

Lecture 23 - Random Walks

6.042 - May 13, 2003

1 Random Walks


A drunkard stumbles out of a bar. Each second, he either staggers one step to the left or staggers one step to the right, with equal probability. His home lies x steps to his left, and a canal lies y steps to his right. This are several natural questions, including:

1. What is the probability that the drunkard arrives safely at home instead of falling into the canal?
2. What is the expected duration of his journey, however it ends?

The drunkard's meandering path is called a *random walk*. Random walks are an important subject, because they can model such a wide array of phenomena. For example, in physics, random walks in three-dimensional space are used to model gas diffusion. In computer science, the Google search engine uses random walks through the graph of web links to determine the relative importance of websites. In finance theory, random walks can serve as a model for the fluctuation of market prices. And, in this lecture, we'll explore some more palatable applications.

2 Pass the Candy

Pass the Candy is game involving n students, one professor, and a bowl of candy. The students are numbered 1 to n , and the professor is numbered 0. Everyone stands in a circle as shown below.



Initially, the professor has the candy bowl. He withdraws a piece of candy and then passes the bowl either left or right, with equal probability. Each person who receives the bowl thereafter does the same thing: he or she takes a piece of candy from the bowl and then passes it either left or right, with equal probability. In effect, the bowl goes for a random walk around the circle of players. The *last* person to receive a piece of candy is declared the winner and gets to keep the entire bowl. Which player is most likely to win?

A natural guess is player $n/2$. She seems most likely to receive the bowl last and thus to win the game, because she is farthest from the professor. On the other hand, players 1 and n seem almost certain to receive the bowl far too early in the process to win the game. Let's see if this intuition is right or wrong!

2.1 A Simpler Problem

Let's begin by looking at a simpler problem. Suppose that the players are arranged in a line, rather than a circle:

$$\begin{array}{ccccccc} A & S_1 & S_2 & \dots & S_k & B \\ & \uparrow & & & & \\ & \text{candy} & & & & \end{array}$$

The players are named A, S_1, S_2, \dots, S_k , and B , as shown. Initially, player S_1 has the candy bowl. As before, whenever a player gets the bowl, he or she takes a piece of candy and then passes the bowl either left or right, with equal probability. What is the probability that A gets candy before B ?

Let P_k be the probability that A gets candy before B . First, suppose that $k = 1$:

$$\begin{array}{ccc} A & S_1 & B \\ & \uparrow & \\ & \text{candy} & \end{array}$$

Here, either A or B gets the bowl on the next step, with equal probability. Thus, $P_1 = \frac{1}{2}$.

Now suppose that $k > 1$. In the first step, there are two possibilities: the bowl either moves left to player A or moves right to player S_2 . We can break up the evaluation of P_k into these two cases using the law of total probability:

$$\begin{aligned}
P_k &= \Pr \{ \text{first step is left} \} \\
&\quad \cdot \Pr \{ A \text{ gets candy before } B \mid \text{first step is left} \} \\
&+ \Pr \{ \text{first step is right} \} \\
&\quad \cdot \Pr \{ A \text{ gets candy before } B \mid \text{first step is right} \} \\
&= \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \Pr \{ A \text{ gets candy before } B \mid \text{first step is right} \}
\end{aligned}$$

To evaluate the last term, we must find the probability that A gets candy before B starting from this configuration:

$$\begin{array}{ccccccc}
A & S_1 & S_2 & \dots & S_k & B \\
& & \uparrow & & & \\
& & \text{candy} & & &
\end{array}$$

This is a little tricky. From here, the candy must eventually reach either S_1 or B . In fact, we know that S_1 gets the candy before B with probability P_{k-1} . (In effect, we are considering a smaller version of the problem, in which S_1 plays the role of A .) If this happens, then we are back in the original configuration, and so A goes on to get the candy before B with probability P_k . Therefore, starting from the configuration above, A gets the candy before B with probability $P_{k-1} \cdot P_k$. Substituting this result into the equation above gives:

$$P_k = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot P_{k-1} \cdot P_k$$

Solving for P_k and adding in the base case gives a complete recurrence:

$$\begin{aligned}
P_1 &= \frac{1}{2} \\
P_k &= \frac{1}{2 - P_{k-1}} \quad (k \geq 0)
\end{aligned}$$

2.2 Solving the Recurrence Equation

We now have a recurrence for P_k , the probability that player A gets the candy before player B . A recurrence is good, but an explicit formula for P_k would be better. The simplest technique for obtaining an explicit formula from a recurrence is called *guess and verify*. The name says it all: we compute a few terms, guess a general pattern, and then verify the result.

Let's apply the guess-and-verify method to our problem. First, we compute a few terms using the recurrence:

$$\begin{aligned} P_1 &= \frac{1}{2} \\ P_2 &= \frac{1}{2 - P_1} \\ &= \frac{2}{3} \\ P_3 &= \frac{1}{2 - P_2} \\ &= \frac{3}{4} \end{aligned}$$

The general pattern appears to be:

$$P_k = \frac{k}{k+1}$$

All that remains is to verify that our guess is correct. We can use induction on k with the hypothesis that $P_k = k/(k+1)$. This holds for $k = 1$, because $P_1 = \frac{1}{2}$ is the base case of the recurrence. Next, we assume $P_k = k/(k+1)$ in order to prove $P_{k+1} = (k+1)/(k+2)$. We can reason as follows:

$$\begin{aligned} P_{k+1} &= \frac{1}{2 - P_k} \\ &= \frac{1}{2 - \frac{k}{k+1}} \\ &= \frac{k+1}{k+2} \end{aligned}$$


The first step uses the recurrence equation, the second uses the induction hypothesis, and the third uses only algebra. Thus, by induction, $P_k = k/(k+1)$ for all $k \geq 1$.

Therefore, player A receives the candy before player B with probability $k/(k+1)$. We'll use this conclusion repeatedly in the next section.

2.3 Analyzing the Game

Now let's return to the original problem. There are n players and a professor standing in a circle. If the professor has the candy initially, which player is most likely to get the candy

last and thus win the game?



Let's begin by considering player n , who is standing just to the right of the professor. The only way that player n can win the game is if the bowl travels clockwise all the way around the circle to player $n - 1$ before n ever touches it. How likely is that? Let's get another perspective by "cutting" the circle between players $n - 1$ and n and arranging them in a line:

$$\begin{array}{cccccccc} n & 0 & 1 & 2 & \dots & n-2 & n-1 \\ & \uparrow & & & & & \\ & \text{candy} & & & & & \end{array}$$

We are asking for the probability that $n - 1$ gets candy before n . But this is precisely the problem that we already solved! Here player n takes the role of A , player $n - 1$ takes the role of B , and we have $k = n - 1$ players between them. Therefore, player n gets the candy before player $n - 1$ with probability $(n - 1)/n$. In complementary terms, player $n - 1$ gets the candy before player n with probability $1/n$. Thus, player n wins the game with probability $1/n$.

This is a startling conclusion! At the beginning of lecture, we guessed that player n had little chance of winning the game, because he was standing too close to the professor. But we've just shown that player n has a 1-in- n chance of winning. Not bad, since there are n players!

What are the winning chances for the other players? By symmetry, player 1 also wins with probability $1/n$. So let's consider the probability that player i wins the game, where $1 < i < n$. There are two cases: either player $i + 1$ gets the bowl before player $i - 1$ or vice versa. Applying the law of total probability, we get:

$$\begin{aligned}
\Pr \{i \text{ wins}\} &= \Pr \{i+1 \text{ gets candy before } i-1\} \\
&\quad \cdot \Pr \{i \text{ wins} \mid i+1 \text{ gets candy before } i-1\} \\
&+ \Pr \{i-1 \text{ gets candy before } i+1\} \\
&\quad \cdot \Pr \{i \text{ wins} \mid i-1 \text{ gets candy before } i+1\}
\end{aligned}$$

Let's begin by determining the probability that i wins, given that player $i+1$ gets candy before player $i-1$. (This is the second line in the formula.) At the moment that player $i+1$ first gets candy, this must be the configuration of the game:

$$\begin{array}{cccccccc}
i & i+1 & \dots & n & 0 & 1 & \dots & i-1 \\
& \uparrow & & & & & & \\
& \text{candy} & & & & & &
\end{array}$$

Now player i wins the game only if player $i-1$ gets the candy before player i . Again, we can invoke our earlier result. Now i takes the role of A , $i-1$ takes the role of B , and there are $k = n-1$ players between them. Therefore, player i gets the candy before player $i-1$ with probability $(n-1)/n$. This means that player i goes on to win the game with probability $1/n$, given that $i+1$ gets candy before $i-1$.

On the other hand, we can use the same argument to show that player i wins with probability $1/n$, given that player $i-1$ gets the candy before player $i+1$. Substituting these results into the equation above gives:

$$\begin{aligned}
\Pr \{i \text{ wins}\} &= \Pr \{i+1 \text{ gets candy before } i-1\} \cdot \frac{1}{n} \\
&\quad + \Pr \{i-1 \text{ gets candy before } i+1\} \cdot \frac{1}{n} \\
&= \frac{1}{n}
\end{aligned}$$

We do not know the probability that player $i+1$ gets candy before player $i-1$ or vice versa, but we do not need to; either way, player i wins with probability $1/n$. Therefore, player i wins with probability $1/n$ overall.

Amazingly, every player is equally likely to win Pass the Candy, regardless of where he or she stands!

3 Chocolate or Broccoli

In the *Chocolate or Broccoli* game, n players are immediately awarded chocolate bars and the remaining m players are awarded broccoli. But the game does not end there! We flip a coin.

If the coin is heads, then one player with chocolate must exchange her prize for broccoli. If the coin is tails, then a player with broccoli must exchange her prize for a chocolate bar. This process of coin-flipping and prize-exchanging continues until either all the players have chocolate or all the players have broccoli. At that point, the players can take their prizes home. How long does the game last?

Let the random variable $T_{c,b}$ be the length of the game if c players have chocolate and b players have broccoli. If either variable is zero, then the length of the game is also zero:

$$\begin{aligned} T_{c,0} &= 0 \\ T_{0,b} &= 0 \end{aligned}$$

Otherwise, we flip a coin and a prize is exchanged, expending one unit of time. There are now two possibilities:

1. With probability $\frac{1}{2}$, there are $c - 1$ chocolate bars and $b + 1$ broccolis.
2. With probability $\frac{1}{2}$, there are $c + 1$ chocolate bars and $b - 1$ broccolis.

The total duration of the game is 1 (to flip the coin and exchange one prize) plus the length of the rest of the game, which we can express using the law of total expectation:

$$\text{Ex}[T_{c,b}] = 1 + \frac{1}{2} \cdot \text{Ex}[T_{c-1,b+1}] + \frac{1}{2} \cdot \text{Ex}[T_{c+1,b-1}]$$

Now we have a recurrence equation for the expected duration of the game. But this is not a very satisfying answer; even with the recurrence in hand, there is still no obvious way to compute, say, $T_{3,1}$. We can still apply a form of guess and verify, however. In this case, we might need to use simulation results as the basis for a guess:

$$T_{c,b} = cb$$

We can then verify that this is a solution to the recurrence by plugging our guess into the right side and showing that we get the left side.

$$\begin{aligned} 1 + \frac{1}{2} \cdot \text{Ex}[T_{c-1,b+1}] + \frac{1}{2} \cdot \text{Ex}[T_{c+1,b-1}] &= 1 + \frac{1}{2} \cdot (c-1)(b+1) + \frac{1}{2} \cdot (c+1)(b-1) \\ &= cb \\ &= \text{Ex}[T_{c,b}] \end{aligned}$$

(There is a technical consideration that we are leaving aside: we have not shown that this solution is unique.)

3.1 An Surprising Implication

Suppose that you have \$1 and I have \$1,000,000. We repeatedly make fair \$1 bets. What is the expected number of bets until one or the other of us goes broke? Intuitively, it seems that the game should be rather short, since there is a half-chance that you go bankrupt after just one bet. However, this is equivalent to playing the Chocolate or Broccoli game where 1 player starts with chocolate and a million players start with broccoli. Therefore, the expected number of bets is actually a million!