CS221 Problem Session Solutions

Week 7

1) Problem 1: Markov Networks

This problem will give you some practice on computing probabilities given a Markov network. Specifically, given the Markov network below, we will ask you questions about the probability distribution $\mathbb{P}(X_1, X_2, X_3)$ over the **binary** random variables $X_1, X_2,$ and X_3.

X_1	X_2	f_3
0	0	1
0	1	1
1	0	2
1	1	1

X_2	X_3	f_4
0	0	1
0	1	1
1	0	2
1	1	1

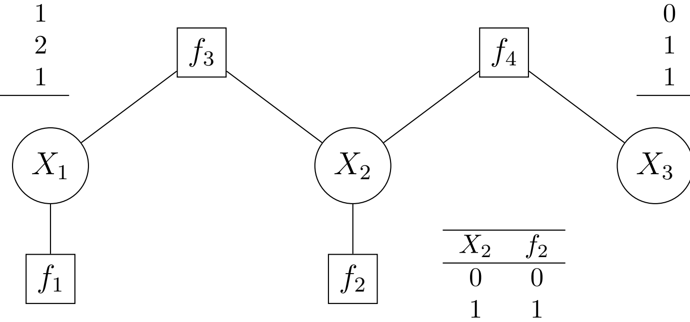

X_1	f_1
0	2
1	1

X_2	f_2
0	0
1	1

(a) What is the normalization constant Z (i.e. the total of all possible weights)?

Solution

$$
\begin{aligned}
Z &= \sum_{x \in X} \prod_{i=1}^{4} f_i(x) \\
&= \sum_{x_1 \in \{0,1\}} \sum_{x_2 \in \{0,1\}} \sum_{x_3 \in \{0,1\}} F_1(x_1) F_2(x_2) F_3(x_1, x_2) F_4(x_2, x_3) \\
&= [X = (0,1,0)]\ 4 + [X = (0,1,1)]\ 2 + [X = (1,1,0)]\ 2 + [X = (1,1,1)]\ 1 \\
&= 9
\end{aligned}
$$

(b) What is $\mathbb{P}(X_1 = 0, X_2 = 0, X_3 = 0)$?

Solution

$$\begin{aligned}\mathbb{P}(X_1=0,X_2=0,X_3=0) &= \frac{\text{Weight}(\mathbf{0})}{\sum_{x'\in\{0,1\}^3}\text{Weight}(x')}\\ &= \frac{f_1(0)f_2(0)f_3(0,0)f_4(0,0)}{Z}\\ &= \frac{2\cdot 0\cdot 1\cdot 1}{9}\\ &= 0\end{aligned}$$

(c) What is $\mathbb{P}(X_1=0,X_2=1,X_3=0)$?

Solution

$$\begin{aligned}\mathbb{P}(X_1=0,X_2=1,X_3=0) &= \frac{f_1(0)f_2(1)f_3(0,1)f_4(1,0)}{Z}\\ &= \frac{2\cdot 1\cdot 1\cdot 2}{9}\\ &= \frac{4}{9}\end{aligned}$$

(d) What is $\mathbb{P}(X_2=0)$?

Solution

$$\mathbb{P}(X_2) = \frac{1}{Z}\sum_{x_1,x_3\in\{0,1\}} f_1(x_1)f_2(0)f_3(x_1,0)f_4(0,x_3) = 0$$

Since $f_2(0)=0$.

(e) What is $\mathbb{P}(X_3=0)$?

Solution

$$\begin{aligned}\mathbb{P}(X_3=0) &= \frac{\sum_{x_1,x_2\in\{0,1\}}\mathbb{P}(X_1=x_1,X_2=x_2,X_3=0)}{Z}\\ &= \frac{4+2}{9}\\ &= \frac{2}{3}\end{aligned}$$

2) **Problem 2: The Bayesian Bag of Candies Model**

You have a lot of candy left over from Halloween, and you decide to give them away to your friends. You have four types of candy: **A**pple, **B**anana, **C**aramel, **D**ark-Chocolate. You decide to prepare candy bags using the following process.

- For each candy bag, you first flip a (biased) coin Y which comes up heads ($Y = \mathrm{H}$) with probability λ and tails ($Y = \mathrm{T}$) with probability $1 - \lambda$.
- If Y comes up heads ($Y = \mathrm{H}$), you make a **H**ealthy bag, where you:
  (a) Add one **A**pple candy with probability p_1 or nothing with probability $1 - p_1$;
  (b) Add one **B**anana candy with probability p_1 or nothing with probability $1 - p_1$;
  (c) Add one **C**aramel candy with probability $1 - p_1$ or nothing with probability p_1;
  (d) Add one **D**ark-Chocolate candy with probability $1 - p_1$ or nothing with probability p_1.
- If Y comes up tails ($Y = \mathrm{T}$), you make a **T**asty bag, where you:
  (a) Add one **A**pple candy with probability p_2 or nothing with probability $1 - p_2$;
  (b) Add one **B**anana candy with probability p_2 or nothing with probability $1 - p_2$;
  (c) Add one **C**aramel candy with probability $1 - p_2$ or nothing with probability p_2;
  (d) Add one **D**ark-Chocolate candy with probability $1 - p_2$ or nothing with probability p_2.

For example, if $p_1 = 1$ and $p_2 = 0$, you would deterministically generate: **H**ealthy bags with one **A**pple and one **B**anana; and **T**asty bags with one **C**aramel and one **D**ark-Chocolate. For general values of p_1 and p_2, bags can contain anywhere between 0 and 4 pieces of candy.

Denote A, B, C, D random variables indicating whether or not the bag contains candy of type **A**pple, **B**anana, **C**aramel, and **D**ark-Chocolate, respectively.

(a) Draw the Bayesian network corresponding to process of creating a single bag.

Solution

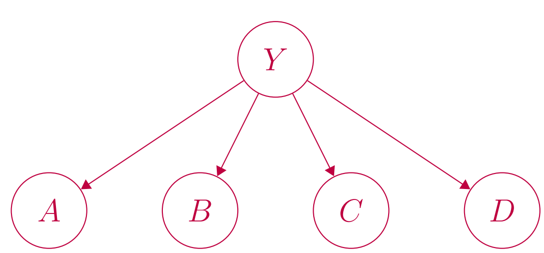

(b) What is the probability of generating a **H**ealthybag containing **A**pple, **B**anana, **C**aramel, and not **D**ark-Chocolate? For compactness, we will use the following notation to denote this possible outcome:

$$(\textbf{H}\text{ealthy}, \{\textbf{A}\text{pple}, \textbf{B}\text{anana}, \textbf{C}\text{aramel}\}).$$

Solution By definition, we create a **H**ealthy bag with probability λ, and include the candies with probability $p_1 p_1 (1 - p_1) p_1$, so the result is

$$\lambda p_1 p_1 (1 - p_1) p_1$$

(c) What is the probability of generating a bag containing **A**pple, **B**anana, **C**aramel, and *not* **D**ark-Chocolate?

Solution The bag could be **H**ealthy or **T**asty. We have computed the probability for the **H**ealthy case above. For a **T**asty one, a similar computation gives

$$(1 - \lambda) p_2 p_2 (1 - p_2) p_2$$

so the result is:

$$\lambda p_1 p_1 (1 - p_1) p_1 + (1 - \lambda) p_2 p_2 (1 - p_2) p_2$$

(d) What is the probability that a bag was a **T**asty one, given that it contains **A**pple, **B**anana, **C**aramel, and *not* **D**ark-Chocolate?

Solution Using the definition of conditional probability, we get:

$$\begin{aligned}
\mathbb{P}(T | A, B, C, \neg D) &= \frac{\mathbb{P}(A, B, C, \neg D | T)\mathbb{P}(T)}{\mathbb{P}(A, B, C, \neg D)} \\
&= \frac{p_2 p_2 (1 - p_2) p_1 (1 - \lambda)}{\lambda p_1 p_1 (1 - p_1) p_1 + (1 - \lambda) p_2 p_2 (1 - p_2) p_2}
\end{aligned}$$

Notice that we are using our result from (c) for the denominator. Take a moment to understand how these quantities connect.