CS 188
Summer 2022

Exam Prep 3A Solutions

Q1. Local Search

(a) Hill Climbing

(i) Hill-climbing is complete. ☐ True ■ False

Consider hill-climbing for 8-queen.

(ii) Hill-climbing is optimal. ☐ True ■ False

no completeness indicates no optimality.

(b) Simulated Annealing

(i) The higher the temperature T is, the more likely the randomly chosen state will be expanded. ■ True ☐ False

The higher T is, the larger $e^{\Delta E/T}$ is given ΔE is negative.

(ii) In one round of simulated annealing, the temperature is 2 and the current state S has energy 1. It has 3 successors: A with energy 2; B with energy 1; C with energy 1-ln 4. If we assume the temperature does not change, What's the probability that these states will be chosen to expand after S eventually?

A, B will be expanded with probability $\frac{2}{5}$, C will be expanded with probability $\frac{1}{5}$.

Proof. First, the problem is asking which node will be expanded **next**, not **in this round**. A, B and C are randomly selected for expansion. If A or B is selected, they will surely be expanded. If C is selected, it has probability of 1/2 to be expanded and 1/2 to restart the random selection. Thus the probability ratio between A, B and C is 2:2:1.

(iii) On a undirected graph, If T decreases slowly enough, simulated annealing is guaranteed to converge to the optimal state. ■ True ☐ False

(c) Local Beam Search

The following state graph is being explored with local beam search with a beam of size K=2. The state's score given, and lower scores are considered better. Which of the following statements are true?

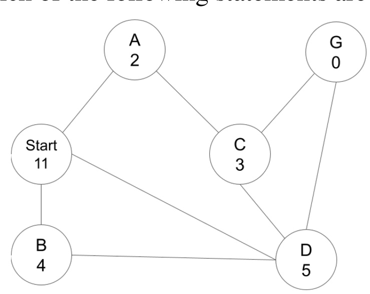

■ States A and B will be expanded before C and D.

☐ States A and D will be expanded before B and C.

☐ States B and D will be expanded before A and C.

☐ None of above.

(d) Genetic Algorithm

(i) In genetic algorithm, cross-over combine the genetic information of two parents to generate new offspring.

■ True ☐ False

(ii) In genetic algorithm, mutation involves a probability that some arbitrary bits in a genetic sequence will be flipped from its original state.

■ True □ False

(e) Gradient Descent

(i) Gradient descent is optimal. □ True ■ False

False. Gradient descent can become trapped in a local minimum.

(ii) For a function $f(x)$ with derivative $f'(x)$, write down the gradient descent update to go from x_t to x_{t+1}. Learning rate is α.

$x_{t+1} = x_t - \alpha f'(x_t)$, where α is the learning rate.

Q2. Games

Alice is playing a two-player game with Bob, in which they move alternately. Alice is a maximizer. Although Bob is also a maximizer, Alice believes Bob is a minimizer with probability 0.5, and a maximizer with probability 0.5. Bob is aware of Alice's assumption.
In the game tree below, square nodes are the outcomes, triangular nodes are Alice's moves, and round nodes are Bob's moves. Each node for Alice/Bob contains a tuple, the left value being Alice's expectation of the outcome, and the right value being Bob's expectation of the outcome.
Tie-breaking: choose the left branch.
The left values are Alice's expectations, and are the only thing Alice can refer to when making decisions.
The right values are Bob's expectations, and they also accurately track the expected outcome of the game according to each choice of branching (regardless of it is Alice's or Bob's decision, since Bob has all the information). Hence the right values are accurate information about the game, and would be what Bob looks at when making his decisions. However, when it is Alice's turn to make decisions, Bob will think about how Alice would maximize the outcome w.r.t to what she believes, and he will update his expectations accordingly.

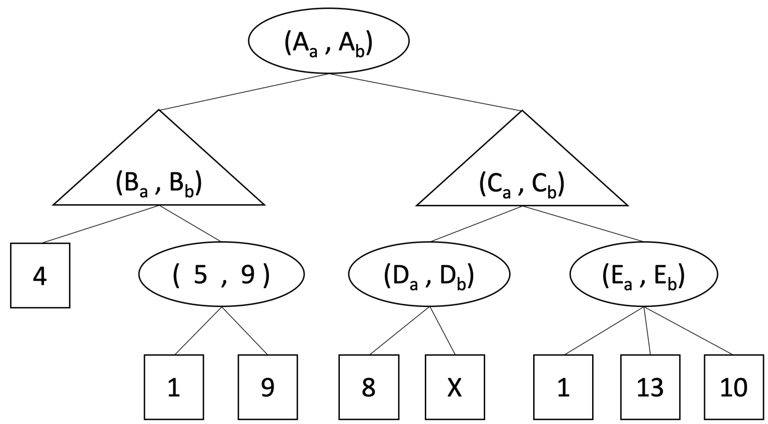

(a) In the blanks below, fill in the tuple values for tuples (B_a, B_b) and (E_a, E_b) from the above game tree.

$(B_a, B_b) = ($ 5 , 9 $)$

$(E_a, E_b) = ($ 7 , 13 $)$

For a square node, its value v means the same to Alice and Bob, i.e., we can think of it as a tuple (v,v).

The left value of Alice's nodes is the maximum of the left values of it's children nodes, since Alice believes that the values of the nodes are given by left values, and it's her turn of action, so she will choose the largest value.
The right value of Alice's nodes is the right value from the child node that attains the maximum left value since Bob's expectation is consistent with how Alice will act.
So for a triangular node, its tuple is the same as its child that has the maximum left value.

The left value of Bob's nodes is the average of the maximum and minimum of the left values of it's children nodes since Alice believes Bob is 50% possible to be adversarial and 50% possible to be friendly.
The right value of Bob's nodes is the maximum of the right values of the immediate children nodes since Bob would choose the branch that gives the maximum outcome during his turn.
So for a round node, left = 0.5(max(children.left) + min(children.left)), and right = max(children.right).

(b) In this part, we will determine the values for tuple (D_a, D_b).

(i) $D_a =$ ○ 8 ○ X ○ 8+X ● 4+0.5X ○ min(8,X) ○ max(8,X)

(ii) $D_b =$ ○ 8 ○ X ○ 8+X ○ 4+0.5X ○ min(8,X) ● max(8,X)

It's a round node, so left = 0.5(max(children.left) + min(children.left)), and right = max(children.right).
Its children: (8,8) and (X,X). So left = 0.5(8+X) = 4+0.5X, and right = max(8, X).

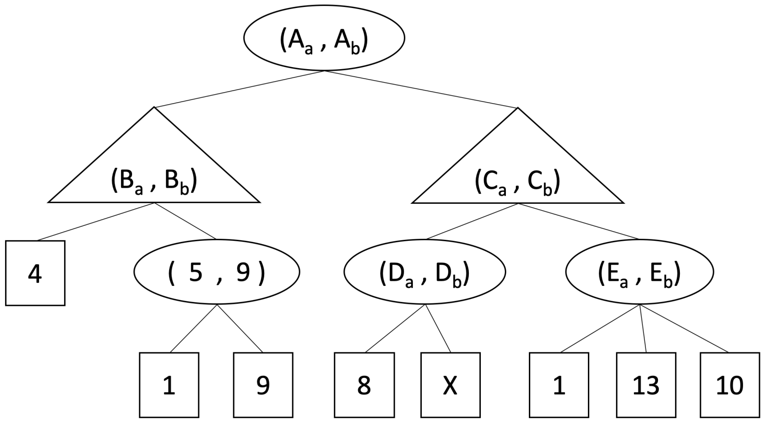

(The graph of the tree is copied for your convenience. You may do problem e on this graph.)

(c) Fill in the values for tuple (C_a, C_b) below. For the bounds of X, you may write scalars, ∞ or $-\infty$.
If your answer contains a fraction, please write down the corresponding **simplified decimal value** in its place. (i.e., 4 instead of $\frac{8}{2}$, and 0.5 instead of $\frac{1}{2}$).

1. If $-\infty <$ X $<$ $\boxed{6}$, $(C_a, C_b) = ($ $\boxed{7}$, $\boxed{13}$)
2. Else, $(C_a, C_b) = ($ $\boxed{4+0.5X}$, max($\boxed{8}$, $\boxed{X}$))

It's a triangular node, so its tuple is the same as its child that has the maximum left value.
Its children: (4+0.5X, max(8,X)) and (7, 13).
So if 4+0.5X < 7, i.e. $-\infty < X < 6$, it's the same as child node (7, 13), and otherwise it's (4+0.5X, max(8,X)).

(d) Fill in the values for tuple (A_a, A_b) below. For the bounds of X, you may write scalars, ∞ or $-\infty$.
If your answer contains a fraction, please write down the corresponding **simplified decimal value** in its place. (i.e., 4 instead of $\frac{8}{2}$, and 0.5 instead of $\frac{1}{2}$).

1. If $-\infty <$ X $<$ $\boxed{6}$, $(A_a, A_b) = ($ $\boxed{6}$, $\boxed{13}$)
2. Else, $(A_a, A_b) = ($ $\boxed{4.5+0.25X}$, max($\boxed{9}$, $\boxed{X}$))

It's a round node, so left = 0.5(max(children.left) + min(children.left)), and right = max(children.right).
Its children: (5,9) and node "Part (c)".
If $-\infty < X < 6$, these children are (5,9) and (7, 13).
left = 0.5(max(children.left) + min(children.left)) = 0.5(5+7) = 6
right = max(children.right) = max(9, 13) = 13.
Otherwise ($6 < X < +\infty$), these children are (5,9) and (4+0.5X, max(8,X)).
left = 0.5(max(children.left) + min(children.left)) = 0.5(5+4+0.5X) = $4.5 + 0.25X$
right = max(children.right) = max(9, max(8,X)) = max(9,X).

(e) When Alice computes the left values in the tree, some branches can be pruned and do not need to be explored. In the game tree graph on this page, put an 'X' on these branches. If no branches can be pruned, mark the "Not possible" choice below.
Assume that the children of a node are visited in left-to-right order and that you should not prune on equality.

● Not possible
It's impossible to determine the average of *min* and *max* until all children nodes are seen, so no pruning can be done for Alice. Leaving "Not possible" unmarked and no 'X' found in the graph is interpreted as 'no conclusion' and will not be given credit.