

- You have approximately 2 hours and 50 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 with *circular bubbles*, you should only mark ONE option; for those with *checkboxes*, you should mark ALL that apply (which can range from zero to all options)

First name	
Last name	
edX username	

For staff use only:

Q1. Potpourri	/13
Q2. More Advanced Problems	/19
Q3. Variable Elimination	/12
Q4. Bayes' Nets: Representation and Independence	/16
Q5. VPI	/13
Q6. Sampling as an MDP	/13
Q7. HMMs	/19
Total	/105

THIS PAGE IS INTENTIONALLY LEFT BLANK

Q1. [13 pts] Potpourri

(a) Probability

	A	B	$P(B A)$	B	C	$P(C B)$	C	D	$P(D C)$	
A	$P(A)$	+a	+b	0.9	+b	+c	0.8	+c	+d	0.25
+a	0.8	+a	-b	0.1	+b	-c	0.2	+c	-d	0.75
-a	0.2	-a	+b	0.6	-b	+c	0.8	-c	+d	0.5
		-a	-b	0.4	-b	-c	0.2	-c	-d	0.5

Using the table above and the assumptions per subquestion, calculate the following probabilities given no independence assumptions. If it is impossible to calculate without more independence assumptions, specify the least number of independence assumptions that would allow you to answer the question (don't do any computation in this case).

(i) [1 pt] $P(+a, -b) = .8 * .1 = .08$

(ii) [1 pt] $P(-a, -b, +c) =$ **Given that C is independent of A given B**

(iii) [1 pt] Now assume C is independent of A given B and D is independent of everything else given C. Calculate $P(+a, -b, +c, +d) =$ or say what other independence assumptions are necessary. **$.8 * .1 * .8 * .25$**

(b) Independence

(i) [2 pts] Mark all expressions which indicate that X is independent of Y given Z.

$P(X, Y | Z) = P(X | Z)P(Y | Z)$

$P(X | Y, Z) = P(X | Z)$

$P(X, Y, Z) = P(X, Z)P(Y)$

None of the above

(ii) [2 pts] Fill in the circles of **all** expressions that are equal to $P(\mathbf{R}, \mathbf{S}, \mathbf{T})$, **given no independence assumptions**:

$P(R | S, T) P(S | T) P(T)$

$P(T, S | R) P(R)$

$P(T | R, S) P(R) P(S)$

$P(T | R, S) P(R, S)$

$P(R | S) P(S | T) P(T)$

None of the above

$P(R | S, T) P(S | R, T) P(T | R, S)$

(c) Bayes Nets

(i) [1 pt] During variable elimination, the ordering of elimination does not affect the final answer.

True

False

(ii) [1 pt] During variable elimination, the ordering of elimination does not affect the runtime.

True


False

(d) [1 pt] **Sampling** For the following descriptions, provide the sampling method that is being described.

- Tally all of the values, but ignore anything that doesn't match the conditional evidence. **Rejection Sampling**
- Tally the values, weighting by the value of actually seeing that evidence based on the parents. **Likelihood Sampling**
- Sample from the original joint distribution, ignoring the evidence. **Prior Sampling**

(e) [3 pts] **Stationary Distributions**

Consider a Markov chain with 3 states and transition probabilities as shown below:



Compute the stationary distribution. That is, compute $P_\infty(A), P_\infty(B), P_\infty(C)$.

$$P_\infty(A) = 0.4$$

$$P_\infty(B) = 0.4$$

$$P_\infty(C) = 0.2$$

Q2. [19 pts] More Advanced Problems


(a) [2 pts] Probability

Assume that Q, R, S, and T are all independent binary random variables. For the following probabilities, assume you are given a table of all values of each probability. Write down what the sum of all of the values in the table would equal (as a number). If it is impossible to tell, write down "Impossible".

- $P(+r | S) =$
 Impossible
 $P(R, T | +s) =$
 1
 $P(R | Q, S, T) =$
 8
 $P(R, T | +s, Q) =$
 2

(b) [3 pts] Sampling

Consider the following Bayes Net and corresponding probability tables.



Fill in the following table with the probabilities of drawing each respective sample given that we are using each of the following sampling techniques. For rejection sampling, we say that a sample has been drawn only if it is not rejected. You may leave your answer in the form of an expression such as $.8 \cdot .4$ without multiplying it out. (Hint: $P(f, m) = .181$)

$P(\text{sample} \text{method})$	$(+r, +e, -w, +m, +f)$	$(+r, -e, +w, -m, +f)$
prior sampling	$.4 * .3 * .1 * .45 * .75 = .00405$	$.4 * .7 * .9 * .65 * .15 = 0.02457$
rejection sampling	$\frac{P(+r, +e, -w, +m, +f)}{P(+m, +f)} = \frac{.00405}{.181} = .0224$	0
likelihood weighting	$P(+r)P(+e +r)P(-w +r) = .4 * .3 * .1 = .012$	0


(c) HMM

(i) [2 pts] Consider an HMM with state variables $\{X_i\}$ and emission variables $\{Y_i\}$. Which of the following assertions are true?

- X_i is always conditionally independent of Y_{i+1} given X_{i+1} .
- There exists an HMM where X_i is conditionally independent of Y_i given X_{i+1} .
- If $Y_i = X_i$ with probability 1, and the state space is of size k , then the most efficient algorithm for computing $p(X_i|y_1 \dots, y_t)$ takes $O(k)$ or less time.
- If we take the Bayes net below for part (ii) and reverse the vertical arrows so that we have edges from each Y_i to X_i , the result is an HMM.
- None of the above

(ii) [7 pts]

Likelihood weighting.



Assume each of the variables $X_1, X_2, X_3, Y_1, Y_2, Y_3$ are binary with domains $\{\pm 1\}$. Assuming a uniform starting distribution $[.5, .5]$, and emission probabilities all equal to:

y	$P_E(y -1)$	$P_E(y 1)$
-1	.2	.7
1	.8	.3

And transition probabilities all equal to:

x	$P_T(x -1)$	$P_T(x 1)$
-1	.4	.6
1	.6	.4

Assume that the samples are $(X_1, X_2, X_3, Y_1, Y_2, Y_3)$. Fill the following table with the samples' likelihood sampling weight (conditioning on $X_2 = 1$ and $Y_3 = 1$) and the probability of drawing the sample during likelihood weighting (You can leave the desired values as products). If a sample is invalid, say so.

Index	Sample	Weight in Likelihood Sampling	Probability of sample $P(x_1, x_2, x_3, y_1, y_2, y_3)$
1.	(1,1,1,1,1,1)		
2.	(1,1,-1,1,-1,1)		
3.	(-1,1,-1,1,-1,1)		
4.	(1,-1,1,1,-1,-1)		
5.	(1,-1,-1,-1,1,1)		

Index	Sample	Weight in Likelihood Sampling	Probability of sample $P(a, b, c, d, e, f)$
1.	(1,1,1,1,1,1)	.4*.3	.5*.4*(.3) ²
2.	(1,1,-1,1,-1,1)	.4*.8	.5*.6*.3*.7
3.	(-1,1,-1,1,-1,1)	.6*.8	0.5*.6*.8*.7
4.	(1,-1,1,1,-1,-1)	Invalid sample	0
5.	(1,-1,-1,-1,1,1)	Invalid sample	0

What is $P(A = 1|B = 1, F = 1)$?


$$\frac{P(A=1)P(B=1|A=1)}{P(A=1)P(B=1|A=1)+P(A=-1)P(B=1|A=-1)} = .4$$

Using Likelihood sampling what is $\hat{P}(A = 1|B = 1, F = 1)$?

$$\frac{.4*.3+.4*.8}{.4*.3+.4*.8+.6*.8} = \frac{3+8}{3+8+12} = \frac{11}{23}$$

Q3. [12 pts] Variable Elimination

The following questions use the Bayes' net below. All variables have binary domains:



- (a) [6 pts] Karthik wants to see the ocean, and so decides to compute the query $P(B, E, A, C, H)$. He wants you to help him run variable elimination to compute the answer, with the following elimination ordering: I, D, G, F . Complete the following description of the factors generated in this process:

He initially has the following factors to start out with:

$$P(A), P(B), P(C|B, G), P(D|I), P(E|A, B), P(F|E), P(G|E), P(H|C), P(I|C)$$

When eliminating I we generate a new factor f_1 as follows:

$$f_1(C, D) = \sum_i P(i|C)P(D|i)$$

This leaves us with the factors:

$$P(A), P(B), P(C|B, G), P(E|A, B), P(F|E), P(G|E), P(H|C), f_1(C, D)$$

When eliminating D we generate a new factor f_2 as follows:

$$f_2(C) = \sum_d f_1(C, D)$$

This leaves us with the factors:

$$P(A), P(B), P(C|B, G), P(E|A, B), P(F|E), P(G|E), P(H|C), f_2(C)$$

When eliminating G we generate a new factor f_3 as follows:

$$f_3(C, B, E) = \sum_g P(C|B, g)P(g|E)$$

This leaves us with the factors:

$$P(A), P(B), P(C|B, G), P(E|A, B), P(F|E), P(G|E), P(H|C), f_2(C), f_3(C, B, E)$$

When eliminating F we generate a new factor f_4 as follows:

$$f_4(E) = \sum_f P(f|E)$$

This leaves us with the following factors. Another acceptable answer involved noting the fact that summing out the above factor yields 1, and so not appending $f_4(E)$ was fine.

$$P(A), P(B), P(C|B, G), P(E|A, B), P(G|E), P(H|C), f_2(C), f_3(C, B, E), f_4(E)$$

- (b) [2 pts] Among f_1, f_2, f_3, f_4 , which is the largest factor generated, and how large is it? Assume all variables have binary domains and measure the size of each factor by the number of rows in the table that would represent the factor.

$f_2(A, +c, E, F)$ is the largest factor generated. It has 3 non-instantiated variables, hence $2^3 = 8$ entries.

- (c) [4 pts] Given a list of all factors in a Bayes net, suppose that there exists a variable V that only occurs once in the entire list. Which of the following statements must be true when running variable elimination?


- The factor containing variable V must have precisely 2 variables.
- Eliminating V produces a factor whose size is lesser than or equal to the largest factor size during the full variable elimination process.
- Variable V must be a leaf node; that is, V cannot have any children nodes.
- The factor containing variable V must contain an even number of variables.
- The factor containing variable V must contain an odd number of variables.
- Variable V must appear on the left hand side of the conditioning bar, i.e. the $|$, in the factor that it appears in.
- There must also exist a different variable W that appears only once in the entire list of factors.
- There must also exist a different variable W that appears more than once in the entire list of factors.
- None of the above

If a variable V appears only once in the list of factors, it cannot have any children (otherwise, it would appear both on the left-hand side and the right-hand side of the conditioning bar). As such, the third and sixth choices must be marked. Eliminating this variable produces a factor equal to one everywhere, so the size of the factor produced must be lesser than or equal to the largest factor generated during the full elimination process.

The number of variables in the factor containing V can be any number; it's just the number of parents of V plus 1 (for V itself). As a counterexample to the last two options, consider the Bayes net where V is the only variable. Then there isn't a different variable W to begin with.

Q4. [16 pts] Bayes' Nets: Representation and Independence

Parts (a), (b), and (c) pertain to the following Bayes' Net.



- (a) [1 pt] Express the joint probability distribution as a product of terms representing individual conditional probabilities tables associated with the Bayes Net.

$$P(A)P(C|A)P(B|A)P(D|B)P(E)P(F|D, E)P(G|D)$$

- (b) [1 pt] Assume each node can take on 4 values. How many entries do the factors at A, D, and F have?

A: 4

D: 4²

F: 4³

- (c) [2 pts] Mark the statements that are guaranteed to be true.

$B \perp\!\!\!\perp C$

$A \perp\!\!\!\perp F$

$D \perp\!\!\!\perp E|F$

$E \perp\!\!\!\perp A|D$

$F \perp\!\!\!\perp G|D$

$B \perp\!\!\!\perp F|D$

$C \perp\!\!\!\perp G$

$D \perp\!\!\!\perp E$


Parts (d) and (e) pertain to the following probability distribution tables. The joint distribution $P(A, B, C, D)$ is equal to the product of these probability distribution tables.

A	$P(A)$	A	B	$P(B A)$	B	C	$P(C B)$	C	D	$P(D C)$
+a	0.8	+a	+b	0.9	+b	+c	0.8	+c	+d	0.25
-a	0.2	+a	-b	0.1	+b	-c	0.2	+c	-d	0.75
		-a	+b	0.6	-b	+c	0.8	-c	+d	0.5
		-a	-b	0.4	-b	-c	0.2	-c	-d	0.5

(d) [2 pts] State all non-conditional independence assumptions that are implied by the probability distribution tables.

From the tables, we have $A \perp\!\!\!\perp B$ and $C \perp\!\!\!\perp D$. Then, we have every remaining pair of variables: $A \perp\!\!\!\perp C, A \perp\!\!\!\perp D, B \perp\!\!\!\perp C, B \perp\!\!\!\perp D$

(e) [3 pts] Circle all the Bayes net(s) that can represent a distribution that is consistent with the tables given.



The question asks for Bayes Nets that **can** represent the distribution in the tables. So, in the nets we circle, the only requirement must be that A and B must not be independent, and C and D must not be independent.

The top left, bottom left, and bottom right nets have arrows between the A-B nodes and the C-D nodes, so we can circle those.

The top middle net has C and D as independent (D is not connected to anything), so we cannot circle it. The bottom middle net has A and B as independent (common cause), so we cannot circle it.


The top right net seems like it could represent the distribution, because D-separation finds that: A and B are not guaranteed to be independent (common effect), and C and D are not guaranteed to be independent (causal chain). However, according to Part D, $A \perp\!\!\!\perp C, A \perp\!\!\!\perp D$, and $B \perp\!\!\!\perp C$, so all of the arrows in the net are vacuous. That means, in this net, A and B are independent, and C and D are independent, so we cannot circle this net.

You are building advanced safety features for cars that can warn a driver if they are falling asleep (A) and also calculate the probability of a crash (C) in real time. You have at your disposal 6 sensors (random variables):

- E : whether the driver's eyes are open or closed
- W : whether the steering wheel is being touched or not
- L : whether the car is in the lane or not
- S : whether the car is speeding or not
- H : whether the driver's heart rate is somewhat elevated or resting
- R : whether the car radar detects a close object or not

A influences $\{E, W, H, L, C\}$. C is influenced by $\{A, S, L, R\}$.

(f) [2 pts] Draw the Bayes Net associated with the description above by adding edges between the provided nodes where appropriate.



(g) [2 pts] Mark all the independence assumptions that must be true.

- | | |
|--|--|
| <input checked="" type="checkbox"/> $E \perp\!\!\!\perp S$ | <input type="checkbox"/> $L \perp\!\!\!\perp R C$ |
| <input checked="" type="checkbox"/> $W \perp\!\!\!\perp H A$ | <input checked="" type="checkbox"/> $W \perp\!\!\!\perp R$ |
| <input checked="" type="checkbox"/> $S \perp\!\!\!\perp R$ | <input type="checkbox"/> $A \perp\!\!\!\perp C$ |
| <input type="checkbox"/> $E \perp\!\!\!\perp L$ | <input type="checkbox"/> $E \perp\!\!\!\perp C L$ |

(h) [2 pts] The car's sensors tell you that the car is in the lane ($L = +l$) and that the car is not speeding ($S = -s$). Now you would like to calculate the probability of crashing, $P(C|+l, -s)$. We will use the variable elimination ordering R, A, E, W, H . Write down the largest factor generated during variable elimination. Box your answer.

Our factors if we don't observe evidence are $P(A), P(S), P(R), P(E|A), P(W|A), P(H|A), P(L|A), P(C|L, A, S, R)$. We observe evidence, and we have: $P(A), P(R), P(E|A), P(W|A), P(H|A), P(C|+l, A, -s, R)$. We first eliminate R , so we select $P(R)$ and $P(C|+l, A, -s, R)$ to get $f_1(C|+l, A, -s)$. Now we eliminate A , so we select $P(A), P(E|A), P(W|A), P(H|A), f_1(C|+l, A, -s)$ and get $f_2(C, E, W, H|+l, -s)$. We see that this must be the largest factor because this is the only factor we have left at this point, and variable elimination is not yet finished.

(i) [1 pt] Write down a more efficient variable elimination ordering, i.e. one whose largest factor is smaller than the one generated in the previous question.

Any ordering of the five variables where at least one of $\{E, W, H\}$ is before A would be more efficient than the previous ordering. As an example, R, E, W, H, A would work.

Q5. [13 pts] VPI

You are the latest contestant on Monty Hall's game show, which has undergone a few changes over the years. In the game, there are n closed doors: behind one door is a car ($U(car) = 1000$), while the other $n - 1$ doors each have a goat behind them ($U(goat) = 10$). You are permitted to open exactly one door and claim the prize behind it.

You begin by choosing a door uniformly at random.

- (a) [2 pts] What is your expected utility?

$$(1000 * \frac{1}{n} + 10 * \frac{n-1}{n}) \text{ or } (10 + 990 * \frac{1}{n})$$

Answer:

We can calculate the expected utility via the usual formula of expectation, or we can note that there is a guaranteed utility of 10, with a small probability of a bonus utility. The latter is a bit simpler, so the answers to the following parts use this form.

- (b) [4 pts] After you choose a door but before you open it, Monty offers to open k other doors, each of which are guaranteed to have a goat behind it.

If you accept this offer, should you keep your original choice of a door, or switch to a new door?

$$10 + 990 * \frac{1}{n}$$

$EU(keep)$:

$$10 + 990 * \frac{(n-1)}{n * (n-k-1)}$$

$EU(switch)$:

switch

Action that achieves MEU :

The expected utility if we keep must be the same as the answer from the previous part: the probability that we have a winning door has not changed at all, since we have gotten no meaningful information.

In order to win a car by switching, we must have chosen a goat door previously (probability $\frac{n-1}{n}$) and then switch to the car door (probability $\frac{1}{n-k-1}$).

Since $n - 1 > n - k - 1$ for positive k , switching gets a larger expected utility.

- (c) [2 pts] What is the value of the information that Monty is offering you?

$$990 * \frac{1}{n} * \frac{k}{n-k-1}$$

Answer:

The formula for VPI is $MEU(e) - MEU(\emptyset)$. Thus, we want the difference between $EU(switch)$ (the optimal action if Monty opens the doors) and our expected utility from part (a).

(It is true that $EU(keep)$ happens to have the same numeric expression as in part (a), but this fact is not meaningful in answering this part.)

(d) [2 pts] Monty is changing his offer!

After you choose your initial door, you are given the offer to choose any other door and open this second door. If you do, after you see what is inside the other door, you may switch your initial choice (to the newly opened door) or keep your initial choice.

What is the value of this new offer?

$$\frac{990}{n}$$

Answer:

Intuitively, if we take this offer, it is as if we just chose two doors in the beginning, and we win if either door has the car behind it. Unlike in the previous parts, if the new door has a goat behind it, it is not more optimal to switch doors.

Mathematically, letting D_i be the event that door i has the car, we can calculate this as $P(D_2 \cup D_1) = P(D_1) + P(D_2) = 1/n + 1/n = 2/n$, to see that $MEU(\text{offer}) = 10 + 990 * \frac{2}{n}$. Subtracting the expected utility without taking the offer, we are left with $990 * \frac{1}{n}$.

(e) [3 pts] Monty is generalizing his offer: you can pay $\$d^3$ to open d doors as in the previous part. (Assume that $U(\$x) = x$.) You may now switch your choice to any of the open doors (or keep your initial choice). What is the largest value of d for which it would be rational to accept the offer?

$$d = \sqrt{\frac{990}{n}}$$

Answer:

It is a key insight (whether intuitive or determined mathematically) that the answer to the previous part is constant for each successive door we open. Thus, the value of opening d doors is just $d * 990 * \frac{1}{n}$. Setting this equal to d^3 , we can solve for d .

Q6. [13 pts] Sampling as an MDP

- (a) (i) [1 pt] You are given a Bayes net with binary random variables A, B, C, D , and E . You want to estimate $P(A, B, C, E | +d)$ using rejection sampling. Which of the following quantities denotes the probability that a sample will be *rejected*? (Mark all that apply.)

- $P(+d)$
 $P(-d)$
 $1 - P(+d)$
 $1 - P(-d)$


- (ii) [1 pt] For the same Bayes net, you would like to estimate $P(A, B, C | +d, +e)$ using rejection sampling. Which of the following quantities denotes the probability that a sample will be *rejected*? (Mark all that apply.)

- $P(+d, +e)$
 $P(-d, -e)$
 $1 - P(+d, +e)$
 $1 - P(-d, -e)$
 $1 - P(+d)P(+e)$
 $1 - P(-d)P(-e)$

- (iii) [1 pt] For the same Bayes net, suppose additionally that $D \perp\!\!\!\perp E$. Which of the following quantities denotes the probability that a sample will be *rejected*?

- $P(+d, +e)$
 $P(-d, -e)$
 $1 - P(+d, +e)$
 $1 - P(-d, -e)$
 $1 - P(+d)P(+e)$
 $1 - P(-d)P(-e)$

- (b) [2 pts] Use the following Bayes Net for this question *only*:



In how many different orders could I sample from the random variables in this Bayes Net? (You may use simple arithmetic operations in your answer.)

4

This is equal to the number of linear orders in the graph. Since we can't sample a node before sampling its parent, we must first sample A and B (but the order of the two doesn't matter), then sample C and D (but again the order doesn't matter) then sample E . So the number of orderings is $2 \times 2 = 4$.

- (c) (i) [1 pt] In a general Bayes net over N random variables, what is the largest possible number of orderings in which I could sample?

$n!$

Consider a Bayes net with no edges. Any ordering of the nodes is a valid sampling ordering.

- (ii) [1 pt] What is the smallest possible number of orderings in which I could sample?

1

Consider a long causal chain. Since every node must be sampled before its children, there is only one valid sampling ordering.

- (d) Recall that rejection sampling is most efficient when we reject as *early* as possible. In general, it might be hard to determine which sample ordering will make this possible. We'd like to formulate the problem as an MDP, and use policy iteration to select an optimal ordering.

- (i) [1 pt] The state space of this MDP will either be some collection of random variables, or (variable, value) pairs. More specifically, which of the following is an appropriate minimal state representation for this MDP? (Mark one.) Hint: it may be helpful to refer to the transition function described below.

- Set of variables that have been sampled so far (e.g. $\{A, B, D, \dots\}$).
- Set of (variable, value) pairs that have been sampled so far (e.g. $\{(A, +a), (B, -b), (D, +d), \dots\}$).
- Ordered list of variables pairs that have been sampled so far (e.g. $\{[A, B, D, \dots]\}$).
- Ordered list of (variable, value) pairs that have sampled so far (e.g. $\{[(A, +a), (B, -b), (D, +d), \dots]\}$).

We need to keep track of values because we can't sample at a node if we don't know the values we've already chosen for its parents. But we can use a set, rather than a list, because the order in which we sampled everything in the past doesn't affect our ability to sample in the future.

- (ii) [1 pt] If the Bayes net has N binary random variables, how big is this state space? (Choose the tightest upper bound out of the answers given.)

- $O(n)$
- $O(2^n)$
- $O(3^n)$
- $O(n!)$
- $O(2^n n!)$
- $O(3^n n!)$

The answer to this question depends on the answer to the previous question. A set of variables takes $O(2^n)$ space (an indicator on every variable for whether it's in the set). A set of (variable, value) pairs takes $O(3^n)$ space (every variable is either assigned a positive value, a negative value, or not). The list-valued answers involve a term that is $O(n(n+1)!)$ (there are $n!$ states with n variables sampled, $(n-1)!$ states with $n-1$ variables sampled, etc.). $O((n+1)!)$ is strictly greater than $O(n!)$, and $O(2^n n!)$ is the tightest upper bound.

The action space and transition function of this MDP are as follows: Every random variable corresponds to an action. When we select a random variable, we sample a value from the corresponding distribution. If this value causes the sampler to reject, we immediately transition to a terminal "sink state". Otherwise, we add the variable (or (variable, value) pair) to the collection chosen above.

- (iii) [2 pts] If $\gamma = 0.5$, which of the following is an appropriate reward? Recall that we want to reward the sampler for rejecting as quickly as possible. (Mark all that apply.)

- 1 per turn, and 0 if the sample is rejected
- 1 per turn, and 0 if the sample is rejected
- 0 per turn, and 1 if the sample is rejected

0 per turn, and -1 if the sample is rejected

We want to reward the agent for rejecting the sample as quickly as possible. We can either do this by punishing it for remaining alive (corresponding to a reward of -1 per turn), or rewarding for entering the sink state (corresponding to a single final reward of 1). The second of these works only with a discount factor less than 1, which means a reward in the near future is better than a reward in the far future.

(iv) [2 pts] If $\gamma = 1.0$, which of the following is an appropriate reward? (Mark all that apply.)

-1 per turn, and 0 if the sample is rejected

1 per turn, and 0 if the sample is rejected

0 per turn, and 1 if the sample is rejected

0 per turn, and -1 if the sample is rejected

As in the previous question. But because $\gamma = 1$, the agent is indifferent between rejecting a sample now and rejecting a sample arbitrarily far in the future, which means this is not a suitable reward.

Q7. [19 pts] HMMs

Consider a process where there are transitions among a finite set of states s_1, \dots, s_k over time steps $i = 1, \dots, N$. Let the random variables X_1, \dots, X_N represent the state of the system at each time step and be generated as follows:

- Sample the initial state s from an initial distribution $P_1(X_1)$, and set $i = 1$
- Repeat the following:
 1. Sample a duration d from a duration distribution P_D over the integers $\{1, \dots, M\}$, where M is the maximum duration.
 2. Remain in the current state s for the next d time steps, i.e., set

$$x_i = x_{i+1} = \dots = x_{i+d-1} = s \tag{1}$$
 3. Sample a successor state s' from a transition distribution $P_T(X_t|X_{t-1} = s)$ over the other states $s' \neq s$ (so there are no self transitions)
 4. Assign $i = i + d$ and $s = s'$.


This process continues indefinitely, but we only observe the first N time steps.

- (a) [2 pts] Assuming that all three states s_1, s_2, s_3 are different, what is the probability of the sample sequence $s_1, s_1, s_2, s_2, s_2, s_3, s_3$? Write an algebraic expression. Assume $M \geq 3$.

$$p_1(s_1)p_D(2)p_T(s_2|s_1)p_D(3)p(s_3|s_2)(1 - p_D(1)) \tag{2}$$

At each time step i we observe a noisy version of the state X_i that we denote Y_i and is produced via a conditional distribution $P_E(Y_i|X_i)$.

- (b) [1 pt] Only in this subquestion assume that $N > M$. Let X_1, \dots, X_N and Y_1, \dots, Y_N random variables defined as above. What is the maximum index $i \leq N - 1$ so that $X_1 \perp\!\!\!\perp X_N | X_i, X_{i+1}, \dots, X_{N-1}$ is guaranteed?
 $i = N - M$
- (c) [3 pts] Only in this subquestion, assume the max duration $M = 2$, and P_D uniform over $\{1, 2\}$ and each x_i is in an alphabet $\{a, b\}$. For $(X_1, X_2, X_3, X_4, X_5, Y_1, Y_2, Y_3, Y_4, Y_5)$ draw a Bayes Net over these 10 random variables with the property that removing any of the edges would yield a Bayes net inconsistent with the given distribution.



(d) In this part we will explore how to write the described process as an HMM with an extended state space. Write the states $z = (s, t)$ where s is a state of the original system and t represents the time elapsed in that state. For example, the state sequence $s_1, s_1, s_1, s_2, s_3, s_3$ would be represented as $(s_1, 1), (s_1, 2), (s_1, 3), (s_2, 1), (s_3, 1), (s_3, 2)$. Answer all of the following in terms of the parameters $P_1(X_1), P_D(d), P_T(X_{j+1}|X_j), P_E(Y_i|X_i), k$ (total number of possible states), N and M (max duration).

(i) [1 pt] What is $P(Z_1)$?

$$P(x_1, t) = \begin{cases} P_1(x_1) & \text{if } t = 1 \\ 0 & \text{o.w.} \end{cases} \quad (3)$$

(ii) [3 pts] What is $P(Z_{i+1}|Z_i)$? Hint: You will need to break this into cases where the transition function will behave differently.

$$P(X_{i+1}, t_{i+1}|X_i, t_i) = \begin{cases} P_D(d \geq t_i + 1 | d \geq t_i) & \text{when } X_{i+1} = X_i \text{ and } t_{i+1} = t_i + 1 \text{ and } t_{i+1} \leq M \\ P_T(X_{i+1}|X_i)P_D(d = t_i | d \geq t_i) & \text{when } X_{i+1} \neq X_i \text{ and } t_{i+1} = 1 \\ 0 & \text{o.w.} \end{cases}$$

Where $P_D(d \geq t_i + 1 | d \geq t_i) = P_D(d \geq t_i + 1) / P_D(d \geq t_i)$.

Being in X_i, t_i , we know that d was drawn $d \geq t_i$. Conditioning on this fact, we have two choices, if $d > t_i$ then the next state is $X_{i+1} = X_i$, and if $d = t_i$ then $X_{i+1} \neq X_i$ drawn from the transition distribution and $t_{i+1} = 1$.

(4)

(iii) [1 pt] What is $P(Y_i|Z_i)$?

$$p(Y_i|X_i, t_i) = P_E(Y_i|X_i)$$

- (e) In this question we explore how to write an algorithm to compute $P(X_N|y_1, \dots, y_N)$ using the particular structure of this process.

Write $P(X_t|y_1, \dots, y_{t-1})$ in terms of other factors. Construct an answer by checking the correct boxes below:

$$P(X_t|y_1, \dots, y_{t-1}) = \underline{\hspace{1cm} \text{(i)} \hspace{1cm}} \quad \underline{\hspace{1cm} \text{(ii)} \hspace{1cm}} \quad \underline{\hspace{1cm} \text{(iii)} \hspace{1cm}}$$

(i) [1 pt]

- $\sum_{i=1}^k \sum_{d=1}^M \sum_{d'=1}^M$ $\sum_{i=1}^k$
 $\sum_{i=1}^k \sum_{d=1}^M$ $\sum_{d=1}^M$

(ii) [1 pt]

- $P(Z_t = (X_t, d)|Z_{t-1} = (s_i, d))$ $P(X_t|X_{t-1} = s_d)$
 $P(X_t|X_{t-1} = s_i)$ $P(Z_t = (X_t, d')|Z_{t-1} = (s_i, d))$

(iii) [1 pt]

- $P(Z_{t-1} = (s_d, i)|y_1, \dots, y_{t-1})$ $P(Z_{t-1} = (s_i, d)|y_1, \dots, y_{t-1})$
 $P(X_{t-1} = s_d|y_1, \dots, y_{t-1})$ $P(X_{t-1} = s_i|y_1, \dots, y_{t-1})$

- (iv) [1 pt] Now we would like to include the evidence y_t in the picture. What would be the running time of each update of the **whole table** $P(X_t|y_1, \dots, y_t)$? Assume tables corresponding to any factors used in (i), (ii), (iii) have already been computed.

- $O(k^2)$ $O(k^2M^2)$
 $O(k^2M)$ $O(kM)$

Note: Computing $P(X_N|y_1, \dots, y_N)$ will take time $N \times$ your answer in (iv).

Just the running time for filtering when the state space is the space of pairs (x_i, t_i) ,

Given $B_{t-1}(z)$, the step $p(z_t|y_1, \dots, y_{t-1})$ can be done in time kM . (size of the statespace for z).

The computation to include the y_t evidence can be done in $O(1)$ per z_t .

Therefore each update to the table per evidence point will take $(Mk)^2$. So it is $O((Mk)^2)$.

Using N steps, the whole algorithm will take $O(Nk^2M^2)$ to compute $P(X_N|Y_1, \dots, Y_N)$.

- (v) [4 pts] Describe an update rule to compute $P(X_t|y_1, \dots, y_{t-1})$ that is faster than the one you discovered in parts (i), (ii), (iii). **Specify its running time.** Hint: Use the structure of the transitions $Z_{t-1} \rightarrow Z_t$.

Answer is $O(k^2M + kM)$.

The answer from the previous section is:

$$P(X_t|y_1, \dots, y_{t-1}) = \sum_{i=1}^k \sum_{d=1}^M \sum_{d'=1}^M P(Z_t = (X_t, d')|Z_{t-1} = (s_i, d))P(Z_{t-1} = (s_i, d)|y_1, \dots, y_{t-1}) \quad (5)$$

To compute this value we only really need to loop through those transitions $P(Z_t = (X_t, d')|Z_{t-1} = (s_i, d))$ that can happen with nonzero probability.

For all $X_t = s$ we need to sum over all factors of the form $P(Z_t = (s, d')|Z_{t-1} = (s_i, d))P(X_{t-1} = s_i|y_1, \dots, y_{t-1})$. For a fixed s the factor $P(Z_t = (X_t, d')|Z_{t-1} = (s_i, d))$ can be nonzero only when $s_i = s$ and $d' = d + 1$ (M tuples). And when $s_i \neq s$ and $d' = 1$ and $d = 1, \dots, M$ (kM tuples).

Since this needs to be performed for all k possible values of s , the answer to update the whole table is $O(k^2M + kM)$.

THIS PAGE IS INTENTIONALLY LEFT BLANK