1 Convergence to the Invariant Distribution

In this section, we will state (without proof) the so-called Fundamental Theorem of Markov Chains, which says that, under mild conditions, a finite Markov chain will always converge, as time tends to infinity, to a unique invariant distribution, regardless of its initial state. This implies that, in the long term, the probability of finding the chain in any given state has a certain fixed value that doesn't change over time.

To guarantee this nice behavior, we need two conditions. The first is known as *irreducibility*.

Definition 26.1 (Irreducible)**.** *A Markov chain is* irreducible *if it can go from every state i to every other state j in a finite number of steps.*

We can picture this condition graphically as follows. For a Markov chain to be irreducible, its state transition diagram must be "strongly connected", i.e., there must exist a directed path of transitions (each with non-zero probability) from every state i to every state j. For example, the two-state Markov chain in Figure 1 is irreducible for all $0 < a \leq 1$, but not for $a = 0$. And the Markov chain in Figure 2 is irreducible. Note that irreducibility is a necessary condition for convergence, since otherwise there will be states that are unreachable from some initial states, which will therefore never be visited. It turns out that irreducibility guarantees that a stationary distribution π with $\pi(i) > 0$ for all i not only exists but is also unique.

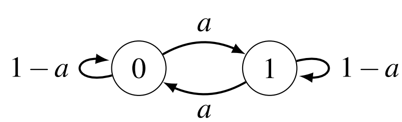

Figure 1: A simple Markov chain

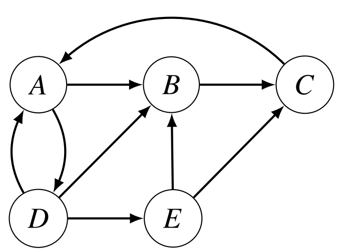

Figure 2: A five-state Markov chain. The outgoing arrows are equally likely.

For the remainder of this section, since we are focusing on convergence, we will restrict our attention to irreducible Markov chains. In later sections, when we shift our focus to hitting times, we will naturally also consider non-irreducible chains that have absorbing states.

The second condition we need for convergence requires a short detour into the notion of *periodicity*.

Definition 26.2 (Period)**.** *For any state i in an irreducible finite Markov chain, the* period *of i is defined as*

$$d(i) := \gcd\{n > 0 : P^n(i,i) > 0\}.$$

Let's parse this definition. The set of values in the gcd expression are exactly the lengths of all possible paths in the Markov chain from state i back to itself: that's because the entry $P^n(i,j)$ is the probability that the chain goes from state i to state j in exactly n steps, so $P^n(i,i) > 0$ if and only if there is at least one length-n path from i back to itself. (Note that we're considering *all* paths here, not just simple paths.) The value $d(i)$ determines whether there is any kind of periodic behavior in these path lengths: if $d(i) > 1$ then the return visits to i, starting from i, can only occur at intervals of at least $d(i)$ steps; on the other hand, if $d(i) = 1$ there is no such periodic behavior. A very simple example of a situation where $d(i) > 1$ is the 2-state example in Figure 1 when $a = 1$. In this case the chain just bounces back and forth between the two states (the self-loops have probability 0 so should be removed from the graph), so all paths from state 0 back to itself have even length and $d(i) = 2$.

An important fact about periods is that, in an irreducible Markov chain, they are the same for all states!

Proposition 26.1. *In an irreducible finite Markov chain, $d(i) = d(j)$ for all pairs of states i, j.*

Proof. Since the chain is irreducible, there is a path $\mathscr{P}_{ij}$ from i to j and a path $\mathscr{P}_{ji}$ from j to i. Call the lengths of these paths ℓ_{ij} and ℓ_{ji}, respectively.

Now consider the following path from i back to itself: take path $\mathscr{P}_{ij}$ to j, then *any* path (of length k, say) from j back to itself, and finally path $\mathscr{P}_{ji}$ back to i. (Note that paths from j back to itself certainly exist, since e.g. $\mathscr{P}_{ji}$ followed by $\mathscr{P}_{ij}$ is one such path.) The total length of this path is $\ell_{ij} + \ell_{ji} + k$. The following two observations follow from Definition 26.2:

- $d(i) \mid (\ell_{ij} + \ell_{ji} + k)$ (since this is the length of the above path from i back to itself);
- $d(i) \mid (\ell_{ij} + \ell_{ji})$ (since $\mathscr{P}_{ij}$ followed by $\mathscr{P}_{ji}$ is also a path from i back to itself.

Putting these together immediately implies that $d(i) \mid k$. But since k was the length of an *arbitrary* path from j back to itself, $d(i)$ must in fact divide *all* such path lengths, and therefore $d(i) \mid d(j)$ (since $d(j)$ is the gcd of all such path lengths).

By a symmetrical argument, we also see that $d(j) \mid d(i)$. But if $d(i) \mid d(j)$ and $d(j) \mid d(i)$, then we must have $d(i) = d(j)$, as claimed. □

We are now ready to define the second property required for convergence.

Definition 26.3 (Aperiodic)**.** *An finite irreducible Markov chain is* aperiodic *if the period $d(i) = 1$ for all states i.*

For an example of an aperiodic Markov chain, consider Figure 2. Note that there are paths of length 2 and 3 from A back to itself, so $d(A) = 1$. Since the chain is irreducible, this implies by Proposition 26.1 that all states have period 1, so the chain is aperiodic. Another useful fact to bear in mind is that if an irreducible Markov chain contains a self-loop on *any* state, then it is aperiodic. (Why?) This applies, for example, to the chain in Figure 1 for any $a \in (0,1)$.

When a chain is periodic, it can't converge to an invariant distribution from a given starting state, since the probability of being at that state will oscillate between zero and non-zero values according to the period. You can see this in the very simple example of Figure 1 when $a = 1$. However, it turns out that this is the *only* obstacle to convergence for an irreducible chain, as we now state.

Theorem 26.1 (Fundamental Theorem of Markov Chains)**.** *For any finite, irreducible, aperiodic Markov chain, the probability distribution at time n for any initial state X_0 converges as $n \to \infty$ to π, where π is* *the unique invariant distribution and $\pi(i) > 0$ for all states i. In other words, for any X_0 and any state i, $\mathbb{P}[X_n = i] \to \pi(i)$ as $n \to \infty$.*

There are several alternative proofs of this theorem, all of which are beyond the scope of the course. However, we will briefly discuss some sample applications.

First, recall that the chain in Figure 1 with $0 < a < 1$ is aperiodic and irreducible with invariant distribution $\begin{bmatrix} \frac{1}{2} & \frac{1}{2} \end{bmatrix}$. Thus $\mathbb{P}[X_n = 0] \to \frac{1}{2}$ as $n \to \infty$, and similarly for $\mathbb{P}[X_n = 1]$, regardless of where we start. Second, we have seen that the chain in Figure 2 is aperiodic and irreducible with invariant distribution $\frac{1}{39}\begin{bmatrix} 12 & 9 & 10 & 6 & 2 \end{bmatrix}$. Thus, e.g., as $n \to \infty$, $\mathbb{P}[X_n = D] \to \frac{6}{39}$.

For a more interesting example, consider the problem of shuffling a deck of 52 cards. Here our goal is to construct a *uniform* distribution over all 52! permutations of the deck (a huge number, around 8×10^{67}). How do we achieve this feat? We can model the shuffling process as a Markov chain whose states are these 52! permutations, and whose transitions correspond to the kind of "riffle shuffles" that card players perform. To model real-life shuffling, we would need a mathematical model of these riffles, which does exist but is a little tricky to describe. Instead, we'll consider a "slow" shuffle in which, at each step, we pick two random cards (with replacement) in the current deck and switch their positions. We claim that this shuffle is aperiodic and irreducible and that its invariant distribution is uniform.

The fact that it's irreducible follows from the standard fact that *any* permutation can be written as a sequence of transpositions (i.e., switches of two cards); thus we can get from any permutation to any other one via a sequence of our simple card switching operations. To see that it's aperiodic, note that $P(i,i) > 0$ for any state i (because we could pick the same card, meaning that nothing happens); this means that $\gcd\{n > 0 : P^n(i,i) > 0\} = 1$ for all i, so the chain is aperiodic. Finally, to see that the invariant distribution is uniform, we note that the transition matrix P is *symmetric*, i.e., $P(i,j) = P(j,i)$ for all i,j; this follows because either the permutations i,j differ by the transposition of just two cards, c,c' (say), in which case $P(i,j) = P(j,i) = \frac{2}{52^2}$ because both transitions are effected by picking cards c and c' (in either order); or they don't, in which case $P(i,j) = P(j,i) = 0$.

Now *any* irreducible, aperiodic Markov chain that is symmetric has uniform invariant distribution, as we can easily check from the balance equations, as follows. We need to verify that the uniform distribution $\pi(i) = \frac{1}{N}$ satisfies the balance equations $\pi(j) = [\pi P](j)$, where in this case $N = 52!$. But

$$[\pi P](j) = \sum_i \pi(i)P(i,j) = \sum_i \frac{1}{N} P(j,i) = \frac{1}{N} = \pi(j),$$

where in the second equality we used the symmetry of P and in the third equality we used the fact that $\sum_i P(j,i) = 1$.

Putting all this together gives us the following perhaps surprising conclusion: if you start from any ordering of the deck, and perform enough random switches of pairs of cards, then you will eventually reach an (almost) perfectly shuffled deck! The same holds for mathematical models of the shuffles used in casinos (namely, the "riffle" or "dovetail" shuffle), which reach the uniform distribution after fewer steps. (For professional dealers, it is generally accepted that 7 riffle shuffles achieve a deck that is shuffled well enough that even professional card players cannot exploit any remaining structure.)

2 Bonus Content: Random Walk on an Undirected Graph

A class of Markov chains that arises very often in practice are known as random walks on graphs. Let $G = (V,E)$ be an undirected graph, as we have seen earlier in the class. Define a Markov chain as follows:

the state space is V (the vertices of the graph), and at each step, if the process is at v, it moves to a neighbor u of v chosen uniformly at random.

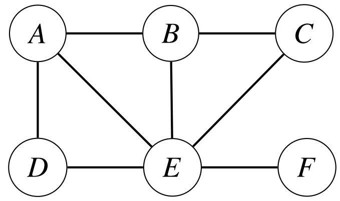

Figure 3: An undirected graph on six vertices. To interpret this as the transition diagram of a Markov chain, view each undirected edge as a pair of directed edges, one in each direction. Transition probabilities for all edges out of a given vertex are equal.

Figure 3 shows a simple example. For this graph, you should check that the transition matrix of the random walk is

$$P = \begin{bmatrix} 0 & \frac{1}{3} & 0 & \frac{1}{3} & \frac{1}{3} & 0 \\ \frac{1}{3} & 0 & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ 0 & \frac{1}{2} & 0 & 0 & \frac{1}{2} & 0 \\ \frac{1}{2} & 0 & 0 & 0 & \frac{1}{2} & 0 \\ \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & \frac{1}{5} & 0 & \frac{1}{5} \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}.$$

What can we say about random walk on a graph? First, it should be clear that the process is irreducible if and only if G is connected (since that corresponds precisely to the condition that there is a path from every vertex to every other vertex). Second, the random walk is aperiodic if and only if the graph is *not* bipartite. (You may like to prove this: the "only if" direction is easy, while the "if" direction follows from the fact that a non-bipartite graph must contain a cycle of odd length.) So, assuming G is connected and not bipartite, the random walk on G converges to a unique invariant distribution.

What is that distribution? We claim that it has a very nice form, namely

$$\pi(v) = \frac{\deg(v)}{D}, \tag{1}$$

where $\deg(v)$ denotes the degree of vertex v and $D = \sum_{v \in V} \deg(v)$. ($D$ here is just a normalizing factor to make the probabilities sum to 1. Note that in fact $D = 2|E|$.) To prove this we just have to verify that π as defined in (1) satisfies the balance equations $[\pi P](v) = \pi(v)$ for all v. This follows from

$$[\pi P](v) = \sum_{u:\{u,v\} \in E} \pi(u) P(u,v) = \sum_{u:\{u,v\} \in E} \frac{\deg(u)}{D} \times \frac{1}{\deg(u)} = \sum_{u:\{u,v\} \in E} \frac{1}{D} = \frac{\deg(v)}{D} = \pi(v).$$

In the second equality here, we used the fact that $P(u,v) = \frac{1}{\deg(u)}$ for all neighbors v of u.

Thus we see that, for random walk on a (connected, non-bipartite) graph, the probability of finding the walk in a given vertex v after many steps is proportional to $\deg(v)$. For the example graph in Figure 3, these proportions are

$$\pi(A) = \frac{3}{16}; \quad \pi(B) = \frac{3}{16}; \quad \pi(C) = \frac{2}{16}; \quad \pi(D) = \frac{2}{16}; \quad \pi(E) = \frac{5}{16}; \quad \pi(F) = \frac{1}{16}.$$