LE/EECS 4401 3.0 Artificial Intelligence
GS/EECS 5326 3.0 Topics in Artificial Intelligence

Dept. of Electrical Eng. & Computer Sci.
York University
Winter 2022

Final Examination — April 18, 2022

Duration: 120 minutes
Total marks: 75

Name:

Student Number:

1)	/5
2)	/5
3)	/5
4)	/5
5)	/5
6)	/15
7)	/10
8)	/10
9)	/15
Total	/75

1. [5 points] Reasoning in classical logic (e.g., first-order logic) is *monotonic*. Explain what this mean. Also explain why default reasoning is *nonmonotonic*.

2. [5 points] Explain how STRIPS deals with the frame problem.

3. [5 points] Q-learning is called a *model-free* reinforcement learning method. Explain what this mean.

4. [5 points] Consider an undiscounted MDP with states s_0, s_1, s_2, s_3, and s_4, where s_0 is the starting state, and s_3 and s_4 are terminal states. The reward in s_4 is $+1$, in s_3 it is -1, and in all other states it is 0. Action a is only possible in s_0 and leads to s_1 with probability 0.5 and to s_2 with probability 0.5. Action b is possible in s_1 and leads to s_4 with probability 0.9 and to s_3 wtth probability 0.1; it is also possible in s_2 and leads to s_3 with probability 0.9 and to s_4 wtth probability 0.1. Action c is possible in s_2 and leads to s_4 with probability 0.8 and to s_3 wtth probability 0.2; it is also possible in s_1 and leads to s_3 with probability 0.9 and to s_4 wtth probability 0.1. (No actions are possible in the terminal states.)

   What is the expected value of the policy π where $\pi(s_0) = a$, $\pi(s_1) = b$, and $\pi(s_2) = c$? Show how you compute this value.

5. [5 points] Show using logical interpretations that $\forall x \exists y Loves(x, y)$ does *not* logically entail $\exists y \forall x Loves(x, y)$.

6. [15 points] Assume the following facts:

   (1) Steve likes all easy courses.

   (2) Science courses are hard.

   (3) All the courses in the basketweaving department are easy.

   (4) BK301 is a basketweaving course.

   (a) For each fact, give its translation in first-order logic.

   (b) Convert all of the above facts into clausal form.

(c) Use resolution to answer the question “What course would Steve like?”. Indicate clearly which clauses are resolved and what substitutions are used.

7. [10 points] Suppose that we have the following belief network (or Bayes net):

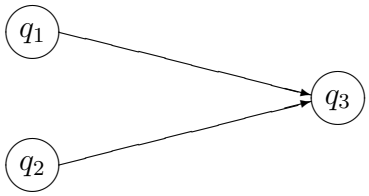

(a) Give an example of an independence assumption that is implicit in this network (write it formally).

(b) What are the 6 conditional probabilities that need to be specified to fully determine the joint probability distribution?

(c) Express $Pr(q_3|q_2)$ in terms of the 6 conditional probabilities given in your answer to the previous question.

8. [10 points] Consider the following STRIPS actions:

Action Name	Preconditions	Add Effects	Delete Effects
A	p	q	r
B	p	r	-
C	-	p	-
D	o	-	o

Suppose that the start state is $S_0 = \{o\}$ and the goal is $G = \{p, q, r, o\}$.

(a) Which actions are applicable in the start state S_0? What is the new state for each of the applicable actions?

(b) Suppose we want to do backward/regression planning. Which actions are consistent with the goal G? What is the regressed goal for all the consistent actions?

(c) Give a sequence of actions that achieves the goal G when executed in the starting state S_0.

9. [15 points] Consider the following dynamic domain, to be specified in the situation calculus. Suppose that we have a fluent $LightOn(x, s)$ that is true if and only if light x is on in situation s and a fluent $PowerOn(s)$ that is true if and only if lthe power is on in situation s. Suppose also that we have an action $flipSwitch(x)$ that flips the switch of light x; the only effects of $flipSwitch(x)$ are (1) that it will turn light x on if x is currently off (i.e. not on) and the power is on, as well as (2) that it will turn light x off if x is currently on. Finally, assume that $flipSwitch$ is the only action that affects the fluent $LightOn$.

   a) Write effect axioms for the action $flipSwitch$; your axioms should capture all the effects of the action.

   b) Write frame axiom(s) for the action $flipSwitch$ and the fluent $LightOn$; your axioms should handle all cases where the fluent does not change when $flipSwitch$ is performed.

   c) Write a successor state axiom for the fluent $LightOn$.

d) Suppose that the power is on and that lights $L1$ and $L2$ are both off initially. Write a sentence in the situation calculus that represents the *planning task* for the goal of having both $L1$ and $L2$ on.

e) Write a ground situation term that represents a plan that is achieves the goal of having both $L1$ and $L2$ on in an initial situation as described above.

f) This domain cannot be represented directly in STRIPS. Briefly explain why. Also describe how you could change the language to allow the domain to be represented in STRIPS.