Homework #1: Chapters 1, 2, 3

The following exercises are due at the beginning of class on January 30. Please type your answers or neatly write them on your own paper. Each exercise will be graded for correctness, so start early and be sure you are confident in your answers. Also, remember that all work should be your own. Note, the last problem is on the back of this sheet.

1. *[15 points]* Develop a PEAS description for the following task environments:
   a) A grocery store scanner that digitally scans a fruit or vegetable and identifies it.
   b) A GPS system for an automobile. Assume that the destination has been preprogrammed and that there is no ongoing interaction with the driver. However, the agent might need to update the route if the driver misses a turn.
   c) A credit card fraud detection agent that monitors an individual's transactions and reports suspicious activity.
2. *[15 points]* For each of the agents described above, categorize it with respect to the six dimensions of task environments as described on pages 41-45 (you can omit known vs. unknown since it does not directly refer to the environment itself). Be sure that your choices accurately reflect the way you have designed your sensors and actuators. Give a short justification for each property.
3. *[10 points]* One of the problems with the table-driven agent is that the tables can get enormous. One way to reduce the table's size is to only do lookup based on the current percept, as opposed to the entire percept history. Under what conditions would this result in a rational agent? When is it better to use the entire percept history?
4. *[25 points]* The missionaries and cannibals problem is usually stated as follows. Three missionaries and three cannibals are on one side of a river, along with a boat that can hold one or two people. Find a way to get everyone to the other side without ever leaving a group of missionaries in one place outnumbered by the cannibals in that place. Assume there is no way to send the boat across without at least one person in it. Choose a suitable representation for states and then give the initial state, goal test, actions, transition model, and path cost function for this problem, making only those distinctions necessary to ensure a valid solution. Draw a diagram of the complete state space.
5. *[25 points]* Consider the 8-puzzle with the initial and goal states shown below. Use breadth-first graph search to solve this problem. Recall, that in a graph search, any state that has already been expanded will not be added to the frontier. Show your search tree, explicitly showing the puzzle grid at each node, and labeling each node with the order in which it is expanded. To save yourself some unnecessary work, you may stop as soon as you have generated the goal state (i.e., you don't need to expand any other nodes after you have found the goal state).

Initial State

2		3
1	8	4
7	6	5

Goal State

1	2	3
8		4
7	6	5

6. *[10 points]* Now consider a tree-based version of depth-first tree search that does **not** check for repeated states. Also, assume that when there are multiple deepest nodes, that one is selected randomly for expansion. If a solution is "found" as soon as it is generated, then what is the minimum number of expansions before the algorithm will find the solution to the problem in exercise #5? What is the maximum number of expansions before a solution is found?