CS 188 Spring 2017 | Introduction to Artificial Intelligence | **Midterm 2**

- You have approximately 110 minutes.
- The exam is closed book, closed calculator, and closed notes except your one-page crib sheet.
- Mark your answers ON THE EXAM ITSELF. If you are not sure of your answer you may wish to provide a *brief* explanation. All short answer sections can be successfully answered in a few sentences AT MOST.
- For multiple choice questions,
  - ☐ means mark **all options** that apply
  - ◯ means mark a **single choice**
  - When selecting an answer, please fill in the bubble or square **completely** (● and ■)

First name	
Last name	
SID	
Student to your right	
Student to your left	

Your Discussion/Exam Prep* TA (fill all that apply):

☐ Brijen (Tu)
☐ Peter (Tu)
☐ David (Tu)
☐ Nipun (Tu)
☐ Wenjing (Tu)
☐ Aaron (W)
☐ Mitchell (W)
☐ Abhishek (W)
☐ Caryn (W)
☐ Anwar (W)
☐ Aarash (W)
☐ Daniel (W)
☐ Yuchen* (Tu)
☐ Andy* (Tu)
☐ Nikita* (Tu)
☐ Shea* (W)
☐ Daniel* (W)

For staff use only:

Q1.	Probability Warm-Up	/9
Q2.	Dressed to Impress	/12
Q3.	Independence	/12
Q4.	Markov Model Jambalaya	/12
Q5.	Bayes Nets: Elimination	/16
Q6.	Bayes Nets: Sampling	/16
Q7.	Mini Forward	/10
Q8.	HMM: Human-Robot Interaction	/8
Q9.	A Not So Random Walk	/5
	Total	/100

THIS PAGE IS INTENTIONALLY LEFT BLANK

Q1. [9 pts] Probability Warm-Up

(a) [4 pts] You have three coins in your pocket:

- Coin 1 is a fair coin that comes up heads with probability $1/2$.
- Coin 2 is a biased coin that comes up heads with probability $1/4$.
- Coin 3 is a biased coin that comes up heads with probability $3/4$.

Suppose you pick one of the coins uniformly at random and flip it three times. If you observe the sequence HHT (where H stands for heads and T stands for tails), what is the probability that you chose Coin 3?

$P(C_3 \mid HHT) =$ ____________________

(b) [5 pts] Suppose X and Y are **independent** random variables over the domain $\{1, 2, 3\}$ with $P(X = 3) = 1/6$. Given the following partially specified joint distribution, what are the remaining values? Write your answers as simplified fractions in the blanks.

$P(X = 1, Y = 1) = 1/4$	$P(X = 2, Y = 1) = 1/6$	$P(X = 3, Y = 1) =$ ________
$P(X = 1, Y = 2) = 1/16$	$P(X = 2, Y = 2) = 1/24$	$P(X = 3, Y = 2) =$ ________
$P(X = 1, Y = 3) =$ ________	$P(X = 2, Y = 3) =$ ________	$P(X = 3, Y = 3) =$ ________

Q2. [12 pts] Dressed to Impress

(a) [12 pts] Alice is invited to a party tonight which is said to be once-in-a-lifetime. However, this mysterious party doesn't publicize who is going and thus Alice has no idea whether the size S of the party will be be large ($S = +s$) or tiny ($S = -s$). The size can affect the noise N outside the party, and it will either be noisy ($N = +n$) or quiet ($N = -n$). Alice has three dresses: blue, red and yellow. Each dress will have a different utility for her depending on the size of the party. Let's help her decide which will be best!

We have the following decision network, where circles are chance nodes, squares are decision nodes, and diamonds are utility nodes:

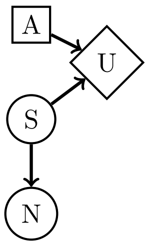

S	$P(S)$
+s	0.5
-s	0.5

S	N	$P(N \mid S)$
+s	+n	0.7
+s	-n	0.3
-s	+n	0.1
-s	-n	0.9

S	A	U
+s	blue	80
-s	blue	60
+s	red	40
-s	red	100
+s	yellow	60
-s	yellow	40

(i) [4 pts] What is the expected utility of wearing each dress, with both S and N unobserved?

- EU(A=blue) = ________
- EU(A=red) = ________
- EU(A=yellow) = ________

What is Alice's maximum expected utility?

- MEU({}) = ________

(ii) [6 pts] Suppose Alice observes that the party is quiet, $N = -n$. Compute the following conditional probabilities with this observation:

- $P(+s|-n) =$ ________
- $P(-s|-n) =$ ________

What is the expected utility of wearing each dress?

- EU(A=blue | $N = -n$) = ________
- EU(A=red | $N = -n$) = ________
- EU(A=yellow | $N = -n$) = ________

What is Alice's maximum expected utility given that $N = -n$?

- MEU({N=-n}) = ________

(iii) [2 pts] Construct a formula for $VPI(N)$ for the given network. To decouple this problem from your work above, use any of the symbolic terms from the following list (rather than plugging in numeric values):
$P(+n|+s)$, $P(+n|-s)$, $P(-n|+s)$, $P(-n|-s)$, $P(+n)$, $P(-n)$, $P(+s)$, $P(-s)$,
$MEU(\{\})$, $MEU(\{N = +n\})$, $MEU(\{N = -n\})$

- $VPI(N) =$________

Q3. [12 pts] Independence

In each part of this question, you are given a Bayes' net where the edges do not have a direction. Assign a direction to every edge (by adding an arrowhead at one end of each edge) to ensure that the Bayes' Net structure implies the assumptions provided. You cannot add new edges. The Bayes' nets can imply more assumptions than listed, but they *must* imply the ones listed. There may be more than one correct solution.

(a) [4 pts]

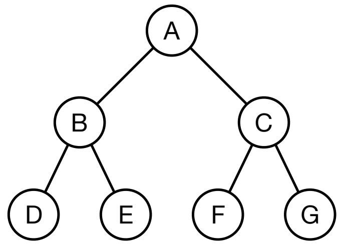

Assumptions:

- $A \perp\!\!\!\perp G$
- $D \perp\!\!\!\perp E$
- $E \perp\!\!\!\perp F$
- $F \perp\!\!\!\perp G \mid C$

(b) [4 pts]

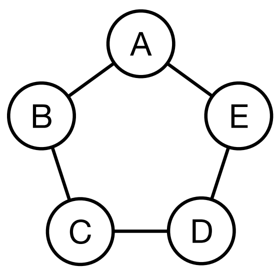

Assumptions:

- $B \perp\!\!\!\perp E$
- $E \perp\!\!\!\perp C \mid D$

(c) [4 pts]

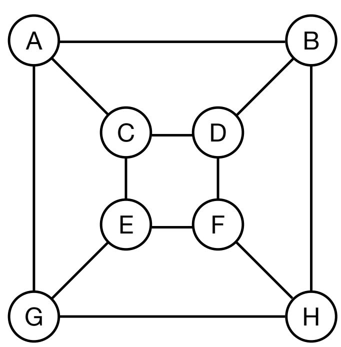

Assumptions:

- $F \perp\!\!\!\perp G$
- $F \perp\!\!\!\perp B \mid G$
- $D \perp\!\!\!\perp E \mid F$

Q4. [12 pts] Markov Model Jambalaya

(a) [8 pts] Consider a Markov chain for X specified by the following transition diagram. Please express all final answers as simplified fractions.

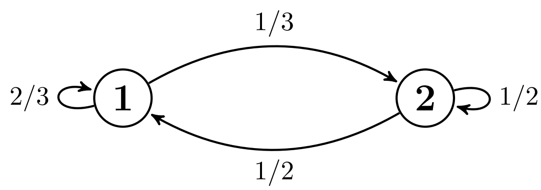

(i) [4 pts] Given that $X_0 = 1$, find $P(X_1)$ and $P(X_2)$.

$P(X_1 = 1) =$ ______________________

$P(X_1 = 2) =$ ______________________

$P(X_2 = 1) =$ ______________________

$P(X_2 = 2) =$ ______________________

(ii) [4 pts] Find $P(X_\infty)$, the stationary distribution of our Markov Chain.

$P(X_\infty = 1) =$ ______________________

$P(X_\infty = 2) =$ ______________________

(b) [4 pts] Consider a Markov chain for Y specified by the following transition diagram

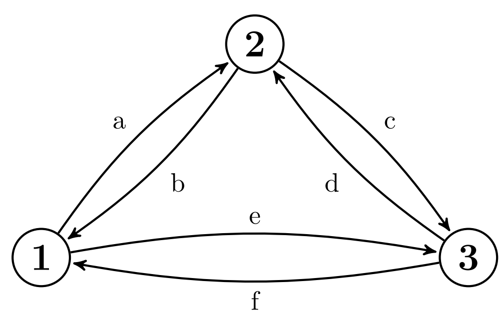

Given that $Y_0 = 1$, find $P(Y_1 = 3)$, $P(Y_2 = 3)$, $P(Y_3 = 3)$.

$P(Y_1 = 3) =$ ______________________

$P(Y_2 = 3) =$ ______________________

$P(Y_3 = 3) =$ ______________________

Q5. [16 pts] Bayes Nets: Elimination

(a) [5 pts] Consider running variable elimination on the Bayes Net shown below.

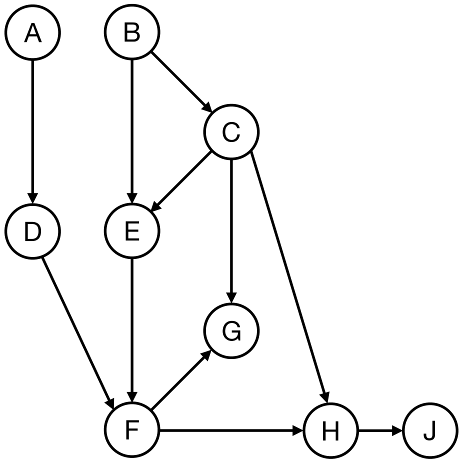

First, we eliminate D to create a factor f_1
Next, we eliminate E to create a factor f_2
Next, we eliminate H to create a factor f_3

From the list below, select all factors that remain after D, E and H have been eliminated.

☐ f_1	☐ $P(A)$	☐ $P(D\|A)$	☐ $P(G\|C)$
☐ f_2	☐ $P(A\|F)$	☐ $P(D\|A,F)$	☐ $P(G\|C,F)$
☐ f_3	☐ $P(B)$	☐ $P(E\|B)$	☐ $P(H\|F)$
	☐ $P(B\|A)$	☐ $P(E\|C)$	☐ $P(H\|C,F)$
	☐ $P(B\|C)$	☐ $P(E\|B,C)$	☐ $P(H\|J)$
	☐ $P(C)$	☐ $P(F)$	☐ $P(J\|C)$
	☐ $P(C\|B)$	☐ $P(F\|A,B)$	☐ $P(J\|F)$
		☐ $P(F\|G)$	☐ $P(J\|H)$

(b) [4 pts] Consider the Bayes Net shown below. Each variable in the Bayes Net can take on two possible values.

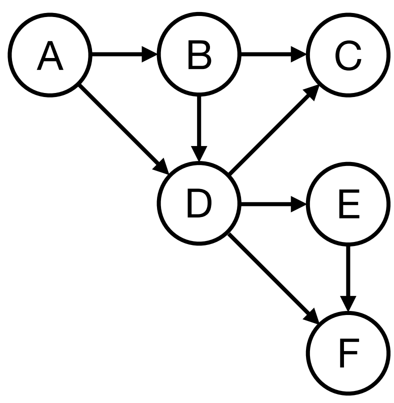

You are given the query $P(C|F)$, which you would like to answer using variable elimination. Please find a variable elimination ordering where the largest intermediate factor created during variable elimination is as small as possible.

Elimination ordering: ____________________

(c) [7 pts] Consider doing inference in an m x n lattice Bayes Net, as shown below. The network consists of mn binary variables $V_{i,j}$, and you have observed that $V_{m,n} = +v_{m,n}$.

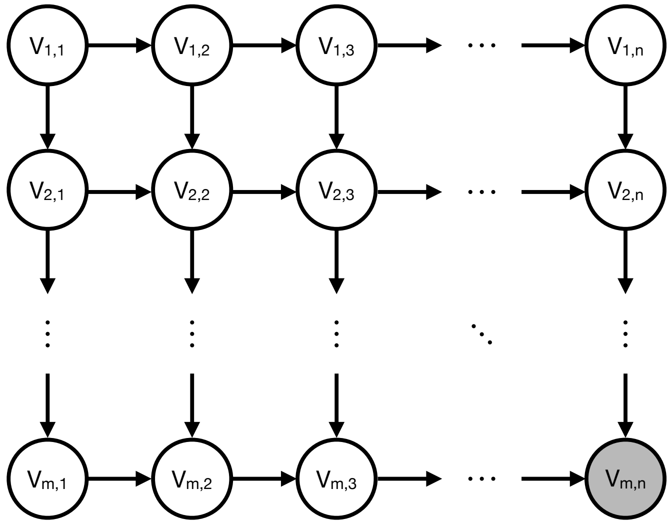

You wish to calculate $P(V_{1,1} | +v_{m,n})$ using variable elimination. To maximize computational efficiency, you wish to use a variable elimination ordering for which the size of the largest generated factor is as small as possible.

(i) [4 pts] First consider the special case where $m = 4$ and $n = 5$. A reproduction of the lattice is shown below, with variable names for non-query variables omitted. Please provide your optimal elimination ordering for this example by numbering the nodes below in the order they will be eliminated (i.e. write a number such as $1, 2, 3, \ldots$ inside every node that will be eliminated.)

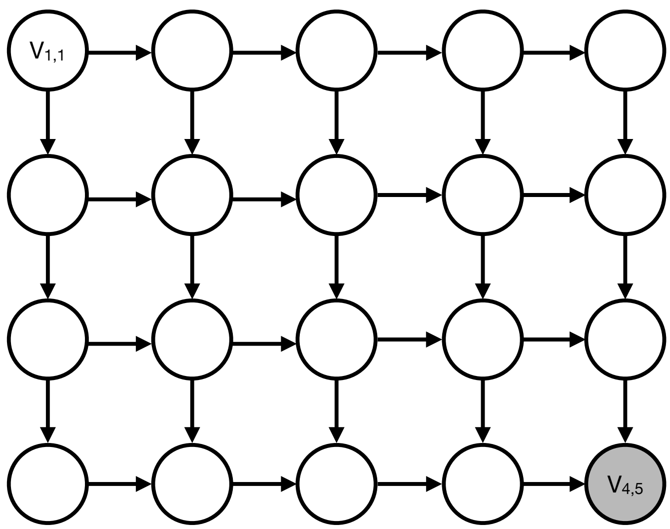

(ii) [3 pts] Now consider the general case (assume $m > 2$ and $n > 2$). What is the size of the largest factor generated under the most efficient elimination ordering? Your answer should be the number of rows in the factor's table, expressed in terms of m and n.

Size (number of rows) of the largest factor: ________________________

Q6. [16 pts] Bayes Nets: Sampling

Consider the following Bayes Net, where we have observed that $B = +b$ and $D = +d$.

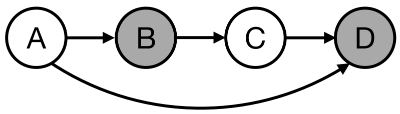

$P(A)$	
$+a$	0.5
$-a$	0.5

$P(B\|A)$		
$+a$	$+b$	0.8
$+a$	$-b$	0.2
$-a$	$+b$	0.4
$-a$	$-b$	0.6

$P(C\|B)$		
$+b$	$+c$	0.1
$+b$	$-c$	0.9
$-b$	$+c$	0.7
$-b$	$-c$	0.3

$P(D\|A,C)$			
$+a$	$+c$	$+d$	0.6
$+a$	$+c$	$-d$	0.4
$+a$	$-c$	$+d$	0.1
$+a$	$-c$	$-d$	0.9
$-a$	$+c$	$+d$	0.2
$-a$	$+c$	$-d$	0.8
$-a$	$-c$	$+d$	0.5
$-a$	$-c$	$-d$	0.5

(a) [4 pts] Consider doing Gibbs sampling for this example. Assume that we have initialized all variables to the values $+a, +b, +c, +d$. We then unassign the variable C, such that we have $A = +a$, $B = +b$, $C = ?$, $D = +d$. Calculate the probabilities for new values of C at this stage of the Gibbs sampling procedure.

$P(C = +c$ at the next step of Gibbs sampling$) =$ ____________________

$P(C = -c$ at the next step of Gibbs sampling$) =$ ____________________

(b) [8 pts] Consider a sampling scheme that is a hybrid of rejection sampling and likelihood-weighted sampling. Under this scheme, we first perform rejection sampling for the variables A and B. We then take the sampled values for A and B and extend the sample to include values for variables C and D, using likelihood-weighted sampling.

(i) [2 pts] Below is a list of candidate samples. Mark the samples that would be rejected by the rejection sampling portion of the hybrid scheme.

☐ $-a$ $-b$
☐ $+a$ $+b$
☐ $+a$ $-b$
☐ $-a$ $+b$

(ii) [4 pts] To decouple from part (i), you now receive a *new* set of samples shown below. Fill in the weights for these samples under our hybrid scheme.

				Weight
$-a$	$+b$	$-c$	$+d$	____________________
$+a$	$+b$	$-c$	$+d$	____________________
$+a$	$+b$	$-c$	$+d$	____________________
$-a$	$+b$	$+c$	$+d$	____________________
$+a$	$+b$	$+c$	$+d$	____________________

(iii) [2 pts] Use the weighted samples from part (ii) to calculate an estimate for $P(+a| +b, +d)$.

The estimate of $P(+a| +b, +d)$ is ____________________

(c) [4 pts] We now attempt to design an alternative hybrid sampling scheme that combines elements of likelihood-weighted and rejection sampling. For each proposed scheme, indicate whether it is valid, i.e. whether the weighted samples it produces correctly approximate the distribution $P(A, C|+b, +d)$.

(i) [2 pts] *First collect a likelihood-weighted sample for the variables A and B. Then switch to rejection sampling for the variables C and D. In case of rejection, the values of A and B and the sample weight are **thrown away**. Sampling then restarts from node **A**.*

◯ Valid ◯ Invalid

(ii) [2 pts] *First collect a likelihood-weighted sample for the variables A and B. Then switch to rejection sampling for the variables C and D. In case of rejection, the values of A and B and the sample weight are **retained**. Sampling then restarts from node **C**.*

◯ Valid ◯ Invalid

Q7. [10 pts] Mini Forward

(a) [10 pts] Let $Z_0, Z_1, Z_2, ...$ be a Markov Chain. You believe that any state Z_i is represented as the union of two states X_i and Y_i. Therefore $P(Z_i) = P(X_i, Y_i)$. The following diagrams illustrate the independence assumptions of the problem.

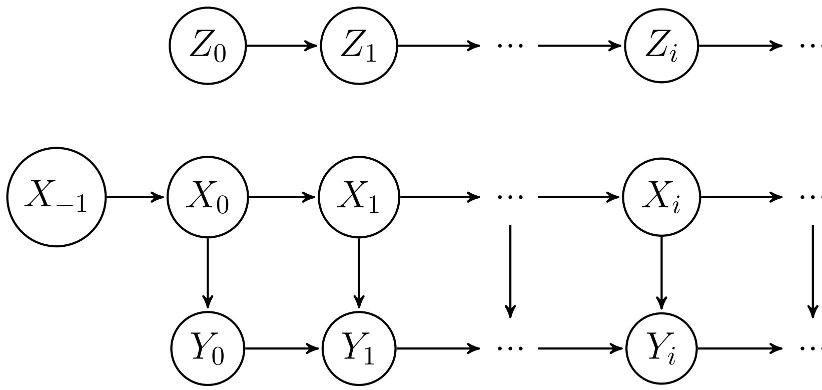

Recall that the mini forward algorithm allows us to recursively calculate the forward simulation of our model when we are given the initial state and transition dynamics. In other words, it gives us a way to find $P(Z_i)$.

Suppose we are given $X_{-1} = x$ and the tables $P(Y_0 \mid X_0)$, $P(X_i \mid X_{i-1})$, and $P(Y_i \mid Y_{i-1}, X_i)$. Derive the mini forward algorithm that tells us $P(Z_i)$ in terms of given distributions.

(i) [3 pts] Write out the base case of the mini forward algorithm for $P(Z_i)$.

(ii) [7 pts] Write out the recursive component of the mini forward algorithm for $P(Z_i)$. It is not necessary to solve the problem, but you can consider Z_i to be equal to (X_i, Y_i).

Q8. [8 pts] HMM: Human-Robot Interaction

In the near future, autonomous robots would live among us. Therefore, it is important for the robots to know how to properly act in the presence of humans. In this question, we are exploring a simplified model of this interaction. Here, we are assuming that we can observe the robot's actions at time t, R_t, and an evidence observation, E_t, directly caused by the human action, H_t. Humans actions and Robots actions from the past time-step affect the Human's and Robot's actions in the next time-step. In this problem, we will remain consistent with the convention that capital letters (H_t) refer to random variables and lowercase letters (h_t) refer to a particular value the random variable can take. The structure is given below:

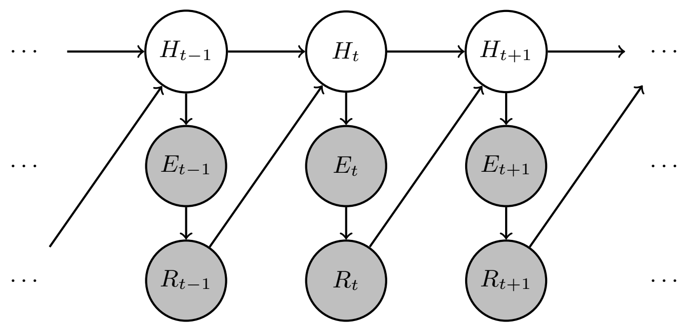

You are supplied with the following probability tables: $P(R_t \mid E_t)$, $P(H_t \mid H_{t-1}, R_{t-1})$, $P(H_0)$, $P(E_t \mid H_t)$.

Let us derive the forward algorithm for this model. We will split our computation into two components, a **time-elapse update** expression and a **observe update** expression.

(a) [2 pts] We would like to incorporate the evidence that we observe at time t. Using the time-lapse update expression we will derive separately, we would like to find the **observe update** expression:

$$O(H_t) = P(H_t | e_{0:t}, r_{0:t})$$

In other words, we would like to compute the distribution of potential human states at time t given all observations up to and including time t. In addition to the conditional probability tables associated with the network's nodes, we are given $T(H_t) = P(H_t \mid E_{0:t-1}, R_{0:t-1})$, which we will assume is correctly computed in the time-elapse update that we will derive in the next part. From the options below, select *all* the options that **both** make valid independence assumptions and would evaluate to the observe update expression.

☐ $\frac{P(H_t|e_{0:t-1},r_{0:t-1})P(e_t|H_t)P(r_t|e_t)}{\sum_{h_t} P(h_t|e_{0:t-1},r_{0:t-1})P(e_t|h_t)P(r_t|e_t)}$

☐ $\frac{P(H_t|e_{0:t-1},r_{0:t-1})P(e_t|H_t)}{\sum_{h_t} P(h_t|e_{0:t-1},r_{0:t-1})P(e_t|h_t)}$

☐ $\frac{\sum_{e_t} P(H_t|e_{0:t-1},r_{0:t-1})P(e_t|H_t)}{\sum_{h_t} P(h_t|e_{0:t-1},r_{0:t-1})P(e_t|r_{t-1},H_{t-1})}$

☐ $\sum_{r_{t-1}} P(H_t \mid e_{0:t-1}, r_{0:t-1})P(r_{t-1} \mid e_{t-1})$

☐ $\sum_{r_t} P(H_t \mid e_{0:t-1}, r_{0:t-1})P(r_t \mid r_{t-1}, e_t)$

☐ $\sum_{h_{t+1}} P(H_t \mid e_{0:t-1}, r_{0:t-1})P(h_{t+1} \mid r_t)$

The structure below is identical to the one in the beginning of the question and is repeated for your convenience.

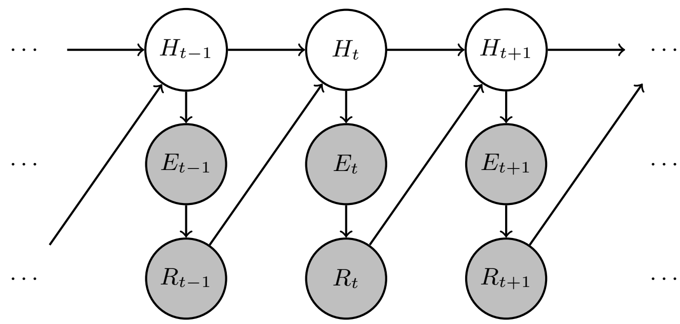

(b) [6 pts] We are interested in predicting what the state of human is at time t (H_t), given all the observations through $t-1$. Therefore, the **time-elapse update** expression has the following form:

$$T(H_t) = P(H_t|e_{0:t-1}, r_{0:t-1})$$

Derive an expression for the given time-elapse update above using the probability tables provided in the question and the observe update expression, $O(H_{t-1}) = P(H_{t-1}|e_{0:t-1}, r_{0:t-1})$. Write your final expression in the space provided at below. You may use the function O in your solution if you prefer.

$P(H_t|E_{0:t-1}, R_{0:t-1}) =$ ____

Q9. [5 pts] A Not So Random Walk

Pacman is trying to predict the position of a ghost, which he knows has the following transition graph:

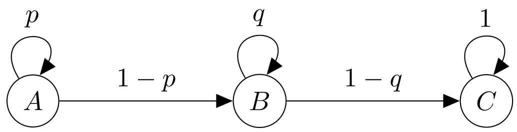

Here, $0 < p < 1$ and $0 < q < 1$ are arbitrary probabilities. It is known that the ghost always starts in state A. For this problem, we consider time to begin at 0. For example, at time 0, the ghost is in A with probability 1, and at time 1, the ghost is in A with probability p or in B with probability $1 - p$.

In all of the following questions, you may assume that n is large enough so that the given event occurs with non-zero probability.

Note: For full credit, you may refer to the correct answer of a **previous** part using the notation $f_k(t)$, where k is a part number and t is a time. For instance, to refer to the correct answer for part (ii) for time $n - 1$, you could write $f_{(\text{ii})}(n - 1)$. You may only refer to previous parts, not future parts.

Please note the low point value for each subproblem and allocate time accordingly. Answers should be simplified to the extent possible.

(i) [1 pt] Suppose $p \neq q$. What is the probability that the ghost is in A at time n?

(ii) [1 pt] Suppose $p \neq q$. What is the probability that the ghost first reaches B at time n?

(iii) [1 pt] Suppose $p \neq q$. What is the probability that the ghost is in B at time n?

(iv) [1 pt] Suppose $p \neq q$. What is the probability that the ghost first reaches C at time n?

(v) [1 pt] Suppose $p \neq q$. What is the probability that the ghost is in C at time n?