Exam Introduction Intelligent Systems, November 12th 2010, 13.30 – 16.30

This exam has 9 questions. Your answers should be given in Dutch or English. With the answers to the questions 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:
   (a) For any two men one can find a woman that likes one of them and dislikes the other.
   (b) You can fool some of the people all of the time, and all of the people some of the time, but you cannot fool all of the people all of the time (Abraham Lincoln) (You can identify 'for some' with 'there is a').

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) What is the definition of 'negation as failure' in terms of the 'cut'?
   (b) What are the differences between negation as failure in Prolog and classical negation in logic?
   (c) What are the other differences between the interpretations of Prolog sentences and first-order logic sentences?

4. (a) Explain, using an example, what a 'unique names axiom' is.
   (b) Explain what a Herbrand Model is.

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.
   (c) Is your answer to (b) the only partial order plan possible as an answer to this problem?

see the other side

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) HTN-planning can be seen as a heuristic for standard partial order planning. Explain this.
(c) Is HTN planning an admissible heuristic? Explain.

8. Your doctor has bad news. You tested positive for a serious disease and the test is 98% accurate. The doctor says that the 98% accuracy means that the probability of testing positive if you have the disease is 0.98, and that the probability of testing negative if you do not have the disease is also 0.98. The doctor says there might also be good news: the disease is rare, striking only 1 out of 100000 people of your age.
(a) You claim that the good news is no good news since the test has a known accuracy. Who is right, and why?
(b) Is the information "the probability of testing negative if you do not have the disease is also 0.99" relevant for determining your chance of having the disease? Explain your answer.
(c) Calculate the chance that you actually have the disease.

9. Below two 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 left 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 right picture shows the utilities associated with the states, for the optimal policy of the left picture.
(a) Calculate the discount factor γ resulting in the optimal policy pictured below, by applying a Belmann equation to room (3,1).
(b) Assume the utilities for the problem of the pictures below are calculated using value iteration where initially all utilities were set to zero. Which are the rooms for which the initial step of the iteration yields a -0.04?