

INTRODUCTION TO ARTIFICIAL INTELLIGENCE

4005-750-01, 20081

MIDTERM EXAMINATION, OCTOBER 9, 2008
INSTRUCTOR: RICHARD ZANIBBI, DURATION: 60 MINUTES

NAME: _____

INSTRUCTIONS

- **The exam questions are worth a total of 100 points.**
- Hand in this sheet with your exam booklet(s).
- Remain seated in the exam room for the last 5 minutes of the exam.
- Answer all questions in the booklet provided. You may use pencil or pen, and write on the back of pages in the booklet.
- If you require clarification of a question or an additional booklet, please raise your hand.
- Write your name on all booklets that you use.

QUESTIONS

1. True or False (3 points each (15 pts)):

You may provide a brief justification for your answer (for part marks)

- (T / F) Training the parameters of a neural network is normally performed using a variation of local search.
- (T / F) Constraint satisfaction problems may be solved more quickly by identifying independent subproblems, solving them separately, and then combining the intermediate solutions to create a solution for the original problem.
- (T / F) For a statement α to be entailed by a knowledge base KB , all models satisfying α must also satisfy KB .
- (T / F) A *rational agent* is one which consistently finds the optimal solution to a problem.
- (T / F) In games of partial information such as trump card games (bridge, euchre, etc.), the optimal action will be the one which minimizes the expected loss over all possible outcomes.


2. History of AI and Agents (21 pts)

- (4 pts) When, where, and by whom was the term *Artificial Intelligence* first defined? Have the goals of Artificial Intelligence research changed significantly since that time?
- (4 pts) The course text summarizes definitions of Artificial Intelligence using two dimensions, each with two values. Provide the four resulting categories for definitions.

- (c) (7 pts) Briefly describe the protocol for Alan Turing's *Imitation Game* - who participates, and how is the game "won?" Which of the four categories for AI definitions from the previous question is the test consistent with?
- (d) (6 pts) Briefly define the following agent types (1-2 sentences each).
 - i. Simple reflex agent
 - ii. Goal-based agent
 - iii. Utility-based agent

3. Uninformed and Heuristic Search (25 pts):

- (a) (4) Name the four components used to define a problem definition that will be solved using search.
- (b) (7) Assuming a finite search space, provide the worst-case run-time and space complexity for the following algorithms, making use of the branching factor (b), solution depth (d) and maximum depth (m) of a search tree. Also, identify which of these searches are optimal.
 - i. Breadth-first
 - ii. Iterative-deepening
 - iii. Depth-first
- (c) (3) How can finite *state spaces* produce infinite *search trees*? How can this problem be avoided?
- (d) (6) Provide the order of node expansions and path cost for finding state 'e' in the graph below, using the following algorithms:
 - i. Depth-first search
 - ii. Uniform cost search
 - iii. A* using tree search and the heuristic $h(n) = 3 - \text{depth}(n)$



- (e) (5) To try and overcome problems with local minima/maxima inherent in hill-climbing, simulated annealing may be used. Describe how the simulated annealing algorithm works, and in particular how the algorithm avoids settling into local minima/maxima.

4. Minimax and Alpha-Beta Pruning (20 pts)

Consider a version of tic-tac-toe where the board is size 2x2, each player may choose to pass during their turn (not marking any square), and X moves first.

- (a) (2) Which uninformed search algorithm is *minimax-search* a variation of?
- (b) (6) Draw the game tree down to depth 2. You do not have to show nodes that are reflections or rotations of sibling nodes already shown (these are redundant).
- (c) (4) Suppose the evaluation function being used for play is the number of rows and columns on which X may win ('make a line') minus the number of rows and columns on which O may win. Mark the minimax value for all leaves and internal nodes in the game tree that you drew for part (b).
- (d) (6) Circle the nodes that would not be evaluated by the alpha-beta pruning algorithm while doing a left-to-right search of your game tree. Also, describe briefly (1-3 sentences) how alpha and beta values are used to prune nodes from the game tree.

5. Constraint Satisfaction (6 pts)

- (a) (2) Which uninformed search algorithm is *backtracking search* a variation of?
- (b) (4) Briefly explain how the *minimum remaining values* and *least constraining value* heuristics may be used to speed up backtracking search.

6. Logic (13 pts):

- (a) (2) What is the simplest (though inefficient) way to show that a statement α is entailed by a knowledge base?
- (b) (7) Convert the following propositional logic statements to conjunc-

tive normal form, and then prove that **a** is true using resolution.

$$\begin{aligned}(b \wedge c) &\rightarrow a \\ b \\ (d \wedge e) &\rightarrow c \\ e \vee f \\ (d \wedge \neg f)\end{aligned}$$

- (c) (3) Write a first-order logic statement for each of the following.
- i. All people living in New York State are friends with a dog named Buddy.
 - ii. The owner of Buddy is Reggie.