CS 440: Introduction to AI

Homework 1

Due: Tuesday September 8th

Your answers must be concise and clear. Explain sufficiently that we can easily determine what you understand. We will give more points for a brief interesting discussion with no answer than for a bluffing answer.

Solutions will be posted no sooner than two days after the due date. Homework will be accepted until that point with a penalty of 10% per day that it is late. No assignments will be accepted after the solutions have been posted. Late homework will only be accepted in class, during office hours, or electronically by email to the TA.

You are expected to do each homework on your own. You may discuss concepts with your classmates, but there must be no interactions about solutions. You may consult the web but the work handed in must be done on your own.

The penalty for cheating on any assignment is straightforward. On the first occurrence, you will receive a zero for the assignment and your course grade will be reduced by one full letter grade. A second occurrence will result in course failure.

1) Suppose that someone wants to drive from his home to Siebel Center and there are three available paths. This person is considering the following alternative models to estimate the driving time:

- **a.** The distance between the home and Siebel Center in each of the paths.
- **b.** The number of stop lights and the time each one is green, yellow and red.
- **c.** Expected time you will have to wait on red lights.
- **d.** Usual time conditions in each path.
- **e.** Probability that there is an accident in the path.
- **f.** Average time it takes to get to Siebel Center using each path.

If you had to implement each of the models *a-f* would you implement it as an analytical or empirical? Explain your reasons. (For some items both answers are acceptable, depending on justification).

2) Consider the search tree below. Node A is the root of the tree and double circled nodes are goal states. The number by the edged indicates the cost of applying an operator and the number below the node label represents the heuristic function *h*(n) for that node.

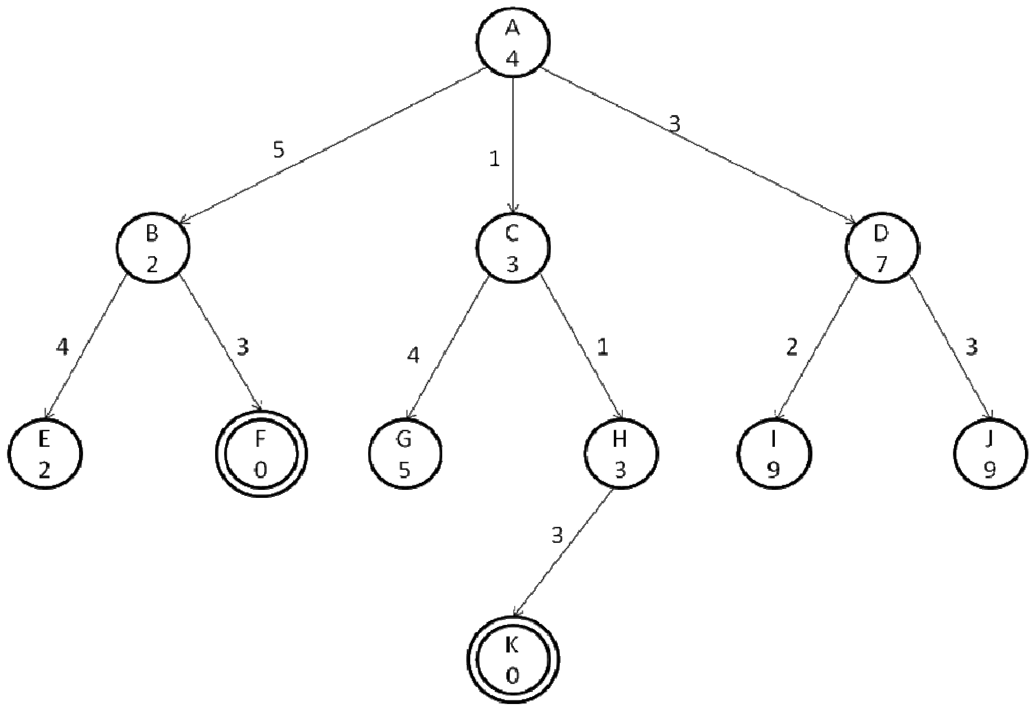

a) For each search strategy listed bellow, give the order the nodes are visited, the goal state found (if any) and the state of the queue at each step of the search (as shown in lecture 3). When nodes are expanded, news nodes are added to the queue in ascending alphabetical order whenever the order is indifferent for the search algorithm.

i. Depth First Search

ii. Breath First Search

iii. Greedy Search/Best First Search

iv. Uniform Cost

v. A*

b) In this example, is the Heuristic *h* admissible? Is it consistent? Explain your answers in a sentence or two.

c) For each of the methods *i-v* determine if it is guaranteed to find the best solution and, in a sentence or two, explain why or why not.

3) Consider the grid world bellow:

	Rocky	Goal
Start		Rocky
	Rocky	

In this grid world an agent is allowed to perform 4 operators: Move North, Move South, Move West and Move East. Moving into a wall causes a "time space wrap", i.e., Moving South from the bottom line causes the agent to remain in the same column but to move to the upper line. Moving West from the leftmost line causes the agent to remain in the same line but move to the rightmost column. The same logic is applied for Move East and Move North.

The objective is to move from the Start space to the goal space. Moving into a normal space costs 1 and moving into a rocky space costs 4. The start and goal spaces are considered normal spaces.

a) **i.** Define a data structure that is adequate to represent the states of this problem.

ii. Define the operator "Move South" by specifying its preconditions and effects.

b) Suppose that the robot is allowed to choose any order in each step to expand the states during the search, i.e., on a depth first search the robot could chose Move North from the first state and then Move West. For the following consider only the luckiest possible order:

i. What is the fewest possible number of nodes that must be visited if depth-first search is performed? Is depth-first search guaranteed to reach the goal state? Would an optimal solution be found in this case?

ii. What is the fewest possible number of nodes that must be visited if breadth-first search is performed? Is breadth-first search guaranteed to reach the goal state? Would an optimal solution be found in this case?

iii. What is the fewest possible number of nodes that must be visited if uniform cost search is performed? Is uniform cost search guaranteed to reach the goal state? Would an optimal solution be found in this case?

c) If we used the Manhattan distance as a heuristic, would A* be admissible? Can you think of any other heuristic that would be admissible?

4) Suppose we have an heuristic $h_1(n)$ such that h $0.9h^*(n) \leq h_1(n) \leq 1.5h^*(n)$ and a heuristic $h_2(n)$ such that $h_2(n) \leq 0.1h^*(n)$

a) Is the heuristic $h_1(n)$ admissible?

b) Is the heuristic $h_2(n)$ admissible?

c) What would be the advantage of using $h_1(n)$ instead of $h_2(n)$? Could there be any disadvantage?

d) If any of the heuristics $h_1(n)$ or $h_2(n)$ are not admissible, can they be made admissible? If so, how and why might it work?