VI51: Virtual Life Simulation - Final Exam P2011

Time: 2h. No documents allowed.
Each part should be written on separate cheats.
English recommended, French accepted.

Part 1: Markov models (4 points)

Let consider two biased coins, which are flipping. Let also consider an observer that can perceive only the result of the coin flip but not which coin is flipped. Coin 1 has got a 2/3 probability for head, and Coin 2 has got a 5/6 probability for tails. The probability of changing coin after a coin 1 flip is 20% and 40% after a coin 2 flip. The coin flip starts always with coin 1.

Question 1.1:

Draw the coin flip Markov graph. Which Markov model suits this example?

Question 1.2:

Let consider the following observation: HHTTTH. Which algorithm should be used to compute the probability of this observation? Compute it.

Question 1.3:

How can we compute the most probable sequence that leads to this observation?

Part 2: Learning Algorithms (6 points)

Question 2.1:

1. Under what conditions should we focus on the algorithm Q-learning algorithm rather than the Q-iteration?
2. Under what conditions should the Q-learning algorithm can be used with a Q-value for each pair (state, action)?
3. Under what conditions the algorithm Q-learning is it supposed to converge?

Exercise 2.2:

Let a robot perform the task described by the MDP $M = (S, A, T, R)$ as follows:

- $S = 0, 1, 2, 3, 4, 5$
- $A = -1, 1$
- $T : S \times A \times S \rightarrow [0, 1]$ as follow (an empty cell corresponds to the probability 0) :

(s,a)	T(s,a,0)	T(s,a,1)	T(s,a,2)	T(s,a,3)	T(s,a,4)	T(s,a,5)
(0,-1)	1					
(1,-1)	0.8	0.15	0.05			
(2,-1)		0.8	0.15	0.05		
(3,-1)			0.8	0.15	0.05	
(4,-1)				0.8	0.15	0.05
(5,-1)						1
(0,1)	1					
(1,1)	.0.5	0.15	0.8			
(2,1)		.0.5	0.15	0.8		
(3,1)			.0.5	0.15	0.8	
(4,1)				.0.5	0.15	0.8
(5,1)						1

- Let $R : S \times A \times S \rightarrow R$ be defined as follow:

$$\begin{cases} R(s,a,s') = 5, & \text{if } s' = 5, s = 4 \text{ and } a = 1 \\ R(s,a,s') = 1, & \text{if } s' = 0, s = 1 \text{ and } a = -1 \\ R(s,a,s') = 0, & \text{otherwise} \end{cases}$$

Question 2.2.1:

Give all the Q-values âĂŇâĂŇafter the application of two iterations of the algorithm Q-iteration (algorithm of type Iteration-Value using Q-values).

Question 2.2.2:

Provide policy π followed by the agent using the Q-values âĂŇâĂŇdetermined above. It is assumed that the agent chooses for each state s the action that maximizes the expected rewards.

Remainder:

The Bellman optimality equation for a stochastic MDP is:
$Q\star(s,a) = \sum_{s'} T(s,a,s')\,[R(s,a,s') + \lambda max_{a'} Q\star(s',a')]$

Part 3: Environment Model and Movement Behaviors (10 points)

Question 3.1:

Explain the principle and the main advantage of the influence/reaction approach.

Question 3.2:

Define the kinematic and steering behaviors and their differences.

Question 3.3:

Provide the algorithm of the "Flee" steering behaviour.

Exercise 3.4:

You want to write the behaviours of several mobile entities that are moving inside a room because they are trying to reach a point. You should avoid collision among them.

Question 3.4.1:

What are the different algorithms required to create this global behaviour? List them and provides the algorithms.

Question 3.4.2:

How can you merge these previously selected algorithm to create a single global behaviour? Provide the algorithm for this global behavior.

Exercise 3.5:

You want to write a perception algorithm for agents with is supporting occlusion culling (occluded objects must not be seen).
Assuming that frustum culling is done by the function: `List<Object> frustumCulling()`, propose an occlusion culling algorithm.
If you have made several design hypothesis, you must quickly explain them.