

- You have approximately 170 minutes.
- The exam is open book, open calculator, and open notes.
- For multiple choice questions,
  - ☐ means mark **all options** that apply
  - ☐ means mark a **single choice**

First name	
Last name	
SID	

**For staff use only:**

Q1.	Potpourri	/20
Q2.	Model-Based RL with Function Approximation	/14
Q3.	Naive Bayes and Perceptron	/18
Q4.	Backpropagation with Activation Checkpointing	/18
Q5.	Ace King Queen	/16
Q6.	Pure Romance	/20
Q7.	Games	/18
Q8.	Hidden Markov Models and Particle Filtering	/15
Total		/139

THIS PAGE IS INTENTIONALLY LEFT BLANK

# Q1. [20 pts] Potpourri

(a) [3 pts] Which of the following statements are always true?

- ☐  $P(X, Y) = \sum_a P(X, a) \sum_b P(Y, b)$   
☐  $P(X) = \sum_a \sum_b \sum_c \sum_d P(X, a, b, c, d)$   
☐  $P(X_1, X_2, \dots, X_n) = P(X_1) \prod_{i=2}^n P(X_i | X_{i-1})$   
☐  $P(X) \propto \sum_Y P(X|Y)$   
☐  $P(X|Y) = \frac{P(X,Y)}{\sum_Y P(X,Y)}$   
☐  $P(X|y) \propto P(y|X)P(X)$

(b) [1 pt] Oski trains a neural network to classify whether or not a student is from Stanford. He notices that his classifier gets high accuracy when he tests it on his friends at Berkeley, but low accuracy when he visits Stanford. Which of the following is the best reason for why this is happening?

- ☐ Oski used a learning rate that was too low which led classifier to be stuck in a local minimum  
☐ Oski added too much regularization when training his model  
☐ Oski's training data has disproportionately more Berkeley examples than Stanford examples  
☐ Oski is incorrectly calculating the accuracy

(c) [1 pt] Regina is trying to perform gradient descent on a function  $f(x)$  using the following update rule:

$$x = x - \frac{\partial f}{\partial x}(x)$$

Is gradient descent guaranteed to converge to the global minimum for any  $f(x)$ ?

- ☐ Yes, since she's updating using the gradient of  $x$ .  
☐ Yes, but not for the reason above.  
☐ No, since she is updating  $x$  in the wrong direction.  
☐ No, but not for the reason above.

(d) [3 pts] Which of the following statements regarding VPI are always true?

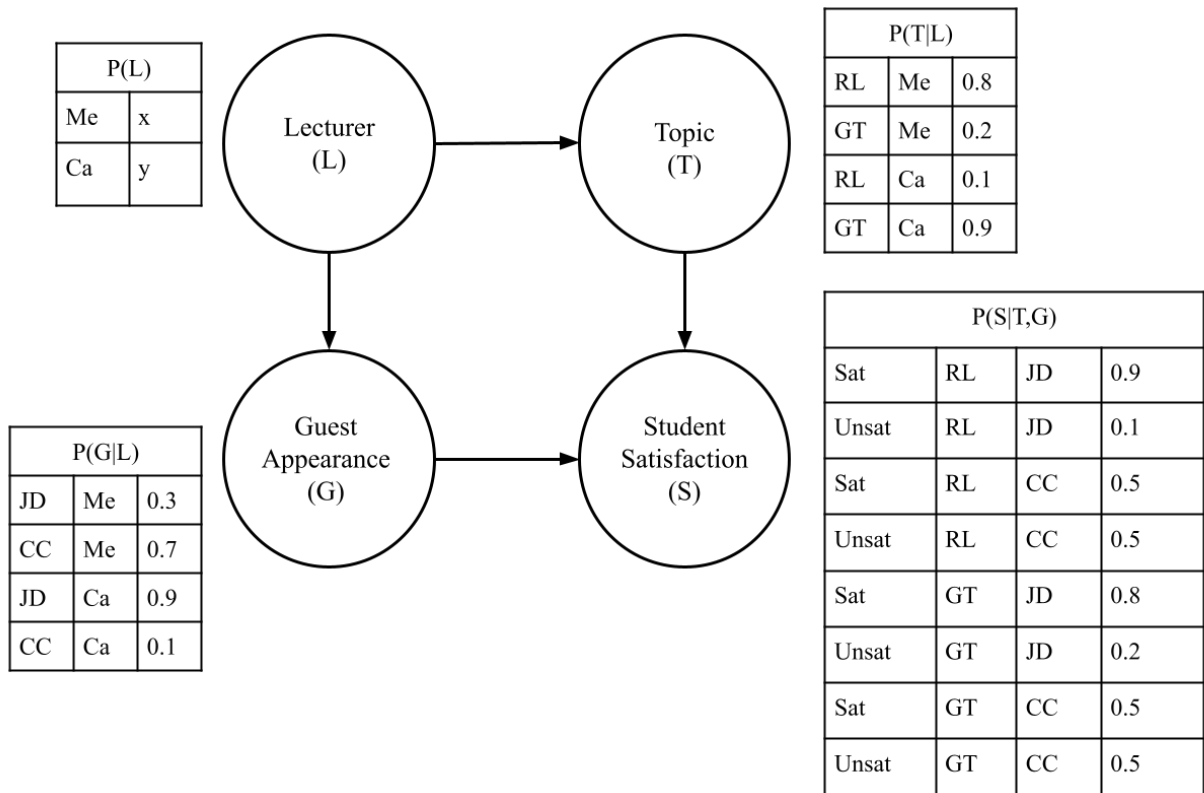
- ☐  $VPI(E'|E=e) - VPI(F'|F=f) \geq 0$   
☐  $VPI(E'|E=e) * VPI(F'|F=f) \geq 0$   
☐  $VPI(E_h, E_i, E_j|E=e) = VPI(E_h|E=e) + VPI(E_i|E=e, E_h) + VPI(E_j|E=e, E_h, E_i)$   
☐  $VPI(E_h, E_i, E_j|E=e) = VPI(E_h|E=e) + VPI(E_i|E=e) + VPI(E_j|E=e)$   
☐  $VPI(E_j, E_k|E=e) = VPI(E_j|E=e) + VPI(E_k|E=e)$

(e) [3 pts] Seth tries to generate samples using a modified version of prior sampling. Half of the time, he follows the normal prior sampling procedure. The other half of the time, he randomly generates a sample where he gives all of the variable assignments an equal chance. What is the probability of a certain sample following this procedure? Let the variables be  $x_1, x_2, \dots, x_n$ , each of which can take on  $k$  possible values.

$$P(\text{sample}) = \frac{1}{2}(A)(B) + \frac{1}{2}(C)$$

- (A): ☐  $\prod_{i=1}^n$  ☐  $\sum_{i=1}^n$  ☐  $\max_{i=1 \dots n}$   
 • (B): ☐  $P(x_i)$  ☐  $P(x_i | \text{parents}(x_i))$  ☐  $P(x_i | \text{children}(x_i))$   
 • (C): ☐  $(\frac{1}{k})^n$  ☐  $(\frac{1}{n})^k$  ☐  $\frac{1}{n}$

(f) Although both great lecturers, Carl (Ca) and Mesut (Me) want to formalize this by seeing student satisfaction based off each of their lectures. They both lecture about two different topics, Reinforcement Learning (RL) or Game Trees (GT). They also both have different guest appearances in lecture, John Denaro (JD), who Carl is closer to, and Carol Christ (CC), who Mesut is closer to.

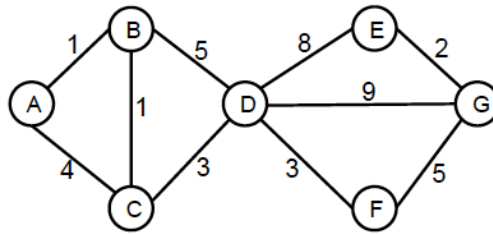


- (i) [2 pts] What's the probability that the students are satisfied with a lecture where Carol Christ makes an appearance?

$$P(Sat \mid CC) = \boxed{\phantom{0.5}}$$

- (ii) [2 pts] Given that the students were satisfied with a lecture, what's the probability that the lecture was given by Carl? You should use  $x = 0.5$  and  $y = 0.5$  for this question if needed.

$$P(Ca \mid Sat) = \boxed{\phantom{0.5}}$$



Node	$h_1$	$h_2$
A	9.5	10
B	9	12
C	8	10
D	7	8
E	1.5	1
F	4	4.5
G	0	0

- (g) Consider the state space graph shown above. A is the start state and G is the goal state. The costs for each edge are shown on the graph. Each edge can be traversed in both directions. Note that the heuristic  $h_1$  is consistent but the heuristic  $h_2$  is not consistent.

For each of the following graph search strategies (do not answer for tree search), mark which, if any, of the listed paths it could return. Note that for some search strategies the specific path returned might depend on tie-breaking behavior. In any such cases, make sure to mark all paths that could be returned under some tie-breaking scheme.

- (i) [1 pt] DFS

- ☐ A-B-D-G  
☐ A-C-D-G  
☐ A-B-C-D-F-G  
☐ None of the above

- (ii) [1 pt] BFS

- ☐ A-B-D-G  
☐ A-C-D-G  
☐ A-B-C-D-F-G  
☐ None of the above

- (iii) [1 pt] A\* Search with  $h_2$

- ☐ A-B-D-G  
☐ A-C-D-G  
☐ A-B-C-D-F-G  
☐ None of the above

Suppose you are completing the new heuristic function  $h_3$  shown below. All the values are fixed except  $h_3(B)$

Node	A	B	C	D	E	F	G
$h_3$	10	?	9	7	1.5	4.5	0

- (iv) [1 pt] What values of  $h_3(B)$  make  $h_3$  admissible?

$$\alpha \leq h_3 \leq \beta$$

$$\alpha = \beta =$$

- (v) [1 pt] What values of  $h_3(B)$  make  $h_3$  consistent?

$$\gamma \leq h_3 \leq \lambda$$

$$\gamma = \lambda =$$

## Q2. [14 pts] Model-Based RL with Function Approximation

	1	2	3	4
1				
2				
3				
4				

Consider a robot navigating in the above grid world with walls around the edges as shown above. The robot's state is represented by  $(row, column)$ , with the starting position at  $(1, 1)$

The robot cannot leave the grid. The robot can move into the wall, but its state will stay the same after the action. The allowed actions are  $a \in \{\text{up, down, left, right}\}$  in all states except the sink state shaded in green. In this sink state, only the exit action is available, which takes the agent to a terminal state in which it can no longer take actions or receive rewards.

When the robot transitions into the terminal state, it receives a reward of 10. For all other states, the robot receives a living reward of  $-1$  for transitioning into the state.

The robot wants to learn a policy to maximize its reward, but unfortunately does not know the transition model of the MDP exactly. However, the robot does know that any given state  $s$ , there is an associated (unknown) probability  $s_p$  that any action it takes will be flipped (eg. up becomes down, right becomes left).

- (a) [3 pts] To estimate the MDP transition model, the robot executes some policy  $\pi$  in the grid world and collects a set of  $N$  transitions  $\{(s_i, a_i, s'_i, r_i)\}_{i=1}^N$ . The robot then decides to use this to create a dataset  $\mathcal{D} = \{(s_i, y_i)\}_{i=1}^N$  of states in which action flips occur, where  $y_i = 1$  if a flip happened at  $s_i$  and  $-1$  otherwise.

2 The robot could use the number of action flips at each state to estimate the empirical probability of an action flip at each individual state. However, the robot is also considering training a logistic regression classifier  $g : \mathcal{S} \mapsto \{0, 1\}$  which uses  $\mathcal{D}$  to learn whether an action flip will occur at a given state. Which of the following is a possible advantage of learning a classifier? Select all that apply.

- ☐ The classifier will use less memory if the state space dimension is much lower than the number of possible states
- ☐ The classifier will give us information about action flip probabilities even on states not visited in  $\mathcal{D}$ .
- ☐ The classifier will always be more accurate.
- ☐ The classifier will be able to reuse information about the probability of action flips across different states.

- (b) [2 pts] You decide to use logistic regression to estimate the probability of an action flip at a given state  $(P(y_i | s_i, w))$ . Define the logistic function  $\phi(z) = \frac{1}{1+e^{-z}}$ . Which of the following is a correct expression for  $P(y_i | s_i, w)$ ?

- ☐  $1 + \phi(y_i w \cdot s_i)$
- ☐  $1/\phi(y_i w \cdot s_i)$
- ☐  $1/\phi(1 - y_i w \cdot s_i)$
- ☐  $1/\phi(1 + y_i w \cdot s_i)$
- ☐  $\phi(y_i w \cdot s_i)$
- ☐  $\phi(1 + y_i w \cdot s_i)$

- (c) Suppose the robot now wants to update the weights for logistic regression with gradient ascent to maximize the likelihood of the  $N$  transitions in  $\mathcal{D}$ .

Directly maximizing the likelihood of transitions in  $\mathcal{D}$  gives the following update:

$$w \leftarrow w + \eta \nabla_w \prod_{i=1}^N P(y_i | s_i, w)$$

Suppose the dataset has  $N = 64$  transitions and we are computing our weights on a computer which can only store variables  $x$  in memory if  $|x| > 10^{-32}$ .

- (i) [2 pts] If  $P(y_i | s_i, w) \leq \alpha \forall y_i, s_i$ , what is the largest value of  $\alpha$  such that  $\prod_{i=1}^N P(y_i | s_i, w)$  will **not** fit in memory?
- (ii) [2 pts] If  $N = 64$  and  $\alpha = 0.1$ , will value of  $\prod_{i=1}^N P(y_i | s_i, w)$  fit in memory?

- ☐ Yes  
☐ No

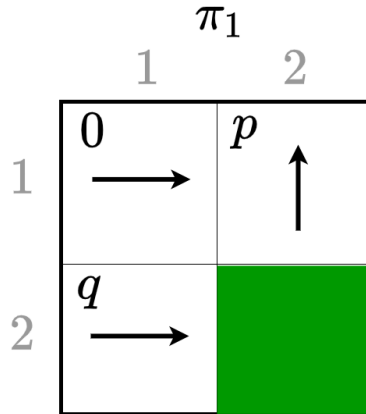
(d) Concerned about memory issues, you decide to instead maximize the log-likelihood of the transitions in  $\mathcal{D}$  as follows:

$$w \leftarrow w + \eta \nabla_w \sum_{i=1}^N \log P(y_i | s_i, w)$$

For this question, assume that log is in base 10.

- (i) [2 pts] For  $N = 64$  and  $\alpha = 0.1$ , will  $\sum_{i=1}^N \log P(y_i | s_i, w)$  fit in memory?

- ☐ Yes  
☐ No



(e) Having estimated  $w$  by running logistic regression, we can now use the resulting classifier to estimate the value of different policies in the environment. For this problem, consider a simplified  $2 \times 2$  grid world where the action flip probability is included in the top left of each grid cell. For policy  $\pi_1$  illustrated above:

- (i) [3 pts] What is  $V^{\pi_1}((1, 1))$  if  $p = q = 0.3$ ? Here  $p$  and  $q$  are the flip probabilities (probabilities of choosing the opposite action) for the respective states.

### Q3. [18 pts] Naive Bayes and Perceptron

Pacman has received a ton of spam lately. He decides to use some machine learning techniques to filter his emails.

- (a) Pacman first tries using Naive Bayes. For some reason, he chooses the words "buy", "discount", and "dollar" as features during classification. Below is the training dataset:

"buy" ( $W_1$ )	"discount" ( $W_2$ )	"dollar" ( $W_3$ )	label ( $E$ )
1	1	0	spam
1	0	1	spam
1	0	0	spam
0	1	0	spam
0	0	0	ham
0	1	1	ham

Under the assumptions of Naive Bayes, work out the following probabilities.

(i) [2 pts]  $\mathbb{P}(E = \text{ham}) =$  \_\_\_\_\_  $\mathbb{P}(E = \text{spam}) =$  \_\_\_\_\_

(ii) [4 pts]  $\mathbb{P}(W_1 = 1|E = \text{spam}) =$  \_\_\_\_\_  $\mathbb{P}(W_1 = 0|E = \text{ham}) =$  \_\_\_\_\_  
 $\mathbb{P}(W_3 = 1|E = \text{spam}, W_2 = 0) =$  \_\_\_\_\_  $\mathbb{P}(W_2 = 1, W_3 = 0|E = \text{ham}) =$  \_\_\_\_\_

- (iii) [2 pts] After filling out the probability table, Pacman found that one probability gets value zero, so he decided to use Laplace smoothing with  $k = 1$ . However, his roommate said that it's better to use  $k = 2$ . What  $k$  value should Pacman choose?

- ☐ He should pick the value that works best in training data  
☐ He should pick the value that works best in validation data  
☐ He should pick the value that works best in testing data  
☐ He should pick the average number of samples per class, which is 3 in this case.

- (iv) [3 pts] Right after implementing Laplace smoothing with  $k = 1$ , Pacman receives an email which includes all three feature words. How would the model classify this email?

$\mathbb{P}(E = \text{spam}, W_1 = 1, W_2 = 1, W_3 = 1) =$  \_\_\_\_\_

$\mathbb{P}(E = \text{ham}, W_1 = 1, W_2 = 1, W_3 = 1) =$  \_\_\_\_\_

- ☐ ham    ☐ spam

- (b) Pacman is unhappy with the performance of Naive Bayes, so he decided to switch to using a Linear Perceptron with bias, with new features: number of "\$" sign, and number of digits. Below is the training dataset:

number of "\$" sign ( $n_1$ )	number of digits ( $n_2$ )	label ( $E$ )
7	3	spam (1)
0	2	ham (-1)
5	5	spam (1)
1	1	ham (-1)

- (i) [2 pts] Pacman is confident that the Perceptron will correctly classify all incoming emails after it converges. Is that true?

- ☐ Yes, because the training data is linearly separable  
☐ Yes, but not the reason above  
☐ No, because no decision boundary that goes through the origin can separate the data  
☐ No, but not the reason above

- (ii) [3 pts] Starting from initial weight  $w = \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}$  (the last entry being the bias weight), determine the weight for the

first few iterations. **Leave all answers in form of "[a,b,c]"**

After seeing the first data: \_\_\_\_\_



After seeing the second data: \_\_\_\_\_

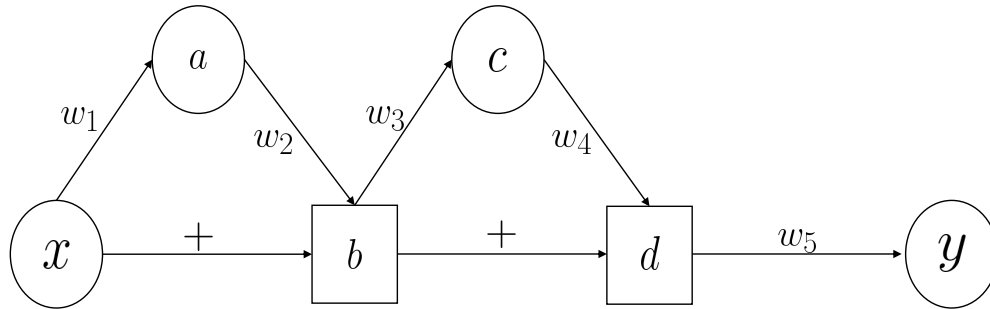
After seeing the third data: \_\_\_\_\_

- (iii) [2 pts] To decouple from the above, suppose Pacman now uses a Linear Perceptron **without bias**, and the current weight is  $\begin{bmatrix} \frac{1}{2} \\ 1 \\ -\frac{1}{2} \end{bmatrix}$ , which is in the fourth quadrant. The next training example has 1 occurrence of "\$" sign, and 1 occurrence of digits. Which quadrant could the weight be in after training using this sample?

- ☐ The first quadrant
- ☐ The second quadrant
- ☐ The third quadrant
- ☐ The fourth quadrant
- ☐ The x or y axis

## Q4. [18 pts] Backpropagation with Activation Checkpointing

Below is a neural network with residual connections (square nodes) whose weights are  $w_1, w_2, w_3, w_4, w_5$ . The neural network takes  $x$  as input and outputs  $y$ .



The outputs at each node are computed as the following:

$$\begin{aligned} o_a &= \text{ReLU}(z_a) \text{ where } z_a = x \cdot w_1 \\ o_b &= \text{ReLU}(z_b) + x \text{ where } z_b = o_a \cdot w_2 \\ o_c &= \text{LeakyReLU}(z_c) \text{ where } z_c = o_b \cdot w_3 \\ o_d &= \text{LeakyReLU}(z_d) + o_b \text{ where } z_d = o_c \cdot w_4 \\ y &= o_d \cdot w_5 \end{aligned}$$

$$\text{Let } \text{ReLU}(z) = \max(z, 0) \text{ while } \text{LeakyReLU}(z) = \begin{cases} z, & \text{if } z > 0 \\ \gamma \cdot z, & \text{otherwise} \end{cases}$$

Suppose the network has input  $x = 2$  and  $\gamma = 0.1$ .

The weight values are  $w_1 = 1, w_2 = 2, w_3 = 1, w_4 = -5, w_5 = 3$

(a) [2 pts] Perform forward propagation on the neural network.

$$\begin{aligned} o_a &= \underline{\hspace{2cm}} \\ o_b &= \underline{\hspace{2cm}} \\ o_c &= \underline{\hspace{2cm}} \\ o_d &= \underline{\hspace{2cm}} \\ y &= \underline{\hspace{2cm}} \end{aligned}$$

(b) [3 pts] Run backpropagation to calculate the following partial derivatives. Express the values of partial derivatives using only input ( $x$ ), activations ( $o_a, o_b, o_c, o_d$ ), and constants. Do not write as a single number (must be an expression using  $x$  and/or  $o_i, i \in \{a, b, c, d\}$ ).

**Input  $o_a, o_b, o_c, o_d$  as “o\_a”, “o\_b”, “o\_c”, “o\_d” respectively.**

$$\begin{aligned} \frac{\partial y}{\partial w_5} &= \underline{\hspace{2cm}} \\ \frac{\partial y}{\partial w_4} &= \underline{\hspace{2cm}} \\ \frac{\partial y}{\partial w_3} &= \underline{\hspace{2cm}} \\ \frac{\partial y}{\partial w_2} &= \underline{\hspace{2cm}} \\ \frac{\partial y}{\partial w_1} &= \underline{\hspace{2cm}} \end{aligned}$$

(c) [1 pt] Let's say storing the value of a single activation  $o_i$ ,  $i \in \{a, b, c, d\}$  costs 1 memory unit. What is the maximum number of memory units used while running the backpropagation above? \_\_\_\_\_

(d) As we try to train this neural network, we get an out-of-memory error and it turns out the main culprit is the cost of storing activations. To address this issue, we explore **activation checkpointing** where *only a subset of activations are stored/checkpointed during forward propagation*. This means that we may need to re-run parts of the forward propagation in order to re-compute activations that are missing, but needed during backpropagation.

During backpropagation, you may *use more memory units* to re-compute and store additional activations needed for computing a particular partial derivative, but they *must be released* either when they are no longer needed for computing that partial derivative or if that partial derivative is successfully computed. Checkpointed activations are never released.

As an example, suppose  $o_a$  is stored/checkpointed after forward propagation (memory:  $[o_a]$ ) and we are interested in knowing the partial derivative which can be computed using just  $o_c$ . The neural network re-computes  $o_b$  based on  $x$  and  $o_a$  (memory:  $[o_a, o_b]$ ). Then, the neural network uses  $o_b$  to re-compute  $o_c$  (memory:  $[o_a, o_b, o_c]$ ).  $o_b$  is released as it's no longer needed in computing the partial derivative of interest (memory:  $[o_a, o_c]$ ). Once the partial derivative in interest is computed using  $o_c$ ,  $o_c$  is released (memory:  $[o_a]$ ).

**Now, suppose we only checkpointed  $o_b$  during forward propagation.**

(i) [1 pt] Can you compute  $\frac{\partial y}{\partial w_5}$ ?

- ☐ Yes, and without needing any additional computation than what is required in the vanilla backpropagation (without activation checkpointing).
- ☐ Yes, but requiring more computation than what is used in the vanilla backpropagation (without activation checkpointing).
- ☐ No

(ii) [1 pt] Can you compute  $\frac{\partial y}{\partial w_4}$ ?

- ☐ Yes, and without needing any additional computation than what is required in the vanilla backpropagation (without activation checkpointing).
- ☐ Yes, but requiring more computation than what is used in the vanilla backpropagation (without activation checkpointing).
- ☐ No

(iii) [1 pt] Can you compute  $\frac{\partial y}{\partial w_3}$ ?

- ☐ Yes, and without needing any additional computation than what is required in the vanilla backpropagation (without activation checkpointing).
- ☐ Yes, but requiring more computation than what is used in the vanilla backpropagation (without activation checkpointing).
- ☐ No

(iv) [2 pts] What is the maximum number of memory units used while computing the above partial derivatives? If a partial derivative cannot be computed, assume no additional memory unit was used. \_\_\_\_\_

**Now, suppose we only checkpoint  $o_c$  during forward propagation.**

(v) [1 pt] Can you compute  $\frac{\partial y}{\partial w_5}$ ?

- ☐ Yes, and without needing any additional computation than what is required in the vanilla backpropagation (without activation checkpointing).
- ☐ Yes, but requiring more computation than what is used in the vanilla backpropagation (without activation checkpointing).
- ☐ No

(vi) [1 pt] Can you compute  $\frac{\partial y}{\partial w_4}$ ?

- ☐ Yes, and without needing any additional computation than what is required in the vanilla backpropagation (without activation checkpointing).
- ☐ Yes, but requiring more computation than what is used in the vanilla backpropagation (without activation checkpointing).
- ☐ No

(vii) [1 pt] Can you compute  $\frac{\partial y}{\partial w_3}$ ?

- ☐ Yes, and without needing any additional computation than what is required in the vanilla backpropagation (without activation checkpointing).
- ☐ Yes, but requiring more computation than what is used in the vanilla backpropagation (without activation checkpointing).
- ☐ No

(viii) [2 pts] What is the maximum number of memory units used while computing the above partial derivatives? If a partial derivative cannot be computed, assume no additional memory unit was used. \_\_\_\_\_

(ix) [2 pts] Which one is a better checkpoint between  $o_b$  and  $o_c$ ?

- ☐ Both are the same.
- ☐  $o_b$  because it has a lower peak memory usage.
- ☐  $o_c$  because it has a lower peak memory usage.
- ☐  $o_b$  because it requires less additional compute.
- ☐  $o_c$  because it requires less additional compute.

## Q5. [16 pts] Ace King Queen

Your friend, Trevor, proposes a simplified game of poker with three cards - an Ace, a King, and Queen (best to worst in that order).

In the game, each player gets a card (drawn from the 3 total cards without replacement) face down and puts in a mandatory \$10 ante. The first player can either "bet" (put in another) \$20 or "check" which ends the game (causing the player with the better card to take the \$20 in the middle).

If the first player bets, the second player can either "call", matching first player's bet (and have the player with the highest card win the now \$60 pot), or "fold" (let the first player win without seeing his card).

An example round would be the first player drawing an Ace and second player drawing a King. The first player could either "bet" or "check". Say the first player "bets", then the second player could either "call" or "fold". If the second player "calls" he would lose \$30 to the first player (they would reveal cards and first player would have a higher card).

(a) Now, let's model this game as a Bayes net. ofa

(i) [2 pts] Is the first player's card independent from the second player's card?

- ☐ Yes, because they are drawn separately.
- ☐ Yes, but for another reason.
- ☐ No, because the card may affect the player's strategy.
- ☐ No, but for another reason.

(ii) [1 pt] Is the first player's card independent from the second player's card given the last (third) card in the deck?

- ☐ Yes
- ☐ No

(b) Suppose you're the second player, and the opponent has just bet, so now it's your turn to decide on an action.

Furthermore, you know that the opponent will bet a third of his Queens, two-thirds of his Kings, and all of his Aces.

(i) [1 pt] If you have a Queen, what is the expected utility of "calling" in your spot?

EU = \_\_\_\_\_

(ii) [1 pt] If you have a Queen, what is the expected utility of "folding" in your spot?

EU = \_\_\_\_\_

(iii) [1 pt] Say you have a Queen as the second player, and the first player bets, what is the optimal move?

- ☐ Fold
- ☐ Call

(iv) [1 pt] If you now have an Ace, what is the expected utility of "calling" in your spot?

EU = \_\_\_\_\_

(v) [1 pt] What is the expected utility of the game for the first player if he chooses to always check?

Value = \_\_\_\_\_

(c) Suppose you're the second player with a King, and the opponent has just bet, so now it's your turn to decide on an action.

You still know that the opponent will bet a third of his Queens, two-thirds of his Kings, and all of his Aces.

(i) [1 pt] What is probability that your opponent is holding an Ace?

P = \_\_\_\_\_

(ii) [1 pt] What is probability that your opponent is holding an King?

P = \_\_\_\_\_

(iii) [1 pt] What is probability that your opponent is holding an Queen?

P = \_\_\_\_\_

(iv) [3 pts] Now, the dealer tells you that if you pay him, he will tell you what card your opponent is holding. The dealer never lies.

What is the maximum expected utility of this state if you were to find out your opponent's card?

MEU = \_\_\_\_\_

(v) [2 pts] Suppose the dealer doesn't always tell you the truth, even if you pay him. Can we still model how much to pay him with value of perfect information?

☐ Yes

☐ No

**(d) Bonus! Only attempt if you have extra time. Worth 0 points. Not on the examtool for confusion reasons.**

(i) [0 pts] As the first player, with what probability should you bet when holding a King?

P = \_\_\_\_\_

(ii) [0 pts] As the first player, with what probability should you bet when holding a Queen?

P = \_\_\_\_\_

(iii) [0 pts] As the mandatory contribution to the pot increases, how does the optimal strategy with a Queen change as the first player?

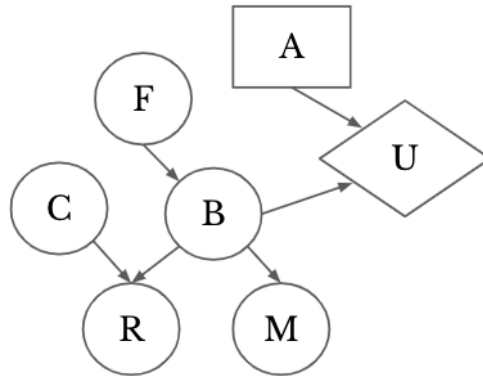
☐ Bet more

☐ Check more

☐ Optimal strategy doesn't change

## Q6. [20 pts] Pure Romance

Andy's got a problem: he has a crush on Brianna and doesn't know if he should ask her out. One day, Andy watches lecture and realizes something — he can model his current worry as a decision net! He drew up the following net with utility function  $U(A, B)$ :



- $A \in \{yes, no\}$ : Whether Andy asks Brianna out
- $B$ : How Brianna feels towards Andy
- $F$ : How Brianna's best friend thinks of Andy
- $M$ : Brianna's mood when Andy is around her
- $R$ : The time it takes for Brianna to reply to Andy's messages
- $C$ : The courseload that Brianna is taking this semester

(a) For the following relations, indicate whether it is always, sometimes, or never true.

(i) [2 pts]  $\forall PI(R|M) > 0$

- ☐ Always true
- ☐ Sometimes true
- ☐ Never true

(ii) [2 pts]  $\forall PI(C) \leq 0$

- ☐ Always true
- ☐ Sometimes true
- ☐ Never true

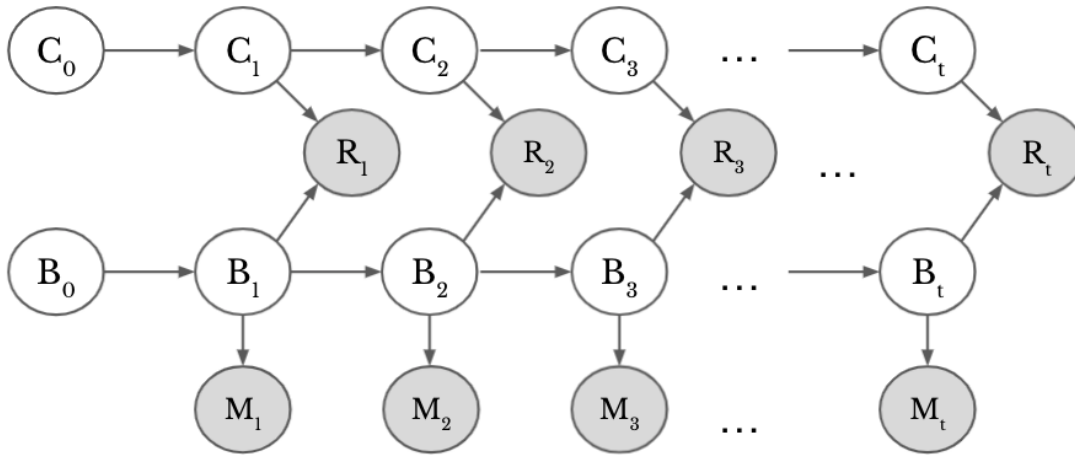
(iii) [2 pts]  $\forall PI(F, B) > \forall PI(B)$

- ☐ Always true
- ☐ Sometimes true
- ☐ Never true

(iv) [2 pts]  $\forall PI(R, M|B) = \forall PI(R|B) + \forall PI(M|B)$

- ☐ Always true
- ☐ Sometimes true
- ☐ Never true

- (b) Andy realizes he can observe  $R$  and  $M$ . He's not sure that his estimates of how Brianna feels about him ( $B$ ) and her courseload ( $C$ ) are close to reality, and so he decides to take his time to estimate these two variables using a Hidden Markov Model. He decides to disregard  $F$  in his HMM.



He plans to spend  $t$  days collecting evidence, and use his belief of  $B$  and  $C$  on day  $t$  to inform his decision net.

- (i) [2 pts] How should Andy solve for these beliefs?

- ☐ Particle filtering, because particle weights are useful in decision nets.
- ☐ Particle filtering, because we don't know how long the time horizon  $t$  will be.
- ☐ Exact inference, since it is always more accurate than particle filtering
- ☐ Exact inference if the time horizon  $t$  is short and particle filtering if the time horizon  $t$  is long
- ☐ Unable to answer with the information provided

- (ii) [2 pts] We want to develop a model for this HMM. Which of the following are equivalent to an observation model  $P(M_i, R_i | B_i, C_i)$ , where  $3 \leq i \leq t$ ?

- ☐  $P(M_i | B_i)P(R_i | B_i, C_i)$
- ☐  $P(M_i | B_i)P(M_i | C_i)P(R_i | B_i)P(R_i | C_i)$
- ☐  $P(M_i | B_i, C_i)P(R_i | B_i, C_i)$
- ☐  $P(M_i | B_i)P(M_i | C_i)P(R_i | B_i)$
- ☐  $P(M_i | R_i, B_i, C_i)P(R_i | B_i, C_i)$
- ☐ None of the above

- (iii) [2 pts] Which of the following are equivalent to  $P(B_i, C_i | B_{i-1}, C_{i-1}, R_{i-1})$ , where  $3 \leq i \leq t$ ?

- ☐  $P(B_i | B_{i-1}, R_{i-1})P(C_i | C_{i-1}, R_{i-1})$
- ☐  $P(B_i | B_{i-1}, R_{i-1})P(C_i | B_i, C_{i-1}, R_{i-1})$
- ☐  $P(B_i | B_{i-1}, C_i, R_{i-1})P(C_i | B_i, C_{i-1}, R_{i-1})$
- ☐  $P(C_i | B_i, C_{i-1})$
- ☐  $P(B_i | B_{i-1}, C_{i-1})$
- ☐ None of the above



- (c) [2 pts] An issue is that running this HMM takes time. And every day Andy puts off the decision to ask Brianna out makes him miserable. What of the following ways can he incorporate this logic into his algorithm?
- ☐ Conduct Laplace Smoothing on  $B_t(B)$  and  $B_t(C)$ , using  $t$  as the hyperparameter  $k$
  - ☐ Replace  $U(A, B)$  with  $U_1(A, B, t)$ , where  $U_1(A, B, t) = U(A, B) + t$
  - ☐ Replace  $U(A, B)$  with  $U_2(A, B, t)$ , where  $U_2(A, B, t) = \frac{U(A, B)}{t}$
  - ☐ None of the above
- (d) [2 pts] Andy's friend Joe thinks he might be able to break this HMM into two separate HMMs, one with state variable  $B$  and evidence variables  $R$  and  $M$ , and one with state variable  $C$  and evidence variable  $R$ . Joe says Andy can then calculate his beliefs of  $C$  and  $B$  separately. Is this approach valid?
- ☐ Yes, because  $C$  doesn't affect the final utility
  - ☐ Yes, but not for the reason above
  - ☐ No, because  $B$  and  $C$  are conditionally independent
  - ☐ No, but not for the reason above
- (e) [2 pts] Andy solved for his belief of his HMM at a time  $t$ . He's ready to make his decision. Does he need any other factors to decide on asking Brianna out?
- ☐ He needs the factor  $P(F)$
  - ☐ He needs the factors  $P(F)$  and  $P(C|R, M)$
  - ☐ He needs factors other than the options above
  - ☐ He doesn't need any additional factors, but he would have to do additional computation before using the belief in the decision net.
  - ☐ He doesn't need any additional factors, and he would not have to do additional computation before using the belief in the decision net

## Q7. [18 pts] Games

(a) In the following problems please choose all the answers that apply. You may pick more than one answer.

(i) [2 pts] In the context of adversarial search,  $\alpha - \beta$  pruning

- ☐ can reduce computation time by pruning portions of the game tree
- ☐ is generally faster than minimax, but loses the guarantee of optimality
- ☐ always returns the same value as minimax for the root of the tree
- ☐ always returns the same value as minimax for all nodes on the leftmost (first to be explored) edge of the tree, assuming successor game states are expanded from left to right
- ☐ always returns the same value as minimax for all nodes of the tree

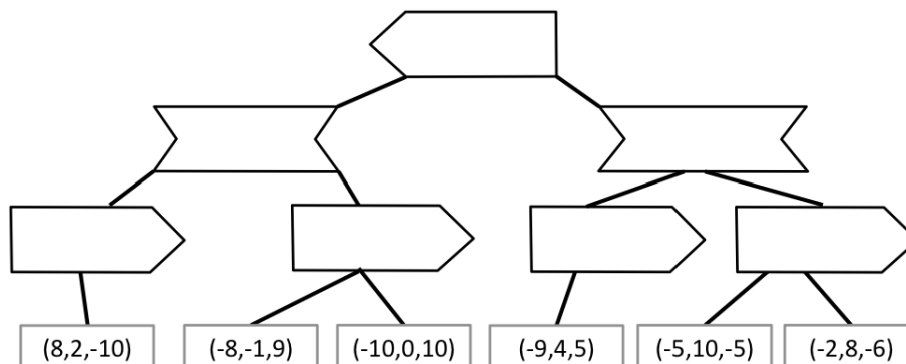
(ii) [2 pts] Consider an adversarial game in which each state  $s$  has minimax value  $v(s)$ . Assume that the maximizer plays according to the optimal minimax policy  $\pi$ , but the opponent (the minimizer) plays according to an unknown, possibly suboptimal policy  $\pi'$ . Which of the following statements are true?

- ☐ The score for the maximizer from a state  $s$  under the maximizer's control could be greater than  $v(s)$
- ☐ The score for the maximizer from a state  $s$  under the maximizer's control could be less than  $v(s)$ .
- ☐ Even if the opponent's strategy  $\pi'$  were known, the maximizer should play according to  $\pi$ .
- ☐ If  $\pi'$  is optimal and known, the outcome from any  $s$  under the maximizer's control will be  $v(s)$ .

(iii) [3 pts] Consider a very deep game tree where the root node is a maximizer, and the complete-depth minimax value of the game is known to be  $v_\infty$ . Similarly, let  $\pi_\infty$  be the minimax-optimal policy. Also consider a depth-limited version of the game tree where an evaluation function replaces any tree regions deeper than depth 10. Let the minimax value of the depth-limited game tree be  $v_{10}$  for the current root node, and let  $\pi_{10}$  be the policy which results from acting according to a depth 10 minimax search at every move. Which of the following statements are true?

- ☐  $v_\infty$  may be greater than or equal to  $v_{10}$
- ☐  $v_\infty$  may be less than or equal to  $v_{10}$
- ☐ Against a perfect opponent, the actual outcome from following  $\pi_{10}$  may be greater than  $v_\infty$ .
- ☐ Against a perfect opponent, the actual outcome from following  $\pi_{10}$  may be less than  $\pi_\infty$ . This assumes that the perfect opponent is playing with infinite depth lookahead.

(b) (i) [2 pts] Consider the 3-player game shown below. The player going first (at the top of the tree) is the Left player, the player going second is the Middle player, and the player going last is the Right player, optimizing the left, middle and right components respectively of the utility vectors shown. Fill in the values at all nodes. Note that all players maximize their own respective utilities.



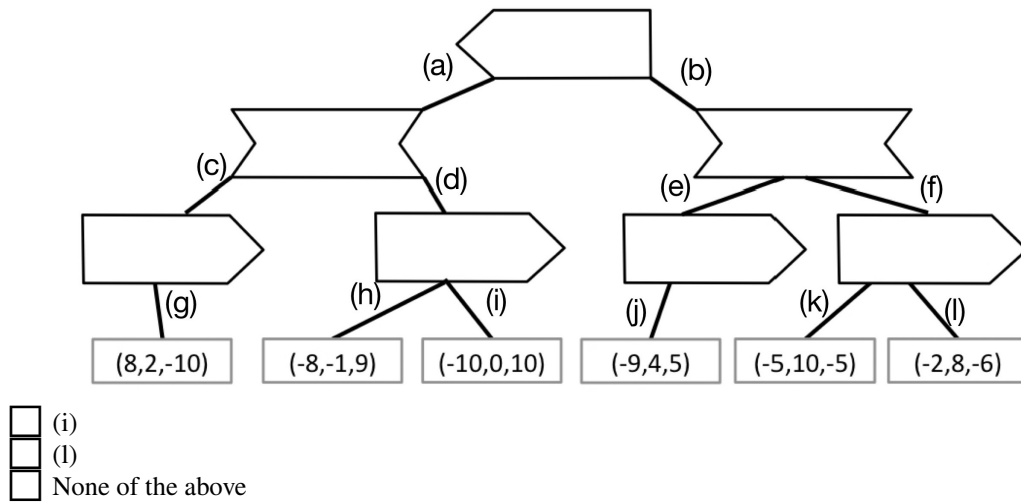
(ii) [6 pts]

**We have the knowledge that the sum of the utilities of all 3 players is always zero.**

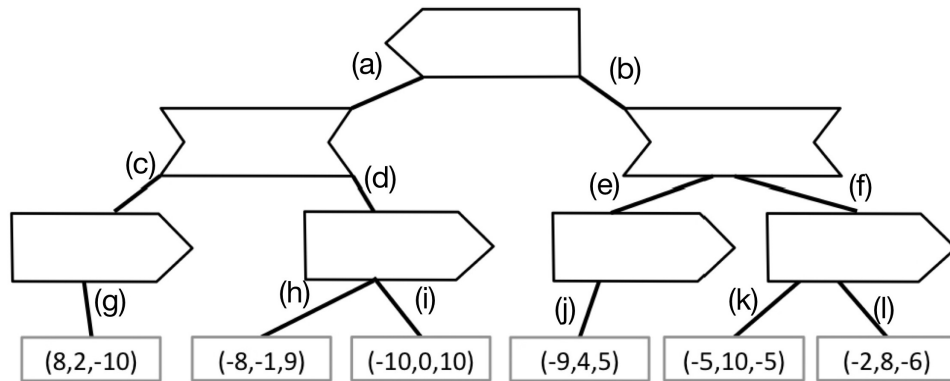
Select all edges for which observing the node will not affect the top-level decision (edges that can be pruned). To clarify, we only care about the left player's value.

If you prune a parent branch, do not mark any downstream children as pruned.

☐ (f)



(iii) [3 pts]



If we assume more about a game, additional pruning may become possible. Now, in addition to assuming that the sum of the utilities of all 3 players is still zero, we also assume that all utilities are in the interval  $[-10, 10]$ . Select all edges that would be pruned under these assumptions.

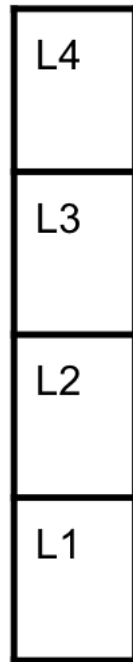
If you prune a parent branch, do not mark any downstream children as pruned.

- ☐ (f)  
☐ (i)  
☐ (l)  
☐ None of the above

## Q8. [15 pts] Hidden Markov Models and Particle Filtering

The elevator in the Tower of Terror moves up and down to the other floors (L1, L2, L3, L4). The location of the elevator at time  $t$  is  $X_t$ . At the beginning of each timestep,

- (i) the elevator goes upwards with a probability of 0.4. It may go to any floor above its current position with equal probability.
- (ii) the elevator goes downwards with a probability of 0.4. It may go to any floor below its current position with equal probability.
- (iii) the elevator stays where it is with a probability of 0.2. If the elevator is on floor L4, it goes down with probability 0.8 and stays in position with probability 0.2. Similarly, if the elevator is on floor L1, it goes up with probability 0.8 and stays in position with probability 0.2.



$X_0$	$P(X_0)$
L4	0.2
L3	0.2
L2	0.3
L1	0.3

- (a) [3 pts] Fill in the table below with the distribution of the elevator's location at time  $t = 1$ .

$X_1$	$P(X_1)$
L4	0.26
L3	
L2	
L1	

(b) Calculate the stationary distribution for the tower states by filling the unknown values in the matrix below

$$\begin{bmatrix} 0.2 & (i) & 0.2 & 0.266 \\ (ii) & (iii) & 0.2 & 0.266 \\ 0.266 & 0.2 & 0.2 & 0.266 \\ 0.266 & (iv) & 0.4 & 0.2 \\ 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} P(X_\infty = L1) \\ P(X_\infty = L2) \\ P(X_\infty = L3) \\ P(X_\infty = L4) \end{bmatrix} = \begin{bmatrix} P(X_\infty = L1) \\ P(X_\infty = L2) \\ P(X_\infty = L3) \\ P(X_\infty = L4) \\ 1 \end{bmatrix}$$

Fill in the missing values in the stationary system of equations for the following subparts.

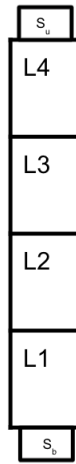
(i) [1 pt]

(ii) [1 pt]

(iii) [1 pt]

(iv) [1 pt]

To keep track of the position of the elevator, a sound sensor  $S_u$  is installed on the top of the tower and a sound sensor  $S_b$  is installed in the basement. Both sensors detect the excited sounds of the passengers, (+s), or no sound at all, -s. The distribution of sensor measurements is determined by  $d$ , the number of floors between the elevator and the respective sensor. For example, if the elevator is on floor L3, then  $d_b = 2$  because there are two floors (L2 and L1) between floor L3 and the bottom and  $d_u = 1$  because there is one floor (L4) between floor L3 and the top. The prior of the both sensors' outputs are identical and listed below.



$S_u$	$P(S_u d_u)$
+s	$0.2/(d_u+1)$
-s	$1 - 0.2/(d_u+1)$

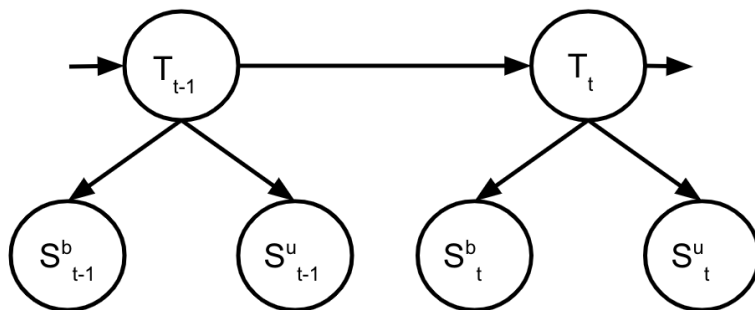
$S_b$	$P(S_b d_b)$
+s	$1 - 0.2*d_b$
-s	$0.2*d_b$

(c) [2 pts] You decide to track the elevator's position by particle filtering with 3 particles. At the end of time  $t = 1$ , the particles are at positions  $p_1 = L1$ ,  $p_2 = L2$  and  $p_3 = L3$ . Without incorporating any sensory information, what is the probability that the particles will be resampled as  $p_1 = L3$ ,  $p_2 = L2$ , and  $p_3 = L4$ , at the end of time  $t=1$ ?

(d) [3 pts] To decouple this from the previous question, assume the particles after time elapsing are  $p_1 = L3$ ,  $p_2 = L2$ ,  $p_3 = L1$ , and the sensors observe  $S_u = +s$  and  $S_b = -s$ . What are the particle weights given these observations?

Particle	Weight
$X_1 = L3$	
$X_2 = L2$	
$X_3 = L1$	

- (e) [3 pts] Note: the u and b subscripts from before will be written here as superscripts. Part of the expression for the forward algorithm update for Hidden Markov Models is given below.  $s_{0:t}^u$  are all the measurements from the roof sensor  $s_0^u, s_1^u, s_2^u, \dots, s_t^u$ .  $s_{0:t}^b$  are all the measurements from the roof sensor  $s_0^b, s_1^b, s_{12}^b, \dots, s_t^b$ . Choose all the correct options for the blank given



$$\begin{aligned}
 P(x_t | s_{0:t}^u, s_{0:t}^b) &\propto P(x_t, s_{0:t}^u, s_{0:t}^b) \\
 &= \sum_{x_{t-1}} P(x_{t-1}, x_t, s_{0:t}^u, s_{0:t}^b) \\
 &= \sum_{x_{t-1}} P(x_{t-1}, x_t, s_{0:t-1}^u, s_{0:t-1}^b, s_t^u, s_t^b) \\
 &\propto \sum_{x_{t-1}} \text{_____}
 \end{aligned}$$

- ☐  $P(s_t^u, s_t^b | x_{t-1}, x_t, s_{0:t-1}^u, s_{0:t-1}^b)$   
☐  $P(s_t^u | x_t) P(s_t^b | x_t)$   
☐  $P(s_t^u | x_{t-1}) P(s_t^b | x_{t-1})$   
☐  $P(s_t^u | s_{t-1}^u) P(s_t^b | s_{t-1}^b)$   
☐ None of the above