CS-171, Intro to A.I. — Midterm Exam — Summer Session 1, 2018

YOUR NAME: _____

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. (15 pts total) PROPOSITIONAL LOGIC.
2. (15 pts total, 3 pts each) SEARCH STRATEGIES.
3. (5 pts total, -1 pt for each error, but not negative) MINI-MAX SEARCH IN GAME TREES.
4. (10 pts total, -1 for each error, but not negative) ALPHA-BETA PRUNING.
5. (6 pts total, 1 pt each) Logic Concepts.
6. (12 pts total, 1 pt each) Properties of task environments.
7. (15 pts total, 5 pts each, -1 for each error, but not negative) BAYESIAN NETWORKS.
8. (7 pts total, 1 pt each) The Knowledge Engineering process.
9. (6 pts total, 1 pt each) Logic-To-English.
10. (9 pts total, 1 pt each) Probability concepts and formulae.

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

1. (15 pts total) PROPOSITIONAL LOGIC. You are a logical agent who has been given the following sentence as your knowledge base (KB): $(A \rightarrow C) \wedge (C \rightarrow A) \wedge (B \rightarrow C)$

1.a. (2 pts) Convert the sentence above into CNF (Conjunctive Normal Form)

$(\neg A \text{ or } C) \wedge (\neg C \text{ or } A) \wedge (\neg B \text{ or } C)$

Clausal form is also OK:
$(\neg A\ C)\ (\neg C\ A)\ (\neg B\ C)$

1.b. (5 pts, -1 each error, but not negative) Fill in the following truth table for your knowledge base. **The first one is done for you as an example.**

A	B	C	KB
T	T	T	T
T	T	F	F
T	F	T	T
T	F	F	F
F	T	T	F
F	T	F	F
F	F	T	F
F	F	F	T

1.c. (2 pts, -1 each error, but not negative) List all the possible worlds that are models of KB. List each possible world as (A=T/F, B=T/F, C=T/F), where T/F are chosen to correspond to one row in the table above. **The first one is done for you as an example.**

(A=T, B=T, C=T), (A=T, B=F, C=T), (A=F, B=F, C=F)

1.d. (2 pts) Does your knowledge base entail $\neg (A \wedge B)$? (Answer Yes or No.)

No. (A=T, B=T, C=T) evaluates the sentence to be false.

1.e. (2 pts, -1 each error, but not negative) You learned a new sentence, $\neg C$, and added it to your existing KB to make a new KB. List all the possible worlds that are models of your new KB.

(A=F, B=F, C=F)

1.f. (2 pts) Does your new knowledge base entail $\neg (A \wedge B)$? (Answer Yes or No.)

Yes. (A=F, B=F, C=F) is the only model of the new KB, and it evaluates the sentence to be true.

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

2. (15 pts total, 3 pts each) SEARCH STRATEGIES. Execute Tree Search through this graph (i.e., do not remember visited nodes, so repeated nodes are possible). It is not a tree, but pretend that you don't know that it is not. S is the start node and G1, G2 are the two goal 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 S are A, B; successors of A are G1, G2; successors of B are A, B; in those orders.**

For each search strategy below, show the order[s] children are generated) . If stuck in a loop, indicate the [s] write "None". **The first one is done for you as an exa[mple].**

Please see the lecture slides for Uninformed Search, topic "When to do Goal-Test? When generated? When popped?" for clarification about exactly what to do in practical cases.

See Chapter 3.

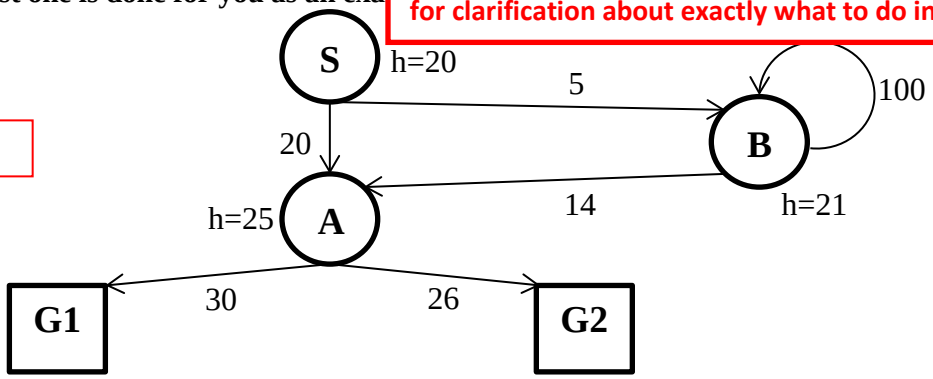

2.a. (example) DEPTH FIRST SEARCH.

DFS does the Goal-test iteratively on each child as generated, keeping the queue on the stack.

See Section 3.4.3 and Fig. 3.17.

Order of node expansion: S A G1

Path found: S A G1

2.b. (3 pts total) BREADTH FIRST SEARCH.

BFS does the Goal-test before the child is pushed onto the queue. The goal is found when A is expanded.

See Section 3.4.1 and Fig. 3.11.

(2 pts) Order of node expansion: S A G1

(1 pt) Path found: S A G1

2.c. (3 pts total) UNIFORM COST SEARCH.

UCS does goaltest when node is popped off queue.

See Section 3.4.2 and Fig. 3.14.

(2 pts) Order of node expansion: S B A G2

(1 pt) Path found: S B A G2

2.d. (3 pts total) GREEDY (BEST-FIRST) SEARCH.

B always has lower h (=21) than any other node on queue.

See Section 3.5.1 and Fig. 3.23.

(2 pts) Order of node expansion: S B B B B B B ... etc.

(1 pt) Path found: None

2.e. (3 pts total) ITERATED DEEPENING SEARCH

IDS does the Goal-test iteratively on each child as generated, keeping the queue on the stack.

See Sections 3.4.4-5 and Figs. 3.18-19.

(2 pts) Order of node expansion: S S A G1

(1 pt) Path found: S A G1

2.f. (3 pts total) A* SEARCH.

A* does goaltest when node is popped off queue.

See Section 3.5.2 and Figs. 3.24-25.

(2 pts) Order of node expansion: S B A G2

(1 pt) Path found: S B A G2

3. (5 pts total, -1 pt for each error, but not negative) MINI-MAX SEARCH IN GAME TREES.
The game tree below illustrates a position reached in the game. Process the tree left-to-right. It is **Max**'s turn to move. At each leaf node (number at bottom) is the estimated score returned by the heuristic static evaluator.

3.a. Fill in each blank square with the proper mini-max search value.

3.b. What is the best move for Max? (write A, B, or C) A

3.c. What score does Max expect to achieve? 5

See Section 5.2.1.

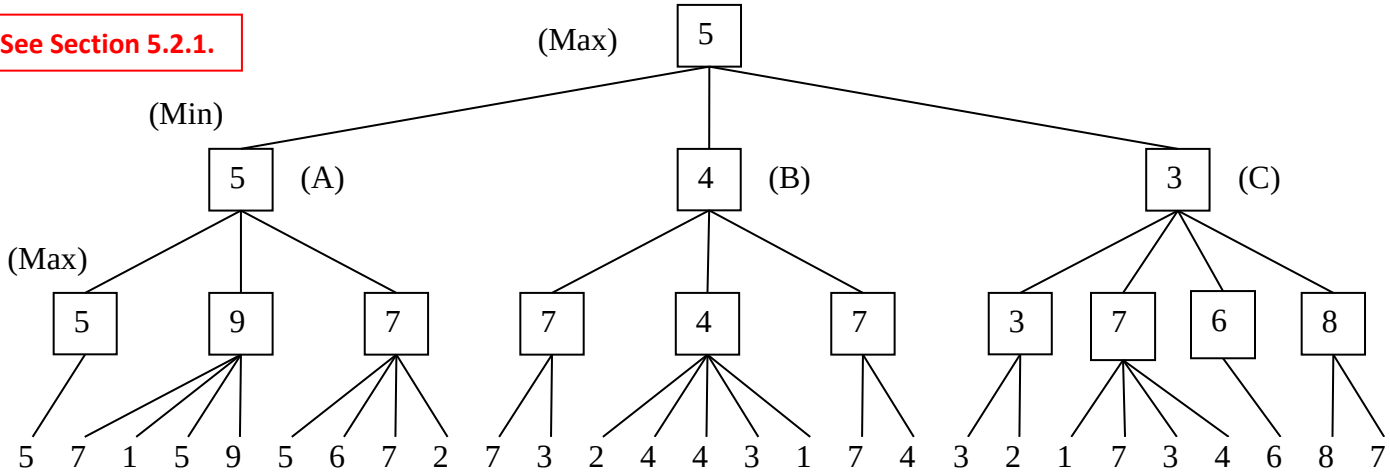

4. (10 pts total, -1 for each error, but not negative) ALPHA-BETA PRUNING. Process the tree left-to-right. This is the same tree as above (3.a). You do not need to indicate the branch node values again (but you may do so if you wish).

<u>Draw X over each leaf node (number at bottom) that will be pruned by Alpha-Beta Pruning.</u>

See Section 5.3.

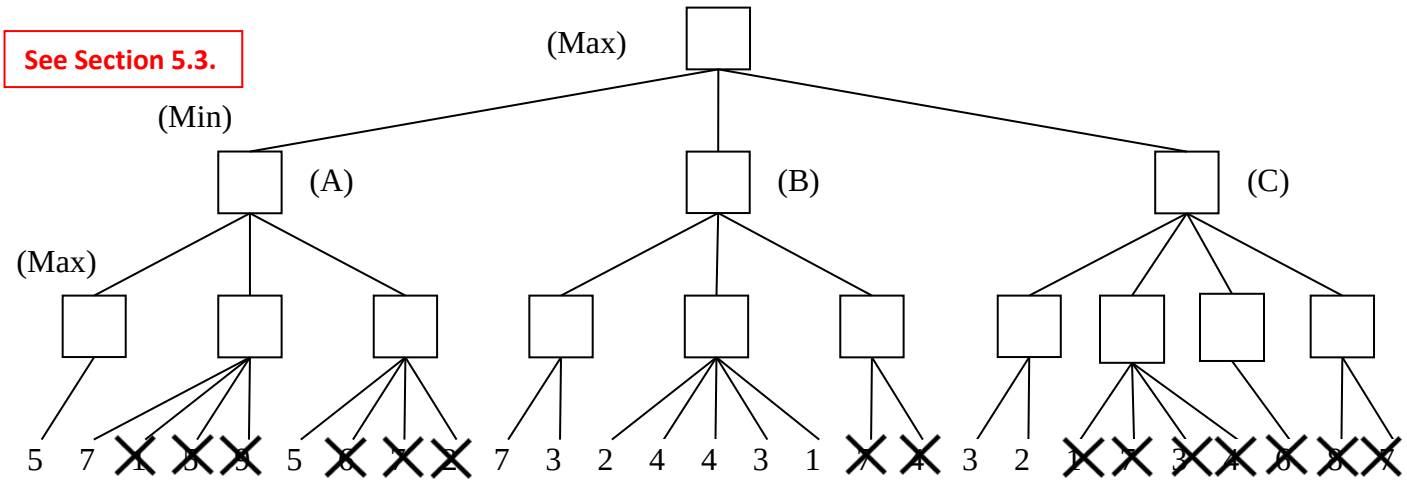

****** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ******

5. (6 pts total, 1 pt each) Logic Concepts. Match each term or phrase on the left to the best correspondence on the right. **The first one is done for you as an example.**

See Chapter 7.

A .	Logic	A	Formal symbol system for representation and inference
C	Valid	B	The idea that a sentence follows logically from other sentences
G	Complete	C	True in every possible world
E	Conjunctive Normal Form	D	True in at least one possible world
F	Sound	E	A sentence expressed as a conjunction of clauses (disjuncts)
D	Satisfiable	F	Inference system derives only entailed sentences
B	Entailment	G	Inference system can derive any sentence that is entailed

6. (12 pts total, 1 pt each) Properties of task environments. For each of the following terms or phrases on the left, write in the letter corresponding to the best answer on the [illegible]on on the right. **The first one is done for you as an example.**

See Section 2.3.2 and Fig. 2.5.

A	Agent	A	Perceives environment by sensors, acts by actuators
K	Episodic	B	Sensors give complete state of environment at each time point
G	Discrete	C	More than one agent in the task environment
I	Static	D	Next state is exactly determined by current state and agent action
E	Sequential	E	The current decision could affect all future decisions
J	Semidynamic	F	Environment can change while the agent is deliberating
D	Deterministic	G	Finite number of states, percepts, and actions
B	Fully observable	H	The outcomes (or probabilities) for all actions are given
M	Uncertain	I	Environment does not change while the agent is deliberating
H	Known	J	Environment does not change while the agent is deliberating, but its performance measure does
C	Multiagent	K	A series of atomic episodes, each independent of prior agent actions
L	Stochastic	L	Next state not exactly determined by current state and agent action
F	Dynamic	M	Not fully observable or not deterministic

7. (15 pts total, 5 pts each) BAYESIAN NETWORKS.

7.a. (5 pts total, -1 for each error, but not negative) Write down the factored conditional probability expression corresponding to this Bayesian Network:

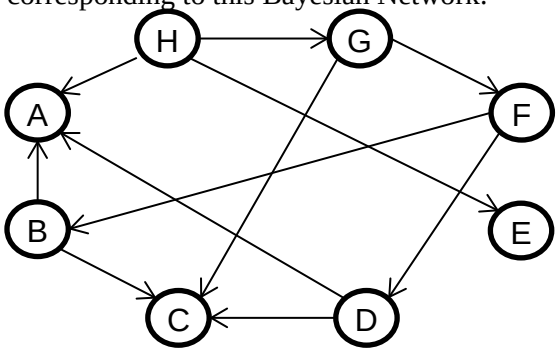

See Section 14.1-4.

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

7.b. (5 pts, -1 for each error, but not negative) Draw the Bayesian Network corresponding to this factored conditional probability expression:

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

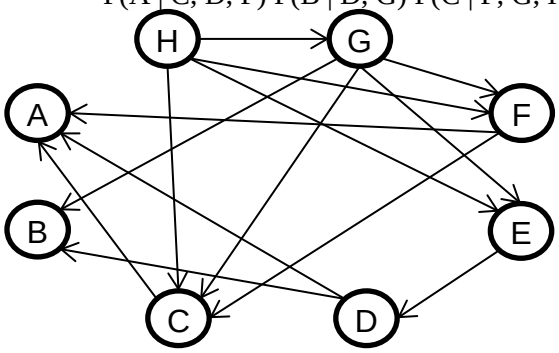

7.c. (5 pts, -1 for each error, but not negative) Shown below is the Bayesian network corresponding to the Burglar Alarm problem, i.e., P(J,M,A,B,E) = P(J | A) P(M | A) P(A | B, E) P(B) P(E). This is Fig. 14.2 in your R&N textbook. The probability tables show the probability that a boolean random variable is true. For example, P(J=t|A=t) is 0.9 and so P(J=f|A=t) is 1 – 0.9 = 0.1.

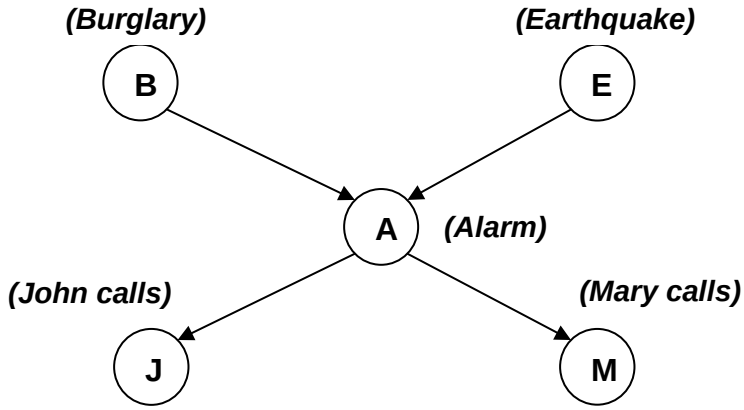

P(B)
.001

P(E)
.002

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

A	P(J)
t	.90
f	.05

A	P(M)
t	.70
f	.01

Write down an expression that will evaluate to P(J=t ∧ M=t ∧ A=f ∧ B=f ∧ 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).

P(J=t ∧ M=t ∧ A=f ∧ B=f ∧ E=t) = .05 * .01 * .71 * .999 * .002

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

8. (7 pts total, 1 pt each) The Knowledge Engineering process.
Your book identifies seven sequential steps in the knowledge engineering process, which steps are below. Unfortunately, the order of the steps has been scrambled. Please, straighten them out.

See Section 8.4

A. Encode a description of the specific problem instance
B. Assemble the relevant knowledge
C. Pose queries to the inference procedure and get answers
D. Encode general knowledge about the domain
E. Debug the knowledge base
F. Identify the task
G. Decide on a vocabulary of predicates, functions, and constants

Fill in the blanks with the letters A, B, C, D, E, F, and G, all in the proper sequence.

F B G D A C E .

9. (6 pts total, 1 pt each) Logic-To-English. For each of the following FOPC sentences on the left, write the letter corresponding to the best English sentence on the right. Use these intended interpretations: (1) "Student(x)" is intended to mean "x is a student." (2) "Quiz(x)" is intended to mean "x is a quiz." (3) "Got100(x, y)" is intended to mean "x got 100 on y."

See Section 8.2.6

B	$\forall s \exists q$ Student(s) $\Rightarrow$ [Quiz(q) $\wedge$ Got100(s, q)]	A	For every quiz, there is a student who got 100 on it.
E	$\exists q \forall s$ Quiz(q) $\wedge$ [Student(s) $\Rightarrow$ Got100(s, q)]	B	For every student, there is a quiz on which that student got 100.
A	$\forall q \exists s$ Quiz(q) $\Rightarrow$ [Student(s) $\wedge$ Got100(s, q)]	C	Every student got 100 on every quiz.
F	$\exists s \forall q$ Student(s) $\wedge$ [Quiz(q) $\Rightarrow$ Got100(s, q)]	D	Some student got 100 on some quiz.
C	$\forall s \forall q$ [Student(s) $\wedge$ Quiz(q)] $\Rightarrow$ Got100(s, q)	E	There is a quiz on which every student got 100.
D	$\exists s \exists q$ Student(s) $\wedge$ Quiz(q) $\wedge$ Got100(s, q)	F	There is a student who got 100 on every quiz.

10. (9 pts total, 1 pt each) Probability concepts and formulae. Match each term or phrase on the left to the best correspondence on the right. **The first one is done for you as an example.**

See Chapter 13.

A .	Probability Theory	A	Assigns each sentence a degree of belief ranging from 0 to 1
H	Conditional independence	B	Degree of belief accorded without any other information
G	Independence	C	Degree of belief accorded after some evidence is obtained
J	Product rule (chain rule)	D	Gives probability of all combinations of values of all variables
C	Conditional probability	E	Takes values from its domain with specified probabilities
B	Unconditional probability	F	$P(A) = \Sigma_B P(A, B)$
F	Sum Rule	G	$P(a \wedge b) = P(a) P(b)$
E	Random variable	H	$P(a \wedge b \mid c) = P(a \mid c) P(b \mid c)$
I	Bayes' rule	I	$P(a \mid b) = P(b \mid a) P(a) / P(b)$
D	Joint probability distribution	J	$P(a \wedge b \wedge c) = P(a \mid b \wedge c) P(b \mid c) P(c)$

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