Computer Science Department
Stanford University
Comprehensive Examination in Artificial Intelligence
97/98

January 21, 1998

READ THIS FIRST!

1. You should write your answers for this part of the Comprehensive Examination in a BLUE BOOK. Be sure to write your MAGIC NUMBER on the cover of every blue book that you use.

2. Be sure you have all the pages of this exam. There are 4 pages.

3. This exam is OPEN BOOK. You may use notes, articles, or books—but no help from other sentient agents such as other humans or robots.

4. Show your work, since PARTIAL CREDIT will be given for incomplete answers. For example, you can get credit for making a reasonable start on a problem even if the idea doesn't work out; you can also get credit for realizing that certain approaches are incorrect. On a true/false question, you might get partial credit for explaining why *you* think something is true when *we* think it is actually false. But no partial credit can be given if you write nothing.

This exam contains 60 points. We have allocated points according to the number of minutes that we believe a student familiar with the material should take to do the questions. If you are somewhat less familiar with the material, a question may well take you longer than the number of points that it's worth. **DON'T FALL INTO THIS TRAP.** If you are taking too long on a question, write down whatever you have and move on. You may well get substantial partial credit for a partial answer, but you will not get any partial credit for a blank answer!

1. **KNOWLEDGE REPRESENTATION [16 points]**

   (a) **[3 points]** Translate each of the following sentences into first-order logic. Use the following vocabulary: the predicate symbols Workstation(x), Monitor(x), Component(x,y) (x is a component of y), Belongs(x,p) (x belongs to p), and At(x,y) (x is at location y); the function symbol Office-of(p); and the constant symbols Tom and Sun.[1]

      i. **[1 points]** All workstations come with a monitor.

      ii. **[1 points]** Any workstation belonging to a person is in that person's office.

      iii. **[1 points]** The workstation Sun belongs to Tom.

   (b) **[2 points]** Can you prove from these axioms that there is a monitor in Tom's office? If not add any (consistent) axioms necessary for proving this fact.

   (c) **[4 points]** Translate each of the above sentences, including the negated query, into clausal form. For brevity, you may use (here and below) the first initial of the various vocabulary symbols to represent them (e.g., you can use $C(x, y)$ instead of *Component(x,y)*). You may use either disjunctive normal form or implicational normal form.

   (d) **[7 points]** Using the clauses you have generated, use refutation resolution to prove that there exists a monitor in Tom's office.

   Answer format: number the clauses in your axioms, as well as any clauses generated during your proof; for each step in your proof, specify the numbers of the clauses resolved, any variable unifications needed, and the resulting clause. For brevity, you may resolve more than two clauses in one resolution step.

2. **PROBABILITY [8 points]**

   Consider the following Bayesian network (influence diagram) over three binary-valued variables:

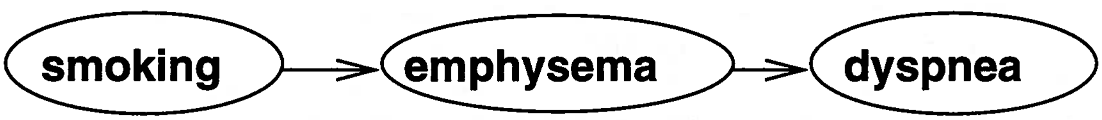

P(smoking) = 0.2

P(emphysema | smoking) = 0.6

P(emphysema | no smoking) = 0.05

P(dyspnea | emphysema) = 0.9

P(dyspnea | no emphysema) =0.15

   (a) **[2 points]** Name one conditional independence assumption which is encoded in the structure of this network.

   (b) **[6 points]** Show how you would compute *P(smoking | dyspnea)* in this network. For brevity, you may use the abbreviations s, e, and d for the events smoking, emphysema, and dyspnea, and the abbreviations $\neg s$, $\neg e$, and $\neg d$ for their negations. A formula for the answer is fine; you do not have to compute the final numerical answer.

[1] If we were being careful, you would also have a Person(x) predicate. However, for simplicity we'll ignore that.

3. **LEARNING [10 points]**

The marketing department of Microsquish Corporation is trying to construct a classifier which will tell them whether a customer will like a piece of software. From responses to survey forms, they have collected the following data:

Form #	customer is computer nerd	s/w has lots of features	customer liked s/w
1	yes	yes	yes
2	yes	yes	yes
3	yes	yes	yes
4	yes	yes	yes
5	yes	yes	yes
6	yes	yes	yes
7	yes	yes	yes
8	yes	no	no
9	yes	no	no
10	no	yes	yes
11	no	yes	no
12	no	yes	no
13	no	no	yes
14	no	no	yes
15	no	no	no

(a) **[2 points]** What Boolean function would do the best job of classifying these examples?

(b) **[5 points]** What decision-tree, including classifications, would be output by an ID3-style decision-tree learning algorithm? Explain or show your computations.

(c) **[3 points]** Is it possible to construct a neural network with a single thresholding element (i.e., a perceptron) which classifies these examples as well as the decision tree? If so, show the parameters of the thresholding unit. If not, explain why not.

4. **SEARCH [14 points]**

(a) **[9 points]** Consider the problem of heuristic search in a space where our heuristic function h' is almost, but not quite, admissible. More precisely, we have that for each node n, $h'(n) \leq h(n) + \epsilon$. In this case, we are not guaranteed that the first goal node returned by the A* algorithm will be the *optimal* goal node. Explain briefly how you could extend the A^* algorithm in order to find the optimal goal node n^*. (Hint: Consider the f-values of n and n^*.)

(b) **[5 points]** Consider the problem of search in the (familiar) blocks world domain. The domain consists of several square blocks of equal size on a (very large) table. Each block can be on the table or on top of exactly one other block. (Locations on the table are not labelled, so we don't care where on the table a block is.) A block can have at most one block on top of it. Operators in this space consist of taking one clear block (one which is not under any other block) and moving it to any other legal location (on top of a different clear block or to the table). Each application of an operator has cost 1.

Assuming your state space consists of complete blocks world configurations, and that the goal is a single fully specified state (e.g., a specific configuration of towers), define a nontrivial *admissible* heuristic function for this domain. Try to make your heuristic as powerful as possible while still making it admissible and relatively easy to compute. Explain why your heuristic is admissible.

5. **SHORT ANSWERS [12 points]** Each of the following questions requires at most one sentence in response. Do not write more.

   (a) **[1 points]** True or false: α-β pruning, although typically more efficient than minimax, can occasionally result in a less desirable move. (No explanation is required.)

   (b) **[3 points]** Consider the action of toggling a light switch, whose effect is to turn the light on if it's currently off, and to turn the light off if it's on. Can you provide a pure STRIPS description (as in Section 14.5 of Ginsberg's book) of the *toggle* action, assuming that the predicates in our language are *On(x)* and *Off(x)*? Either show the description, or explain (in one sentence) why one is impossible.

   (c) **[3 points]** R2D1 is an office cleaning robot which vacuums the floor as it moves around. The following is a STRIPS description of its *move(x,y)* action (where *at(x)* is true if the robot is at location x):

   Preconditions: *at(x)*, *adjacent(x,y)*.

   Add list: *at(y)*, *clean(x)*.

   Delete list: *at(x)*.[2]

   Is the following situation calculus axiom an equivalent description of the effects of this action (yes/no)?

   $$\forall s,x,y((at(x,s) \wedge adjacent(x,y)) \rightarrow \quad (at(y, result(move(x,y),s)) \wedge clean(x, result(move(x,y),s)) \wedge \neg at(x, result(move(x,y),s))))$$

   If not, why not (one sentence)?

   (d) **[3 points]** The following figure contains two consistent Waltz labellings for the same line drawing. The first (a) is essentially the one given by Ginsberg, only slightly simplified. Is the second labelling (b) illegal (yes/no only)? If not, what differences in the physical object do the differences in the labellings represent (one sentence)? (Note: the differences are in the labelling of the construction at the center of the object.)

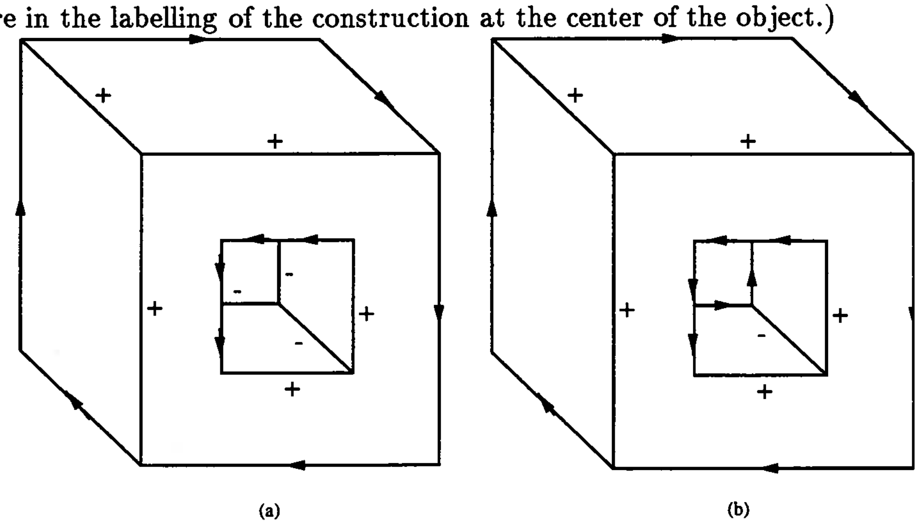

[2]In the notation of Russell & Norvig, the effects of this action are *at(y)*, *clean(x)*, $\neg$*at(x)*.

(e) **[2 points]** In the following nonmonotonic semantic network, which of the following conclusions can we make (using brave extensions, as described in Ginbserg's book)? Mark all that apply; no explanation is required.

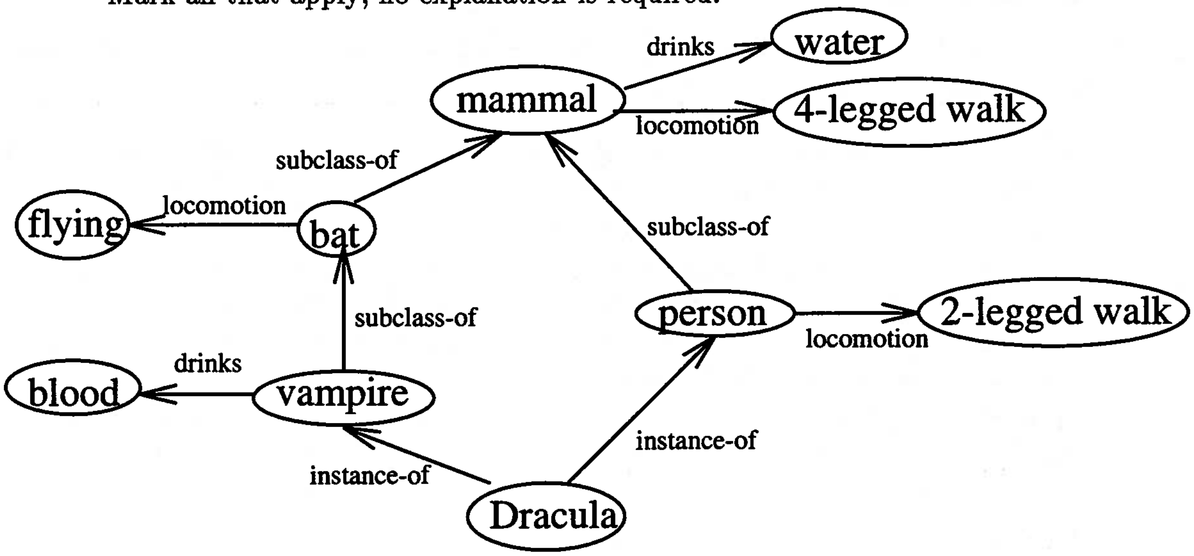

(a) Dracula's locomotion is by walking on two legs.
(b) Dracula's locomotion is by walking on four legs.
(c) Dracula's locomotion is by flying.
(d) Dracula drinks blood.
(e) Dracula drinks water.