EXAMINATIONS — 1999

COMP 307

Introduction to Artificial Intelligence

Time Allowed: 3 Hours

Instructions: Attempt ALL Questions.

The exam will be marked out of 180.
Non-programmable calculators without full alphabet keys are permitted.
Non-electronic foreign language dictionaries are permitted.

Questions

	Marks
1. Goal Stack Planning	[30]
2. Representation for Planning	[25]
3. Rule Systems	[30]
4. Decision Trees	[35]
5. Language and Grammars	[60]

Questions 1 to 4 concern an intelligent car-buying agent that assists a person to buy a car. The agent will need to travel to the different car dealers in the city, test drive the cars, evaluate them, and buy a good one.

Question 1. Goal Stack Planning [30 marks]

In order to find out about the cars at all the dealers, the agent must construct a plan for traveling to each dealer, and test driving their cars. The agent can travel by foot or by bus.

To construct a plan, the agent needs to represent knowledge about the world and the actions the agent can perform:

Ontology:

- The agent's world contains two kinds of locations: bus stops and dealers. The predicates busStop and dealer are true of the two kinds of locations.
  The closeTo predicate specifies whether two locations are close enough to walk between.
- The agent has three actions:
  walk(Loc1,Loc2), which moves the agent between locations (if they are close enough),
  bus(Loc1,Loc2), which moves the agent between two bus stops, and
  testdrive(Dealer), which gets the agent to test drive all the cars at the dealer.
- There are two predicates for describing the agent's state:
  - at(Loc), which specifies the current location of the agent, and
  - knowsAbout(Dealer), which specifies that the agent knows about the cars at the given Dealer.

Actions, represented by STRIPS Operators:

```
walk(From, To):
    pre: at(From), closeTo(From, To)
    del: at(From)
    add: at(To)

bus(From, To):
    pre: at(From), busStop(From), busStop(To)
    del: at(From)
    add: at(To)

testdrive(Dealer):
    pre: at(Dealer), dealer(Dealer)
    del:
    add: knowsAbout(Dealer)
```

Facts about the world.

Assume that the agent's world contains three dealers and two bus-stops:

dealer(wellingtonCars)	closeTo(wellingtonCars, cityAutos)	closeTo(wellingtonCars, stop1)
dealer(cityAutos)	closeTo(cityAutos, wellingtonCars)	closeTo(cityAutos, stop1)
busStop(stop1)	closeTo(stop1, cityAutos)	closeTo(stop1, wellingtonCars)
dealer(huttVehicles)	closeTo(huttVehicles, stop2)	
busStop(stop2)	closeTo(stop2, huttVehicles)	

(Question 1 continued on next page)

(Question 1 continued)

(a) [4 marks] What is an optimal plan (fewest steps) for the agent to find out about the cars at all three dealers, assuming that in the initial state it is at(stop1) and does not know about the cars at any of the dealers?

(b) [6 marks] Suppose the agent used the Goal Stack Planner (as described in the course) to construct a plan to achieve the following goal given the same initial state as in part (a):

goal: [knowsAbout(wellingtonCars), knowsAbout(huttVehicles), knowsAbout(cityAutos)]

What plan would the agent come up with? You may assume that it has a smart rate_operators that will choose the best of the candidate operators for each subgoal.

(c) [2 marks] What initial sequence of the goals would allow the Goal Stack Planner to come up with the optimal plan?

(Question 1 continued on next page)

(Question 1 continued)

(d) [8 marks] Why doesn't the Goal Stack Planner come up with the optimal plan?

(e) [10 marks] Outline a modification to the Goal Stack Planner to enable it to find an optimal plan that finds out about all the dealers, regardless of the order in which the goal is stated. (Your answer should be a modification of the Goal Stack Planning algorithm, not a version of the POP algorithm.)

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 2. Representation for Planning [25 marks]

(a) [7 marks] The STRIPS operators in question 1 are not the only way of representing knowledge about actions. Give Situation Calculus axioms that represent the walk(Loc1, Loc2) action. You should use the predicates described in question 1.

Suppose we wanted the agent to be able to buy a car for us. To work out a plan for buying the car, the agent needs to know about the action of buying.

We might give the agent the following STRIPS operator:

buy(Item, Seller, Buyer, Price):
  pre: owns(Seller, Item), owns(Buyer, Price), fairPrice(Item, Price)
  del: owns(Seller, Item), owns(Buyer, Price)
  add: owns(Seller, Price), owns(Buyer, Item)

An initial state in which cityCars owns a car worth \$1500 and the agent has \$5000 could be represented by:

[owns(cityCars, honda-HC3921), owns(agent, \$5000), fairPrice(honda-HC3921, \$1500)]

(Question 2 continued on next page)

(Question 2 continued)

(b) [8 marks] Why would a planner using the buy STRIPS operator on the facing page **not** be able to construct a plan to achieve the goal owns(agent, honda-HC3921), given the initial state above?

(c) [10 marks] Outline a representation for the buy action that could address this problem and suggest how the Goal Stack Planner could be extended to use this kind of action.

Question 3. Rule systems [30 marks]

(a) [8 marks] Describe how a backward chaining rule system is able to provide two kinds of explanation to a user. In your answer, identify the information that the rule system must record about its reasoning to be able to provide these explanations.

The car-buying agent would need some knowledge about cars in order to evaluate a car that it took on a test drive. Suppose that part of this knowledge was the fact that loose screws under the dashboard can cause a metalic rattling sound when the car is being driven. This fact could be formulated as an if-then rule:

if	Driving a car &
	Metalic rattling sound
then	Loose screws under the dashboard

(Question 3 continued on next page)

(Question 3 continued)

The car-buying agent notices that there is a metalic rattling sound while it is test driving one car and, using this rule, concludes that the car must have loose screws under the dashboard.

(b) [2 marks] What kind of inference is the agent doing?

(c) [6 marks] In fact, the agent's conclusion is wrong. Give at least two different kinds of reason why the agent's conclusion may be wrong.

(d) [4 marks] Suggest how the rule could be modified to make the rule more accurate.

(Question 3 continued on next page)

(Question 3 continued)

(e) [10 marks] When a rule-based expert system is attempting to classify an object, it may ask the user whether some property of the object is true or false. However, the user may not know whether it is true or false, so that the expert system should give the user the option to answer “I don’t know”. Suggest how a backward chaining rule-based system should use a “don’t know” answer.

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 4. Decision Trees [35 marks]

The user of the car-buying agent needs to tell the agent what kind of cars they like. The user has shown the agent a collection of 17 cars and classified them into “OK” cars and “bad” cars. The agent now needs to use these training examples to construct a decision tree for classifying other cars that it looks at.

Here is a table showing the values of three attributes of each car, along with its classification as **“OK”** or **“bad”**. (“s/w” means “station wagon”)

	Class	Model	Age	Colour
1	**OK**	sedan	0–5	white
2	**OK**	sedan	0–5	white
3	**OK**	s/w	0–5	white
4	**OK**	s/w	0–5	white
5	**OK**	s/w	0–5	white
6	**OK**	s/w	0–5	red
7	**OK**	s/w	0–5	red
8	**OK**	s/w	6–20	red

	Class	Model	Age	Colour
9	**bad**	s/w	0–5	gold
10	**bad**	s/w	0–5	gold
11	**bad**	s/w	6–20	white
12	**bad**	sedan	6–20	white
13	**bad**	sedan	6–20	red
14	**bad**	sedan	6–20	red
15	**bad**	sedan	0–5	red
16	**bad**	sedan	0–5	red
17	**bad**	sedan	0–5	red

(a) [9 marks] If the agent uses the greedy decision tree building algorithm that was presented in lectures, which attribute would be chosen for the root of the decision tree? Show your working.

Assume that the algorithm minimises the average impurity of the nodes, and uses $p_{OK} \times p_{bad}$ as the impurity measure of a node, where p_{OK} is the fraction of instances in a node that are **OK** and p_{bad} is the fraction of instances in a node that are **bad**.

(Question 4 continued on next page)

(b) [5 marks] Show the complete decision tree that the algorithm would produce. (You do not need to show your working for the second level nodes.)

(c) [5 marks] Construct a collection of rules that a backward chaining rule system could use to classify

(Question 4 continued)

Suppose the user now wants to distinguish between *three* classes of cars: (*e.g.* "Love to buy", "Might buy", and "Won't buy"). At some point in building the decision tree, it might be that n instances have been classified down to one node in the tree, where n_ℓ of them are "Love to buy", n_m are "Might buy", and n_w are "Won't buy".

(d) [8 marks] Explain why the following formula would not be a good impurity (mixedness) measure for the node.

$$\frac{n_\ell}{n} \times \frac{n_m}{n} \times \frac{n_w}{n}$$

Hint: consider different extreme values of n_ℓ, n_m, and n_w, and compare to the two-class case.

(e) [8 marks] Suggest an alternative impurity measure for the three-class case.

Question 5. Language and Grammars [60 marks]

(a) [6 marks] Most implementations of Prolog support the DCG notation. Explain, in outline, why this is so, and explain briefly what phrase/2 does.

(Question 5 continued on next page)

(Question 5 continued)

The company you work for has decided to expand into New Caledonia, a French-speaking territory in the Pacific. You want to work there so you decide to impress your boss by implementing some tools for the French language. You have discussed the French language with a linguist, who has told you the following facts:

- A sentence consists of a noun phrase followed by a verb phrase.
- A noun phrase is either a proper noun or it is a determiner followed by a common noun.
- A common noun can be a common noun followed by an adjective.
- A verb phrase is either an intransitive verb or it is a transitive verb followed by a noun phrase.
- *André* is a proper name.
- *le* and *un* are both determiners.
- *gâteau* and *garçon* are common nouns.
- *mange* is a verb which can be either transitive or intransitive.
- *jete* is a transitive verb.
- *blanc* is an adjective.

(b) [14 marks] Write a DCG grammar which captures these rules for the formation of sentences.

(Question 5 continued on next page)

(Question 5 continued)

The linguist now informs you that, in French, all nouns have either masculine (male) or feminine (female) gender, and that determiners and adjectives come in masculine and feminine forms.

The data that you now have is that:

- A sentence consists of a noun phrase followed by a verb phrase.
- A noun phrase is either a proper noun or it is a determiner followed by a common noun.
- The determiner and the common noun in a noun phrase agree on gender.
- A common noun can be a common noun followed by an adjective.
- An adjective which follows a common noun agrees with the common noun on gender.
- A verb phrase is either an intransitive verb or it is a transitive verb followed by a noun phrase.
- *André* is a boy's name.
- *Marie* is a girl's name.
- *le* and *un* are both masculine determiners.
- *la* and *une* are both feminine determiners.
- *gâteau* and *garçon* are masculine common nouns.
- *pâtisserie* and *fi lle* are feminine common nouns.
- *mange* is a verb which can be either transitive or intransitive.
- *jete* is a transitive verb.
- *blanc* is a masculine adjective.
- *blanche* is a feminine adjective.

(c) [12 marks] Rewrite your grammar to capture this additional information.

(Question 5 continued on next page)

The linguist now supplies you with the following table of translations between French and English words:

French	English
André	*Andrew*
blanc, blanche	*white*
fi lle	*girl*
garçon	*boy*
gâteau	*cake*
jete	*throws*
le, la	*the*
mange	*eats*
Marie	*Mary*
pâtisserie	*cake*
un, une	*a*

(d) [14 marks] Explain how you could use DCG's to translate sentences between French and English, for example, between:

Andrew eats the cake and *André mange le gâteau,*

(Question 5 continued on next page)

(Question 5 continued)

(e) [14 marks] *Either*:

- Explain how linguistic information can be represented in unification-based grammars, such as PATR;

or

- Discuss the roles of grammar rules and the lexicon in representing linguistic information.

****

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.

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.