Anna Janicka

Probability Calculus 2019/2020
Lecture 14

1. Markov Chains

During the previous lecture, we looked mainly at sequences of independent (or at least uncorrelated) random variables. Here we will briefly introduce quite a different topic: one where we wish to precisely define the relationship between subsequent random variables. In many cases, this scenario – albeit of course limiting due to the need to greatly simplify the relationship – is better fit to describe the real world than models which assume independence of factors, especially for processes observed in time.

We will start by telling a simple story: we track the evening activities of a Mr. X, who has the following customs. Each evening, he randomly chooses one of three types of activities: staying at home, visit to a bar, or a date. He follows the following rule:

- if, on a given evening, he stayed at home, then on the next day he chooses either a visit to a bar, or a date with probabilities $\frac{1}{2}$ and $\frac{1}{2}$;
- if, on a given evening, he went to a bar, then on the next day with probability $\frac{1}{2}$ he has a hangover and stays at home, with probability $\frac{1}{4}$ he visits the bar again, and with probability $\frac{1}{4}$ he goes on a date;
- if, on a given evening, he had a date, then the next day he will choose a visit to the bar with probability $\frac{1}{2}$, and staying at home or a date with probabilities $\frac{1}{4}$ each.

We now have a set of three possible **states**, which we will call "home", "bar" and "date" and denote by 1, 2 and 3, respectively. Each new day brings a new **step**, or move. This move is described by probabilities of changing state from i to j, which we will denote by p_{ij}, and call **transition probabilities**; these probabilities do not depend on the whole history of the process, but only on its current state of it. In our case they, additionally, do not change with time (but this is not a necessary feature of Markov Chains). Therefore, the properties of the process may be described by a matrix of transition probabilities p_{ij}, called the **transition matrix**, which in the case of our study would take the fo llowing form:

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

An alternative description of the same process could be provided by means of a graph, where the vertices denote states, and arrows denote non-zero transition probabilities:

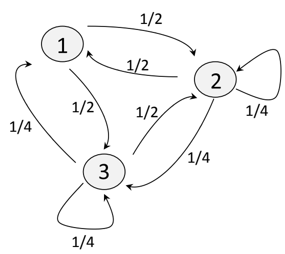

In order to supply full knowledge about the whole process, we also need information on where Mr. X spent his evening on the initial day of the process; then, we can calculate the probabilities of his locations during all subsequent evenings.

Formally, we will introduce the following definition:

Definition 1. *A sequence of random variables $(X_n)_{n=0}^{\infty}$, taking on values in a finite set E is a* **Markov Chain**, *if for any $n = 1, 2, 3, \ldots$ and any sequence $x_0, x_1, \ldots, x_n$ of elements of the set E, we have*

$$\mathbb{P}(X_n = x_n | X_{n-1} = x_{n-1}, X_{n-2} = x_{n-2}, \ldots, X_0 = x_0) = \mathbb{P}(X_n = x_n | X_{n-1} = x_{n-1}),$$

provided that $\mathbb{P}(X_{n-1} = x_{n-1}, X_{n-2} = x_{n-2}, \ldots, X_0 = x_0) > 0$.

If, for each $i, j \in E$, $\mathbb{P}(X_n = j | X_{n-1} = i)$ *does not depend on* n*, the chain is called* **time-homogenous** *or stationary. In this case, we can define the transition matrix* $P = (p_{ij})_{ij \in E}$*, by the formula*

$$p_{ij} = \mathbb{P}(X_1 = j | X_0 = i).$$

Note that the row of matrix P corresponds to the state of departure, and the columns – to the new states. Therefore, the sum of row probabilities in the matrix is always equal to 1. The sum of column probabilities need not be equal to 1.

The definition above may easily be extended to infinite state spaces (imagine the example of a drunkards walk); most of the theorems which are valid for a finite set of states may be extended (albeit sometimes with slight modifications) to infinite. Note also that the condition that the future value of the chain (X_{n+1}) depends only on the current state (X_n), and not on the historical values of the chain, may be modified to incorporate (some of) these past values. For example, it suffices to consider a modified chain with a state space of $E \times E$ instead of E, to allow dependence on the value of X_{n-1} apart from dependence on X_n, etc.

The matrix notation for the Markov chain, although less intuitive than the graph, has an important feature: it allows to calculate the values of different probabilities easily. Note that since all the variables in the Markov chain take on values from the same, finite set, when describing the distributions of these variables we may omit this set of values (since it is obvious), and treat these distributions as simple vectors. For example, in our case of the activities of Mr. X, a vector $\mathbf{x} = (1, 0, 0)$ will describe the situation where Mr. X is at home (a.s.), while a vector $\mathbf{y} = (0, \frac{1}{2}, \frac{1}{2})$ denotes a situation where Mr. X is in a bar or on a date with probabilities $\frac{1}{2}$ and $\frac{1}{2}$, respectively. The distribution of the variable X_0 is denoted the **initial state**. Note that the multiplication of a vector by the transition matrix P gives the distribution for the next step; in particular, if we know that Mr. X is initially at home and we wish to describe his whereabouts the next evening, we will look at $(1, 0, 0) \cdot P$, which is equal to the vector $(0, \frac{1}{2}, \frac{1}{2})$. Note that if this vector is, again, multiplied by the matrix P, we will obtain the probabilities describing Mr. X's whereabouts two days from the initial state: $(\frac{3}{8}, \frac{3}{8}, \frac{1}{4})$, etc. The above property may be summarized in the following:

Theorem 1. *Let* $X_0, X_1, X_2, \ldots$ *be a Markov chain with an initial state* $\mathbf{q}$ *and a transition matrix* P*. Then, the variable* X_n *has a distribution equal to* $\mathbf{q} \cdot P^n$*, and the (so-called) matrix of transition in* n *steps, whose elements are denoted by* $p_{ij}(n)$*, is equal to* P^n*. In other words, for any* $j \in E$ *we have*

$$\mathbb{P}(X_n = j) = \sum_{i_0 \in E} \sum_{i_1 \in E} \cdots \sum_{i_{n-1} \in E} \mathbf{q}_{i_0} p_{i_0 i_1} p_{i_1 i_2} \cdots p_{i_{n-1} j}.$$

Note that the summation in the above theorem is nothing more than the usual chain rule for conditional probabilities (applied to the special case of Markov chains, where the conditional probability depends only on the last state observed).

The features which make Markov chains so useful in practice are the possibilities of describing particular properties of a process (such as the mean time until...) or making predictions, based on properties of specific types of Markov chains. In order to be able to formulate the limit theorems for Markov chains, we will need some more definitions.

Definition 2. *A Markov chain is* **irreducible***, if for any* $i, j \in E$ *there exists an* $n > 0$*, such that* $p_{ij}(n) > 0$*; in other words, it is possible to go from any state to any state (not necessarily in one step) – each two states* **communicate**.

The Markov chain of Mr. X's activities is irreducible.

Definition 3. *A state* i *has a* **period** *equal to* k*, if* $o(i) = GCD(n : p_{ii}(n) > 0) = k$*. A state is* **aperiodic** *if* $o(i) = 1$*, and* **periodic** *if* $o(i) > 1$.

In the case of our basic example, all states have periods equal to 1. In the case of states 2 and 3 this is obvious; in the case of the first state: we may return to state 1 in two steps (for example $1 \to 2 \to 1$) or three steps (for example $1 \to 2 \to 2 \to 1$), etc., so the greatest common divisor of all times of return is equal to 1. Note that it is not necessary that it is possible to return to a given state in a single step for it to be aperiodic.

In the case of a simple Markov chain described by the transition matrix

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

both states are periodic with periods equal to 2. It is not a coincidence that in the above examples all states are of the same type; we have the following:

Theorem 2. *If a Markov chain is irreducible, all states have the same period.*

This theorem validates the following definition:

Definition 4. *An irreducible Markov chain is* **periodic***, if all the states are periodic; it is* **aperiodic***, if all the states are aperiodic.*

Before we can formulate the limit theorem, we need one more definition:

Definition 5. *A distribution (vector) π is a* **stationary distribution** *(or state) of a Markov chain of a transition matrix P, if*

$$\pi \cdot P = \pi.$$

In other words, a stationary distribution has the property that if any X_n has this distribution, then all subsequent X_{n+k} will also have this distribution.

In order to find the stationary distribution, it is usually necessary to solve the equation from the definition. In the case of our basic example, this equation would translate to the following system of equations:

$$\begin{aligned}
0 \cdot \pi_1 + \frac{1}{2} \cdot \pi_2 + \frac{1}{4} \cdot \pi_3 &= \pi_1 \\
\frac{1}{2} \cdot \pi_1 + \frac{1}{4} \cdot \pi_2 + \frac{1}{2} \cdot \pi_3 &= \pi_2 \\
\frac{1}{2} \cdot \pi_1 + \frac{1}{4} \cdot \pi_2 + \frac{1}{4} \cdot \pi_3 &= \pi_3
\end{aligned}$$

This set of equations has infinitely many solutions – the initial equation is invariant to changes of scale of π. It is therefore necessary to add a condition stating that π is a probability distribution: $\pi_1 + \pi_2 + \pi_3 = 1$; in this case, we the obtain a system of four equations with three unknowns. We may then omit any equation (other than the distribution equation), to find a solution; in our basic case, this would be

$$\begin{aligned}
\pi_1 &= \frac{7}{25} \\
\pi_2 &= \frac{2}{5} \\
\pi_3 &= \frac{8}{25}.
\end{aligned}$$

We are now in a position to formulate the important **Ergodic Theorem** for Markov chains:

Theorem 3. *Let $(X_n)_{n \geqslant 0}$ be an aperiodic irreducible Markov chain over a finite set of states. Then, this Markov chain has a single stationary distribution π, which also satisfies the following property: for any $i, j \in E$, we have*

$$\lim_{n \to \infty} p_{ij}(n) = \pi_j > 0.$$

From the theorem, we have that for large n the behavior of variables X_n does not depend (almost) on the initial state, or the initial distribution of X_0; $p_{ij}(n)$ in the limit does not depend on i. The above theorem allows also to assess the probability that in the far future, the Markov chain will be in a given state. For example, if we were to assess the probability that on a given day we will find Mr. X in a bar, if we knew that his behavior had been following the Markov chain rules for a longer period of time, we would approximate it by $\pi_2 = \frac{2}{5}$. We also know that the probabilities in the stationary distribution will all be non-negative. From the theorem, we may also infer the form of the limit of the matrix P_n: it is a matrix whose all rows are the same, and equal to the stationary distribution.

Note that relaxing the assumptions of irreducibility and aperiodicity destroys the property described by the ergodic theorem. A Markov chain which is not irreducible may have many stationary distributions, and even if a limit of $p_{ij}(n)$ exists, it is not clear which of the stationary distributions that would be. On the other hand, a Markov chain which is periodic may have a single stationary state; this state is in general not the limit of $p_{ij}(n)$, however. In particular, the limit of $p_{ii}(n)$ does not exist.

The stationary distribution is also connected to another notion, namely the mean recurrence time. In general, for any states i and j we can propose the following definition:

Definition 6. *The* **mean first passage time** *from state i to j for an irreducible Markov chain is the expected number of steps to reach state j from i for the first time, denoted by m_{ij}.*

By convention, $m_{ii} = 0$, but we can define:

Definition 7. *The* **mean recurrence time** *for state i for an irreducible Markov chain is the expected number of steps to return to state i for the first time, denoted by m_i.*

We have the following properties:

$$m_{ij} = 1 + \sum_{k \neq j} p_{ik} m_{kj},$$

for $i \neq j$, and

$$m_i = 1 + \sum_{k} p_{ik} m_{ki}.$$

We have

Theorem 4. *Let $(X_n)_{n \geqslant 0}$ be an aperiodic irreducible Markov chain over a finite set of states. Then, we have that the stationary distribution π satisfies*

$$\pi_j = \frac{1}{m_j},$$

where m_j is the mean recurrence time for state j.

Therefore, in our basic example, Mr. X returns to the bar every 2.5 days, on average.

A wholly different class of problems may be solved with the use of Markov chains if we consider a class of chains which are not irreducible. The most useful cases are those where there are "ending" states, i.e. states which can not be left when reached.

Definition 8. *A state i is* **absorbing** *if it is impossible to leave the state; in other words, $p_{ii} = 1$ (while $p_{ij} = 0$ for $j \neq i$).*

These types of Markov chains are useful when modeling processes which may end abruptly (which we model as reaching an absorbing state), such as different types of games. Assume, for example, a game where one of the players (A) has an initial capital amounting to a, and the other (B) – b, and the players toss a coin, betting a dollar at a time, until one of them goes bankrupt (his capital is 0). This game may be modeled with the use of a Markov chain of states $\{0, 1, 2, \ldots, a+b\}$, corresponding to the (temporary) capital of, say, player A; the probabilities of transition are such that $p_{i,i+1} = \frac{1}{2} = p_{i,i-1}$ for "internal" i and $p_{00} = p_{a+b,a+b} = 1$ (the states where player A goes bankrupt or wins everything because B went bankrupt are absorbing).

What we are interested in these types of models are, usually, the probabilities of absorbtion by the particular absorbing states. They may be determined with the use of systems of equations, based on the law of total probability. In the case of the above-mentioned game, these probabilities would (not surprisingly) amount to $\frac{b}{a+b}$ and $\frac{a}{a+b}$ for losing (reaching state 0) and winning (reaching state $a+b$) by player A, respectively.