


Q1. Sampling in Bayes Nets

We would like to analyze people’s ice cream eating habits on sunny and rainy days. Suppose we consider the weather, along with a person’s ice-cream eating, over the span of two days. We’ll have four random variables: W_1 and W_2 stand for the weather on days 1 and 2, which can either be rainy R or sunny S, and the variables I_1 and I_2 represent whether or not the person ate ice cream on days 1 and 2, and take values T (for truly eating ice cream) or F. We can model this as the following Bayes Net with these probabilities.



Suppose we produce the following samples of (W_1, I_1, W_2, I_2) from the ice-cream model:

~~R, F, R, F~~ ~~R, F, R, F~~ ~~S, F, S, T~~ ~~S, T, S, T~~ S, T, R, F
~~R, F, R, T~~ ~~S, T, S, T~~ ~~S, T, S, T~~ S, T, R, F ~~R, F, S, T~~

- (a) Using these samples, what is our estimate of $P(W_2 = R)$? 5/10 = 0.5
- (b) Cross off samples above which are rejected by rejection sampling if we’re trying to estimate $P(W_2|I_1 = T, I_2 = F)$

Rejection sampling seems to be wasting a lot of effort, so we decide to switch to likelihood weighting. Assume we generate the following six samples for (W_1, I_1, W_2, I_2) , given the evidence $I_1 = T$ and $I_2 = F$:

(S, T, R, F) (R, T, R, F) (S, T, R, F) (S, T, S, F) (S, T, S, F) (R, T, S, F)

- (c) Calculate the weight of each sample.

In this case, the evidence is $I_1 = T, I_2 = F$.

The weight of the first sample is thus $w = \Pr(I_1 = T|W_1 = S) \cdot \Pr(I_2 = F|W_2 = R) = 0.9 \cdot 0.8 = 0.72$

Similarly, we can calculate the weights of the other samples as follows:

(W_1, I_1, W_2, I_2)	w
S, T, R, F	0.72
R, T, R, F	0.16
S, T, R, F	0.72
S, T, S, F	0.09
S, T, S, F	0.09
R, T, S, F	0.02

- (d) Estimate $P(W_2|I_1 = T, I_2 = F)$ using our likelihood weights from the previous part.

To compute the probability estimate, we normalize the weights and find

$$\hat{P}(W_2 = R|I_1 = T, I_2 = F) = \frac{0.72 + 0.16 + 0.72}{0.72 + 0.16 + 0.72 + 0.09 + 0.09 + 0.02} = 0.889$$

$$\hat{P}(W_2 = S|I_1 = T, I_2 = F) = 1 - 0.889 = 0.111.$$

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? ($1000 * \frac{1}{n} + 10 * \frac{n-1}{n} = 10 + 990 * \frac{1}{n}$)

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?

$$EU(keep): 10 + 990 * \frac{1}{n}$$

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

Action that achieves MEU : switch

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

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? $\frac{990}{n}$

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 $MEU(offer) = 10 + 990 * \frac{2}{n}$. 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? $d = \sqrt{\frac{990}{n}}$

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 .