C1 Artificial Intelligence (25 points)

Uncertainty. (17 points) We are given the following training set and testing set, assuming binary values:

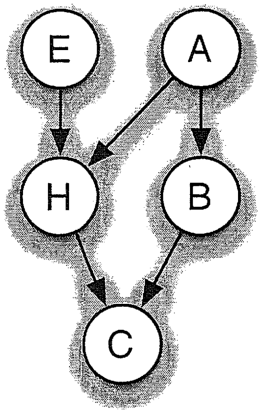

A	P(A)
T	0.7

E	P(E)
T	0.2

A	P(B)
T	0.2
F	0.6

A	E	P(H)
T	T	0.1
T	F	0.2
F	T	0.4
F	F	0.8

B	H	P(C)
T	T	0.1
T	F	0.3
F	T	0.5
F	F	0.8

1a.(17 points) Compute $P(A, \neg E|C)$. Show all of your work.

Search. (8 points)

1b. (4 points) Stochastic beam search is a local search method. Describe precisely and in detail exactly how stochastic beam search differs from steepest ascent hill climbing.

1c. (4 points) What advantage does stochastic beam search have over steepest ascent hill climbing?

C2 Artificial Intelligence (25 points)

Search. (12 points) Iterative lengthening search is an iterative analog of uniform cost search. The idea is to use increasing limits on path cost. If a node is generated whose path cost exceeds the current limit, it is immediately discarded. For each new iteration, the limit is set to the lowest path cost of any node discarded in the previous iteration. Assume that the costs are always positive.

2a. (3 points) Show that this algorithm is optimal for general path costs.

2b. (4 points) Consider a uniform tree with branching factor b, solution depth d, and unit step costs. How many iterations will iterative lengthening require? How many nodes will be generated (big-O notation)?

2c. (5 points) Now consider step costs drawn from a continuous range [0, 1] with a minimum positive cost ε. How many iterations are required in the worst case?

Planning. (13 points) Your ceiling light is controlled by two switches. As usual, changing either switch changes the state of the light. Assume all bulbs work. The light only works if there is a bulb in the socket, but you have no way to add a bulb. Initially the light is off and there is a bulb in the socket.

2d. (5 points) Formalize this situation in situational calculus. (Looks like FOPC; don't plan, just formalize.)

2e. (5 points) Formalize this situation in STRIPS rules. (Looks like STRIPS rules; don't plan, just formalize.)

2f. (3 points) Briefly describe a partial order plan for turning on the light. (Maybe a diagram)

C3 Artificial Intelligence (25 points)

Game Trees. (10 points) Here is a snapshot of a game tree that would be generated by (left-to-right) depth-first search to one ply. The numbers are the evaluations (for Max) of the resulting positions. The arcs across Min's branches are to remind you that Max must consider the AND of Min's choices but the OR of his own.

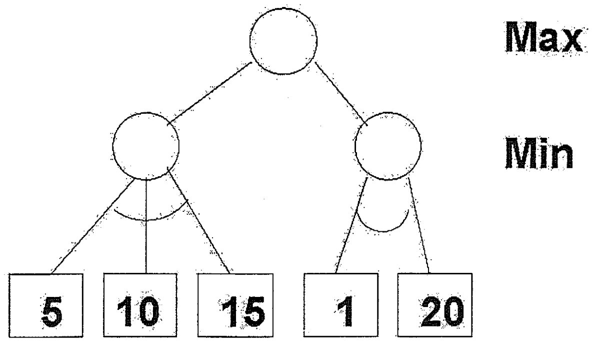

3a. (3 pts) Now say this tree is actually being generated by the search with α-βpruning. Do any nodes get pruned and if so which? What is the final backed up value for Max?

3b. (3 pts) Draw the tree that results if the nodes are generated in an order that yields the minimum amount of pruning. Show which nodes are pruned, if any. What is the final backed up value for Max?

3c. (4 pts) Is there an order for node generation that achieves more pruning?

Planning. (15 points)

3d. (3 points) If GraphPlan terminates with a successful, 3-action plan in the first iteration, what constraints are there on the order in which the actions must be executed?

3e. (9 points) Consider the following domain:

OP: Fizz	OP: Fuzz
Pre:	Pre:
Post: $C \wedge A$	Post: $D \wedge \neg A$

The goal is $C \wedge D$. What solution is returned by POP? What solution is returned by GraphPlan? (If there are multiple answers, give them.) What does this example reveal about the expressive power of the solution descriptions in the two algorithms?

3f. (3 points) It has been suggested that the first phase of GraphPlan be used as a heuristic function for forward search in the following way: Given a state s and goal g, run the graph-construction phases of GraphPlan until all the components are present and not mutex in the last layer. Let n be the number of action layers in the graph. We will let n be the heuristic value for s. Is this an admissible heuristic?

C4 Artificial Intelligence (25 points)

Decision Trees. (11 points) We are given the following training set and testing set, assuming binary values:

Training

A	B	C	Class
1	1	1	-
1	1	0	+
0	1	1	+
0	0	0	-

Testing

A	B	C	Class
1	0	0	+
0	1	1	+
0	1	0	-

4a.(8 points) Use Information Gain to create a Decision Tree for the training set, and classify each example in the testing set using your tree. Show your work for full credit; you do not need to calculate the exact information gain values. In the case of a tie, choose A over B over C, and "+" over "-".

4b. (3 points) Information Gain is not guaranteed to find the smallest tree (i.e. the shallowest tree that does the fewest tests). Find a smaller tree than you found in part a that fully describes the training set.

Search. (14 points) Consider the following maze. A robot starts on a white square and must find a path to a designated goal square that is also white. The robot can move from a white square (1) horizontally to an adjacent white square (2) vertically to an adjacent white square (3) diagonally to an adjacent white square. For example, from the square numbered "3" the robot can move horizontally to 4, vertically to 2 or 6, and diagonally to 1 or 5.

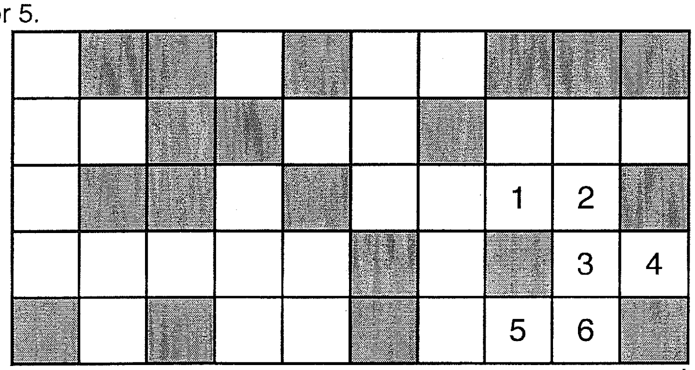

The cost of a horizontal or vertical move is 1 and the cost of a diagonal move is $\sqrt{2}$.

For each of the following, determine whether the heuristic is admissible and prove it.

4c. (7 points) *Manhattan Distance.* The Manhattan Distance between two points (x_1, y_1) and (x_2, y_2) is

$$[|x_2 - x_1| + |y_2 - y_1|]$$

4d. (7 points) L-distance. The L-distance between two points (x_1, y_1) and (x_2, y_2) is

$$\max(|x_2 - x_1|, |y_2 - y_1|)$$