


Assigned: 10/19/21 Due: 10/26/21 at 11:55 PM (submit on gradescope, mark the page(s) associated for each problem when submitting)

Problem 1. (20 points)

Assume you are part-way through a Monte-Carlo Tree Search (MCTS) and have the tree below (ignore that it might not be generated correctly). Each node shows [wins / total picks].



(1) Calculate the UCB values for the whole tree (in practice we'd only do this on the path to the selected node).

(2) Which node would be selected to perform the random rollout upon?

(3) If the random rollout resulted in a loss, show the new values for all UCB values on a tree (in practice, we would not actually re-calculate all of them).

(4) Using your updated values from part (2), which node would be selected to perform the random rollout upon?

Problem 2. (25 points)

Generate a heuristic for each of these two problems. Ensure you clearly detail (1) how you are relaxing the problem, (2) what the optimal solution to this relaxed problem is and (3) clearly identify how you can generate a number for any given state.

(1) A "sliding puzzle" (video: https://www.youtube.com/watch?v=vx25e_9Z3ok). You have different shape blocks and want to get a selected block outside the puzzle (there is a on the edge somewhere). You want to find the least amount of moves possible to get the selected block out. Much like the 8-

puzzle there is a blank square, which you can use to move the various shapes around (no rotating).

(2) The bridge crossing puzzle. Suppose there are a number of people that need to cross a dangerous bridge at night, however there is only one flashlight. This flashlight only provides enough light for 1 or 2 people to cross the bridge safely (the flashlight also needs to be walked to the other side and back, not thrown or something). Each person walks at a different speed and when crossing the bridge together you have to walk the slower speed of the pair (otherwise they would get separated). You want to figure out how to get people across the bridge as fast as possible.

Note: this problem might be more complicated than you think. For example if you have four people: $A=1$ min to cross, $B = 2$ min, $C = 10$ min, and $D=20$ min. The best solution is not to have A always carry the flashlight: $(A+B) \rightarrow A_back \rightarrow (A+C) \rightarrow A_back \rightarrow (A+D) = 2+1+10+1+20 = 34$ minutes. In this example, you can actually get everyone across the bridge in 27 minutes. You **do not** need to actually figure out this solution to generate a heuristic... but understanding why the problem is difficult might help you figure out what you should relax.

Problem 3. (15 points)

Suppose you were running k-beam search on a problem that is best represented by an undirected graph. What is an issue you will run into when performing k-beam search in this setting that did not occur in the tree examples we did in class?

Problem 4. (15 points)

Which of the following are appropriate to phrase as a minimax tree and which are not (ignoring the complexity of the problem)? Write a short sentence justifying the reasoning.

(1) Rock-paper-scissors


(2) Tic-tac-toe

(3) Volley ball (assume the actions are either: set to another player or hit it back immediately)

Problem 5. (25 points)

Run alpha-beta pruning on this tree. At every node show the history of alpha/beta (or up/down arrow) values when running DFS from left-to-right. Clearly indicate what parts of the tree do not need to be searched (i.e. pruned).

(Picture on next page.)



Problem 6. (20 points)

Use the given genetic algorithm in search.py and apply it to the n-queens problem. Analyze this on a number of different configuration to answer the following questions:

- (1) Is the genetic algorithm good in general to solve n-queens? Why or why not?
- (2) What is the purpose of f_ thresh and how did you handle/use it in your testing?
- (3) What is one positive thing about the genetic algorithm in general? One negative aspect?

Note: the current implementation of the genetic algorithm is a bit slow/inefficient. 1000 generations with a population size of 100 will probably take a few minutes to run (depending on your machine).