CS 188 Fall 1999 — Introduction to AI — 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 and using A*.
   (d) (2) Any decision tree with Boolean attributes can be converted into an equivalent feedforward neural network.
   (e) (2) The clause $\neg A \vee \neg B \vee C$ entails the clause $\neg A \vee \neg B$.

2. **(18 pts.) Logic** *True/false*:
   (a) (3) For a propositional logic clause to be valid, it must contain literals P and $\neg P$, for some proposition symbol P.
   (b) (3) $\neg Q(A, F(A))$ is a possible resolvent of the clauses $\neg Q(x, F(x)) \vee Q(A, x)$ and $Q(w, A)$.
   (c) (4) If C_1 and C_2 are clauses in first-order logic, and all the literals in $C_1\sigma$ are contained in C_2, for some σ, then $C_1 \models C_2$.
   (d) (4) "Some dogs hate all cats who eat birds" is a good translation of
   $\exists d\ Dog(d) \wedge [\forall c\ Cat(c) \wedge (\exists b\ Bird(b) \Rightarrow Eats(c, b)) \Rightarrow Hates(d, c)]$
   (e) (4) The following search formulation yields a complete inference procedure for propositional logic KBs:
   initial state = CNF(KB) plus CNF($\neg\alpha$), where α 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, labelled 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.

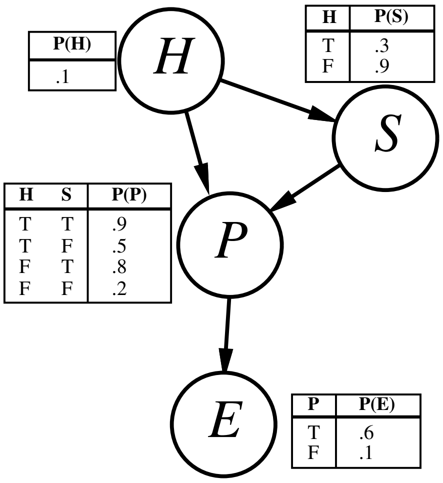

Fig. 1: A simple belief network with Boolean variables $H = Honest$, $S = Slick$, $P = Popular$, $E = Elected$.

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. $\mathbf{P}(H, S) = \mathbf{P}(H)\mathbf{P}(S)$

ii. $\mathbf{P}(E|P, H) = \mathbf{P}(E|P)$

iii. $\mathbf{P}(E) \neq \mathbf{P}(E|H)$

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

(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.

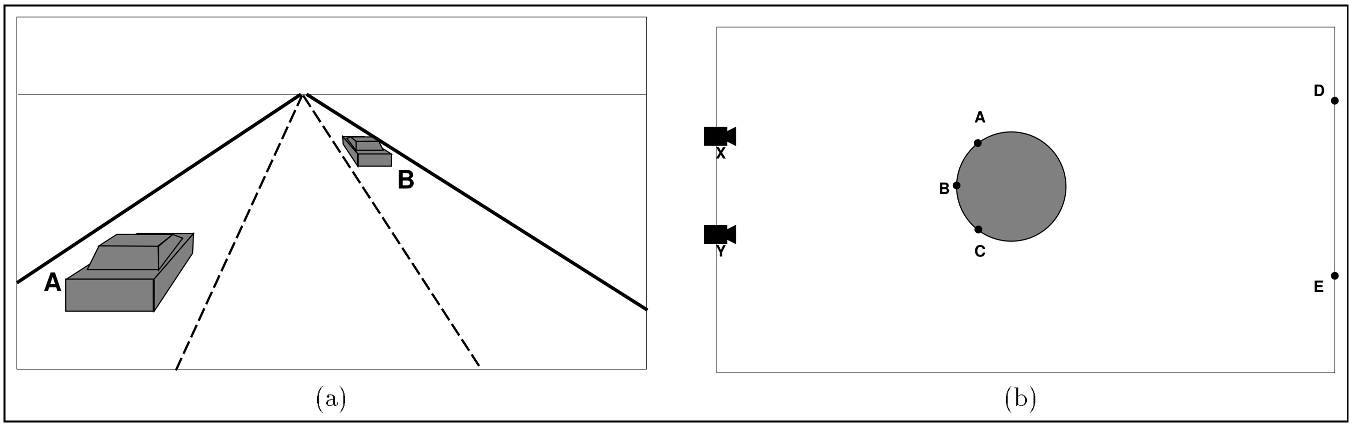

Fig. 2: (a) View from the driver's seat on a freeway. (b) Top view of a two-camera vision system observing a bottle with a wall behind.

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?

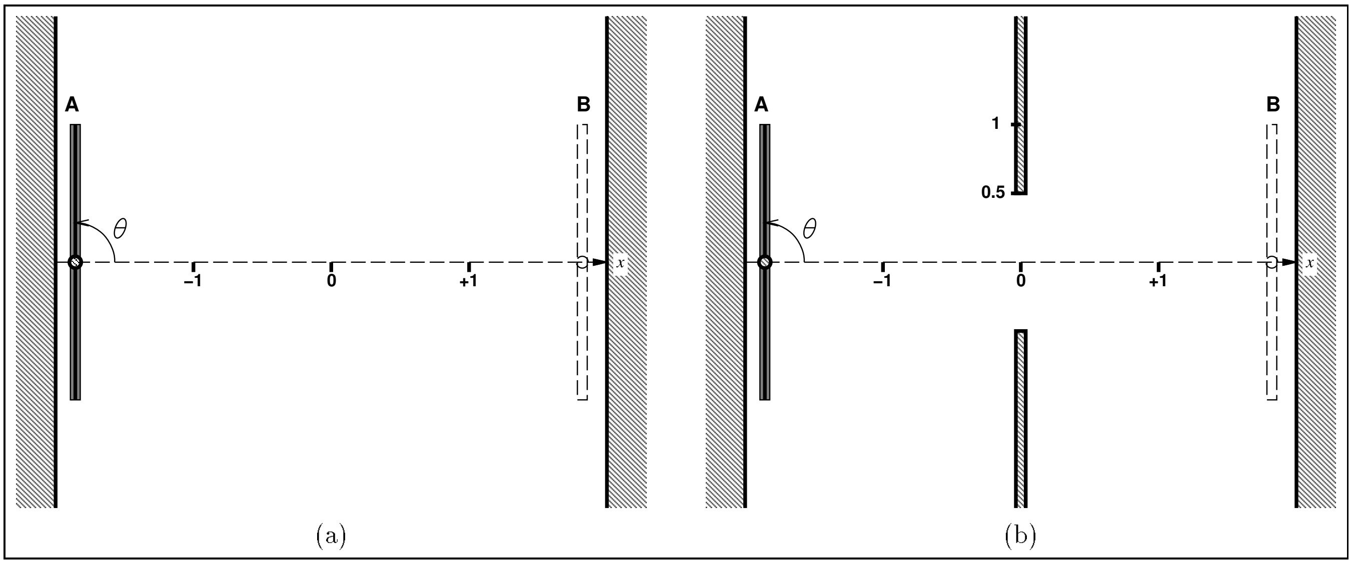

Fig. 3: Two simple physical environments. The arm can rotate around the pivot, which can slide along the x–axis.

6. (12+7 pts.) Robotics

Fig. 3(a) above shows a simple physical setup with a pivoted arm at A whose pivot can slide along the x–axis. The arm can rotate around the pivot, and the angle of rotation is θ, as shown.

(a) (7) Draw the configuration space for the arm, marking on your diagram the points A and B and showing a solution path for the problem of going from A to B. Your drawing need not be exact.

(b) (7 *extra credit*) Draw the configuration space for the situation in Fig. 3(b) and say what you can about the equations of the configuration-space obstacle boundaries.

(c) (5) Suppose that the "doorway" opening in Fig. 3 were very narrow—only a tiny bit wider than the arm itself. How would a real robot, which has rather inaccurate sensors and effectors, get through the door?

EXAM CONTINUES ON NEXT PAGE

7. (20 pts.) **Learning** A *1–decision list* or 1-DL is a decision tree with Boolean inputs in which at least one branch from every attribute test leads immediately to a leaf (obviously the final test leads to two leaves).

(a) (3) Draw a 1-DL that is equivalent to the *disjunction* of three attributes a_1, a_2, a_3.

(b) (3) Now consider the representation of Boolean functions by perceptrons, where "true" is +1 and "false" is –1. Assume a fixed input $a_0 = -1$ and three other inputs a_1, a_2, a_3, and assume the activation function is a step function $g(x) = +1$ if $x > 0$, -1 otherwise. Draw a perceptron that represents the *disjunction* of three attributes a_1, a_2, a_3.

(c) (6) For the following three 1-DLs and three perceptrons, say which is equivalent to which [e.g, A=1 B=2 C=3; hint: this is not the right answer].

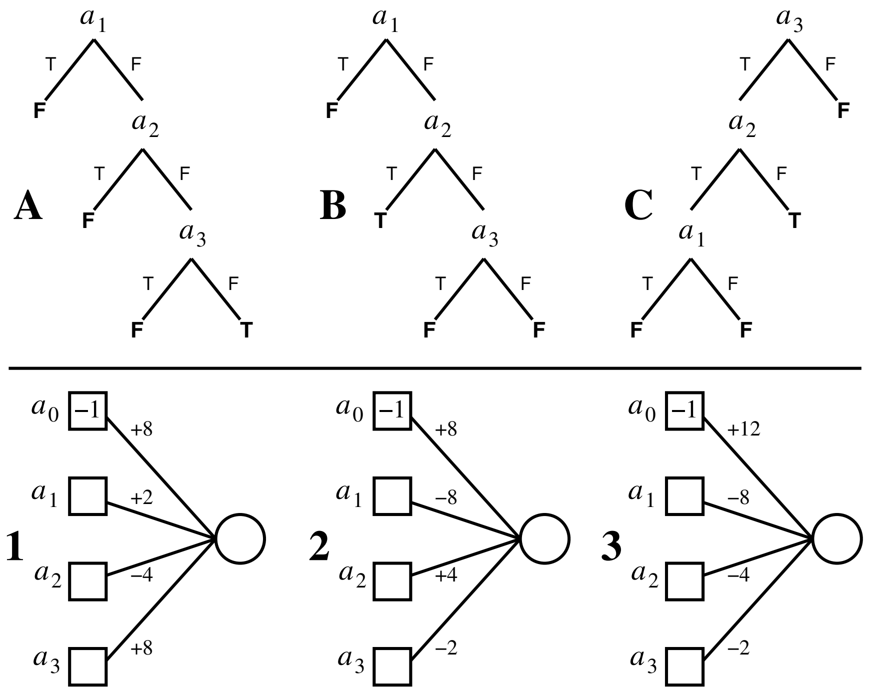

(d) (5, nontrivial) Is it true that perceptron learning will converge to a perfect representation of *any* data set where the target function is a 1-DL? Explain.

(e) (3) Give a simple argument to show that a perceptron cannot represent some data sets generated by decision *trees*.