CS 188 Fall 2002 — Introduction to AI — Stuart Russell — Final Solutions

1. (10 pts.) Some Easy Questions to Start With

(a) (2) False; ambiguity resolution often requires commonsense knowledge, context, etc. Dictionaries often giv multiple definitions of a word, which is one *source* of ambiguity.

(b) (2) True; it's hard to see directly into the mental state of other humans.

(c) (2) False; feedforward NNs have no internal state, hence cannot track a partially observable environment.

(d) (2) False; obviously it depends on how well it's programmed, and we don't know how to do that yet (despite what Kurzweil, Moravec, Joy, and others suggest).

(e) (2) False; odometry errors accumulate over time; the robot must sense its environment.

2. (13 pts.) Search

(a) (1) The branching factor is 4 (number of neighbors of each location).

(b) (2) The states at depth k form a square rotated at 45 degrees to the grid. Obviously there are a linear number of states along the boundary of the square, so the answer is $4k$.

(c) (2) Without repeated state checking, BFS expends exponentially many nodes: counting precisely, we get $((4^{x+y+1}-1)/3)-1$.

(d) (2) There are quadratically many states within the square for depth $x+y$, so the answer is $2(x+y)(x+y+1)-1$.

(e) (2) True; this is the Manhattan distance metric.

(f) (2) False; all nodes in the rectangle defined by $(0,0)$ and (x,y) are candidates for the optimal path, and there are quadratically many of them, all of which may be expended in the worst case.

(g) (1) True; removing links may induce detours, which require more steps, so h is an underestimate.

(h) (1) False; nonlocal links can reduce the actual path length below the Manhattan distance.

3. (6 pts.) Propositional Logic

(a) (1) False; $\{A = false, B = false\}$ satisfies $A \Leftrightarrow B$ but not $A \vee B$.

(b) (1) True; $\neg A \vee B$ is the same as $A \Rightarrow B$, which is entailed by $A \Leftrightarrow B$.

(c) (2) True; the first conjunct has models and entails the second.

(d) (2) True; both have four models in A, B, C.

4. (10 pts.) First-Order Logic

(a) (3) "Every cat loves its mother or father" can be translated as

  i. $\forall x\ \neg Cat(x) \vee Loves(x, Mother(x)) \vee Loves(x, Father(x))$.

(b) (3) "Every dog who loves one of its brothers is happy" can be translated as

  i. $\forall x\ Dog(x) \wedge (\exists y\ Brother(y,x) \wedge Loves(x,y)) \Rightarrow Happy(x)$.

  ii. $\forall x,y\ Dog(x) \wedge Brother(y,x) \wedge Loves(x,y) \Rightarrow Happy(x)$.

(c) (4) (a)(ii) and (a)(iii) both contain disjunctions of two positive literals and hence are not Horn-clause-representable. (b)(i) and (b)(ii) are logically equivalent; (b)(ii) is obviously Horn, so (b)(i) is too.

5. (8 pts.) Logical Inference

(a) (1) False; $Q(a)$ is not entailed.

(b) (2) True; via $P(F(b))$.

(c) (2) True; breadth-first FC is complete for Horn KBs.

(d) (2) False; infinite loop applying the first rule repeatedly.

(e) (1) False; $P(b)$ is an example.

6. (16 pts.) Probability, Bayes Nets, Decision Theory

(a) (3) (ii) and (iii). (iii) is the "correct" model and (ii) is a complete network and can represent any distribution. (i) is incorrect because *Wrapper* and *Shape* are in fact dependent.

(b) (iii) because it is correct and minimal.

(c) True; there is no link, so they are asserted to be independent.

(d) 0.59; $P(W = Red) = P(W = Red|F = s)P(F = s) + P(W = Red|F = a)P(F = a) = (0.7 \times 0.8) + (0.3 \times 0.1)$.

(e) (3) > 0.99; $P(F|S = r, W = r) = \alpha P(S = r, W = r|F)P(F) = \alpha P(S = r|F)P(W = r|F)P(F) =$
$= \alpha\langle 0.8 \times 0.8 \times 0.7, 0.1 \times 0.1 \times 0.3\rangle = \alpha\langle 0.448, 0.003\rangle$.

(f) (2) Strictly speaking, we have to assume risk neutrality, which is reasonable for one candy: $0.7s + 0.3a$.

(g) (3) This is tricky and can be viewed in several ways. Here's one argument: Less than before. An owner of a wrapped candy has now lost a choice (eat or sell, vs. eat, sell, or unwrap) hence the state of owning the candy has lost expected value.

7. (12 pts.) Dynamic Bayes Nets

(a) (3) The only tricky bit is noticing that X and Y evolve independently of each other given the action.

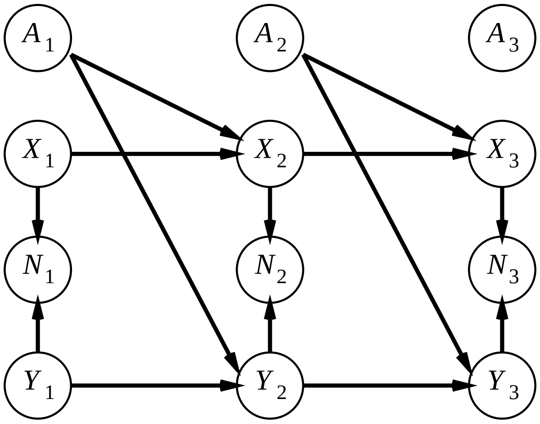

(b) (2) After $N_1 = 1$, the agent could be in any of the four edge squares. Of these, only (2,1) and (2,3) are consistent with $N_2 = 2$ after *Right*.

(c) (2) Could not have been in (1,2) and (2,2).

(d) (2) False; even with no sensors, the agent can execute, e.g., LLDD and know that it is in (1,1).

(e) (3) True; this DBN is not completely connected between layers, so its transition model will increase in size.

8. (15 pts.) Learning

(a) (2) False positives are X_3, X_4, X_6.

(b) (2) No false negatives.

(c) (2) Only $Sunny = T$ has a mixture of positive and negative examples.

(d) (3) The examples at the leaf are X_2, X_3, X_4, X_6. Only $Warm$ provides any information to distinguish these.

(e) (2) True; plot on the 2-D plane to show separability.

(f) (1) True; any Boolean function can be so represented.

(g) (2) True.

(h) (1) True. X_2 and X_6 are inconsistent.

9. (10 pts.) Robotic Path Planning

(a) (3) (i) and (iii) are sound. (ii) may fail by returning a mixed-cell path that is not traversable.

(b) (3) (i) is not complete because there may be a mixed-cell path but no pure-cell path. (iii) is not complete because it may go into an infinite loop of recursive subdivision if the freespace boundary has a local minimum at a coordinate that is not exactly representable in finitely many bits. Hence, none of the three is sound and complete.

(c) (1) True; see the example above, and assume an optimal path in pure free-space cells.

(d) (3) Run (iii) but treat mixed cells of side $\epsilon/\sqrt{d}$ or less as if they were pure obstacle cells.