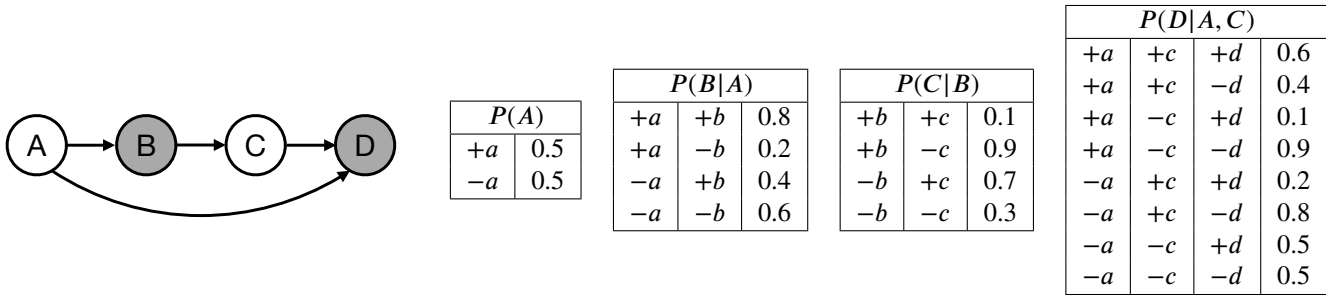


Q1. Bayes Nets: Sampling

Consider the following Bayes Net, where we have observed that $B = +b$ and $D = +d$.



- (a) Consider doing Gibbs sampling for this example. Assume that we have initialized all variables to the values $+a, +b, +c, +d$. We then unassign the variable C , such that we have $A = +a, B = +b, C = ?, D = +d$. Calculate the probabilities for new values of C at this stage of the Gibbs sampling procedure.

$$P(C = +c \text{ at the next step of Gibbs sampling}) = \frac{0.1 \cdot 0.6}{0.1 \cdot 0.6 + 0.9 \cdot 0.1} = \frac{2}{5}$$

$$P(C = -c \text{ at the next step of Gibbs sampling}) = \frac{0.9 \cdot 0.1}{0.1 \cdot 0.6 + 0.9 \cdot 0.1} = \frac{3}{5}$$

- (b) Consider a sampling scheme that is a hybrid of rejection sampling and likelihood-weighted sampling. Under this scheme, we first perform rejection sampling for the variables A and B . We then take the sampled values for A and B and extend the sample to include values for variables C and D , using likelihood-weighted sampling.

- (i) Below is a list of candidate samples. Mark the samples that would be rejected by the rejection sampling portion of the hybrid scheme.

- $-a \quad -b$
 $+a \quad +b$
 $+a \quad -b$
 $-a \quad +b$

- (ii) To decouple from part (i), you now receive a *new* set of samples shown below. Fill in the weights for these samples under our hybrid scheme.

	Weight
$-a \quad +b \quad -c \quad +d$	<u>0.5</u>
$+a \quad +b \quad -c \quad +d$	<u>0.1</u>
$+a \quad +b \quad -c \quad +d$	<u>0.1</u>
$-a \quad +b \quad +c \quad +d$	<u>0.2</u>
$+a \quad +b \quad +c \quad +d$	<u>0.6</u>

- (iii) Use the weighted samples from part (ii) to calculate an estimate for $P(+a | +b, +d)$.

The estimate of $P(+a|+b,+d)$ is $\frac{0.1 + 0.1 + 0.6}{0.5 + 0.1 + 0.1 + 0.2 + 0.6} = \frac{8}{15}$

(c) We now attempt to design an alternative hybrid sampling scheme that combines elements of likelihood-weighted and rejection sampling. For each proposed scheme, indicate whether it is valid, i.e. whether the weighted samples it produces correctly approximate the distribution $P(A, C | +b, +d)$.

(i) First collect a likelihood-weighted sample for the variables A and B . Then switch to rejection sampling for the variables C and D . In case of rejection, the values of A and B and the sample weight are **thrown away**. Sampling then restarts from node A .

Valid Invalid

(ii) First collect a likelihood-weighted sample for the variables A and B . Then switch to rejection sampling for the variables C and D . In case of rejection, the values of A and B and the sample weight are **retained**. Sampling then restarts from node C .

Valid Invalid

The sampling procedure in part (i) is the correct way of combining likelihood-weighted and rejection sampling: any time a node gets rejected, the sample must be thrown out in its entirety. In part (ii), however, the evidence that $D = +d$ has no effect on which values of A are sampled or on the sample weights. This means that values for A would be sampled according to $P(A | +b)$, not $P(A | +b, +d)$.

As an extreme case, suppose node D had a different probability table where $P(+d | +a) = 0$. Following the procedure from part (ii), we might start by sampling $(+a, +b)$ and assigning a weight according to $P(+b | +a)$. However, when we move on to rejection sampling we will be forced to continuously reject all possible values because our evidence $+d$ is inconsistent with our the assignment of $A = +a$. This means that the procedure from part (ii) is flawed to the extent that it might fail to generate a sample altogether!

Q2. VPI

You are the latest contestant on Monty Hall's game show, which has undergone a few changes over the years. In the game, there are n closed doors: behind one door is a car ($U(car) = 1000$), while the other $n - 1$ doors each have a goat behind them ($U(goat) = 10$). You are permitted to open exactly one door and claim the prize behind it.

You begin by choosing a door uniformly at random.

- (a) What is your expected utility?

$$(1000 * \frac{1}{n} + 10 * \frac{n-1}{n}) \text{ or } (10 + 990 * \frac{1}{n})$$

Answer:

We can calculate the expected utility via the usual formula of expectation, or we can note that there is a guaranteed utility of 10, with a small probability of a bonus utility. The latter is a bit simpler, so the answers to the following parts use this form.

- (b) After you choose a door but before you open it, Monty offers to open k other doors, each of which are guaranteed to have a goat behind it.

If you accept this offer, should you keep your original choice of a door, or switch to a new door?

$$10 + 990 * \frac{1}{n}$$

$EU(keep)$:

$$10 + 990 * \frac{(n-1)}{n*(n-k-1)}$$

$EU(switch)$:

switch

Action that achieves MEU :

The expected utility if we keep must be the same as the answer from the previous part: the probability that we have a winning door has not changed at all, since we have gotten no meaningful information.

In order to win a car by switching, we must have chosen a goat door previously (probability $\frac{n-1}{n}$) and then switch to the car door (probability $\frac{1}{n-k-1}$).

Since $n - 1 > n - k - 1$ for positive k , switching gets a larger expected utility.

- (c) What is the value of the information that Monty is offering you?

$$990 * \frac{1}{n} * \frac{k}{n-k-1}$$

Answer:

The formula for VPI is $MEU(e) - MEU(\emptyset)$. Thus, we want the difference between $EU(switch)$ (the optimal action if Monty opens the doors) and our expected utility from part (a).

(It is true that $EU(keep)$ happens to have the same numeric expression as in part (a), but this fact is not meaningful in answering this part.)

(d) Monty is changing his offer!

After you choose your initial door, you are given the offer to choose any other door and open this second door. If you do, after you see what is inside the other door, you may switch your initial choice (to the newly opened door) or keep your initial choice.

What is the value of this new offer?

$$\frac{990}{n}$$

Answer:

Intuitively, if we take this offer, it is as if we just chose two doors in the beginning, and we win if either door has the car behind it. Unlike in the previous parts, if the new door has a goat behind it, it is not more optimal to switch doors.

Mathematically, letting D_i be the event that door i has the car, we can calculate this as $P(D_2 \cup D_1) = P(D_1) + P(D_2) = 1/n + 1/n = 2/n$, to see that $MEU(\text{offer}) = 10 + 990 * \frac{2}{n}$. Subtracting the expected utility without taking the offer, we are left with $990 * \frac{1}{n}$.

(e) Monty is generalizing his offer: you can pay $\$d^3$ to open d doors as in the previous part. (Assume that $U(\$x) = x$.) You may now switch your choice to any of the open doors (or keep your initial choice). What is the largest value of d for which it would be rational to accept the offer?

$$d = \sqrt{\frac{990}{n}}$$


Answer:

It is a key insight (whether intuitive or determined mathematically) that the answer to the previous part is constant for each successive door we open. Thus, the value of opening d doors is just $d * 990 * \frac{1}{n}$. Setting this equal to d^3 , we can solve for d .

Q3. [Timed: 18 mins] Decision Networks and VPI


Valerie has just found a cookie on the ground. She is concerned that the cookie contains raisins, which she really dislikes but she still wants to eat the cookie. If she eats the cookie and it contains raisins she will receive a utility of -100 and if the cookie doesn't contain raisins she will receive a utility of 10 . If she doesn't eat the cookie she will get 0 utility. The cookie contains raisins with probability 0.1 .

- (a) We want to represent this decision network as an expectimax game tree. Fill in the nodes of the tree below, with the top node representing her maximizing choice.



- (b) Should Valerie eat the cookie? Yes No

- (c) Valerie can now smell the cookie to judge whether it has raisins before she eats it. However, since she dislikes raisins she does not have much experience with them and cannot recognize their smell well. As a result she will incorrectly identify raisins when there are no raisins with probability 0.2 and will incorrectly identify no raisins when there are raisins with probability 0.3 . This decision network can be represented by the diagram below where E is her choice to eat, U is her utility earned, R is whether the cookie contains raisins, and S is her attempt at smelling.



Valerie has just smelled the cookie and she thinks it doesn't have raisins. Write the probability, X , that the cookie has raisins given that she smelled no raisins as a simplest form fraction or decimal.

$$X = \boxed{0.04}$$

$$P(+r|-s) = \frac{P(-s|+r)P(+r)}{P(-s)} = \frac{P(-s|+r)P(+r)}{P(-s|+r)P(+r)+P(-s|-r)P(-r)} = \frac{.3*.1}{.3*.1+.8*.9} = \frac{.03}{.75} = .04$$

- (d) What is her maximum expected utility, Y given that she smelled no raisins? You can answer in terms of X or as a simplest form fraction or decimal.

$$Y = \boxed{-100X + 10(1 - X), 5.6}$$

$$MEU(-s) = \max(MEU(eating|-s), MEU(not eating|-s)) = \max(P(+r|-s) * EU(eating, +r) + P(-r|-s) * EU(eating, -r), MEU(not eating)) = \max(X * (-100) + (1 - X) * 10, 0) = X * 100 + (1 - X) * 10$$

- (e) What is the Value of Perfect Information (VPI) of smelling the cookie? You can answer in terms of X and Y or as a

simplest form fraction or decimal.

$$VPI = \boxed{0.75 * Y, 4.2}$$

$$VPI(S) = MEU(S) - MEU(\emptyset)$$

$$MEU(S) = P(-s)MEU(-s) + P(+s)MEU(+s)$$


$$P(-s) = .75 \text{ from part (c), } MEU(-s) = Y$$

$MEU(+s) = 0$ because it was better for her to not eat the raisin without knowing anything, smelling raisins will only make it more likely for the cookie to have raisins and it will still be best for her to not eat and earn a utility of 0. Note this means we do not have to calculate $P(+s)$.

$$MEU(\emptyset) = 0$$

$$VPI(S) = .75 * Y + 0 - 0 = .75 * Y$$

- (f) Valerie is unsatisfied with the previous model and wants to incorporate more variables into her decision network. First, she realizes that the air quality (A) can affect her smelling accuracy. Second, she realizes that she can question (Q) the people around to see if they know where the cookie came from. These additions are reflected in the decision network below.



Choose one for each equation:

	Could Be True	Must Be True	Must Be False
$VPI(A, S) > VPI(A) + VPI(S)$	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>
$VPI(A) = 0$	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
$VPI(Q, R) \leq VPI(Q) + VPI(R)$	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
$VPI(S, R) > VPI(R)$	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
$VPI(Q) \geq 0$	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
$VPI(Q, A) > VPI(Q)$	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
$VPI(S A) < VPI(S)$	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
$VPI(A S) > VPI(A)$	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>