

- You have approximately 110 minutes.
- The exam is open book, open calculator, and open notes.
- In the interest of fairness, we want everyone to have access to the same information. To that end, we will not be answering questions about the content or making clarifications.
- For multiple choice questions,
  - means mark **all options** that apply
  - means mark a **single choice**

First name	
Last name	
SID	

**For staff use only:**

Q1. Potpourri	/15
Q2. MangoBot Human Detector	/15
Q3. Adventurers Assemble	/14
Q4. It's So Logical!	/20
Q5. Ball Games	/16
Q6. Bayes Nets and Inference	/20
Total	/100

THIS PAGE IS INTENTIONALLY LEFT BLANK

# Q1. [15 pts] Potpourri

(a) Decide whether the following statements are *True/False*.

(i) [1 pt] Consider a search problem in which the state space includes variable  $x$ . Then, any quantities that are a deterministic function of  $x$  don't need to be included in the state space.

True  False

This is true because adding the extra variable  $y = f(x)$  would not change the number of states: for every possible state  $(x, y, \dots)$  consistent with  $y = f(x)$  there is exactly one state  $(x, \dots)$  and vice versa.

(ii) [1 pt] Without any constraints on the utility value, it's impossible to perform pruning on an Expectimax tree.

True  False

This is true because both max and chance nodes are monotonically nondecreasing in their argument values. So, while a partially evaluated max node has a lower bound (e.g.,  $u = \max(2, 7, x) \geq 7$ ), it can never have an upper bound, and the same goes for chance nodes (e.g.,  $E(0.5u + 0.5v) \geq 7$  if  $u \geq 7$  and  $v \geq 7$ ). Put another way, if there's a finite probability of reaching a leaf node, that leaf node might have a value so large as to make the action sequence leading to it optimal. So pruning it is impossible.

(iii) [1 pt] If  $N_a > N_b$ , then running Monte Carlo Tree Search with  $N_a$  rollouts will always yield better decisions than with  $N_b$  rollouts.

True  False

Although as  $N \rightarrow \infty$ , MCTS will yield the optimal decision, for finite  $N$  the stochastic nature of MCTS means that a worse outcome may occur.

(iv) [1 pt] If one fixes temperature  $T$  to be 0, simulated annealing will reduce to standard (steepest-ascent) hill climbing.

True  False

This will reduce to first-choice hill climbing, as introduced in discussion. Simulated annealing picks a random successor and then decides whether to accept the move. At  $T=0$ , it will accept the first successor that improves on the current state.

(v) [1 pt] Given that a Bayes Net expresses a joint distribution over its variables, there can be conditional independence relationships implied by the joint distribution that are not asserted by the Bayes net topology.

True  False

There can also be conditional independence relationships that are captured by the CPTs in the Bayes Nets.

(vi) [1 pt] Assume that we have  $P(X_t | e_{1:t})$  for a standard HMM model. Then the cost of running the forward algorithm to calculate  $P(X_{t+1} | e_{1:t+1})$  is  $\mathcal{O}(t|X|)$ , where  $|X|$  is the number of states.

True  False

It should be  $\mathcal{O}(|X|)$  because the cost of running forward algorithm for 1 step does not involve  $t$ .

(b) [4 pts] Below is a list of task environments. For each of the sub-parts, choose all the environments in the list that falls into the specified type.

**A:** The tic-tac-toe game

**B:** Self driving cars

**C:** The "N-queens" search problem (discussed in the local search lecture)

**D:** The "Weather forecast" problem (predict weather from measurements of wind, pressure, humidity, etc.)

Which of the environments are *dynamic*?  A  B  C  D

Which of the environments are *fully observable*?  A  B  C  D

Which of the environments are *stochastic*?  A  B  C  D

Which of the environments have *known physics*?  A  B  C  D

"Dynamic" means we need to make a time-sensitive decision. Driving and weather prediction are time sensitive. "Fully observable" means we can accurately measure all the information we need to make the prediction. Driving and weather prediction are not fully observable. "Stochastic" means the transition and/or sensor model are not deterministic. Driving and weather prediction are stochastic. For "known physics", i.e., does the agent know the transition model, this is certainly true for A and C, and not for B. We allowed both answers for D, weather forecasting, because a decent amount of the physics is known, but there are still gaps in our understanding of important phenomena such as cloud formation.

(c) [2 pts] Consider a search problem with a set of goal states  $S$ . Define  $h^*(x)$  to be the shortest(optimal) distance from state  $x$  to a goal state  $s \in S$ . Using heuristic function  $h$ , which of the following statements are true? Select all that apply.

Assume that  $h(x) \geq 0$  for all  $x$ .

- Given  $h(x) > h^*(x)$  for some state  $x$ , A\* graph search cannot find the optimal solution.
- Given  $h$  is an admissible heuristic, A\* graph search cannot find the optimal solution.
- Given  $h$  is a consistent heuristic, running A\* graph search will find the optimal path to every goal state.
- Given  $h$  is an admissible heuristic, there exists a constant  $c > 0$  such that  $c * h$  becomes a consistent heuristic.
- None of the Above

Even with a bad heuristic (i.e. inadmissible, inconsistent), A\* graph search *might* still find the optimal path. A consistent heuristic will **guarantee** A\* graph search to find the optimal path, but that doesn't mean that a bad heuristic will guarantee a suboptimal path. Therefore, 1,2 are wrong.

3 is wrong because A\* graph search will return the optimal path to the goal it finds, not to every goal.

4 is correct because if there exists an edge  $A \rightarrow C$  where  $h(A) - h(C) > cost(A, C)$ , we can always multiply the heuristic with a small enough constant  $c \leq 1$  such that  $c * (h(A) - h(C)) < cost(A, C)$ . Such a heuristic will remain admissible and will still have value 0 for goal states.

(d) [3 pts] We propose a variant of the depth-first graph search algorithm as follows: we record the number of times we have visited each state in the state space, and we will not expand a state if we have visited the state at least  $k$  times, where  $k$  is a positive integer greater than 1. (When  $k = 1$ , this is equivalent to standard depth-first graph search). Which of the following statements are true? Select all that apply.

- The new algorithm is complete, while standard depth-first tree search is not.
- Given that there are no cycles in the state-space graph, the worst-case time complexity of the new algorithm is the same as that of standard depth-first tree search.
- The worst-case space complexity of the new algorithm is the same as that of standard depth-first tree search.
- The worst-case space complexity of the new algorithm is the same as that of standard breadth-first tree search.
- None of the Above


The proposed algorithm operates in a very similar way to depth-first graph search, reaching the same set of nodes but possibly repeating some node-generation steps up to  $k$  times each.

- Completeness: false. Although the algorithm cannot get into infinite loops like depth-first tree search, it still fails in state spaces where there are infinite paths of non-identical states. (Recall that standard depth-first graph search is incomplete for the same reason.)
- Time complexity with no cycles: false. The absence of cycles does not mean the same state cannot be reached by multiple paths (see section 3.3.3). Redundant paths can cause a tree search to generate exponentially many more nodes than a graph search. The proposed algorithm generates each node at most  $k$  times, so the time complexities are not the same.
- Space complexity compared to depth-first tree search: false. A graph search stores all the nodes it visits, which might be  $\mathcal{O}(b^m)$ , whereas a tree search is linear space  $\mathcal{O}(bm)$ , where  $m$  is the maximum possible depth.
- Space complexity compared to breadth-first tree search: false. A graph search stores all the nodes it visits, which might be  $\mathcal{O}(b^m)$ , whereas breadth-first tree search is  $\mathcal{O}(b^d)$ , where  $d$  is the depth of the shallowest solution, so the worst-case space complexities of the two algorithms are not the same.

So the correct answer is none of the above.

## Q2. [15 pts] MangoBot Human Detector

Your startup company MangoBot wants to build robots that delivers packages on the road. One core module of the robot's software is to detect whether a human is standing in front of it. We model the presence of humans with a Markov model:



where  $H_t \in \{0, 1\}$  corresponds to a human being absent or present respectively. The initial distribution and the transition probabilities are given as follows:

$H_0$	$P(H_0)$
0	$p$
1	$1 - p$

$H_t$	$H_{t+1}$	$P(H_{t+1} H_t)$
0	0	0.9
0	1	0.1
1	0	0.8
1	1	0.2

(a) Express the following quantities in terms of  $p$ :

(i) [1 pt]  $P(H_1 = 1) = -0.1p + 0.2$

$$P(H_1 = 1) = P(H_1 = 1, H_0 = 0) + P(H_1 = 1, H_0 = 1) = P(H_0 = 0)P(H_1 = 1|H_0 = 0) + P(H_0 = 1)P(H_1 = 1|H_0 = 1) = 0.2(1 - p) + 0.1p = 0.2 - 0.1p$$

(ii) [1 pt]  $\lim_{t \rightarrow \infty} P(H_t = 0) = 8/9$

As  $t \rightarrow \infty$ , the system converges to the stationary distribution  $\pi$ , which satisfies  $\pi = T^T \pi$ , where  $T$  is the transition probability matrix  $\begin{bmatrix} 0.9 & 0.1 \\ 0.8 & 0.2 \end{bmatrix}$ . Assume  $\pi = [q \quad (1 - q)]^T$  and solve for  $q$  in  $\pi = T^T \pi$  gives  $q = 8/9$ .

(b) The first-order Markov assumption in the model above can be inaccurate in real-world situations. Some potential ways to improve the model are listed below. For each option, determine whether it is possible to rewrite the process as a first-order Markov process, potentially with a different state representation.

(i) [2 pts]  $H_t$  depends not only on  $H_{t-1}$  but also on  $H_{t-2}$ .  Yes  No

We can make the state  $S_t = \{H_t, H_{t-1}\}$ , then  $S_{t+1}$  is conditionally independent of  $S_{t-1}$  given  $S_t$  for all  $t$ , which means the sequence  $S_t$  satisfies the Markov property.


(ii) [2 pts]  $H_t$  depends not only on  $H_{t-1}$  but also on  $H_{t-2}, H_{t-3}, \dots, H_{t-k}$  for some fixed  $k \geq 3$ .  Yes  No

Same basic argument except now  $S_t$  keeps a history of the  $k$  preceding time steps and the sequence  $S_t$  satisfies the Markov property. Note that the state-space size is now exponential in  $k$ .

(iii) [2 pts]  $H_t$  depends not only on  $H_{t-1}$  but also on  $H_{t-2}, H_{t-3}, \dots, H_1, H_0$ .  Yes  No

Since now  $H_t$  depend on a variable number of previous time steps, the trick we used in the previous two parts doesn't work any more. There is no finite state representation that maintains the Markov property.


To make things simple, we stick to the original first-order Markov chain formulation. To make the detection more accurate, the company built a sensor that returns an observation  $O_t$  each time step as a noisy measurement of the unknown  $H_t$ . The new model is illustrated in the figure, and the relationship between  $H_t$  and  $O_t$  is provided in the table below.



$H_t$	$O_t$	$P(O_t H_t)$
0	0	0.8
0	1	0.2
1	0	0.3
1	1	0.7

(c) Based on the observed sensor values  $o_0, o_1, \dots, o_t$ , we now want the robot to find the most likely sequence  $H_0, H_1, \dots, H_t$  indicating the presence/absence of a human up to the current time.

(i) [2 pts] Suppose that  $[o_0, o_1, o_2] = [0, 1, 1]$  are observed. The following "trellis diagram" shows the possible state transitions. Fill in the values for the arcs labeled A, B, C, and D with the product of the transition probability and the observation likelihood for the destination state. The values may depend on  $p$ .



A:  $0.8p$ ; B:  $0.3(1 - p)$ ; C:  $0.9 * 0.2 = 0.18$ ; D:  $0.2 * 0.7 = 0.14$ .

(ii) [3 pts] There are two possible most likely state sequences, depending on the value of  $p$ . Complete the following (Write the sequence as "x,y,z" (without quotes), where x, y, z are either 0 or 1):  
Hint: it might be helpful to complete the labelling of the trellis diagram above.

• When  $p < \boxed{0.25}$ , the most likely sequence  $H_0, H_1, H_2$  is  $\boxed{1,0,0}$

• Otherwise, the most likely sequence  $H_0, H_1, H_2$  is  $\boxed{0,0,0}$

After filling out the full trellis diagram, we can easily observe that the sequence with largest probability given  $H_0 = 0$  is (0, 0, 0) and the sequence with largest probability given  $H_0 = 1$  is (1, 0, 0). (To see what is the sequence with largest probability, we run search from  $H_0 = 0$  to either  $H_2 = 0$  or  $H_2 = 1$ , but instead of adding the costs we multiply the probabilities.) Therefore the two possible most likely state sequences are (0, 0, 0), with probability  $0.8p * 0.18 * 0.18$ , and (1, 0, 0), with probability  $0.3(1 - p) * 0.16 * 0.18$ . Setting up the equation  $0.8p * 0.18 * 0.18 = 0.3(1 - p) * 0.16 * 0.18$  gives  $p = 0.25$  to be the threshold.

Note that this is a bit counter-intuitive since the observations suggest the exact opposite thing. However, in this problem the transition probabilities for 0 to 0 and 1 to 0 are so large that they dominates the computation.

(d) [2 pts] True or False: For a fixed  $p$  value and observations  $\{o_0, o_1, o_2\}$  in general,  $H_1^*$ , the most likely value for  $H_1$ , is always the same as the value of  $H_1$  in the most likely sequence  $H_0, H_1, H_2$ .  True  False

The maximum likelihood estimation (MLE) for a single variable is in general not the same as the value of that variable in the most likely sequence estimation (MLSE). For example, in this problem, when  $p = 0.3$ , the MLE for  $H_0$  is 1, but the value of  $H_0$  in the MLSE is 0. There exists similar examples for  $H_1$ .

### Q3. [14 pts] Adventurers Assemble

Chuck and Nancy are exploring a cave off the coast of Maine when Chuck gets separated from Nancy. He decides to model his situation as a search problem to try to find Nancy.

Chuck draws out the cave as an  $X$  by  $Y$  grid on his magical map, and keeps track of his current location as  $Z$ . He also marks the locations of  $N$  slime monsters within the cave as  $S_1, \dots, S_N$ . The slimes are stationary but very dangerous, so Chuck must make sure to avoid them as he explores.

Chuck observes Nancy's current location on his map, and marks her location as  $T(s)$  for any given state  $s$ . On each turn Chuck chooses an action from the action set  $\{left, right, up, down\}$  and moves one square in the direction he chooses, although any action that would bump into a wall results in no movement. Similarly, Nancy chooses an action from the same action set and moves in that direction, which can be different from the action that Chuck takes. Once Chuck has reached the same location as Nancy on a given timestep, the search problem ends. Consider the following heuristic function, where  $s$  is the input state:

$$h(n) = \begin{cases} 0 & \text{where } Z(s) = T(s) \\ h_i(s) & \text{otherwise} \end{cases}$$

For the condition in each subpart, determine whether the following options for  $h_i(s)$  would allow  $h(s)$  to guarantee optimality for A\* tree search, A\* graph search, neither of them, or both of them given the condition. We will look at different proposals for what  $h_i(s)$  could be.  $M(A, B)$  is a function that returns the Manhattan distance between two locations in the grid  $A$  and  $B$ .

To be optimal for A\* tree search, the heuristic must be admissible. To be optimal for A\* graph search, the heuristic must be consistent. It's impossible for a heuristic to be optimal for only A\* graph search because consistency implies admissibility.

(a) Assume the cost of taking actions within a single timestep is 1 (i.e.: for a given timestep, Chuck moving one square and Nancy moving another square costs 1 in total).

(i) [2 pts]  $h_1 = M(Z(s), T(s))$

- |   |   |
|---|---|
| <input type="radio"/> optimality for A* tree search only  | <input checked="" type="radio"/> optimality for neither |
| <input type="radio"/> optimality for A* graph search only | <input type="radio"/> optimality for both               |

If Chuck and Nancy travel directly toward each other, they must both travel approximately  $M(Z, T)/2$  to meet each other in the middle. Therefore, the heuristic is not admissible, since the full Manhattan distance is an overestimate to the actual minimal cost of meeting Nancy in the middle.

(ii) [2 pts]  $h_2 = \left\lfloor \frac{M(Z(s), T(s))}{2} \right\rfloor$

- |   |  |
|---|--|
| <input type="radio"/> optimality for A* tree search only  | <input type="radio"/> optimality for neither         |
| <input type="radio"/> optimality for A* graph search only | <input checked="" type="radio"/> optimality for both |

Halving the Manhattan distances between Chuck and Nancy is admissible and consistent as a heuristic. At every location Chuck and Nancy can be at, they must both travel at minimum  $\left\lfloor \frac{M(Z(s), T(s))}{2} \right\rfloor$  to meet each other, making it admissible. The heuristic is also consistent because the heuristic value between states can decrease by at most 1. Using  $\lfloor \cdot \rfloor$  is actually conservative; if they are three steps apart, this gives a value of 1, whereas it takes a minimum of 2 time steps; so  $\lceil \cdot \rceil$  would also work.

(b) Chuck gets more and more tired with each subsequent action. At timestep  $t$ , the cost of Chuck's action  $a_t$  is 2 times the cost of taking action  $a_{t-1}$ , and the cost of taking his first action  $a_1$  is 1. The function  $C(a_t)$  returns the cost of Chuck taking action  $a_t$ . Nancy isn't tired, so the cost of Nancy's actions is always 1. (i.e.: for a given timestep  $t$ , Chuck moving one square and Nancy moving another square costs  $C(a_t) + 1$  in total.)

(i) [2 pts]  $h_3 = \sum_{j=k}^n C(a_j)$  where  $k =$  the current timestep,  $n = k + \left\lfloor \frac{M(Z(s), T(s))}{2} \right\rfloor$

- optimality for A\* tree search only                       optimality for neither  
 optimality for A\* graph search only                       optimality for both

This heuristic is admissible and consistent, since it computes the true cost over the number of moves it takes for Chuck to meet Nancy halfway minus the cost of Nancy moving as well, given this new action cost. Note that Chuck has no option to take a rest, even though this would be the best choice if it were available because Nancy doesn't get tired.

(ii) [2 pts]  $h_4 = n + \sum_{j=k}^n C(a_j)$  where  $k$  = the current timestep,  $n = k + \left\lfloor \frac{M(Z(s), T(s))}{2} \right\rfloor$

- optimality for A\* tree search only                       optimality for neither  
 optimality for A\* graph search only                       optimality for both

The only difference between this and the previous heuristic is the addition of  $n$ , which represents the cost of Nancy moving since the beginning. However, since  $n = k + \left\lfloor \frac{M(Z(s), T(s))}{2} \right\rfloor$ , this becomes an overestimate of the actual cost to proceed from the current state. If  $n = \left\lfloor \frac{M(Z(s), T(s))}{2} \right\rfloor$ , this heuristic would in fact be both admissible and consistent.

(c) Assume the cost of taking actions within a single timestep is 1, as in the very first subpart. Chuck must now destroy all  $N$  slimes in the cave before reuniting with Nancy. He destroys a slime  $S_i$  by moving onto the same square as it, i.e.,  $Z(s) = S_i$  for  $i \in [1, \dots, N]$ . Chuck records the index of slimes that he destroys by adding the index  $i$  to a list  $R$ . Only Chuck can destroy slimes; Nancy does not have that power. Even if he and Nancy move onto the same square, the search problem will not end until all slimes in the cave are eliminated, when  $R$  contains all integers from  $1, \dots, N$ . Note that  $len(R)$  returns the number of elements currently in the list  $R$ .

(i) [3 pts]  $h_5(s) = \begin{cases} 0 & \text{if } len(R) \geq N - 1 \\ \min_{i \notin R} \{M(Z(s), S_i)\} + \max_{j \notin R, j \neq i^*} \{M(S_{i^*}, S_j)\} & \text{otherwise} \end{cases}$

for  $i, j \in [1, \dots, N]$  and  $i^* = \arg \min_i \{M(Z(s), S_i)\}$

- optimality for A\* tree search only                       optimality for neither  
 optimality for A\* graph search only                       optimality for both

This heuristic calculates the Manhattan distance from Chuck's current location to the closest slime, plus the "diameter" of the remaining slimes, or returns 0 if there is only one slime or less remaining. This is always an underestimate to the actual goal, since Chuck must always take a single action to destroy every slime, and needs to at least traverse the Manhattan distance between the two furthest slimes from his location. However, this is an inconsistent heuristic if you consider the edge case where Chuck destroys the second to last slime. For example, suppose Slime 1 and Slime 2 are 10 squares apart from each other, and Chuck is adjacent to Slime 1. The heuristic value at this state would be 11, but the heuristic value of the next state once Chuck destroys Slime 1 is 0. Therefore, the heuristic overestimates the cost of moving between these two states, so it is only admissible and not consistent.

(ii) [3 pts]  $h_6(s) = \begin{cases} \left\lfloor \frac{M(Z(s), T(s))}{2} \right\rfloor & \text{if } len(R) \geq N \\ \max_i \{M(S_i, T(s))\} & \text{otherwise} \end{cases}$

for  $i \in [1, \dots, N]$  and  $i \notin R$

- optimality for A\* tree search only                       optimality for neither  
 optimality for A\* graph search only                       optimality for both

This heuristic returns the maximum Manhattan distance between Nancy and each slime when there are still slimes remaining, and returns the halfway distance between Chuck and Nancy once they've all been eliminated, which is neither admissible nor consistent. The maximum distance between Nancy and any living slime can overestimate the remaining cost because it doesn't take into account how close Chuck is to Nancy.



## Q4. [20 pts] It's So Logical!

- (a) [4 pts] Inspired by learning about propositional logic, you decided to write down some sentences about your two hobbies: **playing go and collecting kettlebells**, using symbols to stand for propositions. You now have this list of logic sentences, and you remember what the English meanings of the sentences were, but you forget the meaning of each symbol!

For each English sentence on the left, there is a corresponding logical sentence on the right, **but it is not necessarily the sentence next to it**. Your goal is to recover the meaning for each symbol. Please write down the English sentence that each logic symbol represents below.

English	Propositional logic
I will not buy a 24 kg kettlebell.	$\neg P \vee S$
I will buy an 8 kg kettlebell or a 24 kg kettlebell, but not both.	$(Q \vee R) \wedge (\neg Q \vee \neg R)$
If I play more go, then I will get better at it.	$(\neg Q \wedge \neg R) \vee S$
If I buy an 8 kg kettlebell or a 24 kg kettlebell, then I will get better at go.	$\neg R$

- (i) [1 pt] P: I will play more go.
- (ii) [1 pt] Q: I will buy an 8 kg kettlebell.
- (iii) [1 pt] R: I will buy a 24 kg kettlebell.
- (iv) [1 pt] S: I will get better at go.

- (b) [8 pts] Winston, Xuan, Yvonne, and Zeke just had an 80's exercise party, where they **did some exercise, wore sweatbands, and listened and danced to some music**. They have the following predicates in their vocabulary:

- $Suggested(p, e)$ : Person  $p$  suggested exercise  $e$ .
- $Tried(p, e)$ : Person  $p$  tried exercise  $e$ .
- $Danced(p, s)$ : Person  $p$  danced to song  $s$ .
- $Likes(p, s)$ : Person  $p$  likes song  $s$ .

They also abbreviate their names by the first letter.

After the party, they each tried to write down some English sentences in first-order logic, but they didn't always succeed. Below, you have the English sentences as well as the attendees' attempts to write them in first-order logic.

For each first-order logic sentence, choose whether it is invalid, valid but not equivalent to the English sentence, or equivalent to the English sentence.

- (i) [4 pts] Xuan tried every exercise that Zeke suggested.

$\exists e Suggested(Z, e) \wedge (\forall e' \vee Tried(X, e'))$	<input checked="" type="radio"/> Invalid	<input type="radio"/> Not equiv.	<input type="radio"/> Equiv.
$\forall e, e' \neg(e = e') \vee (\neg Suggested(Z, e') \vee Tried(X, e))$	<input type="radio"/> Invalid	<input type="radio"/> Not equiv.	<input checked="" type="radio"/> Equiv.
$\forall e Suggested(Z, e) \Rightarrow Tried(X, e)$	<input type="radio"/> Invalid	<input type="radio"/> Not equiv.	<input checked="" type="radio"/> Equiv.
$(\exists e Suggested(Z, e)) \wedge (\forall e Tried(X, e))$	<input type="radio"/> Invalid	<input checked="" type="radio"/> Not equiv.	<input type="radio"/> Equiv.

- (ii) [4 pts] Every song that anybody liked was danced to by someone who tried an exercise.

$\neg \exists s, p (Likes(p, s) \wedge Danced(p, s)) \vee (\forall e Tried(p, e))$	<input type="radio"/> Invalid	<input checked="" type="radio"/> Not equiv.	<input type="radio"/> Equiv.
$\forall s \exists p, e Likes(p, s) \Rightarrow Danced(p, s) \wedge Tried(p, e)$	<input type="radio"/> Invalid	<input checked="" type="radio"/> Not equiv.	<input type="radio"/> Equiv.
$\forall s (\exists p Likes(p, s)) \Rightarrow (\exists e, p' Danced(p', s) \wedge Tried(p', e))$	<input type="radio"/> Invalid	<input type="radio"/> Not equiv.	<input checked="" type="radio"/> Equiv.
$\forall s, p, p' Likes(p \wedge p', s) \Rightarrow ((Danced(p, s) \wedge \exists e Tried(p, e)) \wedge (Danced(p', s) \wedge \exists e Tried(p', e)))$	<input checked="" type="radio"/> Invalid	<input type="radio"/> Not equiv.	<input type="radio"/> Equiv.

- (c) [8 pts] You have the following sentence, and you want to know whether or not it is satisfiable:

$$S : (A \vee \neg B \vee D) \wedge (\neg A \vee B \vee E) \wedge (A \vee \neg C \vee \neg F) \wedge (B \vee C \vee F) \wedge (\neg C \vee \neg D \vee \neg E) \wedge (D \vee E \vee F).$$

First, you must choose an algorithm to use.

- (i) [1 pt] Which algorithm is best suited for your task?  
 DPLL     Propositionalization     Forward chaining

Propositionalization is for inference in first-order logic, and forward chaining only works for conjunctions of definite clauses (and checks entailment, not satisfiability).

Now that you've picked your SAT-solving algorithm, you have to run it. Luckily, you know how to do this, because it's essentially a variant of another algorithm you've studied earlier in the course.

- (ii) [1 pt] Which other algorithm that we've studied is your SAT-solving algorithm a variant of?  
 Alpha-beta pruning     Simulated annealing     Depth-first search  
 Breadth-first search     A\* search

DPLL is a depth-first search: each node represents assignments to some variables, and to go deeper you assign a value to an as-yet-unassigned variable. DPLL pursues each branch to termination before backtracking to try the next one. The primary difference from standard DFS is that at each level of the tree it branches on values for exactly one variable and does not try a different variable. This is because variable assignment is permutable.

During the execution of the algorithm, you have set variable  $A$  to be False and variable  $B$  to be True. Answer the following questions:

- (iii) [1 pt] Which variable is now a pure literal?  C     D     E     F

$C$  only appears in positive form in the clause  $B \vee C \vee F$ . But since  $B$  was set to True, that clause is already satisfied, and  $C$  appears in negative form in all remaining clauses (and no other variable does).

- (iv) [1 pt] What value do you set that variable to in this case?  True     False

Since  $C$  appears only negated, we can satisfy every clause it appears in by setting  $C$  to be False.

- (v) [1 pt] Which of the original clauses is now a unit clause? (Note that it was already a unit clause before you dealt with the pure literal)

- $A \vee \neg B \vee D$   
  $\neg A \vee B \vee E$   
  $A \vee \neg C \vee \neg F$   
  $B \vee C \vee F$   
  $\neg C \vee \neg D \vee \neg E$   
  $D \vee E \vee F$

Since  $A$  was set to False and  $B$  was set to True,  $A \vee \neg B \vee D$  is equivalent to  $\text{False} \vee \text{False} \vee D$ , which is equivalent to  $D$ . In no other clause is only one variable 'remaining' after getting rid of literals that have been set to False.

- (vi) [1 pt] What do we do once we find this unit clause?

- Terminate, since sentence is unsatisfiable  
 Terminate, return satisfying assignment for the sentence  
 Satisfy the clause by setting the remaining variable **in the clause** to True  
 Satisfy the clause by setting the remaining variable **in the clause** to False

The thing you do after finding a unit clause is setting a variable to satisfy it, not terminate. Since the unit clause is  $D$ , we set  $D$  to be True.

- (vii) [1 pt] Is there an assignment to the variables that satisfies  $S$  where  $A$  is set to False and  $B$  is set to True?

- Yes     No

Yes:  $A$  False,  $B$  True,  $C$  false, and  $D$  true is enough to satisfy  $S$ , however we set  $E$  and  $F$ .

- (viii) [1 pt] Reflecting on your journey, you notice that resolving the pure literal didn't add any more unit clauses. Can resolving a pure literal ever add a new unit clause?  Yes     No

No. Resolving a pure literal just eliminates clauses from consideration (since they're satisfied by the value we set the pure literal to). It never sets any literal to be False, and therefore never reduces the number of literals left in any clause (except by satisfying that clause), which would be needed to turn a non-unit clause into a unit clause.

## Q5. [16 pts] Ball Games

Alice is playing a ball game with Bob. There are 3 boxes in front of them, each containing 3 balls, and each ball has a score. Alice first selects a box from the 3 boxes, and then Bob takes a ball from the box selected by Alice. In Figure 1, nodes  $B_1$ ,  $B_2$  or  $B_3$  are boxes, and nodes labeled C through K represent individual balls and their scores. Unless otherwise specified, Alice's objective is to maximize the score of the ball that is eventually chosen, and Bob's objective is to minimize Alice's score. Assume both players always act optimally for their goals.




Figure 1

- (a) (i) [2 pts] In the blanks below, fill in the **labels** (not the numerical values) of the balls selected for nodes A,  $B_1$ ,  $B_2$  and  $B_3$ . For example, if Alice's optimal move is to select the box  $B_1$ , and Bob selects the ball C, then  $A=C$ ,  $B_1=C$ .

A = ,  $B_1$  = ,  $B_2$  = ,  $B_3$  = .

Bob is a minimizer, so he will choose E, H, K for  $B_1$ ,  $B_2$ ,  $B_3$  respectively. Alice is a maximizer, so she will choose the maximum among E, H, K, which is E.

- (ii) [1 pt] Let's apply alpha-beta pruning for the tree. Please choose all the **child nodes** of the branches that could be pruned in the search tree.   $B_1$    $B_2$    $B_3$   C  D  E  F  G  H  I  J  K

We cannot prune anything in the  $B_2$  branch because  $\min(F, G) = 8 > \min(C, D, E) = 6$ . We can prune K because after visiting J we know that the value at  $B_3$  is  $\leq 3$ , which must be smaller than the value at  $B_1$  ( $= 6$ ) and will never get picked by Alice.

- (iii) [3 pts] Assume that you can re-allocate all the balls to different boxes as long as in the end, each box has exactly 3 balls in it. Could you find a way to re-allocate the balls so that when we apply alpha-beta pruning, the right branch of  $B_2$  as well as the middle and the right branch of  $B_3$  will be pruned?

Please list all balls in the order of their new positions in your solution. For example, if your solution is not to move any ball, then your answer is "C,D,E,F,G,H,I,J,K". When multiple solutions exist, break ties alphabetically. For example, if there are two solutions starting with "C,D,..." and "D,C,..." , then your answer should be "C,D,...". If there is no solution, please write down "None".

C,D,E,F,H,G,J,I,K

The reasoning is similar to that of the previous part. After visiting C, D, E, we know the value at  $B_1 = \min(C, D, E) = 6$ . Then after visiting  $H = 1 < 6$  we can prune G, and after visiting  $J = 3 < 6$  we can prune I and K. Among all the feasible solutions, this solution must come first in alphabetical order, since an earlier solution needs to either start with C, D, E, F, G, ... or C, D, E, F, H, G, I, ..., which do not satisfy the requirement.




Figure 2

- (b) In this part, Alice still chooses a box of balls, but now Bob only gets to choose a ball uniformly at random from that box, as shown in Figure 2. After Alice chooses a box and Bob randomly picks out a ball out of the box, Alice now has two

options: (1) keep that score, and the game ends normally, or (2) abandon that score and start over from choosing a box (she may choose the same box as the one she previously chose). The latter score will always override the previous score, even if it is lower.

(i) [2 pts] If Alice first chooses box  $B_3$  and get a score of 3, should Alice keep that score? If she choose to abandon the score, what is the optimal action for the next choice? If there are multiple optimal actions with the same expected score, choose all of them.

- Alice should keep the score.  
 Alice should abandon the score and choose  $B_1$ .  
 Alice should abandon the score and choose  $B_2$ .  
 Alice should abandon the score and choose  $B_3$ .

Choosing  $B_1$  at the next step gives an expected reward of  $(6 + 7 + 8)/3 = 7$ , which is greater than the current score (3), so we should abandon our current score. Choosing  $B_2$  or  $B_3$  gives smaller expected rewards compared to  $B_1$ , so we should choose  $B_1$  in the next step.

(ii) [2 pts] What is Alice's optimal action for the first choice she makes, given that she is aware of the potential second chance? If there are multiple optimal actions with the same expected score, choose all of them.

- $B_1$    $B_2$    $B_3$

Similar to the logic in the previous part, if the score we get in the first round is  $\geq 7$ , then we keep the score. Otherwise we abandon the score and re-choose  $B_1$ , which has an expected score of 7. Based on this strategy we can calculate the expected score for choosing different boxes in the first round:

Choosing  $B_1$ : keep the score when getting 8 or 7; abandon and re-choose  $B_1$  when getting 6. Expected score is  $(8 + 7 + 7)/3 = 22/3$ .

Choosing  $B_2$ : keep the score when getting 9 or 8; abandon and re-choose  $B_1$  when getting 1. Expected score is  $(9 + 8 + 7)/3 = 8$ .

Choosing  $B_3$ : keep the score when getting 10; abandon and re-choose  $B_1$  when getting 3 or 2. Expected score is  $(10 + 7 + 7)/3 = 8$ .

So  $B_2$  and  $B_3$  are both optimal first choices.

(iii) [5 pts] Suppose now that Alice has the chance to get the previous score overridden twice (if she is not satisfied with her second score, she start over a third time). Given that she is aware of this rule, what is her optimal policy? A description of the optimal policy is given below. Fill in the blanks. Ties are broken with the order  $B_1, B_2, B_3$ .

First, choose box   $B_1$    $B_2$    $B_3$ .

If the score  $\geq$  , then keep the score and end the game.

Otherwise, abandon the previous score and choose box   $B_1$    $B_2$    $B_3$ .

If the score  $\geq$  , then keep the score and end the game.

Otherwise, abandon the previous score again and choose box   $B_1$    $B_2$    $B_3$  at the final chance.

Similar to the logic in the previous part. In the previous part we calculated that the maximum expected score when we have two choices left is 8. Thus, when the score we get in the first round is  $\geq 8$ , then we keep the score. Otherwise we abandon the score and re-choose  $B_2$  (due to tie-breaking), which has an expected score of 8. Based on this strategy we can calculate the expected score for choosing different boxes in the first round:

Choosing  $B_1$ : keep the score when getting 8; abandon and re-choose  $B_2$  when getting 7 or 6. Expected score is  $(8 + 8 + 8)/3 = 8$ .

Choosing  $B_2$ : keep the score when getting 9 or 8; abandon and re-choose  $B_2$  when getting 1. Expected score is  $(9 + 8 + 8)/3 = 25/3$ .

Choosing  $B_3$ : keep the score when getting 10; abandon and re-choose  $B_2$  when getting 3 or 2. Expected score is  $(10 + 8 + 8)/3 = 26/3$ .

So choosing  $B_3$  is the optimal strategy, and its expected score is  $26/3$ .

(iv) [1 pt] What is the expected score Alice will get if she follows the sequence of choices described in the previous problem?  See the explanation to the previous question.

# Q6. [20 pts] Bayes Nets and Inference

“The greatest progress that the human race has made lies in learning how to make correct inferences”,  
*Friedrich Nietzsche.*

And after such a pompous quote let's dig into our problem.

(a) (i) [4 pts] Assume  $P(C|B, A) = P(C|B)$  for random variables  $A, B$  and  $C$ . Which of the following expressions hold?

$P(A|B, C) = P(A|B)$

$\frac{P(B, C|A)P(A)}{P(B)} = \frac{\sum_b P(A, b, C)}{P(B)}$

$P(A, C|B) = P(A|B)P(C|B)$





$P(B|A, C) = \frac{P(B)P(A|B)P(C|B)}{P(A)P(A|C)}$

Independence is symmetric.

$$P(A|B, C) = \frac{P(B, C|A)P(A)}{P(B, C)} = \frac{P(C|A, B)P(B|A)P(A)}{P(C|B)P(B)} = \frac{P(C|B)P(B|A)P(A)}{P(C|B)P(B)} = \frac{P(B|A)P(A)}{P(B)} = P(A|B)$$


The third equality holds because of the assumption given in the question.

(ii) [4 pts] For which of the following Bayes Nets is  $P(A)P(C|A) = P(C) * \sum_c P(A, c|B)$  guaranteed to hold?



for  $P(A)P(C|A) = P(C) * \sum_c P(A, c|B)$  to hold, we need  $P(A)P(C|A) = P(C) * P(A|B)$ . Thus, we need  $A \perp\!\!\!\perp B$  and  $A \perp\!\!\!\perp C$ . These can be checked easily, remembering that a node is conditionally independent of its non-descendants given its parents (which means absolutely independent if it has no parents).

(iii) [2 pts]



Now consider the Bayes' net depicted above. We want to infer exactly  $P(A|b, e)$ , where  $A$  is the query variable,  $B, E$  the evidence variables,  $C, D$  the hidden variables and  $\alpha$  a normalizing constant which value you can choose. Which of the following expressions is the correct value of the query  $P(A|b, e)$ ?

- $\alpha \sum_c \sum_d P(c|b)P(b|A)P(b|e)P(A|d)P(e|d)P(d)$ 
  $\alpha \sum_c \sum_d P(c|b)P(b|A, e)P(A|d)P(e|d)$   
  $\alpha \sum_c \sum_d P(c|b)P(b|A, e)P(A|d)P(e|d)P(d)$ 
  $\alpha \sum_b \sum_e P(c|b)P(A|b, e)P(A|d)P(e|d)P(d)$

We compute  $P(A|b, e)$  as  $\alpha \sum_c \sum_d P(A, b, c, d, e) = \alpha \sum_c \sum_d P(c|b)P(b|A, e)P(A|d)P(e|d)P(d)$

(iv) [2 pts] For the same network as in (iii) which of the following expressions is the most computationally efficient expression for the query  $P(A|b, e)$ ?

- $\alpha P(b|A, e) \sum_d \sum_c P(A|d)P(e|d)P(c|b)P(d)$ 
  $\alpha P(b|A, e)P(A|d)P(e|d) \sum_c P(c|b) \sum_d P(d)$   
  $\alpha P(b|A, e) \sum_d P(A|d)P(e|d)P(d) \sum_c P(c|b)$ 
  $\alpha P(b|A, e) \sum_c P(A|d)P(e|d) \sum_d P(c|b)P(d)$

Top right and bottom right are wrong as probabilities that depend on  $d$  are left outside the summations. The top left requires  $4 \cdot 4 + 1$  multiplications and additions while the bottom left requires  $2 + 2 \cdot 4 + 1$  additions and multiplications.

(v) [2 pts] We call a variable irrelevant if its value has no effect on the value of the query we want to compute. For the same network as in (iii) which of the following variables, if any, is irrelevant to the query  $P(A|b, e)$ ?

- A
  C
  E  
 B
  D
  None

From the answer above,  $\alpha P(b|A, e) \sum_d P(A|d)P(e|d)P(d) \sum_c P(c|b) = \alpha P(b|A, e) \sum_d P(A|d)P(e|d)P(d)$ . So we can see that  $C$  can be "summed out" and is thus irrelevant. Generally speaking, any leaf node with no evidence can be ignored.

(vi) [2 pts] If you were to perform variable elimination on the the Bayes Net depicted above to estimate  $P(A|b, e)$ , what would the size of the largest generated probability table be under the most efficient elimination ordering? (We define "most efficient elimination ordering" as the order that yields the minimally-sized largest probability table.) Assume that all variables  $A, B, C, D, E$  are binary.

- 2
  8  
 4
  16

$\alpha P(b|A, e) \sum_d P(A|d)P(e|d)P(d) = \alpha f_1(A) \sum_d f_2(A, D)f_3(D)f_4(D)$  and the largest factor is the  $f_2(A, D)$  with 4 elements.


(vii) [1 pt] This question is unrelated to the previous questions. Let  $A, B, C, D, E, F$  and  $G$  be binary random variables. What is the size of the probability table generated by the pointwise product of the factors  $P(A, B, C) \times P(A, F, G)$ ?

- $2^3$ 
  $2^5$   
  $2^4$ 
  $2^6$

The answer is 2 to the power of the number of elements in the union of the variables that appear in the two factors.

(b) In large networks exact inference can become intractable. For this reason we frequently employ sampling methods for inference.

(i) [1 pt]



Consider the seven-node Bayesian network above. Which set of variables comprise the Markov blanket of variable  $D$ ?

- F  
 A, B

- A, B, F  
 A, B, F, C

The Markov blanket of a node consists of the parents, the children, and children's other parents.

- (ii) [2 pts] Assume for convenience that the variables in the above network are binary and that we have observed  $A = 1$ ,  $B = 0$ ,  $F = 0$ . We will be using Gibbs sampling to estimate the probability  $P(D|A = 1, B = 0, F = 0)$ . The initial values of the non-evidence variables are  $C = D = E = G = 1$ . We give you four sets of the first three updates performed by Gibbs sampling. The notation  $X, P(X|Y), X = x$  means that at this update step, variable  $X$  was chosen, it was sampled from distribution  $P(X|Y)$ , and the realized value of the variable was  $x$ . Which of these sets of updates could NOT have resulted from a correct implementation of Gibbs sampling?

- $E, P(E|B = 0, G = 1), E = 1$   
  $E, P(E|B = 0, G = 1), E = 0$   
  $E, P(E|B = 0, G = 1), E = 0$   
  $E, P(E|B = 0, G = 1), E = 1$   
  $G, P(G|E = 1), G = 0$   
  $E, P(E|B = 0, G = 1), E = 0$

- $C, P(C|A = 1, F = 0), C = 1$   
  $G, P(G|E = 1), G = 1$   
  $C, P(C|A = 1, F = 0), C = 0$   
  $C, P(C|A = 1, D = 1, F = 0), C = 1$   
  $G, P(G|E = 1), G = 0$   
  $E, P(E|B = 0, G = 0), E = 1$

The top left choice can plausibly arise from Gibbs sampling. We just keep picking  $E$  to sample, which is perfectly possible for three iterations. The top right is wrong because it fails to include  $D$  in the Markov blanket of  $C$ . The bottom left is also wrong since the third step does not include the updated value of  $G = 0$ . The sequence at bottom right can plausibly arise from Gibbs sampling as now  $C$  has the correct Markov blanket.