CS 540: Introduction to Artificial Intelligence

Midterm Exam: 7:15-9:15 pm, March 12, 2008
Room 1240 CS & Stats

CLOSED BOOK
(one sheet of notes and a calculator allowed)

Write your answers on these pages and show your work. If you feel that a question is not fully specified, state any assumptions you need to make in order to solve the problem. You may use the backs of these sheets for scratch work.

Write your name on this and all other pages of this exam. Make sure your exam contains ***five*** problems on ***nine*** pages.

Name _____

Student ID _____

Problem	**Score**	**Max Score**
1	_____	**33**
2	_____	**28**
3	_____	**13**
4	_____	**14**
5	_____	**12**
TOTAL	_____	**100**

Name: _____

Problem 1 – Decision Trees and Experimental Methodology (33 points)

Assume that you are given the set of labeled *training examples* below, where each of three features has four possible values: *a*, *b*, *c*, or *d*. You choose to apply the ID3 algorithm to this data.

	F1	*F2*	*F3*	*Output*
ex1	*b*	*d*	*c*	+
ex2	*c*	*d*	*b*	-
ex3	*c*	*a*	*c*	+
ex4	*b*	*a*	*b*	-
ex5	*a*	*b*	*a*	+
ex6	*a*	*b*	*a*	-
ex7	*d*	*c*	*c*	+

a) What score would the *information gain* calculation assign to each of the features, when deciding which feature to use as the root?
Be sure to show all your work (use the back of this or the previous sheet if needed).

Which feature would be chosen? _____

b) Regardless of your answer to Part a, assume that F1 is chosen as the root node. Show the recursive calls to ID3, if any, that would result from this choice. Be sure to show all the arguments in these recursive calls (but do NOT perform the calculations needed in them).

Name: _____

c) Describe how pruning the tree ID3 produces, as done in HW1, can be viewed as performing a standard AI state-based search (the components of which appear below).
Be sure to briefly explain your answers.

State Representation

Start State

Operators (or Actions)

Goal Test (if any)

Search Strategy

Heuristic Function Used (if any)

Name: _____

d) Describe one (1) good way to create an *ensemble* of decision trees.

How would we use this ensemble to label new examples?

e) Briefly explain how and why each of these is used when learning a good decision tree.

A training set

A tuning set

A testing set

Name: _____

Problem 2 – Search (28 points)

a) Consider the search space below, where *S* is the start node and *G1*, *G2*, and *G3* satisfy the goal test. Arcs are labeled with the cost of traversing them and the estimated cost to a goal is reported inside nodes.

For each of the following search strategies, indicate which goal state is reached (if any) and list, *in order*, <u>all</u> the states *popped off of the OPEN list*. When all else is equal, nodes should be removed from OPEN in alphabetical order.

Breadth First

Goal state reached: _____ States popped off OPEN: _____

Beam (using f = h and a beam width of 1)

Goal state reached: _____ States popped off OPEN: _____

Iterative Deepening

Goal state reached: _____ States popped off OPEN: _____

A*

Goal state reached: _____ States popped off OPEN: _____

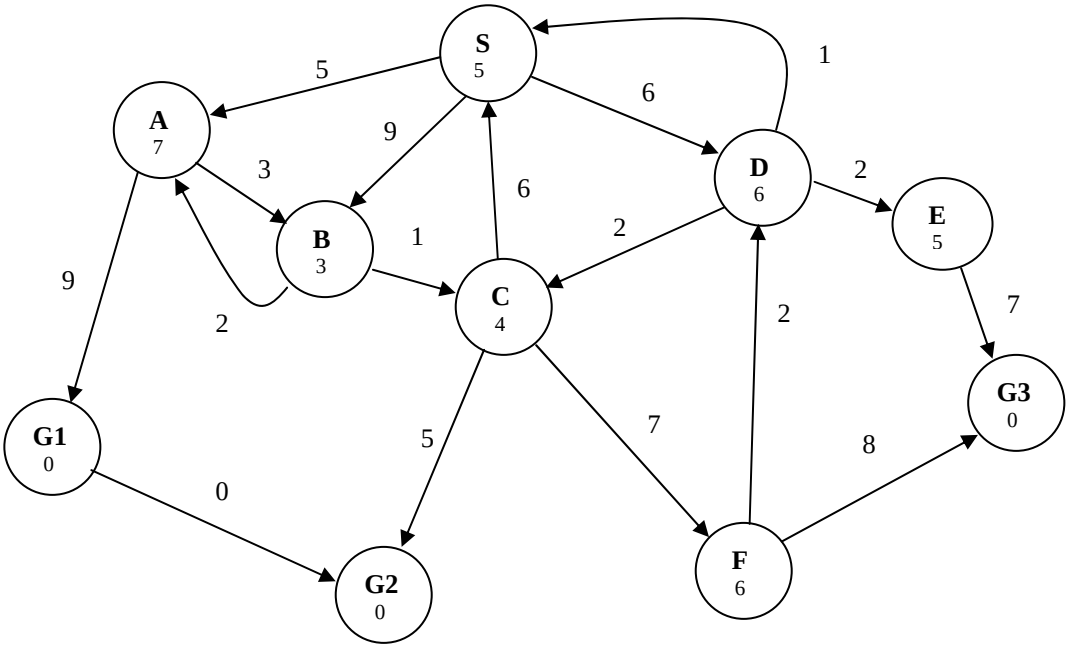

Name: _____

b) Using the same search space as in Part a, consider using *Simulated Annealing* as your search strategy. Assume the current *temperature* is 10.

   i. If you are at Node *S* and simulated annealing has randomly selected node *B* for consideration, what is the probability this node is accepted?

   ii. If you are at Node *S* and simulated annealing has randomly selected node *A* for consideration, what is the probability this node is accepted?

c) Now imagine that you wish to run a *Genetic Algorithm* on Part a's search space. What is the *most* important extra information that you would need to provide? Justify your answer.

Name: _____

Problem 3 – Game Playing and Expected Values (13 points)

a) Apply the *minimax* algorithm to the partial game tree below, where it is the maximizer's turn to play and the game does not involve randomness. The values estimated by the static-board evaluator (SBE) are indicated in the leaf nodes (higher scores are better for the maximizer). Write the estimated values of the intermediate nodes inside their circles, and indicate the proper move of the maximizer by circling one of the root's outgoing arcs.

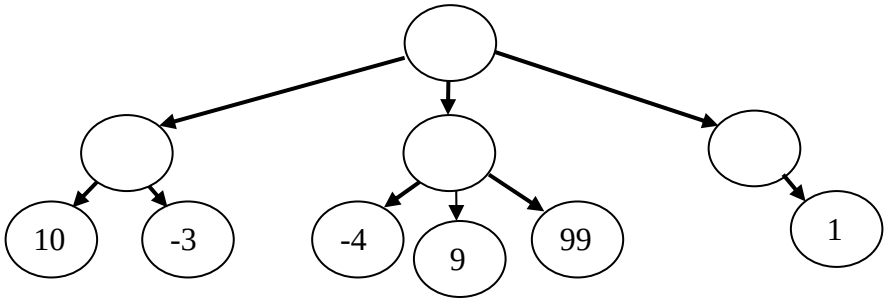

b) List one (1) leaf node in the above game tree whose SBE score need not be computed. Explain why.

c) Imagine that a large bag contains a total of 1000 red, white, or blue poker chips. The red are worth $1, the white $5, and the blue $25. The person holding the bag shakes it thoroughly, and then lets a blindfolded volunteer pull out 100 chips. Fifty (50) are red, 40 are white, and 10 are blue.

   i. What is the most you would pay in order to get one (1) randomly selected chip from this bag? Show your work.

   ii. According to Shannon's definition, how much *information* would you expect to get by being told the color of a randomly selected chip?

Name: _____

Problem 4 – Miscellaneous Questions (14 points)

What is one (1) strength of *beam* search compared to *breadth-first* search? One (1) weakness? Briefly explain your answers.

Strength (of beam search compared to breadth-first search)

Weakness

Show a specific search space where a non-admissible heuristic prevents A* search from finding an optimal solution (hint: limit yourself to at most four (4) nodes). Explain your answer.

Explain what *learning curves* are and give one (1) good reason why they are of interest.

Name: _____

Problem 5 – Miscellaneous Short Answers (12 points)

Briefly describe each of the following AI concepts and explain each's significance.

Turing Test

Description:

Significance:

Crossover

Description:

Significance:

Feature Space

Description:

Significance:

Random Restarts

Description:

Significance:

Have a good spring break!