CS 188 Fall 1992 — Introduction to AI — Stuart Russell — Final examination

You have 2 hours 50 min. The exam is open-book, open-notes. 100 points total
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. In the true/false and multiple choice questions, negative points will be awarded for incorrect answers in such a way as to render random guessing valueless.

1. **(12 pts.)** **Definitions** Provide brief, *precise* definitions of the following:
   (a) Grammar
   (b) Stereopsis
   (c) Information value
   (d) Nonlinear planning
   (e) Inheritance
   (f) Qualification problem

2. **(14 pts.)** **Logic** *True/false*:
   (a) (2) Backward-chaining using Extended Modus Ponens is complete for Horn clause databases.
   (b) (2) Forward-chaining and backward-chaining together are complete for first-order logic databases.
   (c) (2) By controlling the choice of which clauses to resolve, resolution can simulate backward-chaining.
   (d) (2) $Hates(f(x), f(f(x)))$ unifies with $Loves(y, f(y))$.
   (e) (2) "A potatoe is always larger than a tomatoe" is a good translation of
   $\forall xy Potato(x) \wedge Tomato(y) \Rightarrow > (Volume(x), Volume(y))$
   (f) (2) $[\exists x EmployeeOf(x, UC)] \Rightarrow [\exists y Earns(x, y) \wedge < (y, \$100,000)]$
   is a good translation of "If someone works for UC they earn less than $100,000."
   (g) (2) $\exists d Dog(d) \Rightarrow [\forall c Cat(c) \Rightarrow > (Weight(d), Weight(c))]$
   is a good translation of "There is a dog that is heavier than any cat."

3. **(6 pts.)** **AI** *True/false*:
   (a) The Turing test provides a formal definition of intelligence.
   (b) Most AI capabilities already exist — we just need faster computers.
   (c) In principle (ignoring speed and memory), a giant belief net could pass an unlimited-length Turing test.

4. **(12 pts.)** **Planning** *Multiple choice*:
   The presence of other, hostile agents in the environment has the following implications for planning systems:
   (a) Decision theory must be used
   (b) One can only plan for the near future
   (c) Plans with long sequences of interdependent actions are inappropriate
   (d) Execution monitoring is useful
   (e) Conditional plans are useful
   (f) A guaranteed plan can never be constructed

5. **(16 pts.)** **Search** *Multiple choice*:
   In this question we consider the behaviour of search algorithms on a finite-memory machine (assume the memory is large). For each of the following properties, say which of the eight search methods given below exhibits the property:

(a) exhaustive in a small finite space (ie will eventually expand any node of finite depth or find a goal state)

(b) admissible

(c) exhaustive in an infinite space

(d) heuristic

i) Best-first ii) A* iii) SMA* iv) hill-climbing
v) Breadth-first vi) Depth-first vii) Iterative-Deepening-Depth-first
viii) hill-climbing with random restart, remembering best-so-far.

6. (21 pts.) Probability and decision theory

(a) (2) Define conditional probability and write down Bayes' theorem

(b) (3) In the Chernoclinton nuclear power station, there is an alarm which senses when a temperature gauge exceeds a given threshold. The gauge measures the core temperature. Let $A =$ "alarm sounds"; $G = t =$ "Gauge reads t"; $T = t =$ "core temp is t"; $F_A =$ "alarm is faulty"; $F_G =$ "gauge is faulty". Draw a belief net for this domain, given that the gauge is more likely to fail when the core temperature gets too high.

(c) (1) Is your network solvable by a local propagation algorithm?

(d) (2) Suppose there are just two possible actual and measured temperatures, Normal and High; and that the gauge gives the incorrect temperature $x\%$ of the time when it is working, but $y\%$ of the time when it is faulty. Give the conditional probability table associated with G.

(e) (2) Suppose the alarm works unless it is faulty, in which case it never goes off. Give the conditional probability table associated with A.

(f) (4) Suppose the alarm and gauge are working, and the alarm sounds. Calculate the probability that the core temperature is too high. (Hint: odds/likelihood is probably easiest.)

(g) (3) In a given time period, the probability that the temperature exceeds threshold is p. The cost of shutting down the reactor is c_s, the cost of not shutting it down when the temperature is in fact too high is c_m (m stands for ...). Assuming the gauge and alarm to be working normally, calculate the maximum allowable value for x (i.e., if x is any higher than this we have to shut down the reactor all the time).

(h) (2) Suppose we add a second temperature gauge H, and connect the alarm so it goes off when either gauge reads High. Should G be connected to H in the belief network?

(i) (2) Are there circumstances under which adding a second gauge would mean that we would need *more* accurate (i.e., more likely to give the correct temperature) gauges? Why (not)?

7. (12 pts.) Learning
A version space learning system uses a concept language, a background knowledge base (KB), and instance descriptions to arrive at good inductive hypotheses concerning the definition of a goal concept. In this question we will logically analyse its behaviour.

(a) (1) Give an informal definition of what it means for an instance to be a *false positive* for a concept.

(b) (2) What does the VS algorithm use the KB for?

(c) (3) Recall that a *hypothesis* (H, say) is a statement that the goal concept Q has a particular definition C_j. An *instance description* is a sentence giving the description and classification of some object or situation a_i. Give a *formal* logical definition of what it means for a hypothesis to be contradicted by an instance.

(d) (2) Suppose one believes that the correct hypothesis is contained in a hypothesis space consisting of hypotheses $H_1 \ldots H_6$. To what sentence S in *propositional* calculus does this belief correspond?

(e) (4) Suppose 3 instances (with their classifications) are described by the sentences I_1, I_2, I_3. We run the version space algorithm on the instances with the hypothesis space from part d), and come up with only one hypothesis H remaining. Is H logically entailed by S, KB and $I_1 \wedge I_2 \wedge I_3$? Why (not)?

8. (7 pts.) Vision, NLP What do vision and natural language understanding have in common? (Be reasonably brief!)