

CS 486/686 Assignment 3

Due Date: 11:59pm on Thursday, November 15, 2018

Instructions

- Submit the assignment in the Dropbox labeled Assignment 3 Submissions on Learn.
- No late assignment will be accepted.
- This assignment is to be done individually.
- For any programming question, you may use the language of your choice. We highly recommend using Python if you are familiar with it.
- Lead TA: Colin Vandenhof (cm5vande@uwaterloo.ca). Colin's office hours will be posted on the course website.

1 Inference in Bayesian Networks (93 points)

Your dog Fido has been howling for the last three hours. In the past, Fido may howl when he is genuinely sick, but he also howls sometimes when there is a full moon or when the neighbour's dog is howling. You are trying to figure out if you should take him to the vet. You will implement a system to help you decide what to do. First, you will implement the variable elimination algorithm. Second, you will define a Bayesian network to model the reasons for Fido to howl and to compute the probability that Fido is genuinely sick.

You should perform the computations in your Bayesian network by using your variable elimination algorithm. However, if you are not able to complete your implementation, you will get partial marks if you perform the computations by hand.

1. (5 points) Implement the variable elimination algorithm. The implementation is only worth 5 points if it is complete and well-documented. However, you will earn 1/3 of the marks in part 2 (b, c, d, and e) if you perform the computations using your implementation rather than by hand.

You may assume that all the random variables are Boolean. You may use the language of your choice, but we will give you tips on the implementation if you use Python.

Please include documentation for every function in your program. If a function is long, include documentation within the function. In addition, please include detailed instructions for the TA to run your program. You will earn 5 points if your program is complete and well-documented.

Factors are essentially multi-dimensional arrays. Hence, use numpy multi-dimensional arrays as your main data structure. If you are unfamiliar with numpy, go through the following tutorial: <https://docs.scipy.org/doc/numpy-1.15.1/user/quickstart.html>.

Please implement the following five functions.

[Tip: Test each function individually using examples from the lecture slides. If you wait till the end to test your entire program, it will be much harder to debug.]

- (a) **restrictedFactor = restrict(factor, variable, value)**: This function restricts the given variable in the given factor to the given value. [Tip: use slicing operations to implement this function.]
- (b) **productFactor = multiply(factor1, factor2)**: This function performs point-wise multiplication of the given two factors. [Tip: take advantage of the numpy broadcasting rules to multiply multi-dimensional arrays of different shapes.]
- (c) **resultFactor = sumout(factor, variable)**: This function sums out a variable in a given factor. [Tip: use the sum operation to implement this function.]

- (d) **normalizedFactor = normalize(factor)**: This function normalizes a factor by dividing each entry by the sum of all the entries. This function is useful when the factor is a distribution (i.e. the sum of the probabilities must be 1.).
- (e) **resultFactor = inference(factorList, queryVariables, orderedListOfHiddenVariables, evidenceList)**: This function computes $Pr(\text{queryVariables}|\text{evidenceList})$ by the variable elimination algorithm. This function should do the following.
- i. Restrict the factors in factorList according to the evidence in evidenceList.
 - ii. Based on the order of the hidden variables in orderedListOfHiddenVariables, sum out each hidden variable from the product of the factors in factorList.
 - iii. If the result is a product of multiple factors, then multiply them to produce a single factor.
 - iv. If the resulting factor should denote a probability distribution (i.e. the probabilities should sum up to 1), then normalize the resulting factor.
2. (88 points) To decide on whether to take Fido to the vet or not, you have the following information available to you.
- Fido often howls if he is genuinely sick. However, Fido is a healthy dog and is only sick 5% of the time. If Fido is really sick, then he probably has not eaten much of his dinner and has food left in his bowl. In the past, when Fido is sick then he does not eat his dinner 60% of the time. However, when Fido is healthy he still does not eat his dinner 10% of the time.
 - Fido often howls if there is a full moon or your neighbour's dog howls. Your neighbour's dog sometimes howls at the full moon and sometimes howls when your neighbour is away. However, your neighbour's dog is never affected by Fido's howls. Since you live on Earth, there is a full moon once out of every twenty-eight days. You have observed that your neighbour travels three days out of every ten.
 - If there is no full moon and your neighbour is at home then your neighbour's dog never howls. If there is no full moon and your neighbour is away then your neighbour's dog howls half of the time. If there is a full moon and your neighbour is at home then your neighbour's dog howls 40% of the time. Finally, if there is a full moon and your neighbour is away then your neighbour's dog howls 80% of the time.
 - If all three triggers are there (Fido sick, full moon, and your neighbour's dog howling), Fido howls 99% of the time. If none of the three triggers are there, Fido will not howl.
 - If Fido is sick he is more likely to howl. If being sick is the only trigger present, Fido will howl half of the time. If Fido is sick and your neighbour's dog is howling then Fido will howl 75% of the time. If Fido is sick and there is a full moon, then Fido will howl 90% of the time.

- If Fido is not sick then he is less likely to howl. The full moon and your neighbour's dog will cause Fido to howl 65% of the time. If there is only a full moon then Fido will howl 40% of the time. If there is no full moon, but your neighbour's dog is howling, then Fido will howl 20% of the time.

- (a) (40 points) Construct a Bayesian network based on the above information about Fido.

Tip: You should be able to directly take the probabilities given in the text above and insert them into the corresponding conditional probability tables. There is no need to perform any computation to obtain the conditional probabilities.

What to hand in: Show the graph defining the network and the conditional probability tables associated with each node in the graph. The network should encode all the information stated above. It should contain six nodes corresponding to the following binary random variables.

- FH : Fido howls.
- S : Fido is sick.
- B : There is food left in Fido's food bowl.
- M : There is a full moon.
- NA : Your neighbour's away.
- NH : Your neighbour's dog howls.

The arcs in your Bayesian network should accurately capture the probabilistic dependencies between the variables.

- (b) (15 points) What is the prior probability that Fido howls?

What to hand in: Indicate what queries (i.e., $\text{Pr}(\text{variables}|\text{evidence})$) you used to compute those probabilities. Whether you answer the queries by hand or using the program you wrote part 1, show or provide a printout of the factors computed at each step of variable elimination (as done in the lecture notes). **Use the following variable ordering when summing out variables in the variable elimination algorithm: B, FH, M, NA, NH, and S.**

For parts (b), (c), (d), and (e), a maximum of two thirds of the marks are earned if you answer the question correctly by doing the computations by hand instead of using your program.

- (c) (12 points) Fido is howling. You look out the window and see that the moon is full. What is probability that Fido is sick?

What to hand in: Same as Question 2b.

- (d) (12 points) (Fido is howling and the moon is still full.) You walk to the kitchen and see that Fido has not eaten and his food bowl is full. Given this new information, what is the probability that Fido is sick?

What to hand in: Same as Question 2b.

- (e) (9 points) (Fido is howling, the moon is full and Fido has not eaten his dinner.) You decide to call your neighbour to see if they are home or not. The phone rings and rings so you conclude that your neighbour is away. Given this information, what is the probability that Fido is sick?

What to hand in: Same as Question 2b.

2 Constructing Bayesian Networks (24 points)

Consider the Bayesian network for the Holmes scenario in Figure 1 below (without the Radio node). The conditional probabilities are different from the ones shown in class.

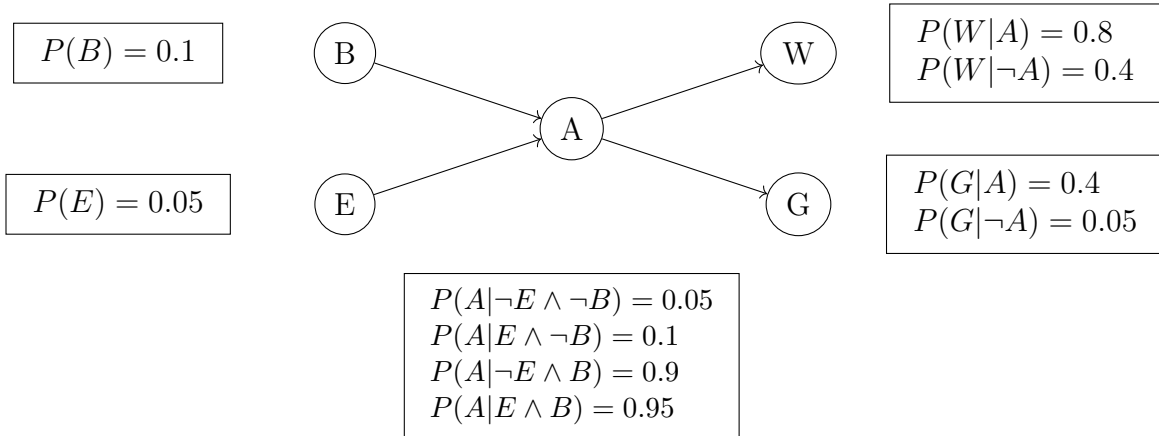


Figure 1: Bayesian Network 1

If we consider the variables in the order of W, G, A, B, E , we will derive the Bayesian network in Figure 2 below.




Figure 2: Bayesian Network 2

You will compute a few probabilities to explain why the structure of the network in Figure 2 is the way it is. You may do the calculations by hand or use the variable elimination algorithm you implemented in Question 1.

- (6 points) We chose W as the only parent of G because G is not independent of W . Calculate $P(G|W)$ and $P(G|\neg W)$ separately and show that they are not equal.
- (6 points) We chose A as the only parent of B because B is conditionally independent of W and G given A . Calculate $P(B|W \wedge G \wedge A)$ and $P(B|A)$ separately and show that they are equal.

3. (6 points) We could not have chosen W as the only parent of B because B is not conditionally independent of A and G given W . Calculate $P(B|A \wedge G \wedge W)$ and $P(B|W)$ separately and show that they are not equal.
4. (6 points) We could not have chosen A as the only parent of E because E is not conditionally independent of B given A . Calculate $P(E|A \wedge B)$ and $P(E|A)$ separately and show that they are not equal.