- You have approximately 2 hours and 50 minutes.
- The exam is closed book, closed calculator, and closed notes except your three crib sheets.
- Mark your answers ON THE EXAM ITSELF. If you are not sure of your answer you may wish to provide a brief explanation or show your work.
- For multiple choice questions,
$-\square$ means mark all options that apply
$-\bigcirc$ means mark a single choice
- There are multiple versions of the exam. For fairness, this does not impact the questions asked, only the ordering of options within a given question.

| First name |  |
| :--- | :--- |
| Last name |  |
| SID |  |
| edX username |  |


| First and last name of student to your left |  |
| :--- | :--- |
| First and last name of student to your right |  |

For staff use only:

| Q1. | Agent Testing Today! | $/ 1$ |
| :--- | :--- | :--- |
| Q2. | Potpourri | $/ 21$ |
| Q3. | Bayes Nets and Sampling | $/ 6$ |
| Q4. | Deep Learning | $/ 15$ |
| Q5. | MDPs: Reward Shaping | $/ 11$ |
| Q6. | Zero Sum MDP's | $/ 6$ |
| Q7. | Planning ahead with HMMs | $/ 11$ |
| Q8. | Naïve Bayes | $/ 7$ |
| Q9. | Beyond Ordinary Pruning | $/ 12$ |
| Q10. | Iterative Deepening Search | $/ 10$ |
|  | Total | $/ 100$ |

## Q1. [1 pt] Agent Testing Today!

It's testing time! Circle your favorite robot below. We hope you have fun with the rest of the exam!


## Q2. [21 pts] Potpourri

(a) (i) [1 pt] Suppose we have a multiclass perceptron with three classes $A, B, C$ and with weights initially set to $w_{A}=[1,2], w_{B}=[2,0], w_{C}=[2,-1]$. Write out the vectors $w_{A}, w_{B}, w_{C}$ of the perceptron after training on the following two dimensional training example once.

| $x_{0}$ | $x_{1}$ | label |
| :---: | :---: | :---: |
| 1 | 1 | A |

$w_{A}=\left[\ldots, \quad w_{B}=\left[\ldots, \quad w_{C}=[\ldots, \quad]\right.\right.$
(ii) [1 pt] Suppose we have a different multiclass perceptron with three classes $A, B, C$ and with weights initially set to $w_{A}=[2,4], w_{B}=[-1,0], w_{C}=[2,-2]$. Write out the vectors $w_{A}, w_{B}, w_{C}$ of the perceptron after training on the following two dimensional training example once.

| $x_{0}$ | $x_{1}$ | label |
| :---: | :---: | :---: |
| -2 | 1 | C |

$$
\begin{aligned}
& w_{A}=[\ldots] \quad w_{B}=[\ldots] \\
& w_{C}=[\ldots, \quad]
\end{aligned}
$$

(iii) [3 pts] Suppose we have a different multiclass perceptron with three classes $A, B, C$ and with weights initially set to $w_{A}=[1,0], w_{B}=[1,1], w_{C}=[3,0]$. After training on the following set of training data an infinite number of times, select which of the following options must be True given no additional information. Convergence indicates that the values do not change even within a pass through the data set.

| training example $i$ | $x_{0}$ | $x_{1}$ | label |
| :---: | :---: | :---: | :---: |
| 0 | 1 | 1 | A |
| 1 | -1 | 1 | B |
| 2 | 1 | -1 | C |
| 3 | -1 | -1 | A |All of the weight vectors $w_{A}, w_{B}, w_{C}$ converge.

$\bigcirc$
Only two of the weight vectors $w_{A}, w_{B}, w_{C}$ converge.
$\bigcirc$
Only one of the weight vectors $w_{A}, w_{B}, w_{C}$ converge.
$\bigcirc$
None of the weight vectors $w_{A}, w_{B}, w_{C}$ converge.None of the above.
(b) You are given a constraint graph for a Constraint Satisfaction Problem as follows. The domains of all variables are indicated in the table, and the binary constraints are as follows:

- $\mathrm{A}>\mathrm{B}$
- $\mathrm{A} \neq \mathrm{C}$
- $\mathrm{C}>\mathrm{B}$
- $\mathrm{D}<\mathrm{B}$


| $A$ | 0 | 1 | 2 | 3 |
| :--- | :--- | :--- | :--- | :--- |
| $B$ | 0 | 1 | 2 | 3 |
| $C$ | 0 | 1 | 2 | 3 |
| $D$ | 0 | 1 | 2 | 3 |

(i) [3 pts] Enforce arc consistency on this graph and indicate what the domains of all the variables are after arc consistency is enforced, in the table below by crossing out eliminated values from the domains.

| $A$ | 0 | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: | :---: |
| $B$ | 0 | 1 | 2 | 3 |
| $C$ | 0 | 1 | 2 | 3 |
| $D$ | 0 | 1 | 2 | 3 |

(ii) [2 pts] Now suppose you are given a different CSP with variables still being $A, B, C, D$, but you are not given the constraints. The domains of variables remaining after enforcing arc consistency for this CSP are given to you below. Select all of the following options which can be inferred given just this information.

| $A$ |  |  | 2 | 3 |
| :---: | :--- | :--- | :--- | :--- |
| $B$ |  |  | 2 | 3 |
| $C$ | 0 | 1 | 2 |  |
| $D$ |  |  | 2 | 3 |The CSP may have no solution.The CSP must have a solution.The CSP must have exactly one solution.The CSP may have more than one solution.The CSP must have more than one solution.

None of the above.
(c) [3 pts] Your assistant gives you the probability distributions for 4 mysterious binary variables: W, X, Y, and Z. Circle the Bayes net(s) amongst those given, that can represent a distribution that is consistent with the tables below using the fewest edges. If there is more than one such minimal net, circle all of them.

| X | $P(X)$ |
| :---: | :---: |
| 0 | 0.75 |
| 1 | 0.25 | | X | W | $P(W \mid X)$ |
| :---: | :---: | :---: |
| 0 | 0 | 0.4 |
| 0 | 1 | 0.6 |
| 1 | 0 | 0.4 |
| 1 | 1 | 0.6 |


| X | Y | $P(Y \mid X)$ |
| :---: | :---: | :---: |
| 0 | 0 | 0.3 |
| 0 | 1 | 0.7 |
| 1 | 0 | 0.1 |
| 1 | 1 | 0.9 |


| Z | W | $P(W \mid Z)$ |
| :---: | :---: | :---: |
| 0 | 0 | 0.2 |
| 0 | 1 | 0.8 |
| 1 | 0 | 0.8 |
| 1 | 1 | 0.2 |




(d) Triangle is a rational agent in the world below, where it gains or loses Utility from moving and picking up money. Triangle can move deterministically Up, Down, Left or Right or Stay still. Black squares indicate that the Triangle cannot traverse them. The squares marked with $L(\$ 100, \$ 0)$ indicate lotteries of $[0.5, \$ 100 ; 0.5$ $\$ 0]$. Taking a step onto a blank square gives Triangle no utility, but stepping onto a lottery square gives it the utility of the lottery, and the lottery disappears.
Additionally, taking a step in any direction has a probability $p$ of giving Triangle pain in addition to whatever money it might earn upon landing on a spot. If Triangle chooses to stay still, it will not feel pain. The utilities are not discounted in this problem so $\gamma=1$.


In both of the problems below, Triangle's starting position is as shown in the figure above.
(i) $[1 \mathrm{pt}]$ For this part, Triangle's utility is as follows(where $k>0$ ):

$$
U(\text { pain })=-k ; \quad U(\$ m)=m
$$

What is the expected utility of going to the closest lottery and staying in that spot forever? Express your answer in terms of numerical constants, $p, k$.
(ii) [2 pts] Now, triangle's utility function is as follows:

$$
U(\text { pain })=-k ; \quad U(\$ m)=\sqrt{m}
$$

For what range of $k$ (where $k>0$ ) will triangle always go to both lotteries. Express your answer in terms of numerical constants and $p$. If no such range exists, write None in the blank below.
(e) $[5 \mathrm{pts}]$ For each of the branches in the game tree below, put an ' X ' on the branches if there exists an assignment of values to leaf nodes, for which that branch could be pruned. The max nodes are upward pointing triangles $(\Delta)$, the min nodes are downward pointing triangles $(\nabla)$, and the chance nodes are circles $(\bigcirc)$. Assume that the children of a node are visited in left-to-right order.
Explicitly write down "Not possible" below if no branches can be pruned, in which case any ' X ' marks above will be ignored. Any ' X ' on the nodes and leaves will be ignored.


## Q3. [6 pts] Bayes Nets and Sampling

You are given a bayes net with the following probability tables:


| E | D | F | $P(F \mid E, D)$ |
| :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 0.6 |
| 0 | 0 | 1 | 0.4 |
| 0 | 1 | 0 | 0.7 |
| 0 | 1 | 1 | 0.3 |
| 1 | 0 | 0 | 0.2 |
| 1 | 0 | 1 | 0.8 |
| 1 | 1 | 0 | 0.7 |
| 1 | 1 | 1 | 0.3 |


|  |  | A | B | $P(B \mid A)$ | A | C | $P(C \mid A)$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| A | $P(A)$ | 0 | 0 | 0.1 | 0 | 0 | 0.3 | E | $P(E)$ |
| 0 | 0.75 | 0 | 1 | 0.9 | 0 | 1 | 0.7 | 0 | 0.1 |
| 1 | 0.25 | 1 | 0 | 0.5 | 1 | 0 | 0.7 | 1 | 0.9 |
|  |  | 1 | 1 | 0.5 | 1 | 1 | 0.3 |  |  |


| E | C | D | $P(D \mid E, C)$ |
| :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 0.5 |
| 0 | 0 | 1 | 0.5 |
| 0 | 1 | 0 | 0.2 |
| 0 | 1 | 1 | 0.8 |
| 1 | 0 | 0 | 0.5 |
| 1 | 0 | 1 | 0.5 |
| 1 | 1 | 0 | 0.2 |
| 1 | 1 | 1 | 0.8 |

You want to know $P(C=0 \mid B=1, D=0)$ and decide to use sampling to approximate it.
(a) $[2$ pts $]$ With prior sampling, what would be the likelihood of obtaining the sample $[\mathrm{A}=1, \mathrm{~B}=0, \mathrm{C}=0, \mathrm{D}=0$, $\mathrm{E}=1, \mathrm{~F}=0]$ ?
$\bigcirc$
$0.25^{*} 0.1^{*} 0.3^{*} 0.9 * 0.8^{*} 0.7$
$0.25 * 0.5 * 0.7 * 0.5 * 0.9 * 0.2$
$\bigcirc$
$0.75 * 0.1 * 0.3^{*} 0.9^{*} 0.5^{*} 0.8$
$\bigcirc 0.25 * 0.5 * 0.3 * 0.2 * 0.9 * 0.2$
$\bigcirc .25^{*} 0.9 * 0.7^{*} 0.1^{*} 0.5 * 0.6$
$0.75^{*} 0.1^{*} 0.3^{*} 0.9^{*} 0.5^{*} 0.2+0.25^{*} 0.5^{*} 0.7^{*} 0.5^{*} 0.9^{*} 0.2$Other $\qquad$
(b) $[2 \mathrm{pts}]$ Assume you obtained the sample $[\mathrm{A}=1, \mathrm{~B}=1, \mathrm{C}=0, \mathrm{D}=0, \mathrm{E}=1, \mathrm{~F}=1]$ through likelihood weighting. What is its weight?$0.25^{*} 0.5^{*} 0.7^{*} 0.5^{*} 0.9^{*} 0.8$
$\bigcirc$
$0.25^{*} 0.7^{*} 0.9{ }^{*} 0.8+0.75 * 0.3^{*} 0.9 * 0.8$
$0.5 * 0.5$$0.25^{*} 0.5^{*} 0.7^{*} 0.5^{*} 0.8$
$\bigcirc 0.9 * 0.5+0.1 * 0.5$Other $\qquad$
(c) $[2 \mathrm{pts}]$ You decide to use Gibb's sampling instead. Starting with the initialization $[\mathrm{A}=1, \mathrm{~B}=1, \mathrm{C}=0, \mathrm{D}=0$, $\mathrm{E}=0, \mathrm{~F}=0]$, suppose you resample F first, what is the probability that the next sample drawn is $[\mathrm{A}=1, \mathrm{~B}=1$, $\mathrm{C}=0, \mathrm{D}=0, \mathrm{E}=0, \mathrm{~F}=1]$ ?
$\bigcirc 0.4$
$\bigcirc 0.6$
$\bigcirc 0.6^{*} 0.1^{*} 0.5$
$\bigcirc 0$
$\bigcirc 0.25^{*} 0.5^{*} 0.7^{*} 0.5^{*} 0.1 * 0.3$
$\bigcirc 0.9 * 0.5+0.1 * 0.5$Other $\qquad$

## Q4. [15 pts] Deep Learning

(a) [3 pts] Perform forward propagation on the neural network below for $x=1$ by filling in the values in the table. Note that (i), ..., (vii) are outputs after performing the appropriate operation as indicated in the node.

| (i) | (ii) | (iii) | (iv) | (v) | (vi) | (vii) |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  |  |  |  |  |


(b) [6 pts] Below is a neural network with weights $a, b, c, d, e, f$. The inputs are $x_{1}$ and $x_{2}$.

The first hidden layer computes $r_{1}=\max \left(c \cdot x_{1}+e \cdot x_{2}, 0\right)$ and $r_{2}=\max \left(d \cdot x_{1}+f \cdot x_{2}, 0\right)$.
The second hidden layer computes $s_{1}=\frac{1}{1+\exp \left(-a \cdot r_{1}\right)}$ and $s_{2}=\frac{1}{1+\exp \left(-b \cdot r_{2}\right)}$.
The output layer computes $y=s_{1}+s_{2}$. Note that the weights $a, b, c, d, e, f$ are indicated along the edges of the neural network here.

Suppose the network has inputs $x_{1}=1, x_{2}=-1$.
The weight values are $a=1, b=1, c=4, d=1, e=2, f=2$.
Forward propagation then computes $r_{1}=2, r_{2}=0, s_{1}=0.9, s_{2}=0.5, y=1.4$. Note: some values are rounded.


Using the values computed from forward propagation, use backpropagation to numerically calculate the following partial derivatives. Write your answers as a single number (not an expression). You do not need a calculator. Use scratch paper if needed.
Hint: For $g(z)=\frac{1}{1+\exp (-z)}$, the derivative is $\frac{\partial g}{\partial z}=g(z)(1-g(z))$.

| $\frac{\partial y}{\partial a}$ | $\frac{\partial y}{\partial b}$ | $\frac{\partial y}{\partial c}$ | $\frac{\partial y}{\partial d}$ | $\frac{\partial y}{\partial e}$ | $\frac{\partial y}{\partial f}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |
|  |  |  |  |  |  |

(c) [6 pts] Below are two plots with horizontal axis $x_{1}$ and vertical axis $x_{2}$ containing data labelled $\times$ and $\bullet$. For each plot, we wish to find a function $f\left(x_{1}, x_{2}\right)$ such that $f\left(x_{1}, x_{2}\right) \geq 0$ for all data labelled $\times$ and $f\left(x_{1}, x_{2}\right)<0$ for all data labelled •
Below each plot is the function $f\left(x_{1}, x_{2}\right)$ for that specific plot. Complete the expressions such that all the data is labelled correctly. If not possible, mark "No valid combination".

$f\left(x_{1}, x_{2}\right)=\max (\underline{(\mathrm{i})}+\underline{(\mathrm{ii})}, \underline{(\mathrm{iii})}+\underline{(\mathrm{iv})})+\underline{(\mathrm{v})}$

$f\left(x_{1}, x_{2}\right)=(\underline{\mathbf{v i})}-\max (\underline{(\mathbf{v i i})}+(\underline{\text { viii}}),(\underline{\text { (ix })}+\underline{(\mathbf{x})})$


## Q5. [11 pts] MDPs: Reward Shaping

PacBot is in a Gridworld-like environment E. It moves deterministically Up, Down, Right, or Left except that it cannot move onto squares which are blackened. PacBot must move at every step or exit. The reward for any of these actions is always zero. Additionally, from a numbered square, PacBot can choose to exit to a terminal state and collect reward equal to the number on the square. PacBot is not required to exit on a numbered square; it can also move in any direction off that square.
(a) [3 pts] Draw an arrow in each square (including numbered squares) in the following board on the right to indicate the optimal policy PacBot will calculate with the discount factor $\gamma=0.5$ in the board on the left. (For example, if PacBot would move Down from the square in the middle on the left board, draw a down arrow in that square on the right board.) If PacBot's policy would be to exit from a particular square, draw an X instead of an arrow in that square.


PacBot now operates in a new environment $E^{\prime}$ with an additional reward function $F\left(s, a, s^{\prime}\right)$, which is added to the original reward function $R\left(s, a, s^{\prime}\right)$ for every ( $s, a, s^{\prime}$ ) triplet.
(b) [4 pts] Consider an additional reward $F_{1}$ that favors moving toward numbered squares. Let $d(s)$ be defined as the Manhattan distance from $s$ to the nearest numbered square. If $s$ is numbered, $d(s)=0$.

$$
F_{1}\left(s, a, s^{\prime}\right)= \begin{cases}0 & s^{\prime} \text { is a terminal state } \\ 10 & d\left(s^{\prime}\right)<d(s) \text { i.e. } s^{\prime} \text { is closer to a numbered square than } s \text { is, } \\ 0 & d\left(s^{\prime}\right) \geq d(s)\end{cases}
$$

Fill in the diagram on the right as in (a) to indicate the optimal policy PacBot will calculate with the discount factor $\gamma=0.5$ and the modified reward function $R_{1}^{\prime}\left(s, a, s^{\prime}\right)=R\left(s, a, s^{\prime}\right)+F_{1}\left(s, a, s^{\prime}\right)$ in the board on the left.

(c) [4 pts] Consider a different artificial reward that also favors moving toward numbered squares in a slightly different way:

$$
F_{2}\left(s, a, s^{\prime}\right)= \begin{cases}0 & s^{\prime} \text { is a terminal state } \\ 10\left(d(s)-\frac{1}{2} d\left(s^{\prime}\right)\right) & \text { otherwise }\end{cases}
$$

Fill in the diagram on the right as in (a) to indicate the optimal policy PacBot will calculate with the discount factor $\gamma=0.5$ and the modified reward function $R_{2}^{\prime}\left(s, a, s^{\prime}\right)=R\left(s, a, s^{\prime}\right)+F_{2}\left(s, a, s^{\prime}\right)$ in the board on the left.


## Q6. [6 pts] Zero Sum MDP's

Consider a Markov Decision Process where it is not just Pacman in the environment, but there is also a ghost. Pacman plays one turn, then the ghost plays one turn and they continue alternating, each of their actions transitioning the state forward using the same transition function T. At any one time step, only one of Pacman and ghost can play a turn. Let $A$ be Pacman's action set can take and $B$ be the ghost's action set. The game is infinite horizon, with discount factor $\gamma$ applied at every turn no matter which agent is taking the turn. $|A|$ is the size of A's action set and $|B|$ is the size of B 's action set. R indicates the utility received by Pacman.
(a) [2 pts] Let us first consider the situation where Pacman tries to maximize his expected utility, while the ghost tries to minimize Pacman's utility, thus playing adversarially. Both Pacman and the ghost try to play optimally and they are aware of this. Given the standard notation for an MDP, choose which of the following updates is the correct one for Q-Value Iteration under this formulation, given that $Q_{p a c}^{*}$ is the infinite horizon Q-function for Pacman.

$$
\begin{aligned}
\bigcirc Q_{p a c}^{*}(s, a) & =\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \sum_{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
Q_{p a c}^{*}(s, a) & =\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime}, a^{\prime}\right)\right] \\
Q_{p a c}^{*}(s, a) & =\sum_{s^{\prime}} R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime}, a^{\prime}\right) \\
Q_{p a c}^{*}(s, a) & =\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \sum_{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \frac{1}{|B|} \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
Q_{p a c}^{*}(s, a) & =\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \min _{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
Q_{p a c}^{*}(s, a) & =\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right]
\end{aligned}
$$

None of the above.
(b) [2 pts] For this part, let us suppose that instead of having a ghost which is adversarial, the ghost is a friendly ghost who is also trying to maximize Pacman's utility. Both Pacman and the ghost know this arrangement, and are aware of the others knowledge. Given the standard notation for an MDP, choose which of the following updates is the correct one for Q -Value Iteration under this formulation, given that $Q_{p a c}^{*}$ is the Q -function for Pacman.

$$
\begin{aligned}
& \bigcirc Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \min _{b \in B} Q_{p a c}^{*}\left(s^{\prime}, b\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \sum_{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \frac{1}{|B|} \sum_{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \frac{1}{|B|} \max _{b \in B} Q_{p a c}^{*}\left(s^{\prime}, b\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \min _{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right]
\end{aligned}
$$

None of the above.
(c) [2 pts] For this part let us suppose that instead of having a ghost which is friendly, the ghost is a confused ghost who takes random actions, with uniform probability in the environment. Given the standard notation for an MDP, choose which of the following updates is the correct one for Q-Value Iteration under this formulation, given that $Q_{p a c}$ is the Q-function for Pacman.

$$
\begin{aligned}
& \bigcirc Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \frac{1}{|B|} \max _{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\sum_{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{b \in B} Q_{p a c}^{*}\left(s^{\prime}, b\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \frac{1}{|B|} \sum_{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \frac{1}{|B|} \min _{b \in B} \sum_{s^{\prime \prime}}\left(T\left(s^{\prime}, b, s^{\prime \prime}\right)\left[R\left(s^{\prime}, b, s^{\prime \prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime \prime}, a^{\prime}\right)\right]\right)\right] \\
& Q_{p a c}^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime} \in A} Q_{p a c}^{*}\left(s^{\prime}, a^{\prime}\right)\right] \\
& \text { None of the above. }
\end{aligned}
$$

## Q7. [11 pts] Planning ahead with HMMs

Pacman is tired of using HMMs to estimate the location of ghosts. He wants to use HMMs to plan what actions to take in order to maximize his utility. Pacman uses the HMM (drawn to the right) of length $T$ to model the planning problem. In the HMM, $X_{1: T}$ is the sequence of hidden states of Pacman's world, $A_{1: T}$ are actions Pacman can take, and $U_{t}$ is the utility Pacman receives at the particular hidden state $X_{t}$. Notice that there are no evidence variables, and utilities are not discounted.

(a) The belief at time $t$ is defined as $B_{t}\left(X_{t}\right)=p\left(X_{t} \mid a_{1: t}\right)$. The forward algorithm update has the following form:

$$
B_{t}\left(X_{t}\right)=\quad(\mathbf{i}) \quad \text { (ii) } \quad B_{t-1}\left(x_{t-1}\right)
$$

Complete the expression by choosing the option that fills in each blank.
(i) $[1 \mathrm{pt}]$
$\sum_{x_{t-1}}$
$\max _{x_{t-1}}$
$\bigcirc \quad \sum_{x_{t}}$ 1
(ii) $[1 \mathrm{pt}]$
$p\left(X_{t} \mid x_{t-1}\right) p\left(X_{t} \mid a_{t}\right)$
$\bigcirc p\left(X_{t} \mid x_{t-1}\right)$
$\bigcirc p\left(X_{t}\right)$
$\bigcirc p\left(X_{t} \mid x_{t-1}, a_{t}\right)$
1

None of the above combinations is correct
(b) Pacman would like to take actions $A_{1: T}$ that maximizes the expected sum of utilities, which has the following form:

$$
\mathrm{MEU}_{1: T}=(\mathrm{i}) \quad \text { (ii) (iii) (iv) (v) }
$$

Complete the expression by choosing the option that fills in each blank.

| (i) $[1 \mathrm{pt}]$ | $\bigcirc$ | $\max _{a_{T}}$ | $\bigcirc$ | $\max _{a_{1: T}}$ | $\bigcirc$ | $\sum_{a_{1: T}}$ | $\bigcirc$ | $\sum_{a_{T}}$ |
| ---: | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| (ii) $[1 \mathrm{pt}]$ | $\bigcirc$ | $\prod_{t=1}^{T}$ | $\bigcirc$ | $\max _{t}$ | $\bigcirc$ | $\min _{t}$ | $\bigcirc$ | 1 |
| (iii) $[1 \mathrm{pt}]$ | $\bigcirc$ | $\sum_{x_{t}}$ | $\bigcirc$ | $\sum_{x_{t}, a_{t}}^{T}$ | $\bigcirc$ | $\sum_{a_{t}}$ | $\bigcirc$ | $\sum_{x_{T}}$ |
| (iv) $[1 \mathrm{pt}]$ | $\bigcirc$ | $p\left(x_{t}\right)$ | $\bigcirc$ | $p\left(x_{t} \mid x_{t-1}, a_{t}\right)$ | $\bigcirc$ | $B_{T}\left(x_{T}\right)$ | $\bigcirc$ | 1 |
| (v) $[1 \mathrm{pt}]$ | $\bigcirc$ | $\frac{1}{U_{t}}$ | $\bigcirc$ | $U_{T}$ | $\bigcirc$ | $U_{t}\left(x_{t}\right)$ | $\bigcirc$ | 1 |
| (vin |  |  |  |  |  |  |  |  |

None of the above combinations is correct
(c) [2 pts] A greedy ghost now offers to tell Pacman the values of some of the hidden states. Pacman needs your help to figure out if the ghost's information is useful. Assume that the transition function $p\left(x_{t} \mid x_{t-1}, a_{t}\right)$ is not deterministic. With respect to the utility $U_{t}$, mark all that can be True:$\operatorname{VPI}\left(X_{t-1} \mid X_{t-2}\right)>0$
$\operatorname{VPI}\left(X_{t-2} \mid X_{t-1}\right)>0$ $\square$ $\operatorname{VPI}\left(X_{t-1} \mid X_{t-2}\right)=0$$\operatorname{VPI}\left(X_{t-2} \mid X_{t-1}\right)=0$

None of the above
(d) [2 pts] Pacman notices that calculating the beliefs under this model is very slow using exact inference. He therefore decides to try out various particle filter methods to speed up inference. Order the following methods by how accurate their estimate of $B_{T}\left(X_{T}\right)$ is? If different methods give an equivalently accurate estimate, mark them as the same number.


## Q8. [7 pts] Naïve Bayes

You are given a naïve bayes model, shown below, with label Y and features $X_{1}$ and $X_{2}$. The conditional probabilities for the model are parametrized by $p_{1}, p_{2}$ and $q$.


| $X_{1}$ | $Y$ | $P\left(X_{1} \mid Y\right)$ |
| ---: | ---: | ---: |
| 0 | 0 | $p_{1}$ |
| 1 | 0 | $1-p_{1}$ |
| 0 | 1 | $1-p_{1}$ |
| 1 | 1 | $p_{1}$ |


| $X_{2}$ | $Y$ | $P\left(X_{2} \mid Y\right)$ |
| ---: | ---: | ---: |
| 0 | 0 | $p_{2}$ |
| 1 | 0 | $1-p_{2}$ |
| 0 | 1 | $1-p_{2}$ |
| 1 | 1 | $p_{2}$ |


| $Y$ | $P(Y)$ |
| ---: | ---: |
| 0 | $1-q$ |
| 1 | $q$ |

Note that some of the parameters are shared (e.g. $\left.P\left(X_{1}=0 \mid Y=0\right)=P\left(X_{1}=1 \mid Y=1\right)=p_{1}\right)$.
(a) [2 pts] Given a new data point with $X_{1}=1$ and $X_{2}=1$, what is the probability that this point has label $Y=1$ ? Express your answer in terms of the parameters $p_{1}, p_{2}$ and $q$ (you might not need all of them).

$$
P\left(Y=1 \mid X_{1}=1, X_{2}=1\right)=
$$

$\qquad$

The model is trained with the following data:

| sample number | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | :--- |
| $X_{1}$ | 0 | 0 | 1 | 0 | 1 | 0 | 1 | 0 | 1 | 1 |
| $X_{2}$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 |
| $Y$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 |

(b) [5 pts] What are the maximum likelihood estimates for $p_{1}, p_{2}$ and $q$ ?

$$
p_{1}=\ldots \quad p_{2}=\square \quad q=
$$

## Q9. [12 pts] Beyond Ordinary Pruning

Important: For all following parts, assume that the children of a node are visited in left-to-right order. You should not prune on equality (This also applies to any bound on utilities, if any. For example, given all utilities are less than or equal to 10 , you should not prune after seeing a node with utility of 10 .)
(a) [3 pts] Consider a two-player game in which both players alternate moves and each player seeks to maximize its own utility. At a leaf node $s$, utilities are represented as a tuple $U(s)=\left(U_{1}(s), U_{2}(s)\right)$, with the $i$-th component corresponding to the utility of the $i$-th player.

For the following special cases of two-player games, select all of the following in which pruning is never possible, given just this information about the relationship between utilities $U_{1}$ and $U_{2}$. Select "None of the above" if none of the options apply.

| $\square$ |
| :--- |
| $\square<U_{1}(s), U_{2}(s)<M$ for all terminal states $s$, where $M$ is a positive constant |
| $\square$ |
| $\square$ |
| $\square$ |
| $\square$ |
| $\square$ |$(s)+U_{2}(s)=M$ for all terminal states $s$, where $M \neq 0$ is a constant $\quad U_{2}(s)$ for all terminal states $s$.

None of the above
(b) Now we consider a three-player game similarly defined as in part (a). Then at a leaf node $s$, utilities are represented as a 3-tuple $U(s)=\left(U_{1}(s), U_{2}(s), U_{3}(s)\right)$, where the player going first (at the top of the tree) maximizes $U_{1}$, the player going second maximizes $U_{2}$, and the player going last maximizes $U_{3}$.
(i) $[2 \mathrm{pts}]$ Fill in the values at all nodes. Note that all players maximize their own respective utilities.

(ii) [3 pts] Without any further information, select all terminal nodes that can be pruned. Or check "None" if no node can be pruned.
Reminder: A node can be pruned only if the node's utilities can have no effect on the utilities at the root, irrespective of the node's utilities, and the utilities of nodes not yet visited by the left-to-right depth-first traversal.

(iii) [4 pts] Now we are given that for all terminal states $s$ the following holds true:

- $U_{i}(s) \geq 0 \quad \forall i=1,2,3$
- $\sum_{i=1}^{3} U_{i}(s) \leq 9$

Select all terminal nodes that can be pruned. Or check "None" if no node can be pruned.$a$$b$c $\square$ d $\square e$$f$ $\square$ None

## Q10. [10 pts] Iterative Deepening Search

Pacman is performing search in a maze again! The search graph has a branching factor of $b$, a solution of depth d, a maximum depth of $m$, and edge costs that may not be integers. Although he knows breadth first search returns the solution with the smallest depth, it takes up too much space, so he decides to try using iterative deepening. As a reminder, in standard depth-first iterative deepening we start by performing a depth first search terminated at a maximum depth of one. If no solution is found, we start over and perform a depth first search to depth two and so on. This way we obtain the shallowest solution, but use only $\mathrm{O}(\mathrm{bd})$ space.

But Pacman decides to use a variant of iterative deepening called iterative deepening $\mathbf{A}^{*}$, where instead of limiting the depth-first search by depth as in standard iterative deepening search, we can limit the depth-first search by the $f$ value as defined in A* search. As a reminder $f[$ node $]=g[$ node $]+h[$ node $]$ where $g[$ node $]$ is the cost of the path from the start state and $h[n o d e]$ is a heuristic value estimating the cost to the closest goal state.

In this question, all searches are tree searches and not graph searches.
(a) [7 pts] Complete the pseudocode outlining how to perform iterative deepening $\mathrm{A}^{*}$ by choosing the option from the next page that fills in each of these blanks. Iterative deepening $A^{*}$ should return the solution with the lowest cost when given a consistent heuristic. Note that cutoff is a boolean and new-limit is a number.

```
function Iterative-Deepening-Tree-Search (problem)
    start-node \(\leftarrow\) Make-Node(Initial-State \([p r o b l e m])\)
    limit \(\leftarrow f\) [start-node]
    loop
        fringe \(\leftarrow\) MAKE-Stack (start-node)
        new-limit \(\leftarrow\) (i)
        cutoff \(\leftarrow \quad\) (ii)
        while fringe is not empty do
            node \(\leftarrow\) Remove-Front(fringe)
            if Goal-Test(problem, State[node]) then
                return node
            end if
            for child-node in EXPAND(State[node], problem) do
                if \(f[\) child-node \(] \leq\) limit then
                fringe \(\leftarrow\) INSERT(child-node, fringe)
                    new-limit \(\leftarrow \square\) (iii)
                        cutoff \(\leftarrow \square(\mathrm{iv})\)
                else
                        new-limit \(\leftarrow \square\)
                        cutoff \(\leftarrow \square\)
            end if
            end for
        end while
        if not cutoff then
            return failure
        end if
        limit \(\leftarrow \quad\) (vii)
    end loop
end function
```

$\mathbf{A}_{1}-\infty$
$\mathbf{A}_{2} \square$
$\mathbf{A}_{3} \infty$
$\mathbf{A}_{4}$ limit
$\mathbf{B}_{1} \xrightarrow{\text { True }}$

$\mathbf{B}_{4}$ not cutoff
$\mathbf{C}_{4}$ new-limit + f[child-node]
$\mathbf{C}_{\mathbf{5}} \begin{aligned} & \text { MIN }(\text { new-limit }, \\ & f[\text { node }])\end{aligned}$

$\mathrm{C}_{8}$
MAX (new-limit, f[child-node])

| (i) $[1 \mathrm{pt}]$ | $\bigcirc \mathbf{A}_{1}$ | $\bigcirc \mathbf{A}_{\mathbf{2}}$ | $\bigcirc \mathbf{A}_{3}$ | $\bigcirc \mathrm{A}_{4}$ |
| :---: | :---: | :---: | :---: | :---: |
| (ii) $[1 \mathrm{pt}]$ | $\bigcirc \mathrm{B}_{1}$ | $\bigcirc \mathbf{B}_{2}$ | $\bigcirc \mathrm{B}_{3}$ | $\bigcirc \mathrm{B}_{4}$ |
| (iii) $[1 \mathrm{pt}]$ | $\begin{aligned} & \bigcirc \mathbf{C}_{1} \\ & \bigcirc \mathbf{C}_{5} \end{aligned}$ | $\begin{aligned} & \bigcirc \mathbf{C}_{2} \\ & \bigcirc \mathbf{C}_{6} \end{aligned}$ | $\begin{aligned} & \mathrm{C}_{3} \\ & \mathrm{C}_{7} \end{aligned}$ | $\begin{aligned} & \mathbf{C}_{4} \\ & \mathbf{C}_{8} \end{aligned}$ |
| (iv) $[1 \mathrm{pt}]$ | $\bigcirc \mathrm{B}_{1}$ | $\bigcirc \mathrm{B}_{2}$ | $\bigcirc \mathrm{B}_{3}$ | $\mathrm{B}_{4}$ |
| (v) $[1 \mathrm{pt}]$ | $\bigcirc \begin{aligned} & \mathbf{C}_{1} \\ & \mathbf{C}_{5} \end{aligned}$ | $\begin{aligned} & \bigcirc \mathbf{C}_{2} \\ & \bigcirc \mathrm{C}_{6} \end{aligned}$ | $\bigcirc \mathbf{C}_{\mathbf{3}} \bigcirc \mathbf{C}_{7}$ | $\begin{aligned} & \mathbf{C}_{4} \\ & \mathbf{C}_{8} \end{aligned}$ |
| (vi) $[1 \mathrm{pt}]$ | $\bigcirc \mathrm{B}_{1}$ | $\bigcirc \mathrm{B}_{2}$ | $\bigcirc \mathrm{B}_{3}$ | $\mathrm{B}_{4}$ |
| (vii) $[1 \mathrm{pt}]$ | $\begin{aligned} & \bigcirc \mathbf{C}_{1} \\ & \bigcirc \mathbf{C}_{5} \end{aligned}$ | $\begin{aligned} & \bigcirc \mathbf{C}_{2} \\ & \mathbf{C}_{6} \end{aligned}$ | $\begin{aligned} & \mathbf{C}_{3} \\ & \mathbf{C}_{7} \end{aligned}$ | $\begin{aligned} & \mathbf{C}_{4} \\ & \mathbf{C}_{8} \end{aligned}$ |

(b) [3 pts] Assuming there are no ties in $f$ value between nodes, which of the following statements about the number of nodes that iterative deepening A* expands is True? If the same node is expanded multiple times, count all of the times that it is expanded. If none of the options are correct, mark None of the above.

The number of times that iterative deepening $A^{*}$ expands a node is greater than or equal to the number of times $\mathrm{A}^{*}$ will expand a node.
$\bigcirc$
The number of times that iterative deepening $A^{*}$ expands a node is less than or equal to the number of times A* will expand a node.We don't know if the number of times iterative deepening $A^{*}$ expands a node is more or less than the number of times $\mathrm{A}^{*}$ will expand a node.

None of the above

