AI Homework 3 – Probability and Bayes Nets

1. Suppose Professor Kirlin is visiting an amusement park where the following events might happen: riding a roller coaster (R), eating a soft pretzel (P), feeling sick (probably because of eating pretzels or riding the roller coaster) (S), and feeling hungry (H).

   You are given the following full joint probability distribution over these four events, each of which either happens or doesn't happen.

Roller Coaster (R)	Eat Pretzel (P)	Feel Sick (S)	Feel Hungry (H)	Joint probability
T	T	T	T	0.144
T	T	T	F	0.216
T	T	F	T	0.016
T	T	F	F	0.024
T	F	T	T	0.096
T	F	T	F	0.024
T	F	F	T	0.224
T	F	F	F	0.056
F	T	T	T	0.008
F	T	T	F	0.012
F	T	F	T	0.032
F	T	F	F	0.048
F	F	T	T	0.008
F	F	T	F	0.002
F	F	F	T	0.072
F	F	F	F	0.018

   Use the table to calculate the following probabilities. Show all your work.

   (a) P(S, P, R)
   (b) P(¬S | R, P)
   (c) P(R, P | ¬S)
   (d) P(R | P, ¬H)

2. Punxsutawney Phil is a groundhog that lives in Pennsylvania, who supposedly can predict how long winter will last. On Groundhog Day, February 2, Phil emerges from his burrow and (according to tradition) if he sees his shadow, there will be six more weeks of winter. If he doesn't see his shadow, there will be an early spring.

   One year, a new groundhog, Punxsutawney Jill, comes to town, and she claims that she makes better predictions than Punxsutawney Phil. Here are their statistics:

   When an early spring is going to happen, Phil detects it correctly 90% of the time. (In other words, *given that* it's an early spring, Phil predicts it correctly 90% of the time.) However, he has a 15% false positive rate (when he predicts an early spring *given that* it doesn't actually happen). On the other hand, Jill correctly recognizes true early springs only 85% of the time, but she only has a 10% false positive rate. Furthermore, early springs only happen 20% of the time.

Answer the following questions. Show your work. Hint: Use three events here: E=there is an early spring. PE=Phil predicts an early spring. JE=Jill predicts an early spring. **I suggest beginning by just translating the description above into marginal and/or conditional probabilities using E, PE, and JE, and writing them down.** Also, this is not a Bayes net problem, just use the rules of probability directly.

(a) If Phil predicts an early spring, what is the probability that the prediction is correct? In other words, *given that* Phil predicts an early spring, what is the probability that the early spring occurs?

(b) If Jill predicts an early spring, what is the probability that the prediction is correct?

(c) What is Phil's overall prediction accuracy? (What is the probability he makes a correct prediction, regardless of the weather?) Hint: this requires some thinking, because there are two ways that a groundhog can make a correct prediction.

(d) What is Jill's overall prediction accuracy? Which groundhog is better at predicting the weather?

3. Consider the following Bayes network:

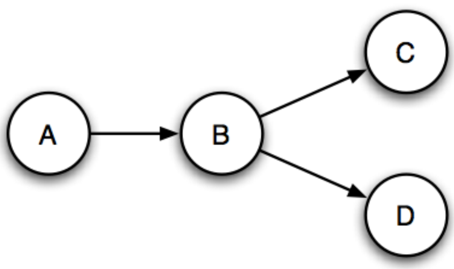

Here are the CPTs (conditional probability tables) for this network:

(I follow the book's convention of using uppercase letters to stand for a random variable, and lowercase letters to for a specific assignment of a value to the random variable. For instance "A" is a random variable, but "a" is the specific setting of
"A = true" and "~a" means "A = false.")

P(a) = 0.4

CPT for P(B | A)

A	P(b | A)
true	0.7
false	0.3

CPT for P(C | B)

B	P(c | B)
true	0.2
false	0.6

CPT for P(D | B)

B	P(d | B)
true	0.5
false	0.8

Suppose we know the value of random variables C and D; specifically, assume C is true and D is false. Use the Bayes net exact inference algorithm to calculate
P(A | c, ~d). (This means calculate the probability of A being true [and then being false] given the values of C and D). Show all of your work, including the steps involving the definition of conditional probability, where you introduce the normalization constant, the marginalization step, the re-arrangement of the summations to make the calculation as efficient as possible, drawing the tree to show your calculations, and the normalization step at the end.

4. In this problem we will study a simplified version of how being left- or right-handed might be passed from parents to children via genetics. Current research suggests handedness is affected by genetics, biology, and the environment, but we will assume for the moment it is solely inherited from your parents, with a small chance of mutation.

   Let us invent three random variables, called H_{mom}, H_{dad}, and H_{child}. These are binary random variables, with values of either *left* or *right*. Each variable represents the probability that a (particular) mother, father, and child are left or right-handed. Furthermore, suppose the way we will model handedness is that there is a gene that exerts most of the control over whether someone will be right- or left-handed. We will invent three more binary random variables, G_{mom}, G_{dad}, and G_{child}, again with possible values of *left* or *right*, denoting the probability that the mother, father, or child has the gene set a certain way. (Again, these are binary random variables, so we're discounting the possibility of being ambidextrous and such.)

   Suppose that the gene for handedness is equally likely to be inherited from each parent (but it only comes from one of them). There is also a small nonzero probability *m* that a mutation will happen and the handedness gene will be flipped in the child. Also assume that once the handedness gene is set in a person, the actual handedness of that person will match the gene 90% of the time.

   Here are three possible bayes nets that could represent this situation:

(1)

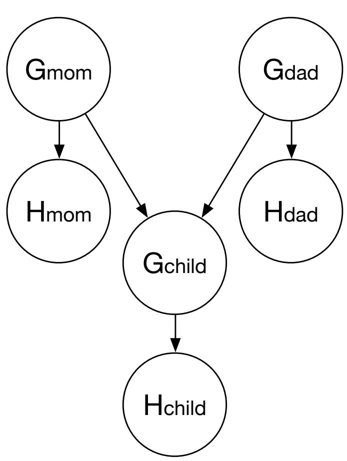

(2)

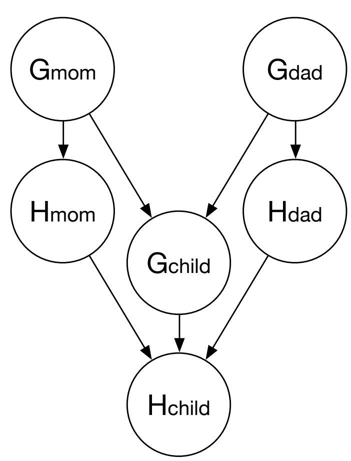

(3)

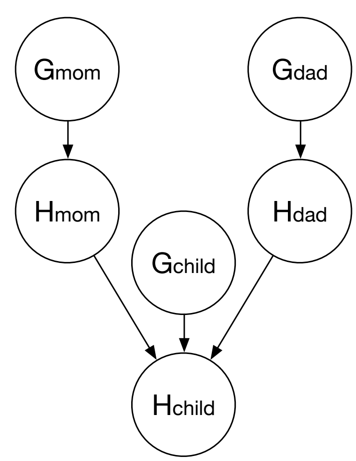

a. Which one(s) of the three bayes nets has the property that $P(G_{mom}, G_{dad}, G_{child}) = P(G_{mom})P(G_{dad})P(G_{child})$?

b. Which one of the three networks is the best description of the situation described in the question setup? Explain why you chose your answer.

c. Write down the CPT for the G_{child} node in bayes net (1). You will need to use the variable *m*, which represents the probability of a mutation (see description above). Since we are using left/right instead of true/false, you can fill in this CPT, which might be easier to read:

G_{mom}	G_{dad}	G_{child}
left	left	$P(G_{child} = left \mid G_{mom} = left, G_{dad} = left) =$ $P(G_{child} = right \mid G_{mom} = left, G_{dad} = left) =$
left	right	$P(G_{child} = left \mid G_{mom} = left, G_{dad} = right) =$ $P(G_{child} = right \mid G_{mom} = left, G_{dad} = right) =$
right	left	$P(G_{child} = left \mid G_{mom} = right, G_{dad} = left) =$ $P(G_{child} = right \mid G_{mom} = right, G_{dad} = left) =$
right	right	$P(G_{child} = left \mid G_{mom} = right, G_{dad} = right) =$ $P(G_{child} = right \mid G_{mom} = right, G_{dad} = right) =$

d. Write down the CPT for the H_{child} node in bayes net (1). You can follow the model above (though there's only one parent for H_{child}).

5. You have a bag containing three biased coins, called coin a, coin b, and coin c, with probabilities of coming up heads of 30%, 60%, and 70% respectively. You reach in and pick a coin randomly from the bag, but you can't tell which coin you picked (they all look the same to you). You flip that same coin three times and observe whether you got heads or tails each time.

   a. Define a complete Bayesian network for this situation, showing the structure of the network and the CPTs.

      Hint: you will need four random variables, one for which coin you chose, and three for the flips. The three coin flips are Boolean random variables (two-valued), but the coin-choice random variable is three-valued.

   b. Calculate which coin was most likely to have been drawn from the bag if the observed flips were heads, heads, and tails. Show all of your work. In other words, calculate the probabilities that given that sequence of flips, the original coin was coin a versus coin b versus coin c. (Note that this last problem is not a ML/MAP problem, it's a direct probability calculation using the Bayes net.)