CMPT 310 (2001-1)
Artificial Intelligence Survey
Instructor: Philip W. L. Fong
Midterm Exam Solution Key

Part I: True-or-False Questions (20%)

1. A pathless search procedure is as powerful as a regular search procedure.

   True.

2. Depth-First Search is incomplete when applied to finite graphs.

   True or False. This question is flawed. You will get 2 marks as long as you attempt it.

3. If $h_1(n)$ and $h_2(n)$ are two admissible heuristics, then $h_1(n) + h_2(n)$ is also admissible.

   False. Although $h_1(n) \leq h^*(n)$ and $h_2(n) \leq h^*(n)$, $h_1(n) + h_2(n)$ could be larger than $h^*(n)$.

4. If a heuristic function $h(n)$ never over-estimates $h^*(n)$, the actual optimal cost from search node n to a goal node, then A* is optimal.

   True.

5. Suppose A* is given a perfectly informed heuristics (i.e. $h(n) = h^*(n)$). Suppose further that the node store orders search nodes with the same f-value in a first-in-first-out (FIFO) manner. Then the number of nodes expanded by A* is linear to the solution depth.

   False. Consider a complete binary tree of depth m so that all leaves are goals. All nodes in this tree have the same f-value. A* will expand all of them.

6. A consistent heuristic is also monotonic.

   True.

7. The Manhattan-distance heuristics is more informed than the relaxed-adjacency heuristics.

   False. Although empirical results tell us that Manhattan-distance heuristics is more powerful than relaxed-adjacency heuristics, this, however, is not true in the general case. Consider the following state in an 8-puzzle problem:

	2	1
3	4	5
6	7	8

Manhattan-distance heuristics gives an estimate of 2, while relaxed-adjacency gives an estimate of 3.

8. The resolution rule is incomplete for propositional CNF formulas.

   True. You won't be able to deduce $p \vee \neg p$ using the resolution rule alone.

9. Let S be a set of clauses containing a tautology clause. It is impossible to deduce the empty clause ($\Box$) from S using only the resolution rule.

   False. The following is an unsatisfiable set of clauses containing a tautology clause: $\{\{p, \neg p\}, \{q\}, \{\neg q\}\}$.

10. Since random walk is in general less informed than hill-climbing, WALKSAT is less likely to find a satisfying truth assignment for a given CNF formula than GSAT.

    False.

Part II: Multiple-Choice Questions (20%)

1. Suppose you are to search for a goal node in a random instance of the following class of search trees: (i) Each search tree has a branching factor of 2 and a maximum depth of 20; (ii) The probability that a search node at level k is a goal node is $1/(k+1)$. On average, which of following search strategies expands the least number of nodes?

   (a) Breadth-First Search

   (b) Depth-First Search

   (c) Depth-Limited Search with depth limit 20

   (d) Iterative-Deepening Search

   Answer: (a). It is more likely to hit a goal node at a shallow level than at a deep level, so we search the shallower level first.

2. Which of the following is **NOT** a sufficient condition for $h(n)$ to be admissible?

   (a) $h(n)$ is consistent, and $h(G) = 0$ for every goal node G.

   (b) $h(n)$ is dominated by another admissible heuristics $h'(n)$.

   (c) $h(n) + g(n)$ is monotonic.

   (d) $h(n) = 0$.

   Answer: (c). If $h(n)$ is dominated by another admissible heuristics, we then have

   $$h(n) \leq h'(n) \leq h^*(n).$$

   So $h(n)$ is admissible. However, the monotonicity of $h(n) + g(n)$ does not imply the admissibility of $h(n)$.

3. Which of the operator preconditions for the sliding-block puzzle problem has/have to be removed in order to derive the relaxed adjacency heuristics?

   (a) A tile can only be moved to an adjacent cell.

   (b) A tile can only be moved to an unoccupied cell.

   (c) Both (a) and (b).

   (d) None of the above.

   Answer: (a).

4. Which of the following modifications to a given search problem does **NOT** result in a relaxed problem with less costly solutions?

   (a) Deleting some preconditions of the operators.

   (b) Deleting some conditions from the goal test.

   (c) Deleting some states from the state space graph.

   (d) Introducing additional operators into the problem.

   Answer: (c). If we delete some preconditions of the operators, we are introducing additional edges into the state space, thereby reducing the cost of a cheapest path from the initial state to a goal. The same effect can be obtained if we introduce additional operators into the problem. Similarly, if we relax the goal test, then there will be more goal nodes, and the cost of a cheapest path from the initial state to a goal state will be reduced. However, if we delete nodes from the state space, than the cost of a cheapest solution path will actually increase.

5. In which of the following aspects is IDA* more superior to A*?

   (a) Completeness.

   (b) Optimality.

   (c) Asymptotic time complexity.

   (d) Asymptotic space complexity.

   Answer: (d).

6. Under which of the following conditions will Hill-Climbing Search fail to return a state with maximum f-value?

   (a) Every state has the same f-value.

   (b) Every state has a different f-value, and there is only one local maxima.

   (c) There are only two possible f-values, 0 and 1, and only one state has f-value equal to 0.

(d) There are only three possible f-values, 0, 1 and 2, only one state has f-value equal to 1, and only one has f-value equal to 2.

Answer: (d). There are two local maximas if condition (d) holds.

7. Which of the following properties does **NOT** apply to a semantic tree of an *unsatisfiable* CNF formula P?

   (a) If the semantic tree has more than one node, then at least one leaf node has a failure node sibling.

   (b) There exists a leaf node that is not a superset of some clauses in P.

   (c) If P does not contain the empty clause ($\Box$), then every non-leaf node is either an inference node or an ancestor of an inference node.

   (d) If the semantic tree has no inference node, then P contains an empty clause.

   Answer: (b). Property (a) is true by definition of an inference node. Property (c) is true because the number of failure nodes in any subtree of a closed semantic tree is one plus the number of non-failure nodes in that subtree, and thus every non-trivial subtree contains a failure node. Property (d) is true since only a semantic tree with one node could contain no inference node.

8. Which of the following satisfiability testing procedures is **no good** for constructing refutation proofs of propositional sentences?

   (a) Truth-table method

   (b) Resolution method

   (c) Davis-Putnam procedure

   (d) WALKSAT procedure

   Answer: (d). WALKSAT is incomplete.

9. Which of the following interpretations is **NOT** a model of $\forall x.\exists y.p(x,y)$?

   (a) The domain is the set of states in the 15-puzzle problem, and $p(x,y)$ means that "x can be transformed to y in one move."

   (b) The domain is $\aleph$, the set of natural numbers, and $p(x,y)$ is interpreted as "$x < y$."

   (c) The domain is $\aleph$, the set of natural numbers, and $p(x,y)$ means that "$x > y$."

   (d) The domain is $\Re$, the set of real numbers, and $p(x,y)$ means "The absolute value of y is smaller than the absolute value of x."

   Answer: (c or d). There is no natural number smaller than 1, and there is no real number with absolute value smaller than 0. This question is flawed. You will get 2 marks if you answer either (c) or (d).

10. Suppose a planning problem is to be formulated as a situation calculus with A actions and F fluents. What is the total number of precondition axioms and successor-state axioms?

   (a) $F + A$

   (b) $2 \times F + A$

   (c) $F \times A$

   (d) $2 \times F \times A$

   Answer: (b).

Part III: Essay Questions (60%)

1. [20%] The Tower of Hanoi problem can be stated as follows:

   > There are n disks $D_1, D_2, \ldots, D_n$, of increasing sizes, and three pegs 1, 2, and 3. Initially all the disks are stacked on peg 1, with D_1, the smallest, on top and D_n, the largest, at the bottom. The goal is to transfer all the disks from peg 1 to peg 3 given that only one disk can be moved at a time. Specifically, the top disk D_k of a stack on peg i can be moved to the top of the stack on peg j if D_k is smaller than the top disk of the stack on peg j.

   (a) State the three steps involved in the systematic derivation of admissible heuristics by problem relaxation.

   Answer: A disciplined methodology for designing admissible heuristic functions:

   I/ [2%] Relax the given problem by *precondition deletion*, thereby obtaining a more "abstract" problem.

   II/ [2%] Find an *efficient* way of solving the relaxed problem.

   III/ [2%] Use the *solution cost* of the relaxed problem as a *heuristic estimate* for the solution cost of the original problem.

   Heuristic functions obtained by such process is *guaranteed* to be admissible.

   (b) Enumerate the preconditions and effects of moving a disk from peg i to peg j.

   Answer: The preconditions and effects of moving D_k from peg i to peg j are the following:

   Preconditions:

   - [2%] D_k has to be a top disk.
   - [2%] D_k must be smaller than the top disk on peg j.

   Effects:

   - [2%] D_k becomes the top disk on peg j.

(c) A disk D_i is considered "*misplaced*" if it is not on peg 3 or if some D_j with $j > i$ is misplaced. The *misplaced-disks* heuristics uses the following formula as a heuristic estimate $h(n)$:

$$(\# \text{ of misplaced disks not on peg 3}) + 2 \times (\# \text{ of misplaced disks on peg 3})$$

Demonstrate the admissibility of this heuristics by problem relaxation.

Answer: [8%] We follow the three steps outlined in part (a).

I/ If we delete both of the preconditions listed in part (b), then we obtain a relaxed problem in which any disk on peg i can be moved to the *top* of the stack on peg j.

II/ The optimal solution for this relaxed problem is to (1) move all misplaced disks *on peg 3* to another peg, and then (2) move all misplaced disks back to peg 3, in decreasing order of disk size.

III/ The misplaced disks not on peg 3 are moved exactly once in (2), and the misplaced disks on peg 3 are moved away from peg 3 in (1) and then back on peg 3 in (2). Therefore, the solution cost is:

$$(\# \text{ of misplaced disks not on peg 3}){+}2{\times}(\# \text{ of misplaced disks on peg 3})$$

If we use this as a heuristic estimate, we obtain the *misplaced disks* heuristics. By (a), this is guaranteed to be admissible.

2. [25%] Given the following premises:

$\neg P \Rightarrow (Q \vee R)$
$Q \Rightarrow S$
$R \Rightarrow S$
$S \Rightarrow \neg(R \vee Q)$

use the Davis-Putnam procedure to construct a refutation proof that P follows from the premises.

Answer: (a) [3%] We attempt to refute the following sentence:

$$(\neg P \Rightarrow (Q \vee R)) \wedge (Q \Rightarrow S) \wedge (R \Rightarrow S) \wedge (S \Rightarrow \neg(R \vee Q)) \wedge (\neg P)$$

(b) [12%] We convert the above sentence to CNF:

$$\begin{aligned}
& (\neg P \Rightarrow (Q \vee R)) \wedge (Q \Rightarrow S) \wedge (R \Rightarrow S) \wedge (S \Rightarrow \neg(R \vee Q)) \wedge (\neg P) \\
\longrightarrow\ & (\neg\neg P \vee (Q \vee R)) \wedge (\neg Q \vee S) \wedge (\neg R \vee S) \wedge (\neg S \vee \neg(R \vee Q)) \wedge (\neg P) \\
\longrightarrow\ & (P \vee Q \vee R) \wedge (\neg Q \vee S) \wedge (\neg R \vee S) \wedge (\neg S \vee (\neg R \wedge \neg Q)) \wedge (\neg P) \\
\longrightarrow\ & (P \vee Q \vee R) \wedge (\neg Q \vee S) \wedge (\neg R \vee S) \wedge (\neg S \vee \neg R) \wedge (\neg S \vee \neg Q) \wedge (\neg P) \\
\longrightarrow\ & \{\{P, Q, R\}, \{\neg Q, S\}, \{\neg R, S\}, \{\neg S, \neg R\}, \{\neg S, \neg Q\}, \{\neg P\}\}
\end{aligned}$$

(c) [10%] We apply Davis-Putnam procedure to demonstrate that the CNF formula is unsatisfiable:

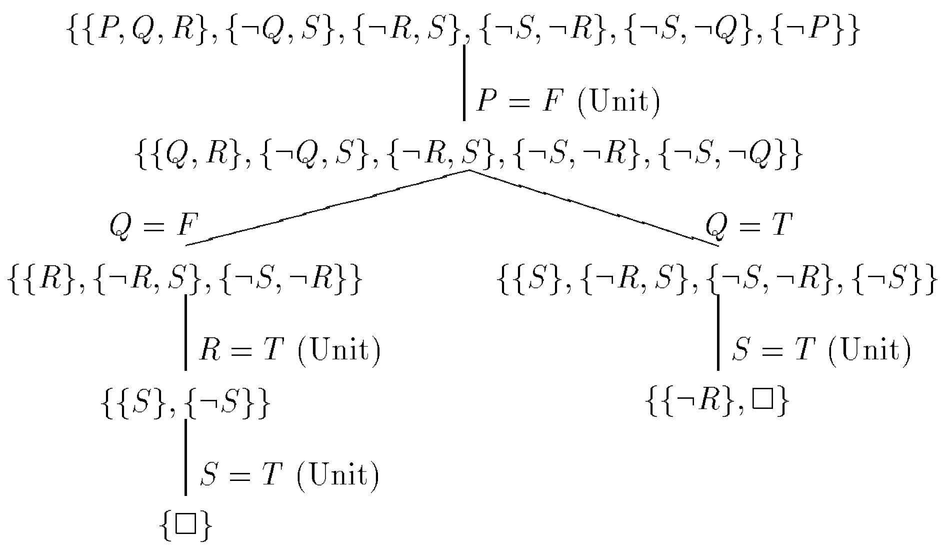

3. [15%] Answer **ONLY ONE** of the following questions.

(a) Design a search tree of branching factor 2 and maximum depth m so that Dijkstra's Uniform Cost Search expands exponentially many more nodes than Breadth-First Search. Give a **ROUGH** estimate of the number of nodes expanded by each of the search strategies using the Big-O notation. You **DO NOT** need to perform a detail time-complexity analysis.

Answer: [12%] We define the following search tree with height m:

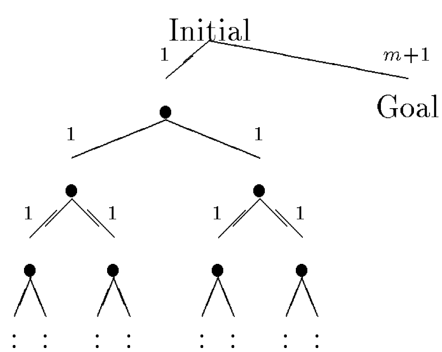

[3%] Dijkstra's algorithm and BFS will expand $O(2^m)$ and $O(1)$ nodes respectively.

(b) Consider the following puzzle. A garden is represented by an $N \times N$ grid of square cells. Rose bushes are to be planted in the garden so that the following constraints are observed:

i. The cells on the border of the grid are occupied by fences, and no roses should be planted there.

ii. For aesthetic reasons, roses are planted in clusters. Specifically, if roses are planted in a cell, then roses should be planted in at least one of the four neighboring cells.

iii. To avoid overcrowdedness, if roses are planted in a cell, then at least one of the four neighboring cells should be free of roses.

iv. Roses should be planted in at least M of the cells.

Reduce the above problem into an instance of propositional CNF satisfiability problem. You **DO NOT** need to justify your answer.

Answer: [3%] Let $R_{i,j}$ represent the fact that roses are planted in the call at row i and column j, for $1 \leq i, j \leq N$.

i. [3%] $(\neg R_{i,j})$ for $i = 1$ or $j = 1$ or $i = N$ or $j = N$.

ii. [3%] $(\neg R_{i,j} \vee R_{i-1,j} \vee R_{i+1,j} \vee R_{i,j-1} \vee R_{i,j+1})$ for $1 < i, j < N$.

iii. [3%] $(\neg R_{i,j} \vee \neg R_{i-1,j} \vee \neg R_{i+1,j} \vee \neg R_{i,j-1} \vee \neg R_{i,j+1})$ for $1 < i, j < N$.

iv. [3%] $(R_{i_1,j_1} \vee R_{i_2,j_2} \vee \ldots \vee R_{i_{(N^2-M+1)},j_{(N^2-M+1)}})$ for every size-$(N^2 - M + 1)$ subset $\{\langle i_1, j_1\rangle, \langle i_2, j_2\rangle, \ldots \langle i_{(N^2-M+1)}, j_{(N^2-M+1)}\rangle\}$ of cell coordinates.