

INSTRUCTIONS

- You have 3 hours.
- The exam is closed book, closed notes except a two-page crib sheet.
- Please use non-programmable calculators only.
- Mark your answers ON THE EXAM ITSELF. If you are not sure of your answer you may wish to provide a *brief* explanation. All short answer sections can be successfully answered in a few sentences AT MOST.


First name	
Last name	
SID	
Login	

For staff use only:

Q1. Search Traces	/15
Q2. Multiple-choice and short-answer questions	/16
Q3. Minimax and Expectimax	/12
Q4. n -pacmen search	/10
Q5. CSPs: Course scheduling	/12
Q6. Cheating at cs188-Blackjack	/19
Q7. Markov Decision Processes	/16
Total	/100

Q1. [15 pts] Search Traces

Each of the trees (G1 through G5) was generated by searching the graph (below, left) with a *graph search* algorithm. Assume children of a node are visited in alphabetical order. Each tree shows *only the nodes that have been expanded*. Numbers next to nodes indicate the relevant “score” used by the algorithm’s priority queue. The start state is A, and the goal state is G.



For each tree, indicate:

- Whether it was generated with depth first search, breadth first search, uniform cost search, or A^* search. Algorithms may appear more than once.
- If the algorithm uses a heuristic function, say whether we used
 - H1** = $\{h(A) = 3, h(B) = 6, h(C) = 4, h(D) = 3\}$
 - H2** = $\{h(A) = 3, h(B) = 3, h(C) = 0, h(D) = 1\}$
- For all algorithms, say whether the result was an optimal path (assuming we want to minimize sum of link costs). If the result was *not* optimal, state why the algorithm found a suboptimal path.

Please fill in your answers on the next page.

(a) [3 pts] **G1:**

1. Algorithm:
2. Heuristic (if any):
3. Did it find least-cost path? If not, why?

(b) [3 pts] **G2:**

1. Algorithm:
2. Heuristic (if any):
3. Did it find least-cost path? If not, why?

(c) [3 pts] **G3:**

1. Algorithm:
2. Heuristic (if any):
3. Did it find least-cost path? If not, why?

(d) [3 pts] **G4:**

1. Algorithm:
2. Heuristic (if any):
3. Did it find least-cost path? If not, why?

(e) [3 pts] **G5:**

1. Algorithm:
2. Heuristic (if any):
3. Did it find least-cost path? If not, why?


Q2. [16 pts] Multiple-choice and short-answer questions

In the following problems please choose **all** the answers that apply, if any. You may circle more than one answer. You may also circle no answers (none of the above)


- (a) [2 pts] Consider two consistent heuristics, H_1 and H_2 , in an A^* search seeking to minimize path costs in a graph. Assume ties don't occur in the priority queue. If $H_1(s) \leq H_2(s)$ for all s , then
- (i) A^* search using H_1 will find a lower cost path than A^* search using H_2 .
 - (ii) A^* search using H_2 will find a lower cost path than A^* search using H_1 .
 - (iii) A^* search using H_1 will not expand more nodes than A^* search using H_2 .
 - (iv) A^* search using H_2 will not expand more nodes than A^* search using H_1 .
- (b) [2 pts] Alpha-beta pruning:
- (i) May not find the minimax optimal strategy.
 - (ii) Prunes the same number of subtrees independent of the order in which successor states are expanded.
 - (iii) Generally requires more run-time than minimax on the same game tree.
- (c) [2 pts] Value iteration:
- (i) Is a model-free method for finding optimal policies.
 - (ii) Is sensitive to local optima.
 - (iii) Is tedious to do by hand.
 - (iv) Is guaranteed to converge when the discount factor satisfies $0 < \gamma < 1$.
- (d) [2 pts] Bayes nets:
- (i) Have an associated directed, acyclic graph.
 - (ii) Encode conditional independence assertions among random variables.
 - (iii) Generally require less storage than the full joint distribution.
 - (iv) Make the assumption that all parents of a single child are independent given the child.
- (e) [2 pts] True or false? If a heuristic is admissible, it is also consistent.
- (f) [2 pts] If we use an ϵ -greedy exploration policy with Q-learning, the estimates Q_t are guaranteed to converge to Q^* only if:
- (i) ϵ goes to zero as t goes to infinity, or
 - (ii) the learning rate α goes to zero as t goes to infinity, or
 - (iii) both α and ϵ go to zero.
- (g) [2 pts] True or false? Suppose X and Y are correlated random variables. Then
- $$P(X = x, Y = y) = P(X = x)P(Y = y|X = x)$$
- (h) [2 pts] When searching a zero-sum game tree, what are the advantages and drawbacks (if any) of using an evaluation function? How would you utilize it?

Q3. [12 pts] Minimax and Expectimax


- (a) [2 pts] Consider the following zero-sum game with 2 players. At each leaf we have labeled the payoffs Player 1 receives. It is Player 1's turn to move. Assume both players play optimally at every time step (i.e. Player 1 seeks to maximize the payoff, while Player 2 seeks to minimize the payoff). Circle Player 1's optimal next move on the graph, and state the minimax value of the game. Show your work.




- (b) [2 pts] Consider the following game tree. Player 1 moves first, and attempts to maximize the expected payoff. Player 2 moves second, and attempts to minimize the expected payoff. Expand nodes left to right. Cross out nodes pruned by alpha-beta pruning.



- (c) [2 pts] Now assume that Player 2 chooses an action uniformly at random every turn (and Player 1 knows this). Player 1 still seeks to maximize her payoff. Circle Player 1's optimal next move, and give her expected payoff. Show your work.



Consider the following modified game tree, where one of the leaves has an unknown payoff x . Player 1 moves first, and attempts to maximize the value of the game.



- (d) [2 pts] Assume Player 2 is a minimizing agent (and Player 1 knows this). For what values of x does Player 1 choose the left action?
- (e) [2 pts] Assume Player 2 chooses actions at random (and Player 1 knows this). For what values of x does Player 1 choose the left action?
- (f) [2 pts] For what values of x is the minimax value of the tree worth more than the expectimax value of the tree?

Q4. [10 pts] n -pacmen search

Consider the problem of controlling n pacmen simultaneously. Several pacmen can be in the same square at the same time, and at each time step, each pacman moves by at most one unit vertically or horizontally (in other words, a pacman can stop, and also several pacmen can move simultaneously). The goal of the game is to have all the pacmen be at the same square in the minimum number of time steps. In this question, use the following notation: let M denote the number of squares in the maze that are not walls (i.e. the number of squares where pacmen can go); n the number of pacmen; and $p_i = (x_i, y_i) : i = 1 \dots n$, the position of pacman i . Assume that the maze is connected.

- (a) [1 pt] What is the state space of this problem?
- (b) [1 pt] What is the size of the state space (not a bound, the exact size).
- (c) [2 pts] Give the tightest upper bound on the branching factor of this problem.
- (d) [2 pts] Bound the number of nodes expanded by uniform cost *tree* search on this problem, as a function of n and M . Justify your answer.
- (e) [4 pts] Which of the following heuristics are admissible? Which one(s), if any, are consistent? Circle the corresponding Roman numerals and briefly justify all your answers.
1. The number of (ordered) pairs (i, j) of pacmen with different coordinates: $h_1 : \sum_{i=1}^n \sum_{j=i+1}^n (p_i \neq p_j)$
(i) Consistent? (ii) Admissible?
 2. $h_2 : \frac{1}{2} \max(\max_{i,j} |x_i - x_j|, \max_{i,j} |y_i - y_j|)$
(i) Consistent? (ii) Admissible?

Q5. [12 pts] CSPs: Course scheduling

An incoming freshman starting in the Fall at Berkeley is trying to plan the classes she will take in order to graduate after 4 years (8 semesters). There is a subset R of required courses out of the complete set of courses C that must all be taken to graduate with a degree in her desired major. Additionally, for each course $c \in C$, there is a set of prerequisites $\text{Prereq}(c) \subset C$ and a set of semesters $\text{Semesters}(c) \subseteq S$ that it will be offered, where $S = \{1, \dots, 8\}$ is the complete set of 8 semesters. A maximum load of 4 courses can be taken each semester.

- (a) [5 pts] Formulate this course scheduling problem as a constraint satisfaction problem. Specify the set of variables, the domain of each variable, and the set of constraints. **Your constraints need not be limited to unary and binary constraints.** You may use any precise and unambiguous mathematical notation.

Variables:

Constraints:

- (b) [4 pts] The student managed to find a schedule of classes that will allow her to graduate in 8 semesters using the CSP formulation, but now she wants to find a schedule that will allow her to graduate in as few semesters as possible. With this additional objective, formulate this problem as an uninformed search problem, using the specified state space, start state, and goal test.

State space: The set of all (possibly partial) assignments x to the CSP.

Start state: The empty assignment.

Goal test: The assignment is a complete, consistent assignment to the CSP.

Successor function: $\text{Successors}(x) =$

Cost function: $\text{Cost}(x, x') =$

- (c) [3 pts] Instead of using uninformed search on the formulation as above, how could you modify backtracking search to efficiently find the least-semester solution?

Q6. [19 pts] Cheating at cs188-Blackjack

Cheating dealers have become a serious problem at the cs188-Blackjack tables. A cs188-Blackjack deck has 3 card types (5,10,11) and an honest dealer is equally likely to deal each of the 3 cards. When a player holds 11, cheating dealers deal a 5 with probability $\frac{1}{4}$, 10 with probability $\frac{1}{2}$, and 11 with probability $\frac{1}{4}$. You estimate that $\frac{4}{5}$ of the dealers in your casino are honest (H) while $\frac{1}{5}$ are cheating ($\neg H$).

Note: You may write answers in the form of arithmetic expressions involving numbers or references to probabilities that have been directly specified, are specified in your conditional probability tables below, or are specified as answers to previous questions.

- (a) [3 pts] You see a dealer deal an 11 to a player holding 11. What is the probability that the dealer is cheating?

The casino has decided to install a camera to observe its dealers. Cheating dealers are observed doing suspicious things on camera (C) $\frac{4}{5}$ of the time, while honest dealers are observed doing suspicious things $\frac{1}{4}$ of the time.

- (b) [4 pts] Draw a Bayes net with the variables H (honest dealer), D (card dealt to a player holding 11), and C (suspicious behavior on camera). Write the conditional probability tables.

- (c) [2 pts] List all conditional independence assertions made by your Bayes net.

- (d) [4 pts] What is the probability that a dealer is honest given that he deals a 10 to a player holding 11 and is observed doing something suspicious?

You can either arrest dealers or let them continue working. If you arrest a dealer and he turns out to be cheating, you will earn a \$4 bonus. However, if you arrest the dealer and he turns out to be innocent, he will sue you for -\$10. Allowing the cheater to continue working will cost you -\$2, while allowing an honest dealer to continue working will get you \$1. Assume a linear utility function $U(x) = x$.

- (e) [3 pts] You observe a dealer doing something suspicious (C) and also observe that he deals a 10 to a player holding 11. Should you arrest the dealer?
- (f) [3 pts] A private investigator approaches you and offers to investigate the dealer from the previous part. If you hire him, he will tell you with 100% certainty whether the dealer is cheating or honest, and you can then make a decision about whether to arrest him or not. How much would you be willing to pay for this information?

Q7. [16 pts] Markov Decision Processes

Consider a simple MDP with two states, S_1 and S_2 , two actions, A and B , a discount factor γ of $1/2$, reward function R given by

$$R(s, a, s') = \begin{cases} 1 & \text{if } s' = S_1; \\ -1 & \text{if } s' = S_2; \end{cases}$$

and a transition function specified by the following table.

s	a	s'	$T(s, a, s')$
S_1	A	S_1	$1/2$
S_1	A	S_2	$1/2$
S_1	B	S_1	$2/3$
S_1	B	S_2	$1/3$
S_2	A	S_1	$1/2$
S_2	A	S_2	$1/2$
S_2	B	S_1	$1/3$
S_2	B	S_2	$2/3$

- (a) [2 pts] Perform a single iteration of value iteration, filling in the resultant Q-values and state values in the following tables. Use the specified initial value function V_0 , rather than starting from all zero state values. Only compute the entries not labeled “skip”.

s	a	$Q_1(s, a)$
S_1	A	
S_1	B	
S_2	A	skip
S_2	B	skip

s	$V_0(s)$	$V_1(s)$
S_1	2	
S_2	3	skip

- (b) [2 pts] Suppose that Q-learning with a learning rate α of $1/2$ is being run, and the following episode is observed.

s_1	a_1	r_1	s_2	a_2	r_2	s_3
S_1	A	1	S_1	A	-1	S_2

Using the initial Q-values Q_0 , fill in the following table to indicate the resultant progression of Q-values.

s	a	$Q_0(s, a)$	$Q_1(s, a)$	$Q_2(s, a)$
S_1	A	$-1/2$		
S_1	B	0		
S_2	A	-1		
S_2	B	1		

- (c) [4 pts] Assuming that an ϵ -greedy policy (with respect to the Q-values as of when the action is taken) is used, where $\epsilon = 1/2$, and given that the episode starts from S_1 and consists of 2 transitions, what is the probability of observing the episode from part b? State precisely your definition of the ϵ -greedy policy with respect to a Q-value function $Q(s, a)$.

- (d) [4 pts] Given an arbitrary MDP with state set S , transition function $T(s, a, s')$, discount factor γ , and reward function $R(s, a, s')$, and given a constant $\beta > 0$, consider a modified MDP (S, T, γ, R') with reward function $R'(s, a, s') = \beta \cdot R(s, a, s')$. Prove that the modified MDP (S, T, γ, R') has the same set of optimal policies as the original MDP (S, T, γ, R) .

- (e) [4 pts] Although in this class we have defined MDPs as having a reward function $R(s, a, s')$ that can depend on the initial state s and the action a in addition to the destination state s' , MDPs are sometimes defined as having a reward function $R(s')$ that depends only on the destination state s' . Given an arbitrary MDP with state set S , transition function $T(s, a, s')$, discount factor γ , and reward function $R(s, a, s')$ that *does depend* on the initial state s and the action a , define an *equivalent* MDP with state set S' , transition function $T'(s, a, s')$, discount factor γ' , and reward function $R'(s')$ that depends only on the destination state s' .

By *equivalent*, it is meant that there should be a one-to-one mapping between state-action sequences in the original MDP and state-action sequences in the modified MDP (with the same value). **You do not need to give a proof of the equivalence.**

States: $S' =$

Transition function: $T'(s, a, s') =$

Discount factor: $\gamma' =$

Reward function: $R'(s') =$