CS 188 Spring 1996 — Introduction to AI — Stuart Russell — Final examination

You have 2 hours 50 min. The exam is open-book, open-notes.
There a total of 100 points available.
Write your answers in blue books. Hand them all in.
Several of the questions on this exam are true/false or multiple choice.
In all the multiple choice questions more than one of the choices may be correct. Give all correct answers. Each multiple choice question will be graded as if it consisted of a set of true/false questions, one for each possible answer.

1. (12 pts.) True/False
Decide if each of the following is true or false. If you are not sure you may wish to provide some explanation to follow your answer.

(a) (2) A feedforward neural network connecting sensors to effectors can implement a reflex agent with state.

(b) (2) No internal state is *required* for an agent to be successful in a perfectly accessible environment.

(c) (2) In a finite state space containing no goal state, A* will always explore all states.

(d) (2) $\forall x\ (x = x) \lor (x > x) \lor (x < x)$ is a valid sentence.

(e) (2) Any MDP model can be translated directly into situation calculus.

(f) (2) Since the value of information is nonnegative, the outcome of acting on more information will always be at least as good as the outcome of acting on less information.

2. (12 pts.) Game-playing

(a) (2) For a game with branching factor b and search depth d, what are the space requirements of alpha-beta search? Explain (briefly).

(b) (4) Roughly speaking, Go is played by placing stones anywhere on a 19×19 board until the board is full, and takes up to six hours. Make a very rough estimate of how deep alpha-beta can search in Go, stating all your assumptions.

(c) (2) Given a choice between moves A and B in the following tree, which would minimax choose (assuming MAX is to move at the root)?

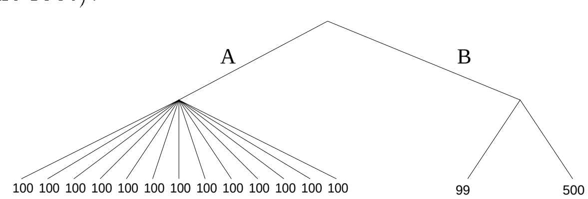

(d) (4) It has often been argued that minimax backup is designed on the assumption that leaf node evaluations are correct. Clearly, leaf node evaluations (except for terminal states) are subject to significant errors—that is, states with high values may be bad, and states with low values may be good, etc. Given the above choice, where all the leaves are nonterminals, which would *you* choose? Why?

3. (15 pts.) Logic *True/false*:

(a) (2) Resolution can always be used to decide if two first-order sentences are consistent.

(b) (2) The results of exhaustive backward-chaining (to find all solutions) are independent of the search order.

(c) (2) A universally quantified sentence can be proved using a sufficient number of examples.

(d) (2) $>(sqrt(x), sqrt(z))$ unifies with $>(y, sqrt(y))$.

(e) (2) "Newt makes more royalties than anyone" is a good translation of
$\exists r\ Royalty(r) \wedge Makes(Newt, r) \wedge \forall x, s\ \ \neg(x = Newt) \Rightarrow Royalty(s) \wedge Makes(x, s) \wedge r > s$

(f) (2) $\forall a, x, s\ Sane(x, s) \wedge \neg(a = TakeCS162) \Rightarrow Sane(x, Result(a, s))$
is a correctly formed effect axiom in situation calculus.

(g) (3) The STRIPS operator

$$Op(\textsc{Action}{:}Go(there), \textsc{Precond}{:}At(here) \wedge Path(here, there),$$
$$\textsc{Effect}{:}At(there) \wedge \neg At(here))$$

is *entirely* equivalent to the successor-state axiom

$$\forall a, p, q, s\ At(q, Result(a, s)) \Leftrightarrow$$
$$[\ \ (a = Go(q) \wedge At(p, s) \wedge Path(p, q))$$
$$\vee\ (At(q, s) \wedge \neg\exists r\ a = Go(r))\,]$$

4. (11 pts.) Probabilistic inference

Consider the belief network shown in the following figure:

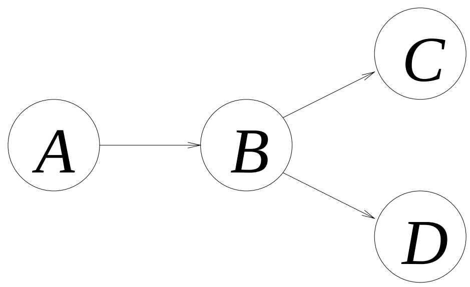

(a) (3) Which of the following statements are *implied* by the network structure?

i. $P(B|A, C, D) = P(B|A)$

ii. $P(D|B) = P(D|B, C)$

iii. $P(B|A) \neq P(B)$

(b) (3) Suppose we have evidence at B. Use Bayes' rule to solve the query $P(A|B)$ in terms of probabilities directly available in the network. (You may assume the usual normalization constant α.)

(c) (5) Suppose instead that we have evidence at C and D. Use Bayes' rule followed by *conditioning* on B to solve the query $P(A|C, D)$ in terms of probabilities directly available in the network. [Given an expression $P(X|Z)$, conditioning on Y gives $P(X|Z) = \sum_y P(X|Y = y, Z)P(Y = y|Z)$.]

5. (10 pts.) Markov decision problems

In this question we will look at how to describe MDPs in the language of probability theory. Consider an agent moving in a 4×4 rectangular grid with locations defined as [1,1] through [4,4]. Let S_t be the actual location at time step t. Let E_t be the location that the agent perceives it is in at time t. Let A_t be the action taken at time t.

(a) (1) State precisely the condition required for this to be an MDP (rather than a POMDP).

(b) (3) Let us assume that the agent has four actions (up, down, left, right), and that they are deterministic. They can be described by a collection of probability statements describing the probability of specific values of S_{t+1} given specific values of S_t and A_t. Write down one such statement for a specific location pair and action. How many such statements are needed in all?

(c) (3) Suppose the location is defined by the variables X_t and Y_t, each of which takes on values 1 through 4. Show by writing down a probability statement for the probability of a specific value X_{t+1} given X_t and A_t that this decomposes the problem of describing the MDP. How many such statements do we need now?

(d) (3) Draw a belief net structure relating the variables X_t, Y_t, A_t, E_t, X_{t+1}, Y_{t+1}, E_{t+1}

6. (12 pts.) Inductive learning

(a) (3) Explain why no representation scheme for Boolean functions can provide a compact representation for all possible functions.

(b) (3) Draw a decision tree to represent the "two or more" function for three inputs.

(c) (3) Let us consider neural nets with inputs in the range $[0, 1]$ and with g a step function. A network is defined by the weights on the links and the threshold value of g at each node. Draw a network to represent the "two or more" function for three inputs.

(d) (3) The standard learning algorithm for NNs uses hillclimbing in weight space. Explain what practical problem would make it difficult to use best-first search instead.

7. (13 pts.) Natural language

Consider the following context-free grammar:

S $\rightarrow$Question ?	Pronoun $\rightarrow$I |you|he| him
Question $\rightarrow$Aux NP VP	QPronoun $\rightarrow$who | what
Question $\rightarrow$QNP VP	Determiner $\rightarrow$a | the | some
Question $\rightarrow$QNP Aux NP Verb	QDeterminer$\rightarrow$which | what
NP $\rightarrow$Determiner NC | Pronoun	Adjective $\rightarrow$big
NC $\rightarrow$Adjective* Noun	Noun $\rightarrow$dog | cat | table
NC $\rightarrow$Adjective* Noun RelClause	Verb $\rightarrow$see | eat | is
QNP $\rightarrow$QDeterminer NC| QPronoun	Aux $\rightarrow$did | will | can
RelClause $\rightarrow$that VP	
RelClause $\rightarrow$that NP Verb	
VP $\rightarrow$Verb NP	

(a) (4) *Multiple choice*: Which of the following sentences are generated by the grammar?

i. Did you see the dog eat the cat ?

ii. Can the table is me ?

iii. Which big big table that the cat see will you eat ?

iv. What cat eat the table that I see ?

(b) (2) Write down at least one other *English* sentence generated by the grammar above. It should be significantly different from the above sentences, and should be at least five words long. Do not use any of the open-class words from the above sentences; instead, add grammatical rules of your own, of the form (grammatical category) $\rightarrow$(specific word)—for instance, Noun $\rightarrow$ bottle.

(c) (2) Show the parse tree for your sentence.

(d) (3) The grammar as given does not accept sentences such as "Is the cat a dog?"? Suggest a (sensible) modification so that it will.

(e) (2) Suppose that the rule for NP is replaced by
NP $\rightarrow$Determiner Adjective* Noun RelClause
Explain why the resulting grammar contains no sentences.

8. (5 pts.) Perception

Briefly list the reasons why it is hard to wreck a nice beach (recognize speech).

Exam continues overleaf ...

9. (10 pts.) Robotics

Consider the very simple robot arm shown in the following figure, where the shaded regions represent obstacles and the arm rotates around the fixed pivot:

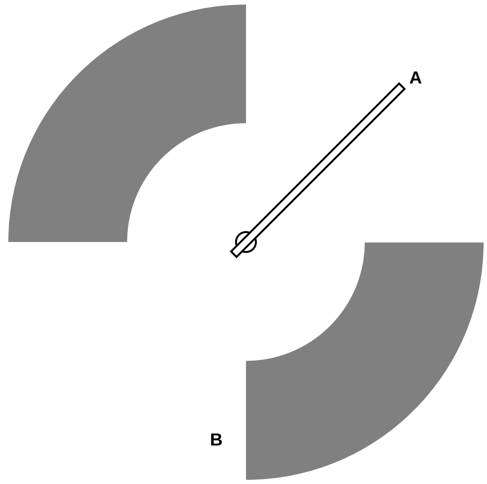

(a) (1) How many degrees of freedom does the robot have?

(b) (5) Draw the configuration space of the robot, labelling each axis with the corresponding degree of freedom and showing the configuration space obstacles.

(c) (4) Suppose that the robot arm can telescope, i.e., shorten or elongate itself to an arbitrary degree. Draw the new configuration space, and show how the robot gets its end-effector from A to B.