

You have 2 hours 50 min. The exam is open-book, open-notes.

100 points total

You should be able to finish all questions, but do your best ones first.

Write your answers in blue books. Hand them all in.

Several of the questions on this exam are true/false or multiple choice.

All other questions are short-answer unless otherwise indicated.

In all the multiple choice questions more than one of the choices may be correct. Give all correct answers.

In the true/false and multiple choice questions, negative points will be awarded for incorrect answers in such a way as to render random guessing valueless.

Panic not.

1. (12 pts.) Definitions

Provide brief, *precise* definitions of the following:

- (a) Consistency (of a knowledge base)
- (b) Partial-order planning
- (c) Perfect rationality
- (d) Quiescence search
- (e) Context-free language
- (f) Qualification problem

2. (10 pts.) Logic True/false:

- (a) Using Extended Modus Ponens in the forward direction, one can generate all logical consequences of a set of sentences.
- (b) The resolution inference rule is inadequate to show consistency of a set of sentences.
- (c) An existentially-quantified sentence with one variable can be instantiated several times as long as the skolem constants used are all distinct.
- (d) When each property is applicable to only a small number of classes, and each object has many properties, inheritance is best implemented as forward chaining, given that we must use a logical system.
- (e) “All goblins are taller than all gnomes” is a good translation of
(ALL x y (> (Height (Goblin x)) (Height (Gnome y))))
- (f) (EXISTS x (AND (Number x) (ALL y (IF (Number y) (* x y x))))))
is a good translation of “There is a multiplicative identity”.

3. (8 pts.) KR Multiple choice:

If you fall into the Representation Trap, which of the following problems will arise?

- (a) Nothing can be solved because the system doesn't know the meaning of the long predicate names
- (b) Too many facts will be needed to construct a general-purpose system
- (c) Rules for your particular problems will become too complex

- (d) The TA will take points off

4. (15 pts.) Search

- (a) *True/false*: The heuristic “0.5 x (Number-of-misplaced-tiles + Total-Manhattan-Distance)” is admissible for the 8-puzzle.
- (b) The IDA* algorithm has the following description:

Beginning with an f-bound equal to the f-value of the initial state, perform a depth-first search bounded by the f-bound instead of a depth bound. Unless the goal is found, increase the f-bound to the lowest f-value found in the previous search that exceeds the previous f-bound, and restart the depth-first search.

Give the sequence of nodes visited by IDA* on the problem with initial state ((8 1 3) (2 0 4) (7 6 5)) and the standard final state ((1 2 3) (8 0 4) (7 6 5))

assuming that successors are generated in the order left, right, up, down for the motions of the blank tile.

- (c) Suppose A* is known to generate k nodes on a given problem, and *every node has a distinct f value*. Roughly how many nodes will IDA* generate on the same problem?

5. (14 pts.) Planning Multiple choice:

Unpredictability of action outcomes in planning has the following implications:

- (a) Decision theory cannot be used
- (b) Probabilities must be used
- (c) Planning for the distant future is computationally infeasible
- (d) Plans with long sequences of interdependent actions are inappropriate
- (e) Execution monitoring is useful
- (f) Conditional plans are useful
- (g) Observation actions cannot be included in plans

6. (5 pts.) NLP

- (a) (2) How does having a grammar aid communication?
- (b) (3) Write a STRIPS-style description for the action (**utter x s p**) of person x uttering a sentence s that informs some previously ignorant individual p of a fact; include any preconditions that are relevant.

7. (13 pts.) Logic/NLP

- (a) Give a context-free grammar for predicate calculus (as used in class; don't forget quantifiers, but don't worry about free variables). Include at least one example for each open syntactic class.
- (b) Show the parse tree for
(ALL x y (> (Height (Goblin x)) (Height (Gnome y))))

8. (10 pts.) Concept learning True/false:

In this question, we are dealing with the problem of learning a correct definition for a goal concept Q from a set of instances. A *hypothesis* is a candidate definition for Q.

- (a) A hypothesis is consistent with an instance if it predicts that the instance will be a member of Q.
- (b) A current-best-hypothesis algorithm always specializes its current hypothesis when the hypothesis disagrees with a positive example.
- (c) The version-space algorithm need not store previous examples.
- (d) When the version space collapses, at least one example must have been incorrectly classified.
- (e) If the concept language allows disjunctions, the S set always has one member.

9. (13 pts.) Probabilistic reasoning

- (a) (3) Draw a suitable network topology for the following set of (two-valued) nodes:
FrozenBattery IcyWeather CarWon'tStart NoGas
where **FrozenBattery** means *really* frozen, and **NoGas** means *really* empty.
- (b) (1) *True/false*: Once the values of a node's parents are fixed, the value of the node is fixed.
- (c) (2) Give reasonable conditional probability tables associated with the **FrozenBattery** and **CarWon'tStart** nodes in your network.
- (d) (1) How many independent values are contained in the joint probability distribution for four two-valued nodes, assuming no conditional independence relations are known to hold among them?
- (e) (2) How many independent probability values do your network tables contain (include priors as appropriate)?
- (f) (2) Suppose we add another node for a broken starter motor, called **BrokenSM**. Where would it go? Say briefly how your existing table for **CarWon'tStart** should be changed.
- (g) (2) The conditional probability table for **CarWon'tStart** is a canonical one. Describe, in English, what its structure will be in general.