

- You have approximately 2 hours and 50 minutes.
- The exam is closed book, closed calculator, and closed notes except your one-page crib sheet.
- 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	
edX username	
Name of person on your left	
Name of person on your right	

For staff use only:

Q1. Game Trees	/7
Q2. CSP Futoshiki	/11
Q3. Search, logic, and learning	/19
Q4. Probability and Bayes Nets	/20
Q5. Decision Networks and MDPs	/17
Q6. MDPs/RL	/8
Q7. Machine Learning	/18
Total	/100

Q1. [7 pts] Game Trees

(a) Alpha-beta pruning true/false For each true/false question, circle the correct answer. Missing choices and wrong choices with no explanation are worth zero.

(i) [1 pt] [*true* or *false*] Minimax search with alpha-beta pruning may not find a minimax optimal strategy.


(ii) [1 pt] [*true* or *false*] Alpha-beta pruning prunes the same number of subtrees independent of the order in which successor states are expanded.

(iii) [1 pt] [*true* or *false*] Minimax search with alpha-beta pruning generally requires more run-time than minimax without pruning on the same game tree.


(b) [4 pts] For each of the following minimax game trees (max is at the root), fill in the leaf nodes with 0 or 1 as utility values. For the left tree (i), fill in the leaves so that as many leaves will be pruned by alpha-beta as possible. For the right tree (ii), fill in the leaves so that no leaves are pruned. Note that the two trees are slightly different.

Assume (1) left to right traversal while pruning, and (2) leaf values are potentially drawn from $[-\infty, \infty]$ (not from $[0, 1]$).

(i) Maximum Pruning




(ii) No Pruning



Q2. [11 pts] CSP Futoshiki

Futoshiki is a Japanese logic puzzle that is very simple, but can be quite challenging. You are given an $n \times n$ grid, and must place the numbers $1, \dots, n$ in the grid such that every row and column has exactly one of each. Additionally, the assignment must satisfy the inequalities placed between some adjacent squares.

To the right is an instance of this problem, for size $n = 4$. Some of the squares have known values, such that the puzzle has a unique solution. (The letters mean nothing to the puzzle, and will be used only as labels with which to refer to certain squares). Note also that inequalities apply only to the two adjacent squares, and do not directly constrain other squares in the row or column.



Let's formulate this puzzle as a CSP. We will use 4^2 variables, one for each cell, with X_{ij} as the variable for the cell in the i th row and j th column (each cell contains its i, j label in the top left corner). The only unary constraints will be those assigning the known initial values to their respective squares (e.g. $X_{34} = 3$).

- (a) [3 pts] Complete the formulation of the CSP using only binary constraints (in addition to the unary constraints specified above). In particular, describe the domains of the variables, and all binary constraints you think are necessary. You do not need to enumerate them all, just describe them using concise mathematical notation. You are not permitted to use n -ary constraints where $n \geq 3$.

- (b) [2 pts] After enforcing unary constraints, consider the binary constraints relating X_{14} and X_{24} . Enforce arc consistency on just these constraints and state the resulting domains for the two variables.

- (c) [2 pts] Suppose we enforced unary constraints and ran arc consistency on this CSP, pruning the domains of all variables as much as possible. After this, what is the maximum possible domain size for any variable? [*Hint*: consider the least constrained variable(s); you should *not* have to run every step of arc consistency.]

- (d) [2 pts] Suppose we enforced unary constraints and ran arc consistency on the initial CSP in the figure above. What is the maximum possible domain size for a variable adjacent to an inequality?

- (e) [2 pts] By inspection of column 2, we find it is necessary that $X_{32} = 1$, despite not having found an assignment to any of the other cells in that column. Would running arc consistency find this requirement? Explain why or why not.

Q3. [19 pts] Search, logic, and learning

In this question we consider the problem of searching for the smallest propositional logic sentence ϕ that satisfies some condition $G(\phi)$. (E.g., “find the smallest unsatisfiable sentence containing two distinct symbols.”) Recall that a propositional logic sentence is a proposition symbol, the negation of a sentence, or two sentences joined by \wedge , \vee , \Rightarrow , or \Leftrightarrow . The proposition symbols are given as part of the problem. The size of a sentence is defined as the sum of the sizes of its logical connectives, where \neg has size 1 and the other connectives have size 2. (It is helpful to think of the sentence as a syntax tree with proposition symbols at the leaves; then the size is the number of edges in the tree. Figure 1(a) shows an example.)




Figure 1: (a) The sentence $A \wedge (B \vee \neg A)$ drawn as a syntax tree with 5 edges. (b) Part of the “parsimonious” search space for sentences with symbols A and B .

(a) [2 pts] Hank Harvard proposes the following problem formulation to explore the search space of sentences:

- There is a dummy start state λ , which is an empty sentence that never satisfies G . The actions applicable in this state simply replace λ by one of the proposition symbols.
- For all other states, the actions take any occurrence of a proposition symbol (call it p) in the sentence and replace it with either $\neg q$ for any symbol q , or $r * s$ where r and s are any symbols and “*” is any binary connective.

Buzzy Berkeley agrees with Hank that this set of actions will generate all and only the syntactically valid sentences, but proposes her own formulation:

- For all other states, the actions take any occurrence of a proposition symbol (call it p) in the sentence and replace it with either $\neg p$, or $p * q$ where q is any symbol and “*” is any binary connective.

Part of the search space generated by Buzzy’s formulation is shown in Figure 1(b). Buzzy claims that her formulation is more efficient than Hank’s but still generates all and only the syntactically valid sentences. Is she right? Explain.

(b) [1 pt] Define the step cost function for Buzzy’s formulation.

- (c) [1 pt] Assuming there are n symbols, give a $O()$ -expression for the branching factor at depth d of Buzzy's search tree.
- (d) [1 pt] Using your $O()$ answer for the branching factor, give a $O()$ -expression for the number of nodes at depth d of Buzzy's search tree.
- (e) [2 pts] We will say that G is a *semantic* condition if $G(\phi)$ depends only on the *meaning* of ϕ , in the following sense: if two sentences ϕ and ψ are logically equivalent, then $G(\phi) = G(\psi)$. Not wishing to be outdone by Buzzy, Hank now makes the following claim: Whenever Buzzy's search space contains a solution for a semantic G then so does the reduced search space using only \neg , \wedge , and \vee . Is he right? Explain.
- (f) [2 pts] Suppose we are running a uniform cost graph search, which maintains the property that $g(n)$ for every node in the frontier is the optimal cost to reach n from the root. In a standard graph search, we discard a newly generated successor state without adding it to the frontier if and only if the *identical* state is already in the frontier set or explored set. Assuming we have a semantic condition G , when is it possible to discard a newly generated successor state? Check all that apply.
- Sentences that are identical to a sentence already in the frontier or explored set
 - Sentences that logically entail a sentence already in the frontier or explored set
 - Sentences that are logically equivalent to a sentence already in the frontier or explored set
 - Sentences that are logically entailed by a sentence already in the frontier or explored set

Now we will apply this general search machinery for sentences to supervised machine learning: the problem of finding concise hypotheses that are consistent with a set of labeled examples and predict labels for unlabeled examples.

- Let X_1, \dots, X_n be the Boolean input variables and Y be the "output" Boolean property we are trying to predict. A hypothesis H asserts that Y is some function of the inputs, i.e., H is a sentence $Y \Leftrightarrow \phi$, where ϕ is a propositional formula containing only X_1, \dots, X_n . For example, let H_0 be the sentence $Y \Leftrightarrow (X_1 \vee X_2)$; then H_0 asserts that the output is true exactly when either of the inputs X_1, X_2 is true.
- A labeled example E gives the output value for particular values of the inputs; i.e., it is a sentence $\psi \Rightarrow \chi$, where ψ is a conjunction of literals, one for each input symbol, and χ is a literal containing the output symbol. Thus if there are just two inputs, three examples might be

$$\begin{aligned}
 E_1 : & \quad X_1 \wedge X_2 \Rightarrow Y \\
 E_2 : & \quad X_1 \wedge \neg X_2 \Rightarrow Y \\
 E_3 : & \quad \neg X_1 \wedge \neg X_2 \Rightarrow \neg Y
 \end{aligned}$$

(g) [2 pts] Show, in general, that every labeled example sentence is *false in exactly one model* for the symbols X_1, \dots, X_n, Y .

(h) [4 pts] Show, by completing the following truth table and marking the relevant rows, that the hypothesis H_0 given above logically entails all three examples.

X_1	X_2	Y	E_1	E_2	E_3	H_0
F	F	F				
F	F	T				
F	T	F				
F	T	T				
T	F	F				
T	F	T				
T	T	F				
T	T	T				

(i) [2 pts] For supervised machine learning, then, we want to find an H , by searching among formulas ϕ , that entails all the examples. Suppose you have available a solver $SAT(\alpha)$ that returns *true* if α is satisfiable and *false* otherwise. Which of the following expressions correctly implements $G(\phi)$?

- $SAT(H \wedge E_1 \wedge E_2 \wedge E_3)$
 $\neg SAT(H \wedge \neg(E_1 \wedge E_2 \wedge E_3))$
 $SAT(H \wedge \neg(E_1 \wedge E_2 \wedge E_3))$
 $\neg SAT(H \wedge E_1 \wedge E_2 \wedge E_3)$

(j) [2 pts] Now suppose we have an unlabeled test example U_4 described only by the ψ -part:

$$U_4 : \quad \neg X_1 \wedge X_2$$

Now we want to predict the label for U_4 given a hypothesis H . Which of the following does this correctly?

- $SAT(H \wedge U_4 \wedge Y)$
 $\neg SAT(H \wedge (U_4 \rightarrow Y))$
 $\neg SAT(H \vee (U_4 \rightarrow Y))$
 $\neg SAT(H \wedge \neg(U_4 \rightarrow Y))$

Extra Credit [2] (Extra credit, HARD) Any ideas for a *nontrivial* admissible heuristic for the supervised machine learning problem in this setting?

Q4. [20 pts] Probability and Bayes Nets

- (a) [3 pts] $A, B, C,$ and D are Boolean random variables. How many entries are in the following probability tables and what is the sum of the values in each table? Write “?” if there is not enough information given.

Table	Size	Sum
$P(a B)$		
$P(A B, c, D)$		
$P(B, C \neg a)$		

- (b) For each of the following expressions for the joint distributions $P(A, B, C, D)$ write the all of the independence assumptions used:

(i) [2 pts] $P(A, B, C, D) = P(A | B, C)P(D | B, C)P(B)P(C | B)$

(ii) [2 pts] $P(A, B, C, D) = P(B)P(C)P(A | B, C)P(D | A, B, C)$

(iii) [2 pts] $P(A, B, C, D) = P(B)P(C | B)P(A | B)P(D | B, C)$

- (c) Write each of the following expressions in its simplest form, i.e., a single term (factor). Make no independence assumptions.

- (i) [2 pts]

$$\sum_{a'} P(a' | D)P(b | a', D)$$


- (ii) [2 pts]

$$\sum_{b', c', d'} P(A)P(b' | A)P(c' | A, b')P(d' | A, b', c')$$

- (iii) [2 pts]

$$\frac{\sum_{b'} P(A | b', D)P(b' | D)P(D)}{\sum_{a', b'} P(a' | b', D)P(b' | D)P(D)}$$

- (d) [2 pts] Given the following Bayes net, write an expression for the probability distribution $P(A|e)$ using only the conditional probability distributions associated with this Bayes net, e.g. $P(A)$, $P(C|B)$, etc. Normalization constants, such as α or Z are not permitted.



$P(A|e) =$

- (e) Sampling: Our new droid BB-8 is using likelihood weighted sampling to answer various queries on the Bayes net from part (d).

- (i) [1 pt] [true or false] For the query $P(E|b)$, the value sampled for variable A will have no effect on the weight of the complete sample.

BB-8 has implemented a simpler version of likelihood weighted sampling to answer the query $P(E|b)$. BB-8's method skips sampling variable A and skips incorporating the weight associated with B and proceeds to sample values for C , D , and then E from $P(C|B)$, $P(D|B)$, and $P(E|C, D)$, respectively.

- (ii) [1 pt] [true or false] BB-8's simpler sampling method will converge to the same answer for $P(E|b)$ as standard likelihood weighted sampling.

- (iii) [1 pt] Is it possible to make a similar simplification to likelihood weighted sampling to make it more efficient (but still accurate in the limit) when answering the query $P(E|c)$. If so, briefly describe the simplification.

Q5. [17 pts] Decision Networks and MDPs

Typically, we define a Markov Decision Process (MDP) as a set of states S , a set of actions A , a reward function $R(s)$ (note that we don't use $R(s, a, s')$ here), a starting state s_0 , and a transition model $T(s, a, s')$. Assume that our discount factor $\gamma = 1$. In this question, we'll explore whether or not we can use *decision networks* to solve MDPs. Recall that a decision network is a Bayes net augmented with action and utility variables to model decision problems.


Here's the idea of how we'll try to solve an MDP using decision networks: our random variable nodes will model the state over the first $t + 1$ timesteps, i.e., S_0, \dots, S_t , and our action nodes will model the action over the first t timesteps, i.e., A_0, \dots, A_{t-1} . Using this formulation, we can then just use the decision algorithm to figure out the sequence of actions to take from time 0 to $t - 1$. We can repeat this to continue to find actions forever.

First, let's introduce the specific, very simple MDP that we'll work with for the rest of the question. In this MDP, we have two states *init* and *goal*, and two actions *stay* and *move*. The starting state is *init*. The reward and transition functions are as follows:

s	$R(s)$
<i>init</i>	0
<i>goal</i>	1

s	a	s'	$T(s, a, s')$
<i>init</i>	<i>stay</i>	<i>init</i>	1
<i>init</i>	<i>stay</i>	<i>goal</i>	0
<i>init</i>	<i>move</i>	<i>init</i>	0.2
<i>init</i>	<i>move</i>	<i>goal</i>	0.8
<i>goal</i>	<i>stay</i>	<i>init</i>	0
<i>goal</i>	<i>stay</i>	<i>goal</i>	1
<i>goal</i>	<i>move</i>	<i>init</i>	0.8
<i>goal</i>	<i>move</i>	<i>goal</i>	0.2

Let's use $t = 2$ in this question to keep things simple. For this MDP, we draw the decision network below. Note that we do not write out the probability table for S_2 , because it is the same as S_1 shifted over.



s_0	$P(s_0)$
<i>init</i>	1
<i>goal</i>	0

s_0	a_0	s_1	$P(s_1 s_0, a_0)$
<i>init</i>	<i>stay</i>	<i>init</i>	1
<i>init</i>	<i>stay</i>	<i>goal</i>	0
<i>init</i>	<i>move</i>	<i>init</i>	0.2
<i>init</i>	<i>move</i>	<i>goal</i>	0.8
<i>goal</i>	<i>stay</i>	<i>init</i>	0
<i>goal</i>	<i>stay</i>	<i>goal</i>	1
<i>goal</i>	<i>move</i>	<i>init</i>	0.8
<i>goal</i>	<i>move</i>	<i>goal</i>	0.2

s_0	s_1	s_2	$U(s_0, s_1, s_2)$
<i>init</i>	<i>init</i>	<i>init</i>	
<i>init</i>	<i>init</i>	<i>goal</i>	
<i>init</i>	<i>goal</i>	<i>init</i>	
<i>init</i>	<i>goal</i>	<i>goal</i>	
<i>goal</i>	<i>init</i>	<i>init</i>	
<i>goal</i>	<i>init</i>	<i>goal</i>	
<i>goal</i>	<i>goal</i>	<i>init</i>	
<i>goal</i>	<i>goal</i>	<i>goal</i>	

- (a) [4 pts] For the given decision network, fill out the utility table above. *Hint:* Think intuitively about how we measure utility in an MDP.

- (b) [5 pts] Using this decision network that is now complete, let's calculate the expected utilities of all possible action sequences $[A_0, A_1]$. As a starting point, we've calculated the expected utility of $[A_0, A_1] = [move, stay]$ as 1.6, and the expected utility of $[move, move]$ as 1.12.
- (i) [2 pts] What is the expected utility of $[stay, stay]$?
- (ii) [2 pts] What is the expected utility of $[stay, move]$?
- (iii) [1 pt] What is the optimal action sequence $[A_0, A_1]$?
- (c) [4 pts] Does this optimal action sequence give us the same optimal policy that we would obtain if we solved the MDP using value or policy iteration? Why, or why not?
- (d) [4 pts] Suppose we fix $A_0 = move$. Calculate $VPI(S_1)$. Clearly show and label your work for potential partial credit. *Hint:* you may need some of the expected utilities used in part (b).

Q6. [8 pts] MDPs/RL

(a) [5 pts] **Multiple Choice.** Select the single best answer for each question. We are given an MDP (S, A, T, γ, R) , where R is only a function of the current state s . We are also given an arbitrary policy π .

i) If $f(s) = R(s) + \sum_{s'} \gamma T(s, \pi(s), s') f(s')$, then f computes

- V^*
 Q^*
 π^*
 V^π
 Q^π
 None of these

ii) If $g(s) = \max_a \sum_{s'} T(s, a, s') [R(s) + \gamma \max_{a'} Q^*(s', a')]$, then g computes

- V^*
 Q^*
 π^*
 V^π
 Q^π
 None of these

iii) If $h(s, a) = \sum_{s'} T(s, \pi(s), s') [R(s, \pi(s), s') + \gamma h(s', a)]$, then h computes

- V^*
 Q^*
 π^*
 V^π
 Q^π
 None of these

iv) Which of the following iterative MDP-solving techniques typically converges in the fewest number of iterations?

- Asynchronous Value Iteration
 Batch Value Iteration
 Policy Iteration

v) Which of the following reinforcement learning techniques sometimes diverges?

- Exact Q-Learning
 Q-Learning with linear function approximations
 Exact TD-Learning

(b) [3 pts] Consider policy evaluation in a setting where the reward R is a function of s, a, s' , instead of just s . Suppose we have n states, s_1 through s_n . Then for any s , we have the following policy evaluation equation:

$$V^\pi(s) = \sum_{s'} T(s, \pi(s), s') [R(s, \pi(s), s') + \gamma V^\pi(s')].$$

Now, suppose the policy $\pi(s)$ that we are evaluating behaves as follows. At each timestep, it picks one out of m different “local” policies $\pi_1(s), \pi_2(s), \dots, \pi_m(s)$ with corresponding probabilities p_1, p_2, \dots, p_m of being picked. (Note that $p_1 + p_2 + \dots + p_m = 1$.) For this timestep, it acts according to the chosen policy. Write down the policy evaluation equation for $V^\pi(s)$ in terms of the local policies $\pi_1(s), \pi_2(s), \dots, \pi_m(s)$.

Q7. [18 pts] Machine Learning

(a) **Overfitting:** For each case, select the scenario that is more likely to **overfit** to training data.

(i) [2 pts] Decision Tree

- Decision Tree with height 10
 Decision Tree with height 100

(ii) [2 pts] Classifying CS188 students against Non-CS188 student

- Using student id as a feature
 Using GPA as a feature


(iii) [2 pts] Statistical Learning

- Laplace smoothing with $k = 2$
 Laplace smoothing with $k = 5$

(iv) [2 pts] Linear Regression: Using loss function

- $\sum_i^N (y_i - w^T x_i)^2$
 $\sum_i^N (y_i - w^T x_i)^2 + \lambda \sum_i^d w_i^2$

(b) Suppose we have a following training data with class 1, 0 with x_1, x_2 features.



For each true/false question, circle the correct answer. Missing choices and wrong choices with no explanation are worth zero.


(i) [1 pt] [*true* or *false*] We can use one perceptron to perfectly classify this training data.

(ii) [1 pt] [*true* or *false*] We can use a decision tree of depth 2 to perfectly classify this training data.

(iii) [1 pt] [*true* or *false*] According to information gain, x_1 is a better feature to split than x_2 in decision tree.

(c) Multilayer Perceptron


As you can see from the diagram below, we can use x_1 XOR x_2 to perfectly classify the four data points. We aim to construct a neural net to represent this function.



(i) [1 pt] First, express x_1 XOR x_2 in conjunctive normal form (CNF) with x_1 and x_2 as symbols.

Another way to express x_1 XOR x_2 is in DNF as $(x_1 \wedge \neg x_2) \vee (\neg x_1 \wedge x_2)$. We will use this logical expression to represent our neural net.

We create the following neural net representation:



with Σ_1 , Σ_2 , and Σ_3 defined as the dot product of the input values and the weights (e.g., $\Sigma_1 = 1 \cdot -0.5 + x_1 w_1 + x_2 w_3$),

and g_1 , g_2 , and g_3 defined by a threshold function:

$$g(x) = \begin{cases} 1 & x \geq 0 \\ 0 & x < 0 \end{cases}$$

x_1 , x_2 are Boolean (0/1) inputs. The weights for the bias inputs for g_1 and g_2 are already defined as -0.5.

(ii) [6 pts] Using this information, find values for weights $w_1, w_2, w_3, w_4, w_5, w_6$ so that:

- the output of g_1 , a_1 , represents $x_1 \wedge \neg x_2$,
- the output of g_2 , a_2 , represents $\neg x_1 \wedge x_2$, and
- the output of g_3 , a_3 , represents $a_1 \vee a_2 = (x_1 \wedge \neg x_2) \vee (\neg x_1 \wedge x_2) = x_1$ XOR x_2 .

[Hint: it suffices to consider weights in the set $\{+1, -1\}$.]

$w_1 =$	$w_3 =$	$w_5 =$
$w_2 =$	$w_4 =$	$w_6 =$

THIS PAGE IS INTENTIONALLY LEFT BLANK