

For the Mid-term Exam, “Perfect” gives the percentage of students who received full credit, “Partial” gives the percentage who received partial credit, and “Zero” gives the percentage who received zero credit.

(Due to rounding, etc., values below may be only approximate estimates.)

### **Problem 1**

Perfect: ~13% (~26 students), Partial: ~87% (~170 students), Zero: ~0% (~0 students)  
A common mistake was to write the wrong search tree in problem 1.c.

### **Problem 2**

Perfect: ~63% (~124 students), Partial: ~28% (~56 students), Zero: ~9% (~18 students)  
A common mistake was to confuse the goal nodes on the queue between G=8 (the long cheap path) and G=11 (the short expensive path).

### **Problem 3**

Perfect: ~67% (~132 students), Partial: ~23% (~46 students), Zero: ~10% (~19 students)

### **Problem 4**

Perfect: ~95% (~189 students), Partial: ~2% (~4 students), Zero: ~3% (~5 students)

### **Problem 5**

Perfect: ~59% (~117 students), Partial: ~38% (~75 students), Zero: ~3% (~6 students)

Common mistakes were:

(1) To do Alpha-Beta Pruning before Mini-Max or otherwise to conflate the two, resulting in confusion. In general, we were fairly generous in grading this question, and if your answers showed that you understood the material and knew what was going on, in general you received full credit.

(2) In branch (B), to believe that left-most leaf node 5 (resp. 4) in the middle (resp. right) branch under (B) needed to be examined. As soon as 2 is seen in the left-most branch under (B), all are pruned.

For 5(3), many students just drew pruning lines, even though the instructions stated clearly:

**Cross out each leaf node that would be pruned by Alpha-Beta Pruning. Do not just draw pruning lines.**

This time, you received full credit for drawing pruning lines. In the future, please follow instructions.

### **Problem 6**

Perfect: ~60% (~119 students), Partial: ~33% (~65 students), Zero: ~7% (~13 students)

Common mistakes were to multiply instead of exponentiating, and to confuse d with m.

### **Problem 7**

Perfect: ~60% (~118 students), Partial: ~39% (~77 students), Zero: ~1% (~1 students)

### **Problem 8**

Perfect: ~75% (~144 students), Partial: ~24% (~47 students), Zero: ~1% (~1 students)

Common mistakes were to confuse D with A, or D with F.

### **Problem 9**

Perfect: ~91% (~179 students), Partial: ~8% (~15 students), Zero: ~1% (~2 students)

Common mistakes were:

P = Percept(s), Perceptor(s), Process, Program, Pieces

A = Agent(s), Agility

S = State

# CS-171, Intro to A.I. — Mid-term Exam — Fall 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. No electronics are allowed at any point of the exam.

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. (15 pts total) Problem Solving by search.
2. (15 pts total, -1 for each error, but not negative) A\* search showing queue.
3. (10 pts total, -1 for each error, but not negative) Conversion to CNF.
4. (10 pts total, -1 for each error, but not negative) Resolution Theorem Proving.
5. (10 pts total, -1 for each error, but not negative) Mini-Max, Alpha-Beta Pruning.
6. (10 pts total, -1 pt each wrong answer, but not negative) Search Properties.
7. (15 pts total, 5 pts each, -1 each error, but not negative) Bayesian Networks.
8. (10 pts total, 1 pt each) Adversarial (Game) Search Concepts.
9. (5 pts total, -2 pts for each error, but not negative) Task Environment.


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



**1. (15 pts total) Problem Solving by search.**

We are now trying to solve the miniature of Tower of Hanoi problem. There are 3 towers (A, B, C), and 2 disks (small one and large one). The purpose of this problem is to move both disks from the tower A to tower C (as illustrated in the figure below), subject to following two conditions:

- You can move only one disk at a time.
- You can move only the top disk in a stack of disks.
- You cannot put the large disk on top of the small disk.



The possible states can be denoted as follows (A = tower 1, B = tower 2, C = tower 3):

a:(b c) where a is the state number, b is the tower number for the large disk, and c is the tower number for the small disk. E.g., 3:(1 3) implies that in state 3 the large disk is on tower 1 and the small disk is on tower 3.

If we use this notation, the following nine states are possible.

1:(1 1), 2:(1:2), 3:(1 3), 4:(2 1), 5:(2 2), 6:(2 3), 7:(3 1), 8:(3 2), 9:(3 3)

**1.a (4 pts total, 2 pts each) Which states are the initial state and goal state?**

**1.a.1 (2 pts) Initial state: 1 or 1:(1 1)**

**1.a.2 (2pts) Goal state: 9 or 9:(3 3)**

**1.b. (6 pts total, -1 for each error, but not negative) Enumerate all one-move transitions between states.**

For example, 1->2, 3 means that it is possible to make a transition from state 1 to state 2 or state 3 in one move.

**The first one is done for you as an example.**

- 1-> 2, 3
- 2-> 1, 3, 8
- 3-> 1, 2, 6
- 4-> 5, 6, 7
- 5-> 4, 6
- 6-> 3, 4, 5
- 7-> 4, 8, 9
- 8-> 2, 7, 9
- 9-> 7, 8

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


1.c. (5 pts total, -1 for each error, but not negative) Problem Solving by search (continued).

Find a solution using breath-first search and draw the search tree. Stop when you reach goal state 9.

Use Tree Search, i.e., do not remember visited nodes. Children of a node are returned in increasing numerical order by state; i.e., the children of a node are ordered left-to-right by increasing state number.

Assume that cycles are detected and eliminated by never expanding a node containing a state that is repeated on the path back to the root. That is, if a newly-generated child node state also occurs on the path from the parent back to the root, then that child is pruned immediately. This pruning strategy is a “light-weight” way to avoid many cycles and repeated states in Tree Search without the full memory burden of remembering all visited nodes. However, it is an approximate strategy and does not fully solve the repeated node problem.

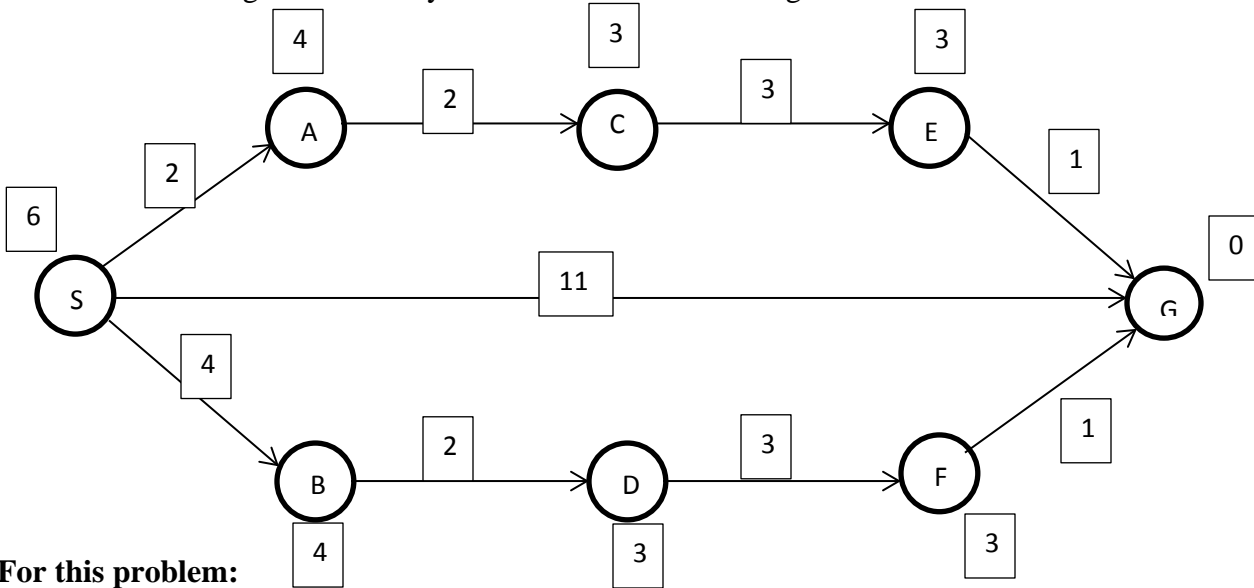
The first expansion is done for you as an example.



**2. (15 pts total, -1 for each error, but not negative) A\* search showing queue.**

Use the A\* algorithm on the graph below to find a path from node S to node G.

- Each node is labeled by a capital letter and the value of a heuristic function at that node.
- Each edge is labeled by the cost to traverse that edge.



**For this problem:**

- Perform the A\* algorithm on this graph, filling in the table below. Indicate the f values of each node on the queue as shown in the first two rows of the table. You need not to order the contents of the (priority) queue in the table. You should not need all the lines provided in the table.
- Assume that you are doing Tree Search, i.e., do not remember visited nodes.
- Write the priority queue as <node-label>=<f-value>
- The first one is done for you, as an example.

iteration	Node expanded	Priority queue at the end of this iteration
0		S=6
1	S=6	A=6; B=8; G=11
2	A=6	B=8; C=7; G=11
3	C=7	B=8; E=10; G=11
4	B=8	D=9; E=10; G=11
5	D=9	E=10; F=12; G=11
6	E=10	F=12; G=8; G=11
7	G=8	F=12; G=11
8		
9		
10		

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

**3. (10 pts total, -1 for each error, but not negative) Conversion to CNF.**

Convert the following Propositional Logic expression to Conjunctive Normal Form (CNF).

**SHOW YOUR WORK. Correct work = full credit, correct answer without work = zero credit.**

$$( X \Leftrightarrow ( Y \wedge Z ) )$$

:: eliminate equivalence by  $( A \Leftrightarrow B ) = [( A \Rightarrow B ) \wedge ( B \Rightarrow A )]$

$$[( X \Rightarrow ( Y \wedge Z ) ) \wedge ( ( Y \wedge Z ) \Rightarrow X )]$$

:: eliminate implication by  $( A \Rightarrow B ) = ( \neg A \vee B )$

$$[( ( \neg X ) \vee ( Y \wedge Z ) ) \wedge ( ( \neg ( Y \wedge Z ) ) \vee X )]$$

:: use DeMorgan's laws to move NOT all the way inside to literals

$$[( ( \neg X ) \vee ( Y \wedge Z ) ) \wedge ( ( \neg Y \vee \neg Z ) \vee X )]$$

:: simplify

$$[( ( \neg X ) \vee ( Y \wedge Z ) ) \wedge ( ( \neg Y ) \vee ( \neg Z ) \vee X )]$$

:: distribute

$$[( ( \neg X ) \vee Y ) \wedge ( ( \neg X ) \vee Z ) \wedge ( ( \neg Y ) \vee ( \neg Z ) \vee X )]$$

:: final answer (same as above)

$$( ( \neg X ) \vee Y ) \wedge ( ( \neg X ) \vee Z ) \wedge ( ( \neg Y ) \vee ( \neg Z ) \vee X )$$

:: final answer in clausal form (optional)

$$( ( \text{NOT } X ) Y )$$

$$( ( \text{NOT } X ) Z )$$

$$( ( \text{NOT } Y ) ( \text{NOT } Z ) X )$$

**4. (10 pts total, -1 for each error, but not negative) Resolution Theorem Proving.**

Your Knowledge Base (KB) contains the following statements:

- $(\neg A)$
- $((\neg A) \Leftrightarrow ((\neg B) \wedge (\neg C)))$
- $((\neg B) \Rightarrow ((\neg D) \wedge (\neg E)))$
- $((\neg C) \Rightarrow (D \vee E \vee (\neg F)))$

**You are asked to prove the goal statement  $(\neg F)$ .**

After converting your KB to CNF and adjoining the negated goal statement you have (in clausal form):

- (NOT A)                      (A (NOT B))                      (A (NOT C))                      (B (NOT D))
- (B (NOT E))                      (C D E (NOT F))                      F

**Produce a resolution proof that  $(\neg F)$  is true. I.e., run resolution on the statements above to yield  $(\ )$ .**

*Think about your problem, and find a proof that mirrors how you think. You know that  $(\neg A)$  so it is easy to prove  $((\neg B) \wedge (\neg C))$ . From  $(\neg B)$  it is easy to prove  $((\neg D) \wedge (\neg E))$ . From  $(\neg C)$  it is easy to prove  $(D \vee E \vee (\neg F))$ . You just have proved  $((\neg D) \wedge (\neg E))$ , so only  $(\neg F)$  remains.*

The shortest proof I know if is only eight lines long (Bonus Point for a shorter proof).

- Resolve       (NOT A)       with       (A (NOT B)      ) to produce:       (NOT B)
- Resolve       (NOT A)       with       (A (NOT C))       to produce:       (NOT C)
- Resolve       (NOT B)       with       (B (NOT D))       to produce:       (NOT D)
- Resolve       (NOT B)       with       (B (NOT E))       to produce:       (NOT E)
- Resolve       (NOT C)       with       (C D E (NOT F))       to produce:       (D E (NOT F))
- Resolve       (NOT D)       with       (D E (NOT F))       to produce:       (E (NOT F))
- Resolve       (NOT E)       with       (E (NOT F))       to produce:       (NOT F)
- Resolve       (NOT F)       with       F       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: \_\_\_\_\_
- Resolve \_\_\_\_\_ with \_\_\_\_\_ to produce: \_\_\_\_\_

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



**Quite a few students found proofs shorter than eight lines and got bonus points.**

An example six-line proof is:

Resolve F with (C D E (NOT F)) to produce: (C D E)

Resolve (C D E) with (B (NOT E)) to produce: (B C D)

Resolve (B C D) with (B (NOT D)) to produce: (B C)

Resolve (B C) with (A (NOT B)) to produce: (A C)

Resolve (A C) with (A (NOT C)) to produce: (A)

Resolve (A) with (NOT A) to produce: ( )

An example seven-line proof is:

Resolve (A (NOT B)) with (B (NOT D)) to produce: (A (NOT D))

Resolve (A (NOT D)) with (C D E (NOT F)) to produce: (A C E (NOT F))

Resolve (A C E (NOT F)) with (A (NOT C)) to produce: (A E (NOT F))

Resolve (A E (NOT F)) with (B (NOT E)) to produce: (A B (NOT F))

Resolve (A B (NOT F)) with (A (NOT B)) to produce: (A (NOT F))

Resolve (A (NOT F)) with F to produce: (A)

Resolve (A) with (NOT A) to produce: ( )

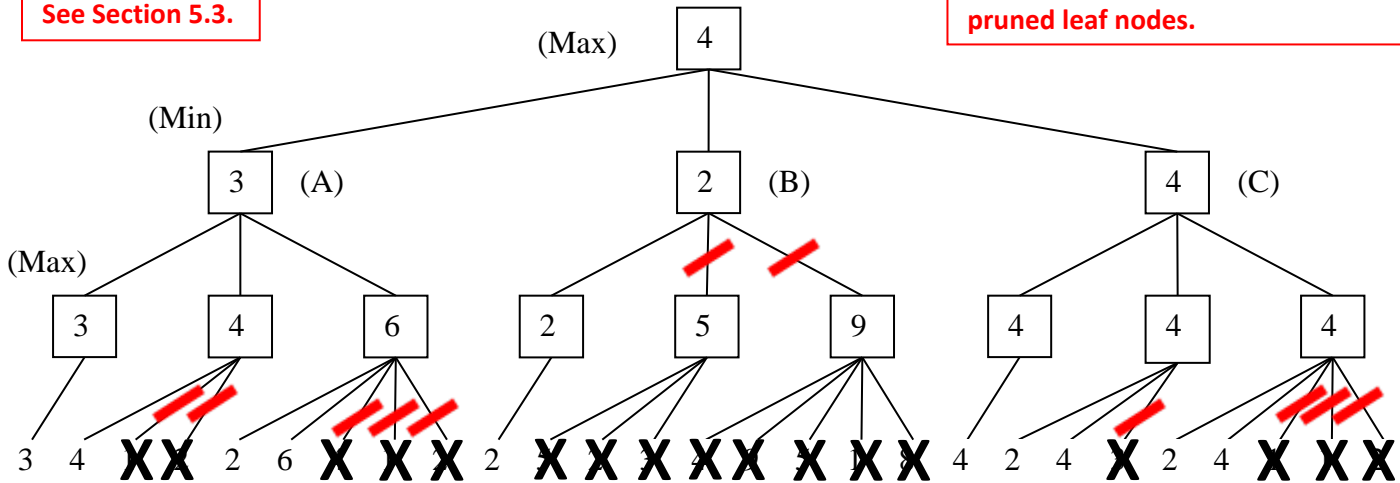
5. (10 pts total, -1 for each error, but not negative) Mini-Max, Alpha-Beta Pruning. In the game tree below it is **Max's** turn to move. At each leaf node is the estimated score of that resulting position as returned by the heuristic static evaluator.

- (1) Perform Mini-Max search and label each branch node with its value.
- (2) Cross out each leaf node that would be pruned by alpha-beta pruning.
- (3) What is Max's best move (A, B, or C)?     C

Red lines indicate where in the tree pruning occurred. You are not obliged to provide the red lines — only to cross out pruned leaf nodes.

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

See Section 5.3.



6. (10 pts total, -1 pt each wrong answer, but not negative) Search Properties.

Fill in the values of the four evaluation criteria for each search strategy shown. Assume a tree search where  $b$  is the finite branching factor;  $d$  is the depth to the shallowest goal node;  $m$  is the maximum depth of the search tree;  $C^*$  is the cost of the optimal solution; step costs are identical and equal to some positive  $\epsilon$ ; and in Bidirectional search both directions use breadth-first search.

Note that these conditions satisfy all of the footnotes of Fig. 3.21 in your textbook.


See Figure 3.21.

Criterion	Complete?	Time complexity	Space complexity	Optimal?
Breadth-First	Yes	$O(b^d)$	$O(b^d)$	Yes
Uniform-Cost	Yes	$O(b^{(1+\lfloor C^*/\epsilon \rfloor)})$ $O(b^{(d+1)})$ also OK	$O(b^{(1+\lfloor C^*/\epsilon \rfloor)})$ $O(b^{(d+1)})$ also OK	Yes
Depth-First	No	$O(b^m)$	$O(bm)$	No
Iterative Deepening	Yes	$O(b^d)$	$O(bd)$	Yes
Bidirectional (if applicable)	Yes	$O(b^{(d/2)})$	$O(b^{(d/2)})$	Yes

7. (15 pts total, 5 pts each, -1 each error, but not negative) Bayesian Networks.


7.a. (5 pts) Write down the factored conditional probability expression that corresponds to the graphical Bayesian Network shown.

$P(A | C,D,H) P(B | D,E,G) P(C | F) P(D | G,H) P(E | H) P(F | G) P(G | H) P(H)$




7.b. (5 pts) Draw the Bayesian Network that corresponds to this conditional probability:

$P(A | B,C,E) P(B | D,G) P(C | E,F,H) P(D | G) P(E | G,H) P(F | H) P(G | H) P(H)$



7.c. (5 pts) Shown below is the Bayesian network corresponding to the Burglar Alarm problem,

$P(J | A) P(M | A) P(A | B, E) P(B) P(E)$ .



P(E)
.002

A	P(M)
t	.70
f	.01

B	E	P(A)
t	t	.95
t	f	.94
f	t	.29
f	f	.001

P(B)
.001

A	P(J)
t	.90
f	.05

The probability tables show the probability that variable is True, e.g., P(M) means P(M=t).

Write down an expression that will evaluate to  $P(J=t \wedge M=f \wedge A=t \wedge B=f \wedge E=t)$ . Express your answer as a series of numbers (numerical probabilities) separated by multiplication symbols. You do not need to carry out the multiplication to produce a single number (probability). **SHOW YOUR WORK.**

$$\begin{aligned} &P(J=t \wedge M=f \wedge A=t \wedge B=f \wedge E=t) \\ &= P(J=t | A=t) * P(M=f | A=t) * P(A=t | B=f \wedge E=t) * P(B=f) * P(E=t) \\ &= .90 * .30 * .29 * .999 * .002 \end{aligned}$$

Many students simply provided a list of numbers, but did not show the intermediate symbolic conditional probabilities. Such an answer will be accepted for this Mid-term Exam only. In the future, **SHOW YOUR WORK** means show your work.

**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.

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

**9. (5 pts total, -2 pts for each error, but not negative) TASK ENVIRONMENT.** Your book defines a task environment as a set of four things, with the acronym PEAS. Fill in the blanks with the names of the PEAS components.

Performance (measure)                      Environment                      Actuators                      Sensors

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