CS-171, Intro to A.I. — Mid-term Exam — Fall Quarter, 2018

NAME:____ UCI NetID: ______________________

YOUR ID#: _______________ ID# TO RIGHT: ________________ ROW:______ SEAT:______

Please turn off all cell phones now.

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

When told to begin, check first to ensure that your copy has all the pages, as numbered 1-12 in the bottom-right corner of each page. We will supply a new exam for any copy problems.

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

Clear your desk except for pen, pencil, eraser, & water bottle. Put backpacks under your seat. Please do not detach the provided scratch paper from the exam.

After you first stand up from your seat, your exam is over and must be turned in immediately. You may turn in your Midterm exam early and leave class when you are finished.

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

1. (16 pts total, 2 pts each) A* SEARCH AND HEURISTICS.

2. (10 pts total, 2 pts each) STATE-SPACE SEARCH.

3. (14 pts total, 2 pt each) The heuristic function f(n) = g(n)+h(n) and its components.

4. (10 pts total, 2 pts each) LOCAL SEARCH STATEMENTS.

5. (5 pts total) AGENT.

6. (14 pts total, 2 pt each) LOCAL SEARCH ACTIONS.

7. (10 pts total) ITERATED DEEPENING SEARCH (IDS).

8. (11 pts total) PROBABILITY.

9. (10 pts total) BAYESIAN NETWORKS.

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

1. (16 pts total, 2 pts each) A* SEARCH AND HEURISTICS. Consider the following graph. The initial state is S and the goal state is G. The numbers next to each arc give the step cost for traversing that arc.

The table under the graph defines two different heuristic functions, h1() and h2(). **<u>These are NOT the h1() and h2() that we used in the sliding tile puzzle.</u>** Here, h1() and h2() are defined only by the table below.

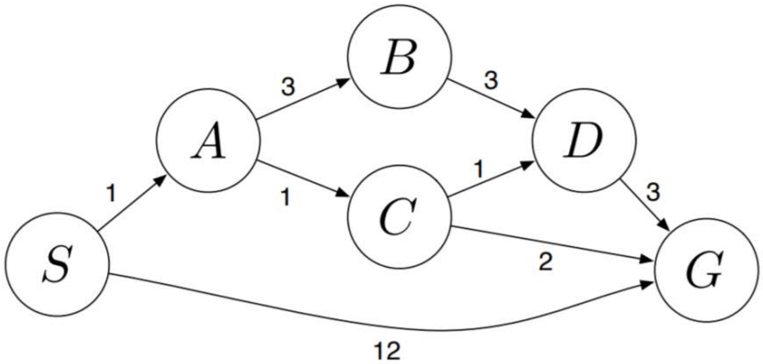

State n	**S**	**A**	**B**	**C**	**D**	**G**
h1(n)	4	3	6	2	3	0
h2(n)	4	2	5	2	2	0

1.a. (2 pts) Does h1() dominate h2()? Answer Y (= Yes) or N (= No) ____Y_____

1.b. (2 pts) Does h2() dominate h1()? Answer Y (= Yes) or N (= No) ____N_____

1.c. (2 pts) Give the tree search order of node expansion for A* with h1(). S A C (G)______________________

1.d. (2 pts) Give the tree search order of node expansion for A* with h2(). S A C (G)______________________

1.e. (2 pts) Is h1() admissible? Answer Y (= Yes) or N (= No) ____Y_____

1.f. (2 pts) Is h1() consistent? Answer Y (= Yes) or N (= No) _____Y_____

1.g. (2 pts) Is h2() admissible? Answer Y (= Yes) or N (= No) ____Y_____

1.h. (2 pts) Is h2() consistent? Answer Y (= Yes) or N (= No) ____N_____

2. (10 pts total, 2 pts each) STATE-SPACE SEARCH. Label the following statements as T (true) or F (false).

2.a. (2 pts) __T_ Breadth-First Search guarantees to find the shallowest goal node, if a goal node exists.

2.b. (2 pts) __F_ Depth-First search guarantees to find the cheapest path to a goal node.

2.c. (2 pts) __T_ Greedy Best First is a type of heuristic search.

2.d. (2 pts) __F_ Both Uniform Cost Search and Iterative Deepening Search use priority queues to order the frontier nodes.

2.e. (2 pts) __T_ The heuristic function is a way to estimate cost of the optimal path from the current state to a goal state.

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

3. (14 pts total, 2 pts each) The heuristic function f(n) = g(n)+h(n) and its components.

3.a. (8 pts total, 2 pts each) Fill in each blank with the letter of the best phrase or concept.

. C .	f(n)	A	Estimate of optimal cost from n to goal
D	g(n)	B	True optimal cost from n to goal (unknown to agent)
A	h(n)	C	Estimate of optimal cost from initial state to goal through n
B	h*(n)	D	Known path cost so far to reach n

3.b. (6 pts total, 2 pts each) Fill in each blank with the letter of the search that would result if the function specified on the left were to be the sort function of the priority queue used by the search.

. B .	f(n)	A	Greedy Best First Search
C	g(n)	B	A* Search
A	h(n)	C	Uniform Cost Search

4. (10 pts total, 2 pts each) LOCAL SEARCH STATEMENTS.
Label the following statements as T (true) or F (false).

4.a. (2 pts) ___F______ Local Search guarantees to find a global optima, i.e., it is an optimal search method.

4.b. (2 pts) ___T______ Local Search suffers from problems such as local minima/maxima, plateaus, and ridges.

4.c. (2 pts) ___F______ Hill-Climbing can escape from local minima/maxima but Simulated Annealing cannot.

4.d. (2 pts) ___T______ Local Search keeps one or a few states, and so has bounded controllable memory needs.

4.e. (2 pts) _________ [Accidentally omitted. Everyone automatically gets full credit.]

5. (5 pts total) AGENT. The textbook defines an agent as something that perceives and acts in an environment. Consider the PEAS acronym in the task environment specification when answering the questions below.

5.a. (4 pts total, 2 pts each) Consider the figure below. Define A, B to show how an agent interacts with an environment.

A ______Sensors__________________ B_______Actuators__________________

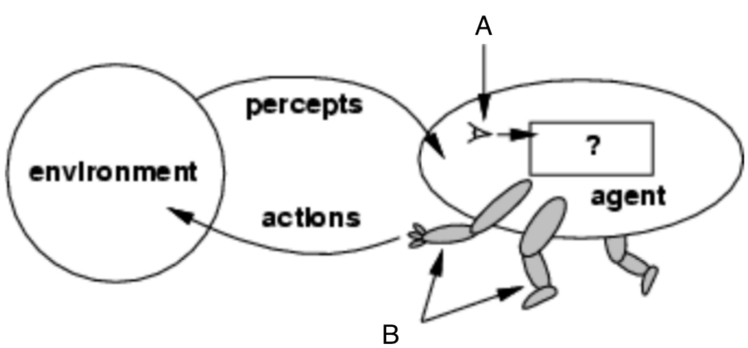

5.b. (1 pt) A rational agent acts so as to maximize the expected value of its ___performance (measure)______, based on the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

6. (14 pts total, 2 pts each) LOCAL SEARCH ACTIONS. Consider the crossword puzzle shown in (A).

(A) Blank Puzzle

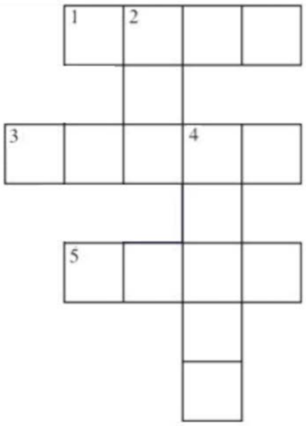

(B) Random Initial State

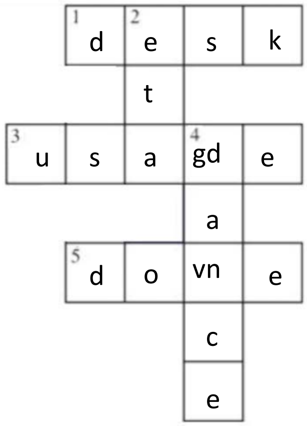

Each number N corresponds to a position where a word is supposed to go. Wherever two words cross each other, the letter in the intersection square of the two words must be the same letter in each word. Each number N is also associated with a variable XN. The numbers 1 to 5 correspond to variables X1, X2, X3, X4, and X5. Note: The same word may be used multiple times in the same state. Variable bindings are independent.

6.a. (10 pts total, 2 pts each) Suppose that you create a random initial state by assigning X1=desk, X2=eta, X3=usage, X4=dance, and X5=dove. **Fill in the template** for (B) Random Initial State using these assignments. If two letters from two words conflict in the square where the words intersect, write both letters in that square.

6.b. (2 pts) Your objective function for the cost of a state is the number of pairs of letters that conflict where words intersect, and so lower values are better. **What is the total cost of the resulting (B) Random Initial State from 6.a. using those assignments?** (You will get full credit if your answer to 6.b is correct given your answer to 6.a, even if your answer to 6.a was wrong.)

2

6.c. (2 pts) Let the domain of each variable XN be { dance, desk, dives, dove, eta, given, hats, kind, usage }. Your available actions are to select a variable and replace its current binding with a new value from its domain.

What possible one-action steps might reduce the cost? Write "None" if no one action can reduce the cost. Otherwise, list all one-action steps that would reduce the cost in one action. Use the following format: (variable=new-word, resulting-cost). (You will get full credit if your answer to 6.c is correct given your answer to 6.a, even if your answer to 6.a was wrong.)

(V4=given, 0) (V4=dives, 1) (V5=kind, 1)

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

7. (10 pts total) ITERATED DEEPENING SEARCH (IDS). Consider the search space below. S is the start node and G is the goal node. Children are returned in left-to-right order. For example, the children of B are E and F in that order, and the children of D are I and J in that order. Perform Iterated Deepening Search (IDS).

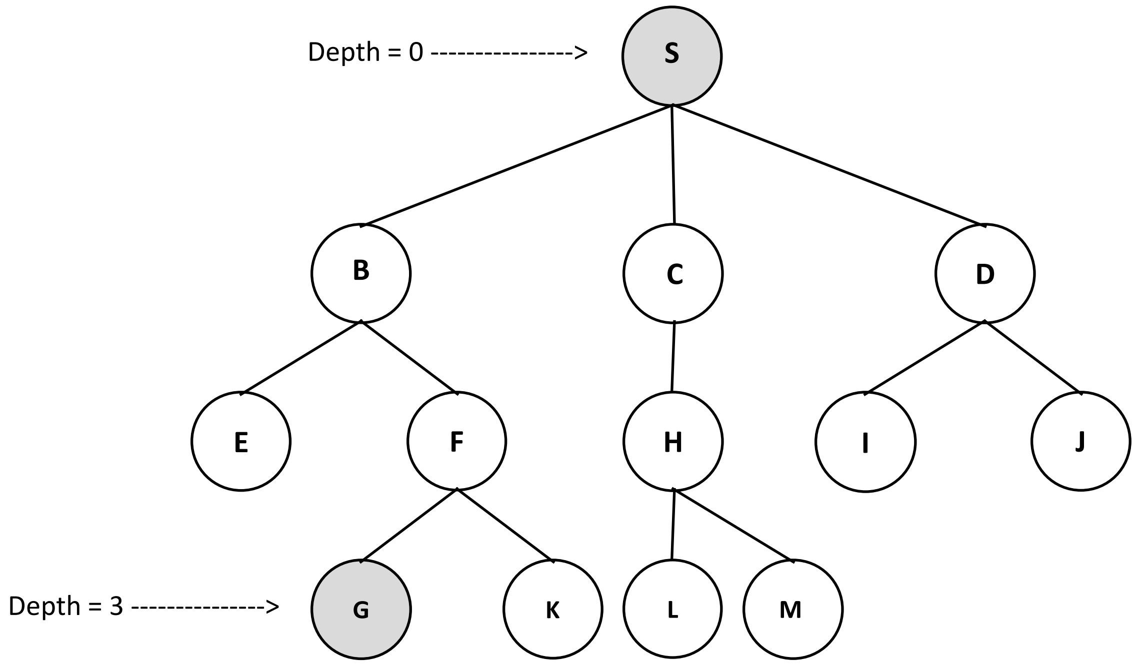

7.a. (7 points total, -1 for each mistake, but not negative) List the order in which nodes are expanded at each depth limit while performing IDS. Recall that to expand a node means to generate its children. Stop at the goal node. **For each depth limit in the table below, write "None" if no nodes are expanded at that depth limit. Otherwise, fill in the expanded nodes at that depth limit, in the order of their expansion.**

Depth Limit	Order of Node Expansion at that Depth Limit
0	None
1	S
2	S B C D
3	S B E F (G)

7.b. (3 points total, 1 pt each). Consider the search space below. S is the start node and G is the goal node. Children are returned in left-to-right order. Fill in the blanks with the numbers {10, 20, 2000} for step costs along the associated arc, using each number exactly once, so that IDS will find a worse than optimal solution.

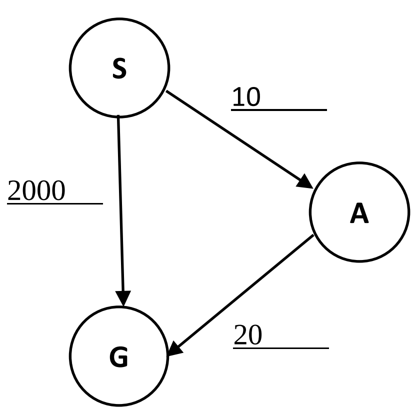

OR

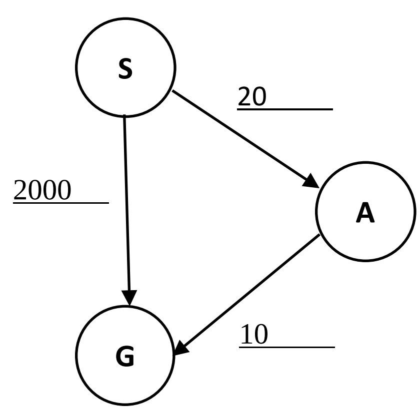

8. (11 pts total) PROBABILITY. Your friend has a game night every week. For fun, he uses a fake 6-sided die on certain nights, and on other nights he uses a true 6-sided die. [For non-native English speakers: a "die" is a gambling device, a small cube with random outcomes 1, 2, 3, 4, 5, or 6 depending on which side of the cube is on top when the die is rolled onto the table. Any true die has the probability of each outcome = 1/6. A fake die may have different probabilities associated with each outcome, which property makes it fake.]

In this question, D is a boolean random variable such that D = f denotes a fake die and D = t denotes a true die. Also, R is a discrete random variable that takes values in the domain {1, 2, 3, 4, 5, 6}. R denotes the outcome of rolling a single die. For example R = 1 means that a 1 was rolled.

8.a. (3 pts) Suppose you know the probability of rolling a 1, 2, 3, 4, or 5 given that the die is fake. In other words, you know $P(R = 1|D = f)$, $P(R = 2|D = f)$, $P(R = 3|D = f)$, $P(R = 4|D = f)$, and $P(R = 5|D = f)$. Given this information, find $P(R = 6|D = f)$. **Write your answer as a symbolic expression that involves $P(R = 1|D = f)$, $P(R = 2|D = f)$, $P(R = 3|D = f)$, $P(R = 4|D = f)$, and $P(R = 5|D = f)$.**

$$P(R = 6|D = f) = 1 - [\ P(R = 1|D = f) + P(R = 2|D = f) + P(R = 3|D = f) + P(R = 4|D = f) + P(R = 5|D = f)\ ]$$

8.b. (3 pts) This night, you are the first to roll the die. You do not know whether you have the fake die (where D = f) or the true die (where D = t). **What is the probability that you roll a 6 (i.e. $P(R = 6)$)?**

Use the sum rule and the product rule. Write your answer as a symbolic expression in terms of $P(R = 6|D = f)$, $P(R = 6|D =t)$, $P(D = f)$, and $P(D = t)$. **Show your work.**

$$P(R = 6) = P(R = 6 \wedge D = f) + P(R = 6 \wedge D = t)$$

You will receive full credit if you omitted the intermediate step shown here.

$$= P(R = 6|D = f)P(D = f) + P(R = 6|D = t)P(D = t)$$

8.c. (3 pts) Suppose you rolled the die first tonight, and the outcome was a 6 (R1 = 6). The next turn, your friend rolls the die and the outcome was a 1 (R2 = 1). **Given all that, what is the probability of the die being fake (= that D = f)?** That is, what is $P(D = f|R1 = 6 \wedge R2 = 1)$?

Use Baye's Rule to write your answer as a symbolic expression in terms of $P(R1 = 6 \wedge R2 = 1|D = f)$, $P(D = f)$, and $P(R1 = 6 \wedge R2 = 1)$.

$$P(D = f|R1 = 6 \wedge R2 = 1) = P(R1=6\ \ R2=1|D=f)P(D=f) / P(R1=6 \wedge R2=1)$$

8.d. (2 pts total, 1 pt each) The probability of each roll is conditionally independent given D. Given this, use conditional independence assumptions to rewrite $P(R1 = 6 \wedge R2 = 1|D = f)$ and $P(R1 = 6 \wedge R2 = 1|D = t)$ more simply in terms of $P(R1 = 6|D = f)$, $P(R2 = 1|D = f)$, $P(R1 = 6|D = t)$, and $P(R2 = 1|D = t)$.

$$P(R1 = 6 \wedge R2 = 1|D = t) = P(R1 = 6|D = t)P(R2 = 1|D = t)$$

$$P(R1 = 6 \wedge R2 = 1|D = f) = P(R1 = 6|D = f)P(R2 = 1|D = f)$$

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

9. (10 pts total) BAYESIAN NETWORKS. Here we create a probabilistic model for turning in a late assignment.

Let the boolean variable L represent that the student turns in a late assignment. L = 1 means the student turns in an assignment late. L = 0 means the student turns in an assignment on time.

Let the boolean variable P represent that the student has a habit of procrastinating. P = 1 means the student procrastinates. P = 0 means the student does not procrastinate.

Let the boolean variable S represents that the student is sick. S = 1 means the student is sick. S = 0 means the student is healthy.

Normally, we observe that variables P and S above have influence on the variable L. If P = 1 or S = 1, then it is more likely that L = 1. Normally, getting sick is independent of a habit of procrastinating. If L is unobserved, then P and S are independent. The following joint distribution illustrates these relationships.

P	**S**	**L**	**P(P,S,L)**
0	0	0	0.3
0	0	1	0.05
0	1	0	0.05
0	1	1	0.1
1	0	0	0.15
1	0	1	0.1
1	1	0	0.0
1	1	1	0.25

9.a. (4 pts) Based on the information above, draw a Bayesian Network diagram to demonstrate the relationships between P, S, L. Ignore the associated probability tables for now. Just draw the network. Your answer should be a Directed Acyclic Graph wherein nodes are variables and arcs represent direct influence.

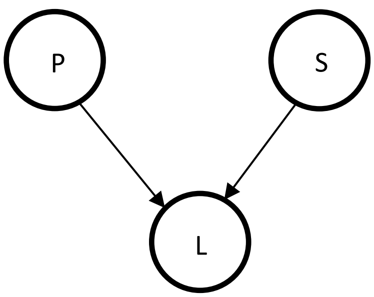

9.b. (6 pts total, -1 for each error, but not negative) The full joint distribution requires a large table. **Fill in the probability tables below associated with each node above.** An arithmetic expression contains only numbers, parentheses, and +, –, *, and /. Write your answer as an arithmetic expression that will evaluate to the correct number.

P(Procrastination)
0.15 + 0.1 +0.0 + 0.25

P(Sick)
0.05 + 0.1 + 0.0 +0.25

P(Procrastination) = Σ P(P = 1, S, L)
= 0.15 + 0.1 +0.0 + 0.25 = 0.5

P(Sick) = Σ P(P, S = 1, L)
= 0.05 + 0.1 + 0.0 +0.25 = 0.4

P	S	P(L=1|P,S)
0	0	0.05 /(0.3 + 0.05) or 1/7
0	1	0.1 /(0.05 + 0.1) or 2/3
1	0	0.1 / (0.15 + 0.1) or 2/5
1	1	0.25 / (0.0 + 0.25) or 1

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