

- You have approximately 110 minutes.
- The exam is closed book, closed calculator, and closed notes except your one-page crib sheet.
- Mark your answers ON THE EXAM ITSELF. If you are not sure of your answer you may wish to provide a *brief* explanation.
- For multiple choice questions with *circular bubbles*, you should only mark ONE option; for those with *checkboxes*, you should mark ALL that apply (which can range from zero to all options). FILL in your answer COMPLETELY.

First name	
Last name	
SID	
edX username	
Name of person on your left	
Name of person on your right	

Your Discussion/Exam Prep* TA (fill all that apply):

- | | | | |
|--------------------------------------|---------------------------------------|---------------------------------------|---------------------------------------|
| <input type="checkbox"/> Brijen (Tu) | <input type="checkbox"/> Wenjing (Tu) | <input type="checkbox"/> Caryn (W) | <input type="checkbox"/> Andy* (Tu) |
| <input type="checkbox"/> Peter (Tu) | <input type="checkbox"/> Aaron (W) | <input type="checkbox"/> Anwar (W) | <input type="checkbox"/> Nikita* (Tu) |
| <input type="checkbox"/> David (Tu) | <input type="checkbox"/> Mitchell (W) | <input type="checkbox"/> Aarash (W) | <input type="checkbox"/> Daniel* (W) |
| <input type="checkbox"/> Nipun (Tu) | <input type="checkbox"/> Abhishek (W) | <input type="checkbox"/> Yuchen* (Tu) | <input type="checkbox"/> Shea* (W) |

For staff use only:

Q1. CSP: Air Traffic Control	/15
Q2. Utility	/10
Q3. State Representations and State Spaces	/14
Q4. Informed Search and Heuristics	/14
Q5. Strange MDPs	/16
Q6. Reinforcement Learning	/15
Q7. MedianMiniMax	/16
Total	/100

THIS PAGE IS INTENTIONALLY LEFT BLANK. NOTHING ON THIS PAGE WILL BE GRADED.

Q1. [15 pts] CSP: Air Traffic Control

We have five planes: A, B, C, D, and E and two runways: international and domestic. We would like to schedule a time slot and runway for each aircraft to **either** land or take off. We have four time slots: $\{1, 2, 3, 4\}$ for each runway, during which we can schedule a landing or take off of a plane. We must find an assignment that meets the following constraints:

- Plane B has lost an engine and must land in time slot 1.
 - Plane D can only arrive at the airport to land during or after time slot 3.
 - Plane A is running low on fuel but can last until at most time slot 2.
 - Plane D must land before plane C takes off, because some passengers must transfer from D to C.
 - No two aircrafts can reserve the same time slot for the same runway.
- (a) [3 pts] Complete the formulation of this problem as a CSP in terms of variables, domains, and constraints (both unary and binary). Constraints should be expressed implicitly using mathematical or logical notation rather than with words.

Variables: A, B, C, D, E for each plane.

Domains: a tuple $(runway\ type, time\ slot)$ for runway type $\in \{international, domestic\}$ and time slot $\in \{1, 2, 3, 4\}$.

Constraints:

$$B[1] = 1$$

$$D[1] \geq 3$$

$$A[1] \leq 2$$

$$D[1] < C[1]$$

$$A \neq B \neq C \neq D \neq E$$

- (b) For the following subparts, we add the following two constraints:

- Planes A, B, and C cater to international flights and can only use the international runway.
- Planes D and E cater to domestic flights and can only use the domestic runway.

- (i) [2 pts] With the addition of the two constraints above, we completely reformulate the CSP. You are given the variables and domains of the new formulation. Complete the constraint graph for this problem given the original constraints and the two added ones.

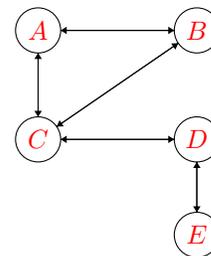
Variables: A, B, C, D, E for each plane.

Constraint Graph:

Domains: $\{1, 2, 3, 4\}$

Explanation of Constraint Graph:

We can now encode the runway information into the identity of the variable, since each runway has more than enough time slots for the planes it serves. We represent the non-colliding time slot constraint as a binary constraint between the planes that use the same runways.



- (ii) [4 pts] What are the domains of the variables after enforcing arc-consistency? Begin by enforcing unary constraints. (Cross out values that are no longer in the domain.)

A	1	2	3	4
B	1	2	3	4
C	1	2	3	4
D	1	2	3	4
E	1	2	3	4

- (iii) [4 pts] Arc-consistency can be rather expensive to enforce, and we believe that we can obtain faster solutions using only **forward-checking** on our variable assignments. Using the Minimum Remaining Values heuristic, perform backtracking search on the graph, breaking ties by picking lower values and characters first. List the $(variable, assignment)$ pairs in the order they occur (including the assignments that are reverted upon reaching a dead end). Enforce unary constraints before starting the search.

(You don't have to use this table, it won't be graded.)

A	1	2	3	4
B	1	2	3	4
C	1	2	3	4
D	1	2	3	4
E	1	2	3	4

Answer: (B, 1), (A, 2), (C, 3), (C, 4), (D, 3), (E, 1)

- (c) [2 pts] Suppose we have just one runway and n planes, where no two planes can use the runway at once. We are assured that the constraint graph will always be tree-structured and that a solution exists. What is the runtime complexity in terms of the number of planes, n , of a CSP solver that runs arc-consistency and then assigns variables in a topological ordering?

- $O(1)$
- $O(n)$
- $O(n^2)$
- $O(n^3)$
- $O(n^n)$
- None of the Above

Modified AC-3 for tree-structured CSPs runs arc-consistency backwards and then assigns variables in forward topological (linearized) ordering so we that we don't have to backtrack. The runtime complexity of modified AC-3 for tree-structured CSPs is $O(nd^2)$, but note that the domain of each variable must have a domain of size at least n since a solution exists

Q2. [10 pts] Utility

- (a) [4 pts] Pacwoman is enticed by a lottery at the local Ghostway. Ghostway is offering a lottery which rewards three different outcomes at equal probability: 2 dollars, 5 dollars and 6 dollars. The price to purchase a lottery ticket is 2 dollars. The ghost at the cash register, however, is now offering Pacwoman an option: give him a bribe of 1 dollar more and he'll manipulate the lottery so that she will not get the worst outcome (the rest of the options have equal likelihood). Pacwoman's utility is dependent on her net gain, p . For which of the following utility functions, should Pacwoman decide to bribe the ghost?

- $U(p) = p$
 $U(p) = p^2$
 None of the Above

Explanation: Compare $\frac{1}{3} * U(2 - 2) + \frac{1}{3} * U(5 - 2) + \frac{1}{3} * U(6 - 2)$ and $\frac{1}{2} * U(5 - 3) + \frac{1}{2} * U(6 - 3)$ for different options.

- (b) You are in a pack of 9 Pacpeople about to embark on a mission to steal food from the ghosts. In order to proceed, the pack needs a Pacleader to direct the mission.

Being a leader is very tiring, so the Pacleader would need to spend 5 food pellets, to consume for energy to lead. If the mission is successful, however, the Pacleader gets 1/5th of the food that is scavenged. The utility of the Pacleader $U_L(l, t)$ depends on their **individual food net gain** (l) and on the **total food** collected for the pack (t).

The remaining non-leading Pacmembers get an equal share of the remaining 4/5ths of the food. A Pacfollower gains utility $U_f(f)$, which is dependent on only their **individual food net gain** (f).

Let X be the probability of going home with no food, otherwise the mission is successful and brings back 100 pellets. For what value of X , should you, a rational Pacperson, decide to step up and lead the pack? Assuming that if you don't, some other Pacmember will.

- (i) [4 pts] Express your answer in terms of X, U_L, U_f (You do not have to simplify for X):

$$(1 - X)U_L\left(\frac{1}{5} * 100 - 5, 100\right) + XU_L(-5, 0) > (1 - X)U_f\left(\frac{4}{5} * 100 * \frac{1}{8}\right) + XU_f(0)$$

$$(1 - X)U_L(15, 100) + XU_L(-5, 0) > (1 - X)U_f(10) + XU_f(0)$$

- (ii) [2 pts] Calculate an explicit expression for X given:

$$U_L(l, t) = 5l + \frac{t}{5}$$

$$U_f(f) = f^2 - 4f$$

where,
 l is the individual food net gain of the Pacleader
 t is the total amount of food scavenged for the pack
 f is the individual food net gain of a Pacfollower

$$(1 - X)(5 * 15 + 20) - X * (25) > (1 - X)(100 - 40)$$

$$95 - 95X - 25X > 60 - 60X$$

$$95 - 120X > 60 - 60X$$

$$95 - 60 > 120X - 60X$$

$$35 > 60X$$

$$X < \frac{7}{12}$$

Q3. [14 pts] State Representations and State Spaces

For each part, state the size of a minimal state space for the problem. Give your answer as an expression that references problem variables. Below each term, state what information it encodes. For example, you could write $2 \times MN$ and write “whether a power pellet is in effect” under the 2 and “Pacman’s position” under the MN . State spaces which are complete but not minimal will receive partial credit.

Each part is independent. A maze has **height** M and **width** N . A Pacman can move *NORTH*, *SOUTH*, *EAST*, or *WEST*. There is initially a **pellet in every position** of the maze. The **goal is to eat all of the pellets**.

(a) [4 pts] Personal Space

In this part, there are P Pacmen, numbered $1, \dots, P$. Their turns cycle so Pacman 1 moves, then Pacman 2 moves, and so on. Pacman 1 moves again after Pacman P . Any time two Pacmen enter adjacent positions, the one with the lower number dies and is removed from the maze.

$$(MN)^P \times 2^{MN} \times P$$

$(MN)^P$: for each Pacman (the dead Pacmen are “eaten” by the alive Pacman, we encapsulate this in the transition function so that Pacmen in the same position must move together and can only move during the turn of the higher numbered Pacman)

2^{MN} : for each position, whether the pellet there has been eaten

P : the Pacman whose turn it is

$(MN + 1)^P \times 2^{MN} \times P$ is also accepted for most of the points. Where $(MN + 1)$ also includes DEAD as a position

(b) [4 pts] Road Not Taken

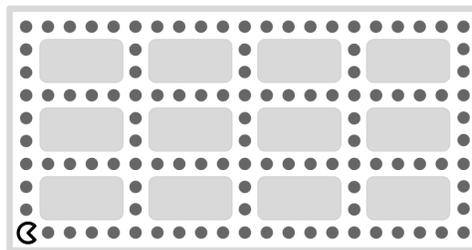
In this part, there is one Pacman. Whenever Pacman enters a position which he has visited previously, the maze is reset – each position gets refilled with food and the “visited status” of each position is reset as well.

$$MN \times 2^{MN}$$

MN : Pacman’s position

2^{MN} : for each position, whether the pellet there has been eaten (and equivalently, whether it has been visited)

(c) [6 pts] Hallways



In this part, there is one Pacman. The walls are arranged such that they create a grid of H hallways **total**, which connect at I intersections. (In the example above, $H = 9$ and $I = 20$). In a single action, Pacman can move from one intersection into an adjacent intersection, eating all the dots along the way. Your answer should only depend on I and H .

(note: H = number of vertical hallways + number of horizontal hallways)

$$I \times 2^{2I-H}$$

I : Pacman's position in any one of the intersections

2^{2I-H} : for each path between intersections, whether the pellets there have been eaten. The exponent was calculated via the following logic:

Approach 1:

Let v be the number of vertical hallways and h be the number of horizontal hallways. Notice that $H = v + h$ and $I = v * h$

Each vertical hallway has $h - 1$ segments, and each horizontal hallway has $v - 1$ segments.

Together, these sum to a total of $v(h - 1) + h(v - 1) = 2vh - v - h = 2I - H$ segments, each of which is covered by a single action.

Approach 2:

Let $H = v + h$ where v is the number of vertical hallways and h is the number of horizontal hallways

$4I$ = Every intersection has 4 paths adjacent to it

$-2v$ = The top and bottom intersection of each vertical hallway has one less vertical path

$-2h$ = The right-most and left-most intersection of each horizontal hallway has one less horizontal path

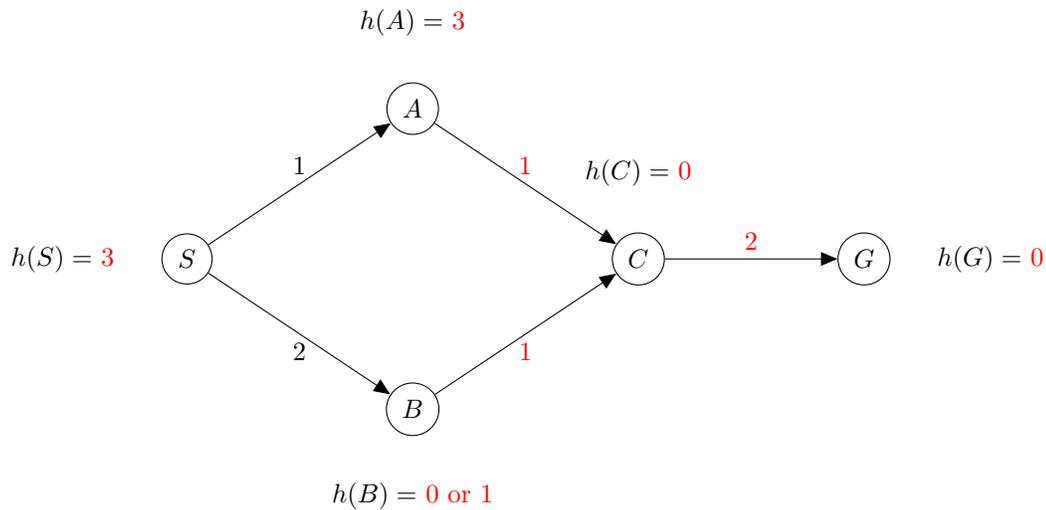
$4I - 2v - 2h$ must be divided by 2 since you count every path twice in each direction from the intersections which gives you

$$\frac{4I - 2v - 2h}{2} = 2I - v - h = 2I - H$$

Q4. [14 pts] Informed Search and Heuristics

(a) [6 pts] Consider the state space shown below, with starting state S and goal state G . Fill in a cost from the set $\{1, 2\}$ for each blank edge and a heuristic value from the set $\{0, 1, 2, 3\}$ for each node such that the following properties are satisfied:

- The heuristic is admissible but not consistent.
- The heuristic is monotonic non-increasing along paths from the start state to the goal state.
- A* graph search finds a suboptimal solution.
- You will never encounter ties (two elements in the fringe with the same priority) during execution of A*.

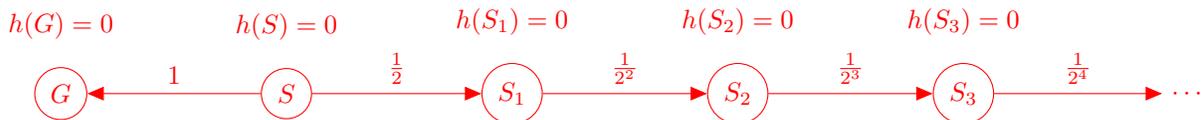


(b) [8 pts] **Don't spend all your time on this question.** As we saw in class, A* graph search with a consistent heuristic will always find an optimal solution when run on a problem with a finite state space. However, if we turn to problems with infinite state spaces, this property may no longer hold. Your task in this question is to provide a concrete example with an infinite state space where A* graph search fails to terminate or fails to find an optimal solution.

Specifically, you should describe the state space, the starting state, the goal state, the heuristic value at **each node**, and the cost of **each transition**. Your heuristic should be consistent, and all step costs should be strictly greater than zero ($cost \in \mathbb{R}_{>0}$) to avoid trivial paths with zero cost. To keep things simple, each state should have a finite number of successors, and the goal state should be reachable in a finite number of actions from the starting state.

You may want to start by drawing a diagram.

A* may not terminate if the costs are not bounded away from zero. Consider the following example with starting state S , goal state G , and the trivial heuristic $h(s) = 0$ for all s :



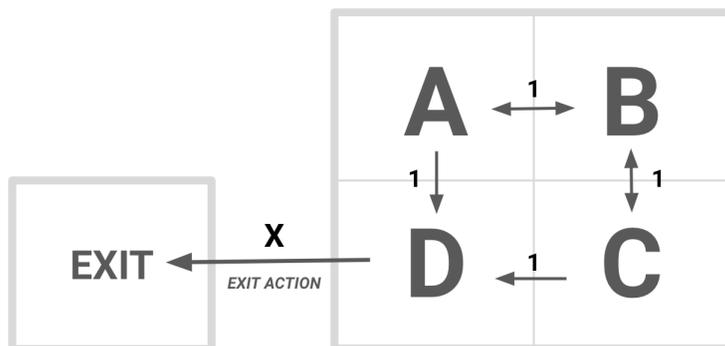
After expanding the starting state S , A* will expand S_1, S_2, \dots and will never expand the goal G , since the path cost of $S \rightarrow S_1 \rightarrow \dots \rightarrow S_k$ is

$$\sum_{j=1}^k \frac{1}{2^j} = 1 - \frac{1}{2^k},$$

which is strictly less than 1 for all k .

Q5. [16 pts] Strange MDPs

In this MDP, the available actions at **state A, B, C** are *LEFT*, *RIGHT*, *UP*, and *DOWN* unless there is a wall in that direction. The only action at **state D** is the *EXIT ACTION* and gives the agent a **reward of x** . The **reward for non-exit actions is always 1**.



- (a) [6 pts] Let all actions be deterministic. Assume $\gamma = \frac{1}{2}$. Express the following in terms of x .

$$V^*(D) = x$$

$$V^*(C) = \max(1 + 0.5x, 2)$$

$$V^*(A) = \max(1 + 0.5x, 2)$$

$$V^*(B) = \max(1 + 0.5(1 + 0.5x), 2)$$

The 2 comes from the utility being an infinite geometric sum of discounted reward = $\frac{1}{(1-\frac{1}{2})} = 2$

- (b) [6 pts] Let any non-exit action be successful with probability = $\frac{1}{2}$. Otherwise, the agent stays in the same state with reward = 0. The *EXIT ACTION* from the **state D** is still deterministic and will always succeed. Assume that $\gamma = \frac{1}{2}$.

For which value of x does $Q^*(A, DOWN) = Q^*(A, RIGHT)$? Box your answer and justify/show your work.

$$Q^*(A, DOWN) = Q^*(A, RIGHT) \text{ implies } V^*(A) = Q^*(A, DOWN) = Q^*(A, RIGHT)$$

$$V^*(A) = Q^*(A, DOWN) = \frac{1}{2}(0 + \frac{1}{2}V^*(A)) + \frac{1}{2}(1 + \frac{1}{2}x) = \frac{1}{2} + \frac{1}{4}(V^*(A)) + \frac{1}{4}x \quad (1)$$

$$V^*(A) = \frac{2}{3} + \frac{1}{3}x \quad (2)$$

$$V^*(A) = Q^*(A, RIGHT) = \frac{1}{2}(0 + \frac{1}{2}V^*(A)) + \frac{1}{2}(1 + \frac{1}{2}V^*(B)) = \frac{1}{2} + \frac{1}{4}V^*(A) + \frac{1}{4}V^*(B) \quad (3)$$

$$V^*(A) = \frac{2}{3} + \frac{1}{3}V^*(B) \quad (4)$$

Because $Q^*(B, LEFT)$ and $Q^*(B, DOWN)$ are symmetric decisions, $V^*(B) = Q^*(B, LEFT)$.

$$V^*(B) = \frac{1}{2}(0 + \frac{1}{2}V^*(B)) + \frac{1}{2}(1 + \frac{1}{2}V^*(A)) = \frac{1}{2} + \frac{1}{4}V^*(B) + \frac{1}{4}V^*(A) \quad (5)$$

$$V^*(B) = \frac{2}{3} + \frac{1}{3}V^*(A) \quad (6)$$

Combining (2), (4), and (6) gives us:

$$x = 1 \quad (7)$$

There is also a shortcut which involves you noticing that the problem is highly symmetrical such that $Q^*(A, DOWN) = Q^*(A, RIGHT)$ is the same as solving the equivalence of $V^*(A)$ in the previous part and the utility of an infinite cycle with reward scaled by half (to account for staying) and discount = 0.5. That leads us to conclude $0.5 + 0.5x = \frac{0.5}{1-0.5} = 1$ so $x = 1$

- (c) [4 pts] We now add one more layer of complexity. Turns out that the reward function is not guaranteed to give a particular reward when the agent takes an action. Every time an agent transitions from one state to another, once the agent reaches the new state s' , a fair 6-sided dice is rolled. If the dices lands with value x , the agent receives the reward $R(s, a, s') + x$. The sides of dice have value 1, 2, 3, 4, 5 and 6.

Write down the new bellman update equation for $V_{k+1}(s)$ in terms of $T(s, a, s')$, $R(s, a, s')$, $V_k(s')$, and γ .

$$\begin{aligned} V_{k+1}(s) &= \max_a \sum_{s'} T(s, a, s') \left[\frac{1}{6} (\sum_{i=1}^6 R(s, a, s') + i) + \gamma V_k(s) \right] \\ &= \max_a \sum_{s'} T(s, a, s') (R(s, a, s') + 3.5 + \gamma V_k(s)) \end{aligned}$$

Q6. [15 pts] Reinforcement Learning

Imagine an unknown environments with four states (A, B, C, and X), two actions (\leftarrow and \rightarrow). An agent acting in this environment has recorded the following episode:

s	a	s'	r	Q-learning iteration numbers (for part b)
A	\rightarrow	B	0	1, 10, 19, ...
B	\rightarrow	C	0	2, 11, 20, ...
C	\leftarrow	B	0	3, 12, 21, ...
B	\leftarrow	A	0	4, 13, 22, ...
A	\rightarrow	B	0	5, 14, 23, ...
B	\rightarrow	A	0	6, 15, 24, ...
A	\rightarrow	B	0	7, 16, 25, ...
B	\rightarrow	C	0	8, 17, 26, ...
C	\rightarrow	X	1	9, 18, 27, ...

- (a) [4 pts] Consider running model-based reinforcement learning based on the episode above. Calculate the following quantities:

$$\hat{T}(B, \rightarrow, C) = \frac{\boxed{2}}{\boxed{3}}$$

$$\hat{R}(C, \rightarrow, X) = \frac{\boxed{1}}{\boxed{1}}$$

- (b) [5 pts] Now consider running Q-learning, repeating the above series of transitions in an infinite sequence. Each transition is seen at multiple iterations of Q-learning, with iteration numbers shown in the table above.

After which iteration of Q-learning do the following quantities first become nonzero? (If they always remain zero, write *never*).

$$Q(A, \rightarrow)? \quad \frac{\boxed{14}}{\boxed{1}}$$

$$Q(B, \leftarrow)? \quad \frac{\boxed{22}}{\boxed{1}}$$

- (c) [6 pts] True/False: For each question, you will get positive points for correct answers, zero for blanks, and negative points for incorrect answers. Circle your answer **clearly**, or it will be considered incorrect.

- (i) [1.5 pts] [true or *false*] In Q-learning, you do not learn the model.

Q learning is model-free, you learn the optimal policy explicitly, and the model itself implicitly.

- (ii) [1.5 pts] [*true* or false] For TD Learning, if I multiply all the rewards in my update by some nonzero scalar p , the algorithm is still guaranteed to find the optimal policy.

If p is positive then yes, the discounted values, relative to each other, are just scaled. But if p is negative, you will be computing negating the values for the states, but the policy is still chosen on the max values.

- (iii) [1.5 pts] [*true* or false] In Direct Evaluation, you recalculate state values after each transition you experience.

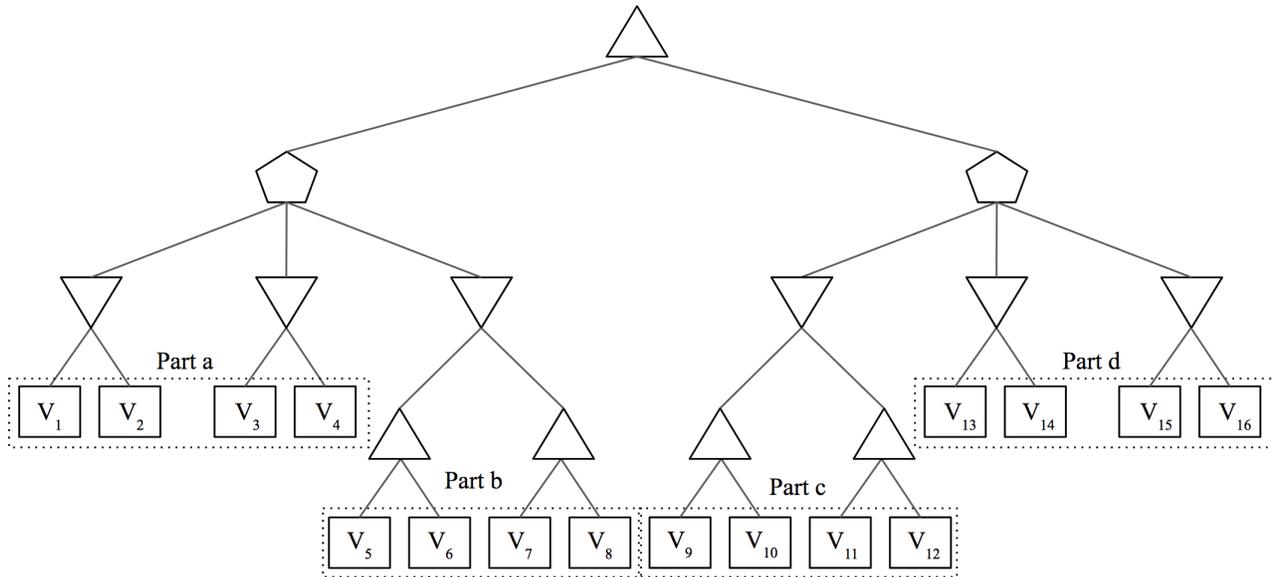
In order to estimate state values, you calculate state values from episodes of training, not single transitions.

- (iv) [1.5 pts] [*true* or false] Q-learning requires that all samples must be from the optimal policy to find optimal q-values.

Q-learning is off-policy, you can still learn the optimal values even if you act suboptimally sometimes.

Q7. [16 pts] MedianMiniMax

You're living in utopia! Despite living in utopia, you still believe that you need to maximize your utility in life, other people want to minimize your utility, and the world is a 0 sum game. But because you live in utopia, a benevolent social planner occasionally steps in and chooses an option that is a compromise. Essentially, the social planner (represented as the pentagon) is a median node that chooses the successor with median utility. Your struggle with your fellow citizens can be modelled as follows:



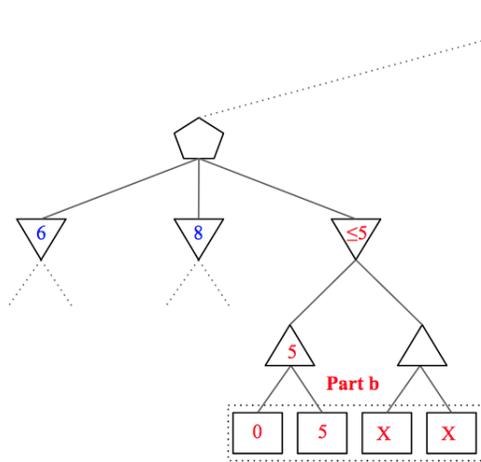
There are some nodes that we are sometimes able to prune. In each part, mark all of the terminal nodes such that **there exists a possible situation** for which the node **can be pruned**. In other words, you must consider **all** possible pruning situations. Assume that evaluation order is **left to right** and all V_i 's are **distinct**.

Note that as long as there exists ANY pruning situation (does not have to be the same situation for every node), you should mark the node as prunable. Also, alpha-beta pruning does not apply here, simply prune a sub-tree when you can reason that its value will not affect your final utility.

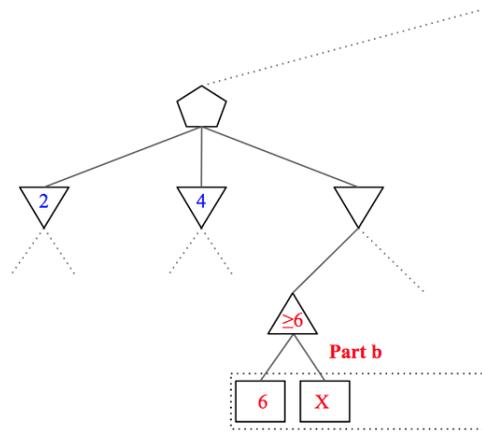
- | | | | |
|--|--|--|--|
| <p>(a)</p> <p><input type="checkbox"/> V_1</p> <p><input type="checkbox"/> V_2</p> <p><input type="checkbox"/> V_3</p> <p><input type="checkbox"/> V_4</p> <p><input checked="" type="checkbox"/> None</p> | <p>(b)</p> <p><input type="checkbox"/> V_5</p> <p><input checked="" type="checkbox"/> V_6</p> <p><input checked="" type="checkbox"/> V_7</p> <p><input checked="" type="checkbox"/> V_8</p> <p><input type="checkbox"/> None</p> | <p>(c)</p> <p><input type="checkbox"/> V_9</p> <p><input type="checkbox"/> V_{10}</p> <p><input checked="" type="checkbox"/> V_{11}</p> <p><input checked="" type="checkbox"/> V_{12}</p> <p><input type="checkbox"/> None</p> | <p>(d)</p> <p><input type="checkbox"/> V_{13}</p> <p><input checked="" type="checkbox"/> V_{14}</p> <p><input checked="" type="checkbox"/> V_{15}</p> <p><input checked="" type="checkbox"/> V_{16}</p> <p><input type="checkbox"/> None</p> |
|--|--|--|--|

Part a:

For the left median node with three children, at least two of the children's values must be known since one of them will be guaranteed to be the value of the median node passed up to the final maximizer. For this reason, none of the nodes in part a can be pruned.



The value of this subtree will only get smaller.
 The median node will **NOT** choose the value of this subtree. 6 is the median.



The value of this subtree will only get bigger.

If the value of this subtree is chosen by the minimizer*, it will **NOT** be chosen by the median node.

**It is possible that the median is the value of the subtree to the right that we haven't looked at yet*

Part b (pruning V_7, V_8):

Let min_1, min_2, min_3 be the values of the three minimizer nodes in this subtree.

In this case, we may not need to know the final value min_3 . The reason for this is that we may be able to put a bound on its value after exploring only partially, and determine the value of the median node as either min_1 or min_2 if $min_3 \leq \min(min_1, min_2)$ or $min_3 \geq \max(min_1, min_2)$.

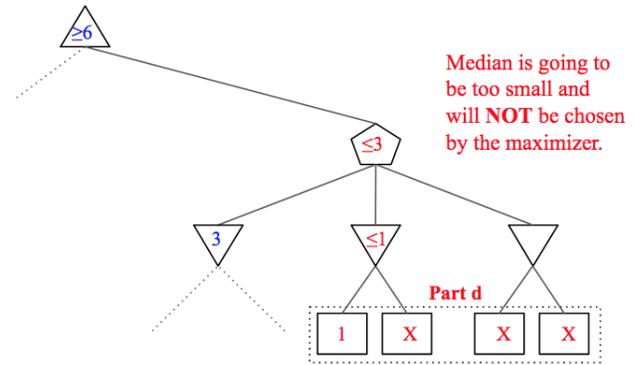
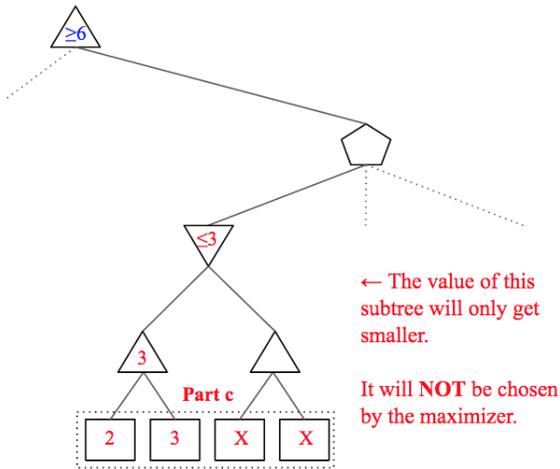
We can put an upper bound on min_3 by exploring the left subtree V_5, V_6 and if $\max(V_5, V_6)$ is lower than both min_1 and min_2 , the median node's value is set as the smaller of min_1, min_2 and we don't have to explore V_7, V_8 in Figure 1.

Part b (pruning V_6):

It's possible for us to put a lower bound on min_3 . If V_5 is larger than both min_1 and min_2 , we do not need to explore V_6 .

The reason for this is subtle, but if the minimizer chooses the left subtree, we know that $min_3 \geq V_5 \geq \max(min_1, min_2)$ and we don't need V_6 to get the correct value for the median node which will be the larger of min_1, min_2 .

If the minimizer chooses the value of the right subtree, the value at V_6 is unnecessary again since the minimizer never chose its subtree.



Part c (pruning V_{11}, V_{12}):

Assume the highest maximizer node has a current value $max_1 \geq Z$ set by the left subtree and the three minimizers on this right subtree have value min_1, min_2, min_3 .

In this part, if $min_1 \leq \max(V_9, V_{10}) \leq Z$, we do not have to explore V_{11}, V_{12} . Once again, the reasoning is subtle, but we can now realize if either $min_2 \leq Z$ or $min_3 \leq Z$ then the value of the right median node is for sure $\leq Z$ and is useless.

Only if both $min_2, min_3 \geq Z$ will the whole right subtree have an effect on the highest maximizer, but in this case the exact value of min_1 is not needed, just the information that it is $\leq Z$. Clearly in both cases, V_{11}, V_{12} are not needed since an exact value of min_1 is not needed.

We will also take the time to note that if $V_9 \geq Z$ we do have to continue the exploring as V_{10} could be even greater and the final value of the top maximizer, so V_{10} can't really be pruned.

Part d (pruning V_{14}, V_{15}, V_{16}):

Continuing from part c, if we find that $min_1 \leq Z$ and $min_2 \leq Z$ we can stop.

We can realize this as soon we explore V_{13} . Once we figure this out, we know that our median node's value must be one of these two values, and neither will replace Z so we can stop.