

6.4110/16.420
Representation, Inference and Reasoning in AI


Extra practice problems

Danger—Warning—Not carefully debugged

Old homework and exam questions from the archive

1 Bayesian Network

Consider the following Bayesian network:



- (a) Is it a polytree?
- (b) Is A independent of C?
- (c) Is C independent of E?
- (d) Is D independent of C?
- (e) Name a variable that, if it were an evidence variable, your answer to the question in part (b) would be different, or say that there is no such variable. (So, if your answer to (b) was that they are independent, then name a variable X for which A is not conditionally independent of C given X .)
- (f) Name a variable that, if it were an evidence variable, your answer to part (c) would be different, or say that there is none.
- (g) Name a variable that, if it were an evidence variable, your answer to part (d) would be different, or say that there is none.
- (h) If all the nodes are binary, how many parameters would be required to specify all the CPTs in this network? (Remember that if p is specified then it is not necessary to specify $1 - p$ as well.)
- (i) Give an expression for $\Pr(D | C)$ given probabilities that are stored in the CPTs. Don't include any unnecessary terms.

Solution:

- (j) What factor is created if we eliminate B first in the course of using variable elimination to compute $\Pr(A | G)$?

Solution:

- (k) What is the Markov blanket of B ?
- (l) Imagine that you're doing likelihood weighting to compute $\Pr(E = e | A = a)$. What weight would you have to assign to sample $\langle a, b, c, d, f, g \rangle$?

2 Surveillance

You are performing surveillance, trying to decide which destination in a harbor a particular submarine is headed toward. At each time step, the situation can be characterized by the following variables:

destination Which destination the submarine is headed to.

location The submarine's current location.

observation Your noisy observation (via underwater sensing) of the submarine's location

action The submarine's actions (speed and steering)

In addition, there is a variable, **type**, which encodes the type of the submarine (which is useful to know, because it affects the sub's speed and maneuverability).

- (a) Draw a dynamic Bayesian network diagram that describes this system. Only the **observation** variable is directly observable. Show how the values of the variables at the current time step depend on their values in the previous time step. Your model should be able to encode these relationships:
- Submarines tend to stick with the same destination, and don't frequently change which one they're aiming at.
 - The choice of action depends on relative position of the submarine and its destination.

Solution:

- (b) What would you change in your diagram in order to model the idea that some types of submarines have different ranges and that the distance to a location might affect its selection as a destination?

Solution:

- (c) If you have made three observations, O_1, O_2 , and O_3 , give an expression for the probability distribution over the destination at step 3:

$$\Pr(D_3 \mid O_1, O_2, O_3)$$

using only probabilities stored in the CPTs of the network (in part (a)).

Solution:

3 Logic

1. Entailment

Let S_1 and S_2 be sentences in propositional logic, let I_1 be the set of interpretations that make S_1 true, and let I_2 be the set of interpretations that make S_2 true. Assume that I_1 is a subset of I_2 . Mark all of the statements which *must* be true in this case:

- S_1 entails S_2 .
- S_2 entails S_1 .
- S_2 can be proven from S_1 with a sound proof system.

- S_2 can be proven from S_1 with a complete proof system.
- The sentence $(S_1 \rightarrow S_2)$ is satisfiable.
- The sentence $(S_1 \rightarrow S_2)$ is valid.
- The sentence $(S_2 \rightarrow S_1)$ is satisfiable.
- The sentence $(S_2 \rightarrow S_1)$ is valid.

2. **Propositional Proof**

Prove that the sentence

$$(P \vee Q \vee R) \wedge (\neg P \vee Z) \wedge (\neg Q \vee W) \wedge \neg R \wedge \neg Z \wedge \neg W$$

is unsatisfiable.

1. Explain the method you are going to use.

Solution:

2. Do the proof.

Solution:

4 Itinerary

We're all looking forward to when we can travel again! So, let's think about how to make a flight search engine... just like Travelocity only (much, much) smaller.

A flight is specified with a starting city and time and an ending city and time. We want to be able to answer queries that are specified with a starting city, starting time, final destination city and deadline time. The objective is to find a sequence of flights, starting in the starting city, leaving sometime at or after the starting time, and arriving at the final destination before (or at) the deadline.

Assume you have a database D consisting of tuples (c_o, t_o, c_d, t_d) specifying the city and time where the flight originates and the site and time where it arrives at its destination.

How should we formulate this as a path-search problem? Please specify:

1. State space

Solution:

2. Action space

Solution:

3. Successor function

Solution:

4. Goal test

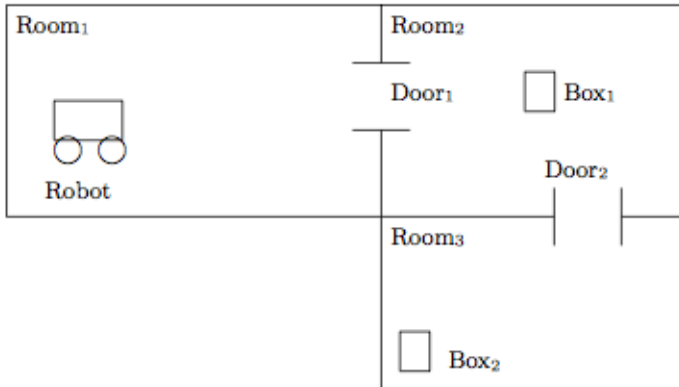
Solution:

Pat wants to find a way to get from Rome at time 1 to Istanbul at the earliest time possible. They propose to repeatedly call a path search algorithm on the formulation above, with different deadlines, to accomplish this goal. Sketch a strategy that Pat could use. As an additional challenge, try to minimize the number of times your strategy calls the search procedure.

Solution:

5 Shakey the Robot

Consider the following planning domain consisting of a robot pushing boxes between connected rooms.



We will represent this domain with the following symbols

- $b1, b2$ the two boxes; $r1, r2, r3$ the three rooms, $d1, d2$ the two doors.
- $open(?X)$ – door $?X$ is open.
- $in(?X, ?Y)$ – box $?X$ is in room $?Y$.
- $robin(?X)$ – the robot is in room $?X$.
- $join(?X, ?Y, ?Z)$ – door $?X$ joins rooms $?Y$ and $?Z$.

1. Specify the initial state shown in the figure. Assume that the doors start out closed.

Solution:

2. Give a STRIPS representation of the following action. Use the above symbols to specify sensible pre-conditions and effects for the action $pushthru(?B, ?R1, ?R2, ?D)$, meaning the robot pushes box $?B$ from room $?R1$ to room $?R2$ via door $?D$.


Solution:

3. We want an action to open a door $?D$. This action should work when the robot is in either of the rooms that the door joins. Give a STRIPS representation for this action; make sure that you have all the necessary facts in your initial state.

Solution:

6 The Deciding Factor

Consider the following directed graphical model:



Assume the variables are all binary and the CPTs are as follows:

| A | $P(B = 1)$ |
|-----|------------|
| 0 | 0.3 |
| 1 | 0.6 |

| A | $P(C = 1)$ |
|-----|------------|
| 0 | 0.9 |
| 1 | 0.2 |

| B | C | $P(D = 1)$ |
|-----|-----|------------|
| 0 | 0 | 0.9 |
| 0 | 1 | 0.1 |
| 1 | 0 | 0.1 |
| 1 | 1 | 0.9 |

| $P(A = 1)$ |
|------------|
| 0.3 |

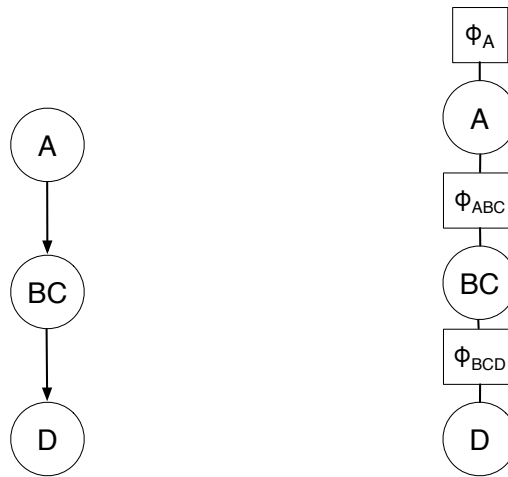
1. Draw its associated factor graph and specify the factors in terms of the CPTs given above.

Solution:

2. What algorithm is appropriate for exact inference on this model?

Solution:

3. Jody suggests converting the original directed graph to the following one, where BC is a random variable that can take on four possible values: 00, 01, 10, 11 which correspond to joint assignments to variables B and C from the original model. We also show its associated factor graph.



How does this directed graph compare in expressive power to the original one?
 More Less Same Briefly explain your answer.

Solution:

4. Provide tables for any factors in the factor graph from part 3 that differ from the one in part 1.

Solution:

5. Show how to use belief propagation on the factor graph in part 3 to compute $P(A \mid D = 0)$, by supplying formulas for each of the messages that is computed. Your expressions may use factor values and values of any previously computed messages. You do not need to do numeric computation.

- Message $\mu_{D \rightarrow \phi_{BCD}}(d)$

Solution:

- Message $\mu_{\phi_{BCD} \rightarrow BC}(b, c)$

Solution:

- Message $\mu_{BC \rightarrow \phi_{ABC}}(b, c)$

Solution:

- Message $\mu_{\phi_{ABC} \rightarrow A}(a)$

Solution:

- Message $\mu_{\phi_A \rightarrow A}(a)$


Solution:

- Final result: $\Pr(A = a \mid D = 0)$

Solution:

7 Bayesian network inference

Consider the Bayesian network below where all variables are binary.



1. What is the size of the largest CPT in this network?
2. What nodes can be ignored while computing $\Pr(H|M)$?
3. What is the time complexity of the problem of finding the elimination order that generates the smallest-size largest factor?
4. If you were computing $\Pr(P|B = b)$ for a very unlikely value of b , would you prefer importance sampling or Gibbs sampling? Why?
5. If you were computing $\Pr(B|P = p)$ for a very unlikely value of p , would you prefer importance sampling or Gibbs sampling? Why?

8 Planning


We have seen the standard “blocks world” domain in PDDL. It assumes that objects are all the same size and only one object can be on another object. In this problem, we want to extend the blocks world domain so that there are objects of two sizes: **big** and **small**. In this new domain, the stacking rules are as follows:

- One **big** object can be stacked on a **big** object, leaving no extra room.
- One **small** object can be stacked on another **small** object, leaving no extra room.
- No **big** object can be stacked on a **small** object.
- Two **small** objects can be stacked (in separate **stack** operations) on a **big** object.

We will use the following new predicates:

- (**big** ?x): object ?x is **big**
- (**small** ?x): object ?x is **small**
- (**zero-on** ?x): object ?x has no other objects on top of it
- (**one-on** ?x): object ?x has exactly one object directly on top of it
- (**two-on** ?x): object ?x has exactly two objects directly on top of it

In this problem, we will ignore the **on-table** predicate.



1.

Write a PDDL description of the state shown above, using predicates **on**, **arm-empty**, **holding**, **big**, **small**, **zero-on**, **one-on**, **two-on**.

Solution:

2. Here is a PDDL operator description and a state description.

| Operator description | State description |
|---|--|
| <pre> (:action unstack-one-big :parameters (?ob ?underob) :precondition (and (arm-empty) (two-on ?underob) (on ?ob ?underob) (zero-on ?ob)) :effect (and (holding ?ob) (not (arm-empty)) (one-on ?underob) (not (two-on ?underob)) (not (on ?ob ?underob)))) </pre> | <pre> (and (on a b) (on c b) (on e d) (small a) (small c) (small e) (big b) (big d) (arm-empty) (zero-on a) (zero-on c) (zero-on e) (one-on d) (two-on b)) </pre> |

How many ways are there to apply this operator to this state? For each one, provide a binding of the parameters in the operator.

Solution:

3. For one of the applicable operator instances you identified in the previous question, what is the state that results from applying the operator?


Solution:

4. Write a PDDL operator description that describes the operation of stacking a second block onto a big block that currently has a single small block directly placed on it.

Solution:

9 Bayesian networks and probabilistic inference


9.1 Bayesian network definitions



1. If each node is binary, how many parameters are needed to specify the distribution encoded by this network? (Assume we only need one parameter to define a distribution over two values). 9
2. Write a simple expression for $P(A = 1, B = 0, C = 1, D = 0)$, using only quantities explicitly represented in the conditional probability tables of the network.

$$P(A=1)P(C=1|A=1)P(D=0|C=1)P(B=0|A=1, D=0)$$

9.2 Factor graph



| A | B | C | $\phi_1(A, B, C)$ |
|---|---|---|-------------------|
| 0 | 0 | 0 | .3 |
| 0 | 0 | 1 | .2 |
| 0 | 1 | 0 | .7 |
| 0 | 1 | 1 | .8 |
| 1 | 0 | 0 | .1 |
| 1 | 0 | 1 | .3 |
| 1 | 1 | 0 | .9 |
| 1 | 1 | 1 | .7 |

| A | $\phi_2(A)$ |
|---|-------------|
| 0 | 1 |
| 1 | 3 |

| C | $\phi_3(C)$ |
|---|-------------|
| 0 | 3 |
| 1 | 3 |

1. Draw a Bayesian network graph that specifies the same distribution as this factor graph with the same number of parameters (but not necessarily the same parameter values).


Solution:

2. Specify parameters for the Bayesian network so that it specifies the same distribution as this factor graph.

Solution:

9.3 Computing messages

Consider performing belief propagation on the following simple factor graph:




| A | $\phi_1(A)$ |
|---|-------------|
| 0 | .3 |
| 1 | .7 |

| A | B | $\phi_2(A, B)$ |
|---|---|----------------|
| 0 | 0 | .1 |
| 0 | 1 | .6 |
| 1 | 0 | .1 |
| 1 | 1 | .2 |

What are the values of the following messages? Specify the entire message, including all values for the relevant variable(s).

1. $\mu_{1 \rightarrow A}$
2. $\mu_{A \rightarrow 2}$
3. $\mu_{2 \rightarrow B}$

9.4 Message passing structure



Our goal is to compute $P(C)$ in this factor graph **using the belief propagation algorithm**. List a legal sequence of messages that would be generated by belief propagation, specifying which variable or factor originates the message and which receives it. You do not need to calculate the actual numerical values.


Solution:

10 In search of...

The following assumptions apply to all the search problems

- Assume that the data structure implementations and successor state orderings are all such that *ties are broken alphabetically*. For example, a partial plan $S \rightarrow X \rightarrow A$ would be expanded before $S \rightarrow X \rightarrow B$; similarly, $S \rightarrow A \rightarrow Z$ would be expanded before $S \rightarrow B \rightarrow A$.
- All search algorithms are *graph* search (as opposed to tree search). In particular, when a node is expanded (popped from the agenda/frontier), if a node with this state has been previously expanded, then the node is pruned (discarded). Note that this is not the same as never *visiting* (adding to the agenda/frontier) a state twice.
- All the algorithms terminate when a node corresponding to a goal state is **expanded**, not when the node is added to the agenda/frontier.
- All edge costs are positive.
- When asked for a path, give your answers as a list of node names ‘*SADG*.’

Consider the state space search problem shown to the right. A is the start state and the shaded states are goals. Arrows encode possible state transitions, and numbers by the arrows represent action costs. Note that state transitions are directed; for example, $A \rightarrow B$ is a valid transition, but $B \rightarrow A$ is not. Numbers shown in diamonds are heuristic values that estimate the optimal (minimal) cost from that node to a goal.



For each of the following search strategies, give the path that would be returned, or write *none* if no path will be returned. Also, give all the states that are expanded in the order that they are expanded.

1. Depth-first graph search (ignores costs)

Solution:

2. Breadth-first graph search (ignores costs)

Solution:

3. Uniform cost graph search

Solution:

4. Greedy graph search


Solution:

5. A* graph search

Solution:

11 Admit it!

For the following questions, consider the search problem shown on the left. It has only three states, and three directed edges. A is the start node and G is the goal node. To the right, four different heuristic functions are defined, numbered I through IV.



| | $h(A)$ | $h(B)$ | $h(G)$ |
|-----|--------|--------|--------|
| I | 4 | 1 | 0 |
| II | 5 | 4 | 0 |
| III | 4 | 3 | 0 |
| IV | 5 | 2 | 0 |

For each heuristic function, circle whether it is admissible for the search problem given above.

| | Admissible? | | Consistent? | |
|-----|-------------|----|-------------|----|
| | Yes | No | Yes | No |
| I | | | | |
| II | | | | |
| III | | | | |
| IV | | | | |

Solution: II is the only inadmissible heuristic, as it overestimates the cost from B : $h(B) = 4$, when the actual cost to G is 3.

To check whether a heuristic is consistent, ensure that for all paths, $h(N) - h(L) \leq \text{path}(N \rightarrow L)$, where N and L stand in for the actual nodes. In this problem, $h(G)$ is always 0, so making sure that the direct paths to the goal ($A \rightarrow G$ and $B \rightarrow G$) are consistent is the same as making sure that the heuristic is admissible.

The path from A to B is a different story.

Heuristic I is not consistent: $h(A) - h(B) = 4 - 1 = 3 > \text{path}(A \rightarrow B) = 2$.


Heuristic III is consistent: $h(A) - h(B) = 4 - 3 = 1 \leq 2$

Heuristic IV is not consistent: $h(A) - h(B) = 5 - 2 = 3 > 2$

12 Hive mind

This hive of insects needs your help. You control an insect in a rectangular maze-like environment with dimensions $M \times N$, as shown to the right. At each time step, the insect can move into a free adjacent square or stay in its current location. All actions have cost 1.

In this particular case, the insect must pass through a series of partially flooded tunnels. Flooded squares are lightly shaded in the example map shown. The insect can hold its breath for A time steps in a row. Moving into a flooded square requires your insect to expend 1 unit of air, while moving into a free square refills its air supply.




1. Give a minimal state space for this problem (i.e. do not include extra information). You should answer for a general instance of the problem, not the specific map shown.

Solution:

2. Give the maximum size of your state space, in the general case of an $M \times N$ grid. $M \times N \times A$

13 Banana Nut

In the following Bayesian network



| A | P(A) |
|---|------|
| 0 | .9 |
| 1 | .1 |

| B | A | P(B A) |
|---|---|----------|
| 0 | 0 | .7 |
| 1 | 0 | .3 |
| 0 | 1 | .6 |
| 1 | 1 | .4 |

| C | A | P(C A) |
|---|---|----------|
| 0 | 0 | .9 |
| 1 | 0 | .1 |
| 0 | 1 | .5 |
| 1 | 1 | .5 |

1. What is $P(A = 1, B = 1)$

Solution:

2. What is $P(A = 1, B = 1, C = 1)$

Solution:


3. What is $P(B = 1 | A = 1)$

Solution:

4. What is $P(B = 1 | C = 1)$

Solution:

14 The Reverend Thomas



For each of the Bayesian networks above:

1. Write a simple expression for $P(A = 1, B = 0, C = 1, D = 0)$, using only quantities explicitly represented in the conditional probability tables of the network.

Solution:


2. Draw the corresponding factor graph.

Solution:

3. Indicate whether belief propagation on each of these corresponding factor graphs will correctly compute marginal probabilities.

Solution:


15 Factor it out



In the factor graph above, assuming the variables are all binary,

1. What are the variable involved in the potential φ_1 ?
2. What are the variable involved in the potential φ_3 ?
3. Mark all that are true:
 - The values in a factor have to sum to 1.
 - The values in a factor have to represent a conditional probability distribution.
 - The values in a factor have to be less than 1.
 - The values in a factor have to be non-negative.

16 Factor Fiction



1. Our goal is to compute $P(B)$ in this factor graph. List a legal sequence of message, specifying which variable or factor originates the message and which receives it. You do not need to calculate the actual numerical values (but it's good practice!).

Solution:

2. Repeat for computation of $P(C)$.


Solution:

3. Repeat for computation of $P(C|A = 1)$.

Solution:

17 Belief propagation

Consider the sum-product algorithm on this factor graph:



We will refer to factors by the set of variables they are connected to.

$$1. \phi_{ADE} = \begin{array}{|c|c|c|c|} \hline A & D & E & \\ \hline 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 2 \\ 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 3 \\ 1 & 0 & 0 & 2 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 4 \\ 1 & 1 & 1 & 3 \\ \hline \end{array} \quad \phi_{AB} = \begin{array}{|c|c|c|} \hline A & B & \\ \hline 0 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 0 & 2 \\ 1 & 1 & 4 \\ \hline \end{array} \quad \phi_A = \begin{array}{|c|c|} \hline A & \\ \hline 0 & 2 \\ 1 & 1 \\ \hline \end{array}$$

What message does A send to ϕ_{AB} ? Your answer should be a table of numbers.

Solution: