Final Exam

CS4630

Out: April 26, 2002
Due: Midnight May 3, 2002

April 26, 2002

1. (10 points) Name the three primitives for expressing the components of a robotics paradigm. Your answer should be three words.

2. (10 points) You have purchased a new robot and you conduct some experiments to evaluate its vision sensor. You discover that the robot always mis-estimates the range to detected objects, but with a seemingly random error of about 10%. You want to use this sensor to help the robot localize. Name and describe two ways that you could represent the sensor error and explain how you could use this representation to help the robot localize. What are the strengths and weaknessesof each?

3. (20 points) Design a reinforcement learning system to control a two-link robot arm (see figure).

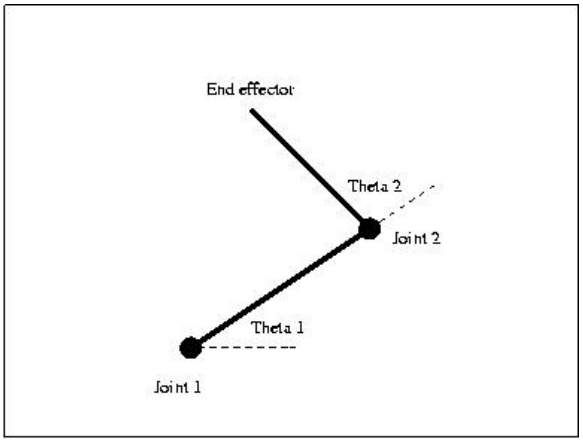

The arm moves only in two dimensions. It is anchored to a table at joint 1; the first link extends to joint 2, which connects the two links together. Motors control both joints, which can be set to any angle. There are four

commands or "actions" that can be sent to the arm: rotate joint 1 CW one step; rotate joint 1 CCW one step; rotate joint 2 CW one step; rotate joint 2 CCW one step. The control system is able to read the angles at each joint and also to detect the X,Y location of the tip of the arm (via an overhead camera).

The task is to move the tip of the arm to a given location (GoalX,GoalY) from a randomly initialized starting location.

(a) Design a reinforcement learning system that will learn to reach the goal location. Describe the algorithm and the training process. Describe the data structures you would use. What reward function would you use? Why?

(b) After learning successfully to move to the goal location, suppose the system is given the task of moving to a different goal location. If you knew that this might occur (that the robot might have to learn different goals) would you prefer to program the system with a model-based or a model-free learning system? Why?

4. (10 points) Describe the frame problem.

5. (10 points) How does the reactive paradigm handle the frame problem?

Consider the following diagram in the following questions.

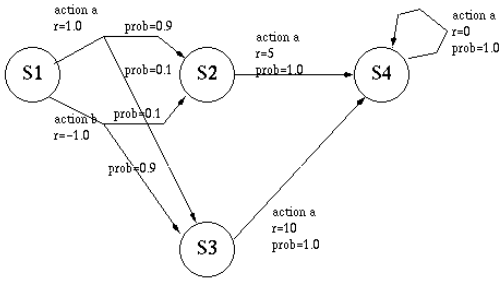

6. (10 points) What is the value of $T(S_2, a, S_4)$?

7. (10 points) What is the value of $R(S_3, a)$?

8. (20 points) We are using Q-learning to learn a policy. There are two possible actions (a and b). We have set $\gamma = 0.8$ and $\alpha = 0.2$. At present portions of our Q table have the following values:

$Q[S_1, a]$=1.0
$Q[S_1, b]$=1.0
$Q[S_2, a]$=2.0
$Q[S_2, b]$=1.0

Now, suppose we were in state S_1, we took action b, we got reward 1.0 and ended up in state S_2. Which item in the Q table will change and what is its new value?