CIS 716.5 Spring 2010 Homework 1

1. The term "agent" in the context of computer systems has no universally accepted meaning — the term is usually defined by listing attributes that it claimed an agent has. Briefly describe what you understand to be the most commonly mentioned attributes of agency.

   (10 points)

2. Identify an intelligent agent from a film, book, or TV show. With reference to your answer to Qestion 1, describe why you think that the agent is intelligent.

   (10 points)

3. For the agent in your answer to Question 2 identify its environment, and classify its environment as:

   - Accessible *or* inaccessible
   - Deterministic *or* non-deterministic
   - Episodic *or* non-episodic
   - Static *or* dynamic
   - Discrete *or* continuous

   Explain why you classify the environment this way.

   (10 points)

4. Recall the coffee-collecting robot that I talk about in class. Classify its environment as:

   - Accessible *or* inaccessible
   - Deterministic *or* non-deterministic
   - Episodic *or* non-episodic
   - Static *or* dynamic
   - Discrete *or* continuous

   Explain why you classify the environment this way.

   (10 points)

5. Figure 1 (next page) shows a an agent situated in an enviroment. Using the notation from Lecture #2, we can think of the environment E being made up of 36 states:

   $$e_{0,0}, e_{0,1}, \ldots$$

   where the subscript of each e indicates a square in the grid.

   Thus the agent, sitting in the bottom lefthand corner, is in state $e_{0,0}$, while if the agent were at the goal, in the top righthand corner, it would be in state $e_{5,5}$. $e_{0,0}$ is the initial state of the environment.

   The filled squares indicate obstacles — the agent cannot be in these states.

   The agent can move north, south, east or west, which we write as:

   $$\alpha_n, \alpha_s, \alpha_e, \alpha_w$$

   and these have the effects you would expect. If the agent is in state $e_{0,0}$ and takes action α_n, it will end up in state $e_{0,1}$, while if the agent is in state $e_{3,2}$ and takes action α_e it will end up in $e_{4,2}$. If the agent tries to move outside the grid then it does not move (for example if the agent is in $e_{0,5}$ and tries to do α_n then it stays in $e_{0,5}$).

   If the agent enters the state $e_{5,5}$, marked with the word goal, then it gets a reward of 10. If it enters state $e_{1,4}$, marked with a dark circle, it gets a reward of -10 (ie it takes a loss).

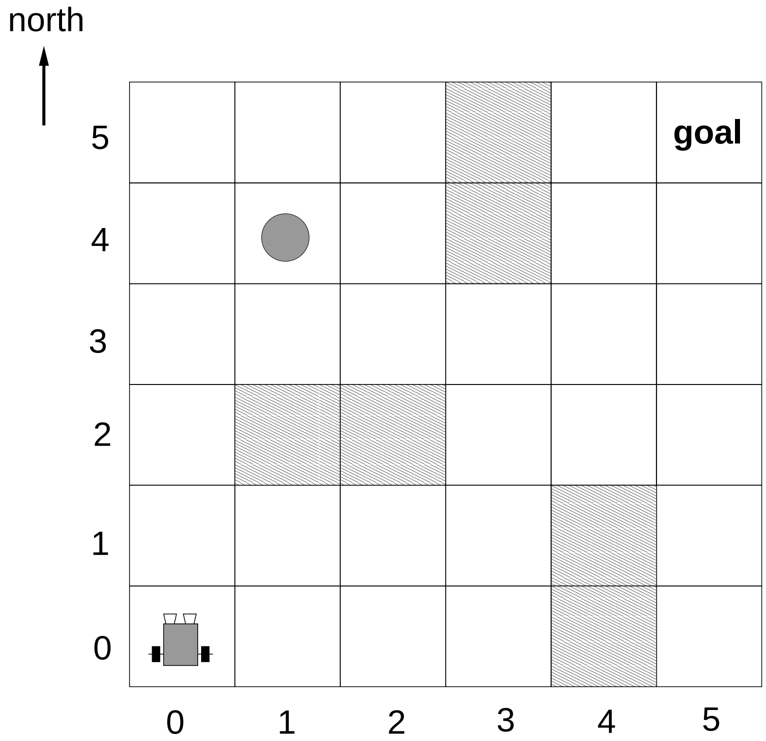

Figure 1:

(a) Give one *maintenance* task for the agent in this environment and one *achievement* task.
(10 points)

(b) Write down a run of the agent that takes it from $e_{0,0}$ to $e_{1,4}$.
(10 points)

(c) Consider the following control program:

while(not in state $e_{5,5}$){
  randomly pick either α_n or α_e (each with probability 0.5)
  execute the action that was selected
}

Write down two runs that the agent might carry out when executing this program.
(10 points)

(d) If the agent executes the program, can it ever reach $e_{5,0}$? Why?
(10 points)

(e) What is the maximum and minimum reward that the agent can get running this program? Why?
(10 points)

(f) How likely is it that when the agent runs the above program it will get the reward of -10?
(Hint: There is a precise probability that the agent will get to the relevant state. That is what I want you to calculate).
(10 points)