T-106.5900 Software Technology Course with a Varying Content (Introduction to Artificial Intelligence)

Examination April 7, 2015

Examiner: Kary Främling

No written material is allowed in this exam. Calculators are not allowed. Submit at least one answer sheet, even if an empty one! Write on each answer sheet you submit the code of the course, the date, your name, and your student ID number.

1 (20 points)

Provide a short explanation of the difference between the following concepts:

(a) Linear versus *Non-linear* Regression?

(b) Discrete versus *Continuous* environment?

(c) Exploitation versus *Exploration* in Reinforcement learning?

(d) Benign versus *Adversarial* environment?

2 (20 points)

(a) Describe the principles of A* search

(b) Apply A* search to a problem, where the following state transitions are possible:

Transition	Its cost
S → A	2
S → B	1
S → C	2
A → D	2
B → E	1
C → F	1
D → H	2
E → H	8
E → I	7
F → I	2
H → G	1
I → G	2

The problem is to find a path from the state S to the state G. The values of the function h in the different states are the following $h(A) = 3$, $h(B) = 3$, $h(C) = 3$, $h(D) = 2$, $h(E) = 2$, $h(F) = 2$, $h(H) = 1$, and $h(I) = 1$. Which one of the two paths giving the minimum cost will your algorithm find? Why?

3 (20 points)

What does the Occam's razor principle say? What implications does it have in the case of *supervised learning*? How can you minimize the *generalization error* using the Occam's razor principle in supervised learning?

4 (20 points)

Dynamic programming: What is the value of the start state in the world shown in the figure below? The start state is in the lower left corner with the robot facing north. The goal state is in the indicated position, with the robot heading north. Reaching the goal state gives a +100 reward. Going straight costs -1, turning right costs -10 and turning left costs -2. Please justify your answer rather than just providing a number!

PS: Actions are deterministic (state transition function is deterministic), no discounting (γ=1).

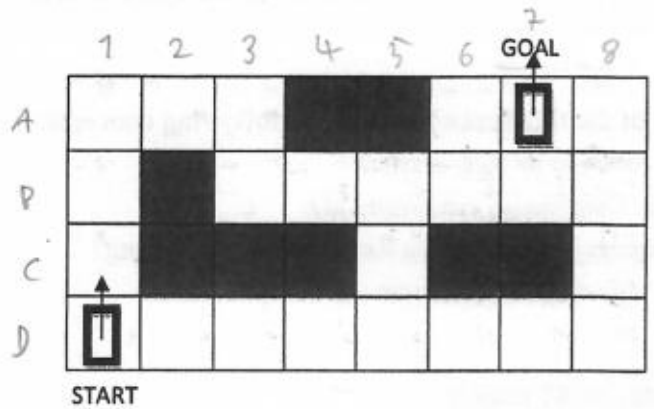

5 (20 points)

Hidden Markov Models (HMM): Describe what they are and what are they used for? What is meant by "hidden"? Provide an example of at least one application that they are used for and describe why and how they are used in that application. Are HMMs applicable to "continuous state" / "continuous space" applications? How and why (or why not)?

Good Luck!