

IA51: AI for Serious Games - Final Exam P2018

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 different sheets.

Part 1: Agent Simulator (S. Galland, 16 points)

Exercise 1.1: General Questions (5 points)

Question 1.1.1: (1 pts)

Agents may interact together by using the stigmergy's principles. What is stigmergy?

Question 1.1.2: (1 pts)

Considering the part of the simulation process during which an agent is acting within the environment, what is the difference between an influence and an action?

Question 1.1.3: (3 pts)

Excluding the agent's body, what are the three typical layers of an agent's architecture? Explain briefly each of them.

Exercise 1.2: Environment Model (5 points)

Consider a continuous environment for modeling the indoor of a building, whatever the usage of this building.

Question 1.2.1: Topological Data Structure (2 pts)

The answer to this question could be merge with the answer to the next question. Write a UML class diagram that defines the data structure for representing the topology of the space within the environment. An indoor space is composed of rooms and corridors, that are connected by doors.

Question 1.2.2: Object Container Data Structure (3 pts)

The answer to this question could be merge with the answer to the previous question. A building must contain objects (chairs, tables, etc.). These objects are located in the rooms or the corridors. Each object is located inside a single space (room or corridor). Upgrade the UML class diagram that you have provided for the previous question in order to define the data structure (list, set, tree, graph) that permits to store the objects.

Exercise 1.3: Warrior Behavior (5 points)

You have to write the behavior of a warrior within a First-Person Shooter Game. The warrior has the goal to kill any other player on the map. He has a gun that could be recharged with bullets picked on the ground. He has a life level that could be increased by picking up a medkit on the ground.

Question 1.3.1: Strategic Actions (1 pts)

What are the strategic actions that could be selected in this behavior? You must define at least 4 different actions.

What are the tactical or operational behaviors that are corresponding to the defined strategic actions?

Question 1.3.2: State Machine (4 pts)


Write the state machine that corresponds to the expected behavior of the warrior. You must indicate if the warrior gives priority to its life or to kill the enemy, and write the state machine accordingly.

Exercise 1.4: Agent Oriented Programming (1 point)

What are the differences between an agent in agent-oriented programming and an object in object-oriented programming?

Part 2: Machine Learning (F. Lauri, 4 points)

Exercise 2.1: Turn Around



Consider the circle of states above, which depicts the 8 states of an MDP. The agent can choose between two actions: going one step clock-wise or one step counter-clock-wise. When $s = 1$ (the green state) the next state will be $s' = 4$, independent of action a .

Description of the reward $R(s, a)$:

- The agent receives a reward of $r = 8$ when $s = 1$ (the green state).
- The agent receives a reward of $r = 1$ when $s = 2$ (the red state).

To simplify the computations, we assume $\gamma = 0.5$.

Question 2.1.1:

Recall the algorithm Q-Iteration.

Question 2.1.2:

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

Question 2.1.3:

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