15-381 Artificial Intelligence: Representation and Problem Solving
Homework 4

Out: 10/30/08
Due: 11/13/08

Instructions

This assignment is due on Thursday, Nov 13, 2008. The written portion must be turned in at the beginning of class at noon on Nov 13. Type or write legibly; illegible submissions will not receive credit. Write your name and Andrew ID clearly at the top of the assignment.

You are to complete this assignment individually. However, you are encouraged to discuss the general algorithms and ideas in the class in order to help each other answer homework questions. You are also welcome to give each other examples that are not in the assignment in order to demonstrate how to solve problems. But we require you to:

- not explicitly tell each other answers;
- not copy answers;
- not allow your answers to be copied.

In those cases where you work with one or more other people on the general discussion of the assignment and surrounding topics, we ask that you specifically record on the assignment the names of the people you were in discussion with (or "none" if you did not talk to anyone else). This will help resolve the situation where a mistake in general discussion led to a replicated error among multiple solutions. This policy has been established in order to be fair to everyone in the class. We have a grading policy of watching for cheating and we will follow up if it is detected.

Refer to the web page for policies regarding collaboration, due dates, and extensions.

1 [12 pts] Local search

Give the name of the search algorithm that results from each of the following special cases:

1. Local beam search with $k = 1$
2. Local beam search with one initial state and no limit on the number of states retained
3. Simulated annealing with $T = 0$ at all times (and omitting the termination test)
4. Genetic algorithm with population size $N = 1$

2 [30 pts] Genetic algorithm

In this question, you need to propose crossover and mutation for genetic algorithms that result in "valid" states.

1. The 8-queens problem, as discussed in class, asks you to place 8 queens on an 8x8 chessboard such that no two queens can attack each other (i.e. share the same row, column, or diagonal). The chromosome (representing an individual) and crossover for the 8-queens problem shown in the class generate chromosomes with exactly one queen per column, but often more than one queens (or none) per row. Propose a chromosome representation and crossover and mutation mechanisms that generate result in exactly one queen per column and one queen per row.

   Hint: Look at the chromosome in the slides. What makes a chromosome valid (one queen per column and per row)? Think about how you can keep the chromosome valid. For crossover, you do not have to cut chromosome strings and concatenate them together. For mutation, you do not have to change one character in the chromosome strings.

2. The traveling salesman problem (TSP) is the problem of finding the shortest route to visit a set of cities exactly once and return to the starting city. In class, we showed that TSP can be solved by a genetic algorithm, but did not present the chromosome, crossover, and mutation. Propose a chromosome representation and the corresponding crossover and mutation, such that the chromosomes generated after crossover and mutation are still valid (visiting each city exactly once).

   Hint: It is likely that you can use a solution similar to your 8-queens solution. Think about what the valid chromosome representations for these two problems have in common?

3 [40 pts] Reinforcement learning

Consider reinforcement learning for the card game *blackjack*, or *twenty-one*. The following web page contains a description of the game and a demo of blackjack AI trained by reinforcement learning. You can experiment with this demo, or play blackjack against the dealer and trained AI.

`http://lslwww.epfl.ch/~anperez/BlackJack/classes/RLJavaBJ.html`

Note: Please do not run your experiments at the last moment, because we do not know if there is any download bottleneck on this server (even though applets run locally on your computer).

1. Propose a state space, including actions and rewards. Describing the actions in plain English is enough. There is no need to draw a transition diagram or table. Note that this is a simplified version of blackjack: The deck of 52 cards is shuffled before each hand. An Ace is automatically valued as 11, unless it would cause a bust. In that case it is valued as 1. Compare your state space to the real (unsimplified) blackjack game in the casinos, and describe any limitation of your state space.

2. Simulate Q-learning using your state space by hand for three epochs (three games). Set alpha to be 0.1, gamma to be 0.9, and epsilon to be 0.1. Start with a $Q(s, a)$ matrix initialized to all zeros except for terminal states s_{term}, which get $Q(s_{term}, a) = R(s_{term})$. When there is a tie, hit by default. Assume the following games (note that it is 6+Q now instead of 4+Q in previous version).

   ```
   Game 1: start with 3 + 8, hit, given 6, hit, given 10, burst.
   Game 2: start with 6 + Q, hit, given 2, hit, given 5, burst.
   Game 3: start with J + 4, hit, given 2, stand, lose to the dealer.
   ```

3. Experiment the effect of changing the alpha (learning rate), gamma (discount factor), and epsilon (greediness) parameters on the demo program. Start with the default values, except set reward for win to be +1 and run 500 episodes instead of 1000 to save time. Modify one parameter at a time. Try a smaller value, a larger value, and the minimal and maximal values; report the winning rate of each parameter set (1 default and 12 modified parameter sets). Use 500 episodes instead of 1000 episodes to save time. What effects do these parameters have on the winning rate? Discuss potential explanation of these effects.

4. Describe how to use genetic algorithm for blackjack. Propose a chromosome representation, the crossover and mutation mechanisms, and the fitness function (or how to select the best individuals).

5. Discuss the advantages and disadvantages of Q-learning and genetic algorithm for blackjack.

4 [18 pts] Property space

A man wishes to take a wolf, a goat and a cabbage across a broad river, from the south side to the north side, in a rowboat. The boat is very small and the poor man has room in the boat for only one of these items at a time. What's worse, he cannot leave the wolf and the goat on the same bank, whether coming or going, for the wolf will surely eat the goat. Nor can he leave the goat and cabbage on the same bank because the goat will eat the cabbage. He does not know what to do and turns to the smart students in 15-381 for help.

1. Describe this problem domain in a property space representation. Write the properties or literals that capture the problem space, the operators needed, and the operators' preconditions and postconditions.

2. Write the initial state set I and goal state set G using your literals, and describe the solution path through property space by writing the operators and literals describing the journey from I to G.