CPE/CSC 480 ARTIFICIAL INTELLIGENCE

FINAL EXAM

FALL 2006

SECTION 3

PROF. FRANZ J. KURFESS

CAL POLY, COMPUTER SCIENCE DEPARTMENT

This is the Fall 2006 midterm exam for the CPE/CSC 480 class. You may use textbooks, course notes, or other material, but you must formulate the text for your answers yourself. The use of calculators or computers is allowed for viewing documents or for numerical calculations, but not for the execution of algorithms or programs to compute solutions for exam questions. The exam time is two hours, 50 minutes.

Student Name:

Signature: **Date:**

PART 1: MULTIPLE CHOICE QUESTIONS

Mark what you think is the best answer for the questions below. I tried to formulate the questions and answers such that there is only one suitable answer. Each question is worth 3 points, for a total of 30 in this part.

a) Why are many of the games discussed in the context of AI also referred to as *adversarial search*?
- ❑ because the player makes a move, and then the opponent (or adversary) must conduct a search for the next possible move
- ❑ because the two (or more) adversaries involved usually cooperate in the search to find the solution
- ❑ because search methods can be adapted to consider a zero-sum utility function where an advantage for one player results in a disadvantage for the other player
- ❑ it is basically a variation of bi-directional search, where the two opponents start from different initial states, and compete to find the goal

b) What is the *utility function* in the context of game-playing programs?
- ❑ a degree of uncertainty, introduced by the presence of an opponent or by chance elements
- ❑ It returns an estimate of the expected utility from a given game configuration.
- ❑ It calculates a numeric value for the outcome of a game.
- ❑ It determines when the game is over.

c) In the general formulation of the *minimax strategy*, which of the following steps is the most unrealistic one in practice?
- ❑ the generation of the whole game tree
- ❑ the calculation of the utility values for the terminal states
- ❑ the stepwise calculation of the utility values of higher-level nodes from the next lower level
- ❑ the assumpation that max (min) will select the node with the highest (lowest) value

d) What is the main difference between the minimax and alpha-beta methods?
- ❑ only minimax is guaranteed to find the best move for Max
- ❑ alpha-beta pruning only works for randomly ordered successor nodes
- ❑ alpha-beta pruning typically examines fewer nodes than minimax
- ❑ minimax only works if the whole search tree can be computed before the respective best moves for Max and Min are determined

e) In which of the following games do players have to deal with *imperfect* (or incomplete) information?
- ❑ Backgammon
- ❑ Go.
- ❑ Minesweeper
- ❑ Tic-Tac-Toe

f) Which of the following statements characterizes *propositional logic* (in contrast to predicate logic)?

- ❑ The world is described through sentences specifying individual objects with properties, and relations between the objects.
- ❑ The world is described through sentences consisting of constants, symbols, connectives, and parentheses.
- ❑ A simple logic in which truth tables are the only way of proving sentences.
- ❑ A logic that relies on resolution as the only sound inference rule.

g) What is the role of a *universal quantifier* ∀ in a predicate logic sentence?

- ❑ It allows statements about some objects in a collection.
- ❑ It allows general statements about every object in a collection.
- ❑ It is used in the specification of the semantics for terms.
- ❑ It can be used to make statements about quantitative aspects of objects, such as length, weight, temperature, etc.

h) Is the concept of *Probably Approximately Correct (PAC)* learning suitable for an agent trying to learn the answers for multiple-choice questions?

- ❑ Yes, as long as there are enough sample questions to train the agent.
- ❑ No, because there are no "approximately correct" answers for multiple-choice questions: Each answer is either correct or incorrect.
- ❑ This answer is probably approximately correct.
- ❑ This question cannot be answered in a multiple-choice manner.

i) What is the main reason that it is so much harder to define a learning algorithm for multi-layer neural networks than for perceptrons?

- ❑ The weights for the connections leading to the hidden units are fixed, and cannot be changed by the learning algorithm.
- ❑ It is difficult to determine how to change the weights for the connections leading to the hidden layer units.
- ❑ They are constructed from different types of neurons.
- ❑ They usually have more weights that need to be adjusted.

j) Which of the following statements describes the term of "back-propagation" for learning in multi-layer, feed-forward neural networks?

- ❑ Activation values are propagated from the input nodes through the hidden layers to the output nodes.
- ❑ Activation values determined by the difference between the desired and the calculated output are propagated from the output nodes through the hidden layers towards the input nodes.
- ❑ Weights on the arcs are modified based on values propagated from input nodes to output nodes.
- ❑ Weights on the arcs are modified based on values propagated from output nodes to input nodes.
- ❑ Arcs in the network are modified, gradually shortening the path from input nodes to output nodes.
- ❑ Weights on arcs from the input nodes are compared to the weights on the arcs coming into the output nodes, and then these weights are modified to reduced the difference.

PART 2: SHORT QUESTIONS

In this part of the exam, you should answer the questions in about one or two paragraphs. Please note that the number of points differs for the questions.

1. For most games, the search space is so large that it is beyond the scope of current computers to expand the whole tree, and the search has to be cut off at some level. Is it possible to use alpha-beta pruning in combination with such a cut-off? In other words, is it necessary to explore individual branches to the end points (leaves) for alpha-beta pruning to be applied during the game? Explain your answer!

   10 points

   ❑ Yes, alpha-beta pruning can be combined with cut-offs.
   ❑ No; it requires the full expansion of at least some branches.

   Explanation:

2. Search vs. Reasoning for Problem Solving: The methods we discussed in the context of search, and those related to reasoning are suitable for different types of problems. Identify types of problems and example domains for which the respective methods are better suited, and describe their strenghts and weaknesses.

10 points

Note: This is a relatively broad question. If you want to, you can concentrate on a specific problem domain (e.g. the Wumpus World), but if possible, you should also consider the more general aspects.

Aspect	*Search*	*Reasoning*
Problem types and example domains	•	
Strengths		
Weaknesses	•	•

3. What is the search space and what is the search method used by the *back-propagation* algorithm for training neural networks? What is a frequently used analogy to visually represent the search space and search method?

10 points

a) Search Space:

b) Search Method:

c) Visualization:

PART 3: DECISION TREES FOR HOLIDAY GIFTS

At this time of the year, many people are spending a lot of money and energy on selecting and buying holiday gifts for friends, relatives, and associates. Instead of agonizing over this, one person decided to collect data and construct a decision tree for this. It is a first attempt, so there is room for improvement.

The data are based on last year's decisions, and the person uses the following attributes and outcome to record the decision for the set of potential recipients P_1, …, P_{15}:

- *Gift-Received* {Great | So-So | Crummy |None}: Did I receive a gift from that person the year before, and how much did I like their gift?
- *Closeness* {Close | Medium | Distant}: How close am I to this person?
- *Gift-Given* {Expensive | Moderate | Cheap | Recycled | None}: What kind of gift did I give them the year before?
- *Still-Alive* {Yes | No}: No point in giving a gift to that great-granduncle who died last spring …
- *Outcome* {Gift | No-Gift}: Suggests what to do this year for that person.

Please note that the *Gift-Received* and *Gift-Given* attributes refer to the year before the outcome was recorded: Assuming you started collecting the data last year (2005), those attributes are about gifts from the year before that (2004).

Person	*Gift-Received*	*Closeness*	*Gift-Given*	*Still-Alive*	*Outcome*
1	None	Close	Expensive	Yes	**No-Gift**
2	Great	Distant	None	Yes	**No-Gift**
3	Crummy	Close	Recycled	Yes	**Gift**
4	So-So	Medium	Moderate	Yes	**Gift**
5	Great	Distant	Expensive	No	**No-Gift**
6	Crummy	Medium	Expensive	Yes	**Gift**
7	Crummy	Close	Recycled	Yes	**Gift**
8	None	Distant	None	Yes	**No-Gift**
9	Crummy	Distant	None	No	**No-Gift**
10	Great	Close	Moderate	Yes	**Gift**
11	So-So	Close	Moderate	Yes	**Gift**
12	None	Medium	Recycled	No	**No-Gift**
13	None	Close	None	Yes	**No-Gift**
14	So-So	Distant	Recycled	Yes	**Gift**
15	Great	Medium	None	Yes	**No-Gift**

1. Construct the decision tree based on the sample set. Identify the attributes you select below, and list the outcomes for the respective values of the attribute.

   a) Identify the attributes in the sequence they are selected, and describe the reason for the selection of each attribute (e.g. "complete agreement of Value-1 with *Outcome-Gift*").

   5 points

Attribute	*Reason:*
1	
2	
3	
4	

   b) Draw the decision tree based on the attributes identified above. Write the samples that need to be considered next to each node in the decision tree, and separate them according to the outcome ("Yes" on top, "No" below). The root node is given below; replace "Attrib-1" with your first attribute, and complete the tree. If necessary, draw the tree on the empty opposite page.

   10 points

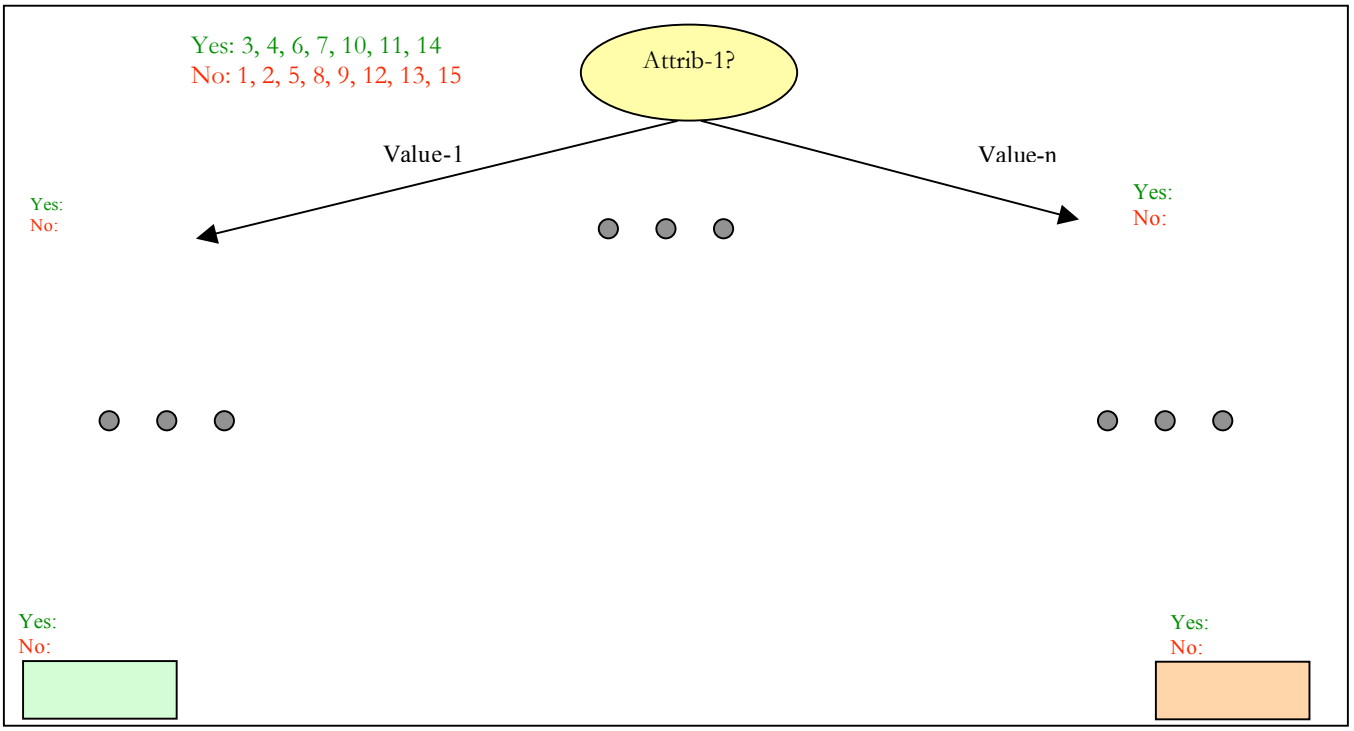

To test the decision tree before it goes into production for this season, here are a few samples. Based on your decision tree, what is the predicted outcome? What is your explanation for the prediction?

10 points

Test Case	*Gift-Received*	*Closeness*	*Gift-Given*	*Still-Alive*	*Outcome*
1	None	Medium	Recycled	Yes	?
2	Great	Close	Expensive	Yes	?
3	So-So	Medium	Recycled	Yes	?
4	Crummy	Distant	Moderate	Yes	?
5	So-So	Distant	Moderate	No	?

Explanation Test Case 1:

Explanation Test Case 2:

Explanation Test Case 3:

Explanation Test Case 4:

Explanation Test Case 5:

2. The example given above is rather simplistic. It does not distinguish between types of gifts to be given this year, for example. If you were to use the same values for the outcome as for the Gift-Given (i.e. {Expensive | Moderate | Cheap | Recycled | None}), would it still be possible to use a decision tree? Explain your answer!

5 points

Extra Credit Question: When the suggested outcome is a "recycled" gift, this recommendation needs to be taken with a grain of salt. What is the potentialy serious problem with that outcome? Can decision trees be used to prevent such problems?

Extra Credit: 3 points

3. The art of gift giving is often based on the history of gifts given to and received from a particular person. Are decision trees suitable for dealing with such information? Explain your answer!

5 points

4. Overall, do you think it is a good idea to use a decision tree for this problem? Explain your answer!

5 points

Total Points: