 in Durret book)

If a DTMC is irreducible and a stationary distribution π exists, then

$$\pi(y) = \frac{1}{\mathbb{E}_y[T_y]},$$

it is unique, and the chain is positive recurrent.

Comments

- the mean waiting time for the occurrence of an event is the inverse of its probability (cf. the mean of a geometric rv).
- it is a law of large numbers, since we have an average $\frac{N_n(y)}{n} = \frac{1}{n}\sum_{n=1}^{n} \mathbb{1}(\{X_n = y\}))$ converging to its expectation
- the formula if often used to calculate $\mathbb{E}_y[T_y]$, that is often hard to calculate more directly

Ergodic theorem (asymptotic reward)

Theorem (Theorem 1.23 in the textboook)

Suppose a DTMC is irreducible and a stationary distribution π exists. Assume that for some function f, $\sum_x |f(x)|\pi(x) < \infty$. Then, a.s.

$$\frac{1}{n}\sum_{m=1}^{n} f(X_m) \to \sum_x f(x)\pi(x).$$

I.e. time averages (left) equal space averages in stationary regime (right).

You can interpret $f(i)$ as the reward (or cost) you get for visiting i.

The theorem is as a **generalization of the law of large numbers** for discrete random variables, when the random variables we are averaging are not necessarily independent.

Examples

Class Mobility, weather and gambler

1d Random Walk

Consider the discrete time Markov chain on $S = \mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ where for some $0 < p < 1$ we have

$$p(i, i+1) = p, \quad p(i, i-1) = q, \quad \text{with} \quad q \stackrel{\text{def}}{=} 1 - p.$$

It is irreducible, meaning that from any state we can move forward as we like, and backwards as we like. Every state communicates with every other. We claim the following:

1. all states are recurrent if $p = 1/2$;
2. all states are transient if $p \neq 1/2$;
3. In any case, there exists no stationary distribution (without proof).

Hence, in this example we see that if S is infinite, then the fact that S is closed and irreducible does not imply recurrence

1d Random Walk

Since it is an irreducible chain, either all state are recurrent or all states are transient. We prove that 0 is recurrent. showing that

$$E_0[N(0)] = \sum_{n=1}^{\infty} p^{(n)}(0,0) = \infty$$

Note that $p^{(n)}(0,0) \neq 0$ only if n is even. Moreover to be at zero again after $2n$ the process needs to have dome the same amount of steps (n) forward and backward. Hence

$$p^{(2n)}(0,0) = \binom{2n}{n} p^n q^n = \frac{(2n)!}{n!n!}(pq)^n.$$

1d Random Walk, $p = 1/2$

Now, the latter is a bit complicated. Stirling's formula states that

$$\frac{n!}{n^n e^{-n}\sqrt{2\pi n}} \xrightarrow[n\to\infty]{} 1.$$

Hence, for large n

$$p^{(2n)}(0,0) = \frac{(2n)!}{n!n!}(pq)^n \approx \frac{\sqrt{4\pi n}(2n)^{2n}e^{-2n}}{2\,\pi\,n\,n^{2n}\,e^{-2n}}(pq)^n = \frac{1}{\sqrt{\pi n}}(4pq)^n.$$

Now, remember that $p = 1/2$, so $4pq = 4(1/2)(1/2) = 1$. Hence, for n big enough, say for $n \geq M$,

$$p^{(2n)}(0,0) \approx \frac{1}{\sqrt{\pi n}} > \frac{1}{2\sqrt{\pi n}}.$$

Hence,

$$\sum_{n=1}^{\infty} p^{(2n)}(0,0) \geq \sum_{n=M}^{\infty} \frac{1}{2\sqrt{\pi n}} = \infty,$$

1d Random Walk, $p \neq 1/2$

We have already seen that

$$E_0[N(0)] = \sum_{n=1}^{\infty} p^{(n)}(0,0) = \sum_{n=1}^{\infty} p^{(2n)}(0,0)$$

and for large n

$$p^{(2n)}(0,0) \approx \frac{1}{\sqrt{\pi n}}(4pq)^n.$$

In particular, for n big enough, say for $n \geq M$,

$$p^{(2n)}(0,0) \leq (4pq)^n.$$

1d Random Walk, $p \neq 1/2$

If $p \neq 1/2$ the function $4pq = 4p(1-p)$ is strictly less than one (this function is a parabola with maximum 1, obtained at $p = 1/2$). Hence,

$$\sum_{n=1}^{\infty} p^{(2n)}(0,0) \leq \sum_{n=1}^{M-1} p^{(2n)}(0,0) + \sum_{n=M}^{\infty} (4pq)^n \leq M - 1 + \frac{1}{1-4pq} < \infty,$$

where in the second last passage we used that the series starting from $n = M$ is less than the series starting from $n = 0$, and we used the formula for geometric series.