Artificial Intelligence

Qualifying Exam. Fall 2005.

This is a closed book closed notes exam.

1. Search.

1.1. (15 percent) For the state-space below, the path costs are shown on the links. The start state is S and the goal state is G. All search algorithms are assumed to obey the following rules:

- when a node is expanded in the search tree, the parent of the node is not added as a child
- when multiple nodes have the same cost on the queue, they are placed in the priority queue in alphabetical order (i.e., A goes in front of B, etc).

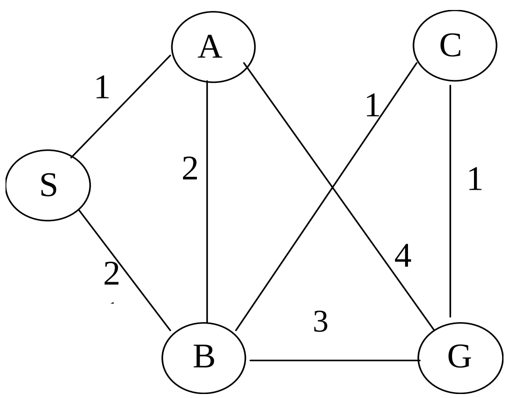

Suppose that the following heuristic is defined for the state space above:
h(S) = 2; h(A) = 2; h(B) = 2; h(G)=h(C) = 0.

Write down the sequence of nodes as they will be expanded by

(a) Depth First Search
(b) Iterative Deepening Search
(c) Greedy Best First Search
(d) A* Search
(e) Hill Climbing (Gradient descent) Search

Note that you only need to mention nodes that are being popped from the priority queue but not the whole fringe.

1.2. (10 percent) Are the following assertions true? Explain your answer.

(a) For the A* algorithm, if h1(node) is always less than h2(node), where both h1 and h2 are admissible heuristics, then A* with h1 will find the goal at least as fast as will A* with h2. (By at least as fast we mean with the same or less number of nodes expanded)

(b) The best possible heuristic for the A* search is to have h(node) = the true (minimum) path cost of going from start S to goal G via that node, i.e. h will be admissible AND A* will expand the minimum possible number of nodes before hitting the goal.

2. Resolution Refutation Procedure in First Order Logic (FOL).

Consider the following sentences:

(1) John likes all kinds of food.
(2) Apples are food.
(3) Anything anyone eats and isn't killed by is food.
(4) Bill eats peanuts and isn't killed by them.

(a) (10 percent) Convert these statements to expressions in first-order logic.

(b) (15 percent) Prove using the Resolution Refutation procedure that John likes peanuts.

3. Bayesian Networks.

Fire in a building produces smoke and turns on the alarm. When the alarm turns on people evacuate the building. The probability of fire is 0.01, the probability of smoke when there is fire is 0.99. The probability of people leaving the building when the alarm turns on is 0.95. The probability that the alarm turns on when there is fire is 0.7.

(a) (7 percent) Show the corresponding Bayesian network and conditional probability tables.

(b) (6 percent) Suppose you see people leaving the building. What is the probability that there is smoke? Show the derivations, not just the numerical value.

(c) (6 percent) Suppose instead that you see smoke. What is the probability that people will evacuate the building? Show the derivations, not just the numerical value

(d) (6 percent) Finally, suppose that the alarm turns on. What is the probability that there is smoke? Show the derivations, not just the numerical value.

4. Planning.

Consider the following STRIPS operators:

Operator	Preconditions	Effects
X	(none)	B
Y	(none)	C, ¬E, ¬B
Z	B	E
W	C	A

The goal is to satisfy sentences A and E.

(a) (10 percent) Use the partial-order algorithm to construct a plan. Be sure to make it clear in your diagram which arcs are of which types (causal or temporal). (Make sure that each step in your plan has all of its preconditions and effects listed, and all preconditions have causal links pointed to them).

(b) (6 percent) Is your plan a total plan? How many linearizations of your partially-ordered plan are there?

(c) (4 percent) List out the steps from one linearization of your plan.