CS 188 Introduction to AI

Fall 1999 Stuart Russell Final Examination

You have 2 hours 50 min. The exam is open-book, open-notes. There are 100 points available.

Write your answers in blue books. Hand them all in.

1. (10 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) The human brain can be described, to a first approximation, as a very large multilayer feedforward neural network.

(b) (2) Depth-first iterative deepening always returns the same solution as breadth-first search if *b* is finite and the successor ordering is fixed.

(c) (2) Any deterministic MDP with a fixed start state can be solved by converting into an equivalent search problem using A*.

(d) (2) Any decision tree with Boolean attributes can be converted into an equivalent feedforward neural network.

(e) (2) The clause A OR B OR C entails the clause A OR B.

2. (18 pts.) Logic *True/False*:

(a) (3) For a propositional clause to be valid, it must contain literals P and P, for some proposition symbol P.

(b) (3) Q(A, F(A)) is a possible resolvent of the clauses Q(x, F(x)) OR Q(A, x) and Q(w, A).

(c) (4) If C_1 and C_2 are clauses in first-order logic, and all the literals in $C_1 o$ are contained in C_2, for some *o*, then C_1 |= C_2.

(d) (4) Some dogs hate all cats who eat birds is a good translation of :

there exists *d* *Dog*(*d*) AND [for all *c* *Cat*(*c*) AND (there exists *b* *Bird*(*b*) => *Eats*(*c*, *b*)) => *Hates*(*d*, *c*)]

(e) (4) The following search formulation yields a complete inference procedure for propositional logic KBs:

initial state = CNF(KB) plus CNF(*a*), where *a* is the query

successor function = return a single KB augmented with all *new* resolvents of all pairs of clauses, if any, else no successor

goal test = KB contains the empty clause

search algorithm = depth-first search

3. (14 pts.) Planning and MDPs

Consider a register of 3 bits, labeled 1 through 3. A bit can be on or off. A bit can be turned on if it is off, and can be turned off if it is on. In the initial state, all the bits are off. In a goal state, bits 2 and 3 are on (bit 1 can be on or off).

(a) (2) Write the *TurnOn* and *TurnOff* actions as STRIPS operators.

(b) (2) Draw the initial partially ordered planning problem using the pictorial notation from the book.

(c) (5) Explain what happens when the POP algorithm is applied to this problem. (You may use a diagram to show the final plan but it is not required.)

(d) (4) Explain precisely how you would set this problem up as an MDP.

(e) (1) Describe an optimal policy for the MDP.

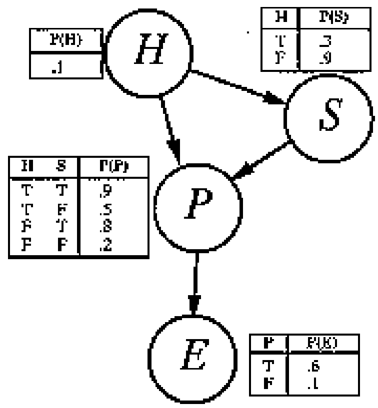

Fig. 1: A simple belief network with Boolean variables H = Honest, S = Slick, P = Popular, E = Elected.

4. (16 pts.) Probabilistic Inference

Consider the belief network shown in Fig. 1.

(a) (3) Which, if any, of the following are asserted by the network *structure* (ignoring the CPTs for now)?

[Note: any subset of these may be correct]

i. **P**(H, S) = **P**(H)**P**(S)

ii. **P**(E | P,H) = **P**(E | P)

iii. **P**(E) != **P**(E | H)

(b) (3) Calculate the value of P(h, s, p, e).

(c) (4) Calculate the probability that someone is elected given that they are honest.

(d) (6) Suppose we want to add the variable L = *LotsOfCampaignFunds* to the network; describe, with justifications, all the changes you would make to the network.

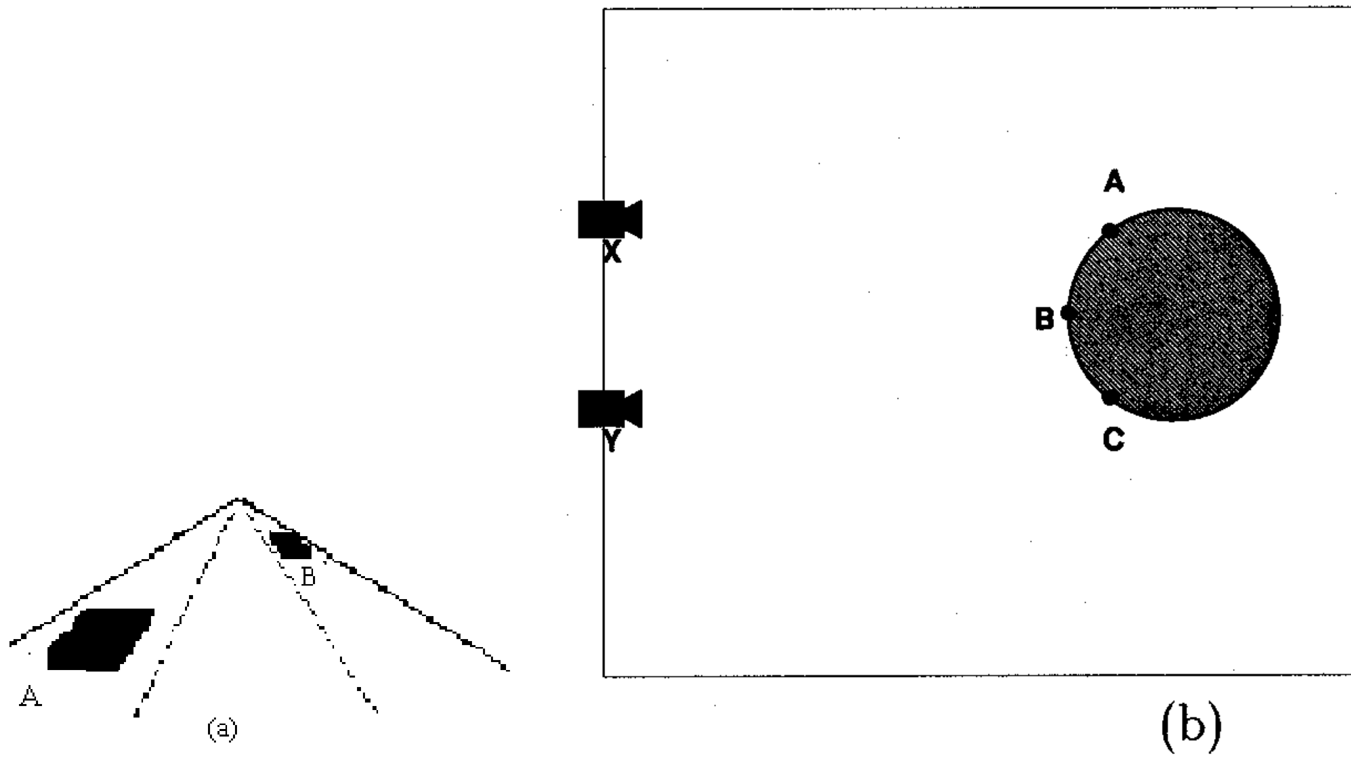

Fig. 2: (a) View from the drivers seat on a freeway. (b) Top view of a two-camera vision system observing a bottle with a wall behind.

5. (10 pts.) Vision

(a) (4) In Fig. 2(a) above, showing two cars on a flat road, what reasons does the viewer have to conclude that car A is closer than car B?

(b) (6) In Fig. 2(b) above, showing two cameras at X and Y observing a scene, what can the viewer conclude about the relative distances of points A, B, C, D, and E from the camera baseline, and on what basis?