Homework #1: Chapters 1, 2, 3, and 4

The following exercises are due at the beginning of class on February 6.

1. *[15 points]* Develop a PEAS description for the following task environments: a) NASA's Mars Rover, b) the mySimon internet shopping agent, and c) a RoboCup soccer playing agent (plays a computer-simulated soccer game)
2. *[10 points]* Create a table, like Figure 2.6 in the book (p. 43), that gives the characteristics of the following task environments: a) a solitaire card game in which you deal all cards facing up at the start, and b) a robot playing a baseball.
3. *[15 points]* Do exercise 3.7(b) from the book (p. 90). You can give the successor function by describing each action formally (i.e., precisely describe what kinds of states each action can be used in and how the state is changed when it is applied).
4. *[10 points]* Suppose you are using a search algorithm to solve the 15-puzzle problem. What are the benefits and drawbacks of using a depth first search? What are the benefits and drawbacks of using a breadth first search?
5. *[10 points]* The book shows how to derive heuristics for the 8-puzzle problem from relaxed versions of the problem (p. 107-108). Consider a new heuristic h3, which is the minimum number of moves to solve the puzzle if any tile can move directly to the blank square. Explain why h3 is at least as accurate as h1 (the number of misplaced tiles). Give an example of a state where h3 is more accurate than both h1 and h2 (the sum of Manhattan distances of each tile to its goal location).
6. *[40 points]* A hurried traveler is seeking an efficient route across Pennsylvania, from Philadelphia to Erie. Develop one route using a greedy best-first search and another route using an A* search with the heuristic estimates given in the table below. For each algorithm, show your search tree and label each node with the order in which it is expanded. For the greedy search, show the h(n) value for each node. For A*, show the f(n), g(n), and h(n) values for each node. In order to reduce unnecessary search, you can ignore moves that return you to the state you just came from, however you must show any other repeated states. How do the two solutions compare? What benefits and drawbacks of each search algorithm did you observe? Why does the A* search still work even though we have expanded repeated states?

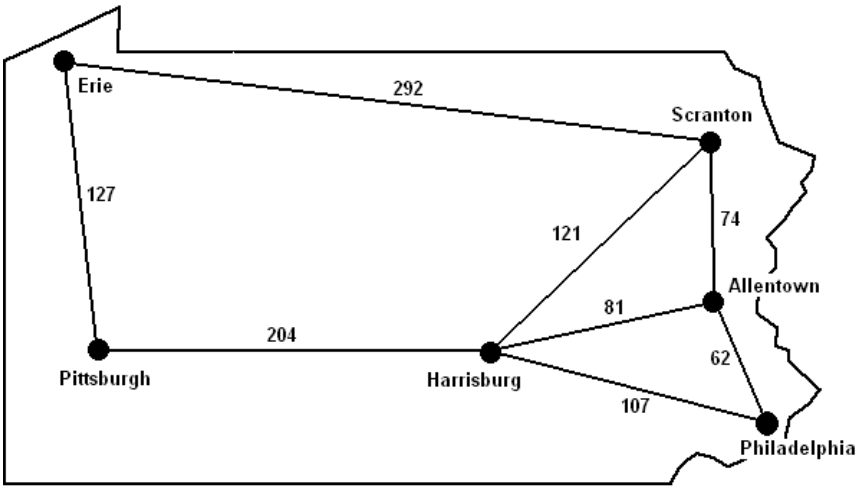

Heuristic Estimates		
Erie	Pittsburgh	100
Erie	Harrisburg	225
Erie	Scranton	250
Erie	Allentown	300
Erie	Philadelphia	350
Pittsburgh	Harrisburg	180
Pittsburgh	Scranton	225
Pittsburgh	Allentown	200
Pittsburgh	Philadelphia	225
Harrisburg	Scranton	100
Harrisburg	Allentown	60
Harrisburg	Philadelphia	85
Scranton	Allentown	65
Scranton	Philadelphia	100
Allentown	Philadelphia	55