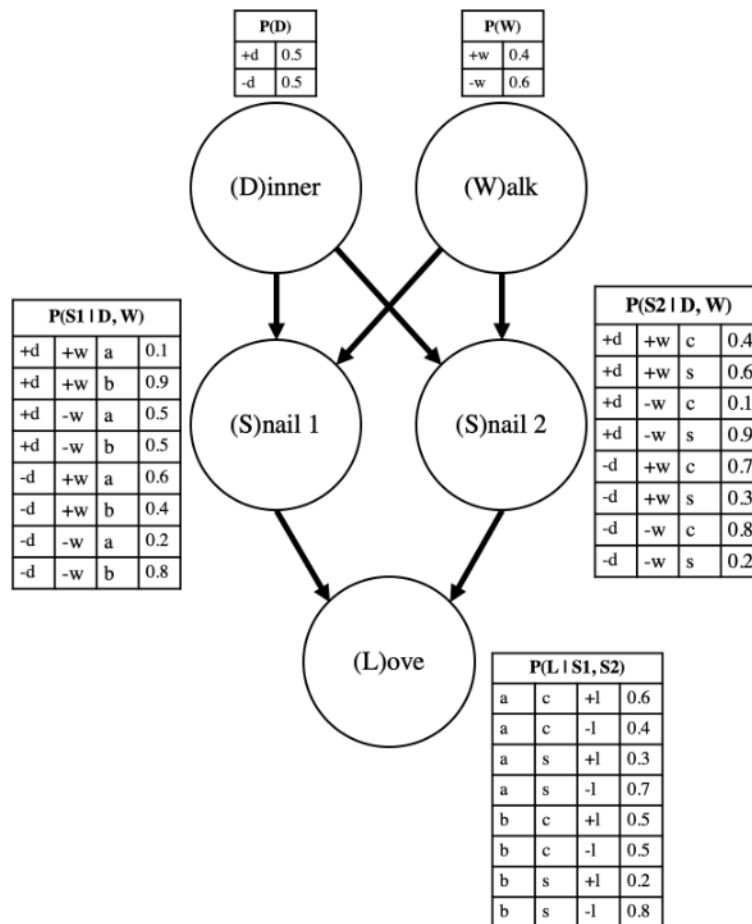


Q1. Snail Bayes

Celebrating the near-end of the semester, the CS188 TAs have gathered around the staff aquarium to check up on the snails and their search for love. To our excitement, two snails decided to go on a date! We don't know who the snails are, but we spent enough time around the terrarium to know that the first one (S_1) is either Alex (a) or Bubbles (b), and the second one (S_2) is either Cuddles (c) or Scorplborg (s). On the date, the snails will eat some dinner (D), which can be a beautiful flower (+d) or a poisonous mushroom (-d), and they will walk (W) around wonderful rocks (+w) or some treacherous puddle (-w). The snails are in the quest for love (L), which, depending on how the date goes, they can find (+l) or not (-l).



- (a) What is the probability of an outcome ($S_1 = a, S_2 = c, D = -d, W = +w, L = +l$), the probability that Cuddles and Alex are on a date, where they share a poisonous mushroom, walk around the wonderful rocks and find love?

- $0.5 * 0.4 * 0.7 * 0.5 * 0.4$
 $0.4 * 0.6 * 0.7 * 0.5 * 0.4$
 $0.6 * 0.6 * 0.7 * 0.5 * 0.4$
 None of the above

$P(a, c, -d, +w, +l) = P(+l | a, c) * P(a | -d, +w) * P(c | -d, +w) * P(-d) * P(+w) = 0.6 * 0.6 * 0.7 * 0.5 * 0.4 =$

0.0504

The date is about to start and people are making guesses for what's going to happen. One TA notices how adorable it would be if the snails were Bubbles and Cuddles.

(b) If Bubbles and Cuddles are on the date, we want to compute the probability of them eating a beautiful flower and walking around the wonderful rocks.

(i) What is the equivalent expression for this probability?

- $P(b, c, +d, +w)$
- $P(b, c \mid +d, +w)$
- $P(+d, +w \mid b, c)$
- $P(+d, +w)$

(ii) What minimal set of probability tables will we use in calculating this probability?

- $P(D)$
- $P(W)$
- $P(S_1 \mid D, W)$
- $P(S_2 \mid D, W)$
- $P(L \mid S_1, S_2)$
- None of the above

We need all the conditional probability table containing D, W, S_1, S_2 on the unconditioned side, which are the first four. The last table is left out because L is not part of the query and none of D, W, S_1, S_2 is conditioned on L

The snails are starving, so the date begins with a delightful dinner. The snails start sharing a mushroom as their dish of choice.

(c) Given their choice of dinner, what is $P(S1 \mid -d)$, the belief over which snail is S1? Please answer in **decimal** numbers. You should not need a calculator.

(i) $P(S_1 = a \mid D = -d) =$ 0.36

$P(S1 = a \mid -d) = \sum_w P(a \mid -d, w) * P(w) = 0.6 * 0.4 + 0.2 * 0.6 = 0.36.$

The snails found love! We are now trying to find the probability that the other snail was Bubbles given all evidence so far, $P(b \mid c, -d, +l)$. The TAs are tired of multiplying probabilities, so they instead try another way. The late TA actually wrote down memories of previous dates he has witnessed in a notebook. He can sample some of his memories from the notebook and help us learn probabilities.

(d) If the TA uses prior sampling, what is the probability of obtaining the sample $[D = -d, W = +w, S_1 = b, S_2 = c, L = -l]$?

- $0.5*0.4*0.6*0.3*0.2$
- $0.4*0.4*0.9*0.2*0.8$
- $0.6*0.1*0.7*0.1*0.2$
- $0.5*0.4*0.4*0.7*0.5$
- $0.25*0.24*0.9*0.1*0.5$
- $0.4*0.5*0.24*0.21*0.25$
- None of the above

Prior sampling samples without taking the evidence into account, so the probability of the sample is $P(-d)P(+w)P(b \mid -d, +w)P(c \mid -d, +w)P(-l \mid b, c)$.

(e) If the TA samples $[D = -d, W = +w, S_1 = b, S_2 = c, L = -l]$, would rejection sampling discard the memory?

- Yes
- No

Rejection sampling discards samples inconsistent with the evidence. In this case, $-l$ is inconsistent with the fact that our snails did in fact find love.

(f) Assuming that the TA actually sampled using likelihood weighing and obtained $[D = -d, W = +w, S_1 = b, S_2 = c, L = +l]$, what is the weight of this sample?

- $0.5*0.5*0.5$ $0.5*0.7*0.5$
- $0.4*0.9*0.6$ $0.5*0.3*0.5$
- $0.4*0.24*0.6$ $0.6*0.3*0.6$
- None of the above

The weight of a sample in likelihood weighing is the probability of the evidence given their parents: $P(-d)P(c | -d, +w)P(+l | b, c)$.

- (g) Sampling using likelihood weighing will systematically underestimate the probability of a variable conditioned on one of its ancestors.
- Yes, because likelihood weighing does not sample all the variables, and thus creates a bias
 - Yes, but not for the reason above
 - No, because likelihood weighing is unbiased
 - No, but not for the reason above

Likelihood weighing is an unbiased sampling procedure.

- (h) To estimate $P(b | c, -d, +l)$, the TA samples five memories in a row:

$[D = -d, W = +w, S_1 = b, S_2 = c, L = +l]$,
 $[D = -d, W = +w, S_1 = b, S_2 = c, L = +l]$,
 $[D = -d, W = +w, S_1 = a, S_2 = c, L = +l]$,
 $[D = -d, W = +w, S_1 = a, S_2 = c, L = +l]$,
 $[D = -d, W = +w, S_1 = b, S_2 = c, L = +l]$.

Could these memories have been generated using Gibbs sampling?

- Yes, because all evidence variables are consistent with their values in the query $P(b | c, -d, +l)$.
- Yes, but the reason above is incorrect because there exist other samples sequences that fit the condition in the previous choice but cannot be generated by Gibbs sampling.
- No, because the sequence of samples only differs by S_1 , the query variable. The values of W , a variable that is not part of the query, never changes throughout the sequence.
- No, but the reason above is incorrect because there exist other samples sequences that fit the condition in the previous choice but cannot be generated by Gibbs sampling.

Since each neighboring sample differs by at most one variable value AND the evidence variables are not change, this sequence could be generated via Gibbs sampling, and thus eliminate choice 3 and 4. [0- i 1: no change, 1- i 2: no change, 2- i 3: only changes S_1 , 3- i 4: no change, 4- i 5: no change]

But choice one "all evidence variables are consistent with their values in the query $P(b | c, -d, +l)$ " is insufficient, because the following sequence fits the condition, yet cannot be generated by Gibbs Sampling:


$[D = -d, W = +w, S_1 = b, S_2 = c, L = +l]$,
 $[D = -d, W = -w, S_1 = a, S_2 = c, L = +l]$,
 $[D = -d, W = -w, S_1 = b, S_2 = c, L = +l]$.

Because 0- i 1: changes both W and S_1

Just a quick note about choice 3: it is totally possible that when it is W 's turn to be re-sampled from $P(W|\text{everything else})$, the result just so happens to remain $+w$ all times, throughout the process of the 5 consecutive samples.

Q2. Bayes Nets: Elimination

(a) Consider running variable elimination on the Bayes Net shown below.




First, we eliminate D to create a factor f_1
 Next, we eliminate E to create a factor f_2
 Next, we eliminate H to create a factor f_3

From the list below, select all factors that remain after D , E and H have been eliminated.

- | | | | |
|---|--|--------------------------------------|---|
| <input type="checkbox"/> f_1 | <input checked="" type="checkbox"/> $P(A)$ | <input type="checkbox"/> $P(D A)$ | <input type="checkbox"/> $P(G C)$ |
| <input checked="" type="checkbox"/> f_2 | <input type="checkbox"/> $P(A F)$ | <input type="checkbox"/> $P(D A, F)$ | <input checked="" type="checkbox"/> $P(G C, F)$ |
| <input checked="" type="checkbox"/> f_3 | <input checked="" type="checkbox"/> $P(B)$ | <input type="checkbox"/> $P(E B)$ | <input type="checkbox"/> $P(H F)$ |
| | <input type="checkbox"/> $P(B A)$ | <input type="checkbox"/> $P(E C)$ | <input type="checkbox"/> $P(H C, F)$ |
| | <input type="checkbox"/> $P(B C)$ | <input type="checkbox"/> $P(E B, C)$ | <input type="checkbox"/> $P(H J)$ |
| | <input type="checkbox"/> $P(C)$ | <input type="checkbox"/> $P(F)$ | <input type="checkbox"/> $P(J C)$ |
| | <input checked="" type="checkbox"/> $P(C B)$ | <input type="checkbox"/> $P(F A, B)$ | <input type="checkbox"/> $P(J F)$ |
| | | <input type="checkbox"/> $P(F G)$ | <input type="checkbox"/> $P(J H)$ |

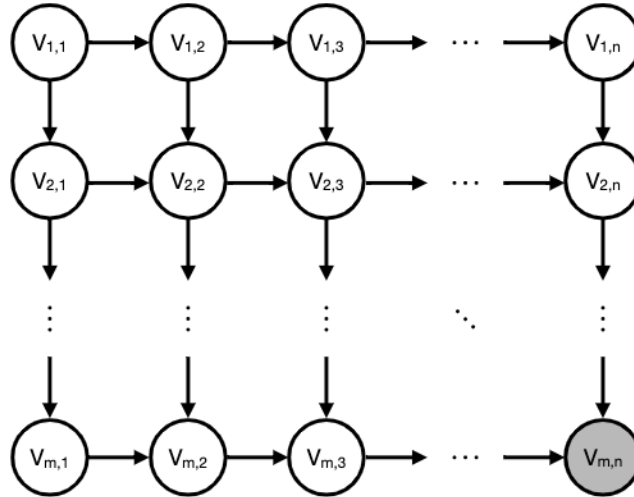
(b) Consider the Bayes Net shown below. Each variable in the Bayes Net can take on two possible values.



You are given the query $P(C|F)$, which you would like to answer using variable elimination. Please find a variable elimination ordering where the largest intermediate factor created during variable elimination is as small as possible.

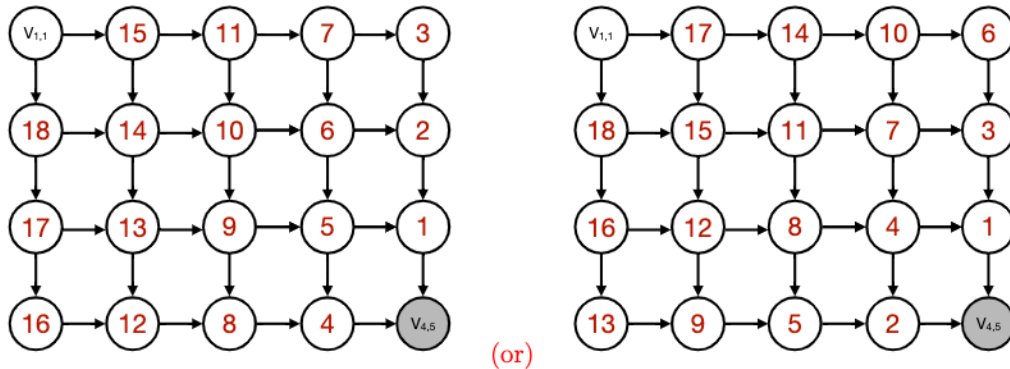
Elimination ordering: ABED, AEBD, or EABD

- (c) Consider doing inference in an $m \times n$ lattice Bayes Net, as shown below. The network consists of mn binary variables $V_{i,j}$, and you have observed that $V_{m,n} = +v_{m,n}$.




You wish to calculate $P(V_{1,1} | +v_{m,n})$ using variable elimination. To maximize computational efficiency, you wish to use a variable elimination ordering for which the size of the largest generated factor is as small as possible.

- (i) First consider the special case where $m = 4$ and $n = 5$. A reproduction of the lattice is shown below, with variable names for non-query variables omitted. Please provide your optimal elimination ordering for this example by numbering the nodes below in the order they will be eliminated (i.e. write a number such as 1, 2, 3, ... inside every node that will be eliminated.)



Note that there is actually more than one correct ordering, and that a few minor variations on the orderings given above are possible. However, it's important to start near the same corner as the evidence variable and to never create a factor that involves more than 4 non-evidence variables. However, the ordering shown below is suboptimal (eliminating node 6 will create a size 2^5 factor involving the five nodes highlighted in blue):



(ii) Now consider the general case (assume $m > 2$ and $n > 2$). What is the size of the largest factor generated under the most efficient elimination ordering? Your answer should be the number of rows in the factor's table, expressed in terms of m and n .

Size (number of rows) of the largest factor: $2^{\min(m,n)}$