

1 Kirby's Pet Spider

Every day when Kirby wakes up, he spends an hour observing his pet spider that he keeps in a cage. The spider can take on three discrete locations within the cage, denoted Left, Middle, and Right. The location of the spider is denoted by X_i for timestep i ($\forall i = 0, 1, \dots$). Every timestep, Kirby notes down which location that the spider is occupying, in relation to its location in the previous timestep.

After studying the spider for many weeks on end, he concludes that the spider's behavior can be modeled as a Markov Model following the initial and transition distributions described in the tables below:

X_0	$P(X_0)$	x	$P(X_{i+1} = x X_i = L)$	$P(X_{i+1} = x X_i = M)$	$P(X_{i+1} = x X_i = R)$
L	0.2	L	0.8	0.1	0
M	0.6	M	0.2	0.5	0.3
R	0.2	R	0	0.4	0.7

Note that each column of the transition matrix sums to 1.

Help Kirby answer the following questions about his spider friend for a new day.

(a) What is $P(X_1 = M|X_0 = L)$?

0.2. This is a quantity directly provided in our transition probabilities table.

(b) What is $P(X_1 = M)$?

$$\begin{aligned}
 P(X_1 = M) &= \sum_{x_0} P(X_1 = M, X_0 = x_0) \\
 &= \sum_{x_0} P(X_1 = M|X_0 = x_0)P(X_0 = x_0) \\
 &= P(X_1 = M|X_0 = L)P(X_0 = L) + P(X_1 = M|X_0 = M)P(X_0 = M) + P(X_1 = M|X_0 = R)P(X_0 = R) \\
 &= 0.2 * 0.2 + 0.5 * 0.6 + 0.3 * 0.2 \\
 &= \mathbf{0.4}
 \end{aligned}$$

One helpful idea for gaining intuition on this these calculations is the law of total probability.

(c) What is the stationary distribution of this Markov Model? Recall that the stationary distribution is a distribution over the possible states that remains the same over the passage of time: $P(X_i) = P(X_{i+1})$.

$$\begin{aligned}
 P(X_{i+1} = L) &= P(X_{i+1} = L|X_i = L)P(X_i = L) + P(X_{i+1} = L|X_i = M)P(X_i = M) + P(X_{i+1} = L|X_i = R)P(X_i = R) \\
 &= 0.8P(X_i = L) + 0.1P(X_i = M) + 0P(X_i = R)
 \end{aligned}$$

Once we derive similar equations for $P(X_{i+1} = M)$ and $P(X_{i+1} = R)$, set $P(X_{i+1}) = P(X_i) = P(X_\infty)$, and simplify, we come to the following system of equations:

$$\begin{aligned}
 P(X_\infty = L) &= 0.8P(X_\infty = L) + 0.1P(X_\infty = M) \\
 P(X_\infty = M) &= 0.2P(X_\infty = L) + 0.5P(X_\infty = M) + 0.3P(X_\infty = R) \\
 P(X_\infty = R) &= 0.4P(X_\infty = M) + 0.7P(X_\infty = R)
 \end{aligned}$$

We also have one more equation from the fact that any probability distribution must sum to 1: $P(X_\infty = L) + P(X_\infty = M) + P(X_\infty = R) = 1$

Solving this system gives us the following for our stationary distribution:

$$P(X_\infty = L) = \mathbf{3/17}$$

$$P(X_\infty = M) = \mathbf{6/17}$$

$$P(X_\infty = R) = \mathbf{8/17}$$

Kirby and his friend Waddle Dee want to turn this spider's movements into a bargaining game. Each game involves Kirby walking into the room at timestep $t = k$ and predicting X_{k+1} , the spider's location at timestep $t = k + 1$. If he predicts X_{k+1} correctly, he wins 15 apples from Waddle Dee. Consider the following scenarios and help advise Kirby in what decisions he should make.

- (d) Kirby walks into the room blindfolded, meaning he does not know the value of X_k . Waddle Dee offers to inform Kirby of the value of X_k as a hint in exchange for 5 apples. If Kirby rejects the hint, he randomly guesses X_{k+1} based on the stationary distribution derived in part c. Given that you know $X_k = L$, should Kirby accept Waddle Dee's new offer?

If Kirby rejects the hint, he samples his prediction from the stationary distribution, while the spider moves according to the transition matrix:

$$\begin{aligned} \text{Expected utility} &= 15 * P(\text{correct}) + 0 * P(\text{incorrect}) \\ &= 15 * (P(\text{predicts } L, X_{k+1} = L) + P(\text{predicts } M, X_{k+1} = M)) \\ &= 15 * (3/17 * 0.8 + 6/17 * 0.2) \\ &= \mathbf{54/17} \end{aligned}$$

If Kirby accepts the hint, he sees that the spider is most likely to stay Left, so he chooses that as his prediction:

$$\begin{aligned} \text{Expected utility} &= 10 * P(\text{correct}) + -5 * P(\text{incorrect}) \\ &= 10 * (P(\text{predicts } L, X_{k+1} = L) + P(\text{predicts } M, X_{k+1} = M)) \\ &\quad - 5 * (P(\text{predicts } L, X_{k+1} = M) + P(\text{predicts } M, X_{k+1} = L)) \\ &= 10 * (1 * 0.8 + 0 * 0.2) - 5 * (1 * 0.2 + 0 * 0.8) \\ &= 10 * 0.8 - 5 * 0.2 \\ &= \mathbf{7} \end{aligned}$$

(we assume that $P(X_{k+1} = R) = 0$)


Kirby **should take the hint**, since Kirby's expected utility would increase as a result. Note that the solution above relies on independence between the spider's transitions and Kirby's predictions.

- (e) Kirby takes off his blindfold, noting the value of X_k . Waddle Dee offers, again for the price of 5 apples, to tell Kirby the value of X_{k-1} as a hint to help with predicting X_{k+1} . Should Kirby accept Waddle Dee's new offer?

The Markov/memoryless property of a Markov Model tells us that the future is independent of the past conditioned on the present. Since Kirby knows the value of X_k already, X_{k-1} tells us no additional information about the distribution of X_{k+1} . Thus, **Kirby should not take the hint** in exchange for any positive number of apples.

2 HMMs

Consider the following Hidden Markov Model.



W_1	$P(W_1)$
0	0.3
1	0.7

W_t	W_{t+1}	$P(W_{t+1} W_t)$
0	0	0.4
0	1	0.6
1	0	0.8
1	1	0.2

W_t	O_t	$P(O_t W_t)$
0	a	0.9
0	b	0.1
1	a	0.5
1	b	0.5

Suppose that we observe $O_1 = a$ and $O_2 = b$.

Using the forward algorithm, compute the probability distribution $P(W_2|O_1 = a, O_2 = b)$ one step at a time.

(a) Compute $P(W_1, O_1 = a)$.

$$P(W_1, O_1 = a) = P(W_1)P(O_1 = a|W_1)$$

$$P(W_1 = 0, O_1 = a) = (0.3)(0.9) = 0.27$$

$$P(W_1 = 1, O_1 = a) = (0.7)(0.5) = 0.35$$

(b) Using the previous calculation, compute $P(W_2, O_1 = a)$.

$$P(W_2, O_1 = a) = \sum_{w_1} P(w_1, O_1 = a)P(W_2|w_1)$$

$$P(W_2 = 0, O_1 = a) = (0.27)(0.4) + (0.35)(0.8) = 0.388$$

$$P(W_2 = 1, O_1 = a) = (0.27)(0.6) + (0.35)(0.2) = 0.232$$

(c) Using the previous calculation, compute $P(W_2, O_1 = a, O_2 = b)$.

$$P(W_2, O_1 = a, O_2 = b) = P(W_2, O_1 = a)P(O_2 = b|W_2)$$

$$P(W_2 = 0, O_1 = a, O_2 = b) = (0.388)(0.1) = 0.0388$$

$$P(W_2 = 1, O_1 = a, O_2 = b) = (0.232)(0.5) = 0.116$$

(d) Finally, compute $P(W_2|O_1 = a, O_2 = b)$.

Renormalizing the distribution above, we have

$$P(W_2 = 0|O_1 = a, O_2 = b) = 0.0388/(0.0388 + 0.116) \approx 0.25$$

$$P(W_2 = 1|O_1 = a, O_2 = b) = 0.116/(0.0388 + 0.116) \approx 0.75$$