

COMPUTER SCIENCE AND ENGINEERING 415
Introduction to Artificial Intelligence

Steve Tanimoto, Instructor
Spring Quarter, 2016

Sample Midterm Exam

Instructions: This is a closed-book exam. Do not use any notes, books, calculators, or electronic devices.

When possible, write your answers within the rectangles provided. If you need to use the margins or the back of a page, you are welcome to do so, but clearly indicate on or near the rectangle where to find your answer.

1. (5 points) Name the AI program that became famous after winning a game of Jeopardy on national television.

2. (5 points) Name the AI program that became famous after beating a top-rated Go player.

3. (5 points) What does it mean for a game such as Checkers to be "solved" in the context of artificial intelligence?

4. (10 points) What are the pros and cons of using iterative-deepening depth-first search to solve a puzzle such as the Towers-of-Hanoi puzzle?

(a) What are the positive aspects?

(b) What are the negative aspects?

5. (10 points) Describe, using either Python code or pseudocode, an operator for Tic-Tac-Toe to place an X in the center of the board. Follow the same format for operators used in Assignments 3 and 4.

6. (10 points) Suppose a maze is represented by a graph with one node per room and an edge between nodes n_1 and n_2 provided that n_1 and n_2 are adjacent and there is no wall between them. Assume the rooms are square cells in a rectangular grid. Let's assume that the cost of each edge is 1. Consider the following heuristic evaluation functions for solving this maze using the A* method: Euclidean distance between n and the goal, Manhattan distance between n and the goal. We assume that the distance is taken from the center point of the room, and that the goal is also a specific room. For example, the Euclidean distance from (4,7) to (14, 2) is $\text{sqrt}(125)$, and the Manhattan distance is 15.

(a) Which of these heuristics is admissible?

(b) Which is likely to prove more efficient in terms of the number of nodes opened by A* and why?

7. (25 points) Zobrist hashing can be an effective technique for speeding up the kinds of tree search that occur in applications such as playing games like Checkers and Chess.

(a) (10 points) In this part and part (b) explain in English how Zobrist hashing computes hash values. In this part, explain what must be computed at start-up time or offline in advance.

(b) (5 points) What must be computed at run time (i.e., when the hash function is called)?

(c) (5 points) Why is it fast to compute the hash value of a state in this way?

(d) (5 points) Why are hash table collisions relatively unlikely between two similar states?

8. (15 points) Suppose we are applying Genetic Search to the Traveling Salesman Problem, and we take as the "individual" a list of cities representing a path in the TSP graph. We use two kinds of operators: crossovers (that split parent paths and recombine the pieces) and mutations (that simply pick two list elements and interchange them). The initial population is, say, 10 copies of a list of the cities in alphabetical order.

(a) (5 points) Describe a fitness function that would give preference to low-cost paths for this kind of formulation.

(b) (10 points) Suppose that the formulation above is changed in two ways. First, the initial population is changed such that we have 10 copies of a list that consists of a copies of the first city. (Here n is the number of cities in the TSP.) Second, a second type of mutation operator is added that randomly selects a position in the list and then replaces the city there with a random city from the problem. Describe a fitness function that would not only give preference to low-cost paths but would greatly prefer paths that are tours.

9. (20 points) (a) Using the pre-drawn nodes below, draw the graph of the ISA relationships corresponding directly to the following sentences. (Do not attempt to exploit the partial order properties yet.) There should be one directed graph edge per sentence.

(b) Circle any group of nodes that must be equivalent according to the antisymmetry property.

(c) Mark each redundant edge (due to any of the three partial order properties) by placing an "R" on it.

(d) Using dotted lines, show the missing edges in the graph that are implied by transitivity. (If a group of equivalent nodes is involved, due to antisymmetry, you may use one node as the representative for these missing edges. Do not draw extra dotted edges within a group of equivalent nodes.)

1. An android is an automaton.
2. An android is a bot.
3. An automaton is a contraption.
4. A bot is an automaton.
5. A car is a machine.
6. A contraption is a device.
7. A device is a machine.
8. A machine is a contraption.
9. A bot is a machine.

