


Practice Qualifying Examination 2022 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.

PhD students must answer all questions. MS students do not have to answer questions 3d, 4c, 5b, 5d, 6.

1. Consider the 8-Puzzle. In this puzzle there are eight sliding tiles and one open blank spot. Each move consists of sliding one tile into the blank spot. The initial state and goal state are given below.




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.
- (b) Write an action schema for moving a tile.
- (c) Using your schema from part (b), show how resolution is used to move two tiles, for example how to move the 5 tile to the right and then the 7 tile down.
- (d) Give an admissible A* heuristic to guide the search, so as to minimize the total moves to solve the problem.

2. Bayes Nets

You are given the Bayesian network structure below, consisting of 5 binary random variables A, B, C, D, E. Each variable corresponds to a gene, whose expression can be either “ON” or “OFF”.



Part 1:

We covered the chain rule of probability for Bayes Nets, which allows us to factor the joint probability over all the variables into terms of conditional probabilities. For each of the following cases, factor $P(A,B,C,D,E)$ according to the independencies specified and give the minimum number of parameters required to fully specify the distribution.

- 1: A,B,C,D,E are all mutually independent
- 2: A,B,C,D,E follow the independence assumptions of the above network

Part 2:

You are given the conditional probabilities below for the network.

$$\begin{aligned}
 P(A = ON) &= 0.6 \\
 P(B = ON | A) &= \begin{cases} 0.1, & A = OFF \\ 0.95, & A = ON \end{cases} \\
 P(C = ON | A) &= \begin{cases} 0.8, & A = OFF \\ 0.5, & A = ON \end{cases} \\
 P(D = ON | A, B) &= \begin{cases} 0.1 & A = OFF, B = OFF \\ 0.9 & A = ON, B = OFF \\ 0.3 & A = OFF, B = ON \\ 0.95 & A = ON, B = ON \end{cases} \\
 P(E = ON | D) &= \begin{cases} 0.8, & D = OFF \\ 0.1, & D = ON \end{cases}
 \end{aligned}$$

Calculate the following:

1: $P(A=ON, B=ON, C=ON, D=ON, E=ON)$

2: $P(E=ON \mid A=ON)$

3. Suppose you were locked in a windowless room for several days so that you could not see the weather outside. There are three types of weather: sunny, foggy and rainy. The only clue you have to the weather is whether the caretaker who comes into the room carrying your daily meal has an umbrella.

The probabilities of tomorrow's weather given today's weather are:

		Tomorrow's Weather		
		Sunny	Rainy	Foggy
	Sunny	0.8	0.05	0.15
	Rainy	0.2	0.6	0.2
Today's Weather	Foggy	0.2	0.3	0.5

And the probabilities of seeing an umbrella based on the weather are:

	Probability of Umbrella
Sunny	0.1
Rainy	0.8
Foggy	0.3

The prior probability of the caretaker carrying the umbrella on any day is 0.5.

- Suppose the day you were locked in was sunny. The next day the caretaker carried an umbrella. What is the probability that the second day was rainy?
- Suppose the day you were locked in was sunny. The next day the caretaker carried an umbrella but on the third day he did not. What is the probability that the third day was foggy?
- Is a Markov Chain a polytree?
- Explain the key advantages of the Markov assumption. (In other words, explain how using the Markov assumption helps us analyze a problem.)

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. The table below specifies a Markov Decision Process. There are four states. In each state there are one or two actions possible. The table shows the results of executing each action and the probabilities of each result, as well as the initial utilities of the states. For example, action A1 is possible in state A and has a 50% chance of ending in state B and 50% chance of staying in state A.

The value of gamma (the discount factor) is 0.9

State	Utility	Actions	Result	Probability
A	12	A1	A	.5
			B	.5
		A2	C	.7
			D	.3
B	8	B1	A	.4
			D	.6
C	-3	C1	A	.5
			C	.5
D	4	D1	B	.4
			C	.6
		D2	B	.8
			D	.2

Using the value iteration algorithm, perform two update steps to the utilities. Show all work.

- b. Does value iteration always converge to the optimal policy?
- c. Suppose that the Markov assumption doesn't actually hold for a particular system. Is the Bellman equation still applicable? Describe specifically (in terms of the policy, rewards and utilities) how the situation has changed.

5. Consider the following set of training examples to train a robot janitor to predict when an office contains a recycling bin:

STATUS	FLOOR	DEPT	OFFICE SIZE	RECYCLING BIN
Faculty	4	CS	medium	yes
Faculty	4	EE	medium	yes
Student	4	CS	small	no
Faculty	5	CS	medium	yes

Assume that each of the attributes can have just the values that appear in the table.

- a – How many possible instances are there for this example?
- b – What is the size of the hypothesis space?
- c – Write the sequence of S and G boundary sets computed by the candidate elimination algorithm if it is given the above examples in the order in which they are listed in the table.
- d – Write a query (an example) that is guaranteed to reduce the size of the final version space above (S4 and G4) regardless of whether it is classified as positive or negative.

6. Q-learning

$$Q(s, a) \leftarrow Q(s, a) + \alpha(R(s) + \gamma \max_{a'} Q(s', a') - Q(s, a))$$

- a. α is the learning parameter. Discuss the difference between when α is small (near zero) and when it is large (near one). (Discuss how it affects Q learning.)
- b. γ is the discount factor. Discuss the difference between when γ is small (near zero) and when it is large (near one). (Discuss how it affects Q learning.)
- c. Suppose we create a very optimistic prior distribution of values over the states. How does this affect the behavior of a Q-learning agent?