CPSC 502 — Fall 2013

Assignment 3

Solution

Question 1

Solution (a), (b) see `http://www.cs.ubc.ca/~poole/cs502/2013/as3/plumbing.xml` for an AIspace representation.

(c) Give some examples of the interesting things you can do!

(d) there are many more arcs, in particular there are arcs between any pair of nodes X and Y in the new ordering if there were arcs in the original network or if they have common ancestors in the original (causal) belief network that are lower in the total ordering than both X and Y or where there are common descendents (in the original graph) of X and Y that are also parents in the new graph. [This is a convoluted way of using D-separation in the original causal network to determine parents in the inverse network.] You just need to give some examples.

Question 2

Solution

(a) To compute $P(e)$, you can sum out D and F, and the factors created are just 1's (that is why they can be pruned). I'll ignore these factors.
You start off with the factors $f_0(A), f_1(B), f_2(A,B,C), f_3(C,E)$.
Eliminating A, you multiply f_0 and f_2, and sum out A, creating a factor $f_4(B,C)$.
Eliminating B, you multiply f_1 and f_4, sum out B, and create a factor $f_5(C)$.
Eliminating C, you multiply f_5 and f_3, sum out C, and the resulting factor $f_6(E)$ represents $P(E)$.

(b) To compute $P(e|\neg f)$, you create a factor (lets call it f_7 to avoid confusion with the previous part) $f_7(C)$.
You can prune D. Eliminating A and B acts exactly as before, creating the same factors.
Eliminating C, you multiply f_5, f_3 and f_7, sum out C, with the resulting factor $f_8(E)$. To get the answer, you need to divide each element of f_8 by $\sum_E f_8(E)$. This creates $f_9(E)$ which represents $P(E|\neg f)$.

(c) Everything can be pruned! Or there are lots of tables with just ones, and a constant factor. $P(a|d) = P(a)$.

(d) E and D can be pruned.

(e) What I wanted you to notice here is that D is not irrelevant when F is also observed.

(f) When C, E or F are observed and B isn't observed.

Question 3

Solution In all of these I assume you remove E is it is irrelevant.

(a) Rejection sampling: Sample B, then sample D, then reject the sample if it doesn't assign $D = \mathit{false}$. Then sample the rest of the variables. Return the proportion of the non-rejected samples where $F = \mathit{true}$.

(b) Importance sampling: Sample B. Weight those samples where $B = true$ by 0.9 and those samples with $B = false$ by 0.2. [You can actually do something smarter if you don't sample B by its prior, but use a different proposal distribution.]. Then sample the other variables. Return the weighted average of the samples with $F = true$.

(c) Particle filtering: create lots of samples (here I will use 1000 as an example). Sample B 1000 times. Weight each sample using D as in importance sampling. Now generate a new 1000 samples, choosing each sample with probability proportional to its weight. Sample the rest of the variables with these new samples, and return the proportion of the samples where $F = true$.

Question 4

Solution a) & b) see `http://www.cs.ubc.ca/~poole/cs502/2013/as3/cab.xml`

c) One solution is: `http://www.cs.ubc.ca/~poole/cs502/2013/as3/cab_collusion.xml`

d) one solution is: `http://www.cs.ubc.ca/~poole/cs502/2013/as3/cab_accident.xml` The rate of accidents was not given and cannot be inferred from the description. I arbitrarily chose $P(accident) = 0.1$.

Question 5

Solution There is no right answer for this.

Question 6

Solution See the wiki for solutions.
`http://wiki.ubc.ca/Course:CPSC:Artificial_Intelligence`

Question 7