Homework #1: Chapters 1–3.4

The following exercises are due at the beginning of class on Thursday, February 9. Each exercise will be graded for correctness, so please start early and be sure you are confident in your answers. Also, remember that all work should be your own. Note this homework is continued on the reverse side.

1. *[15 points]* Develop a PEAS description for the following task environments:
   a) A flying package delivery drone that delivers small packages to people's houses.
   b) A software agent that can play a computerized version of the following traditional game of Freecell: a standard 52 playing card deck is shuffled and eight *cascade* piles of cards are created (four with seven cards and four with six cards), with all cards face-up and visible. There are four *foundation* piles that must be built up by suit from Ace to King. The cascade piles can be built in descending order by alternate colors (e.g, a black 7 can go on a red 8) however cards can only be moved one-at-a-time and only if they are uncovered (i.e. the topmost card in a cascade). Additionally there are four *free cells* that can each store a single card. If a cascade pile is emptied, a card may be moved into the empty space to begin a new cascade. The agent wins if it moves all cards into the foundation piles.
   c) An agent that attempts to acquire an item via on an online auction site such as eBay. On eBay, buyers submit public bids on items. Each item has a timer, and when the timer expires the highest bidder pays their bid amount and receives the item.
2. *[15 points]* For each of the agents described above, categorize it with respect to the six dimensions of task environments as described in Section 2.3.2. Be sure that your choices accurately reflect the way you have specified your environment, especially the sensors and actuators. Give a short justification for each property.
3. *[15 points]* Consider the vacuum-cleaner world depicted in Figure 2.2, where the agent perceives which square it is in and whether there is dirt in the square. Consider the following three questions assuming that the agent's performance measure is reduced each time it moves left or right.
   a) Can a simple reflex agent be perfectly rational for this environment? Explain.
   b) What about a model-based reflex agent? Explain.
   c) How do your answers to **a** and **b** change if the agent's percepts give it the clean/dirty status of *every* square in the environment at each time step?
4. *[25 points]* Using the road map of Romania from the book (Figure 3.2), find a path from Pitesti to Sibiu using breadth-first graph-search. Assume that when all else is equal, cities are chosen in alphabetical order. Show your search tree and indicate the order in which nodes were expanded. Is the resulting path optimal?
5. *[10 points]* Does a finite state space always have a finite search tree? Explain your answer.

6. *[20 points]* You are standing on a river bank with a wolf, a goat, some cabbage, and a boat. You want to get the wolf, goat, and cabbage across the river however the boat can only hold yourself and one other item. Additionally, the wolf will eat the goat if left alone and the goat will eat the cabbage if left alone. Choose a suitable state representation for this problem and then give the initial state, goal test, actions, transition model, and path cost function making only those distinctions necessary to ensure a valid solution. You do not need to find such a solution.