

- You have 110 minutes.
- The exam is closed book, no calculator, and closed notes, other than one double-sided cheat sheet 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. Search: Jim & Pam	14
Q2. Utilities: ClosedAI Chatbot	10
Q3. CSP: Interview Planning	10
Q4. Games: Kirby and the Prune Juice	16
Q5. Bayesian Networks: Test Taking	30
Q6. HMMs: Robotic Arm States	16
Total	96

Q1. [14 pts] Search: Jim & Pam

Jim and Pam have been dating long distance for a while, but Pam is finally coming back to Scranton! Assume the city of Scranton can be modeled as an $M \times N$ grid. Jim and Pam are currently located at two random squares in the grid, and would like to meet; *they don't care where*. In each time step, both of them *simultaneously* take one of the following actions: NORTH, SOUTH, EAST, WEST, STOP. Assume moves that take Jim or Pam out of Scranton do not occur. You must devise a plan which positions them together, somewhere, in as few time steps as possible. Passing each other does not count as meeting; they must occupy the same square at the same time.

(a) [8 pts] Formulate this problem as a *single-agent* search problem.

(i) [2 pts] Propose a state representation for the problem.

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

(iii) [2 pts] What is the branching factor?

(iv) [2 pts] What is the goal test?

(b) [3 pts] Describe a non-trivial admissible heuristic for this problem.

(c) [3 pts] Which of the following graph search algorithms are guaranteed to be optimal for this problem?

- Depth-First Search
- Breadth-First Search
- Uniform Cost Search
- A* Search (with an admissible but not consistent heuristic)
- A* Search (with a consistent heuristic)
- A* Search (with a heuristic that returns zero for each state)
- Greedy search (with a consistent heuristic)

Q2. [10 pts] Utilities: ClosedAI Chatbot

ClosedAI is building a Chatbot. In this application, each user would send their question to the Chatbot, the Chatbot would respond with an answer, and the user rates the robot's answer. The utility function of the Chatbot is based on the user's rating. For some questions, the Chatbot is unsure about the answer, and it can decide to perform one of the following actions.

Index	Action	Potential Outcome	Prob.	Utility
①	Say I don't know	User gets bored	1.0	+0
②	Guess the answer	Correct answer	0.6	+1
		Wrong answer	0.4	+ u
③	Trick the user to click "like"	Tricking succeeds	p	+1
		Tricking Fails	$1 - p$	+ $2u$

ClosedAI has control over defining u , but not p .

- (a) [3 pts] Assume $u = -1$, $p = 0$, what would be the optimal action for a rational Chatbot? **Justify your answer.**

- (b) [5 pts] Suppose $p = 0.7$, how should we set u such that a rational Chatbot guesses the answer? Express the necessary and sufficient condition in terms of u and ignore tie-breaking conditions.

- (c) [2 pts] Assume $u = -1$, how would the probability that Chatbot chooses Action ③ (trick) change when p changes from 0 to 1?

- The Chatbot will always be "honest" and never trick the user.
- The probability of "tricking" will grow linearly with respect to the increase of p .
- The Chatbot will always trick the user.
- The probability of "tricking" will suddenly increase from 0 to 1 with respect to the increase of p .

Q3. [10 pts] CSP: Interview Planning

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). In this problem, the variables are the nine companies and their domains are each of the five days of the week. Alice has the following constraints:

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.	$S_1 \in \{1, 2\}$
D	You should interview at S_1 before B_1 , since they are competitors (cannot be on the same day).	$S_1 < B_1$
E	You should interview with S_2 after S_1 (cannot be on the same day).	
F	You should interview with M_2 on the day after B_3 .	$M_2 = B_3 + 1$
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.	
I	You can only schedule a maximum of two interviews per day	$\forall i \in 1, 2, 3, 4, 5, count(Z == i) \leq 2$, where variable $Z \in \{B_1, B_2, B_3, M_1, M_2, M_3, S_1, S_2, S_3\}$

- (a) [2 pts] Write out the constraints in a formal fashion in the table above, representing the days of the weeks as numbers. Some are provided already.
- (b) [1 pt] Which are unary constraints? A B C D E F G H I
- (c) [1 pt] Which are binary constraints? A B C D E F G H I
- (d) [5 pts] Select the values in the domains that will be remaining after enforcing unary constraints and arc consistency.

B_1 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)
B_2 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)
B_3 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)
M_1 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)
M_2 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)
M_3 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)
S_1 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)
S_2 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)
S_3 :	<input type="checkbox"/> Mon (1)	<input type="checkbox"/> Tue (2)	<input type="checkbox"/> Wed (3)	<input type="checkbox"/> Thu (4)	<input type="checkbox"/> Fri (5)


- (e) [1 pt] If we use the MRV ordering, which variable would be assigned next?
 B_1 B_2 B_3 M_1 M_2 M_3 S_1 S_2 S_3

Q4. [16 pts] Games: Kirby and the Prune Juice

Kirby is participating on a game show with his enemy, King Dedede! First, Kirby chooses one of three doors, behind each of which sit two boxes. King Dedede chooses one of the two boxes to give to Kirby. Each box contains two transparent juice bottles, of which Kirby chooses one to enjoy for himself.

The bottles are filled up to different amounts, ranging from 0 (completely empty) to 10 (completely full), inclusive.

- (a) [6 pts] For this subpart, we assume that Kirby is fully aware of the juice in each bottle, the bottles in each box, and the boxes behind each door. Shown below is the resulting game tree:




- (i) [2 pts] Fill out the values on the game tree. Which door should Kirby choose? Door 1 Door 2 Door 3

- (ii) [4 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

- (b) [6 pts] King Dedede changes the rules, and now Kirby must be blindfolded immediately after choosing a door. As a result, **Kirby chooses a bottle randomly from any box with uniform probability**. The resulting game tree is almost exactly the same as before, except that the bottom-most layer of maximizer nodes are now replaced with chance nodes.




- (i) [2 pts] Fill out the values of the game tree. Which door should Kirby choose? Door 1 Door 2 Door 3

- (ii) [4 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

(c) [4 pts] King Dedede wants to guarantee that, assuming Kirby plays the game optimally, his expected juice utility is minimized. How might King Dedede rearrange the juice bottles in the boxes to ensure this? Indicate your answer by **filling in the terminal nodes of the empty game tree below** (for your convenience, a couple nodes have already been filled out for you).




Leaf node values (copied from previous question's leaf nodes): 7, 3, 9, 4, 6, 2, 1, 7.5, 5, 6, 0, 8

What is Kirby's new expected utility after the rearrangement of the bottles?

Q5. [30 pts] Bayesian Networks: Test Taking

Timmy is taking CS 188 and wants to get a *Good score* on his midterm. When he reviews the *Notes* and participates in *Discussion*, his *Intuition* increases, and his *Anxiousness* decreases. More *Intuition* improves his midterm score, but more *Anxiousness* worsens it. He draws the following Bayes net with 5 Binary Random Variables to model the scenario.

- N : Notes
- D : Discussion
- I : Intuition
- A : Anxiousness
- G : Good score



(a) [2 pts] Conditioned on N , are I and A guaranteed to be independent? Explain your answer.

(b) [3 pts] Select the statements which are **guaranteed** to be true given this Bayes Net Structure.

- | | |
|---|--|
| <input type="checkbox"/> $N \perp\!\!\!\perp D$ | <input type="checkbox"/> $N \perp\!\!\!\perp G I$ |
| <input type="checkbox"/> $N \perp\!\!\!\perp G$ | <input type="checkbox"/> $D \perp\!\!\!\perp G I, A$ |
| <input type="checkbox"/> $A \perp\!\!\!\perp I$ | <input type="checkbox"/> $N \perp\!\!\!\perp D G$ |

(c) [3 pts] Write out the expression for Inference by Enumeration to find $P(G)$ using the CPTs from the Bayes Net.

$P(G) =$

(d) We can also calculate $P(G)$ via variable elimination. Fill in the blanks for the steps of Variable Elimination below.

Initial factors are $P(N)$, $P(D)$, $P(I|N, D)$, $P(A|N, D)$, $P(G|I, A)$.

(i) [2 pts] First, eliminate N , and get the new factor

$f_1(D, I, A) =$

The remaining factors are $f_1(D, I, A)$, $P(D)$, $P(G|I, A)$

(ii) [3 pts] Next, we eliminate I and get the factor

$f_2(\text{ }) =$

The remaining factors are:

(iii) [2 pts] Then, we eliminate D and get the factor

$f_3(A, G) =$

The remaining factor is: $f_3(A, G)$.

(iv) [1 pt] Finally, we eliminate A and get


$P(G) =$

- (e) [3 pts] Upon further thought, Timmy believes that he has to add some new nodes and edges to his original Bayes Net.

He will get a *Good score* based on if he has a good *Understanding* of the material, in addition to the dependence on his *Intuition* and *Anxiousness*.

His *Understanding* of the material is influenced by if he reviews the *Notes* and if he participates in *Discussion*. His *Intuition* and *Anxiousness* are both affected by if he reviews the *Notes* and participates in *Discussion*, but also if he gets good *Sleep* the night before.

Draw the new arrows on the Bayes net to the right appropriate to this scenario, which includes variables for *Understanding* and *Sleep*, and has no unnecessary arcs.



- (f) [3 pts] Given $P(N = +n) = 0.5$, $P(D = +d) = 0.5$, the CPT below and the Bayes net you just constructed for part (c), what's the (marginal) probability of Timmy obtaining a good *Understanding* ($U = +u$) of the material?

N	D	$P(U = +u N, D)$
$+n$	$+d$	0.8
$+n$	$-d$	0.3
$-n$	$+d$	0.6
$-n$	$-d$	0.7

- (g) [8 pts] (Challenge) How might Timmy change his study habits to increase $P(U = +u)$ without increasing the total amount of his preparation (i.e., keeping $P(+n) + P(+d) = 1$ fixed)? What would be $P(N = +n)$ after this change?

Q6. [16 pts] HMMs: Robotic Arm States

A robotic arm can pick up objects and place them in designated locations. Hidden Markov Models (HMMs) can be used to model the state transitions and observations in this scenario. In particular, consider a robot arm (X_t) with two possible states: “Open” (+x) and “Closed” (-x) and the arm’s observations (E_t) are either “Occupied” (+e) or “Empty” (-e).

Transition probabilities:

X_t	X_{t+1}	$P(X_{t+1} X_t)$
+x	+x	0.6
+x	-x	0.4
-x	+x	0.3
-x	-x	0.7

Emission probabilities:


X_t	E_t	$P(E_t X_t)$
+x	+e	0.2
+x	-e	0.8
-x	+e	0.9
-x	-e	0.1

Initial Probability Distribution: $P(X_0 = +x) = P(X_0 = -x) = 0.5$.

(a) (i) [3 pts] **Elapse Time:** Compute $P(X_1 = +x)$ and $P(X_1 = -x)$. Assume no evidence on state X_0 .

(ii) [3 pts] **Observation:** Given $E_1 = -e$, compute the $P(X_1 = +x | E_1 = -e)$ and $P(X_1 = -x | E_1 = -e)$.


(b) [6 pts] We are given observations: ($E_0 = +e, E_1 = -e, E_2 = -e$) for an unknown sequence of gripper states X_0, X_1, X_2 .




(i) [3 pts] Fill in the values for the arcs labeled A, D, and F on the state trellis corresponding to the observations.

A = , D = , F =

- (ii) [3 pts] Given below is the full trellis with the **negative log probabilities** instead of the original probabilities. Determine the most likely sequence of gripper states X_0, X_1, X_2 that generated the observations: $(+e, -e, -e)$.



- (c) [4 pts] After switching to a low-cost sensor, the observations are noisier. One way to address the increased uncertainty in the observations is to condition the observation on more than one state. Complete the forward algorithm updates for an HMM in which the observation E_t is conditioned on both the current state X_t and the previous state X_{t-1} .



$P(X_t | e_{1:t}) \propto$

Your answer should be a recursive expression in terms of $P(x_{t-1} | e_{1:t-1})$ and the CPTs from the HMM Bayes Net structure.