EE 126 — Probability and Random Processes
February 27, 2018 — Discrete-Time Markov Chains I

1 Reading

W1. Markov Chains

A **Markov chain** is a sequence $\{X(n), n \geq 0\}$ that goes from state i to state j with transition probability $P(i, j)$ (independently of the states it visited before). We call $X(n)$ the state of the Markov chain at time n, where $X(0)$ is the initial state. We refer to $\mathcal{X}$ as the set of states.

It follows that

$$P(X(n+1) = j \mid X(n) = i,\ X(m) \text{ for all } m < n) = P(i, j)$$

for all $i, j \in \mathcal{X}$ and $n \geq 0$.

The **Markov property** revolves around the idea that the probability of moving from state i to state j does not depend on any state before i. (Each state contains all the information necessary to predict the future of the process.)

If a Markov chain is in state j with probability $\pi_n(j)$ at step n for some $n \geq 0$, then it will be in state i at step $n+1$ with probability $\pi_{n+1}(i)$ where

$$\pi_{n+1}(i) = \sum_{j \in \mathcal{X}} \pi_n(j) P(j, i)$$

for $i \in \mathcal{X}$.

Note: $\pi_n(i) = \pi_0(i)$ for all $n \geq 0$ and all $i \in \mathcal{X}$ iff $\pi_0 = \pi_0 P$ [where we treat π as a row vector with components $\pi_0(i)$ and P as a square matrix with entries $P(i, j)$]. We also have $\pi_{n+1} = \pi_n P$ and $\pi_n = \pi_0 P^n$. Incidentally, if the first result is true, π_0 is an **invariant distribution** – a nonnegative solution π of $\pi = \pi P$ whose components sum to one.

A Markov chain is **irreducible** if it can go from any state to any other state (even if it takes many time steps). If a Markov chain is finite and irreducible, it has a unique invariant distribution π, and $\pi(i)$ is the **long-term fraction of time** that $X(n)$ is *almost surely* equal to i.

$$\pi(i) = \lim_{N \to \infty} \frac{1}{N} \sum_{n=0}^{N-1} \mathbb{1}[X(n) = i]$$

W2.4. Law of Large Numbers for Markov Chains

An **invariant probability** of a state is its probability under the invariant distribution of the Markov chain. For example, the Markov chain $X(n)$ on $[0, 1]$ with $P(0, 1) = P(1, 0)$ spends half of the time in state 0, and half in state 1. Therefore the invariant probability of state 0 would be $1/2$. (Note: the *invariant probability* is equivalent to the *long-term fraction of time* that a finite irreducible Markov chain spends in a given state.)

Long-term fraction of time is justified by the strong law of large numbers.

W13.3. Infinite Markov Chains

Let us study Markov chains on a countably infinite space, i.e. where $\mathcal{X} = \{0, 1, ...\}$ instead of (e.g.) $\{1, 2, ..., N\}$. We will assume we are given an initial distribution $\pi = \{\pi(x), x \in \mathcal{X}\}$, where $\pi(x) \geq 0$ and $\sum_{x \in \mathcal{X}} \pi(x) = 1$. We will also assume we are given a set of nonnegative numbers $\{P(x, y)\}$ for $x, y \in \mathcal{X}$ such that $\sum_{y \in \mathcal{X}} P(x, y) = 1$ for all $x \in \mathcal{X}$.

Then the sequence $\{X(n)\}$ for $n \geq 0$ is a Markov chain with initial distribution π and probability transition matrix P if

$$P(X(0) = x_0, ..., X(n) = x_n) = \pi(x_0)P(x_0, x_1) \cdots P(x_{n-1}, x_n)$$

7.1. Discrete-Time Markov Chains

We're interested in models for which some notion of a **state** encodes all of the information necessary to predict the future from the past. The state should take on only a finite set of values, and should change over time according to time-independent transition probabilities. Such models can be applied to an incredible number of dynamical systems whose evolution over time involves uncertainty.

One model that meets our criteria is the **discrete-time Markov chain**, in which the state changes at discrete time instants indexed by the variable n. At each time step, the state is denoted X_n and belongs to a finite set $\mathcal{S} = \{1, ..., m\}$ of potential states (this is the **state space**). There are also transition probabilities: if the current state is i, there is probability p_{ij} that the next state will be j. More specifically,

$$p_{ij} = P(X_{n+1} = j \mid X_n = i)$$

for $i, j \in \mathcal{S}$. The key assumption is that the transition probabilities are the same across all time steps; p_{ij} will retain its value no matter how state i was reached. This can be formalized as the **Markov property**:

$$P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, ..., X_0 = i_0) = P(X_{n+1} = j \mid X_n = i) = p_{ij}$$

The Markov property tells us that the probability law of the next state X_{n+1} depends on the past *only* through the value of the present state X_n.

Specification of Markov Models
A Markov chain is specified by a **set of states** $\mathcal{S} = \{1, ..., m\}$, a **set of possible transitions** (pairs (i, j) for which $p_{ij} > 0$), and the **values of all positive** p_{ij}. It is then a sequence of random variables $X_0, X_1, X_2, ...$ that take values in $\mathcal{S}$ and satisfy

$$P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, ..., X_0 = i_0) = p_{ij}$$

for all possible sequences of states.

A Markov chain model can thus be entirely encoded by a **transition probability matrix**, i.e. a 2D array where the p_{ij} is the element in the ith row and jth column. We can also visualize the model as a transition probability graph, in which the nodes are the states and the edges are the possible transitions (weighted by the p_{ij} values).

If we have a Markov chain model, we can easily compute the probability of any sequence of states:

$$P(X_0 = i_0, X_1 = i_1, ..., X_n = i_n) = P(X_0 = i_0)p_{i_0 i_1} p_{i_1 i_2} \cdots p_{i_{n-1} i_n}$$

We might also be interested in the **n-step transition probability**, i.e. the probability that the state after n time steps will be j, given that the current state is i:

$$r_{ij}(n) = P(X_n = j \mid X_0 = i)$$

This can be computed via the recursion of the Chapman-Kolmogorov equation:

$$\begin{aligned} r_{ij}(n) &= \sum_{k=1}^{m} P(X_{n-1} = k \mid X_0 = i)P(X_n = j \mid X_{n-1} = k, X_0 = i) \\ &= \sum_{k=1}^{m} r_{ik}(n-1) \cdot p_{kj} \end{aligned}$$

for $n > 1$ and all i, j, starting with $r_{ij}(1) = p_{ij}$.

In Markov chains, there is a variety of types of states and asymptotic occupancy behavior.

7.2. Classification of States

A state j is **accessible** from state i if the n-step transition probability $r_{ij}(n)$ is positive for some n (i.e. there is a possibility of reaching j after some number of time periods). A state i is **recurrent** if for every j that is accessible from i, i is also accessible from j. If a state is not recurrent, it is **transient** (i.e. one of the states accessible from i does not permit the possibility of ever getting back to i).

If i is a recurrent state, then the set of states $A(i)$ that are accessible from i make up a **recurrent class**. The states in $A(i)$ will all be accessible from each other, and no state outside of $A(i)$ will be accessible from them. Mathematically, if i is recurrent then $A(i) = A(j)$ for all $j \in A(i)$.

A Markov chain can be decomposed into one or more recurrent classes, potentially along with some transient states. Note that at least one recurrent state will be accessible from every transient state.

If the initial state is transient, the state trajectory will contain an initial portion of transient states and a final portion of recurrent states from the same class. Once the state ends up in a class of recurrent states, it will stay in that class forever, and all states in the class will be visited an infinite number of times.

7.2.1. Periodicity

A recurrent class is **periodic** if its states can be grouped into $d > 1$ disjoint subsets $S_1, ..., S_d$ such that all transitions from one subset lead to the next subset. In other words, if $i \in S_k$ and $p_{ij} > 0$, then $j \in S_1$ if $k = d$ and $j \in S_{k+1}$ if $k = 1, ..., d-1$ (where j is the next state).

Given a periodic recurrent class, a positive time n, and a state i in the class, there must be at least one state j for which $r_{ij}(n) = 0$. Conversely, if a class is aperiodic, there exists a time n such that $r_{ij}(n) > 0$ for all i, j in the class.

7.3. Steady-State Behavior

We're often interested in long-term state occupancy behavior, i.e. the n-step transition probabilities $r_{ij}(n)$ when n is large. These probabilities may converge to steady-state values which are independent of the initial state. Such limiting values, denoted by π_j, are defined so that $\pi_j \approx P(X_n = j)$ when n is large. π_j is called the **steady-state probability** of j. The **convergence theorem** is as follows:

Say we have a Markov chain with a single aperiodic recurrent class. Then the steady-state probabilities π_j for each state j have the following properties:

- For each j, we have $\lim_{n\to\infty} r_{ij}(n) = \pi_j$ for all i.
- π_j is the unique solution to the equations $\pi_j = \sum_{k=1}^{m} \pi_k p_{kj}$ and $1 = \sum_{k=1}^{m} \pi_k$.
- $\pi_j = 0$ for all transient states, and $\pi_j > 0$ for all recurrent states.

The steady-state probabilities form a probability distribution, the **stationary distribution**, on the state space. The following equations, which arise from the first part of the convergence theorem, are known as the **balance equations**.

$$\pi_j = \sum_{k=1}^{m} \pi_k p_{kj} \quad \text{for } j = 1, ..., m$$

Steady-State Probabilities as Expected State Frequencies
For a Markov chain with a single aperiodic class, the steady-state probabilities π_j satisfy

$$\pi_j = \lim_{n\to\infty} \frac{v_{ij}(n)}{n}$$

where $v_{ij}(n)$ is the expectation of the number of visits to state j within the first n transitions, starting from state i. Hence π_j is the long-term expected fraction of time that the state is equal to j.

Expected Frequency of a Particular Transition
Consider n transitions of a Markov chain with a single aperiodic class, starting from a given initial state. Let $q_{jk}(n)$ be the expected number of such transitions that go from state j to state k. Then, irrespective of the initial state, we have

$$\lim_{n\to\infty} \frac{q_{jk}(n)}{n} = \pi_j p_{jk}$$

By the frequency interpretation, the balance equation expresses the fact that the expected frequency π_j of visits to j is simply the sum of the expected frequencies $\pi_k p_{kj}$ of transitions that lead to j.

7.4. Absorption Probabilities and Expected Time to Absorption

Let's imagine every recurrent class as an amorphous blob that *absorbs* things and never lets them go. In this context, a recurrent state k is **absorbing** and we have $p_{kk} = 1$ and $p_{kj} = 0$ for $j \neq k$. Once we're in an absorbing state we'll never get out!

Now, choose an absorbing state s. We will call a_i the probability that s is eventually reached, given that we start from state i. We can determine a_i by solving the linear equations $a_s = 1$, $a_i = 0$ for all absorbing $i \neq s$, and $a_i = \sum_{j=1}^m p_{ij} a_j$ for transient i. (We assume that all states are either transient or absorbing.)

We call μ_i the **expected time to absorption** for state i (i.e. the expected number of transitions until we enter a recurrent state). μ_i is defined as $\mathbb{E}[\min\{n \geq 0 \mid X_n \text{ is recurrent}\} \mid X_0 = i]$, i.e. the unique solution to the linear equations $\mu_i = 0$ (for recurrent states i) and $\mu_i = 1 + \sum_{j=1}^m p_{ij}\mu_j$ (for transient states i). In other words, the time to absorption from a transient state should be 1 plus the expected time to absorption starting from the next state.

We call t_i the **mean first passage time from state i to state s**, i.e. the expected number of transitions to reach s for the first time if we start from i. s should be a recurrent state. We can obtain this by solving the equations $t_s = 0$ and $t_i = 1 + \sum_{j=1}^m p_{ij} t_j$ (for $i \neq s$).

Finally, we can compute the **mean recurrence time** of recurrence state s, i.e. t_s^* the number of transitions up to the first return to s, given that we start from s. This is given by $t_s^* = 1 + \sum_{j=1}^m p_{sj} t_j$. A translation: the time to return to s is just 1 plus the expected time to reach s from the next state. The next state is j with probability p_{sj}.

The previous two paragraphs assume a Markov chain with only one recurrent class.

2 Lecture

Finite Markov Chains

Markov chains have the memoryless property:

$$P(X_{n+1} = x_{n+1} \mid X_n = x_n, X_{n-1} = x_{n-1}, ..., X_0 = x_0) = P(X_{n+1} = x_{n+1} \mid X_n = x_n)$$

The past is independent of the future given the present. Formally, we have a finite set of states $\mathcal{X} = \{1, 2, ..., k\}$, a probability distribution on the state space π_0, and the transition probabilities P_{ij} (defined for all $i, j \in \mathcal{X}$). Note that P_{ij} (the ijth entry of our transition probability matrix) is equivalent to $P(i, j)$.

We have $P(X_0 = i) = \pi_0(i)$ for all $i \in \mathcal{X}$. We also have $P(X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, ..., X_0 = i_0) = P_{ij}$. (*Time homogeneity*: at any time, the probability that we're going from i to j is the same.)

Because of the Markov property,

$$\begin{aligned} &P(X_0 = i, X_1 = j, X_2 = i_2, X_3 = i_3, ..., X_n = i_n) \\ &= P(X_0 = i) \cdot P(X_1 = j \mid X_0 = i) \cdot P(X_2 = i_2 \mid X_1 = j, X_0 = i) \cdots P(X_n = i_n \mid X_{n-1} = i_{n-1}, X_{n-2} = i_{n-2}, ..., X_0 = i) \\ &= P(X_0 = i) \cdot P(X_1 = j \mid X_0 = i) \cdot P(X_2 = i_2 \mid X_1 = j) \cdots P(X_n = i_n \mid X_{n-1} = i_{n-1}) \\ &= \pi_0(i) \cdot P_{ij} \cdot P_{j,i_2} \cdots P_{i_{n-1},i_n} \end{aligned}$$

Example: Prof. Ramchandran's office hours. We will assume that there are three states: zero people in OH, one person in OH, or two people in OH. Our state sequence is $X_0, X_1, X_2, X_3, \ldots$ where $X_i \in \{S_0, S_1, S_2\}$ for $i = 0, 1, \ldots$. (Each X_i corresponds to a time.) We define the state transition matrix P as

$$P = \begin{bmatrix} P_{00} & P_{01} & P_{02} \\ P_{10} & P_{11} & P_{12} \\ P_{20} & P_{21} & P_{22} \end{bmatrix} = \begin{bmatrix} 0.8 & 0.2 & 0 \\ 0 & 0.3 & 0.7 \\ 0.6 & 0.4 & 0 \end{bmatrix}$$

where the ijth entry is "from i to j." Note that the rows sum to one; this is therefore called a *row-stochastic matrix*. This formulation might help us answer many questions of interest, e.g. average number of students in office hours or $P(X_{10} = 2 \mid X_0 = 0)$. However, we need a framework, and this is where Markov chains come in.

Note: we define (for example) $\pi_0(0) = P(X_0 = 0)$ and $\pi_1(0) = P(X_1 = 0)$. One question of interest: given π_0, what is π_n? In row vector form, we have $\pi_0 = \begin{bmatrix}\pi_0(0) & \pi_0(1) & \pi_0(2)\end{bmatrix}$ and $\pi_n = \begin{bmatrix}\pi_n(0) & \pi_n(1) & \pi_n(2)\end{bmatrix}$. Again, we also define P_{ij} as $P(S_i \to S_j)$.

Irreducible and Aperiodic Markov Chains

A Markov chain is **irreducible** if we can go from any state to any other state, possibly in many steps. If a Markov chain is irreducible, let $d(i) = \gcd(n \geq 1 \mid P^n_{ii} > 0)$. P^n_{ii} is the probability that we go from state i to state i in n steps. So this is the GCD of all of the ns greater than 1 for which $P^n_{ii} > 0$.

$d(i)$ is the same (say $d(i) = d$) for all states i in an irreducible Markov chain. **Periodic** with period d means we can only return in multiples of d, i.e. $P^n_{ii} > 0$ only when n is a multiple of d.

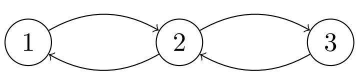

This is irreducible but not aperiodic. If we start at state 1 we can only come back in 2 steps, 4 steps, 6 steps, etc.... i.e. $d = 2$, while aperiodicity requires $d = 1$.

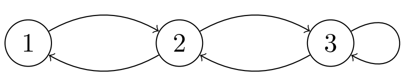

This is irreducible and aperiodic. We can go from any state to any other state, and in this case $d = 1$ (because if we start at state 3, we can come back in 1, 2, 3, ... steps and the GCD of this is 1).

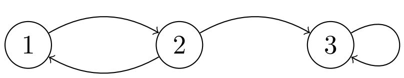

This is *not* irreducible. Once we end up in state 3, we're stuck; we can't go from state 3 to any other state.

Example: OH again. Let's find π_n, the distribution of X_n. We know that

$$\begin{aligned}\pi_{m+1}(j) &= P(X_{m+1} = j) \\ &= \sum_i P(X_m = i) P(X_{m+1} = j \mid X_m = i) \\ &= \sum_i \pi_m(i) P_{ij}\end{aligned}$$

Hence $\pi_{m+1}(j) = \sum_{i \in \mathcal{X}} \pi_m(i) P_{ij}$ for all $j \in \mathcal{X}$.

With π_m and π_{m+1} as row vectors, it must then be the case that $\pi_{m+1} = \pi_m \cdot P$. (Recall that P is the $|\mathcal{X}| \times |\mathcal{X}|$ matrix of transitions from i to j.) Then $\pi_1 = \pi_0 P$, $\pi_2 = \pi_1 P = \pi_0 P^2$, and $\pi_n = \pi_0 P^n$.

Balance Equations

Question: is there a starting state distribution π_0 such that $\pi_m = \pi_0$ for all m? (In such a case, our probability of being in a state will be exactly what we started with; we will never move away from the initial distribution.)

This kind of distribution π_0 is called an **invariant (stationary)** distribution, and is defined formally via the rule "π_0 is invariant iff $\pi_0 P = \pi_0$." If π_0 is invariant, the distribution of X_n is the same as the distribution of X_0.

When we have a stationary distribution, "FLOW IN" = "FLOW OUT" for all states in the Markov chain. Recall that $\pi_{m+1}(j) = \sum_{i \in \mathcal{X}} \pi_m(i) P_{ij}$ for all $j \in \mathcal{X}$. The balance equations are

$$\sum_j \pi(j) P_{ji} = \pi(i)$$

i.e.

$$\pi(i) P_{ii} + \sum_{j \neq i} \pi(j) P_{ji} = \pi(i)$$

Then $\sum_{j \neq i} \pi(j) P_{ji} = \pi(i)(1 - P_{ii}) = \pi(i) \sum_{j \neq i} P_{ij}$. In other words, the amount of probability flow coming into i is equivalent to the amount of probability flow coming out of i.

The balance equation

$$\sum_{j \neq i} \pi(j) P_{ji} = \pi(i) \sum_{j \neq i} P_{ij}$$

is simply saying "flow into state i" = "flow out of state i." Flow conservation! By isolating each state and looking at the flow going in and coming out, we can *almost* solve for the π_i's. The missing component is the equation $\sum_i \pi_i = 1$ (the probabilities have to sum to one).

Big Theorem

If the Markov chain is finite and irreducible, then it has a unique invariant distribution. $\pi(i)$ is the long-term fraction of time that the chain spends in state i, i.e. that X_n is equal to i.

The long-term fraction of time that X_n is equal to i is

$$\lim_{n \to \infty} \left[\frac{1}{N} \sum_{i=0}^{N-1} \mathbb{1}\{X_n = i\} \right]$$

Intuitively, this is the fraction of time we're spending in a certain state (simply a counter over all time steps).

Furthermore, if the Markov chain is irreducible and aperiodic, then the distribution π_n converges to π as $n \to \infty$.

References

[1] D.P. Bertsekas and J.N. Tsitsiklis. *Introduction to Probability.* Athena Scientific books. Athena Scientific, 2002.

[2] Jean Walrand. *Probability in Electrical Engineering and Computer Science: An Application-Driven Course.* Quorum Books, Westport, CT, USA, 2014.