

Final Examination — December 7, 2018

Duration: 180 minutes

No aids allowed.

Total marks: 135.

Name:

Student Number:

1)	/15
2)	/10
3)	/8
4)	/4
5)	/8
6)	/10
7)	/20
8)	/20
9)	/10
10)	/10
11)	/10
12)	/10
Total	/135

1. [15 points] For each of the following statements indicate whether it is *true* or *false*:
- (a) Hill climbing search may get stuck in a local maximum.
 - (b) A Horn clause has at most one positive literal.
 - (c) When we make the closed world assumption about a 2-place predicate/relation r , if the knowledge base does not entail that $r(a, b)$, then $r(a, b)$ must be false.
 - (d) Iterative deepening depth-first search uses more space than breadth-first search on problems with a large search space.
 - (e) A constraint satisfaction problem that has been made arc-consistent can always be solved without backtracking.
 - (f) In forward checking, we check constraints that have only one uninstantiated variable and delete values that violate the constraint.
 - (g) If a heuristic is admissible, it must be monotone/consistent.
 - (h) If q_1 is conditionally independent from q_2 , then q_2 is conditionally independent from q_1 .
 - (i) When A^* uses a heuristic that is only admissible, it may not find the optimal path.
 - (j) Markov decision processes assume that the state is fully observable.
 - (k) The value iteration algorithm can be used to solve a Markov decision process.
 - (l) In Q-learning, one learns the transition model $P(s'|s, a)$.
 - (m) Decision-tree learning produces models that are more understandable than neural network learning does.
 - (n) A single-layer feedforward neural network (perceptron) can learn any continuous function.
 - (o) A learned model that displays overfitting generalizes well.

2. [10 points] Suppose that we have a search algorithm that keeps unexpanded nodes in a sorted list, the frontier (or fringe), and that always selects the first node from the frontier. What kind of search do we have if...
- (a) we always place the successors of the current node at the front of the frontier?
 - (b) we always place the successors of the current node at the back of the frontier?
 - (c) we insert the successors of the current node so as to keep the frontier sorted by path cost $g(n)$ (i.e. the cost of the path from the start node to node n)?
 - (d) we insert the successors of the current node so as to keep the frontier sorted by the value of the heuristic function $h(n)$?
 - (e) we insert the successors of the current node so as to keep the frontier sorted by the value of the sum of the path cost and the heuristic function $g(n) + h(n)$?

3. [8 points] Implement the following description of an *above* relation (in a blocks world) using one or more Prolog rules:

if X is on Y, then X is above Y;
if X on Y and Y is above Z, then X is above Z;
nothing else is in the above relation.

You can assume that the `on(X,Y)` predicate has already been defined.

4. [4 points] What is the output of the following Prolog program, when we invoke the query `?- q.`

```
q:- writeln(qa), r.  
q:- writeln(qb), !, r.  
q:- writeln(qc), fail.  
  
r:- writeln(ra), !, fail.  
r:- writeln(rb), fail.
```

5. [8 points] The predicate `bagof(T, Q, S)` constructs a list `S` of all instances of `T` for which `Q` is true. A textbook describes the computation as follows: the query `Q` is satisfied in all possible ways, and for each instantiation of the variables in the arguments of `Q`, `T` is instantiated and recorded as such in the list `S`.

(a) Give the value computed for `L` by the query

```
?- bagof(X, (member(X, [0, 8, 5, 4]), not(X >= 7)), L).
```

(b) The following query will not work:

```
?- bagof(X, member(0, X), L).
```

Explain briefly why.

6. [10 points] Convert the following formulas into clausal form (recall that the 8 steps are: eliminate implications, move negations inward, standardize variables, skolemize, convert to prenex, distribute \vee over \wedge , flatten conjunctions and disjunctions, and convert to clauses):

(a) $((\exists X.p(X)) \vee (\exists X.q(X))) \rightarrow (\exists X.p(X) \vee q(X))$

(b) $\forall X.(p(X) \rightarrow \exists Y.(\forall Z.(q(X, Y, Z)) \wedge r(X, Y, Z)))$

7. [20 points] Consider the problem of finding a path in a grid using *heuristic search*. The problem is to find a path from square $S = s12$ to square $G = s13$, given that you can move only horizontally and vertically, one square at a time, and no step may be made into a shaded square. Each step *has cost one*.

$s1$	$s2$	$s3$	$s4$
$s5$	$s6$		
$s7$	$s8$	$s9$	$s10$
$s11$	$\textcircled{S} s12$		$\textcircled{G} s13$

Suppose that we want to use A^* search with *cycle-checking* and the *Manhattan distance* as the heuristic function h to solve this problem. The Manhattan distance between two cells is the sum of the horizontal and vertical distances (counting shaded squares as well). Note that this heuristic is *monotonic*.

- (a) Write down the h values of squares $s1$ through $s13$ in the grid below (the values for the start node $s12$ and the goal node $s13$ have already been given):

$s1$	$s2$	$s3$	$s4$
h=	h=	h=	h=
$s5$	$s6$		
h=	h=		
$s7$	$s8$	$s9$	$s10$
h=	h=	h=	h=
$s11$	$\textcircled{S} s12$		$\textcircled{G} s13$
h=	h=2		h=0

- (b) Now, apply the A* search: continue writing down the frontier for each step of the algorithm in the format shown below (it must be sorted on f values, and f written as $f = g + h$). In the case of equal f values, break ties as follows: the square with the *larger number* comes first (e.g. s_{11} before s_{10} and s_5). Underline the node selected for expansion at each step. *Remember to enforce cycle-checking.*

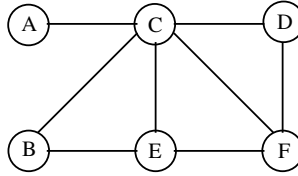
1. $\{ \langle \underline{s_{12}}, 2 = 0 + 2 \rangle \}$

2. $\{ \langle \underline{s_{11}}, 4 = 1 + 3 \rangle, \quad \quad \quad \}$

- 3.

$$\vdots$$

8. [20 points] Consider the following *constraint satisfaction problem*. We have the graph shown below with 6 nodes A, B, C, D, E , and F . We want to color each node such that no two nodes connected by an edge get the same color. There are only three colors b, g and r (i.e. blue, green, red) available for nodes C, D, E, F . Moreover, node A must be colored red and node B must be colored blue.

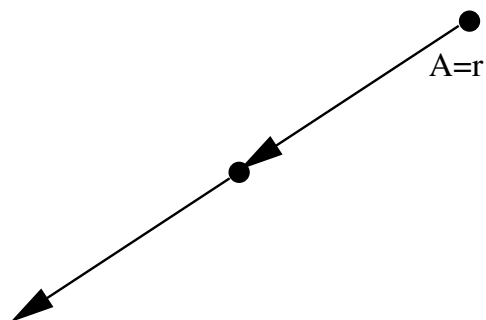
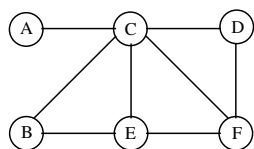


- (a) Write down the domain of each of the 6 variables (do not apply forward checking yet).

$Dom(A) = \{ \quad \quad \quad \}, \quad Dom(B) = \{ \quad \quad \quad \}, \quad Dom(C) = \{ \quad \quad \quad \},$
 $Dom(D) = \{ \quad \quad \quad \}, \quad Dom(E) = \{ \quad \quad \quad \}, \quad Dom(F) = \{ \quad \quad \quad \}.$

- (b) Complete the search tree on the next page to show the tree that would be explored by the *backtracking search with forward checking* algorithm. Remember that in this algorithm, constraints are checked when all but one of their variables have been assigned and values that violate the constraint are removed from the domain of the unassigned variable. You are to use a static variable ordering whereby each branch considers assignments to the variables in the sequence A, B, C, D, E, F . The values are also tried in the sequence b, g , and r when applicable. Annotate each edge with what variable is assigned what value and the updated domains of the remaining variables. Use the symbol DWO to mark whenever a domain wipe-out happens.

Q2b: Complete the search tree. Remember to try values in the order b, g , and then r . The graph on the left is shown for your convenience.



9. [10 points] Consider the following dynamic domain, to be specified in the situation calculus. Suppose that we have a fluent $isOpen(X, S)$ that is true if and only if door X is open in situation S and another fluent $isLocked(X, S)$ that is true if and only if door X is locked in situation S . Suppose also that we have an action $open(X)$ whose only effect is to open a closed (i.e. not open) door X provided that X is not locked, as well as an action $close(X)$ whose only effect is to close an open door X . Finally, assume that these are the only actions that affect the fluent $isOpen$.

a) Write effect axioms for the actions $open$ and $close$; your axioms should capture all the effects of these actions.

b) Write a frame axiom for the action $open$ and the fluent $isOpen$.

c) Write a successor state axiom for the fluent $isOpen$.

10. [10 points] Consider the following STRIPS actions:

Action Name	Preconditions	Add Effects	Delete Effects
A	-	p	-
B	o	-	o
C	p	q	r
D	p	r	-

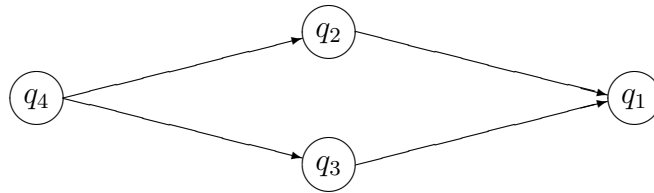
Suppose that the start state is $S_0 = \{o\}$ and the goal is $G = \{p, q, r, o\}$.

- (a) Which actions are applicable in the start state S_0 ? What is the new state for each of the applicable actions?

- (b) Suppose we want to do backward/regression planning. Which actions are consistent with the goal G ? What is the regressed goal for all the consistent actions?

- (c) Give a sequence of actions that achieves the goal G when executed in the starting state S_0 .

11. [10 points] Suppose that we have the following belief network (or Bayes net):



- (a) True or false, given this belief network, is it the case that $Pr(q_1|q_2, q_3, q_4) = Pr(q_1|q_2, q_3)$?
- (b) What are the conditional probabilities that need to be specified to fully determine the joint probability distribution?
- (c) Express $Pr(q_3|q_2)$ in terms of the conditional probabilities given in your answer to the previous question.

12. [10 points] In this question, you must write a Prolog program that solves the N queens problem, that is, figures out how to place N queens on a N by N rectangular board so that no queen attacks another, i.e., no pair of queens is on the same row, column, or diagonal. We can assume that each row contains exactly one queen. So we can represent a potential solution, a placement of the queens, as a list $[c_1, c_2, \dots, c_N]$ where c_1 is the number of the column that the queen on row 1 is placed, c_2 is the number of the column that the queen on row 2 is placed, \dots , and c_N is the number of the column that the queen on row N is placed. We can assume that columns are numbered from 1 to N. Since we can't have more than one queen in any given column, we can represent a potential solution as a permutation of the list of the positive integers from 1 to N. For example $[2, 4, 1, 3]$ is a solution to the 4 queens problem. We want to implement a solution to this problem using a “generate and test” approach.
- (a) Implement a Prolog predicate **safe(L)** that succeeds if and only if the list L representing a placement of the queens (a permutation of the first N positive integers as discussed above) is *safe*, i.e., no pair of queens is on the same row, column, or diagonal. You may introduce auxiliary predicates, with suitable documentation. You must provide definitions for all non-built-in predicates that you use.

- (b) Implement a Prolog predicate `queens(N,L)` that holds if and only if the list `L` represents a placement of `N` queens (as discussed earlier) that is a solution to the `N` queens problem. You may use the predicate `safe(L)` defined in (a) in your solution. You may also use (without defining it) a predicate `permutation(L,P)` that holds if and only if list `P` is a permutation of the list `L`. You may introduce auxiliary predicates, with suitable documentation. You must provide definitions for all non-built-in predicates that you use.