16 Markov Chains: Reversibility

Assume that you have an irreducible and positive recurrent chain, started at its unique invariant distribution π. Recall that this means that π is the p. m. f. of X_0, and all other X_n as well. Now suppose that, for every n, $X_0, X_1, \ldots, X_n$ have the same joint p. m. f. as their time-reversal $X_n, X_{n-1}, \ldots, X_0$. Then we call the chain *reversible* — sometimes it is, equivalently, also said that its invariant distribition π is reversible. This means that a recorded simulation of a reversible chain looks the same if the "movie" is run backwards.

Is there a condition for reversibility that can be easily checked? The first thing to observe is that if the chain is started at π, reversible or not, the time-reversed chain has the Markov property. This is not completely intuitively clear, but can be checked:

$$
\begin{aligned}
&P(X_k = i | X_{k+1} = j, X_{k+2} = i_{k+2}, \ldots, X_n = i_n) \\
&= \frac{P(X_k = i, X_{k+1} = j, X_{k+2} = i_{k+2}, \ldots, X_n = i_n)}{X_{k+1} = j, X_{k+2} = i_{k+2}, \ldots, X_n = i_n)} \\
&= \frac{\pi_i P_{ij} P_{j i_{k+2}} \cdots P_{i_{n-1} i_n}}{\pi_j P_{j i_{k+2}} \cdots P_{i_{n-1} i_n}} \\
&= \frac{\pi_i P_{ij}}{\pi_j},
\end{aligned}
$$

an expression only dependent on i and j. For reversibility, this expression must be the same as the forward transition probability $P(X_{k+1} = i | X_k = j) = P_{ji}$. Conversely, if both original and time-reversed chain have the same transition probabilities (and we already know that the two start at the same invariant distribution, and that both are Markov), then their p. m. f.'s must agree. We have proved the following useful result.

Theorem 16.1. *Reversibility condition.*

A Markov chain with invariant measure π is reversible if and only if

$$\pi_i P_{ij} = \pi_j P_{ji},$$

for all states i and j.

Another useful fact is that once reversibility is checked, invariance is automatic.

Proposition 16.2. *Reversibility implies invariance. If a probability mass function, π_i satisfies the condition in the previous theorem, then it is invariant.*

Proof. We need to check that, for every j, $\pi_j = \sum_i \pi_i P_{ij}$, and here is how we do it:

$$\sum_i \pi_i P_{ij} = \sum_i \pi_j P_{ji} = \pi_j \sum_i P_{ji} = \pi_j.$$

□

We now proceed to describe the *random walks on weighted graphs*, the most easily recognizable examples of reversible chains. Assume that every undirected edge between vertices i and j in a complete graph has a weight $w_{ij} = w_{ji}$; we think of edges with zero weight as not present at all. When in i, the walker goes to j with probability proportional to w_{ij}, so that

$$P_{ij} = \frac{w_{ij}}{\sum_k w_{ik}}.$$

What makes such random walks easy to analyze is existence of a simple reversible measure. Let

$$s = \sum_{i,k} w_{ik}$$

be the sum of all weights, and then let

$$\pi_i = \frac{\sum_k w_{ik}}{s}.$$

To see why this is a reversible measure, call the double sum s and compute

$$\pi_i P_{ij} = \frac{\sum_k w_{ik}}{s} \cdot \frac{w_{ij}}{\sum_k w_{ik}} = \frac{w_{ij}}{s},$$

which clearly remains the same if we switch i and j.

We should observe that this chain is irreducible exactly when the graph with present edges (those with $w_{ij} > 0$) is connected. The graph can only be periodic, and the period can only be 2 (because the walker can always return in two steps), when it is *bipartite*: the set of vertices V is divided into two sets V_1 and V_2 with every edge connecting a vertex from V_1 to a vertex from V_2. Finally, we note that there is no reason to forbid self-edges: some of the weights w_{ii} may well be nonzero. (However, w_{ii} appear only once in s, while each w_{ij}, $i \neq j$ appears there twice.)

By far the most common examples have no self-edges and all nonzero weights equal to 1 — we already have a name for these cases: random walks on graphs. The number of neighbors of a vertex is commonly called its *degree*. Then the invariant distribution is

$$\pi_i = \frac{\text{degree of } i}{2 \cdot (\text{number of all edges})}.$$

Example 16.1. Consider random walk on the graph below.

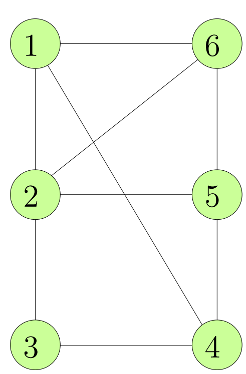

What is the proportion of time the walk spends at vertex 2?

The reversible distribution is

$$\pi_1 = \frac{3}{18}, \pi_2 = \frac{4}{18}, \pi_3 = \frac{2}{18}, \pi_4 = \frac{3}{18}, \pi_5 = \frac{3}{18}, \pi_6 = \frac{3}{18},$$

and thus the answer is $\frac{2}{9}$.

Assume now that the walker may stay at a vertex with probability P_{ii}, but when she does move she moves to a random neighbor as before. How can we choose P_{ii} so that π becomes uniform, $\pi_i = \frac{1}{6}$ for all i?

We should choose weights of self-edges so that the sum of weights for all edges emanating from any vertex is the same. Thus $w_{22} = 0$, $w_{33} = 2$, and $w_{ii} = 1$ for all other i.

Example 16.2. *Ehrenfest chain.* You have M balls, distributed in two urns. Each time, pick a ball at random, move it from the urn where it currently resides to the other urn. Let X_n be the number of balls in urn 1. Prove that this chain has a reversible distribution.

The nonzero transition probabilities are

$$\begin{aligned} P_{0,1} &= 1, \\ P_{M,M-1} &= 1, \\ P_{i,i-1} &= \frac{i}{M}, \\ P_{i,i+1} &= \frac{M-i}{M}. \end{aligned}$$

Time for some inspiration: the invariant measure puts each ball at random into one of the two urns, as switching any ball between the two urns does not alter this assignment. Thus π is Binomial$(M, \frac{1}{2})$,

$$\pi_i = \binom{M}{i}\frac{1}{2^M}.$$

Let's check that this is a reversible measure. The following equalities need to be verified:

$$\begin{aligned} \pi_0 P_{01} &= \pi_1 P_{10}, \\ \pi_i P_{i,i+1} &= \pi_{i+1} P_{i+1,i}, \\ \pi_i P_{i,i-1} &= \pi_{i-1} P_{i-1,i}, \\ \pi_M P_{M,M-1} &= \pi_{M-1} P_{M-1,M}, \end{aligned}$$

and it is straightforward to so so. Note that this chain is irreducible, but not aperiodic (it has period 2).

Example 16.3. *Markov chain Monte Carlo.* Assume that you have a very large probability space, say some subset of $S = \{0,1\}^V$, where V is a large set of n *sites*. Assume also that you have a probability measure on S given via the *energy* (sometimes called *Hamiltonian*) function $E : S \to \mathbb{R}$. The probability of any configuration $\omega \in S$ is

$$\pi(\omega) = \frac{1}{Z}\epsilon^{-\frac{1}{T}E(\omega)}.$$

Here, $T > 0$ is the *temperature*, a parameter, and Z is the normalizing constant that makes $\sum_{\omega \in S} \pi(\omega) = 1$. Such distributions frequently occur in statistical physics, and are often called *Maxwell-Boltzmann* distributions. They have found numerous other applications, however, especially in optimization problems, and yielded an optimization technique called *simulated annealing*.

If T is very large, the role of energy is diminished and states are almost equally likely. On the other hand, if T is very small, the large energy states have a much lower probability than small energy ones, thus the system is much more likely to be found in close to minimal energy states. If we want to find states with small energy, we merely choose some small T and generate at random, according to P, some states, and we have a reasonable answer. The only problem is that although E is typically a simple function, π is very difficult to evaluate exactly, as Z is some enormous sum. (There are a few celebrated cases, called *exactly solvable systems*, in which exact computations are difficult but possible.)

Instead of generating a random state directly, then, we design a Markov chain, which has π as its invariant distribution. It is very common that the convergence to π is quite fast, and the necessary number of steps of the chain to get close to π is some small power of n. This is in startling contrast to the size of S, which is typically exponential in n.

We will illustrate this on the *Knapsack problem*. Assume that you are a burglar and have just broke into a jewellery store. You see a large number n of items, with weights w_i and values v_i. Your backpack (knapsack) has a weight limit b. You are faced with a question of how to fill in your backpack, that is, you have to maximize the combined value of items you will carry out

$$V = V(\omega_1, \dots, \omega_n) = \sum_{i=1}^{n} v_i \omega_i$$

subject to the constraints that $\omega_i \in \{0, 1\}$ and that the combined weight does not exceed the backpack capacity,

$$\sum_{i=1}^{n} w_i \omega_i \le b.$$

This problem is known to be NP-hard; there is no good algorithm to solve it quickly.

The set S of *feasible solutions* $\omega = (\omega_1, \dots, \omega_n)$ that satisfy the constraints above will be our state space, and the energy function E on S is given as $E = -V$, as we want to maximize V. The temperature T measures how good a solution we are happy with — the idea of simulated annealing is in fact a gradual lowering of the temperature to improve the solution. There is give and take: higher temperature improves the speed of convergence and lower temperature improves the quality of the result.

Finally, we are ready to specify the Markov chain (sometimes called *Metropolis algorithm*, in honor of N. Metropolis, a pioneer in computational physics). Assume that the chain is at state

ω at time t, i.e., $X_t = \omega$. Pick a site i, uniformly at random. Let ω^i be the same as ω except that its i'th coordinate is flipped: $\omega^i_i = 1 - \omega_i$. (This means that the status of the i'th item is changed from in to out or from out to in.) If ω^i is not feasible, then $X_{t+1} = \omega$, the state is unchanged. Otherwise, evaluate the difference in energy $E(\omega^i) - E(\omega)$, and proceed as follows:

- if $E(\omega^i) - E(\omega) \le 0$, then make the transition to ω^i, $X_{t+1} = \omega^i$;
- if $E(\omega^i) - E(\omega) > 0$, then make the transition to ω^i with probability $e^{\frac{1}{T}(E(\omega)-E(\omega^i))}$, or else stay at ω.

Note that in the second case the new state has higher energy, but, in physicist's terms, we tolerate the transition because of temperature, which corresponds to energy input from the environment.

We need to check that this chain is irreducible on S: to see this note that we can get from any feasible solution to empty backpack by removing object one by one, and then back by reversing steps. Thus the chain has a unique invariant measure, but is it the right one, that is, π? In fact, the measure π on S is reversible. We need to show that for any pair $\omega, \omega' \in S$

$$\pi(\omega)P(\omega, \omega') = \pi(\omega')P(\omega', \omega),$$

and this is enough to do with $\omega' = \omega^i$, for arbitrary i, and assume that both are feasible (as only such transitions are possible). Note first that the normalizing constant Z cancels out (the key feature of this method) and so does the probability $\frac{1}{n}$ that i is chosen. If $E(\omega^i) - E(\omega) \le 0$, then the equality reduces to

$$e^{-\frac{1}{T}E(\omega)} = e^{-\frac{1}{T}E(\omega^i)} e^{\frac{1}{T}(E(\omega^i)-E(\omega))},$$

and similarly in the other case.

Problems

1. Determine the invariant distribution for random walk in Examples 12.4 and 12.10.

2. A total of m white and m black balls are distributed into two urns, with m balls per urn. At each step, a ball is randomly selected from each urn and the two balls are interchanged. The state of this Markov chain can be thus described by the number of black balls in urn 1. Guess the invariant measure for this chain and prove that it is reversible.

3. Each day, your opinion on a particular political issue is either positive, neutral, or negative. If it is positive today it is neutral or negative tomorrow with equal probability. If it is neutral or negative, it stays the same with probability 0.5, and otherwise it is equally likely to be either of the other two possibilities. Is this a reversible Markov chain?

4. A king moves on the standard 8×8 chessboard. Each time, it makes one of the available legal moves (to a horizontally, vertically or diagonally adjacent square) at random. (a) Assuming that the king starts at one of the four corner squares of the chessboard, compute the expected number of steps before it returns to the starting position. (b) Now you have two kings, they both start at the same corner square, and move independently. What is now the expected number of steps before they simultaneously occupy the starting position?

Solutions to problems

1. Answer: $\pi = \left[\frac{1}{5}, \quad \frac{3}{10}, \quad \frac{1}{5}, \quad \frac{3}{10}\right]$.

2. If you choose m balls to put into urn 1 at random, you get

$$\pi_i = \frac{\binom{m}{i}\binom{m}{m-i}}{\binom{2m}{m}},$$

and the transition probabilities are

$$P_{i,i-1} = \frac{i^2}{m^2}, \; P_{i,i-1} = \frac{(m-i)^2}{m^2}, \; P_{i,i} = \frac{2i(m-i)}{m^2}.$$

Reversibility check is routine.

3. If the three states are labeled, in the order given, 1, 2, and 3, then we have

$$P = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{2} \end{bmatrix}.$$

The only way to check reversibility is to compute the invariant distribution π_1, π_2, π_3, form the diagonal matrix D with π_1, π_2, π_3 on the diagonal, and check that DP is symmetric. We get $\pi_1 = \frac{1}{5}$, $\pi_2 = \frac{2}{5}$, $\pi_3 = \frac{2}{5}$, and DP indeed is symmetric, so the chain is reversible.

4. This is a random walk on a graph with 64 vertices (squares) and degrees 3 (4 corner squares), 5 (24 side squares), and 8 (36 remaining squares). If i is a corner square, $\pi_i = \frac{3}{3 \cdot 4 + 5 \cdot 24 + 8 \cdot 36}$, so the answer to (a) is $\frac{420}{3}$. In (b), you have two independent chains, so $\pi_{(i,j)} = \pi_i \pi_j$ and the answer is $\left(\frac{420}{3}\right)^2$.