# CS 188 Summer 2023 — Discussion 4A

## 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:

| $X_0$ | $P(X_0)$ |
|---|---|
| L | 0.2 |
| M | 0.6 |
| R | 0.2 |

| $x$ | $P(X_{i+1} = x \mid X_i = L)$ | $P(X_{i+1} = x \mid X_i = M)$ | $P(X_{i+1} = x \mid X_i = R)$ |
|---|---|---|---|
| L | 0.8 | 0.1 | 0 |
| M | 0.2 | 0.5 | 0.3 |
| 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 $P(X_1 = M \mid X_0 = L)$?

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

**(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: $P(X_i) = P(X_{i+1})$.

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?

**(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?

## 2 HMMs

Consider the following Hidden Markov Model.

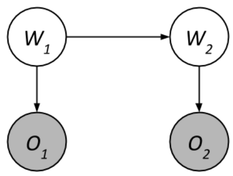


| $W_1$ | $P(W_1)$ |
|---|---|
| 0 | 0.3 |
| 1 | 0.7 |

| $W_t$ | $W_{t+1}$ | $P(W_{t+1}\|W_t)$ |
|---|---|---|
| 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 $P(W_1, O_1 = a)$.

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

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

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