....

STAT 535 Final Exam

Wednesday December 16, 2009, 4:30-6:20

Student name: ....

- notes and books are allowed
- electronic devices are not allowed

- *Do Well!*

Problem 1

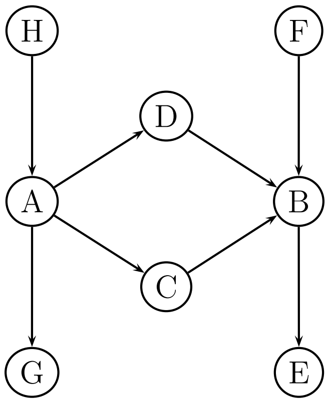

3 points

Which of the conditional independence statements below are true? For those which are not, give an unblocked path.

unblocked path, if any:

$H \perp E$	TRUE	FALSE	
$H \perp F$	TRUE	FALSE	
$H \perp D \mid A, G$	TRUE	FALSE	
$C \perp D \mid A$	TRUE	FALSE	
$C \perp D \mid A, B$	TRUE	FALSE	
$D \perp F \mid E$	TRUE	FALSE	

Problem 2

In the graphical model below all variables are discrete, with values $\{0, 1\}$.

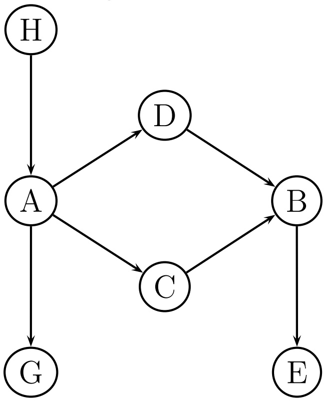

1 point

2.1 Moralize and triangulate this graphical model.

2 points

2.2 Draw a junction tree for this graphical model.

4 points

2.3 Give, step by step, an *efficient* algorithm to compute the probability $P(C = 0|A = 0, B = 0)$ in this graphical model. It is assumed that $P_{ABCDEGH}$ is represented by conditional probability tables in the standard way.

Details:

1. The algorithm should be *specific* for computing $P(C = 0|A = 0, B = 0)$.
2. An algorithm for this is *efficient* if it does not do useless calculations, and if it does not create tables or potentials larger than it needs to. Your algorithm must be efficient to be acceptable as a solution, or only a little worse. Note that an algorithm that solves general inferences in this graphical model may not be efficient according to the present definition.
3. There are several possible methods to compute this probability. Give the name of the method you decide to use (e.g Variable elimination, Junction Tree, Message Passing, Sum-Product, etc).
4. Give the size of the table/potential/message/separator/clique on which every step of the algorithm operates, and an approximate number of operations for each step.

(extra white space)

Problem 3: Rudolf the Reindeer marginalized

Characters: RUDOLF THE REINDEER, SANTA'S LITTLE HELPER
RUDOLF THE REINDEER: Look Santa's Little Helper, I decorated my horns this season using a Markov Ranodom field. The variables A, B, C, D are light bulbs which are on and off (that is 1 and 0 to you) randomly, with the joint distribution given by

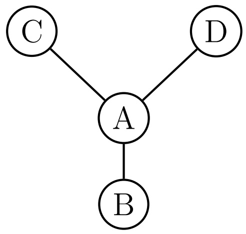

$$P_{ABCD}(a,b,c,d) = \frac{1}{Z}\phi_{AB}(a,b)\phi_{AC}(a,c)\phi_{AD}(a,d)$$

$$\text{with } \phi_{AX}(a,x) = 2^{-|a-x|}$$

SLH: Brilliant! Do you realize that your MRF is a tree, and as such it has a factorization in terms of clique and separator marginals?

RTR: A tree! That's amazing, SLH. I hadn't thought of that! But I wonder: could it be true that A, B, C and D all have the same marginal probability of being on?

2 points

3.1 Prove or disprove: $P_A(x) = P_B(x)$ for $x = 0, 1$.

SLH: I don't know how to answer your question, but I can see that if you fix A, you can write P_{ABCD} in a simpler form. Maybe this will let me compute $P_{A,B}$ and then I will be able to answer...

2 points

3.2 SLH doesn't notice it, but with his observation it's easier to calculate $P_{B|A}$. Show him how, and obtain the tables $P_{B|A=0}$, $P_{B|A=1}$ (numerical values required). Remember the brute force solution is not acceptable.

RtR: Look, SLH, now that you have $P_{B|A}$, $P_{C|A}$, $P_{D|A}$, I see a way to compute Z, the normalization constant, by moving the sums inside the products!

2 points

3.3 Do what RtR says he can: calculate the value of Z (no brute force please).

2 points

3.4 Now find P_A and P_{AB}, using your results from 3.2,3.3.

Problem 4

Santa has enlisted n elves, $X_{1:n}$ to deliver presents on Christmas Eve. The elves are spread all over the world, but each of them can communicate with a set of "neighbors". The pattern of connections forms an undirected graph with no cycles. Santa himself is node X_0 in this graph.

3 points

4.1 By midnight, each elf X_i has delivered k_i presents. Design an efficient algorithm based on communication only between neighbors, that will allow Santa to find out the total number of presents that were delivered until midnight?

Assume that Santa and the elves can only know who their neighbors are (the local structure of the graph) but none of them knows the global structure of the graph, except that it is connected, and that it has no cycles. They can send any kind of data in their messages, but for efficiency they try not to send too much useless data. Santa's neighborhood is not special, but you can assume Santa is able to send special messages if you think this is necessary.

2 points

4.2 Can you design an efficient message passing algorithm, to be run at midnight, so that *each elf* can know how many presents were delivered? (If your algorithm in 4.1 already does this, explain why.)

2 points

4.3 Can you design an algorithm so that Santa and all the elves can be updated regularly on the total number of presents delivered?