EXAM TECHNIQUES OF ARTIFICIAL INTELLIGENCE

Monday January 17, 9.00-13.00

Question 1. Maze search

Suppose you are given the following maze and are asked to find the path from state 1 (entry of the maze) to state 30 (exit of the maze). The state numbers are shown in the cells of the maze. Black cells without state numbers are not accessible. In the maze, a subject can perform the following 4 actions: move one cell up, one cell down, one cell to the left, one cell to the right. Any of these actions can only be performed on condition that the resulting state remains within the maze and that the resulting state is not a black cell. Also, actions that bring you back to a previous state are not allowed.

1	2	3	4	5
6	7	8	9	10
11	12	13	14	15
16	17	18	19	20
21	22	23	24	25
26	27	28	29	30

1. Give the sequence of states visited by depth first search, starting from state 1 until state 30 is reached. How many states are visited by the algorithm before the exit of the maze is found?
2. Give the sequence of states visited by breadth first search, starting from state 1 until state 30 is reached. How many states are visited by the algorithm before the exit of the maze is found?
3. Observing the results from question 1 and 2, John concludes that in maze search, one should always use depth first search, because it finds the exit quicker than breadth first search. Is John right or wrong? Please motivate.
4. Provide the sequence of states visited by A*, starting from state 1 until state 30 is reached. Assume that each action taken increases the g cost by 1 and use the Manhattan distance to the exit as a heuristic function. When several states in the priority queue have an equal lowest f score, the state with the smallest state number is preferred.

Question 2. Concept learning

1. How and why is an attribute selected to be the next one to install in the top-down recursive development of a decision tree?
2. Illustrate your answer on a very small toy problem (consisting of 2 or 3 instances, and a few attributes).
3. Is the resulting hypothesis of a decision tree always an element of the version space which results from applying the candidate elimination algorithm? Motivate.
4. Explain in your own words why learning from examples is futile without some form of inductive bias.

Question 3. Neural Networks

1. What is the mathematical principle underlying learning in multilayer neural nets and giving rise to the backpropagation algorithm?

2. Consider a single sigmoid threshold unit with three inputs, x1, x2, and x3. The output function is

$$y = g(w_0 + w_1x_1 + w_2x_2 + w_3x_3)$$

And the activation function g is given by

$$g(z) = \frac{1}{1 + e^{-z}}$$

We input values of either 0 or 1 for each of these inputs. Assign values to weights w_0, w_1, w_2 and w_3 so that the output of the sigmoid unit is greater than 0.5 if and only if (x_1 AND x_2) OR x_3.

3. True or false? The gradient descent weight update rule for a unit whose output is

$$w_0 + w_1(x_1 + 1) + w_2x_2^2$$

is given by:

$$\Delta w_0 = \eta \sum_d (t_d - o_d)$$

$$\Delta w_1 = \eta \sum_d [(t_d - o_d)x_{d1} + (t_d - o_d)]$$

$$\Delta w_2 = \eta \sum_d (t_d - o_d)\, 2x_{d2}$$

Where t_d is the target output for the dth training example,
o_d is the unit output for the dth example,
x_{d1} is the value of x_1 for the dth training example,
x_{d2} is the value of x_2 for the dth training example.

4. True are false? Regardless of the size of the neural network, the backpropagation algorithm can always find the globally optimal weights for the neural network. Justify your answer.

Question 4. Reinforcement Learning

1. When is it appropriate to use RL? Describe a problem where you would suggest using RL.
2. Consider the reinforcement learning problem given below, where the learning rate $\eta = 0.9$:

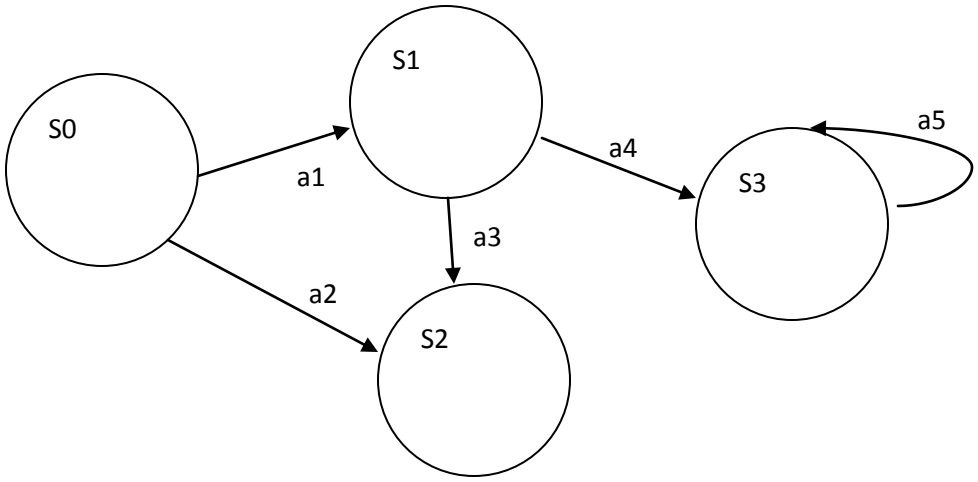

The rewards are defined as follows:

r(s0,a1) = 2
r(s0,a2) = 3
r(s1,a3) = 3
r(s1,a4) = 10
r(s3,a5) = 0

Give all the give all the Q* values (= the Q values for the optimal policy)