1 2004 - 2005

QUESTION START

(a) Explain, using natural language and/or an algorithm, the control loop of a practical reasoning agent.

book page 76, figure 4.3 [Marks: 5]

(b) Describe the three types of commitment an agent can hold.

book page 77, bottom of the page [Marks: 3]

(c) What does it mean for a plan to be correct

book page 74, (1), (2) [Marks: 2]

(d) Consider your agent is in a world formed by a 3-by-3 grid. The world contains a number of lost coins. Your agent is free to move in any direction. The goal of your practical agent is to search for the coins.

(i) Give a description of your initial world.
(ii) Give a description of the goal
(iii) Give the actions your agent can perform to complete the task of fetching all coins from its world.

(i) In(0,0), ToVisit(0,0), ToVisit(0,1), ToVisit(0,2), ToVisit(1,0), ToVisit(1,1), ToVisit(1,2), ToVisit(2,0), ToVisit(2,1), ToVisit(2,2), Coin(1,2), Coin(1,0)

(ii) Empty(X,Y) $\forall 0 \leq X, Y \leq 2$

(iii)
- Clear(x,y)
  pre {In(x,y), ToVisit(x,y)}
  del {ToVisit(x,y)}
  add {Empty(x,y)}
- MoveUp(x,y)
  pre {In(x,y), Empty(x,y), ToVisit(x+1,y)}
  del {In(x,y)}
  add {In(x+1,y)}
- MoveDown(x,y)
  pre {In(x,y), Empty(x,y), ToVisit(x-1,y)}
  del {In(x,y)}
  add {In(x-1,y)}

- MoveLeft(x,y)
  pre {In(x,y), Empty(x,y), ToVisit(x,y-1)}
  del {In(x,y)}
  add {In(x,y-1)}
- MoveRight(x,y)
  pre {In(x,y), Empty(x,y), ToVisit(x,y+1)}
  del {In(x,y)}
  add {In(x,y+1)}

[Marks: 5]

(e) Explain Steels' Mars Explorer

book page 92- 94 [Marks: 5]

QUESTION END

QUESTION START

(i) **State and discuss** the seven desirable properties of mechanism design

bookwork

(a) convergence/guaranteed success: the mechanism should terminate
(b) maximising social welfare: the allocation should be the best for all parties involved
(c) Pareto efficiency: the allocation should be optimal, in that any other allocation would leave some party worse off
(d) Individual rationality: participating parties decisions need only depend on their own rationality
(e) stability:
(f) simplicity:
(g) distribution:

[marks 7]

(ii) **Describe** the parameters defining English, Dutch, first-price sealed bid and Vickrey auctions and where appropriate give the dominant strategy.

bookwork

English: first-price, open outcry, ascending. Dominant strategy to increase price in small increments until agent has wins or it reaches their valuation, when the agent withdraws.

Dutch: first-price, open outcry, descending. There is no dominant strategy

First-price sealed bid: first price, one shot. Dominant strategy is to bid less than true valuation.

Vickrey: second price, one shot. Dominant strategy is to bid true valuation.

[marks 4*2]

(iii) **Outline** the monotonic concession protocol and the Zeuthen strategy.

bookwork pp143-146 of Wooldridge [marks 2+3]

(iv) **Discuss** the problems that arise from the use a valuation function that reduces a multi-attribute offer to a single value. Illustrate your answer with an example valuation function and appropriately chosen offers.

bookwork + synthesis

It is hard to differentiate between offers based on multiple criteria.

For example, if the valuation function is f(x,y) = 0.2x+0.8y and there are two offers (1,1) = 1 and (2,0.75) = 1, they cannot be differentiated.
[marks 2]

QUESTION END

QUESTION START

(a) Find the Nash equilibria of the following three games:

(i)

	A	B
A	$1, 0$	$0, 1$
B	$0, 1$	$1, 0$

(ii)

	A	B
A	$6, 3$	$3, 4$
B	$5, 2$	$4, 3$

(iii)

	A	B
A	$5, 3$	$6, 3$
B	$3, 5$	$6, 4$

(i) None
(ii) (B,B)
(iii) (A,A), (A,B)
[Marks: 3]

(b) Remove the strictly dominated strategies from the following strategic game. Show each step. Determine the Nash equilibria.

	A	B	D	E
A	3,7	1,10	2,5	6,6
B	5,4	3,2	9,3	3,1
D	6,5	9,4	8,6	10,0
E	4,6	2,5	6,3	5,1

	A	B	D
A	3,7	1,10	2,5
B	5,4	3,2	9,3
D	6,5	9,4	8,6
E	4,6	2,5	6,3

	A	B	D
B	5,4	3,2	9,3
D	6,5	9,4	8,6
E	4,6	2,5	6,3

	A	B	D
B	5,4	3,2	9,3
D	6,5	9,4	8,6

	A	D
B	5,4	9,3
D	6,5	8,6

Nash equilibria: none [Marks: 4]

(c) Determine the subgame perfect equilibria for the extensive game depicted below:

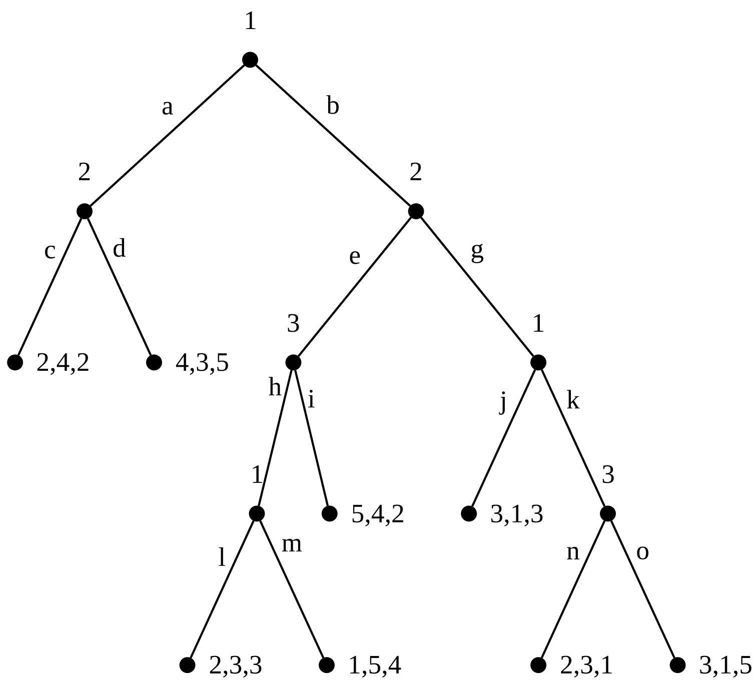

(ajl,ce,ho), (akl,ce,ho), (bjl,ce,ho), (bkl,ce,ho) [Marks: 3]

(d) Give the two model operators if General Normal Modal Logics and explain their meaning

$\Box\varphi$, φ is true in every reachable world and (necessary) $\Diamond\varphi$, φ is true in at least one reachable world (possibly) [Marks: 2]

(e) What is axiom K?

$\models \Box(\varphi \Rightarrow \psi) \Rightarrow (\Box\varphi \Rightarrow \Box psi)$ [Marks: 2]

(f) How can we go from model logic to epistemic logic

Reading $\Box\varphi$ as "it is know that ϕ" [Marks: 2]

(g) Give the definition of an answer set for an logic program.

An interpretation I is an answer set of a program P if it is the deductive closer of P^I, with P^I the positive logic program obtained from P by removing all rules with false negative part of the body and all remaining negations. [Marks: 4]

QUESTION END

QUESTION START

Ontologies are the basis of agent communication while institutional frameworks help simplify interactions and (agent) decision-making.

(a) Develop and justify an ontology for different kinds of learning activities in a university context. Note that the explanation of the concepts you choose to include/exclude is just as important as the ontology. Likewise, the *process* of development is important: your answer should not just be the final ontology, but also the decisions made on the way.

Principles of ontology engineering have been given in the lectures, along with a class discussion of an example ontology and the development of a simple ontology in part of the coursework. The objective of this question is to get the students to demonstrate their understanding of the process by applying in a different but familiar domain.
[marks 10]

(b) Develop and justify, with reference to your ontology an institutional model of the seminar. Identify and describe norms, roles, scenes, and conversations (precise formal specifications are *not* required), using diagrams if you wish, to illustrate your answers.

Other examples of institutions have been given in the lectures, so this part of the question is primarily aimed at getting students to apply knowledge in a new(-ish) situation.

Roles: teacher, learner, collaborator Scenes: presentation, discussion, question, group-working Conversations: a simple protocol situated in one of the above scenes and illustrating a possible interaction Norms: abstract: treat others with respect, as you would wish to be treated yourself; learning is useful concrete: listen while others are speaking; don't put other people down; engage in the topic being taught; ask relevant questions; rule: if you want to ask a question, raise your hand; if your mobile rings then switch it off (or better switch it off before you go in!);

[marks 10]

QUESTION END

QUESTION START

Gaia is an agent-oriented software engineering methodology.

(a) Describe the characteristics of the applications for which the Gaia methodology is supposedly appropriate.

(a) Agents are coarse-grained computational systems, each making use of signicant computational resources (think of each agent as having the resources of a unix process).

(b) It is assumed that the goal is to obtain a system that maximises some global quality measure, but which may be sub-optimal from the point of view of the system components. Gaia is not intended for systems that admit the possibility of true conflict.

(c) Agents are heterogeneous, in that different agents may be implemented using different programming languages, architectures, and techniques. We make no assumptions about the delivery platform.

(d) The organisation structure of the system is static, in that inter-agent relationships do not change at run-time.

(e) The abilities of agents and the services they provide are static, in that they do not change at run-time.

(f) The overall system contains a comparatively small number of different agent types (less than 100).
[marks 6+3]

(b) List the attributes of a role in Gaia

responsibilities, permissions, activities and protocols
[marks 4]

(c) Apply the Gaia methodology to the second piece of coursework. The background information, as supplied in the coursework description, is given below. Not all the information may be relevant to the present question.

4.1 Background Information

The situation is a trading scenario where you must develop client and supplier agents so that a client may purchase a specified quantity of:

- Nuts
- Bolts
- Washers
- Widgets
- Doobries

from a *single* supplier.

Your agent system will consist of a number of supplier agents, and a single client agent implemented as follows. For each run of the scenario:

- Each supplier agent is allocated a supply of a subset of the goods at a particular cost (see the oracle section for information on determining allocation) and a fixed cost for making a whole shipment to the client (i.e. shipment costs are not dependent on quantity).
- The client agent is allocated an order which must be fulfilled (see the oracle for information on allocation).

The client agent and supplier agents should implement the Contract Net Protocol, with the client playing the role of initiator and the suppliers playing the role of responders as follows:

- The client issues a request (call for proposals or CFP) to all of the suppliers indicating their desired bundle.
- The suppliers wait for a CFP and then generate a response to the client, which may be either:
  - a proposal: If the supplier can satisfy the request, the supplier should issue a proposal referring to or containing the original request, and a proposed price. The price should indicate the cost of the goods, and the cost of shipping the goods to the client. *or*
  - a rejection: indicating that the supplier is unable to supply the goods.
- The client waits for responses from the suppliers. If one or more proposal responses is received, then the client must choose a preferred proposal by cost and indicate to the winning supplier that they accept their proposal, and to the losing suppliers that they reject their proposals.
- The winning supplier then acknowledges this to the client.

4.1.1 The Oracle

External information about the supply and demand for goods can be determined by a given supplier or client by interacting with the Oracle agent.

Suppliers receive information about what they can supply by issuing a SupplierInformationRequest message to the oracle, who will then respond with a SupplierInformationResponse message indicating what the supplier can supply.

The client receives information about the order for a given run by issuing a ClientOrderRequest message to the oracle, who responds with a ClientOrderResponse message indicating the required quantities of particular goods desired.

Expect identification of attributes as listed above
[marks 7]

QUESTION END

2 2003-2004

QUESTION START

(i) **State and justify** the four properties of institutions that provide a basis for trust and security in agent interaction.

combination of bookwork and synthesis

(a) decrease uncertainty – because purpose of the institution is to enable a particular kind of interaction, the probability is increased that only agents with co-incident goals will enter the institution.

(b) reduce conflict of meaning – the domain of discourse is largely pre-determined by the shared institutional goals, so that alternative meanings are largely precluded.

(c) create expectations of outcome – the purpose of interacting within an institutional framework is because of a desire to achive a certain goal or set of goals, hence objective is shared

(d) simplify the decision process – the possibility space is pruned through the restricted domain and through the shared goal(s), so there are fewer options to consider.

[marks 4*2]

(ii) **Name and describe** the four classes of norm as they occur in agent-based systems.

bookwork

ABSTRACT NORMS: vague desires, possibly contradictory, probably unenforceable

CONCRETE NORMS: abstract norms grounded in ontology, open to reasoning

RULES: concrete specifications of (un)acceptable actions, and obligations, declarative style

PROCEDURES: concrete specifications of sequences of actions, options, outcomes

[marks 4*1]

(iii) Consider the institutional framework of the lecture:

(a) **Identify** some roles, scenes, and conversations (precise formal specifications are *not* required), but you may use diagrams to illustrate your answers.

Although this situation was used as illustration in a lecture, this part of the question is primarily aimed at getting students to apply knowledge in a new (ish) situation.
Roles: teacher, learner, collaborator Scenes: presentation, discussion, question, group-working Conversations: a simple protocol situated in one of the above scenes and illustrating a possible interaction
[marks 4]

(b) **Propose** an abstract norm, a concrete norm and a rule that might govern behaviour in a lecture.

abstract: treat others with respect, as you would wish to be treated yourself; learning is useful concrete: listen while others are speaking; don't put other people down; engage in the topic being taught; ask relevant questions; rule: if you want to ask a question, raise your hand; if your mobile rings then switch it off;
[marks 4]

QUESTION END

3 2002 - 2003

QUESTION START

(a) Describe Brooks subsumption architecture.

The Brooks subsumption architecture is arguably the best-known reactive agent architecture. It was developed by Rodney Brooks - one of the most vocal and influential critics of the symbolic approach to agency to have emerged in the recent years. Brooks propounded three key theses that have guided his work as follows:

- Intelligent behaviour can be generated without explicit representation of the kind symbolic AI proposes
- Intelligent behaviour can be generated without explicit abstract reasoning of the kind symbolic AI proposes.
- Intelligence is an emergent property of certain complex systems.

Brooks also identifies two key ideas that have informed research:

- Situatedness and embodiment. Real intelligence is situated in the world, not is disembodied systems such as theorem provers or expert systems.
- intelligence and emergence. Intelligent behaviour arises as a result of an agent's interaction with its environment. Also intelligence is in the eye of the beholder - it is not an innate, isolated property.

These ideas were made concrete in the subsumption architecture. There are two defining characteristics of the subsumption architecture. The first is that an agent's decision-making is realised through a set of task-accomplishing behaviours; each behaviour may be thought of as an individual action function, as we defined above, which continually takes perceptual input and maps it to some particular task. In Brook's implementation, the behaviour modules are finite-state machines. An important point to note is that these task-accomplishing modules are assumed to include no complex symbolic representations, and are assumed to do no symbolic reasoning at all. In many implementations, these behaviours are implemented as rules of the form:

$$\text{situation} \rightarrow \text{action},$$

which simply maps perceptual input directly to action. The second characteristic of the subsumption architecture is that many behaviours can fire simultaneously. There must obviously be a mechanism to choose between the different actions selected by these multiple actions. Brooks proposed to arranging the modules into a subsumption hierarchy, with the behaviours arranged into layers. Lower layers in the hierarchy are able to inhibit higher layers: the lower the layer is, the higher is its priority. The idea is that higher layers represent more abstract behaviours. [Marks: 7]

(b) Describe the Blocks World problem. Give an example of the predicates, rules, goal and a corresponding plan.

The Blocks World contains three block (A,B,C) of equal size, a robot arm capable of picking up and moving one block at a time, and a table top. The blocks may be placed on the table top, or one may be placed on top of the other. The general idea is to place to blocks in a given configuration. This configuration will be the goal. The agent has to come up with a plan in order to, starting from an original configuration, obtain the desired goal. A description of the blocks world with an initial configuration could be:

{ Clear(A), On(A,B), OnTable(B), OnTable(C), ArmEmpty }

A goal could look like:

{ OnTable(A),OnTable(B),OnTable(C)}

In other to achieve this we need the following rules:

Stack(x,y)
pre { Clear(y),Holding(x) }
del { Clear(y),Holding(x) }
add { ArmEmpty,On(x,y)}

UnStack(x,y)
pre { On(x,y), Clear(x), ArmEmpty }
del { On(x,y), ArmEmpty }
add { Holding(x),Clear(y)}

Pickup(x)
pre { Clear(x),OnTable(x),ArmEmpty }
del { OnTable(x),ArmEmpty(x) }
add { Holding(x)}

PutDown(x)
pre { Holding(x) }
del { Holding(x) }
add { ArmEmpty,OnTable(x)}

A corresponding plan would be:

UnStack(A,B),PutDown(A)

[Marks: 6]

(c) Give a theoretical description of a planning problem.

Assume a fixed set of actions $AC = \{\alpha_1, \ldots, \alpha_n\}$ that the agent can perform. A descriptor of an action $\alpha \in AC$ is a triple:

$$\langle P_\alpha, D_\alpha, A_\alpha \rangle$$

where:

- P_α is a set of formulae of first-order logic that characterise the precondition of action α;
- D_α is a set of formulae of first-order logic that characterise those facts made false by the performance of α
- A_α is a set of first-order formulae logic that characterise those facts made true by the performance of α.

A planning problem is then determined by a triple:

$$\langle \Delta, O, \gamma \rangle,$$

where

- Δ is the beliefs of the agent about the initial state of the world, presented by a set of first-order formulae

- $O = \{\langle P_\alpha, D_\alpha, A_\alpha \rangle \mid \alpha \in AC\}$ is an indexed set of operator descriptors, one for each available action α; and
- γ is a set of formulae of first-order logic, representing the goal/task/intention to be achieved.

A plan π is a sequence of actions:

$$\pi = (\alpha_1, \ldots \alpha_n),$$

where each α_i is a member of AC. With respect to a planning problem $\langle \Delta, O, \gamma \rangle$, a plan $\pi = (\alpha_1, \ldots \alpha_n)$ is said to be acceptable with respect to the problem $\langle \Delta, O, \gamma \rangle$ if, and only if, the precondition of every action is satisfied in the preceding environment, i.e if $\Delta_{i-1} \models P_{\alpha_i}$, for all $1 \leq i \leq n$. A plan $\pi = (\alpha_1, \ldots \alpha_n)$ is correct with respect to $\langle \Delta, O, \gamma \rangle$ if and only if

(a) it is acceptable; and

(b) $\Delta_n \models \gamma$.

The problem to be solved by a planning system can then be stated as follows:

> Given a planning problem $\langle \Delta, O, \gamma \rangle$, find a correct plan for $\langle \Delta, O, \gamma \rangle$ or else announce that none exists.

[Marks: 7]

QUESTION END

QUESTION START

(a) Speech acts typically have two components: name them and then list *five* different kinds of speech acts along with illustrative examples.

> speech act = performative + proposition kinds are: representative - inform the recipient directive - get recipient to do something commisive - commit speaker to do something expressive - speaker expresses mental state declarative - speaker announces something [Marks: 7]

(b) Negotiation strategy can be characterized by the initial offer an agent makes and the way in which an agent changes its reaction to offers over time. Describe in detail the key aspects of the *monotonic concession* and the *Zeuthen* strategies.

> conventional bookwork pp143-146 of Wooldridge [Marks: 3*2]

(c) Contract net is a very structured form of negotiation used to achieve task sharing. Outline the five phases of the contract net protocol and what problems may arise in practice.

> five phases are bookwork (pp194-200). Problems provide some room for initiative, such task specificaiton, QoS, choosing between competing offers, differentiating between offers baseed on multiple criteria. [Marks: 5+2]

QUESTION END

QUESTION START

(a) Describe, by means of an example game, the differences in Nash equilibria one may obtain when a game is played only once or for a number of rounds. Explain why such a difference occurs. Is there any difference when the iteration has a finite or infinite horizon.

> Consider the Prisoner's Dilemma: two men are arrested for a crime. The have been put is separate cell for questioning. Each of them is asked to betray to other. Assume that the maximum sentence is 4 years. If they both betray each other they are sentenced to 2 years. If the both remain silent they will convicted for only 1 year. If one betrays the other while the other remains silent, the betraying party will be released while the other will serve the maximum sentence. This can be represented by the following payoff matrix:

	Betray	*Silent*
Betray	$2, 2$	$4, 0$
Silent	$0, 4$	$3, 3$

> The payoff represents the year they gain with respect to the maximum sentence. If this game is only played one, the rational outcome or Nash equilibrium will be that both players will betray each other. Neither party is willing to risk the sucker's payoff by remaining silent. When the game is played several times over with the same players the situation can be different. If the players do not know how many games will play (infinite horizon) the Nash equilibrium action for both players will be to remain silent. The reason for this is that they otherwise risk being punished in the next round. In case the players know when the game will end, they will betray in each round. In the last round, round n, they know they can defeat in order the highest payoff without having the risk of punishment. Since both players reason like this, round n -1 will be

the actual ending of the iteration. So we can apply the same reasoning again. This backward induction strategy causes a whole series of betrays. This means that in case of a finite horizon we obtain (Betray, Betray) asa Nash equilirbium just as it was in the single round game. [Marks: 5]

(b) Find the Nash equilibria of the following games:

(i)

	Lake	*Mountains*
Lake	2, 1	0, 0
Mountains	0, 0	1, 2

(ii)

	Head	*Tail*
Head	1, −1	−1, 1
Tail	−1, 1	1, −1

(iii)

	L	*R*
T	3, 0	0, 1
M	0, 0	3, 1
B	1, 1	1, 0

(iv)

	b_1	b_2	b_3	b_4
a_1	0, 7	2, 5	7, 0	0, 1
a_2	5, 2	3, 3	5, 2	0, 1
a_1	7, 0	2, 5	0, 7	0, 1
a_2	0, 0	0, −2	0, 0	10, −1

(i) (Lake,Lake) and (Mountains,Mountains)
(ii) No Nash equilibria
(iii) (R,M)
(iv) (B_2, a_2)
[Marks: 8]

(c) Describe the possible properties of negotiation protocols

- Guaranteed success. A protocol guaranteed success if it ensures that, eventually, agreement is certain to be reached.
- Maximising social welfare. Intuitively, a protocol maximises social welfare if it ensures that any outcome maximises the sum of utilities of negotiations participants. If the utility of an outcome for an agent was simply defined in terms of the amount of money that agents received in the outcome, then the protocol that maximises social welfare would maximise the total amount of money paid out.
- Pareto efficiency. A negotiation outcome is said to be Pareto efficient if there is no other outcome that will make at least one agent better off without making at least one other agent worse off. intuitively, if a negotiation outcome is not Pareto efficient, then there is another outcome that will at least one agent happier while keeping everyone at least as happy.

- Individual rationality. A protocol is said to be individually rational if following the protocol is in the best interest of negotiation participants. Individually rational protocols are essential because without them, there is no incentive for agents to engage in negotiations.
- Stability. A protocol is stable if it provides all agents with an incentive to behave in a particular way. The best-known kind of stability is Nash equilibrium.
- Simplicity. A simple protocol is one that makes the appropriate strategy for a negotiation participant obvious. That is, a protocol is simple if using it, a participant can easily (tractably) determine the optimal strategy
- Distribution. A protocol should ideally be designed to ensure that there is no single point of failure (such as a single arbitrator) and, ideally, so as to minimise communication between agents.

[Marks: 7]

QUESTION END

QUESTION START

A grass-cutting robot can perceive the presence of *grass* or *null* (meaning just something else) and has the following repertoire of actions:

- forward
- cut
- turn

The goal of this robot is to wander around the garden, which comprises grassed areas and paths, cutting grass whereever it finds it. To specify the agent behaviour, you may use the following three domain predicates:

- *In(x,y)* robot is at *(x,y)*
- *Grass(x,y)* there is grass at *(x,y)*

- *Facing(d)* robot is facing in direction *d*

Now answer the following questions:

(i) Draw a diagram of the garden in which the robot is going to operate and then Write down and *explain* the rules you would use (a) to control cutting (b) to control movement You may assume the garden is rectangular with the farthest corner at *(maxx, maxy)*.

Note: this description is a domain change on an example given in the textbook (pp49-54).

Thus the rules are pretty much as in textbook:

```
In(x,y) & Grass(x,y) -> Do(cut)
and
In(0,0) & Facing(north) & not(Grass(0,0)) -> Do(forward)
In(0,0) & Facing(south) & not(Grass(0,0)) -> Do(turn)
In(0,0) & Facing(east) & not(Grass(0,0)) -> Do(forward)
In(0,0) & Facing(west) & not(Grass(0,0)) -> Do(turn)
```

and similarly but differently for other corners and

```
In(0,x) & Facing(west) & not(Grass(0,x)) -> Do(turn)
In(0,x) & Facing(d) & not(Grass(0,x)) -> Do(forward)
```

and similarly but differently for the other boundaries [marks 10]

(ii) Discuss how you would extend the rule set, percepts, actions and/or predicates to prevent the robot from moving on to flower borders.

Looking for general evidence of thinking. Hope no one proposes hard-coding the border edges in the same way as the limits of the garden. Simplest solution is probably to arithmetic in the Grass predicate:

```
In(x,y) & Facing(north) & not(Grass(x,y+1)) -> Do(turn)
```

and similarly but differently for facing other directions. Other solutions might involve introducing a new percept earth and perhaps a new predicate Look. Good student might also remark upon the problem of ordering the rules to get the right results and how this gets increasingly difficult as the size of the ruleset increases

[marks 5]

(iii) Do you feel this is an appropriate architecture for the task or would a subsumption architecture be more suitable? Consider how Steel's Mars explorer model might be adapted for this environment. Statement of rules is not required.

More evidence of thinking required. Could posit the existence of a radio transmitter in the centre of the garden (doesn't have to be exactly in the centre tho, or really even anywhere near the centre!), then robot can use the gradient effect for navigation (and depositing the grass clippings, so best put the transmitter in the compost heap). The radio-active crumbs should be used now not to reinforce a particular path, but instead to make a robot go in a different direction.

[marks 5]

QUESTION END