BITS Pilani, K. K. Birla Goa Campus

Artificial Intelligence (CS F407), Part A

Midsem Exam (03/11/2022)
Total Marks: 32
Time Limit: 1.5 Hours

Part A contains **three** questions carrying **16** marks. All questions are compulsory. Answer each question on a fresh page. Answer all the parts of a question in the same place. Marks will be given **only** if necessary steps and justifications are provided.

Part A

Bandit Algorithm

Question 1 (6 marks)

Consider a k-armed bandit problem with k=4 actions. The actions are denoted by 1, 2, 3, and 4. Suppose the initial action value estimates are $Q(a) = 0$, for all a. The sequence of actions (A_t) and rewards (R_t) till time step $t = 5$ is as follows: $A_1 = 2$, $R_1 = -1$, $A_2 = 4$, $R_2 = 1$,$A_3 = 1$, $R_3 = -2$, $A_4 = 2$, $R_4 = 1$, $A_5 = 3$, $R_5 = -1$. Suppose the Bandit algorithm uses the following update rule for estimating the action values:

$$Q(A) = Q(A) + \frac{1}{N(A)}\left[R - Q(A)\right] \tag{1}$$

It is given that the rewards associated with each action is generated from a stationary probability distribution.

(a) (2 marks) Suppose we use .1-greedy action selection in each time step. What will be the probability of selecting action 3 at time step $t = 6$? Show all the necessary steps. Round your answer to three decimal places.

(b) (2 marks) Suppose we use Optimistic Initial values $Q(a) = 10$, for all a. What will be the action value estimates at time step $t = 6$ if we use the Optimistic Initial values? Let the sequence of actions and rewards be as given above. (Assume that the optimistic initial values is the *only* modification that we make to the Bandit algorithm described above.)

What additional modification (if any) should we make to the Bandit algorithm for the Optimistic Initial values approach to work?

(c) (2 marks) Suppose we use the Upper-Confidence-Bound action selection policy for the 4-armed bandit problem described above. What will be the (asymptotic) probability of selecting the *optimal* action as time step t tends to infinity? Your answer must be rounded to three decimal places. Provide necessary justifications for your answers. (You can assume that the rewards are generated from a stationary distribution and the update rule is as given in Eqn. (1). You can also assume that the Upper-Confidence-Bound action selection policy is as given in the textbook and the exploration parameter c is sufficiently large.)

Propositional Logic

Question 2 (6 marks)

Consider the following statements:
If the unicorn is mythical (T), then it is immortal (I), but if it is not mythical, then it is a mortal mammal (M). If the unicorn is either immortal or a mammal, then it is horned (H). The unicorn is magical (G) if it is horned.

Use the propositional logic symbols shown in the brackets to represent the facts mentioned above. Use the *resolution inference algorithm* to answer the questions given below. Show necessary steps and provide justifications for your answers.

(a) Can we infer that the unicorn is mythical?

(b) Can we infer that the unicorn is magical?

Question 3 (4 marks)

Check whether the following sentences are valid, unsatisfiable or neither. Show the necessary steps and justifications for your answer.

(a) $(Smoke \Rightarrow Fire) \Rightarrow (\neg Smoke \Rightarrow \neg Fire)$

(b) $((Smoke \wedge Heat) \Rightarrow Fire) \Leftrightarrow ((Smoke \Rightarrow Fire) \vee (Heat \Rightarrow Fire))$

BITS Pilani, K. K. Birla Goa Campus
Artificial Intelligence (CS F407), Part B

Midsem Exam (03/11/2022)
Total Marks: 32
Time Limit: 1.5 Hours

Part B contains **three** questions carrying **16** marks. All questions are compulsory. Answer each question on a fresh page. Answer all the parts of a question in the same place. Marks will be given **only** if necessary steps and justifications are provided.

Part B

First-Order Logic

Question 4 (4 marks)

Consider the following knowledge base (KB):

R1. $\forall x \exists y (x = y) \land (y = y)$

R2. $CapitalOf(Delhi, India)$

R3. $\neg(Delhi = Colombo)$

(a) (2 marks) Does $KB \models \neg CapitalOf(Colombo, India)$? If yes, then justify your answer. If no, then describe a model that satisfies the KB, but does not satisfy $\neg CapitalOf(Colombo, India)$.

(b) (2 marks) Suppose we replace rule R1 in KB using the following rule R4:

R4. $\exists y \forall x (x = y) \land (y = y)$

Let KB' be the knowledge base obtained after replacing rule R1 with R4.
Does $KB' \models \neg CapitalOf(Colombo, India)$? Provide necessary justifications for your answer.

Question 5 (7 marks)

Construct a First-order Logic knowledge base containing the logical representations for the following facts:

1. Horses, cats, and dogs are mammals.
2. An offspring of a horse is a horse.
3. Bluebeard is a horse.
4. Bluebeard is Charlie's parent.
5. Offspring and parent are inverse relations.
6. Every mammal has a parent.

(a) (2 marks) Construct a First-order logic knowledge base (KB) that represents the above facts. Your representation must be such that $KB \models Mammal(Charlie)$.

(b) (1 mark) Find the sentence(s) (if any) that *cannot* be represented as a set of definite clauses. Skolemization can be performed if required. (Hint: A definite clause is a clause having exactly one positive literal.)

(c) (2 marks) Represent all the sentences (excluding the ones found in part (b)) using definite clauses. Standardize the variables for each sentence.

(d) (2 marks) Can the Forward Chaining algorithm derive $Mammal(Charlie)$ from the sentences found in part (c)? If yes, draw the proof-tree (And-Or graph) generated by the Forward chaining algorithm for the query $Ask(Mammal(Charlie))$. If no, list all the sentences derived by the Forward chaining algorithm before the fixed point is reached.

Adversarial Search

Question 6 (5 marks)

(a) (1 mark) Is it true that without alpha-beta pruning we will not be able to find the minimax value of the root node. Why do we do alpha-beta pruning?

(b) (2 marks) Consider the partial game tree shown below:

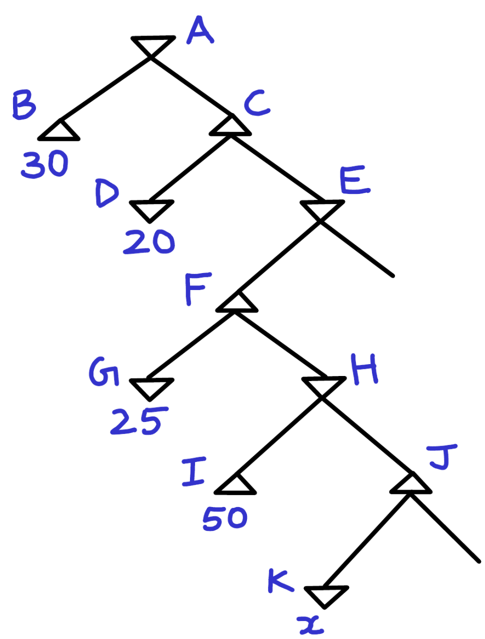

MIN-player is going to make the next move (node A). Nodes B, D, G, I and K have minimax values of 30, 20, 25, 50 and x respectively. If it is given that x is 40, then can we prune the remaining child nodes of MAX-node J? Give an answer based on the information that is given. Provide necessary justifications for your answer. Assume that the actions in each node are considered in a left-to-right manner.

(c) (2 marks) Once again consider the partial game tree shown in part (b). What are the minimum and the maximum values of x that will *ensure* that node K lies on an optimal minimax decision path? Give minimum and maximum values based on the information that is given. Provide necessary justifications for your answer. Assume that the actions are considered in a left-to-right manner.