



EXAMINATIONS — 2003

END-YEAR

COMP 307
ARTIFICIAL INTELLIGENCE

Time Allowed: 3 Hours

Instructions: There are a total of 180 marks on this exam.
Attempt all questions.
Calculators may be used.
Non-electronic foreign language translation dictionaries may be used.

Questions

- | | |
|---------------------------------|------|
| 1. Prolog | [30] |
| 2. Search | [30] |
| 3. Rule Based Systems | [30] |
| 4. Planning | [35] |
| 5. Natural Language Processing. | [25] |
| 6. Machine Learning | [30] |

SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.

Question 1. Prolog

[30 marks]

(a) [12 marks] Give the most general unifier for the following pairs of terms, or say that the terms do not unify. The first answer is given as an example.

property(One, Two)
property(3, 4)

Unifies with One=3, Two=4

date(D, may, Y)
date(17, M, Z)

collection([1,2], X)
collection([H|T], Y)

X
2+3

family(leon, liz, [leon,danya])
family(X, liz, [X])

foo(X, [1, X, 2])
foo(a, Z)

lect(X, room(hu1), [comp307, X])
lect(Y, room(Y), Z)

(Question 1 continued on next page)

(Question 1 continued)

(b) [8 marks] Matching, Backtracking and Cut

Consider the following program.

```
find(X):- p(X), q(X), r(X).
```

```
q(X):- a(X).
```

```
q(X):- b(X), !, c(X).
```

```
q(X):- d(X).
```

```
p(liz).
```

```
p(leon).
```

```
p(andrew).
```

```
a(liz).
```

```
b(liz).
```

```
b(leon).
```

```
c(liz).
```

```
d(liz).
```

```
d(leon).
```

```
d(andrew).
```

```
r(liz).
```

```
r(leon).
```

```
r(andrew).
```

What is the output generated by the following query

```
|-find(X).
```

Give all the possible answers. If an answer is found more than once, show all the occurrences.

(Question 1 continued on next page)

(Question 1 continued)

(c) [10 marks] Lists and Recursion

Write a program for `sublist(+List, +N, ?List1, ?List2)` where `List1` is the first `N` elements of `List` in the same order, and `List2` is the rest of `List` in the same order. For example,

```
sublist([2, 4, 5], 0, X, Y) is true with X=[], Y=[2, 4, 5]
sublist([2, 4, 5], 2, X, Y) is true with X=[2, 4], Y=[5]
sublist([2, 4, 5], 3, X, Y) is true with X=[2, 4, 5], Y=[]
```

Make sure that your program can handle the special cases of `N` being negative or larger than the length of `List`:


```
sublist([2, 4, 5], -5, X, Y) is false and should say no.
sublist([2, 4, 5], 100, X, Y) is true with X=[2, 4, 5], Y=[]
```



Question 2. Search

[30 marks]

Consider the search tree shown below.



The initial state is represented by node A. The goal state is represented by node Q. The number on a link is the cost of the operator that links the two nodes. Each node has an associated number.

In the questions below, assume that the searches consider child nodes from left to right. If you make any additional assumptions, state them clearly.

For the uninformed search strategies in (a) to (d), specify the order in which the nodes will be expanded.

(a) [3 marks] **Breadth first search**

(b) [3 marks] **Depth first search**

(c) [4 marks] **Depth limited search with a limit=2**

(Question 2 continued on next page)

(Question 2 continued)

(d) [4 marks] Iterative deepening search

For the informed searches in (e) to (g), specify the order in which the nodes will be expanded. Assume that the number on a node represents a (lower bound) heuristic estimate of the cost of reaching a goal from the node.

(e) [4 marks] Uniform cost search (best first search using the path cost only).

(f) [4 marks] “Greedy” best first search (best first search using the heuristic estimate only)

(g) [4 marks] A* search

For the hill climbing search below, specify the order in which the nodes will be expanded. In this case, assume that the number on a node is the value of the node, and we want to find the node with the lowest value.

(h) [4 marks] Hill climbing

Question 3. Rule Based Systems

[30 marks]

Consider the following set of rules. Goal1 and Goal2 are goals to be established. H1, H2, H3, H4 are intermediate inferences, and Q1, Q2, . . . , Q6 are the questions that can be asked of the user.

R1: if H1 and H2 then Goal1

R2: if H2 and H3 then Goal2

R3: if H4 then Goal2

R4: if Q1 then H1

R5: if Q2 then H1

R6: if Q3 then H2

R7: if Q4 then H2

R8: if Q5 then H3

R9: if Q6 then H4

Assume the rules are scanned top down, that is, a rule with a smaller number has higher priority.

(a) [8 marks] **Backward chaining**

This question uses backward chaining. If a particular question is asked, the user would respond as follows:

Q1: YES

Q2: No

Q3: No

Q4: No

Q5: YES

Q6: YES

Give the sequence of rules that a backward chaining interpreter would apply to determine whether any goal is true. (List the rule numbers only).

Assume there is **no** caching for intermediate results, and assume backtracking is always allowed (cuts are not used).

(b) [8 marks] **Forward chaining**

This question uses forward chaining. Suppose Q1, Q5, Q6 are true and saved in the working memory.

Give the sequence of rules that a forward chaining interpreter would apply to determine whether any goal is true. (List the rule numbers only).

(c) [6 marks] In general, would you expect forward and backward chaining to give the same conclusions on the same rule set? Why or why not?

(d) [8 marks] The lectures introduced a rule-based system for animal classification that worked fine in a toy world (micro-world). Now consider building a rule-based system for a real world animal classification problem. What are some of the limitations of the rule representation and the reasoning strategy that we introduced?

Question 4. Planning

[35 marks]

(a) [4 marks] What are the four parts of a STRIPS operator?

(b) [2 marks] What constraint on the variables in the four parts must be satisfied to make it a valid STRIPS operator?

(c) [4 marks] Write a STRIPS operator for the action **unlock**. To unlock a door, the robot needs to have a key, the door needs to be locked (and shut), and the key should be the right one for the door.

(Question 4 continued on next page)

(Question 4 continued)

(d) [4 marks] The unlock operator you wrote above almost certainly does not represent all the possible preconditions for successfully unlocking a door. Describe at least four additional preconditions for unlocking a door that a person could reason about.

(Hint: consider situations where the robot could not unlock the door successfully.)

(e) [4 marks] There are also many possible effects of unlocking a door that could occur in some circumstances. For example, the door may blow up if there is a bomb wired to the lock, or a person may walk into the house if they were waiting outside the door. Show how you could represent these two conditional effects using a set of STRIPS operators.

(Question 4 continued on next page)

(Question 4 continued)

Consider the following planning problem. A robot has to put together a kitset table. The goal is for the legs to be coloured black, the table top to be smooth and coloured red, and the top to be joined to the legs. The initial state has the legs and top separate, and both coloured black, and the top is not smooth. The robot has three operators, to join parts together, to paint an object a new colour (black or red), and to sand an object to remove any paint and make it smooth.

Initial State: `separate(legs)`, `separate(top)`, `colour(legs, black)`, `colour(top, black)`

Goal: `joined(top, legs)`, `smooth(top)`, `colour(legs, black)`, `colour(top, red)`

Operators:

`assemble(Obj1, Obj2):`

pre: `separate(Obj1)`, `separate(Obj2)`

del: `separate(Obj1)`, `separate(Obj2)`

add: `joined(Obj1, Obj2)`

`paint(Object, OldColour, NewColour):`

pre: `separate(Object)`, `colour(Object, OldColour)`

del: `colour(Object, OldColour)`

add: `colour(Object, NewColour)`

`sand(Object, OldColour):`

pre: `separate(Object)`, `colour(Object, OldColour)`

del: `colour(Object, OldColour)`

add: `smooth(Object)`, `colour(Object, wood)`

(f) [6 marks] Suppose the robot used a goal-stack planner to construct a plan for this task. On the facing page, show the steps that the goal-stack planner would take, up to the point that it has added the first action to the current plan. Show the items added to and removed from the Goal Stack, the changes to the current state, and the action added to the current plan.

Assume that the subgoals of a conjunction are added to the stack in reverse order, so that the first subgoal is at the top of the stack.

colour(top, black)
colour(legs, black)
separate(top)
separate(legs)

Current State

joined(top,legs) & smooth(top) & colour(top, red) & colour(legs,black)

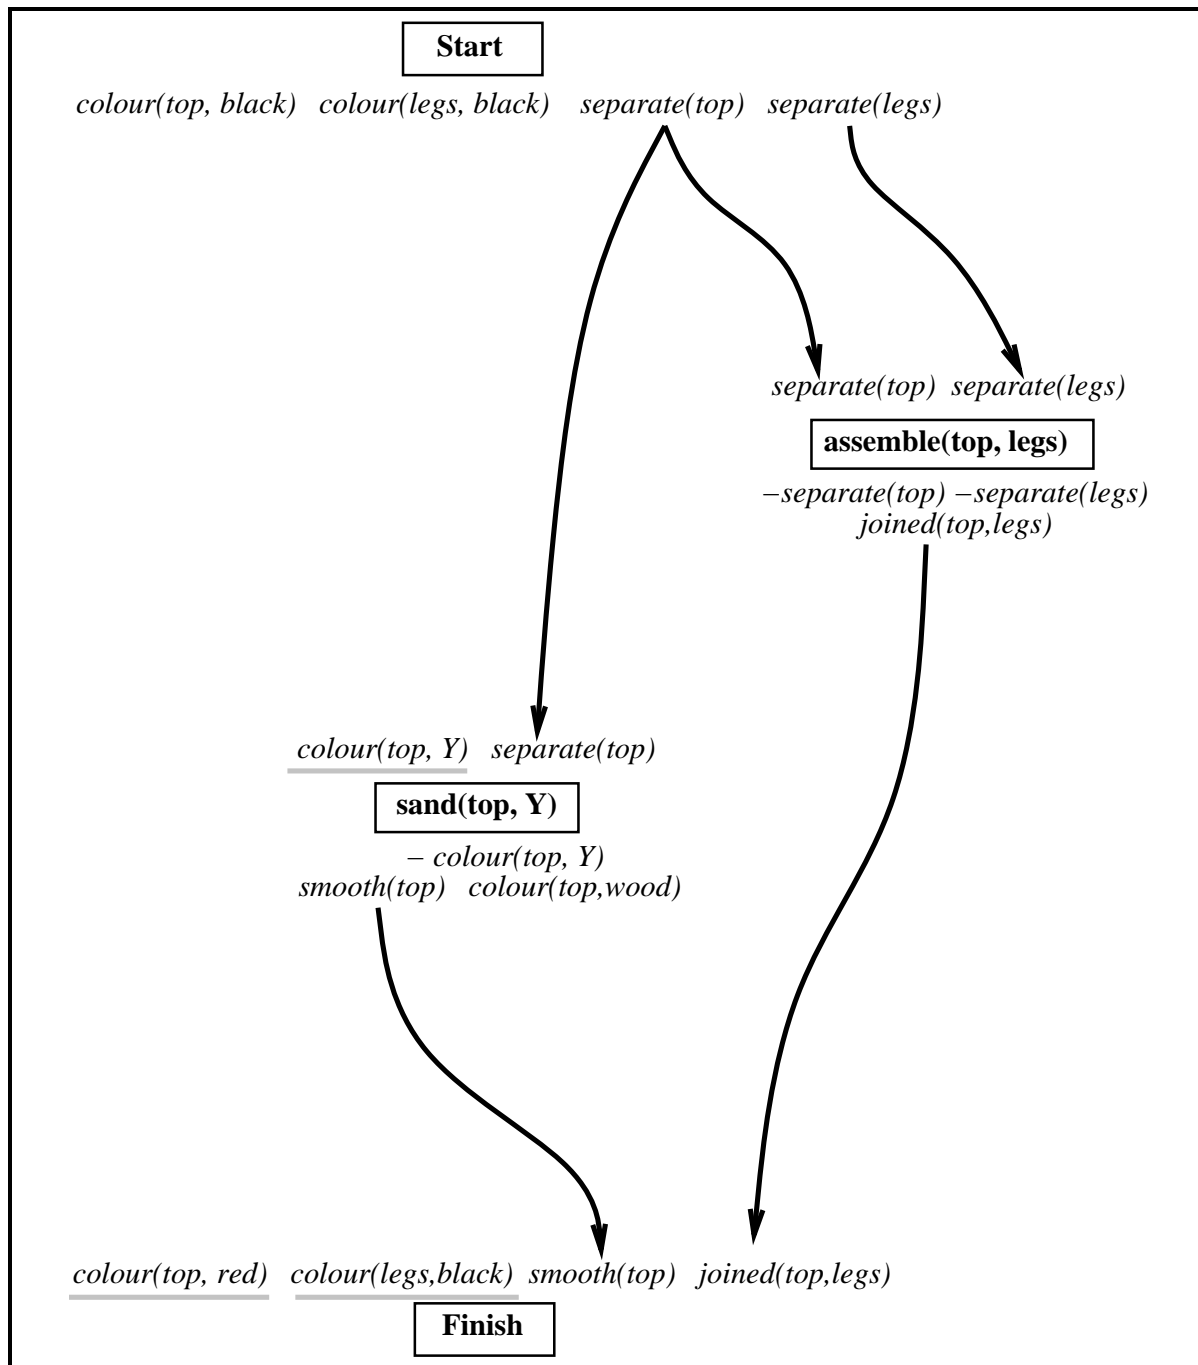
Goal Stack

Current Plan:

(Question 4 continued)

Suppose the robot used a Partial Order Planner (POP) to construct a plan for the same assembly task. The diagram below shows the state of the plan after the planner has added two actions to satisfy some of the subgoals, and has satisfied some of their preconditions.

- (h) [3 marks] Show the threat that exists in the plan, and resolve it.
- (i) [2 marks] Satisfy the other precondition of the `sand(top, Y)` action using an existing action.
- (j) [3 marks] Add a paint action to satisfy the subgoal `colour(top, red)`, and show any threats that are introduced.



SPARE PAGE FOR EXTRA ANSWERS

Cross out rough working that you do not want marked.
Specify the question number for work that you do want marked.

Question 5. Natural Language Processing

[25 marks]

Consider the following grammar and lexicon for simple yes/no questions such as
was the toaster on the bench?
did the robot wash a bowl?
has the toaster eaten the bread and the milk?

Simple yes/no questions like these start with “auxilliary” verbs — a special, restricted set of verbs including “is”, “was”, “were”, “does”, “did”, “has”, “have”, and “had”.

```
s --> aux, np, pp.
s --> aux, np, vp.

pp --> prep, np.

vp --> verb, np, pp.
vp --> verb, np.
vp --> verb.

np --> det, noun, pp.
np --> det, noun.

prep --> [P], {lexicon(P, prep)}.
det --> [D], {lexicon(D, det)}.
noun --> [N], {lexicon(N, noun)}.
aux --> [V], {lexicon(V, aux, _ )}.
verb --> [V], {lexicon(V, verb)}.

lexicon(wash, verb).
lexicon(washing, verb).
lexicon(washed, verb).
lexicon(eat, verb).
lexicon(eating, verb).
lexicon(eaten, verb).

lexicon(robot, noun).
lexicon(toaster, noun).
lexicon(bench, noun).
lexicon(bread, noun).
lexicon(milk, noun).
lexicon(bowl, noun).
lexicon(kitchen, noun).

lexicon(is, aux, be).
lexicon(was, aux, be).
lexicon(were, aux, be).
lexicon(does, aux, do).
lexicon(did, aux, do).
lexicon(has, aux, have).
lexicon(have, aux, have).
lexicon(had, aux, have).

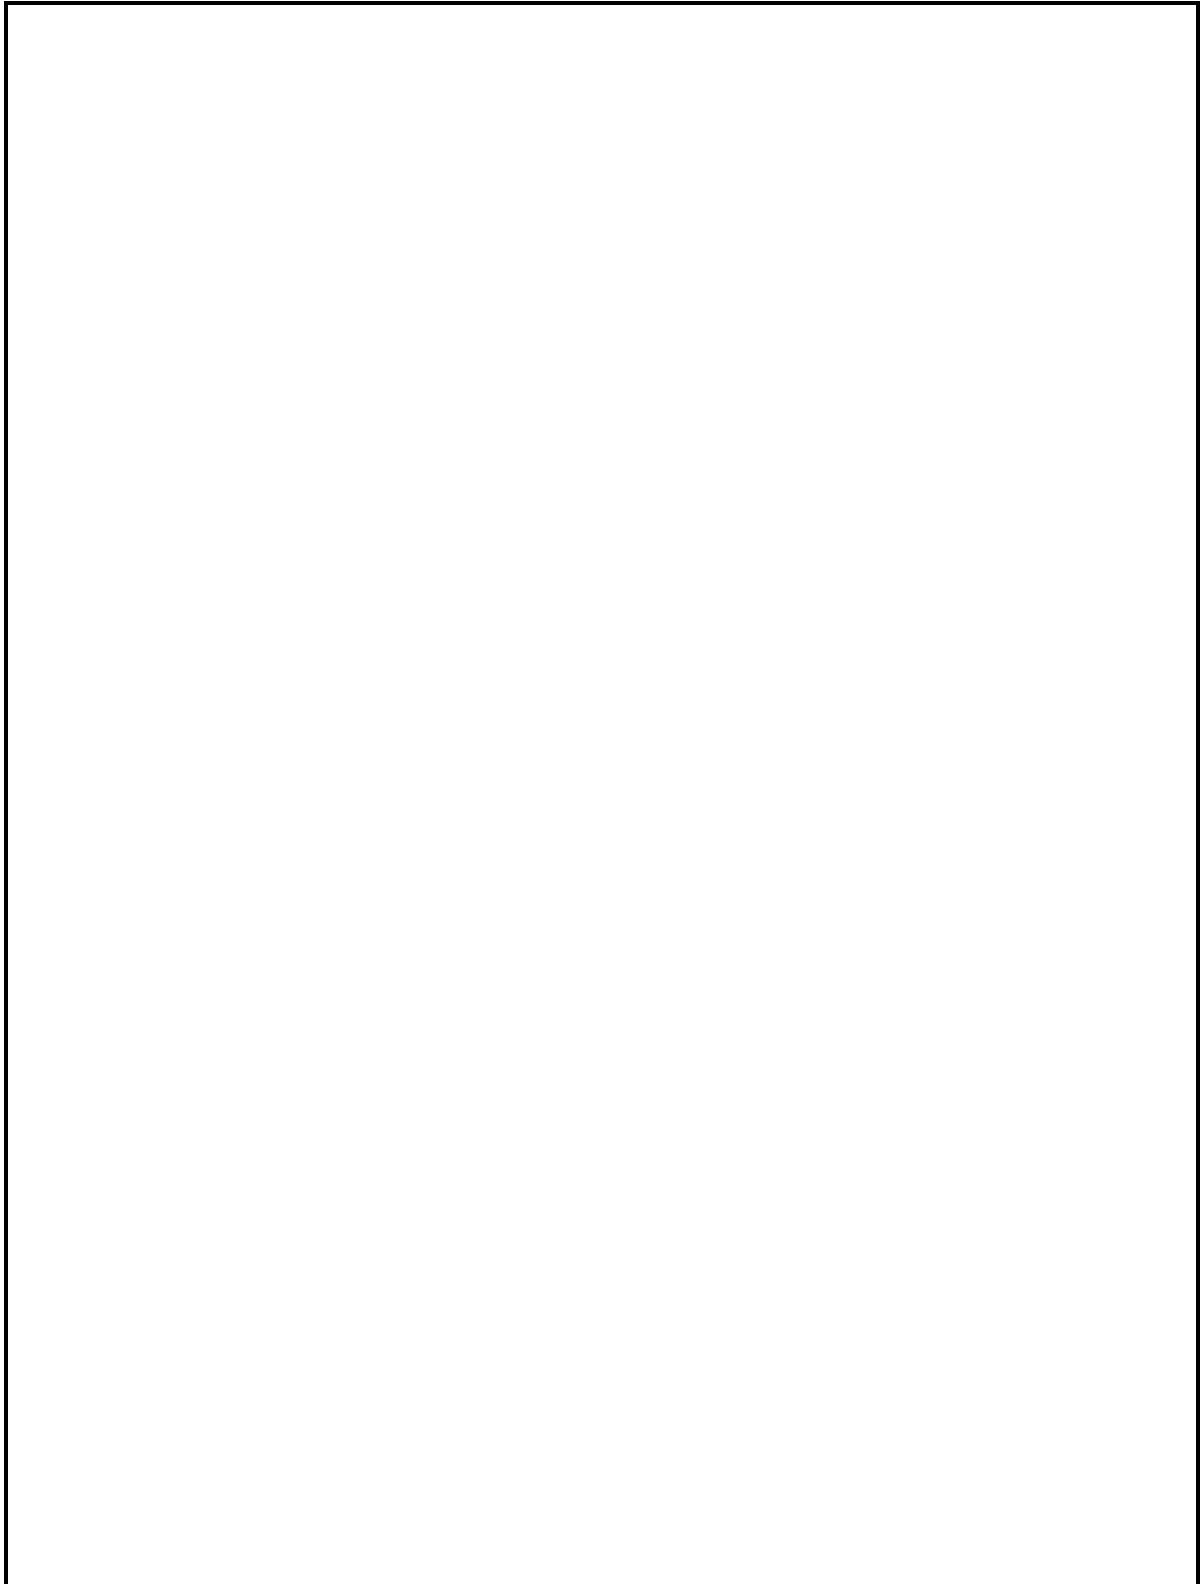
lexicon(on, prep).
lexicon(in, prep).

lexicon(the, det).
lexicon(a, det).

lexicon(and, conj).
lexicon(or, conj).
```

(a) [4 marks] Show the parse tree for the question:
is the bread in the toaster?

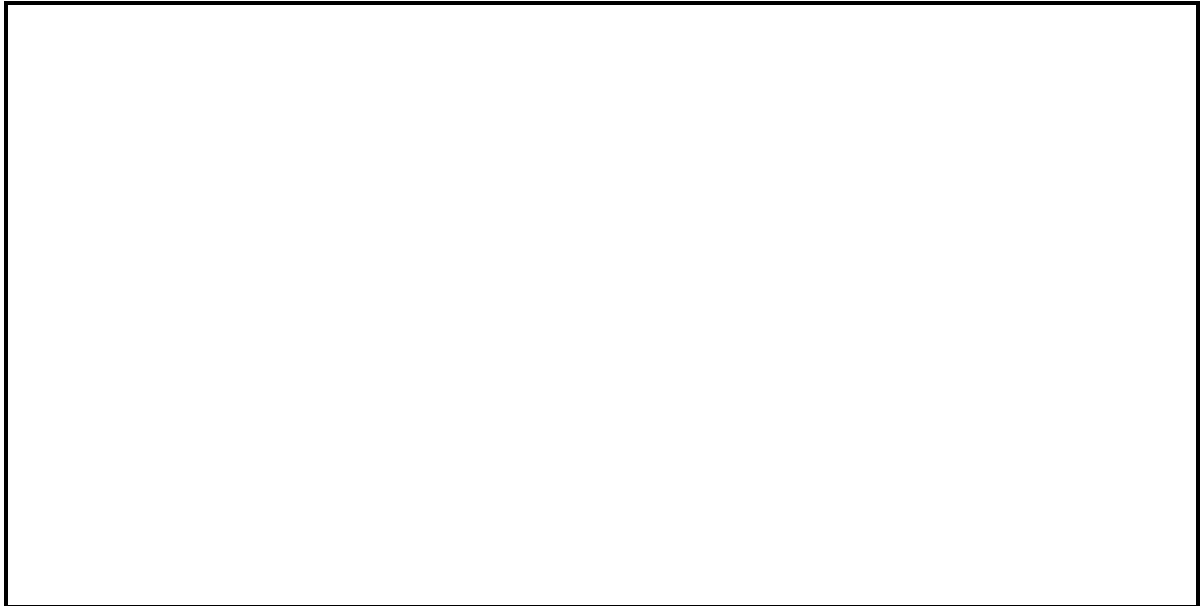
(b) [5 marks] Show three possible parse trees for the question:
did the robot wash the bowl on the bench in the kitchen?



(Question 5 continued on next page)

(Question 5 continued)

(c) [3 marks] Explain the different meanings of the parse trees from part (b).
Hint: which bowl is being referred to and where did the washing happen?



(d) [5 marks] Noun phrases can contain two sub phrases, joined with a conjunction (“and” or “or”).
For example:

Did the robot and the toaster eat the bread?

Is the bread in a bowl on the bench or the toaster?

Modify the grammar to allow a noun phrase to contain two sub phrases connected with conjunctions, as in the sentences above. (Note that the lexicon already contains entries for “and” and “or”.)



(Question 5 continued on next page)

(Question 5 continued)

(e) [8 marks] The original grammar is too general – it allows invalid sentences, like the three on the left below, although the three on the right are valid.

* is the robot washed the bowl	is the robot washing the bowl
* did the toaster eating the bread	did the toaster eat the bread
* did the toaster on the bench	is the toaster on the bench

One problem is that the auxilliary verb at the beginning of the sentence must match the form of the verb in the vp (if there is a vp).

- The “be” type auxilliaries (“is”, “was”, “were”) must have a verb that ends with “ing”.
- The “have” type auxilliaries (“has”, “have”, “had”) must have a verb that ends with “en” or “ed”.
- The “do” type auxilliaries (“do”, “did”, “does”) must have a verb with no ending.

The other problem is that a sentence with no vp (the first grammar rule), can only start with a “be” type auxilliary.

Extend the grammar and the lexicon to enforce these constraints. You do not need to write out rules that are the same as the original grammar and lexicon.

Question 6. Machine Learning

[30 marks]

(a) [6 marks] Suppose that the Naïve Bayes algorithm is used to classify an instance described by a vector of n attributes (A_1, A_2, \dots, A_n) into one of two categories, C_1 and C_2 . Given an unclassified instance described by the vector (x_1, x_2, \dots, x_n) , the Naïve Bayes algorithm computes a score for each category and returns the category with the highest score. The Naïve Bayes learning algorithm analyses its training data to obtain a set of probabilities for computing the scores.

Explain what probabilities the learning algorithm obtains from the training data, and how it computes the two scores for a new instance.

(b) [2 marks] What is the assumption that the Naïve Bayes classifier makes about the data that is almost always wrong?

(c) [4 marks] Decision trees can also be used to classify instances into one of two categories. The central action of the decision tree learning algorithm is to work out what attribute to use for the decision at a node of the tree. Given a node that has some instances from each category, explain how the decision tree chooses an attribute to use for the node. State the formula for the score of an attribute. Call the attribute “A” and assume that it has just two values v_1 and v_2 . Call the classes C_1 and C_2 .

(Question 6 continued on next page)

(Question 6 continued)

(d) [6 marks] Consider the following data set describing 12 species of trees of which 6 are good for timber, and 6 are not. They are described by three binary attributes: the kind of leaf, the hardness of the wood, and the pattern of the branches.

	Leaf	Wood	Branching	Class
G1	needle	soft	irregular	Good
G2	needle	soft	regular	Good
G3	needle	soft	regular	Good
G4	needle	hard	regular	Good
G5	needle	hard	regular	Good
G6	flat	hard	regular	Good
B7	flat	soft	irregular	Bad
B8	flat	hard	regular	Bad
B9	flat	soft	regular	Bad
B10	flat	soft	regular	Bad
B11	needle	hard	regular	Bad
B12	needle	hard	irregular	Bad

Which attribute would the decision tree building algorithm choose for the root of a decision tree for the timber database? Show your working and state the impurity measure you are using.

(Question 6 continued on next page)

(Question 6 continued)

(e) [4 marks] Perceptrons are a classification mechanism that use a set of weights and a threshold. Given a set of examples (positive and negative), the perceptron learning algorithm is able to learn a threshold and set of weights. Explain the major limitation of the perceptron learning algorithm.

(f) [8 marks] Draw a diagram of a multilayer feedforward neural network, and identify the input, hidden, and output nodes. Explain why the network does not suffer from the same limitation as the perceptron, particularly noting the role of the hidden nodes.
