Math 3332: Probability[1]	Mikhail Lavrov

Lecture notes on Markov chains

Fall 2024	*Kennesaw State University*

1 Introduction to Markov chains

1.1 The weather as a random process

A **random process** is a special kind of random experiment: it's a random experiment in which our outcome is a sequence of results. Usually, we think of this sequence as telling a story of how a system behaves over time. For example, when we imagine flipping a coin over and over again (something we've thought about many times this semester), that's a random process: the outcome is a sequence of heads and tails.

To take another example, we can model the weather as a random process: a sequence of results telling us each day's weather. Some people spend their whole lives modeling the weather, so it's not a problem we'll solve in one lecture. To tell a very simplified story, suppose that the weather every day can be in one of two states: "Sunny" or "Rainy".

The simplest random process we could use to describe how the weather changes from Sunny to Rainy would be a process not too different from flipping a coin. For example, maybe we look up some statistics and decide that it rains 70% of the time in Atlanta.[2] The simplest model that would have this behavior is one where every day's weather is chosen independently at random: Sunny with probability 0.3 and Rainy with probability 0.7.

However, you don't have to be very observant to notice that this is not how weather works. There are trends to weather: tomorrow's weather is more likely to be like today's. A slightly more complicated model of the weather would take each day's weather into account to give probabilities for the next day's weather.

- If the weather on day t is Sunny, then the weather on day $t+1$ is Sunny with probability 0.6 and Rainy with probability 0.4.
- If the weather on day t is Rainy, then the weather on day $t+1$ is Sunny with probability 0.2 and Rainy with probability 0.8.

We can visualize this model with a diagram:

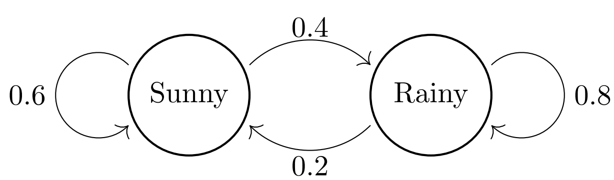

[1]This document comes from the Math 3332 course webpage: https://facultyweb.kennesaw.edu/mlavrov/courses/3332-fall-2024.php

[2]A completely made-up number.

The circles represent the two states the weather can be in: Sunny and Rainy. The arrows between them represent probabilities. For example, the 0.4 on the arrow from Sunny to Rainy says that from the Sunny state, there is a 0.4 probability that the next state will be Rainy.

(Earlier this semester, we solved many problems using branching diagrams. You can think of this diagram as an unusual kind of branching diagram: one that loops back on itself instead of stopping.)

Before we proceed to the general definition of a Markov chain, let's play around with this model for a bit, and see what it can do. For example, suppose that it's Sunny on Monday. What is the probability that it's Rainy on Friday, according to this model?

A very bad way to solve this problem would be to break up the problem into 8 cases for the 8 ways that we could get from Sunny on Monday to Rainy on Friday. For example, one such case is the sequence

$$\text{Sunny} \to \text{Rainy} \to \text{Sunny} \to \text{Sunny} \to \text{Rainy}.$$

The probability of this sequence is the probability that we'll follow all four arrows in the sequence. We follow the initial Sunny-to-Rainy arrow with probability 0.4; we then go back to Sunny with probability 0.2; we then stay in Sunny with probability 0.6; we then go to Rainy with probability 0.4. So the probability of this sequence is $0.4 \cdot 0.2 \cdot 0.6 \cdot 0.4 = 0.0192$. One down, seven to go!

The reason this is not the best approach is that it separates cases that can be combined. You can think of our calculation as adding up the probabilities of eight paths in the branching diagram below. (The calculation we just did is highlighted.)

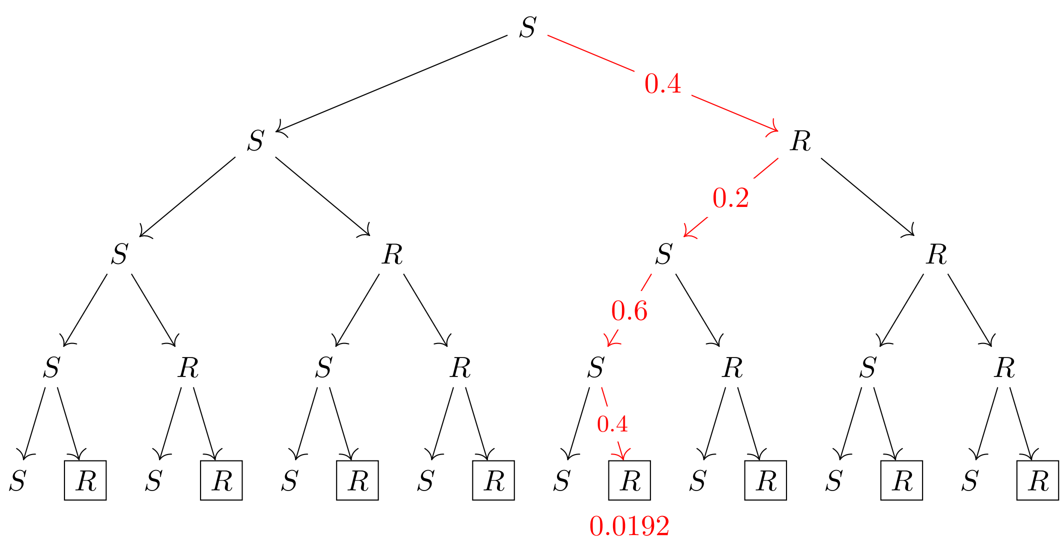

This is wasteful, because our model has the property that once we arrive at a Rainy day on Wednesday, it doesn't matter what the weather was like on Tuesday. If we compute eight separate probabilities, we'll end up duplicating many of our intermediate calculations!

Instead, let's build a table of the probabilities of Sunny and Rainy on each day of the week from Monday to Friday. Initially, for Monday, the probabilities are 1 and 0:

	Monday	Tuesday	Wednesday	Thursday	Friday
Sunny	1				
Rainy	0				

Since we know it's Sunny on Monday, we know how to fill in the Tuesday column: our rules say that the probability of a second Sunny day is 0.6 and the probability of a Rainy day is 0.4.

	Monday	Tuesday	Wednesday	Thursday	Friday
Sunny	1	0.6			
Rainy	0	0.4			

Next, the interesting part happens. What exactly are the numbers on our arrows? They are conditional probabilities: if we write Sunny_2 for the event of a Sunny day on Tuesday, and Sunny_3 for the event of a Sunny day on Wednesday, then 0.6 is the conditional probability $\Pr[\text{Sunny}_3 \mid \text{Sunny}_2]$. By the law of total probability, we have:

$$\begin{aligned}
\Pr[\text{Sunny}_3] &= \Pr[\text{Sunny}_3 \mid \text{Sunny}_2] \cdot \Pr[\text{Sunny}_2] + \Pr[\text{Sunny}_3 \mid \text{Rainy}_2] \cdot \Pr[\text{Rainy}_2] \\
&= 0.6 \cdot \Pr[\text{Sunny}_2] + 0.2 \cdot \Pr[\text{Rainy}_2]. \\
\Pr[\text{Rainy}_3] &= \Pr[\text{Rainy}_3 \mid \text{Sunny}_2] \cdot \Pr[\text{Sunny}_2] + \Pr[\text{Rainy}_3 \mid \text{Rainy}_2] \cdot \Pr[\text{Rainy}_2] \\
&= 0.4 \cdot \Pr[\text{Sunny}_2] + 0.8 \cdot \Pr[\text{Rainy}_2].
\end{aligned}$$

Since we already know $\Pr[\text{Sunny}_2] = 0.6$ and $\Pr[\text{Rainy}_2] = 0.4$, we can compute the next column: $\Pr[\text{Sunny}_3] = 0.6 \cdot 0.6 + 0.4 \cdot 0.2 = 0.44$ and $\Pr[\text{Rainy}_3] = 0.6 \cdot 0.4 + 0.4 \cdot 0.8 = 0.56$.

(We could also have said that $\Pr[\text{Rainy}_3] = 1 - \Pr[\text{Sunny}_3]$, as a shortcut.)

	Monday	Tuesday	Wednesday	Thursday	Friday
Sunny	1	0.6	0.44		
Rainy	0	0.4	0.56		

Now we're in business! The same rule that got us Wednesday's weather from Tuesday's weather can be repeated to predict Thursday's weather from Wednesday's weather. Just adding one to all the indices, we have

$$\begin{aligned}
\Pr[\text{Sunny}_4] &= 0.6 \cdot \Pr[\text{Sunny}_3] + 0.2 \cdot \Pr[\text{Rainy}_3]. \\
\Pr[\text{Rainy}_4] &= 0.4 \cdot \Pr[\text{Sunny}_3] + 0.8 \cdot \Pr[\text{Rainy}_3].
\end{aligned}$$

It's maybe a bit easier to use a calculator at this point, but either way we do it, we'll get $\Pr[\text{Sunny}_4] = 0.6 \cdot 0.44 + 0.2 \cdot 0.56 = 0.376$ and $\Pr[\text{Rainy}_4] = 0.624$. In the same way, we can compute $\Pr[\text{Sunny}_5] = 0.3504$ and $\Pr[\text{Rainy}_5] = 0.6496$, completing our table:

	Monday	Tuesday	Wednesday	Thursday	Friday
Sunny	1	0.6	0.44	0.376	0.3504
Rainy	0	0.4	0.56	0.624	0.6496

We conclude that there is a 64.96% chance of rain on Friday!

1.2 A quick linear algebra aside

I am not assuming that you've seen linear algebra before this class. However, if you have, then there's an important connection to notice here: our calculations can be phrased in terms of probability vectors. Write $\mathbf{x}_2$ for the column vector $(\Pr[\text{Sunny}_2], \Pr[\text{Rainy}_2])$ which summarizes Tuesday's weather, and write $\mathbf{x}_3$ for the column vector $(\Pr[\text{Sunny}_3], \Pr[\text{Rainy}_3])$ which summarizes Wednesday's weather. Then our rule for predicting Wednesday's weather based on Tuesday's weather is a matrix multiplication:

$$\mathbf{x}_3 = \begin{bmatrix} 0.6 & 0.2 \\ 0.4 & 0.8 \end{bmatrix} \mathbf{x}_2.$$

(If you're familiar with linear algebra, you should take a moment to think this over and see if you believe it!)

The calculation we did just now can be reduced to computing the fourth power of this matrix—a problem for which linear algebra has many powerful tools.

1.3 The definition of a Markov chain

This model of the weather is an example of a **Markov chain**.[3] What is a Markov chain in general? We can give two definitions: a concrete one, and an abstract one.

In the concrete definition, a Markov chain is a random process is specified by two pieces of information:

- A **state space** S: the set of possible states the random process can be in at a given time. In our weather example, the state space is $S = \{\text{Sunny}, \text{Rainy}\}$. For simplicity, we will assume S is a finite set, though Markov chains with infinitely large state spaces are also studied.
- Numbers called **transition probabilities**. For every state $a \in S$ and every state $b \in S$ (including even the case $a = b$), there is a probability $p_{a \to b}$. These are the numbers we labeled the arrows in our diagram with: for example, we set $p_{\text{Sunny} \to \text{Rainy}}$ to 0.4.

  The transition probabilities must all lie in the interval $[0, 1]$, and for every state $a \in S$, the sum $\sum_{b \in S} p_{a \to b}$ must equal 1.

Using this information, and an initial state $\mathbf{X}_0 \in S$, we generate a Markov chain $\mathbf{X}_0, \mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots$ one state at a time; $\mathbf{X}_{t+1}$ is chosen with probabilities that depend on $\mathbf{X}_t$. More precisely, if $\mathbf{X}_t = a$, then for each $b \in S$, we set $\mathbf{X}_{t+1}$ to b with probability $p_{a \to b}$.

The abstract definition of a Markov chain is what you will normally see in textbooks. Instead of telling you how to randomly generate the sequence $\mathbf{X}_0, \mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots$, the abstract definition just tells you the properties that their distribution must satisfy:

- The most important one is the **Markov property**. This says that for any time t, and any sequence $a_0, a_1, \dots, a_t, a_{t+1} \in S$, the conditional probability

$$\Pr[\mathbf{X}_{t+1} = a_{t+1} \mid \mathbf{X}_0 = a_0, \mathbf{X}_1 = a_1, \dots, \mathbf{X}_t = a_t]$$

[3] More precisely, a discrete-time Markov chain. We can also consider continuous-time Markov chains, where the weather is able to change from second to second, but we don't know enough about continuous random variables right now to discuss those.

is equal to the simpler conditional probability $\Pr[\mathbf{X}_{t+1} = a_{t+1} \mid \mathbf{X}_t = a_t]$. If we know $\mathbf{X}_t$, then any additional information about $\mathbf{X}_0, \mathbf{X}_1, \dots, \mathbf{X}_{t-1}$ is irrelevant for predicting $\mathbf{X}_{t+1}$.

- We will also assume **time homogeneity** (though not everyone who thinks about Markov chains assumes this property). Time homogeneity is the property that for every pair of states $a, b \in S$, the conditional probability $\Pr[\mathbf{X}_{t+1} = b \mid \mathbf{X}_t = a]$ is the same at every time t.

Time homogeneity lets us define $p_{a\to b}$ to be the conditional probability $\Pr[\mathbf{X}_{t+1} = b \mid \mathbf{X}_t = a]$ and arrive back at the concrete definition. Meanwhile, the Markov property is what tells us that $p_{a\to b}$ is *all* we need to generate the Markov chain: it tells us that the value of $\mathbf{X}_t$ is all we need to know to determine the distribution of $\mathbf{X}_{t+1}$.

In the weather example, what do these assumptions mean?

- The Markov property suggests that weather "doesn't have a long memory": no matter how many days it's been raining, for example, the probability that it clears up the next day will be the same. This is a simplifying assumption to reduce the number of variables in our model—and it turns out to be a very useful assumption, because Markov chains have many good properties.

- Time homogeneity suggests that weather doesn't have long-term trends: all Sunny days are alike, and so are all Rainy days.

  Of course, this is false in reality: real weather has seasonal effects! A more complicated model might drop the assumption of time homogeneity, and have a transition probability from Sunny to Rainy that changes over the course of the year. Our simple model will probably be good enough, though, as long as we only care about short-term trends in the weather.

2 Stationary distributions

2.1 Long-term weather trends

Let's have another look at our probability table for our simple weather model. I've extended it by one more day since the last time we've seen it:

	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Sunny	1	0.6	0.44	0.376	0.3504	0.34016
Rainy	0	0.4	0.56	0.624	0.6496	0.65984

It turns out, though, that adding a column for Saturday didn't make much of a difference; our prediction of Saturday's weather is a lot like our prediction of Friday's weather, with a difference of less than a percentage point.

Can you make a guess about what the numbers will be if we extend this table 100 more days into the future? It turns out that these numbers are converging to a $\frac{1}{3} = 0.33333\dots$ probability of "Sunny", and a $\frac{2}{3} = 0.66666\dots$ probability of "Rainy". These probabilities are called a **limiting distribution** for the Markov chain.

All this is maybe not too surprising, for a weather prediction. If we're only predicting the weather one or two days in advance, then today's weather should be a very useful piece of information.

But if we're trying to make long-term predictions, then today's weather becomes less relevant: the weather could change many times between Monday and Saturday! Our prediction in that case should be mostly based on the overall chance of Rainy weather in this climate, and only slightly affected by the initial conditions.

It's interesting, though, that this is happening for the Markov chain model. Is it guaranteed to happen for all Markov chains?

It turns out that no—there are two situations that give us trouble.

2.2 Periodic or aperiodic?

Here is a weather model with no element of chance in it that can still, technically, be described as a Markov chain. Suppose we just say: every day, the weather is the *opposite* of what it was the previous day. This can be described with the following diagram:

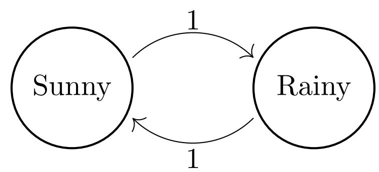

In this model, it is very easy to predict the future: n days from now, the weather will be the same as today's if n is even, and the opposite of today's if n is odd. However, no matter how large n gets, this prediction will still have to depend on what today's weather was like!

We can make this situation more complicated while still maintaining the same behavior, in two ways:

- We could add an element of chance back in by splitting each state (Sunny and Rainy) into several different states, for different ways the weather can be Sunny and different ways it can be Rainy.

  Now we can have a long-term prediction along the lines of "When it's Sunny, it's Sunny_1 with probability tending to 0.4 and Sunny_2 with probability tending to 0.6". However, this won't change the basic Sunny-vs.-Rainy prediction that depends on whether we're looking at a date an even or odd number of days away in the future.

- We could make the cycle longer. Imagine a model in which the weather always cycles Sunny to Cold to Rainy to Stormy to Cloudy to Sunny. Now the prediction for n days ahead no longer depends on whether n is even or odd: instead, it depends on the remainder when n is divided by 5. However, it's still a cyclic pattern.

In general, all fancier examples of this type will have a "period": an integer $d > 1$ for which we can separate the states of our Markov chain into d classes. At every time step, we advance to some random state in the next class (or from the d^{th} class back to the first). Therefore, no matter how far out you go, it still matters which class you started in: $10000d$ steps later, for example, you are guaranteed to be in the same class as the starting class.

Markov chains without such a structure are called **aperiodic**.

2.3 Reducible or irreducible?

Here is a second silly model of the weather:

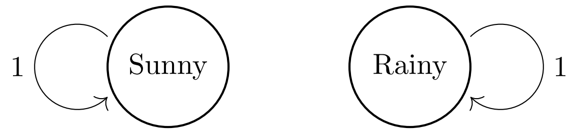

In this model, if the weather happens to be Rainy, it is doomed to be Rainy forever, and if the weather is Sunny today, it will be Sunny every day. (You can imagine that Atlanta is stuck in a Groundhog Day-style time loop, if you like: a single day repeats over and over, and whatever the weather happened to be on that day is what the weather will be like every day of the time loop.) No matter how many days ahead we look, the weather will still depend about as strongly as it possibly could on the weather we had on the first day.

There are more complicated pictures of this type. You could have two (or more) more complicated components with their own long-term trends—but the overall behavior of the Markov chain would still depend on the state you started in.

We say that a Markov chain is **irreducible** if it cannot be split up in this way: if it is possible to get from any state to any other state by following the arrows in the diagram. (This property also excludes Markov chains with a **transient set** of states: a subset of the states with arrows leaving it, but no states coming back.

2.4 The ergodic theorem

Let's state all of this more formally. Given a Markov chain $\mathbf{X}_0, \mathbf{X}_1, \mathbf{X}_2, \dots$ with state space S, the **limiting probability** π_a of a state $a \in S$ is the value

$$\lim_{t \to \infty} \Pr[\mathbf{X}_t = a]$$

(if this limit exists). Together, the limiting probabilities π_a for all $a \in S$ (if they all exist) are called the **limiting distribution** of the Markov chain.[4]

Fact 1. Suppose that the Markov chain is **aperiodic**: there is no integer $d > 1$ such that every path that starts and ends in the same state has a length divisible by d. Under this assumption, the limiting probability π_a exists for all states $a \in S$.

Fact 2. Suppose that, in addition to being aperiodic, the Markov chain is **irreducible**: for every two states $a, b \in S$, there is a path (a sequence of transitions with positive probabilities) that starts at a and ends at b. Then there is *only one* limiting distribution whose values only depend on the transition probabilities of the Markov chain, and not on the way that the initial state is chosen.

Together, these two facts are known as the **ergodic theorem**.

We will not prove this theorem in class. However, we will look at a question which is much more important from a practical point of view: how can we find the limiting distribution?

[4] In this course, we will only think about Markov chains with a finite state space. If S is an infinite set, then it's possible to have limiting probabilities that don't form a coherent limiting distribution together, and the story is more complicated.

2.5 Finding a stationary distribution

Let's return to the example Markov chain from earlier:

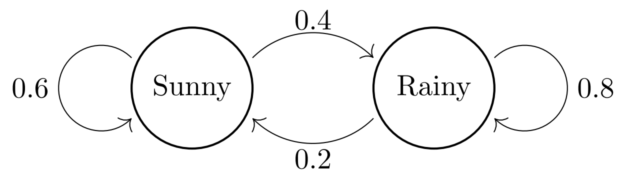

This is aperiodic (the loop transitions Sunny→Sunny and Rainy→Rainy are enough to guarantee this) and irreducible (we can get from Sunny to Rainy and from Rainy to Sunny), so a limiting distribution $(\pi_{\text{Sunny}}, \pi_{\text{Rainy}})$ must exist. How can we find it.

By building a table of probabilities for many days, we can make a guess about the limiting distribution by guessing what the numbers get closer and closer to. However, that's a lot of work. Instead, we can get the same result by solving a system of equations.

Consider: if we know the pair $(\Pr[\mathbf{X}_{101} = \text{Sunny}], \Pr[\mathbf{X}_{101} = \text{Rainy}])$ and we use the law of total probability to solve for the pair $(\Pr[\mathbf{X}_{100} = \text{Sunny}], \Pr[\mathbf{X}_{100} = \text{Rainy}])$, we get the following equations:

$$\Pr[\mathbf{X}_{101} = \text{Sunny}] = 0.6 \cdot \Pr[\mathbf{X}_{100} = \text{Sunny}] + 0.2 \cdot \Pr[\mathbf{X}_{100} = \text{Rainy}]$$
$$\Pr[\mathbf{X}_{101} = \text{Rainy}] = 0.4 \cdot \Pr[\mathbf{X}_{100} = \text{Sunny}] + 0.8 \cdot \Pr[\mathbf{X}_{100} = \text{Rainy}]$$

However, we also don't expect the probabilities to change very much from $t = 100$ to $t = 101$. We expect both $\Pr[\mathbf{X}_{100} = \text{Sunny}]$ and $\Pr[\mathbf{X}_{101} = \text{Sunny}]$ to be approximately equal to π_{Sunny}, and we expect both $\Pr[\mathbf{X}_{100} = \text{Rainy}]$ and $\Pr[\mathbf{X}_{101} = \text{Rainy}]$ to be approximately equal to π_{Rainy}. This suggests that if we started with probabilities exactly equal to the limiting probabilities, they would not change at all: we'd get

$$\pi_{\text{Sunny}} = 0.6 \cdot \pi_{\text{Sunny}} + 0.2 \cdot \pi_{\text{Rainy}}$$
$$\pi_{\text{Rainy}} = 0.4 \cdot \pi_{\text{Sunny}} + 0.8 \cdot \pi_{\text{Rainy}}$$

These are not enough to solve for π_{Sunny} and π_{Rainy}, because these are not independent equations: actually, either one of them implies the other. However, there is a third equation we can write down:

$$\pi_{\text{Sunny}} + \pi_{\text{Rainy}} = 1.$$

(This is true for the limiting probabilities because it is true at every time step.)

Using all three equations, we can find the unique solution: $\pi_{\text{Sunny}} = \frac{1}{3}$ and $\pi_{\text{Rainy}} = \frac{2}{3}$. So this must be the limiting distribution!

In general, we can set up equations like these to identify a probability distribution over the states of a Markov chain that is left unchanged as we apply its transition rule: if $\mathbf{X}_t$ were to have that distribution for some t, then $\mathbf{X}_{t+1}$ would also have that distribution. To be precise, in general, we get an equation

$$\pi_a = \sum_{b \in S} p_{b \to a} \cdot \pi_b$$

for all $a \in S$, as well as the equation

$$\sum_{a \in S} \pi_a = 1.$$

The equations are called the **stationary equations**, and a solution to them is called a **stationary distribution**, because if we choose $\mathbf{X}_0$ according to that distribution, then $\mathbf{X}_1$, $\mathbf{X}_2$, and so on will all have the same probabilities.

It is not quite true that a limiting distribution and a stationary distribution are the same: a stationary distribution can exist even if no limiting distribution exists. However, if a Markov chain is irreducible and aperiodic, then the stationary equations will only have one solution, and that solution will also be the limiting distribution!

3 Hitting times

Suppose that it's raining in Atlanta, and that weather in Atlanta follows the Markov chain we've been discussing. How long do you have to wait until a sunny day?

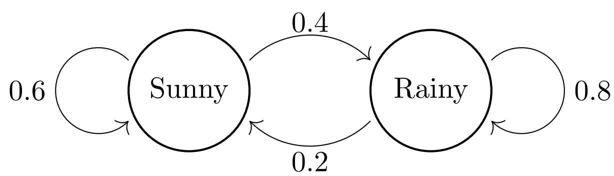

In this scenario, we can identify the problem as one we've already studied in this class. We are waiting for the first transition from Rainy to Sunny, and every day that it's Rainy, the probability of such a transition is 0.2. Therefore the number of days we have to wait until a Sunny day has a Geometric($p = 0.2$) distribution.

In general, we define the **hitting time** $\mathbf{H}(a, b)$ as follows. Begin by initializing the Markov chain so that $\mathbf{X}_0 = a$. Then $\mathbf{H}(a, b)$ is the random variable equal to the least value of t such that $\mathbf{X}_t = b$.

We have $\mathbf{H}(\text{Rainy}, \text{Sunny}) \sim \text{Geometric}(p = 0.2)$, but the familiar Geometric distribution only occurs here because our Markov chain is relatively simple. In a Markov chain with more states and more possible transitions, the distribution also becomes more complicated, and usually won't be one of the ones we have a name for!

3.1 A more complicated example

Suppose that you are flipping a fair coin over and over, and you want to count the number of flips it takes for the coin to land Heads twice in a row. How can we model this as a Markov chain?

What do we need to track, in order to tell that the coin has landed Heads twice in a row? We need to keep track of the length of our "streak" of Heads: the number of Heads that the sequence of coinflips ends with. The streak starts out at 0, and is equal to 0 whenever the last coinflip is Tails. If the streak is 0, and the coin lands Heads, the streak changes to 1. If the streak is 1, and the coin lands Tails, we're back to 0; but if it's heads *again*, the streak changes to 2 and we win!

If we wanted to accurately keep track of the streak length, we'd need infinitely many states, because we can get arbitrarily many results of Heads in a row, given enough time. But for the purposes of solving our problem, it's enough to stop at 2: we'll adopt the convention that "2" means "any streak of two or more consecutive Heads". Now our Markov chain only has three states: $S = \{0, 1, 2\}$. Here is a diagram showing the possible transitions between them:

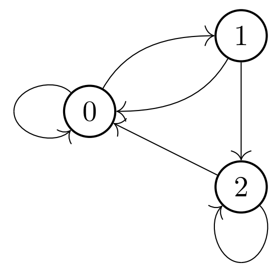

(The diagram does not include transition probabilities, but you should imagine that every arrow is labeled with a probability of $\frac{1}{2}$.)

The number of coinflips until two consecutive Heads is the hitting time $\mathbf{H}(0, 2)$.

Of course, this doesn't tell us the answer to the problem: it just rephrases it as a different problem we don't know how to solve! But now we can use this example as a test case for the techniques we figure out.

3.2 Expected hitting time

When looking at events in Markov chains, we used the law of total probability; now that we are looking at random variables, we can look at the law of total expectation. If we're trying to compute $\mathbf{H}(0, 2)$, then we consider the Markov chain with $\mathbf{X}_0 = 0$; if so, $\mathbf{X}_1$ can be either 0 or 1, each with probability $\frac{1}{2}$. By the law of total probability,

$$\mathbb{E}[\mathbf{H}(0, 2)] = \frac{1}{2}\mathbb{E}[\mathbf{H}(0, 2) \mid \mathbf{X}_1 = 0] + \frac{1}{2}\mathbb{E}[\mathbf{H}(0, 2) \mid \mathbf{X}_1 = 1].$$

What is $\mathbb{E}[\mathbf{H}(0, 2) \mid \mathbf{X}_1 = 0]$? It is the hitting time from 0 to 2 if the first step we take goes from 0 back to 0. That first step is wasted, so this is just equal to $\mathbb{E}[\mathbf{H}(0, 2)] + 1$.

What about $\mathbb{E}[\mathbf{H}(0, 2) \mid \mathbf{X}_1 = 1]$? It is the hitting time from 0 to 2 if the first step we take goes from 0 to 1. In this case, we are left with the easier task of getting from 1 to 2, but we have spent one step getting there, so this conditional expectation is equal to $\mathbb{E}[\mathbf{H}(1, 2)] + 1$. Altogether, we get

$$\begin{aligned}\mathbb{E}[\mathbf{H}(0, 2)] &= \frac{1}{2}\Big(\mathbb{E}[\mathbf{H}(0, 2)] + 1\Big) + \frac{1}{2}\Big(\mathbb{E}[\mathbf{H}(1, 2)] + 1\Big) \\ &= 1 + \frac{1}{2}\mathbb{E}[\mathbf{H}(0, 2)] + \frac{1}{2}\mathbb{E}[\mathbf{H}(1, 2)].\end{aligned}$$

This is a special case of a more general equation, obtained by the same reasoning. In general, for states $a, b \in S$ with $a \neq b$, we have

$$\mathbb{E}[\mathbf{H}(a, b)] = 1 + \sum_{c \in S} p_{a \to c} \cdot \mathbb{E}[\mathbf{H}(c, b)].$$

(We exclude the case $a = b$ because $\mathbf{H}(b, b)$ is just always 0.)

We can apply the general equation one more time to understand $\mathbb{E}[\mathbf{H}(1, 2)]$, which showed up in our previous equation. It will tell us that

$$\begin{aligned}\mathbb{E}[\mathbf{H}(1,2)] &= 1 + \frac{1}{2}\mathbb{E}[\mathbf{H}(0,2)] + \frac{1}{2}\mathbb{E}[\mathbf{H}(2,2)] \\ &= 1 + \frac{1}{2}\mathbb{E}[\mathbf{H}(1,2)].\end{aligned}$$

Here, we use the fact that $\mathbf{H}(2, 2) = 0$.

All this doesn't let us solve for the expected hitting time directly, However, if we set $x = \mathbb{E}[\mathbf{H}(0, 2)]$ and $y = \mathbb{E}[\mathbf{H}(1, 2)]$, our two equations can be summarized as

$$\begin{aligned}x &= 1 + \frac{1}{2}x + \frac{1}{2}y \\ y &= 1 + \frac{1}{2}x\end{aligned}$$

which has the solution $x = 6$, $y = 4$. In other words: on average, we need to flip a fair coin 6 times to see two consecutive results of Heads. After the first time we see Heads, on average there are 4 more coin flips left: that's what $y = 4$ tells us.

3.3 Higher moments

A slightly more challenging problem is to compute the variance of a hitting time. However, this is also a good exercise that combines everything we've done with random variables.

We will use the formula $\mathrm{Var}[\mathbf{X}] = \mathbb{E}[\mathbf{X}^2] - \mathbb{E}[\mathbf{X}]^2$, and the hard part will be understanding the $\mathbb{E}[\mathbf{X}^2]$ term. Here, we can also use the law of total probability.

Recall that our Markov chain for working with $\mathbf{H}(a, b)$ starts with $\mathbf{X}_0 = a$. As before, we can condition on the value of $\mathbf{X}_1$:

$$\mathbb{E}[\mathbf{H}(a,b)^2] = \sum_{c \in S} p_{a \to c} \cdot \mathbb{E}[\mathbf{H}(a,b)^2 \mid \mathbf{X}_1 = c].$$

What is $\mathbb{E}[\mathbf{H}(a, b)^2 \mid \mathbf{X}_1 = c]$? If $\mathbf{X}_1 = c$, then $\mathbf{H}(a, b)$ has the distribution of $1 + \mathbf{H}(c, b)$. So we have

$$\mathbb{E}[\mathbf{H}(a,b)^2] = \sum_{c \in S} p_{a \to c} \cdot \mathbb{E}[(1 + \mathbf{H}(c,b)^2].$$

If we expand out $\mathbb{E}[(1 + \mathbf{H}(c, b)^2]$ as $1 + 2\mathbb{E}[\mathbf{H}(c, b)] + \mathbb{E}[\mathbf{H}(c, b)^2]$, then we get

$$\begin{aligned}\mathbb{E}[\mathbf{H}(a,b)^2] &= \sum_{c \in S} p_{a \to c} + 2\sum_{c \in S} p_{a \to c} \cdot \mathbb{E}[\mathbf{H}(c,b)] + \sum_{c \in S} p_{a \to c} \cdot \mathbb{E}[\mathbf{H}(c,b)^2] \\ &= 1 + 2\Big(\mathbb{E}[\mathbf{H}(a,b)] - 1\Big) + \sum_{c \in S} p_{a \to c} \cdot \mathbb{E}[\mathbf{H}(c,b)^2] \\ &= 2 \cdot \mathbb{E}[\mathbf{H}(a,b)] - 1 + \sum_{c \in S} p_{a \to c} \cdot \mathbb{E}[\mathbf{H}(c,b)^2]\end{aligned}$$

Let's try this out on an example to see how it works: the example in the previous section. Here, we have $\mathbb{E}[\mathbf{H}(0,2)] = 6$ and $\mathbb{E}[\mathbf{H}(1,2)] = 4$, so our system of equations is:

$$\mathbb{E}[\mathbf{H}(0,2)^2] = 2 \cdot 6 - 1 + \frac{1}{2}\mathbb{E}[\mathbf{H}(0,2)^2] + \frac{1}{2}\mathbb{E}[\mathbf{H}(1,2)^2]$$
$$\mathbb{E}[\mathbf{H}(1,2)^2] = 2 \cdot 4 - 1 + \frac{1}{2}\mathbb{E}[\mathbf{H}(0,2)^2] + \frac{1}{2}\mathbb{E}[\mathbf{H}(2,2)^2]$$

or (if we set $x = \mathbb{E}[\mathbf{H}(0,2)^2]$ and $y = \mathbb{E}[\mathbf{H}(1,2)^2]$),

$$x = 11 + \frac{1}{2}x + \frac{1}{2}y,$$
$$y = 7 + \frac{1}{2}x.$$

Solving, we get $x = 58$ and $y = 36$. Finally, the variance of the hitting time we wanted is $\text{Var}[\mathbf{H}(0,2)] = \mathbb{E}[\mathbf{H}(0,2)^2] - \mathbb{E}[\mathbf{H}(0,2)]^2 = 58 - 6^2 = 32$.

3.4 The geometric distribution

We know that a random variable $\mathbf{X}$ has the Geometric distribution with parameter p if the range of $\mathbf{X}$ is $\{1, 2, 3, \dots\}$ and $\Pr[\mathbf{X} = k] = p(1-p)^{k-1}$. We use this to model the number of trials until a success, when the probability of a success is p and the trials are independent.

A similar conditioning approach works for the Geometric distribution, because it it is nothing more than the hitting time in a Markov chain:

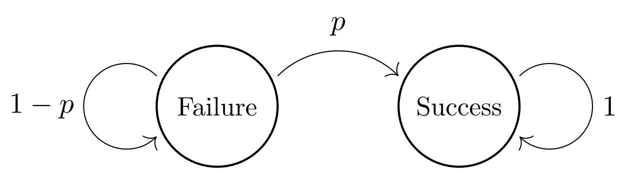

The value of $\mathbf{X}$ is the hitting time from the "Failure" state to the "Success" state. (The behavior at state "Success" doesn't really matter, since we don't need to know what happens after that state.)

Here, if we follow the $1-p$ arrow, the number of steps we have left to wait has the same distribution as $\mathbf{X}$. If we follow the p arrow, we are done. Therefore conditioning on the first step gives us

$$\mathbb{E}[\mathbf{X}] = 1 + p \cdot 0 + (1-p) \cdot \mathbb{E}[\mathbf{X}]$$

which we can solve to get $\mathbb{E}[\mathbf{X}] = \frac{1}{p}$. This should not be surprising; if, say, 1 in 100 trials are successful, then we expect that on average it should take 100 trials to see a success!

Similarly, we have

$$\begin{aligned}\mathbb{E}[\mathbf{X}^2] &= p \cdot 1^2 + (1-p) \cdot \mathbb{E}[(\mathbf{X}+1)^2] \\ &= (1-p)\mathbb{E}[\mathbf{X}^2] + 2(1-p)\mathbb{E}[\mathbf{X}] + 1 \\ &= (1-p)\mathbb{E}[\mathbf{X}^2] + \frac{2(1-p)}{p} + 1\end{aligned}$$

which we can solve to get $\mathbb{E}[\mathbf{X}^2] = \frac{2-p}{p^2}$ and $\text{Var}[\mathbf{X}] = \mathbb{E}[\mathbf{X}^2] - \mathbb{E}[\mathbf{X}]^2 = \frac{1-p}{p^2}$.