CS 188 Summer 2021 — Introduction to Artificial Intelligence — Exam Prep 5A Solutions

Q1. Bayes Nets

(a) For the following graphs, explicitly state the minimum size set of edges that must be removed such that the corresponding independence relations are guaranteed to be true.

Marked the removed edges with an 'X' on the graphs.

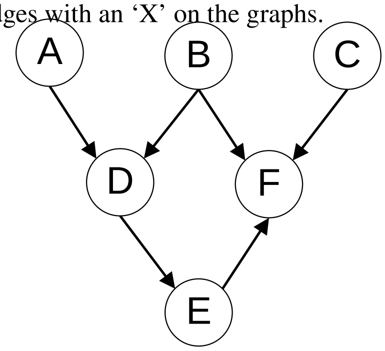

$$A \perp\!\!\!\perp B \mid F$$
$$A \perp\!\!\!\perp F \mid D$$
$$B \perp\!\!\!\perp C$$

(i) AD

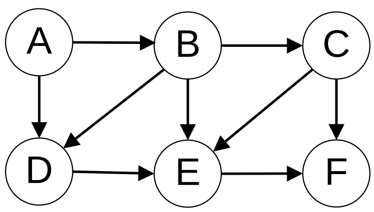

$$A \perp\!\!\!\perp D \mid B$$
$$A \perp\!\!\!\perp F \mid C$$
$$C \perp\!\!\!\perp D \mid B$$

(ii) AD, (EF OR AB)

(b) You're performing variable elimination over a Bayes Net with variables A, B, C, D, E. So far, you've finished joining over (but not summing out) C, when you realize you've lost the original Bayes Net!

Your current factors are $f(A)$, $f(B)$, $f(B, D)$, $f(A, B, C, D, E)$. Note: these are factors, NOT joint distributions. You don't know which variables are conditioned or unconditioned.

(i) What's the smallest number of edges that could have been in the original Bayes Net? Draw out one such Bayes Net below.

Number of edges = 5

The original Bayes net must have had 5 factors, 1 for each node. f(A) and f(B) must have corresponded to nodes A and B, and indicate that neither A nor B have any parents. f(B, D), then, must correspond to node D, and indicates that D has only B as a parent. Since there is only one factor left, f(A, B, C, D, E), for the nodes C and E, those two nodes must have been joined while you were joining C. This implies two things: 1) E must have had C as a parent, and 2) every other node must have been a parent of either C or E.

The below figure is one possible solution that uses the fewest possible edges to satisfy the above.

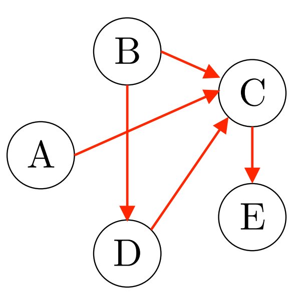

(ii) What's the largest number of edges that could have been in the original Bayes Net? Draw out one such Bayes Net below.

Number of edges = 8

The constraints are the same as outlined in part i). To maximize the number of edges, we make each of A, B, and D a parent of both C and E, as opposed to a parent of one of them.

The below figure is the only possible solution.

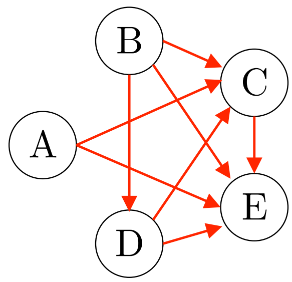

Q2. Bayes Nets: Sampling

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

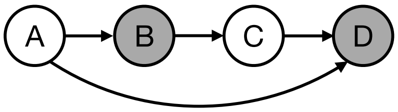

$P(A)$	
$+a$	0.5
$-a$	0.5

$P(B\|A)$		
$+a$	$+b$	0.8
$+a$	$-b$	0.2
$-a$	$+b$	0.4
$-a$	$-b$	0.6

$P(C\|B)$		
$+b$	$+c$	0.1
$+b$	$-c$	0.9
$-b$	$+c$	0.7
$-b$	$-c$	0.3

$P(D\|A,C)$			
$+a$	$+c$	$+d$	0.6
$+a$	$+c$	$-d$	0.4
$+a$	$-c$	$+d$	0.1
$+a$	$-c$	$-d$	0.9
$-a$	$+c$	$+d$	0.2
$-a$	$+c$	$-d$	0.8
$-a$	$-c$	$+d$	0.5
$-a$	$-c$	$-d$	0.5

(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.

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

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

(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.

■ $-a$ $-b$
☐ $+a$ $+b$
■ $+a$ $-b$
☐ $-a$ $+b$

(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
$-a$	$+b$	$-c$	$+d$	0.5
$+a$	$+b$	$-c$	$+d$	0.1
$+a$	$+b$	$-c$	$+d$	0.1
$-a$	$+b$	$+c$	$+d$	0.2
$+a$	$+b$	$+c$	$+d$	0.6

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

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

(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!