6.4110/16.420
Representation, Inference and Reasoning in AI

Midterm Exam

Solutions

October 30, 2023

- Answer the questions in the spaces provided on the question sheets. If you run out of room for an answer, continue on one of the blank pages in the back and label it clearly.
- You are permitted to use a single sheet of paper with notes on (both sides), and a calculator and a timer. If you use your phone for the calculator and timer, please restrict yourself to these functions.
- Try to avoid asking questions. If something seems confusing or ambiguous, write down what assumption you are making and continue.

Name: ______________________________

MIT email: ______________________________

Question	Points	Score
1	20	
2	20	
3	25	
4	35	
Total:	100	

1 Discrete bin packing

1. You are trying to pack several 2D rectangular objects into a 2D box. We'll assume that the objects and the box all have integer dimensions and the objects can only be placed in alignment with the grid, with no rotation.

   - You are given an overall grid of size 5×5. The top left corner is $(0, 0)$.
   - You have to place 4 objects (A, B, C, D), of sizes 1×5, 3×2, 4×1 and 1×1, by choosing the coordinates of their upper-left corner. Each coordinate can be one of $\{0, 1, 2, 3, 4\}$.
   - The objects must be entirely contained in the grid.
   - The objects must not overlap.

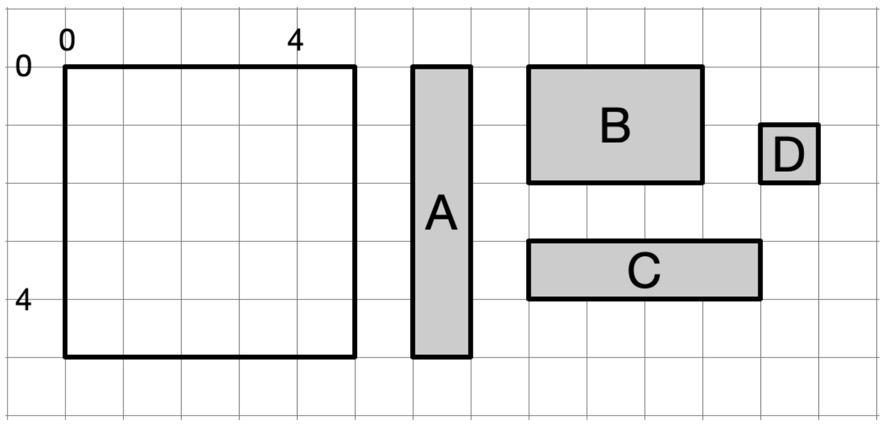

   (a) (6 points) List the constraints for this problem. You don't have to write them mathematically, but make clear how many individual constraint instances there are in a formal description of this problem that you could feed into a solver.

   Solution:

   - There are 4 constraints, one for each object, specifying that it has to be entirely contained in the box.
   - There are 3 + 2 + 1 constraints, one for each distinct pair of objects, specifying that they do not overlap.

   Total of 10 constraints.

(b) (6 points) List the variables and the sizes of their domains after any unary constraints have been applied.

> **Solution:**
>
> - A : $\{0, \ldots, 4\} \times \{0\}$ size = 5
> - B : $\{0, \ldots, 2\} \times \{0, \ldots, 3\}$ size = 3 x 4 = 12
> - C : $\{0, 1\} \times \{0, \ldots, 4\}$ size = 2 x 5 = 10
> - D : $\{0, \ldots, 4\} \times \{0, \ldots, 4\}$ size = 5 x 5 = 25

(c) (2 points) Assuming the unary constraints have been applied, what variable would be selected first for backtracking using the minimum-remaining-values heuristic? Explain.

> **Solution:** Object A has the smallest domain, so it would be placed first.

(d) (6 points) If we were to fix the upper-left corner of object C to be at location $(0, 2)$ and do forward checking, what would the domain of object B be? Provide a list of (x, y) pairs.

> **Solution:** $\{(0, 0), (1, 0), (2, 0), (0, 3), (1, 3), (2, 3)\}$

2 Propositional logic

2. Consider a domain with N propositional variables $P_1, \ldots, P_N$.
   (a) (2 points) How many possible models are there in this domain (as a function of N)?

   Solution: For each of the N propositional variables, you have 2 choices: True or False. Since each variable can be set independently of the others, you multiply the number of choices for each variable to find the total number of models. This gives us 2^N possible models in the domain.

   (b) (10 points) Let α be the following:

$$\begin{aligned} P_1 \wedge P_2 \Rightarrow P_3 \\ \neg P_3 \\ P_1 \end{aligned}$$

   Does α entail $\neg P_2$? Prove your answer using resolution refutation.

   Solution: We can prove α entails $\neg P_2$ with resolution refutation. We will add the negation of $\neg P_2$ to our set of clauses, and convert each clause in α to be in conjunctive normal form. This gives us the following proof:

   1. $\neg P_1 \vee \neg P_2 \vee P_3$
   2. $\neg P_3$
   3. P_1
   4. P_2
   5. $\neg P_1 \vee \neg P_2$ due to (1, 2)
   6. $\neg P_2$ due to (3, 5)
   7. due to (4, 6)

   The empty set gives us a contradiction, proving that the above statement.

(c) (4 points) Now let α_N be a conjunction of $\neg P_3$, P_1, and the formulas

$$P_i \wedge P_{i+1} \Rightarrow P_{i+2}$$

for all $i \in \{1, \ldots, N-2\}$ for some arbitrary positive integer N. Let's consider whether α_N entails $\neg P_2$. Would your previous proof in (b) apply unchanged if we use α_N instead of α? Explain.

Solution: The exact solution to part (b) would apply. We can use the same clauses in the previous solution to yield the empty set, and the additional formulas are not necessary in the proof.

(d) (4 points) Approximately how many legal resolution proof steps are available at the beginning of an automated theorem-proving search process when using α_N?

Solution: $O(N)$, because each clause has the form $\neg P_i \vee \neg P_{i+1} \vee P_{i+2}$. So, there is an opportunity to resolve each such clause with the next one.

3 LineLand

3. Consider a fantastic creature, living on the infinite integer number line $-\infty, \ldots, \infty$. Its state is described by (x, r) where x is the coordinate of its center on the line and r is its radius. So, it occupies the interval $[x - r, x + r]$. It can move around on the line (change x), and depending on where it is and on what actions it takes, it may grow or shrink (change r). Generally, when it moves, it grows, but if it gets into the shrink zone, it will shrink.

   In general, its action can be described with an integer value dx. In any given problem, there is a special "shrink" zone, $[Z_{lo}, Z_{hi}]$, which is an interval of the line. The transition model is as follows:

$$T((x,r), dx) = \begin{cases} (x + dx, \max(0, r - 1)) & \text{if } Z_{lo} \leqslant x \leqslant Z_{hi} \\ (x + dx, r + |dx|) & \text{otherwise} \end{cases}$$

   The creature's goal is specified by a goal region $[g_{lo}, g_{hi}]$, and the goal is for the *entire* creature (the whole interval $[x - r, x + r]$) to be contained in the goal region, so that

$$g_{lo} <= x - r \qquad \text{and} \qquad x + r <= g_{hi} \enspace .$$

   All actions have a cost of 1.

   We'll consider a specific instance of this type of problem in which:

   - The agent's actions (possible values of dx) are $(-2, -1, 0, 1, 2)$.
   - The initial state, s_0, is $(0, 1)$.
   - The goal region is $[3, 7]$ (so $g_{lo} = 3$ and $g_{hi} = 7$).
   - The shrink zone is $[4, 20]$ (so $Z_{lo} = 4$ and $Z_{hi} = 20$).

   (a) (6 points) What is an optimal solution to this problem? If there is none, explain why. If there are multiple optimal solutions, any one is fine. Provide the list of actions.

   > **Solution:** 2, 2, 1, 0, 0
   > (2, 3) ; (4, 5); (5, 4); (5, 3); (5, 2)

   (b) (5 points) Which of the following are relaxations of this problem, which could be applied for <u>any initial state, goal, and shrink zone</u>? In each case, assume that the rules for transitions and goals stay the same except for the stated change. (Mark all that apply)

   √ **The creature's shape never grows.**
   ◯ The creature's shape never shrinks.
   √ **Only the creature's center needs to be inside the goal.**
   ◯ The creature can move at most distance 1 per step.
   √ **The creature can move any number of squares per step.**

(c) (5 points) Provide heuristic that is admissible for this problem instance, and has the property that $h(s) > 1$ for at least one state s.

> **Solution:** $h(x, r) = (1/2) \min(|3-x|, |7-x|)$ is an admissible heuristic that would still guarantee that an optimal solution is found, but encourages the agent to explore directions closer to the goal state.

(d) Pat suggests adding a search-pruning rule, so that if uniform cost search (UCS) has visited a state (x, r) and then it reaches a new state (x, r'), if $r' > r$, we ignore the new state and do not add it to the frontier.

i. (3 points) Given that UCS has visited (x, r) before (x, r'), what can we infer about the cost of the lowest-cost path from the root to reach (x, r) compared to the cost of the lowest-cost path from the root to reach (x, r')?

> **Solution:** The path to (x, r) is cheaper or the same cost.
> Note: We accepted "nothing" as an answer, if an explanation was given detailing that we need to know that a node has been **expanded** to determine that the cost is cheaper. (This question really should have been written with expanded rather than visited.)

ii. (3 points) Would using Pat's pruning rule prevent UCS from finding an optimal solution to this problem?

> **Solution:** It does not prevent finding an optimal solution. There cannot be a cheaper path from (x, r') to a goal state than from (x, r) to a goal state. So, given that the path from the root to (x, r) is cheaper, we will never need to consider paths through (x, r').
> Note: Analogous to part (i), we also accepted an explanation for yes with similar reasoning.

(e) (3 points) There's another interpretation of the model we have been studying. Rather than a creature that shrinks and grows, we could think of our agent as a point robot that has set-based uncertainty about its position on the line (initial uncertainty set is s_0), and a non-deterministic transition function that might have any of a set of possible states as its result. We would like to find a conformant plan (a plan that is guaranteed to reach the goal no matter what happens) for this agent, where the goal for the point robot to be anywhere inside the interval $[g_{lo}, g_{hi}]$.

Explain what, if anything, we would need to change about this planning problem formulation to find a conformant plan.

> **Solution:** Nothing needs to be changed. The "size" of the agent is a representation of its uncertainty. A conformant plan can be found for this problem formulation.

4 Discrete probabilistic graphical models

4. Consider a family of probabilistic graphical models with N random variables and $N-2$ factors of the following form:

v_i	v_{i+1}	v_{i+2}	$\phi(v_i, v_{i+1}, v_{i+2})$
0	0	0	1.0
0	0	1	0.0
0	1	0	1.0
0	1	1	0.0
1	0	0	1.0
1	0	1	0.0
1	1	0	0.0
1	1	1	1.0

(a) (5 points) Draw the associated factor graph for $N = 4$.

Solution:

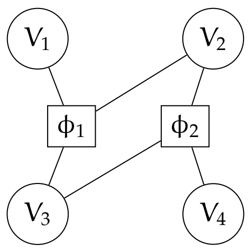

(b) (5 points) This same joint distribution can also be specified using a Bayesian network, where the factors defined above provide some of the required CPTs. Draw such a Bayesian network for $N = 4$.

Solution:

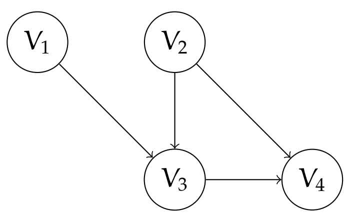

V_2	V_3	$P(V_4 \mid V_2, V_3)$
0	0	0
1	0	0
0	1	0
1	1	1

V_1	V_2	$P(V_3 \mid V_1, V_2)$
0	0	0
1	0	0
0	1	0
1	1	1

(c) (3 points) In a network with $N = 10$ what is the Markov blanket of V_5?

> **Solution:** The Markov blanket are the adjacent variables of V_5. V_5 is involved in the factors $\phi(V_3, V_4, V_5), \phi(V_4, V_5, V_6), \phi(V_5, V_6, V_7)$. Therefore, the blanket is $\{V_3, V_4, V_6, V_7\}$.

(d) (3 points) Can you apply belief propagation to get the exact marginals on this factor graph? Explain briefly.

> **Solution:** The graph is loopy, so we cannot apply standard belief propagation.

(e) (5 points) Gibbs sampling will not give a reasonable approximation of the marginals on this factor graph. Explain, with a concrete example, what the difficulty is.

> **Solution:** No. If we initialize in state $\mathbf{v} \in \{000, 010, 100\}$ then the state 111 is unreachable.

(f) (4 points) What is the maximum likelihood assignment of values to the pair A, B?

> **Solution:** $P(A, B) = \frac{1}{Z}\phi_1(a, b)\phi_2(a)\phi_3(b)$

a	b	$\phi_1(a, b)\phi_2(a)\phi_3(b)$
0	0	5*2*1=10
0	1	5*4*2=40
1	0	3*10*1=30
1	1	3*5*2=30

> So we have that the max likelihood assignment is $A = 0, B = 1$.

(g) (10 points) Let's think about a more concrete factor graph, with binary variables A and B and the following factors:

a	$\phi_1(a)$
0	5
1	3

a	b	$\phi_2(a,b)$
0	0	2
0	1	4
1	0	10
1	1	5

b	$\phi_3(b)$
0	1
1	2

Use the sum-product message passing method to compute the marginal distributions $P(A)$ and $P(B)$. Choose node A as the root. Write down the values of all the messages as well as the final result. It's fine to leave values in fractional form.

Solution:

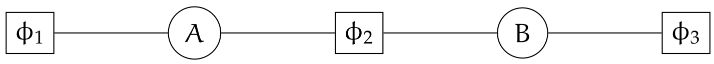

Choose (arbitrarily) A, as the root. We may first calculate $P(A)$ by calculating the messages from leafs to root as follows:

$$P(A) = \frac{1}{Z}\mu_{\phi_1 \to A}(a)\mu_{\phi_2 \to A}(a)$$

$$\mu_{\phi_1 \to A}(a) = \phi_1(a) = \begin{cases} 5 & a = 0 \\ 3 & a = 1 \end{cases}$$

$$\mu_{\phi_2 \to A}(a) = \sum_b \phi_2(a,b)\mu_{B \to \phi_2}(b) = \sum_b \phi_2(a,b)\phi_3(b) = \begin{cases} 10 & a = 0 \\ 20 & a = 1 \end{cases}$$

$$P(A) = \frac{1}{Z}\begin{cases} 5 * 10 & a = 0 \\ 3 * 20 & a = 1 \end{cases}$$

$$= \begin{cases} \frac{5}{11} & a = 0 \\ \frac{6}{11} & a = 1 \end{cases}$$

Now, we can calculate the marginal distribution for variable B as follows:

$$P(B) = \frac{1}{Z}\mu_{\phi_2 \to B}(b)\mu_{\phi_3 \to B}(b)$$

$$\mu_{\phi_3 \to B}(b) = \phi_1(b) = \begin{cases} 1 & b = 0 \\ 2 & b = 1 \end{cases}$$

$$\mu_{\phi_2 \to B}(b) = \sum_a \phi_2(a,b)\mu_{A \to \phi_2}(a) = \sum_a \phi_2(a,b)\phi_1(a) = \begin{cases} 40 & b = 0 \\ 35 & b = 1 \end{cases}$$

$$P(B) = \frac{1}{Z}\begin{cases} 1 * 40 & b = 0 \\ 2 * 35 & b = 1 \end{cases}$$

$$= \begin{cases} \frac{4}{11} & b = 0 \\ \frac{7}{11} & b = 1 \end{cases}$$

Extra space

Extra space

End of exam.