CS 188 Summer 2023 Introduction to Artificial Intelligence Final

- You have 170 minutes.
- The exam is closed book, no calculator, and closed notes, other than two double-sided cheat sheets that you may reference.
- For multiple choice questions,
  - □ means mark **all options** that apply
  - ○ means mark a **single choice**
- For numerical calculation questions, you may leave your answer unsimplified but **show your work**

First name	
Last name	
SID	
Exam Room	
Name and SID of person to the right	
Name and SID of person to the left	
Discussion TAs (or None)	

Honor code: "As a member of the UC Berkeley community, I act with honesty, integrity, and respect for others."

By signing below, I affirm that all work on this exam is my own work, and honestly reflects my own understanding of the course material. I have not referenced any outside materials (other than two double-sided crib sheet), nor collaborated with any other human being on this exam. I understand that if the exam proctor catches me cheating on the exam, that I may face the penalty of an automatic "F" grade in this class and a referral to the Center for Student Conduct.

Signature: ________________________________

Point Distribution

Q1.	Potpourri: Blast from the Past	10
Q2.	Bayes Nets: Across the Spider-Verse	13
Q3.	HMMs: Head TA Kirby	12
Q4.	Decision Networks: Coin Toss Game	15
Q5.	MDPs: Jim & Pam Part 2	13
Q6.	Reinforcement Learning: Island Hopping	12
Q7.	Machine Learning: Hotdog vs. Not Hotdog	13
Q8.	Neural Nets: Quadratic Activation Function	10
	Total	98

Q1. [10 pts] Potpourri: Blast from the Past

(a) Search. Barbie and Ken are in a building of size $H \times W \times L$. Each cell of the building may or may not contain a toy, and their mission is to collect all of the toys. It is okay for Barbie and Ken to be in the same cell of the building.

(i) [2 pts] Propose a minimal state space representation for this problem.

(ii) [2 pts] What is the size of the state space?

(b) [2 pts] **CSPs.**

Alice is scheduling job interviews. Nine different companies reached out to interview her this coming week, and she is panicked trying to schedule all of them in just five days! For the nine companies, three are big (B_1, B_2, B_3), three are medium (M_1, M_2, M_3), and three are small (S_1, S_2, S_3). Write down her constraints formally below:

Index	Explanation	Constraint
A	You should interview with B_2 on Friday.	$B_2 = 5$
B	You should interview with B_3 on Monday.	
C	You should interview with S_1 on either Monday or Tuesday.	$S_1 \in \{1, 2\}$
E	You should interview with S_2 after S_1 (cannot be on the same day).	
G	You should take at least two days break after M_2 before M_3 (If M_2 occurs on Monday, the earliest M_3 can occur is Thursday).	
H	You should interview with M_3 after B_1 (cannot be on the same day), since they have the same interview style.	

(c) [4 pts] **Games.**

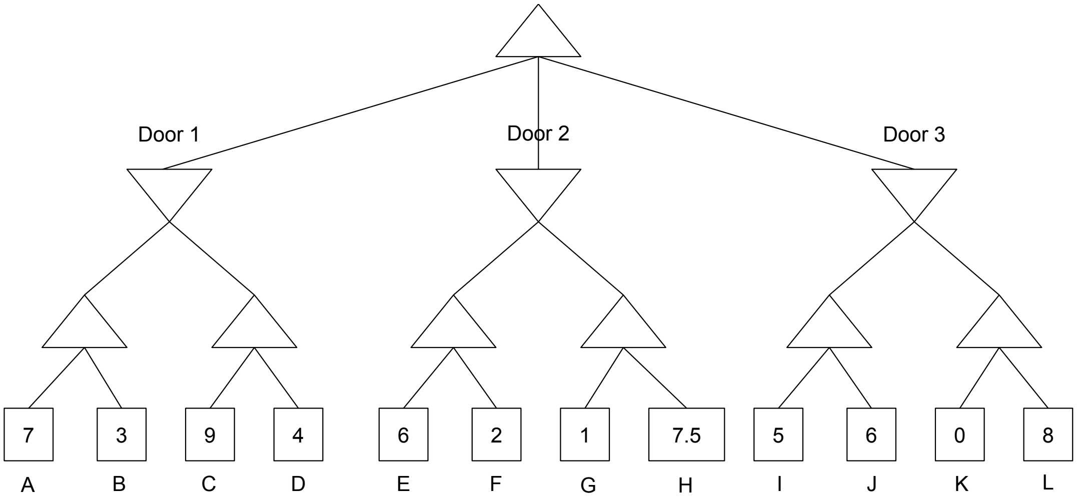

(i) [2 pts] Fill out the values on the game tree. Which door does the top maximizer node choose?

○ Door 1 ○ Door 2 ○ Door 3

(ii) [2 pts] Which terminal nodes are never explored as a consequence of pruning? (Assume that we prune on equality.)

□ A □ B □ C □ D □ E □ F □ G □ H □ I □ J □ K □ L

○ None of the above

Q2. [13 pts] Bayesian Networks: Across the Spider-Verse

Miles is curious about the probability that he arrives to school on time. The factors involved with him arriving to school on time can be represented by the following Bayes Net (assume that each variable is a binary variable):

- A: Sets alarm
- S: Over sleeps
- D: Dad is late
- C: Fighting crime last night
- T: Arrives at school on time

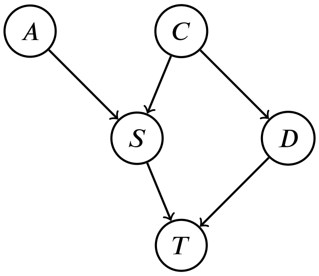

(a) [7 pts] Miles wants to calculate $P(C|T)$ using variable elimination. Assume he eliminates variables in alphabetical order (A, D, S).

(i) [1 pt] What factors does he have available at the start?

(ii) [1 pt] First, he eliminates A, and get the new factor

$f_1(\quad) =$

Write out the remaining factors

(iii) [1 pt] Then, he eliminates D, and get the new factor

$f_2(\quad) =$

Write out the remaining factors

(iv) [1 pt] Then, he eliminates S, and get the new factor

$f_3(\quad) =$

Write out the remaining factors

(v) [1 pt] Finally, join any remaining factors to calculate

$f_4(\quad) =$

(vi) [1 pt] How can he use this to calculate $P(C = +c|T = -t)$? Your answer should be in terms of f_4.

$P(C = +c|T = -t) =$

(vii) [1 pt] Order factors f_1, f_2, and f_3 in increasing order of size. □ < □ < □

The following CPTs correspond to the Bayes Net above:

A	C	S	$P(S\|A,C)$
$+a$	$+c$	$+s$	$11/16$
$+a$	$+c$	$-s$	$5/16$
$+a$	$-c$	$+s$	$1/8$
$+a$	$-c$	$-s$	$7/8$
$-a$	$+c$	$+s$	$9/10$
$-a$	$+c$	$-s$	$1/10$
$-a$	$-c$	$+s$	$1/5$
$-a$	$-c$	$-s$	$4/5$

T	D	S	$P(T\|S,D)$
$+t$	$+d$	$+s$	$1/20$
$+t$	$+d$	$-s$	$2/5$
$+t$	$-d$	$+s$	$1/5$
$+t$	$-d$	$-s$	1
$-t$	$+d$	$+s$	$19/20$
$-t$	$+d$	$-s$	$3/5$
$-t$	$-d$	$+s$	$4/5$
$-t$	$-d$	$-s$	0

C	D	$P(D\|C)$
$+c$	$+d$	$1/4$
$+c$	$-d$	$3/4$
$-c$	$+d$	0
$-c$	$-d$	1

C	$P(A)$
$-c$	$4/5$
$+c$	$1/5$

A	$P(A)$
$-a$	$1/4$
$+a$	$3/4$

Miles is now interested in calculating $P(C = +c|T = -t)$ via sampling. He generates the following random samples (assume the variables were generated from left to right):

Sample	A	C	S	D	T
1	$+a$	$-c$	$-s$	$-d$	$+t$
2	$+a$	$-c$	$+s$	$-d$	$-t$
3	$+a$	$+c$	$+s$	$-d$	$-t$
4	$-a$	$-c$	$-s$	$+d$	$-t$
5	$-a$	$-c$	$+s$	$+d$	$-t$

(b) [3 pts] Assuming Miles uses prior sampling:

(i) [1 pt] Bubble in the samples that Miles uses to calculate the final probability.

□ 1 □ 2 □ 3 □ 4 □ 5

(ii) [2 pts] What is the probability he calculates via prior sampling?

$P(C = +c|T = -t) =$ □

(c) [3 pts] Now assuming Miles uses likelihood weighting:

(i) [1 pt] What weight does Miles assign to each sample?

Sample 1: □ Sample 2: □ Sample 3: □

Sample 4: □ Sample 5: □

What final probability does he calculate?

(ii) [2 pts] What is the probability he calculates via likelihood weighting?

$P(C = +c|T = -t) =$ □

Q3. [12 pts] HMMs: Head TA Kirby

Kirby is serving as a TA. He wants to evaluate his teaching performance after each of his weekly discussion sections, and he does so based on how much his students collaborate during his section. He models this situation using an HMM.

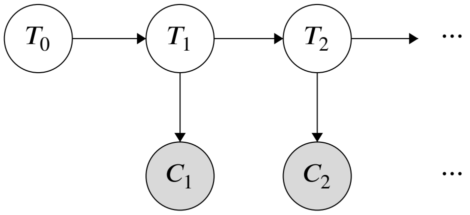

T_{i+1}	T_i	$P(T_{i+1}\|T_i)$
+t	+t	0.7
-t	+t	0.3
+t	-t	0.4
-t	-t	0.6

T_0	$P(T_0)$
+t	0.6
-t	0.4

C_i	T_i	$P(C_i\|T_i)$
+c	+t	0.5
-c	+t	0.5
+c	-t	0.1
-c	-t	0.9

T_i is a binary random variable representing whether Kirby taught sufficiently well during Week i. C_i is another binary random variable representing whether students collaborated during Week i. He does not see his students' collaboration during Week 0.

(a) [8 pts] Using the two steps of the forward algorithm, calculate the distribution of $P(T_1|C_1 = +c)$. For the sake of organization, you may use the first box for the **time elapse** update and the second box for the **observation** update. You may also leave your final answers unsimplified (for example, as fractions).

Time elapse update:

Observation update:

$P(T_1 = +t|C_1 = +c) =$

$P(T_1 = -t|C_1 = +c) =$

(b) [4 pts] In order to save computational resources, Kirby turns to particle filtering to analyze this HMM.

(i) [2 pts] At timestep $t = 3$, Kirby has observed the following evidence: $C_1 = +c$, $C_2 = -c$, and $C_3 = +c$. Following the particle filtering algorithm, assign weights to particles in the following states at $t = 3$:

Particles in state $+t$ will have weight: □

Particles in state -t will have weight: □

(ii) [2 pts] At timestep $t = 6$, we observe 3 particles in state $+t$ and 5 particles in state $-t$, and $C_6 = -c$. Fill in the table describing the distribution that we resample our new particles from for $t = 7$. Show any work in the box on the left.

T_7	$P(T_7)$
+t	
-t	

Q4. [15 pts] Decision Networks: Coin Toss Game

Alice and Bob are participating in a coin toss game. Both players toss two fair coins. Bob's coins are revealed first, after which he must choose to "continue" or "concede". If he concedes, Bob loses \$2 and the game ends. If he continues, Alice's coins are revealed. If Bob's coin toss resulted in strictly more heads than Alice's, he wins \$10. However, if he has an equal or lesser number of heads than Alice, Bob loses \$10. Assume Bob is rational and his utility is the amount of money he wins.

(a) [2 pts] Sketch the decision network for the game, use the following nodes and node types.

Nodes: B represents the number of heads in Bob's coins; A represents the number of heads in Alice's coins; C represents Bob's choice to continue or concede; U represents Bob's utility.

Node types: An elliptical node represents a chance node; a rectangular node represents an action node; a diamond-shaped node represents utility.

C U

B A

(b) [3 pts] For each scenario where Bob gets 0, 1, or 2 heads, what should Bob's decision be: to "continue or "concede"? Justify your answer.

(i) [1 pt] If Bob gets 0 heads:

(ii) [1 pt] If Bob gets 1 heads:

(iii) [1 pt] If Bob gets 2 heads:

(c) [3 pts] What is the expected monetary gain or loss for Bob in a game?

(d) [2 pts] Bob perceives the game as unfair and refuses to participate. To persuade him, Alice proposes an additional rule: Bob can pay Alice c dollars ($c > 0$) to see one of Alice's coins *before seeing his own coins*. **Sketch the decision network for the modified coin toss game.**

New node: R represents the revealed coin, either head ($R = H$) or tail ($R = T$). S represents Bob's decision of whether to see one of Alice's coins.

C U S

B A R

(e) [2 pts] Following the previous part, calculate the conditional distribution of A given R and fill in the table:

R	$\Pr(A = 0 \mid R)$	$\Pr(A = 1 \mid R)$	$\Pr(A = 2 \mid R)$
H			
T			

(f) [3 pts] (**Extra Credit!!!!!**, do this last) Following the previous two parts, Alice secretly insists that this new rule should not impact her expected gain. What should the value of c be to meet this condition? Justify your answer.

Q5. [13 pts] MDPs: Jim & Pam Part 2

Jim and Pam are on a $1x5$ grid, where Jim starts at square 1, and Pam is fixed at square 5.

Jim				Pam

In each time step, Jim chooses to either move right or to rest. Choosing to move succeeds with probability p and fails with probability $1-p$, in which case Jim stays in his original square (Jim received 0 utility regardless of success or failure). Choosing to rest always succeeds, and gives $R(d) = 4^{5-d}$ utility, where d is the distance between Jim and Pam. For example, at the start, where $d = 4$, if Jim decides to rest, he gets 4 utility. We represent this as an infinite horizon MDP with no terminal state.

(a) [3 pts] Jim is considering two policies:

Policy 1: Rest at the start forever.

Policy 2: Attempt to move right once, and then, regardless of success or failure, rest forever.

Assuming that Jim starts in square 1, for what values of p is Policy 1 superior to Policy 2 when the discount factor $\gamma = 0.5$? Hint: the sum S of an infinite geometric series with starting value a and ratio r is $S = \frac{a}{1-r}$

Show your work. $0 \leq \square \leq p \leq \square \leq 1$

(b) [3 pts]

Now assume that $p = 1$. Still assuming that Jim starts in square 1, for what values of γ is Policy 1 superior to Policy 2?

Show your work. $0 \leq \square < \gamma < \square \leq 1$

(c) [7 pts] For the following subparts, assume that $\gamma = 0.5$ and $p = 1$.

(i) [3 pts] Perform two iterations of value iteration, for the following locations of Jim. Show your work.

$States$	$s_J = 1$	$s_J = 2$	$s_J = 3$	$s_J = 4$	$s_J = 5$
V_0	0	0	0	0	0
V_1					
V_2					

(ii) [3 pts] Perform two iterations of policy iteration for the following locations of Jim.

$States$	$s_J = 1$	$s_J = 2$	$s_J = 3$	$s_J = 4$	$s_J = 5$
π_i	$a_J = right$	$a_J = rest$	$a_J = rest$	$a_J = rest$	$a_J = rest$
V^{π_i}					
π_{i+1}					
$V^{\pi_{i+1}}$					
π_{i+2}					

(iii) [1 pt] Assuming that policy iteration has converged, Jim argues it isn't guaranteed values have converged yet, so they need to run value iteration to get the correct values. Pam agrees that policy convergence doesn't guarantee value convergence, but thinks that we don't need to switch to value iteration, as if we continue running policy iteration, eventually the values will converge as well. Who is correct and why?

○ Jim ○ Pam

Q6. [12 pts] Reinforcement Learning: Island Hopping

Bob wants to traverse different islands to reach the island X. He formulates the problem as an MDP where the islands (nodes) represent the states and the arrows represent the possible actions he can take across the seas. Use the direction of the arrows (up, down, left, right) to refer to the specific actions that can be taken.

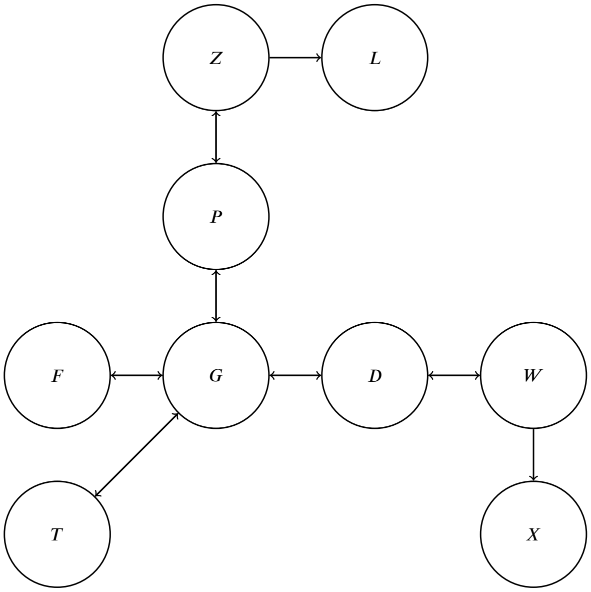

(a) [2 pts] Bob's ship doesn't always move in the direction that he wants it to. Despite this, he wants to use MLE to build an estimate of the transition function $\hat{T}$ and the reward function $\hat{R}$ for model-based reinforcement learning. He follows some specified policy and collects some data in the form of (current state, action, next state, reward) tuples shown below:

State (s)	Action (a)	New State (s')	Reward
F	Right	G	20
D	Right	G	-10
G	Up	T	-30
P	Down	Z	-15
G	Right	D	30
W	Down	D	-25
G	Right	D	30
D	Left	G	-5
G	Right	T	-30
W	Down	X	100

(i) [1 pt] What is $\hat{T}$(G, Right, D)?

(ii) [1 pt] What is $\hat{R}$(W, Down, D)?

(b) [3 pts] Bob looks to use temporal difference learning to learn the values of π, where he has the following initial values:

s	F	G	T	P	Z	L	D	W	X
$V^{\pi}(s)$	5	10	-4	8	2	-10	10	30	50

He performs one update step using the sample (G, Right, D, 30). Assume a discount factor $\gamma = 0.5$ and the learning rate $\alpha = 0.2$. What is the updated value of $V^{\pi}(G)$? Show all your work leading to your answer.

(c) [2 pts] Bob wants to have a greedy policy that will minimize exploration and maximize exploitation. Which of the following functions f will do so? Assume k is a positive real number, $N(s, a)$ represents the number of times that the state-action pair is taken, and ϵ is a very small number greater than 0.

□ $f(s, a) = Q(s, a) + \frac{k}{N(s,a)}$

□ $f(s, a) = Q(s, a) + k \cdot e^{N(s,a)}$

□ $f(s, a) = k \cdot Q(s, a) - log(N(s, a))$

□ $f(s, a) = \frac{\epsilon}{k \cdot Q(s,a) \cdot N(s,a)}$

□ $f(s, a) = \frac{N(s,a) \cdot Q(s,a)}{\epsilon}$

○ None of the above

(d) [5 pts] Bob now switches to Q-learning, where he wants to perform approximate Q-learning for $Q(G, right)$. Assume he has w_i, which denotes the ith value of a weight vector w and $f_i(s, a)$, which denotes the value of the ith feature of the Q-state (s, a). He has the following values and observations:

State (s)	Action (a)	New State (s')	Reward (r)
G	Right	D	10
P	Up	Z	1

w_1	w_2	w_3
2	5	10

$f_1(G, Right)$	$f_2(G, Right)$	$f_3(G, Right)$
6	3	4

State	P	Z	D
Q(State, Up)	2	0	0
Q(State, Down)	5	7	0
Q(State, Left)	0	0	2
Q(State, Right)	0	-8	22

(i) [3 pts] What is the initial value of $Q(G, Right)$ based on the above weights and features?

(ii) [2 pts] What is the resulting weight vector after performing the first iteration of the weight update rule for going right on G? This time, assume a discount factor $\gamma = 0.5$ and learning rate $\alpha = 0.5$.

Q7. [13 pts] Machine Learning: Hotdog vs. Not Hotdog

Bob is building a model to classify whether a picture contains a Hotdog or not. He uses two binary features: whether the picture has brown color in it and whether there is red color in it. He collects this training set:

Brown Color (W_1)	**Red Color (W_2)**	**Label (y)**
1	0	not hotdog
1	0	not hotdog
1	0	not hotdog
1	1	not hotdog
1	1	hotdog
0	0	hotdog

(a) [5 pts] He first builds a Naive Bayes model. Calculate the following probabilities.

(i) [1 pt] $P(y = \text{hotdog}) =$ □ $P(y = \text{not hotdog}) =$ □

(ii) [4 pts] $P(W_1 = 1 \mid y = \text{hotdog}) =$ □ $P(W_1 = 0 \mid y = \text{hotdog}) =$ □

$P(W_2 = 1 \mid y = \text{hotdog}) =$ □ $P(W_2 = 0 \mid y = \text{hotdog}) =$ □

$P(W_1 = 1 \mid y = \text{not hotdog}) =$ □ $P(W_1 = 0 \mid y = \text{not hotdog}) =$ □

$P(W_2 = 1 \mid y = \text{not hotdog}) =$ □ $P(W_2 = 0 \mid y = \text{not hotdog}) =$ □

(b) [3 pts] Next, he uses the model to classify three pictures that are from the test set. Fill in the predicted labels in the table.

Test set

Brown Color (W_1)	**Red Color (W_2)**	**Predicted Label ($\hat{y}$)**
1	1	
0	1	
0	0	

(c) [5 pts] Bob then adds two new examples to his training set, as shown below.

Additional training examples

Brown Color (W_1)	**Red Color (W_2)**	**Label (y)**
0	0	not hotdog
0	1	not hotdog

(i) [3 pts] Now re-classify the test set using the new larger training set (hint: don't forget to update the prior for each class with the new dataset).

Test set

Brown Color (W_1)	**Red Color (W_2)**	**Predicted Label ($\hat{y}$)**
1	1	
0	1	
0	0	

(ii) [2 pts] Did any of the predictions change? Why?

Q8. [10 pts] Neural Networks: Quadratic Activation Function

Consider the following neural network, where the input is $x \in \mathbb{R}^d$ and the output is a real number prediction $\hat{y} \in \mathbb{R}$. We have two neural network weight matrices: $W_1 \in \mathbb{R}^{n_1 \times n_2}$ and $W_2 \in \mathbb{R}^k$, where k, d are known numbers and you need to find out the value of n_1 and n_2. We consider the following neural network:

$$z = W_1 x; \quad \hat{y} := W_2^T g(z) \tag{1}$$

with a quadratic activation function:

$$g(z) := z \odot z \tag{2}$$

where $\odot$ means element-wise multiplication. For example, $g(\begin{bmatrix}1\\2\\3\end{bmatrix}) = \begin{bmatrix}1\\4\\9\end{bmatrix}$.

We consider a quadratic loss function:

$$\mathcal{L}(W_1; W_2) := \frac{1}{2}(y - \hat{y})^2 \tag{3}$$

(a) [2 pts] What's the value of n_1 and n_2 if the above neural network is valid? You can write the answer in terms of k and d

(i) [1 pt] $n_1 =$ □

(ii) [1 pt] $n_2 =$ □

(b) [4 pts] Calculate the derivative of the loss function with respect to the following quantities. Your answer should **NOT** include the activation function g (i.e., you need to explicitly calculate the derivative of g rather than writing g'.) On the other hand, your answer can include x, y, $\hat{y}$, W_1, W_2, and z.

(i) [2 pts] $\frac{\partial \mathcal{L}}{\partial W_2} =$ □

(ii) [2 pts] $\frac{\partial \mathcal{L}}{\partial x} =$ □

(c) [4 pts] We have realized that the model is not getting the accuracy that we were hoping for. For each of the following possible solutions below, answer yes or no if they could possibly improve the model accuracy.

(i) [1 pt] If the model is overfitting, we can try to make it more complex by increasing the number of layers. □

(ii) [1 pt] If the model is overfitting, we can try to make it simpler by decreasing the number of layers. □

(iii) [1 pt] We could try to improve accuracy by getting more training data. □

(iv) [1 pt] We could try out different types of models instead of neural networks (e.g., Naive bayes) and see which one works best on the validation set. □