Stanford University
Computer Science Department

Fall 2005 Comprehensive Exam in Artificial Intelligence

1. **Closed Book - NO laptop. Write only in the Blue Book provided.**
2. **The exam is timed for one hour.**
3. **Write your Magic Number on this sheet and the Blue Book; DO NOT WRITE YOUR NAME.**

The following is a statement of the Stanford University Honor Code:

A. The Honor Code is an undertaking of the students, individually and collectively:

1. that they will not give or receive aid in examinations; that they will not give or receive un-permitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;

2. that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.

B. The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.

C. While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.

By writing my Magic Number below, I certify that I acknowledge and accept the Honor Code.

Magic Number ------------------------------------

2005 Comprehensive Examination
Artificial Intelligence

1. Search. (20 points) Consider a search tree with uniform branching factor 2 and depth d, and consider a search problem for which there is a single solution in the tree at depth. Give expressions for the best and worst case cost of finding the solution, in terms of nodes visited, for (a) breadth-first search, (b) depth-first search, and (c) iterative deepening starting at depth 1 and incrementing by 1 on each iteration. For the purposes of this problem, you should assume that the root of the tree is at depth 1.

2. Logic. (20 points) Let Γ and Δ be sets of closed sentences in first-order logic, and let φ and ψ be individual closed sentences in first-order logic. State whether each of the following statements is true or false. No explanation is necessary.

(a) $\Gamma \models \varphi$ and $\Delta \models \psi$, $\Gamma \cup \Delta \models (\varphi \wedge \psi)$.
(b) $\Gamma \models \varphi$ and $\Delta \models \psi$, $\Gamma \cap \Delta \models (\varphi \vee \psi)$.
(c) $\Delta \models (\varphi \Rightarrow \psi)$ if and only if $\Delta \cup \{\varphi\} \models \psi$.
(d) $\Delta \models \varphi$ or $\Delta \models \psi$ if and only if $\Delta \models (\varphi \vee \psi)$.
(e) If $\Delta \models \varphi$ and $\Delta \models \psi$, then $\Delta \models (\varphi \Rightarrow \psi)$.
(f) If $\Delta \models p(\tau)$ for some ground term τ, then $\Delta \not\models \forall x.\neg p(x)$.
(g) If $\Delta \models p(\tau)$ for every ground term τ, then $\Delta \models \forall x.p(x)$.
(h) If $\Delta \models \exists x.(p(x) \Rightarrow q(x))$, then $\Delta \models \exists x.(p(x) \wedge q(x))$.
(i) If $\Delta \models \varphi$ and $\Delta \models (\varphi \Rightarrow \psi)$, then $\Delta \not\models \neg\psi$.
(j) If $\Gamma \models (\varphi \Rightarrow \psi)$ and $\Delta \models (\psi \Rightarrow \varphi)$, then $\Gamma \cap \Delta \models (\varphi \Rightarrow \psi) \vee (\psi \Rightarrow \varphi)$.

3. Automated Reasoning. (20 points) Use the resolution method and the following premises to prove the conclusion shown below.

Premises:
$\forall x.\forall y.\forall z.(p(x,y,z) \Rightarrow s(x,y,z))$
$\forall x.\forall y.\forall z.(q(x,y,z) \Rightarrow s(x,y,z))$
$\forall x.\exists y.\forall z.(p(x,y,z) \vee q(x,y,z) \vee \neg r(x))$

Conclusion:
$\forall x.(r(x) \Rightarrow \exists y.(s(x,y,b) \wedge s(x,y,c)))$

Note that this is a question about Resolution. You will get zero points (nil, nada, rien, zip, nothing) unless you prove it using the standard resolution procedure.

4. Bayes Nets. (20 points) Consider the following Bayesian.

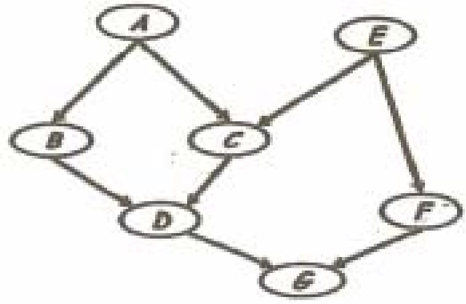

(a) Write $p(G)$ as a sum of values from the full joint distribution for variables A,..., G. Equivalently, what is the mathematical expression computed in this case by the Enumeration-Joint-Ask algorithm described in Russell and Norvig?

(b) Write $p(G)$ as a nested sum of products of conditional probabilities from the tables associated with this Bayes Net. Equivalently, what is the mathematical expression computed in this case by the Enumeration-Ask algorithm described in Russell and Norvig?

(c) Assume we are trying to compute $p(G)$. Fill in the following table to show what the variable elimination algorithm will do on this case on the assumption that we are eliminating variables in the order B, C, D, F (with B being removed first, not last). Note: your generated factors should be written in the form $g_i(X_1, \ldots, X_i)$ that clarifies the variables involved in the factor.. Equivalently, write the factors that would be generated in this case by the Elimination-Ask algorithm described in Russell and Norvig.

Variable	Factors Used	Factors Generated
B		
C		
D		
F		

5. Learning. (20 points) Consider the following training set for a decision-tree learning problem. Here, a, b, c, d, and e are Boolean features, and $x_1, \ldots, x_6$ are samples.

Example	a	b	c	d	e	Goal
x_1	1	0	0	0	0	1
x_2	1	1	1	1	0	0
x_3	0	0	0	1	0	1
x_4	1	1	1	0	0	1
x_5	1	0	0	1	0	0
x_6	0	0	0	0	0	0

(a) Draw a decision tree of minimal depth that correctly classifies the examples in this dataset.

(b) How much information is needed to classify an example in this case? (Reminder: the amount of information needed to classify an example is -(*p*log p*)-(*n*log n*), where *p* is the probability of a positive answer and *n* is the probability of a negative answer.)

(c) How much information is needed to classify an example *given* that a is 1? What if a is 0?

(d) So, what is the information gain from attribute a? What is the information gain from d given a?