


Q1. Search: Heuristic Function Properties

For the following questions, consider the search problem shown on the left. It has only three states, and three directed edges. A is the start node and G is the goal node. To the right, four different heuristic functions are defined, numbered I through IV.



	$h(A)$	$h(B)$	$h(G)$
I	4	1	0
II	5	4	0
III	4	3	0
IV	5	2	0

(a) Admissibility and Consistency

For each heuristic function, circle whether it is admissible and whether it is consistent with respect to the search problem given above.

	Admissible?		Consistent?	
I	<input checked="" type="checkbox"/> Yes	No	Yes	<input checked="" type="checkbox"/> No
II	Yes	<input checked="" type="checkbox"/> No	Yes	<input checked="" type="checkbox"/> No
III	<input checked="" type="checkbox"/> Yes	No	<input checked="" type="checkbox"/> Yes	No
IV	<input checked="" type="checkbox"/> Yes	No	Yes	<input checked="" type="checkbox"/> No

II is the only inadmissible heuristic, as it overestimates the cost from B : $h(B) = 4$, when the actual cost to G is 3.

To check whether a heuristic is consistent, ensure that for all paths, $h(N) - h(L) \leq \text{path}(N \rightarrow L)$, where N and L stand in for the actual nodes. In this problem, $h(G)$ is always 0, so making sure that the direct paths to the goal ($A \rightarrow G$ and $B \rightarrow G$) are consistent is the same as making sure that the heuristic is admissible. The path from A to B is a different story.

Heuristic I is not consistent: $h(A) - h(B) = 4 - 1 = 3 \not\leq \text{path}(A \rightarrow B) = 2$.

Heuristic III is consistent: $h(A) - h(B) = 4 - 3 = 1 \leq 2$

Heuristic IV is not consistent: $h(A) - h(B) = 5 - 2 = 3 \not\leq 2$

(b) Function Domination

Recall that *domination* has a specific meaning when talking about heuristic functions.

Circle all true statements among the following.

1. Heuristic function III dominates IV.
2. Heuristic function IV dominates III.
3. Heuristic functions III and IV have no dominance relationship.
4. Heuristic function I dominates IV.

5. Heuristic function IV dominates I.
6. Heuristic functions I and IV have no dominance relationship.

For one heuristic to dominate another, *all* of its values must be greater than or equal to the corresponding values of the other heuristic. Simply make sure that this is the case. If it is not, the two heuristics have no dominance relationship.

Q2. Search

Each True/False question is worth 1 points. Leaving a question blank is worth 0 points. **Answering incorrectly is worth -1 points.**


- (a) Consider a graph search problem where for every action, the cost is at least ϵ , with $\epsilon > 0$. Assume the used heuristic is consistent.
- (i) [true or false] Depth-first graph search is guaranteed to return an optimal solution.
False. Depth first search has no guarantees of optimality. Further, it measures paths in length and not cost.
 - (ii) [true or false] Breadth-first graph search is guaranteed to return an optimal solution.
False. Breadth first search has no guarantees of optimality unless the actions all have the same cost, which is not the case here.
 - (iii) [true or false] Uniform-cost graph search is guaranteed to return an optimal solution.
True. UCS expands paths in order of least total cost so that the optimal solution is found.
 - (iv) [true or false] Greedy graph search is guaranteed to return an optimal solution.
False. Greedy search makes no guarantees of optimality. It relies solely on the heuristic and not the true cost.
 - (v) [true or false] A* graph search is guaranteed to return an optimal solution.
True, since the heuristic is consistent in this case.
 - (vi) [true or false] A* graph search is guaranteed to expand no more nodes than depth-first graph search.
False. Depth-first graph search could, for example, go directly to a sub-optimal solution.
 - (vii) [true or false] A* graph search is guaranteed to expand no more nodes than uniform-cost graph search.
True. The heuristic could help to guide the search and reduce the number of nodes expanded. In the extreme case where the heuristic function returns zero for every state, A* and UCS will expand the same number of nodes. In any case, A* with a consistent heuristic will never expand more nodes than UCS.

(b) Let $h_1(s)$ be an admissible A* heuristic. Let $h_2(s) = 2h_1(s)$. Then:

- (i) [true or false] The solution found by A* tree search with h_2 is guaranteed to be an optimal solution.
False. h_2 is not guaranteed to be admissible since only one side of the admissibility inequality is doubled.
- (ii) [true or false] The solution found by A* tree search with h_2 is guaranteed to have a cost at most twice as much as the optimal path.
True. In A* tree search we always have that as long as the optimal path to the goal has not been found, a prefix of this optimal path has to be on the fringe. Hence, if a non-optimal solution is found, then at time of popping the non-optimal path from the fringe, a path that is a prefix of the optimal path to the goal is sitting on the fringe. The cost \bar{g} of a non-optimal solution when popped is its f-cost. The prefix of the optimal path to the goal has an f-cost of $g + h_0 = g + 2h_1 \leq 2(g + h_1) \leq 2C^*$, with C^* the optimal cost to the goal. Hence we have that $\bar{g} \leq 2C^*$ and the found path is at most twice as long as the optimal path.
- (iii) [true or false] The solution found by A* graph search with h_2 is guaranteed to be an optimal solution.
False. h_2 is not guaranteed to be admissible and graph search further requires consistency for optimality.

(c) The heuristic values for the graph below are not correct. For which single state (S, A, B, C, D, or G) could you change the heuristic value to make everything admissible and consistent? What range of values are possible to make this correction?

State: Range:



Q3. Formulation: Holiday Shopping

You are programming a holiday shopping robot that will drive from store to store in order to buy all the gifts on your shopping list. You have a set of N gifts $G = \{g_1, g_2, \dots, g_N\}$ that must be purchased. There are M stores, $S = \{s_1, s_2, \dots, s_M\}$ each of which stocks a known inventory of items: we write $g_k \in s_i$ if store s_i stocks gift g_k . Shops may cover more than one gift on your list and will never be out of the items they stock. Your home is the store s_1 , which stocks no items.

The actions you will consider are travel-and-buy actions in which the robot travels from its current location s_i to another store s_j in the fastest possible way and buys whatever items remaining on the shopping list that are sold at s_j . The time to travel-and-buy from s_i to s_j is $t(s_i, s_j)$. You may assume all travel-and-buy actions represent shortest paths, so there is no faster way to get between s_i and s_j via some other store. The robot begins at your home with no gifts purchased. You want it to buy all the items in as short a time as possible and return home.

For this planning problem, you use a state space where each state is a pair (s, u) where s is the current location and u is the set of unpurchased gifts on your list (so $g \in u$ indicates that gift g has not yet been purchased).

(a) How large is the state space in terms of the quantities defined above?

$M \times 2^N$. You are in one of M places (simple index from 1 to M), and have not purchased some subset of N items (binary vector of size N).

(b) For each of the following heuristics, which apply to states (s, u) , circle whether it is admissible, consistent, neither, or both. Assume that the minimum of an empty set is zero.

- (**[neither]** / admissible / consistent / both) The shortest time from the current location to any other store:
 $\min_{s' \neq s} t(s, s')$
- (neither / admissible / consistent / **[both]**) The time to get home from the current location:
 $t(s, s_1)$
- (neither / admissible / consistent / **[both]**) The shortest time to get to any store selling any unpurchased gift:
 $\min_{g \in u} (\min_{s' : g \in s'} t(s, s'))$
- (neither / **[admissible]** / consistent / both) The shortest time to get home from any store selling any unpurchased gift:
 $\min_{g \in u} (\min_{s' : g \in s'} t(s', s_1))$
- (**[neither]** / admissible / consistent / both) The total time to get each unpurchased gift individually:
 $\sum_{g \in u} (\min_{s' : g \in s'} t(s, s'))$
- (**[neither]** / admissible / consistent / both) The number of unpurchased gifts times the shortest store-to-store time:
 $|u| (\min_{s_i, s_j \neq s_i} t(s_i, s_j))$

Remember, a consistent heuristic doesn't decrease from state to state by more than it actually costs to get from state to state. And of course, a heuristic is admissible if it is consistent. If you're confused, remember: the problem defines the minimum of an empty set as 0.

1. This heuristic does not return 0 in the goal state (s_1, \emptyset) , since it gives the minimum distance to any store *other than the current one*.
2. We'll always need to get home from any state; the distance to home from home is 0; and this heuristic does not decrease by more than it costs to get from state to state.
3. We'll always need to get that last unpurchased item, and taking the min distance store guarantees that we underestimate how much distance we actually have to travel. It is consistent because the heuristic never diminishes by more than what is travelled.
4. We'll always need to get home from getting the last unpurchased item, and taking the min underestimates the actual requirement. What makes this heuristic inconsistent is that when we visit the last store to pick up the last unfinished item, the value of the heuristic goes to 0. Let's say the graph looks like this: $s_3 \xrightarrow{-1} s_2 \xrightarrow{-5} s_1$, with s_2 containing the last item. From s_3 , the heuristic is 5, but from s_2 , the heuristic is now 0, meaning that traveling from s_3 to s_2 decreases the heuristic by 5 but the actual cost is only 1.
5. This can overestimate the actual amount of work required.
6. Same.

You have waited until very late to do your shopping, so you decide to send an swarm of R robot minions to shop in parallel. Each robot moves at the same speed, so the same store-to-store times apply. The problem is now to have all robots start at home, end at home, and for each item to have been bought by at least one robot (you don't have to worry about whether duplicates get bought). Hint: consider that robots may not all arrive at stores in sync.

(c) Give a minimal state space for this search problem (be formal and precise!)

We need the location of each robot at each time. At a given time, a robot can either be at one of M stores, or in any of $(T - 1)M$ transition locations, where T is the maximum travel distance between two stores. Thus, the location of each robot takes $(MT)^R$. We also need the set of items purchased (2^N). Therefore, the size of each state is: $(MT)^R \times 2^N$.