

# Qualifying Examination 2019 Artificial Intelligence

There are six questions on this exam. You must answer them all. All answers must be in the test booklets. All questions have the same value for grading.

Your grade depends on the correctness, completeness and style of your answers.

1. Consider the problem of using a truck to move boxes among warehouses in Manhattan. Boxes can be loaded on the truck or unloaded from the truck. The truck can drive from one warehouse to another, and can carry at most 10 boxes at a time.

There is one truck and 100 boxes. There are five warehouses: A, B, C, D and E. The truck starts at warehouse A. Initially there are 25 boxes in each of the warehouses B, C, D, and E (the warehouses can each hold an unlimited number of boxes). The goal is to move all the boxes to warehouse A. The distances between the warehouses are (in kilometers):

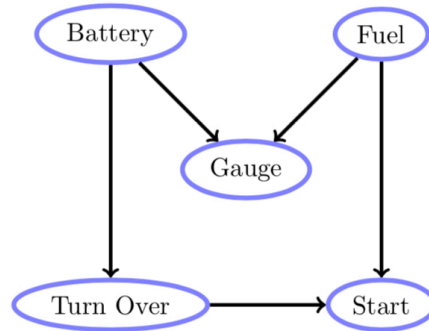
	A	B	C	D	E
A	0	1	2	4	9
B	1	0	3	6	8
C	2	3	0	4	4
D	4	6	4	0	5
E	9	8	4	5	0

You will formulate this as a planning problem using STRIPS (PDDL) action schemas.

- (a) Describe the start and goal states for this problem using STRIPS notation. You don't have to write out *every* necessary literal (all 100 boxes, for example) but indicate what literals are needed and give a few examples.
- (b) Write action schemas for loading a box on the truck, for driving the truck from one warehouse to another, and for unloading a box.
- (c) Using your schemas from part (b), show how resolution is used to load one box on the truck at warehouse C, drive it to warehouse E then unload it.
- (d) Give an admissible A\* heuristic to guide the search, so as to minimize the total distance traveled by the truck.

## 2. Bayes Nets

The network below concerns the probability of a car starting.



You know the following facts ( $b$  = battery,  $s$  = start,  $f$  = fuel,  $t$  = turn over,  $g$  = gauge):

- $p(b = \text{bad}) = 0.02$
- $p(g = \text{empty} \mid b = \text{good}, f = \text{not empty}) = 0.04$
- $p(g = \text{empty} \mid b = \text{bad}, f = \text{not empty}) = 0.1$
- $p(t = \text{false} \mid b = \text{good}) = 0.03$
- $p(s = \text{false} \mid t = \text{true}, f = \text{not empty}) = 0.01$
- $p(s = \text{false} \mid t = \text{false}, f = \text{not empty}) = 1.0$
- $p(f = \text{empty}) = 0.05$
- $p(g = \text{empty} \mid b = \text{good}, f = \text{empty}) = 0.97$
- $p(g = \text{empty} \mid b = \text{bad}, f = \text{empty}) = 0.99$
- $p(t = \text{false} \mid b = \text{bad}) = 0.98$
- $p(s = \text{false} \mid t = \text{true}, f = \text{empty}) = 0.92$
- $p(s = \text{false} \mid t = \text{false}, f = \text{empty}) = 0.99$

Calculate  $P(f = \text{empty} \mid s = \text{false})$ , the probability of the fuel tank being empty given that the car does not start.

3. Consider a village where all villagers are either healthy or have a fever and only the village doctor can determine whether each has a fever. The doctor diagnoses fever by asking patients how they feel. The villagers may only answer that they feel normal, dizzy, or cold. There are two states, "Healthy" and "Fever", but the doctor cannot observe them directly. The prior probability of being Healthy is 0.6.

On each day, there is a certain chance that the patient will tell the doctor he/she is "normal", "cold", or "dizzy", depending on their health condition. If the patient is Healthy, the probability for "normal" is 0.5, for "cold" is 0.4, and for "dizzy" is 0.1. If the patient has a Fever, the probability for "normal" is 0.1, for "cold" is 0.3, and for "dizzy" is 0.6.

A Healthy patient remains Healthy on the next day 70% of the time, and a patient with a Fever has a Fever the next day 60% of the time.

- a. Show how this can be formulated as an HMM, including the probability tables.
- b. A patient visits the doctor three days in a row and tells the doctor that that on the first day he feels normal, on the second day he feels cold, on the third day he feels dizzy. Use the filtering algorithm to determine the probability of being in each state for  $t = 1, 2, 3$ .
- c. What is the time complexity of filtering in an HMM?

#### 4. Markov Decision Processes

The Bellman equation:

$$U(s) = R(s) + \gamma \max_{a \in A(s)} \sum_{s'} P(s' | s, a) U(s')$$

- a. An MDP may have cyclic paths in its states. What keeps us from performing infinite calculations when we solve such MDPs?
- b. What is the time complexity of solving the Bellman equation for a system with  $n$  states?
- c. Is there an optimal policy for every MDP?
- d. Why are iterative solution methods necessary for the Bellman equation?
- e. Name three iterative solution methods for the Bellman equation.

5. In the following, we describe whether a person is ill. We use a representation based on conjunctive constraints (three per subject) to describe each individual person. These constraints are “running nose”, “coughing”, and “reddened skin”, each of which can take the value true (‘+’) or false (‘-’). We say that somebody is ill, if he is coughing and has a runny nose — each single symptom individually does not mean that the person is ill.

- Specify the space of hypotheses that is being managed by the version space approach. To do so, arrange all hypotheses in a graph structure using the more-specific-than relation. If hypotheses are vectors of constraints, denoted by  $\langle N, C, R \rangle$  with  $N, C, R = \{-, +, \emptyset, *\}$ .
- Apply the candidate elimination (CE) algorithm to the sequence of training examples specified in the table and name the contents of the sets  $S$  and  $G$  after each step.

Training Example	N (running nose)	C (coughing)	R (reddened skin)	Classification
$d_1$	+	+	+	positive (ill)
$d_2$	+	+	-	positive (ill)
$d_3$	+	-	+	negative (healthy)
$d_4$	-	+	+	negative (healthy)
$d_5$	-	-	+	negative (healthy)
$d_6$	-	-	-	negative (healthy)

- We now extend the number of constraints used for describing training instances by one additional constraint named “fever”. We say that somebody is ill, if he has a running nose and is coughing (as we did before), or if he has fever.  
What steps does the CE algorithm perform now, and what is the result, given the training examples specified below?

Training Example	N (running nose)	C (coughing)	R (reddened skin)	F (fever)	Classification
$d_1$	+	+	+	-	positive (ill)
$d_2$	+	+	-	-	positive (ill)
$d_3$	-	-	+	+	positive (ill)
$d_4$	+	-	-	-	negative (healthy)
$d_5$	-	-	-	-	negative (healthy)
$d_6$	-	+	+	-	negative (healthy)

- What is the reason for this result? What modification to the hypothesis space will fix things?

## 6. Q-learning

- Q-learning is called “model-free”. Why?
- Q-learning is called “off-policy”. Why?
- Suppose we create a very optimistic prior distribution of values over the states. How does this affect the behavior of a Q-learning agent?