

15-780: Graduate AI

Homework Assignment #1

Out: January 22, 2015
Due: February 4, 2015 5 PM

Collaboration Policy: You may discuss the problems with others, but you must write all code and your writeup independently.

Turning In: Please email your assignment by the due date to shayand@cs.cmu.edu and vdperera@cs.cmu.edu. If your solutions are handwritten, then please take photos or scan them and make sure they are legible and clear. Please submit your code in separate files so we can easily run them.

Note: We use the term nontrivial heuristic in this homework to mean a heuristic function that does not assign a constant value to every node.

1 Streets of Pittsburgh [40pts]

In this question, you will try a variety of search algorithms you learned about in class to be able to find the shortest path between any two intersections in Pittsburgh. You will be given real data exported from OpenStreetMap. OpenStreetMap labels each intersection, called a **node**, with a unique ID. (You can see where each nodeID corresponds to by going to: <http://www.openstreetmap.org/node/nodeID>.) A street is therefore an **edge** between two nodes. All the streets in a large part of Pittsburgh are given to you in `edges.txt`. Each line specifies a single weighted undirected edge in the following format:

```
node1ID node2ID distance_in_km
```

1. **Programming** Write a program that can find a path between any two nodes using breadth-first search, uniform-cost search, iterative deepening, and A*. In particular write four functions as follows:

```
bfs(start_nodeID, goal_nodeID), ucs(start_nodeID, goal_nodeID),  
iterdeep(start_nodeID, goal_nodeID), astar(start_nodeID, goal_nodeID).
```

Each of these functions should read in the file `edges.txt` from the same directory as your code. When choosing which node to visit, please break ties randomly. Your code should output a string that tells us (a) how many nodes were visited (i.e. were popped

from the queue), (b) how many nodes were on the path that was chosen, and (c) the distance of the path chosen (in kilometers). Please format the string to be as follows:
 Num nodes visited: ?
 Num nodes on path: ?
 Distance (km): ?

For A*, your heuristic function will be the Euclidean distance to the goal node. The `astar` function should read in a file `heuristic.txt` from the same directory as your code. Each line in this file gives you the distance from a particular node to a some goal node in the following format:

`nodeID distance_to_goal_in_km`

The `heuristic.txt` file we provide you with the homework provides distances to the Pittsburgh City Hall (node 105012740). You can test your code using that goal node, however we may use another goal node (and another heuristic file) when testing your code.

You may write your code in any commonly used language, but please do not use any graph libraries. Make sure to submit all of your code so we can run it!

2. **Traffic Time!** Now imagine instead of finding the path with the shortest distance, you wanted to find the path that takes the shortest time during current traffic, which you can assume is information you have access to. For this problem you may assume drivers never drive faster than the speed limit. ☺ Give a nontrivial admissible heuristic for this new problem.
3. **Real-World Routing** You shouldn't use the solutions you get from your code for driving just yet! What is one thing that is missing in the way this problem is currently formulated that makes it unusable in the real-world? Suggest an easy way to solve the problem.

2 Heuristics [15pts]

Imagine an $n \times n$ grid where there is a car in each of cells $(1, 1)$ to $(n, 1)$ —the bottom row. Every time step, each car can move up, down, left, right, or stay put. If a car does not move, at most one other car can hop over it. At most one car can be in a cell at any given time. Cars move simultaneously (so a car can move into the place of another car if the other car moves out), and two cars are not allowed to move into the same cell. Moving a car in any direction (including hopping over another car) incurs a cost of 1 whereas staying put has no cost. The total cost at any time step is the sum of the cost for each car. The goal is to move all the cars to the top row but in reverse order so that the car starting in cell $(i, 1)$ ends up in cell $(n - i + 1, n)$.

1. What is the size of the state space in big-O?

2. What is the worst-case branching factor?
3. Suppose there was only one car starting at $(i, 1)$ and the problem was to only move that car to its goal location. Give a nontrivial admissible heuristic $h_i(x_i, y_i)$ for moving car i from cell (x_i, y_i) to its goal at $(n - i + 1, n)$, and explain why it is admissible. Would it be admissible when the other $n - i$ cars are also at various positions on the grid?
4. Which of the following are admissible heuristics for the original problem and why? (Note: h_i is the heuristic you came up with in the previous part.)

(a) $\sum_i h_i$	(b) $\min(h_1, \dots, h_n)$
(c) $\max(h_1, \dots, h_n)$	(d) $n \min(h_1, \dots, h_n)$
(e) $n \max(h_1, \dots, h_n)$	

3 Convex Conundrums [15pts]

This question will ask you to prove or disprove whether certain sets are convex. You should use the definition of convex sets. Please give a formal (but not necessarily lengthy proof) for each. To show a set is not convex, you need to show that there exist two points x and y that are within the set but for some $0 \leq \lambda \leq 1$, $\lambda x + (1 - \lambda)y$ is not in the set.

1. The norm cone, i.e. $\mathcal{C} = \{(x, t) \mid \|x\| \leq t\}$
2. $\mathcal{C} = \{x + y \mid x, y \in \mathcal{D}\}$ for some convex set \mathcal{D}
3. The intersection of m convex sets, i.e. $\mathcal{C} = \bigcap_i \mathcal{C}_i$, where \mathcal{C}_i is convex for $i = 1 \dots m$
4. The union of m convex sets, i.e. $\mathcal{C} = \bigcup_i \mathcal{C}_i$, where \mathcal{C}_i is convex for $i = 1 \dots m$

4 Knowledge Representation [15 pts]

Propositional Logic In the not too distant future, a team of three CoBots is devising to split up delivering mail to NSH and GHC. Here are some possible rules for their deliveries:

- (a) If CoBot A delivers mail to NSH and does not deliver mail to GHC then CoBot B delivers mail to GHC or CoBot C delivers mail to GHC.
- (b) Either CoBot A does not deliver mail to NSH and CoBot B does deliver mail to NSH or CoBot A delivers mail to NSH and CoBot B does not deliver mail to GHC.
- (c) If CoBot B delivers mail to NSH, it is necessary that CoBot A delivers mail to GHC and either CoBot C delivers mail to NSH or CoBot A delivers mail to NSH.
- (d) If CoBot B delivers mail to NSH and CoBot A does not deliver mail to GHC then CoBot C delivers mail to either NSH or GHC and either CoBot B does not deliver mail to GHC or CoBot A delivers mail to NSH.

1. Encode statements (a)-(d) in propositional logic using the following literals:

- A: Cobot A delivers mail to NSH
- B: Cobot A delivers mail to GHC
- C: Cobot B delivers mail to NSH
- D: Cobot B delivers mail to GHC
- E: Cobot C delivers mail to NSH
- F: Cobot C delivers mail to GHC

2. Transform the statements you just wrote to Conjunctive Normal Form (CNF).

3. Using propositional tableaux prove that $\{A \leftrightarrow B, A \vee B\}$ entails $A \wedge B$.

4. (a) Consider the following statements:

- If I leave and go on vacation then I am happy.
- If I leave then I go on vacation.
- I leave.

Can we conclude that: I go on vacation and I am happy?

(b) Consider the following statements

- I am happy if and only if I won the lottery or my friend is with me.
- If it is raining my friend is not with me.
- It is raining and I am happy.

Can we conclude that: I am happy if and only if I won the lottery?

5 Constraint satisfaction problem (CSP) [15pts]

A common problem at universities is to schedule rooms for exams. The basic structure of this problem is to divide the exam period into predefined time slots and then assign courses to the available rooms at particular time slots. The assignments must be consistent with a set of constraints. For this question, you will describe how you would design various elements of a CSP program that would solve a restricted version of this problem.

For this version of the problem, a course is defined as having a name, instructor(s) and students. A room has a name and a capacity. Time slots are positive integers, and we will assume that there are three time slots that are represented as 0, 1, and 2. The constraints that must be observed are the following:

- All courses must be assigned to a room and time slot.
- No assigned room can have less capacity than the number of students in the assigned course.

- No room can be double-booked with more than one course per time slot.
- No instructor or student can be in two different rooms during the same time slot.

Using the backtracking-search algorithm presented in class as a framework for your CSP solver, provide concise answers to the following:

1. Describe the basic data structures you would use to represent the domain objects and the constraints.
2. What data structure would you use as a variable and what would you use as a value?
3. Describe a variable-ordering heuristic that could improve performance. This heuristic would be used in SELECT-UNASSIGNED-VARIABLE.
4. Describe how you would compute the possible values for a variable in ORDER-DOMAIN-VALUES.
5. Given a variable and a value assignment, describe how you would check that the assignment is consistent with the constraints. Note that this question may be straightforward, depending on how you generate the values.
6. How many representations of state do you need in your CSP solver? Remember that backtracking-search is a depth-first search that undoes its assignments when back tracking.