

CS-171, Intro to A.I. — Mid-term Exam — Winter Quarter, 2015

YOUR NAME: _____

YOUR ID: _____ ID TO RIGHT: _____ ROW: _____ SEAT: _____

The exam will begin on the next page. Please, do not turn the page until told.

When you are told to begin the exam, please check first to make sure that you have all ten pages, as numbered 1-10 in the bottom-right corner of each page. We wish to avoid copy problems. We will supply a new exam for any copy problems.

The exam is closed-notes, closed-book. No calculators, cell phones, electronics.

Please turn off all cell phones now.

Please clear your desk entirely, except for pen, pencil, eraser, a blank piece of paper (for scratch pad use), and an optional water bottle. Please write your name and ID# on the blank piece of paper and turn it in with your exam.

This page summarizes the points for each question, so you can plan your time.

1. (14 pts total, 2 pts each) SUDOKU AS A CONSTRAINT SATISFACTION PROBLEM.
2. (8 pts total) ONE FISH, TWO FISH, RED FISH, BLUE FISH. (With apologies to Dr. Seuss.)
3. (5 points total, -1 pt each error, but not negative) CONSTRAINT GRAPH CONSTRUCTION.
4. (12 pts total, 2 pts each) VALID, UNSATISFIABLE, SATISFIABLE.
5. (18 pts total, 1 pt each) GAME SEARCH WITH TIC-TAC-TOE AND WIN-PATHS HEURISTIC FUNCTION.
6. (15 pts total, 3 pts each) STATE SPACE SEARCH.
7. (8 pts total, 1 pt each) LOCAL SEARCH --- SIMULATED ANNEALING.
8. (10 pts total, 1 pt each) ADVERSARIAL (GAME) SEARCH CONCEPTS.
9. (10 pts total, 1 pt each) CONSTRAINT SATISFACTION PROBLEM (CSP) CONCEPTS.

The Exam is printed on both sides to save trees! Work both sides of each page!

1. (14 pts total, 2 pts each) SUDOKU AS A CONSTRAINT SATISFACTION PROBLEM.

A Sudoku board consists of $n \times n$ squares, some of which are initially filled with digits from 1 to n . The objective of the Sudoku puzzle is to fill in all the remaining squares such that no digit appears twice in any row, column, or $\sqrt{n} \times \sqrt{n}$ box. **Consider the case in which $n = 4$.**

The Sudoku puzzle with $n = 4$ can be formulated as a constraint satisfaction problem with $n^2 = 16$ variables, where every variable stands for one of the 16 squares on the 4x4 Sudoku board. If we denote rows with letters A-D and columns with numbers 1-4, we end up with 16 variables: A1, A2, A3, A4, B1, B2, ..., D3, D4. Each variable has domain values $\{1, 2, 3, 4\}$.

What remains to be done is to specify the constraints of the problem. We could use binary constraints for each pair of variables, but might end up with a large number of constraints. A better alternative is to use global *AllDiff* constraints, where *AllDiff*(X_1, X_2, \dots, X_n) means that all n variables, X_1, X_2, \dots, X_n , must have different values.

1.a. (2 pts) How many *AllDiff* constraints are required to specify that no digit may appear twice in any row, column, or 2×2 box? _____

	1	2	3	4
A	1 2 3 4	1 2 3 4	1 2 3 4	1 2 3 4
B	1 2 3 4	1 2 3 4	3	1 2 3 4
C	1 2 3 4	1 2 3 4	1 2 3 4	1 2 3 4
D	1 2 3 4	1 2 3 4	1 2 3 4	1 2 3 4

1.b. (2 pts) FORWARD CHECKING.

Consider the 4×4 Sudoku board on the left. Possible domain values for each variable are shown in each square. Variable B3 has just been assigned value 3 (circled).

Cross out all values from the domains of the remaining unassigned variables that now would be eliminated by Forward Checking.

	1	2	3	4
A	1 2 3 4	1 2 3 4	1 2 3 4	1 2 3 4
B	1 2 3 4	1 2 3 4	3	1 2 3 4
C	1 2 3 4	2	1 2 3 4	1 2 3 4
D	1 2 3 4	1 2 3 4	1 2 3 4	4

1.c. (2 pts) ARC CONSISTENCY.

Consider the 4×4 Sudoku board on the left. Possible domain values for each variable are shown in each square. Variables B3, C2, and D4 have been assigned values, but no constraint propagation has been done.

Cross out all values from the domains of the remaining unassigned variables that now would be eliminated by Arc Consistency (AC-3 in your book).

****** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ******

	1	2	3	4
A	3	1 2 X 4	1 2 X 4	1 2 X X
B	1 2 X 4	1 2 X 4	1 2 3 4	1 2 3 X
C	1 2 X 4	1 2 3 4	1 2 3 X	1 2 3 X
D	1 2 X X	1 2 3 X	1 2 3 X	4

1.d. (2 pts) MINIMUM-REMAINING-VALUES (MRV) HEURISTIC.

Consider the 4×4 Sudoku board on the left. Possible domain values for each variable are shown in each square. Variables A1 and D4 are assigned and constraint propagation has been done.

List all unassigned variables (in any order) that now might be selected by the Minimum-Remaining-Values (MRV) Heuristic: _____.

	1	2	3	4
A	1 2 X 4	1 2 X 4	1 2 3 4	1 2 3 X
B	3	1 2 X 4	1 2 X 4	1 2 X X
C	1 2 X 4	3	1 2 X 4	1 2 X X
D	1 2 X X	1 2 X X	1 2 3 X	4

1.e. (2 pts) LEAST-CONSTRAINING-VALUE (LCV) HEURISTIC.

Consider the 4×4 Sudoku board on the left. Possible domain values for each variable are shown in each square. Variables B1, C2, and D4 are assigned and constraint propagation has been done. A3 (circled) is chosen for assignment next.

List all values (in any order) for variable A3 (circled) that now might be selected by the Least-Constraining-Value (LCV) Heuristic: _____.

	1	2	3	4
A	X 2 X 4	X 2 X 4	1	X 2 3 4
B	1 2 X 4	3	X 2 X 4	1 2 X 4
C	1 2 X 4	1 2 X 4	3	1 2 X 4
D	1 2 3 4	1 2 X 4	X 2 X 4	1 2 X 4

1.f. (2 pts) DEGREE HEURISTIC (DH).

Consider the 4×4 Sudoku board on the left. Possible domain values for each variable are shown in each square. Variables A3, B2, and C3 are assigned and constraint propagation has been done.

List all unassigned variables (in any order) that now might be selected by the Degree Heuristic (ignore MRV for this problem): _____.

	1	2	3	4
A	2	4	1	3
B	3	1	4	2
C	1	2	3	2
D	4	3	?	1

1.g. (2 pts) MIN-CONFLICTS HEURISTIC.

Consider the 4×4 Sudoku board on the left showing a complete but inconsistent assignment. Current values for each variable are shown in each square. Variable D3 (circled) has just been selected to be assigned a new value during local search for a complete and consistent assignment.

What new value for variable D3 (circled) would be chosen now by the Min-Conflicts Heuristic? _____.

2. (8 pts total) ONE FISH, TWO FISH, RED FISH, BLUE FISH. (With apologies to Dr. Seuss.)

Amy, Betty, Cindy, and Diane went out to lunch at a seafood restaurant. Each ordered one fish. Each fish was either a red fish or a blue fish. **Among them they had exactly two red fish and two blue fish.**

You translate this fact into Propositional Logic (in prefix form) as:

/* Ontology: Symbol A/B/C/D means that Amy/Betty/Cindy/Diane had a red fish. */
 (or (and A B (¬ C) (¬ D)) (and A (¬ B) C (¬ D))
 (and A (¬ B) (¬ C) D) (and (¬ A) B C (¬ D))
 (and (¬ A) B (¬ C) D) (and (¬ A) (¬ B) C D)))

Their waiter reported:

“Amy, Betty, and Cindy had exactly one red fish among them; I don’t remember who had what. Betty, Cindy, and Diane had exactly one red fish among them; I don’t remember who had what.”

You translate these facts into Propositional Logic (in prefix form) as:

(or (and A (¬ B) (¬ C)) (and (¬ A) B (¬ C)) (and (¬ A) (¬ B) C))
 (or (and B (¬ C) (¬ D)) (and (¬ B) C (¬ D)) (and (¬ B) (¬ C) D))

Betty’s daughter asked, “Is it true that my mother had a blue fish?”

You translate this query into Propositional Logic as “(¬ B)” and form the negated goal as “(B)”.

Your resulting knowledge base (KB) plus the negated goal (in CNF clausal form) is:

(A B C) (¬ A) (¬ B) (¬ C))
 (A B D) (¬ A) (¬ B) (¬ D))
 (A C D) (¬ A) (¬ C) (¬ D))
 (B C D) (¬ B) (¬ C) (¬ D))
 ((¬ A) (¬ B)) ((¬ A) (¬ C))
 ((¬ B) (¬ C)) ((¬ B) (¬ D))
 ((¬ C) (¬ D)) (B)

Write a resolution proof that Betty had a blue fish.

For each step of the proof, fill in the first two blanks with CNF sentences from KB that will resolve to produce the CNF result that you write in the third (resolvent) blank. The resolvent is the result of resolving the first two sentences. Add the resolvent to KB, and repeat. Use as many steps as necessary, ending with the empty clause. The empty clause indicates a contradiction, and therefore that KB entails the original goal sentence.

The shortest proof I know of is only four lines long. (A Bonus Point is offered for a shorter proof.)

Longer proofs are OK provided they are correct. Think about it, then find a proof that mirrors how you think.

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

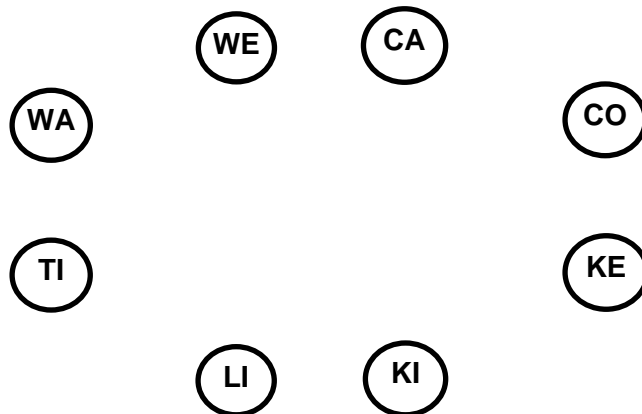

Resolve _____ with _____ to produce: _____

**** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ****

3. (5 points total, -1 pt each error, but not negative) CONSTRAINT GRAPH CONSTRUCTION.

You are a map-coloring robot assigned to color this map of southern Ireland counties. Adjacent regions must be colored a different color (R=Red, B=Blue, G=Green). Please note that all borders between regions are marked by a thin line to highlight the border. This is only for visual clarity, and is not part of any region.

Draw the constraint graph. The nodes are provided. You need draw only the arcs.



4. (12 pts total, 2 pts each) VALID, UNSATISFIABLE, SATISFIABLE.

For each sentence below, write V=valid if the sentence is valid, U=unsatisfiable if the sentence is unsatisfiable, and S=satisfiable if the sentence is satisfiable but is neither valid nor unsatisfiable.

- 4.a. (2 pts) (write V, U, or S) _____ (A OR (\neg A))
- 4.b. (2 pts) (write V, U, or S) _____ (A AND (\neg A))
- 4.c. (2 pts) (write V, U, or S) _____ (A AND (\neg B))
- 4.d. (2 pts) (write V, U, or S) _____ ((A AND (\neg A)) => B)
- 4.e. (2 pts) (write V, U, or S) _____ ((A OR (\neg A)) => B)
- 4.f. (2 pts) (write V, U, or S) _____ ((A OR (\neg A)) => (A AND (\neg A)))

5. (18 pts total, 1 pt each) GAME SEARCH WITH TIC-TAC-TOE AND WIN-PATHS HEURISTIC FUNCTION.

This problem asks about MiniMax Search and Alpha-Beta pruning in Tic-Tac-Toe with the Win-paths static heuristic evaluation function. Recall that the Win-paths heuristic function counts the number of possible win-paths for MAX (= X) and subtracts the number of possible win-paths for MIN (= O). For example:

<table border="1"> <tr><td>X</td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td>O</td><td></td></tr> </table>	X							O		<table border="1"> <tr><td></td><td></td><td></td></tr> <tr><td>X</td><td>O</td><td></td></tr> <tr><td></td><td></td><td></td></tr> </table>				X	O					<table border="1"> <tr><td>X</td><td></td><td></td></tr> <tr><td></td><td></td><td></td></tr> <tr><td></td><td></td><td>O</td></tr> </table>	X								O
X																													
	O																												
X	O																												
X																													
		O																											
MAX (= X) # win paths=6 MIN (= O) # win paths=5 Heuristic value = 1 (= 6-5)	MAX (= X) # win paths=4 MIN (= O) # win paths=6 Heuristic value = -2 (= 4-6)	MAX (= X) # win paths=5 MIN (= O) # win paths=5 Heuristic value = 0 (= 5-5)																											


5.a. (6 pts total, 1 pt each blank branch node [4 in part 1] or answer space [1 each in parts 2&3])

In the game tree below it is **Max's** (= X's) turn to move. At each leaf node is the estimated score of that resulting position as returned by the Win-path heuristic static evaluator (written below as "val=n").

(1) Perform Mini-Max search and label each branch node with its return value (4 branch nodes).

(2) What is Max's best move (A, B, or C)? _____


(3) What value does Max expect to get? _____



5.b. (12 pts total, 1 pt each leaf node)

In the game tree below it is **Max's** (= X's) turn to move (this is the same game tree as in problem 5.a above). At each leaf node is the estimated score of that resulting position as returned by the Win-path heuristic static evaluator (as "val=n"). You do not need to indicate the branch node return values again.

Cross out each leaf value that would be pruned by Alpha-Beta Pruning.




**** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ****

6. (15 pts total, 3 pts each) STATE SPACE SEARCH.

Execute Tree Search through this graph (i.e., do not remember visited nodes). Step costs are given next to each arc. Heuristic values are given next to each node (as $h=x$). The successors of each node are indicated by the arrows out of that node. **Successors are returned in left-to-right order (successors of A are (A, G), and of B are (C, B), in that order).**

For each search strategy below, show the order in which nodes are expanded (i.e., to expand a node means that its children are generated), ending with the goal node that is found. Show the path from start to goal, or write "None". Give the cost of the path found. The first one is done for you as an example.



6.a. (example) DEPTH FIRST SEARCH.

Order of node expansion: S A G

Path found: S A G Cost of path found: 17

6.b. (3 pts) BREADTH FIRST SEARCH.

Order of node expansion: _____

Path found: _____ Cost of path found: _____

6.b. (3 pts) UNIFORM COST SEARCH.

Order of node expansion: _____

Path found: _____ Cost of path found: _____

6.c. (3 pts) GREEDY (BEST-FIRST) SEARCH.

Order of node expansion: _____

Path found: _____ Cost of path found: _____

6.d. (3 pts) ITERATED DEEPENING SEARCH.

Order of node expansion: _____

Path found: _____ Cost of path found: _____

6.e. (3 pts) A* SEARCH.

Order of node expansion: _____

Path found: _____ Cost of path found: _____

Is the heuristic admissible (Yes or No)? _____

7. (8 pts total, 1 pt each) LOCAL SEARCH --- SIMULATED ANNEALING.

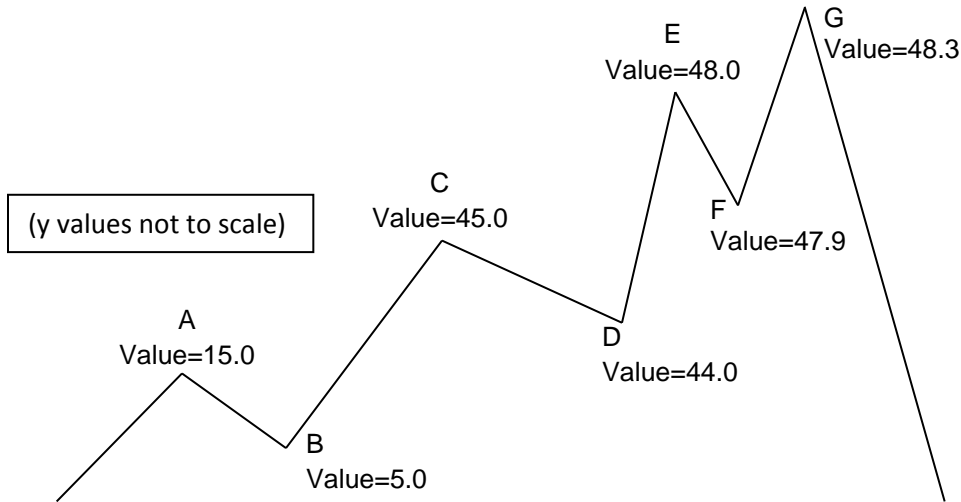
This question asks about Simulated Annealing local search. In the value landscape cartoon below, you will be asked about the probability that various moves will be accepted at different temperatures. Recall that Simulated Annealing always accepts a better move ($\Delta\text{Value} = \text{Value}[\text{next}] - \text{Value}[\text{current}] > 0.0$); but it accepts a worse move ($\Delta\text{Value} < 0.0$) only with probability $e^{(\Delta\text{Value}/T)}$, where T is the current temperature on the temperature schedule.

Please use this temperature schedule (usually, it is a decaying exponential; but it is simplified here):

time (t)	1-100	101-200	201-300
Temperature (T)	10.0	1.0	0.1

You do not need a calculator; the values given have been chosen to follow this table:

x	0.0	-0.1	-0.4	-1.0	-4.0	-40.0
e^x	1.00	≈0.90	≈0.67	≈0.37	≈0.02	≈4.0e-18



Give your answer to two significant decimal places. The first one is done for you as an example.

7.a. (example) You are at Point A and t=23. The probability you will accept a move A -> B = 0.37

7.b. (1 pt) You are at Point B and t=23. The probability you will accept a move B -> C = _____

7.c. (1 pt) You are at Point C and t=123. The probability you will accept a move C -> B = _____

7.d. (1 pt) You are at Point C and t=123. The probability you will accept a move C -> D = _____

7.e. (1 pt) You are at Point E and t=123. The probability you will accept a move E -> D = _____

7.f. (1 pt) You are at Point E and t=123. The probability you will accept a move E -> F = _____

7.g. (1 pt) You are at Point G and t=123. The probability you will accept a move G -> F = _____

7.h. (1 pt) You are at Point G and t=223. The probability you will accept a move G -> F = _____

7.i. (1 pt) With a very, very, very long slow annealing schedule, are you more likely, eventually in the long run, to wind up at point A or at point G? (write A or G) _____.

**** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ****

8. (10 pts total, 1 pt each) ADVERSARIAL (GAME) SEARCH CONCEPTS.

For each of the following terms on the left, write in the letter corresponding to the best answer or the correct definition on the right.

Game Strategy	A	Approximates the value of a game state (i.e., of a game position)
Cut-off Test	B	In all game instances, total pay-off summed over all players is a constant
Alpha-Beta Pruning	C	Tree where nodes are game states and edges are game moves
Weighted Linear Function	D	Function that specifies a player's move in every possible game state
Terminal Test	E	Returns same move as MiniMax, but may prune more branches
ExpectiMiniMax	F	Optimal strategy for 2-player zero-sum games of perfect information, but impractical given limited time to make each move
Game Tree	G	Vector dot product of a weight vector and a state feature vector
Heuristic Evaluation Function	H	Function that decides when to stop exploring this search branch
Zero-sum Game	I	Generalizes MiniMax to apply to games with chance (stochastic games)
MiniMax Algorithm	J	Function that says when the game is over

9. (10 pts total, 1 pt each) CONSTRAINT SATISFACTION PROBLEM (CSP) CONCEPTS.

For each of the following terms on the left, write in the letter corresponding to the best answer or the correct definition on the right.

Minimum Remaining Values Heuristic	A	Set of allowed values for some variable
Degree Heuristic	B	Specifies the allowable combinations of variable values
Min-Conflicts Heuristic	C	Every variable is associated with a value
Solution to a CSP	D	The values assigned to variables do not violate any constraints
Least Constraining Value Heuristic	E	A complete and consistent assignment
Domain	F	Nodes correspond to variables, links connect variables that participate in a constraint
Constraint	G	Chooses the next variable to expand to have the fewest legal values in its domain
Consistent Assignment	H	Chooses the next variable to expand to have the largest number of constraints on other unassigned variables
Complete Assignment	I	Prefers to search next the value that rules out the fewest choices for the neighboring variables in the constraint graph
Constraint Graph	J	Select the value that results in fewest conflicts with other variables

**** THIS IS THE END OF THE MID-TERM EXAM ****