CS 344 (Spring 2017): Class Test 4*

Instructor: Shivaram Kalyanakrishnan

8.30 a.m. – 9.30 a.m., April 11, 2017, 101/103 New CSE Building

Total marks: 15

Note. Provide brief justifications and/or calculations along with each answer to illustrate how you arrived at the answer.

Question 1. Consider a medical test for a virus. The test is 95% effective in recognising the virus when it is present, but has a 10% false positive rate (indicating that the virus is present, when it is not). The virus is carried by 1% of all people. If someone tests positive for the virus, what is the probability that they really carry the virus? [2 marks]

Question 2. X and Y are Boolean random variables. X takes values x and $\bar{x}$, while Y takes values y and $\bar{y}$. The joint probability distribution of X and Y is given below, but two entries—α and β—are left unspecified. Which values of (α, β) would make X and Y independent? If there are multiple such values, list all. [4 marks]

X	Y	$P(X, Y)$
x	y	0.04
x	$\bar{y}$	α
$\bar{x}$	y	β
$\bar{x}$	$\bar{y}$	0.54

Question 3. Consider a random variable A that can take any of exactly three possible values: a_1, a_2, and a_3. For $i \in \{1, 2, 3\}$, A takes value a_i with probability p_i (and hence $p_1 + p_2 + p_3 = 1$). In other words, the "real world" generating i.i.d. samples of A is the tuple (p_1, p_2, p_3).

A Bayesian inferencer, who does not know the real world, assumes a uniform prior over the set of possible worlds (x_1, x_2, x_3) wherein $x_1, x_2, x_3 \in [0, 1]$, subject to $x_1 + x_2 + x_3 = 1$. Thus, for all such (x_1, x_2, x_3), the probability density function $Belief_0(x_1, x_2, x_3) = C$ for some constant C.

The inferencer now observes a sample of A, which happens to be a_2. Assuming the inferencer updates his/her belief distribution in accordance with Bayes' rule, what is the belief distribution after this first sample a_2 has been observed? You must provide a simplified formula for $Belief_1(x_1, x_2, x_3)$, which gives the corresponding probability density function. [3 marks]

*Question 1 is based on an exercise in the textbook by Russell and Norvig (2010).

Question 4. Answer the following questions based on the Bayes Net shown below. Each of the variables is Boolean: X takes values x and $\bar{x}$; Y takes values y and $\bar{y}$; and z takes values z and $\bar{z}$.

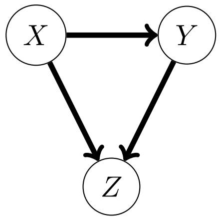

X	$P(X)$
x	0.3
$\bar{x}$	0.7

X	Y	$P(Y\|X)$
x	y	0.1
x	$\bar{y}$	0.9
$\bar{x}$	y	0.5
$\bar{x}$	$\bar{y}$	0.5

X	Y	Z	$P(Z\|X,Y)$
x	y	z	0.6
x	y	$\bar{z}$	0.4
x	$\bar{y}$	z	0.8
x	$\bar{y}$	$\bar{z}$	0.2
$\bar{x}$	y	z	0.2
$\bar{x}$	y	$\bar{z}$	0.8
$\bar{x}$	$\bar{y}$	z	0.1
$\bar{x}$	$\bar{y}$	$\bar{z}$	0.9

4a. What is $P(\bar{x}, y, \bar{z})$? [1 mark]

4b. Draw the table for $P(X, Z|y)$. [2 marks]

4c. Suppose we wish to estimate $P(X, Z|y)$ using Gibbs Sampling. We initialise the variables as: $X = x$, $Y = y$, $Z = z$. At each stage, our particular implementation of Gibbs Sampling picks X or Z uniformly at random, and then randomly sets the variable according to the appropriate conditional probability.

What is the probability that after *two* Gibbs Sampling steps, the variables have values $X = \bar{x}$, $Y = y$, $Z = \bar{z}$? In other words, if we take the initial (zeroth) state to be $s_0 = (x, y, z)$, what it the probability that the second state $s_2 = (\bar{x}, y, \bar{z})$? Provide a numeric answer by substituting all the relevant numbers. Simplify your answer up to the point that it is a sum of proper fractions. There is no need to simplify further. [3 marks]

Solutions

1. $P(\text{Have Virus}|\text{Test Positive}) \propto P(\text{Test Positive}|\text{Have Virus})P(\text{Have Virus}) = 0.95 \times 0.01$.

$P(\text{Not Have Virus}|\text{Test Positive}) \propto P(\text{Test Positive}|\text{Not Have Virus})P(\text{Not Have Virus}) = 0.1 \times 0.99$.

Normalising,

$$P(\text{Have Virus}|\text{Test Positive}) = \frac{0.95 \times 0.01}{0.95 \times 0.01 + 0.1 \times 0.99} = \frac{19}{217}.$$

2. First, we have $\alpha + \beta = 1 - 0.04 - 0.54 = 0.42$. We also have $P(x) = 0.04 + \alpha$, $P(\bar{x}) = 0.54 + \beta$, $P(y) = 0.04 + \beta$, and $P(\bar{y}) = 0.54 + \alpha$. If X and Y are independent, we must have

$$\begin{aligned}
P(x, y) &= P(x)P(y) = (0.04 + \alpha)(0.04 + \beta), \\
P(x, \bar{y}) &= P(x)P(\bar{y}) = (0.04 + \alpha)(0.54 + \alpha), \\
P(\bar{x}, y) &= P(\bar{x})P(y) = (0.54 + \beta)(0.04 + \beta), \text{ and} \\
P(\bar{x}, \bar{y}) &= P(\bar{x})P(\bar{y}) = (0.54 + \beta)(0.54 + \alpha).
\end{aligned}$$

The first equation yields

$$\begin{aligned}
0.04 &= (0.04 + \alpha)(0.04 + \beta) \\
&= 0.0016 + 0.04(\alpha + \beta) + \alpha\beta \\
&= 0.0016 + 0.04(0.42) + \alpha(0.42 - \alpha) \\
&= 0.0184 + 0.42\alpha - \alpha^2.
\end{aligned}$$

In other words, $\alpha^2 - 0.42\alpha + 0.0216 = 0$, or

$$\alpha = \frac{0.42 \pm \sqrt{0.1764 - 0.0864}}{2} = 0.21 \pm 0.15.$$

Thus, to satisfy $P(x, y) = P(x)P(y)$, we can only take $(\alpha, \beta) = (0.06, 0.36)$ or $(\alpha, \beta) = (0.36, 0.06)$. It is easy to verify that the other three conditions are also met by both $(\alpha, \beta) = (0.06, 0.36)$ and $(\alpha, \beta) = (0.36, 0.06)$.

3. Let $Belief_0(x_1, x_2, x_3) = C$. Then

$$Belief_1(x_1, x_2, x_3) \propto Belief_0(x_1, x_2, x_3)P(a_2|x_1, x_2, x_3) = Cx_2.$$

Normalising yields

$$\begin{aligned}
Belief_1(x_1, x_2, x_3) &= \frac{Cx_2}{\int_{y_1=0}^{1}\int_{y_2=0}^{1-y_1} Cy_2 dy_1 dy_2} \\
&= \frac{2x_2}{\int_{y_1=0}^{1}(1 - y_1)^2 dy_1} \\
&= \frac{2x_2}{\int_{y_1=0}^{1}(1 + y_1^2 - 2y_1) dy_1} \\
&= 6x_2.
\end{aligned}$$

4a. Unrolling the Bayes Net gives $P(\bar{x}, y, \bar{z}) = P(\bar{x})P(y|\bar{x})P(\bar{z}|\bar{x}, y) = 0.7 \times 0.5 \times 0.8 = 0.28$.

4b. It becomes easy to answer 4b and 4c once we have $P(X, Y, Z)$ (the fully enumerated joint probability distribution) in front of us. We obtain each row in this table in the same manner we obtained the answer to 4a.

X	Y	Z	$P(X,Y,Z)$
x	y	z	0.018
x	y	$\bar{z}$	0.012
x	$\bar{y}$	z	0.216
x	$\bar{y}$	$\bar{z}$	0.054
$\bar{x}$	y	z	0.07
$\bar{x}$	y	$\bar{z}$	0.28
$\bar{x}$	$\bar{y}$	z	0.035
$\bar{x}$	$\bar{y}$	$\bar{z}$	0.315

Since our query is conditioned on $Y = y$, we calculate $P(y) = 0.018 + 0.012 + 0.07 + 0.28 = 0.38$. We can now obtain the required table.

X	Z	$P(X, Z \mid y)$
x	z	$0.018/0.38 = 9/190$
x	$\bar{z}$	$0.012/0.38 = 3/95$
$\bar{x}$	z	$0.07/0.38 = 7/38$
$\bar{x}$	$\bar{z}$	$0.28/0.38 = 14/19$

4c Since only one of X and Z can get flipped at a time, exactly two routes are possible to reach $s_2 = (\bar{x}, y, \bar{z})$:

1. $s_0 = (x, y, z) \to s_1 = (\bar{x}, y, z) \to s_2 = (\bar{x}, y, \bar{z})$, and
2. $s_0 = (x, y, z) \to s_1 = (x, y, \bar{z}) \to s_2 = (\bar{x}, y, \bar{z})$.

Starting from s_0, the probability of the first route is

$$0.5 \times P(\bar{x}|y, z) \times 0.5 \times P(\bar{z}|\bar{x}, y) = \frac{1}{4} \times \frac{0.07}{0.07 + 0.018} \times \frac{0.28}{0.28 + 0.07} = \frac{1}{4} \times \frac{70}{88} \times \frac{4}{5} = \frac{7}{44}.$$

The probability of the second route is

$$0.5 \times P(\bar{z}|x, y) \times 0.5 \times P(\bar{x}|y, \bar{z}) = \frac{1}{4} \times \frac{0.012}{0.012 + 0.018} \times \frac{0.28}{0.28 + 0.012} = \frac{1}{4} \times \frac{2}{5} + \frac{70}{73} = \frac{7}{73}.$$

The required probability is

$$\frac{7}{44} + \frac{7}{73} = \frac{819}{3212}.$$