Exam Intelligent Systems, April 15th 2013, 16.00 – 19.00

This exam has 9 questions. Your answers should be given in Dutch or English. With the answers you can earn 90 points. You get 10 points for free. 100 points yields a 10.

Exercise	1	2	3	4	5	6	7	8	9
points	10	10	10	10	10	10	10	10	10

Always explain your answers. Good luck!

1. Translate the following assertions into First-Order Logic. Use a membership predicate to denote that points belong to lines or circles. Use the same predicate to denote that lines or circles belong to certain classes of lines or circles. So points, lines, circles, sets of points, sets of lines and sets of circles are all seen as objects and membership relations are denoted using a membership predicate relating these objects.
   (a) Two different circles have either 1, 2 or 0 points in common
   (b) Two circles that have exactly 1 point in common do not have the same center

2. Write a Prolog predicate `swapfl(list1,list2)`, that checks whether the final element of `list1` is equal to the first element of `list2`, whether the first element of `list1` equals the final element of `list2`, and whether the two middle parts of the lists are identical. You can use the standard predicate `append(list1,list2,list3)` that ensures that `list3` is the concatenation of lists `list1` and `list2`.

3. (a) Negation as failure in Prolog is different from classical negation in logic. What are the differences?
   (b) Each Prolog rule can be read as a first-order logic sentence. Yet, Prolog programs are different from first order logic theories. What are the differences in interpretation?

4. (a) Explain, using an example, what a ‘unique names axiom’ is.
   (b) Explain the difference between a Herbrand-model and a standard first-order model.

5. (a) Describe two important choice points used for backtracking in the POP algorithm.
   (b) Explain the difference between an order constraint and a causal link, and explain why we need both in the POP algorithm.

6. We consider a blocks world containing 4 blocks and one robot. The robot knows two types of actions: he has four actions (one for each block) for putting a block on the ground if it is on top of some other block, and twelve actions to pick up a block from the ground and put it on some other block (each of the four blocks can be put on three other blocks). If the robot puts block 'x' on the ground, or picks it up from the ground to put it on some other block, then all the blocks on top of 'x' simply stay on top and move accordingly. In the initial situation, block 1 and 2 are on the ground, block 3 is on block 1, and block 4 is on block 3. The robot's goal is to have block 3 on top of block 2, and nothing on top of block 3.
   (a) Try to give a STRIPS description of this problem. If you do not succeed, give a description is some other planning language (e.g. ADL) and explain why you did not succeed in STRIPS.
   (b) Give a partial order plan as a solution to this problem.

7. If the POP-algorithm is extended with Hierarchical Task Network (HTN) planning capabilities, (as in Russell & Norvig), a new important choice point is added to the algorithm.
   (a) Describe this extra choice point used for backtracking by HTN-planning.
   (b) A plan library in HTN-planning may contain several different plans for a single abstract action. What happens (in the HTN-planning algorithm) if all of these plans fail?

see the other side

8. In inductive logic programming aims we often look for definitions of the form “P(x) ↔ def(x)”, where def(x) is some formula (usually in disjunctive normal form) over the attributes in terms of which the definition has to be learnt. The Least commitment search algorithm for coming to a correct formula def(x) keeps an S-set with the current most specific information and a G-set with the current most general information induced by the examples considered so far.
(a) What are the contents of the S-set and what are the contents of the G-set at the start of the search?
(b) A new example considered by the algorithm says that for certain values of the attributes, P(x) is false. It turns out that this conflicts with information in the current S-set. What kind of example is this (false positive or false negative) and how is the S-set updated (is het specialized or generalized or discarded)?

9. Below three pictures from the book concerning an optimal policy for a Markov decision process discussed during the lecture, where the (dis)reward in each non-terminal state is -0.04. The first picture shows the optimal policy when the chances that as a result of an action one ends up in the room one is heading for are 0.8, and the chances that one ends up in a room to the left or to the right are 0.1 for both possibilities. Bumping into a wall means that one stays in the same room. The rooms with rewards +1 and -1 are terminal states. The second picture shows the utilities associated with the states, for the optimal policy of the first picture.
(a) Calculate the discount factor γ resulting in the optimal policy pictured below, by applying a Belmann equation to room (3,3). Also give your calculation.
(b) Assume the utilities in the second picture are calculated using value iteration where initially all utilities are set to zero. The third picture gives for 5 positions in the problem world a graph of the evolution of the utilities over time, as the value iteration algorithm proceeds. Let us number the graphs 1 to 5, starting with 1 for the left most graph originating at (0,0) in the picture. Say for each of the graphs 1 to 5 to which coordinate in the problem world of the first two figures it may correspond.

3	→	→	→	+1
2	↑		↑	−1
1	↑	←	←	←
	1	2	3	4

3	0.812	0.868	0.912	+1
2	0.762		0.660	−1
1	0.705	0.655	0.611	0.388
	1	2	3	4