CS188 Spring 1993 — Intro to AI — Stuart Russell — Final examination

You have 2 hours 50 min. The exam is open-book, open-notes.
There a total of 110 points available. Only 100 points are required for full credit.
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.) Definitions Provide brief, *precise* definitions of the following:

(a) Turing test
(b) Unifier
(c) Coercible world
(d) Partial-order planner
(e) Singly-connected network
(f) Perceptron

2. (20 pts.) Logical Inference
From "Horses are animals", it follows that "The head of a horse is the head of an animal". Demonstrate that this inference is valid by carrying out the following steps:

(a) (5) Translate the premise and the conclusion into the language of first-order logic. Use three predicates:
- `Head-of`(x, y): x is the head of y.
- `Horse`(x): x is a horse.
- `Animal`(x): x is an animal.

(b) (6) Negate the conclusion, and convert the premise and the negated conclusion into conjunctive normal form.

(c) (7) Use resolution to show that the conclusion follows from the premise.

3. (17 pts.) Search in games

(a) (3) The following code implements minimax search in a two-player game:

```
(defun choose (side state limit)
  (the-biggest #'(lambda (s) (backed-up-value (opponent side) state 1 limit))
               (successors state)))

(defun backed-up-value (side state depth limit)
  (if (= depth limit)
      (evaluate side state)
      (apply (if (oddp depth) #'min #'max)
             (mapcar #'(lambda (s) (backed-up-value (opponent side) s (1+ depth) limit))
                     (successors state)))))
```

Is the search breadth-first or depth-first?

For parts b) through f), use an evaluation function defined as
$f =$ Number of rows, columns or diagonals containing only X's
minus number of rows, columns or diagonals containing only O's.

(b) (4) Show the game tree for tictactoe starting from an empty board down to depth 2, (i.e. one X and one O on the board), taking symmetry into account. You should have 3 positions at level 1 and 12 at level 2.

(c) (3) Mark on your tree the evaluations of all the positions at level 2.

(d) (2) Mark on your tree the backed-up values for the positions at levels 1 and 0, using the minimax algorithm, and use them to choose the best starting move.

(e) (2) Circle the nodes at level 2 that would *not* be evaluated if alpha-beta pruning were applied, with the nodes generated left-to-right in the order you have shown.

(f) (3) (Open ended, don't write for more than 5 minutes unless you have extra time. No BS!) Suppose we play a variant of tictactoe in which each player sees only their own moves. If they make a move on a square occupied by an opponent, the board "beeps" and they get another try. Would the backgammon model suffice for this game, or would we need something more sophisticated? Why?

4. (16 pts.) Decision theory and the value of information

In this question, let A_i be an action, let K be the agent's current information, let U be the agent's utility function on world states W_j. Also, let $argmax_x(f(x))$ denote the value of x that gives the highest value of $f(x)$, and let $max_x(f(x))$ be the highest value of $f(x)$ over all possible values of x.

(a) (3) Give a mathematical expression that evaluates to the action that a rational agent should take.

(b) (4) *Multiple choice*: In many situations, as well as simply choosing an action one also has the option to obtain additional information E, although usually at some cost. Let the possible values of E be denoted by e_k. The value of obtaining evidence E is defined as the difference between the expected value of acting rationally with knowledge of E and the expected value of acting rationally without knowledge of E. Which of the following equations express this, given that the value of the information must be calculated *before* the information is known?

i.

$$V(E) = \sum_k P(E = e_k | A_i, K) max_{A_i} \sum_j U(W_j) P(W_j | A_i, K, E = e_k) - max_{A_i} \sum_j U(W_j) P(W_j | K)$$

ii.

$$V(E) = \sum_k P(E = e_k | K) max_{A_i} \sum_j U(W_j) P(W_j | A_i, K, E = e_k) - max_{A_i} \sum_j U(W_j) P(W_j | A_i, K)$$

iii.

$$V(E) = max_{A_i} \sum_k P(E = e_k | K) \sum_j U(W_j) P(W_j | A_i, K, E = e_k) - max_{A_i} \sum_j U(W_j) P(W_j | A_i, K)$$

(c) (6) *Multiple choice*: Suppose you are asked to choose one of k identical doors, behind only one of which is a prize of C dollars. What is the most you should be willing to pay to sneak a peek behind the first door before having to choose your door?

i. C/k

ii. $C/(k+1)$

iii. $(k-1)C/k^2$

(d) (3) *True/false*: If you are allowed to buy a peek behind doors 1 and 2, it is just as good to buy two peeks at the start as to buy one, then decide whether or not to buy the second after seeing the first.

5. (12 pts.) Learning in agents

(a) (6) Consider the following data, each consisting of three input bits and a classification bit: (111,1) (110,1) (011,1) (010,0) (000,0).

i. (3) Draw a decision tree consistent with this data.

ii. (3) Draw a neural network (using a threshold activation function) consistent with this data.

(b) (6) Consider a learning program LP as taking a set of classified examples as input, and returning a function that is supposed to calculate the appropriate classification given an unclassified example. Show how to use LP to construct a learning agent LA. The agent should learn from percepts that include the correct action to take, as well as doing the action. When a percept arrives that does not include the correct action, it should respond with the action that its past experience indicates might be appropriate. Write LA, either in pseudocode or LISP, being as precise as possible. Don't worry about efficiency. If you're not sure that your code is completely transparent, you might want to include a little comment.

6. (15 pts.) Natural language

(a) (6) Consider the following sentence:

Someone walked slowly to the Safeway.

and the following set of context-free rewrite rules which give the grammatical categories of the individual words of the sentence:

Pronoun →"someone"
V →"walked"
Adv →"slowly"
Prep →"to"
Det →"the"
Noun →"Safeway"

Which of the following three sets of rewrite rules, when added to the above rules, yield context-free grammars which can generate the above sentence?

(A):	(B):	(C):
S →NP VP	S →NP VP	S →NP VP
NP →Pronoun	NP →Pronoun	NP →Pronoun
NP →Det Noun	NP →Noun	NP →Det NP
VP →VP PP	NP →Det NP	VP →V Adv
VP →VP Adv Adv	VP →V Vmod	Adv →Adv Adv
VP →V	Vmod →Adv Vmod	Adv →PP
PP →Prep NP	Vmod →Adv	PP →Prep NP
NP →Noun	Adv →PP	NP →Noun
	PP →Prep NP	

(b) (2) Write down at least one other English sentence generated by Grammar (B) above. It should be significantly different from the above sentence, and should be at least six words long. Do not use any of the words from the above sentence; instead, add grammatical rules of your own, of the form (grammatical category) →(specific word)—for instance, Noun → "bottle".

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

(d) (5) Which of the following:

(A) "Only gentlemen of a certain stripe drink red wine with fish."
(B) "The astronomy professor's presentation exhibited stellar brilliance."
(C) "This classroom is on a major fault line."

exhibit

i. Context-sensitive reference?
ii. Noun-noun modification?
iii. Noncompositional meaning?
iv. Word-level ambiguity?
v. Idiomatic phrases (common constructions that cannot be understood by analysis into parts)?

7. (18 pts.) Perception and robotics using belief networks

In this question we will consider part of the problem of building a robot that plays ping pong.

One of the things it will have to do is find out where the ball is and estimate its trajectory. Let's suppose for a moment that we have a vision system that can return an estimated instantaneous (x,y,z) position for the ball. This position estimate for time t_i will be used as an evidence node O_i in a belief network that infers the desired information. The idea is that nodes X_i and V_i in the network represent the ball's instantaneous position, and velocity. The trajectory of the ball is followed by having multiple copies of these nodes, one for each time step. In addition to these nodes, we also have nodes for the instantaneous friction force F_i acting on the ball due to air resistance, which depends only on the current velocity. The times t_1, t_2, t_3 ... can be assumed to be separated by some small interval δt.

(a) (4) *Multiple choice*: Which of the following belief networks correctly (but not necessarily efficiently) represent the above information? (You may assume that second-order effects — proportional to δt^2 — are ignored.)

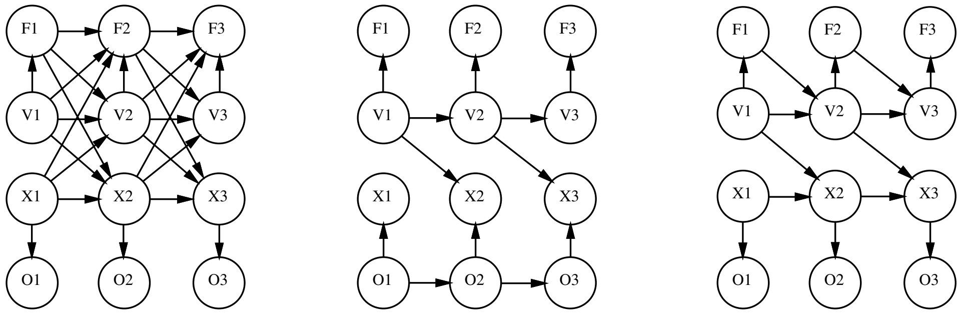

i) ii) iii)

(b) (1) *Multiple choice*: Which is the best network?

(c) (1) *Multiple choice*: Which networks can be solved by local propagation algorithms?

(d) (4) The nodes in the networks shown above refer to past and present times. Explain how the network can be extended and used to predict the position at any future time.

(e) (6) Now consider the vision subsystem, whose job it is to provide the observation evidence. Its input is an array of intensity values from the camera image. Assuming that the ball is white, while most of the room is darker, explain briefly how to calculate the position of the ball given the array.

(f) (2) Does it make more sense to calculate the position in coordinates relative to the robot or relative to the room? Why?