# CS322 Fall 1997
# Practice Final Questions

Here are some questions for the final exam.

The following chapters/sections will be covered: Chapter 1, Chapter 2, Chapter 3 (except Section 3.7), Chapter 4, Chapter 5, Chapter 6 (except 6.7 and 6.8), Sections 7.2, 7.3, 7.4, Chapter 8 (excluding the event calculus and partial-order planning), Sections 11.1 and 11.2.

This practice exam is designed to give you some impression of the types of questions that may be asked. You should also expect some questions based on the assignments. **You should also reread the practice midterm. The final exam covers the whole course, but this practice exam emphasizes the latter part of the course.**

Note that you can bring in one sheet of $8\frac{1}{2} \times 11$ paper, with anything you like written on it. **You must bring your student card to the final examination.** Your student card will be checked.

## Question 1

Given the object-level knowledge base *KB* containing the clauses:

$false \Leftarrow a \& b.$
$false \Leftarrow c \& d.$
$a \Leftarrow p.$
$b \Leftarrow r.$
$b \Leftarrow e.$
$e \Leftarrow p.$
$c \Leftarrow f \& r.$
$f \Leftarrow g \& p.$
$f \Leftarrow q \& s.$
$f \Leftarrow u.$

Suppose that the assumables are $p$, $q$, $r$, $s$, $t$ and $u$.

1. What are all of the minimal conflicts? You will lose marks for omitting minimal conflicts and for giving things that are not minimal conflicts.

2. Give an SLD-resolution showing how how a delaying meta-interpreter can be used to find one of these minimal conflicts.

## Question 2

Consider the data on 4 Boolean attributes $a$, $b$, $c$, and $d$, where $d$ is the target classification.

| | $a$ | $b$ | $c$ | $d$ |
|---|---|---|---|---|
| $e_1$ | true | true | false | false |
| $e_2$ | false | true | false | true |
| $e_3$ | false | true | true | true |
| $e_4$ | false | false | true | false |
| $e_5$ | true | false | false | false |

In this question we will consider decision-tree learning based on this data.

1. What is a good attribute to split on first? Explain why.
2. Draw a decision tree that the top-down myopic decision tree learning algorithm could build. For each node (including the leaves) show which examples are used to determine the classification at that node. (The root note of the tree will be labelled with the list of all of the examples).
3. Explain how the learning bias inherent in learning decision-trees can be used to classify unseen instances. Give an instance that is not in the training data, show how the above tree classifies that instance. Justify why this is an appropriate classification.

## Question 3

Suppose that you are new in a job, and that your boss thinks that, because you did so well in cs322, you are an expert in AI. You want to impress your boss (and don't want to say "That course was really dumb. We didn't learn anything."). The boss has heard of neural networks and decision trees, but doesn't know anything about them, and thinks that neural nets seem so great and decision trees so stupid (or the other way around). The boss wants your informed opinion. (Your informed opinion should be based on facts, which you should make explicit). You need to write a coherent, well written half a page executive summary (using proper English) to impress the boss.

## Question 4

Suppose that someone has suggested using using $A^*$ search for neural network learning. Is this feasible? Either say how it could work, or give three reasons why $A^*$ search is not suitable for neural network learning. You need to use full sentences in your answer.

## Question 5

Suppose that you have the action for the robot to spray everything at its location with paint of a certain color, with the result that everything at its location becomes that color.

- How can such an action be defined in the situation calculus?
- How can such an action be defined in STRIPS?

## Question 6

Suppose you have a base-level knowledge base defined in terms of

- the meta-level predicate symbol $\Leftarrow$ , such that $H \Leftarrow B$ is true if $H \leftarrow B$ is in the base-level knowledge base, and
- the relation $primitive(A)$ that is true if atom $A$ is primitive.

An example knowledge base is

$at(Ag, Pos) \Leftarrow sitting_at(Ag, Pos)$.
$adjacent(P_1, P_2) \Leftarrow$
    $between(Door, P_1, P_2) \& $
    $unlocked(Door)$.
$between(door1, o103, lab2) \Leftarrow true$
$primitive(unlocked(D))$.
$primitive(sitting_at(A, P))$.
$primitive(carrying(A, O))$.

Write a meta-interpreter $reduce_prove(L0, L1)$ that, given a list $L0$ of atoms, reduces them to a list $L1$ of primitive atoms that, given the knowledge base, implies the elements of $L0$.

For example, the query

ask $reduce_prove([at(rob, P), adjacent(P, lab2), carrying(rob, parcel)], L)$

returns with one answer with $L = [sitting_at(rob, o103), unlocked(door1), carrying(rob, parcel)]$. Note that this operation was assumed in the regression planner.

## Question 7

Suppose a regression planner was used with the goal:

$[sitting_at(rob, o109), sitting_at(parcel, lab2)]$

- Draw the first-level of the search tree showing the possible actions and the weakest preconditions to achieve this goal.
- Show one element on the next level of the search tree. (That is one action to achieve one of the goals you constructed in the preceding part, and the corresponding weakest precondition.)
- What happens when the naive STRIPS planner is used with this goal? (A qualitative description of what will happen is what is needed).
- Describe qualitatively what happens when STRIPS is used so that it protects subgoals.