IA51: AI for Games and Vir. Env. - Final Exam P2017

Duration: 2h.
No document nor calculator nor smart phone nor touchpad allowed.
English and French are accepted.
Malus of 1 point for dirty sheets.
Each part must be written on separated sheets.

Part 1: Markov Models (10 points)

Question 1.1: MCHMM (3 points)

What is a MCHMM? What is its main interest as compared to classical HMM? Compare the MCHMM to another approach with similar purpose? (3 points)

Exercise 1.2: Flipping coins (7 points)

Let consider two biased coins, which are flipping. Let also consider an observer that is able to perceive only the result of the coin flip but not which coin is flipped. Coin 1 has got a 2/3 probability for head (H) and Coin 2 has got a 5/6 probability for tails (T). 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 1.2.1: That's much better with a graph

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

Question 1.2.2: The probability of an observation

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

Question 1.2.3: What is most probable?

How can we compute the most probable sequence that leads to this observation? Make the necessary modifications to the computations of question 2? What is the most probable sequence

Question 1.3: Bonus

According to Markov, models that aim at making decisions are always related to a reward. Can this assertion be extended to Humans behaviors? In our current societies, can decisions be made without reward?

Part 2: Machine Learning (4 points)

Consider the following 3x3 gridworld:

	!	
*	T	*
		*

The agent's goal is to collect as many food (represented by a star in the figure above) as possible, before going to the cell T, which ends the task. The agent can move up, down, left and right. Carrying out an action that leads the agent out of the gridworld results in staying in the cell. When an agent reaches a cell containing a food, it takes it from the cell and receives a reward of +1. Taking it from the cell means that the agent can receive a reward of +1 only once in the cell containing a food. It reveives a reward of -10 if it reaches the cell $(2, 1)$ with the !. Otherwise it receives a reward of -1. To simplify the computations, we assume $\gamma = 0.5$.

Question 2.1:

Model this task by an oriented graph, where each node represents a state and each edge a transition between states. On each edge is indicated the action to take to go from one state to another one and the reward received from going from one state to the next one.

Question 2.2:

Recall the algorithm Q-Iteration.

Question 2.3:

Apply 3 iterations of the Q-iteration method and show the final matrix of the Q-values.

Question 2.4:

What is the greedy policy assuming the Q-values obtained after 3 iterations?

Part 3: Agent Simulator (6 points)

You want to write a perception algorithm for agents that is supporting occlusion culling (occluded objects must not be seen by the agents, even if they are in the perception field-of-view).
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.