	1	2	3	4	5	6	7	8	9	total
	10	10	40	30	20	15	30	25	30	210

Name:_____

UMBC CMSC471/671 Final Exam
December 14, 2006

Please write all of your answers on this exam. The exam is closed book and consists of N problems which add up to M points. You have the two hours to work on this exam. Good luck.

0. Warm Up Exercise (0)

The department just purchased a FixBot maintenance robot (model 2.0) and plans to use it to maintain all the department's robots who do not maintain themselves. Is there a problem with this plan?

1. Uninformed search (10)

Briefly discuss the advantages and disadvantages of depth and breadth first search. What sort of problem is each appropriate for?

.

2. Heuristic search (10)

Briefly describe all of the conditions that must hold in order for an instance of graph search (i.e., algorithm A) to be an instance of A* search. Include any relevant constraints on the size and structure of the graph, the costs associated with the arcs, the number of goal states, and the heuristic function employed.

3. True/False (40 points)

Circle either T or an F in the space before each statement to indicate whether the statement is true or false. If you think the answer is simultaneously true and false, quit while you are ahead. There is no penalty for incorrect answers (but then, there are no points for incorrect answers either).

T F Every complete inference procedure is also sound.

T F Every sound inference procedure is also complete.

T F An incomplete inference procedure can produce an incorrect answer.

T F An unsound inference procedure can produce an incorrect answer.

T F The scope of a variable in Prolog is just the clause it is in.

T F Prolog programs are sound but not complete.

T F Hill climbing search algorithms only work for search spaces that are two-dimensional or have solution-preserving projections onto two-dimensions.

T F Iterative deepening is essentially a technique to combine the best features of breadth-first and depth-first search.

T F Overfitting can occur in machine learning when the data contains many irrelevant attributes.

T F The decision tree learning algorithm covered in class finds an optimal decision tree, i.e., one that minimizes the number of questions needed to classify a case.

T F Ockham's razor is a heuristic that says to prefer the simplest consistent explanation .

T F Decision tree learning can not be used if the data is noisy.

T F In a proper Bayesian network, a node is always conditionally independent of its non-descendants given its parents.

T F *Information gain* is used to determine the network structure in a Bayesian Network.

T F Random variables A and B are independent if $p(A \wedge B) = p(A|B)*p(B)$.

T F If two variables are independent, then they are always conditionally independent given an arbitrary third variable.

T F In searching for a plan, a partial order planner searches though a space formed by possible situations.

T F A partial order planner can either search forward from an initial state to a goal, backward from a goal state to the initial state, or in both directions toward the middle.

T F The situation calculus is an approach to reasoning about changes in random variables using first order logic.

T F Resolution is refutation-complete, which means that if a set of sentences is unsatisfiable, then a contradiction can always be derived in a finite amount of time.

4. English to logic (30)

Translate the following statements into a **single sentence** in first order logic, choosing appropriate predicates and functions. Avoid redundancy.

    (a) Good food is not cheap and cheap food is not good.

    (b) If a computer beats Kramnik, then that computer beats every person.

    (c) Every person is either good or bad but not both.

Now rewrite each of these in clausal normal form, i.e. as a set of sentences without quantifiers in which each sentence is a disjunction of atomic literals. Recall that an atomic literal is either a predicate (i.e, loves(john,X)) or the negation of a predicate (i.e., not(loves(john, father(mary)))). Avoid redundancy.

    (d) Good food is not cheap and cheap food is not good.

    (e) If a computer can beat Kramnik in chess, then a computer can beat anyone in chess.

    (f) Every person is either good or bad but not both.

5. Logic to English (20)

(a) Which of the following are correct translations of "No two adjacent countries have the same color"?

    i.   ∀x ∀y ~Country(x) v ~Country(y) v ~Adjacent(x,y) v ~(Color(x) = Color(y))
    ii.   ∀x ∀y (Country(x) ^ Country(y) ^ Adjacent(x,y) ^ ~(x = y))) => ~(Color(x) = Color(y))
    iii.   ∀x ∀y Country(x) ^ Country(y) ^ Adjacent(x,y) ^ ~ (Color(x) = Color(y))
    iv.   ∀x ∀y (Country(x) ^ Country(y) ^ Adjacent(x,y)) => Color(x ≠ y)

(b) Which of the following are correct translations of "No dog bites a child of its owner" ?

    i.   ∀x Dog(x) => ~Bites(x,Child(Owner(x)))
    ii.   ~Ex ~Ey Dog(x) ^ Child(y,Owner(x)) ^ Bites(x, y)
    iii.   ∀x Dog(x) => (∀y Child(y,Owner(x)) => ~Bites(x, y))
    iv.   ~Ex Dog(x) => (Ey Child(y,Owner(x)) ^ Bites(x, y))

6. Failure as negation (15)

(a) Describe the relationship, if any, between Prolog's *failure as negation* and *non-monotonic reasoning*.

(b) Write a Prolog rule encoding the knowledge that "All birds fly, except for penguins, kiwis and dead birds" choosing appropriate names for your predicates.

7. Decision trees (30)

You are constructing a decision tree and are deciding which attribute is best to ask about first using the *information gain* metric. You are given a set of 128 examples, with 64 positively labeled and 64 negatively labeled. There are three attributes, HomeOwner, InDebt, and Rich. For 64 examples, HomeOwner is true with 1/4 negative and 3/4 positive. For 96 examples, InDebt is true with half positive and half negative. For 32 examples, Rich is true with 3/4 positive and 1/4 negative. Recall that the entropy of a distribution P with probabilities $(p_1...p_n)$ is defined as $-\Sigma\, p_i *\ln(p_i)$. Information gain is the difference in entropy before asking the question and the (weighted) entropy after.

(a) Complete the tables to the right to help you with your computations.

(b) What's the entropy of the initial data?

(b) What is the information gain of splitting on the HomeOwner attribute as the root node?

(c) What is the information gain of splitting on the InDebt attribute as the root node?

(d) What is the information gain of splitting on the Rich attribute as the root node?

(e) Which attribute do you split on?

Note	
ln(1/4)	= -2.0
ln(5/12)	= -1.263
ln(1/2)	= -1.0
ln(7/12)	= -.778
ln(3/4)	= -0.415
5/12	= 0.417
7/12	= 0.583

ALL	neg	pos
	64	64

HO	neg	pos
true	16	48
false		

ID	neg	pos
true	48	48
false		

RICH	neg	pos
true	8	24
false		

8. Bayesian networks (25)

You're building a troubleshooting advisor for PCs. Let CF be a boolean random variable representing whether the computer fails (CF=true) or not. Assume there are two possible causes of failure: Electricity-Failure and Malfunction-of-the-Computer represented using the Boolean random variables EF and MC, respectively. The table below lists some of the common formulae. Assume the following probabilities

$P(EF) = 0.1$
$P(MC) = 0.2$
$P(CF \mid {\sim}EF \wedge {\sim}MC) = 0.0$
$P(CF \mid {\sim}EF \wedge MC) = 0.5$
$P(CF \mid EF \wedge {\sim}MC) = 1.0$
$P(CF \mid EF \wedge MC) = 1.0$

rule	not conditioned	conditioned on C
basic	$P(A,B) = P(A\mid B)P(B)$	$P(A,B\mid C) = P(A\mid B,C)P(B\mid C)$
Bayes	$P(B\mid A)P(A) = P(A\mid B)P(B)$	$P(B\mid A,C)P(A\mid C) = P(A\mid B,C)P(B\mid C$
independence	A&B are independent iff $P(A\mid B) = P(A)$	A&B cond. Indep. given C iff $P(A\mid B,C) = P(A\mid C)$
chain rule	$P(A,B,C) = P(A)P(B\mid A)P(C\mid A,B)$	
axiom	$P(A) + P({\sim}A) = 1$	$P(A\mid C) + P({\sim}A\mid C) = 1$
marginalization	$P(A) = $ SUMb $P(A,B)$	$P(A\mid C) = $ SUMb $(P(A,B\mid C)$

(a) Are EF and MC independent? Use given data and make reasonable assumptions about the domain.

(b) Draw the Bayesian Network for this problem.

(b) Compute $P(CF \wedge {\sim}EF \wedge MC)$. Hint: use the chain rule.

(c) Compute $P(MC \mid EF)$

(d) Compute $P(EF \mid CF)$

9. Planning (30 points)

In the minimal block world domain used in HW7 and in class, there are just five predicates describing states (ontable/1, holding/1, handempty, on/2, clear/1) and four operators (pickup, putdown, stack, unstack). Operators are defined in terms of five elements: description, preconditions, adds, deletes, and constraints, as in this example:

stack operator
Description:	stack(X,Y)
Preconditions:	[holding(X), clear(Y)]
Adds:	[on(X,Y), clear(X)]
Deletes:	[holding(X), clear(Y)]
Constraints:	[not(Y=table)]

Define the following strips operators in this format.

(a) The **clone** operation allows the robot to create a copy of an existing object that it is holding. The newly created object will be on the table and have nothing on it. If the robot clones an object named *foo*, then the new copy will be named *copyOf(foo)*.

clone operator
Description:
Preconditions:
Adds:
Deletes:
Constraints:

(b) The **destroy** operation allows the robot to destroy an object that it is holding. The object is removed from the world description.

destroy operator
Description:
Preconditions:
Adds:
Deletes:
Constraints:

(c) The **swap** operation can reverse the order of the top two objects in a stack. That is, if a is on b, and b is on c, and nothing is on a, then after the **swap** operation, b will be on a and a will be on c and nothing will be on b. In order initiate a swap operation, the robot's hand has to be empty. Nothing else in the state will have changed.

swap operator
Description:
Preconditions:
Adds:
Deletes:
Constraints: