

To earn the extra credit, one of the following has to hold true. Please circle and sign.

A I spent 2 or more hours on the practice midterm.

B I spent fewer than 2 hours on the practice midterm, but I believe I have solved all the questions.

Signature: _____

To simulate midterm setting, print out this practice midterm, complete it in writing, and then scan and upload into Gradescope. It is due on Tuesday 11/13, 11:59pm.

Exam Instructions:

- You have approximately 2 hours.
- The exam is closed book, closed notes except your one-page cheat sheet.
- Please use non-programmable calculators only.
- 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.

First name	
Last name	
SID	
First and last name of student to your left	
First and last name of student to your right	


For staff use only:

Q1. Probability and Decision Networks	/13
Q2. MDPs and Utility: Short Questions	/23
Q3. Machine Learning	/7
Q4. Bayes' Net Reasoning	/12
Q5. D-Separation	/8
Q6. Variable Elimination	/19
Q7. Bayes' Nets Sampling	/10
Q8. Modified HMM Updates	/8
Q9. Learning a Bayes' Net Structure	/9
Total	/109

THIS PAGE IS INTENTIONALLY LEFT BLANK

Q1. [13 pts] Probability and Decision Networks

A	P(A)
e	0.5
l	0.5



S	P(S)
e	
l	

A	S	P(S A)
e	e	0.8
e	l	0.2
l	e	0.4
l	l	0.6

A	T	U(A,T)
e	e	600
e	l	0
l	e	300
l	l	600

S	A	P(A S)
e	e	
e	l	
l	e	
l	l	

Your parents are visiting you for graduation. You are in charge of picking them up at the airport. Their arrival time (A) might be early (e) or late (l). You decide on a time (T) to go to the airport, also either early (e) or late (l). Your sister (S) is a noisy source of information about their arrival time. The probability values and utilities are shown in the tables above.

Compute $P(S)$, $P(A|S)$ and compute the quantities below.

$$EU(T = e) =$$

$$EU(T = l) =$$

$$MEU(\{\}) =$$

Optimal action with no observations is $T =$

Now we consider the case where you decide to ask your sister for input.

$$EU(T = e|S = e) =$$

$$EU(T = l|S = e) =$$

$$MEU(\{S = e\}) =$$

Optimal action with observation $\{S = e\}$ is $T =$

$$EU(T = e|S = l) =$$

$$EU(T = l|S = l) =$$

$$MEU(\{S = l\}) =$$

Optimal action with observation $S = l$ is $T =$

$$VPI(S) =$$

Q2. [23 pts] MDPs and Utility: Short Questions

Each True/False question is worth 2 points. Leaving a question blank is worth 0 points. **Answering incorrectly is worth -2 points.**

For the questions that are not True/False, answer as concisely as possible (and no points are subtracted for a wrong answer to these).

(a) Utility.

- (i) [2 pts] [true or false] If an agent has the preference relationship $(A \succ B) \wedge (B \succ C) \wedge (C \succ A)$ then this agent can be induced to give away all of its money.
- (ii) [2 pts] [true or false] Assume Agent 1 has a utility function U_1 and Agent 2 has a utility function U_2 . If $U_1 = k_1 U_2 + k_2$ with $k_1 > 0, k_2 > 0$ then Agent 1 and Agent 2 have the same preferences.

(b) Insurance. Some useful numbers: $\log(101) \approx 4.6151, \log(71) \approx 4.2627$.

PacBaby just found a \$100 bill—it is the only thing she owns. Ghosts are nice enough not to kill PacBaby, but when they find PacBaby they will steal all her money. The probability of the ghosts finding PacBaby is 20%. PacBaby's utility function is $U(x) = \log(1 + x)$ (this is the natural logarithm, i.e., $\log e^x = x$), where x is the total monetary value she owns. When PacBaby gets to keep the \$100 (ghosts don't find her) her utility is $U(\$100) = \log(101)$. When PacBaby loses the \$100 (per the ghosts taking it from her) her utility is $U(\$0) = \log(1 + 0) = 0$.

- (i) [2 pts] What is the expected utility for PacBaby?

- (ii) [4 pts] Pacgressive offers theft insurance: if PacBaby pays an insurance premium of \$30, then they will reimburse PacBaby \$70 if the ghosts steal all her money (after paying \$30 in insurance, she would only have \$70 left). What is the expected utility for PacBaby if she takes insurance? For PacBaby to maximize her expected utility should she take this insurance?

- (iii) [2 pts] In the above scenario, what is the expected *monetary* value of selling the insurance from Pacgressive's point of view?

(c) MDPs.

- (i) [2 pts] [true or false] If the only difference between two MDPs is the value of the discount factor then they must have the same optimal policy.
 - (ii) [2 pts] [true or false] When using features to represent the Q-function (rather than having a tabular representation) it is possible that Q-learning does not find the optimal Q-function Q^* .
 - (iii) [2 pts] [true or false] For an infinite horizon MDP with a finite number of states and actions and with a discount factor γ , with $0 < \gamma < 1$, value iteration is guaranteed to converge.
- (d) [5 pts] Recall that for a deterministic policy π where $\pi(s)$ is the action to be taken in state s we have that the value of the policy satisfies the following equations:

$$V^\pi(s) = \sum_{s'} T(s, \pi(s), s') (R(s, \pi(s), s') + \gamma V^\pi(s')).$$

Now assume we have a stochastic policy π where $\pi(s, a) = P(a|s)$ is equal to the probability of taking action a when in state s . Write the equivalent of the above equation for the value of this stochastic policy.

$$V^\pi(s) =$$

Q3. [7 pts] Machine Learning

(a) **Maximum Likelihood**

(i) [4 pts] **Geometric Distribution**

Consider the geometric distribution, which has $P(X = k) = (1 - \theta)^{k-1}\theta$. Assume in our training data X took on the values 4, 2, 7, and 9.

(a) Write an expression for the log-likelihood of the data as a function of the parameter θ .

(b) What is the value of θ that maximizes the log-likelihood, i.e., what is the maximum likelihood estimate for θ ?


(ii) [3 pts] Consider the Bayes' net consisting of just two variables A, B , and structure $A \rightarrow B$. Find the maximum likelihood estimates and the $k = 2$ Laplace estimates for each of the table entries based on the following data: $(+a, -b), (+a, +b), (+a, -b), (-a, -b), (-a, -b)$.

A	$P^{\text{ML}}(A)$	$P^{\text{Laplace, } k=2}(A)$
$+a$		
$-a$		

A	B	$P^{\text{ML}}(\mathbf{B} \mid \mathbf{A})$	$P^{\text{Laplace, } k=2}(\mathbf{B} \mid \mathbf{A})$
$+a$	$+b$		
$+a$	$-b$		
$-a$	$+b$		
$-a$	$-b$		

Q4. [12 pts] Bayes' Net Reasoning

$P(A D, X)$			
+d	+x	+a	0.9
+d	+x	-a	0.1
+d	-x	+a	0.8
+d	-x	-a	0.2
-d	+x	+a	0.6
-d	+x	-a	0.4
-d	-x	+a	0.1
-d	-x	-a	0.9



$P(D)$	
+d	0.1
-d	0.9

$P(X D)$		
+d	+x	0.7
+d	-x	0.3
-d	+x	0.8
-d	-x	0.2

$P(B D)$		
+d	+b	0.7
+d	-b	0.3
-d	+b	0.5
-d	-b	0.5

(a) [3 pts] What is the probability of having disease D and getting a positive result on test A ?
 $P(+d, +a) =$


(b) [3 pts] What is the probability of *not* having disease D and getting a positive result on test A ?
 $P(-d, +a) =$

(c) [3 pts] What is the probability of having disease D given a positive result on test A ?
 $P(+d | +a) =$

(d) [3 pts] What is the probability of having disease D given a positive result on test B ?
 $P(+d | +b) =$

Q5. [8 pts] D-Separation

- (a) [8 pts] Based only on the structure of the (new) Bayes' Net given below, circle whether the following conditional independence assertions are guaranteed to be true, guaranteed to be false, or cannot be determined by the structure alone.



$U \perp\!\!\!\perp V$ Guaranteed true Cannot be determined Guaranteed false

$U \perp\!\!\!\perp V \mid W$ Guaranteed true Cannot be determined Guaranteed false

$U \perp\!\!\!\perp V \mid Y$ Guaranteed true Cannot be determined Guaranteed false

$U \perp\!\!\!\perp Z \mid W$ Guaranteed true Cannot be determined Guaranteed false

$U \perp\!\!\!\perp Z \mid V, Y$ Guaranteed true Cannot be determined Guaranteed false


$U \perp\!\!\!\perp Z \mid X, W$ Guaranteed true Cannot be determined Guaranteed false

$W \perp\!\!\!\perp X \mid Z$ Guaranteed true Cannot be determined Guaranteed false

$V \perp\!\!\!\perp Z \mid X$ Guaranteed true Cannot be determined Guaranteed false

Q6. [19 pts] Variable Elimination

- (a) [10 pts] For the Bayes' net below, we are given the query $P(Z \mid +y)$. All variables have binary domains. Assume we run variable elimination to compute the answer to this query, with the following variable elimination ordering: U, V, W, T, X .



Complete the following description of the factors generated in this process:

After inserting evidence, we have the following factors to start out with:

$$P(U), P(V), P(W|U, V), P(X|V), P(T|V), P(+y|W, X), P(Z|T).$$

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

$$f_1(W|V) = \sum_u P(u)P(W|u, V).$$

This leaves us with the factors:

$$P(V), P(X|V), P(T|V), P(+y|W, X), P(Z|T), f_1(W|V).$$

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

This leaves us with the factors:

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

This leaves us with the factors:

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

This leaves us with the factor:

When eliminating X we generate a new factor f_5 as follows:

This leaves us with the factor:


- (b) [2 pts] Briefly explain how $P(Z | +y)$ can be computed from f_5 .
- (c) [2 pts] Amongst f_1, f_2, \dots, f_5 , which is the largest factor generated? (Assume all variables have binary domains.) How large is this factor?
- (d) [5 pts] Find a variable elimination ordering for the same query, i.e., for $P(Z | +y)$, for which the maximum size factor generated along the way is smallest. Hint: the maximum size factor generated in your solution should have only 2 variables, for a size of $2^2 = 4$ table. Fill in the variable elimination ordering and the factors generated into the table below.

Variable Eliminated	Factor Generated

Note: in the naive ordering we used earlier, the first line in this table would have had the following two entries: U , $f_1(W|V)$.

Q7. [10 pts] Bayes' Nets Sampling

Assume the following Bayes net, and the corresponding distributions over the variables in the Bayes net:



A	P(A)
+a	1/5
-a	4/5

B	C	P(C B)
+b	+c	1/4
+b	-c	3/4
-b	+c	2/5
-b	-c	3/5

A	B	P(B A)
+a	+b	1/5
+a	-b	4/5
-a	+b	1/2
-a	-b	1/2

B	D	P(D B)
+b	+d	1/2
+b	-d	1/2
-b	+d	4/5
-b	-d	1/5

- (a) [2 pts] Your task is now to estimate $P(+b | -a, -c, -d)$ using rejection sampling. Below are some samples that have been produced by prior sampling (that is, the rejection stage in rejection sampling hasn't happened yet). Cross out the samples that would be rejected by rejection sampling:

-a	-b	+c	+d	-a	-b	-c	-d
+a	-b	-c	+d	-a	+b	+c	+d
-a	-b	+c	-d	+a	-b	-c	-d

- (b) [1 pt] Using those samples, what value would you estimate for $P(+b | -a, -c, -d)$ using rejection sampling?

- (c) [3 pts] Using the following samples (which were generated using likelihood weighting), estimate $P(+b | -a, -c, -d)$ using likelihood weighting, or state why it cannot be computed.

-a	-b	-c	-d
-a	+b	-c	-d
-a	-b	-c	-d

- (d) (i) [2 pts] Consider the query $P(A | -b, -c)$. After rejection sampling we end up with the following four samples: $(+a, -b, -c, +d)$, $(+a, -b, -c, -d)$, $(+a, -b, -c, -d)$, $(-a, -b, -c, -d)$. What is the resulting estimate of $P(+a | -b, -c)$?
- (ii) [2 pts] Consider again the query $P(A | -b, -c)$. After likelihood weighting sampling we end up with the following four samples: $(+a, -b, -c, -d)$, $(+a, -b, -c, -d)$, $(-a, -b, -c, -d)$, $(-a, -b, -c, +d)$, and respective weights: 0.1, 0.1, 0.3, 0.3. What is the resulting estimate of $P(+a | -b, -c)$?


Q8. [8 pts] Modified HMM Updates

(a) Recall that for a standard HMM the Elapse Time update and the Observation update are of the respective forms:

$$P(X_t | e_{1:t-1}) = \sum_{x_{t-1}} P(X_t | x_{t-1})P(x_{t-1} | e_{1:t-1})$$

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

We now consider the following two HMM-like models:



Mark the modified Elapse Time update and the modified Observation update that correctly compute the beliefs from the quantities that are available in the Bayes' Net. (Mark one of the first set of six options, and mark one of the second set of six options for (i), and same for (ii).)

(i) [4 pts]

- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}, z_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t | x_{t-1}, z_{t-1})P(Z_t)$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}, z_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t | x_{t-1}, z_{t-1})$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}, z_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t, Z_t | x_{t-1}, z_{t-1})$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t | x_{t-1}, z_{t-1})P(Z_t)$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t | x_{t-1}, z_{t-1})$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t, Z_t | x_{t-1}, z_{t-1})$

- $P(X_t, Z_t | e_{1:t}) \propto P(X_t, Z_t | e_{1:t-1})P(e_t | X_t, Z_t)$
- $P(X_t, Z_t | e_{1:t}) \propto \sum_{X_t} P(X_t, Z_t | e_{1:t-1})P(e_t | X_t, Z_t)$
- $P(X_t, Z_t | e_{1:t}) \propto \sum_{Z_t} P(X_t, Z_t | e_{1:t-1})P(e_t | X_t, Z_t)$
- $P(X_t, Z_t | e_{1:t}) \propto P(X_t, Z_t | e_{1:t-1})P(e_t | X_t)P(e_t | Z_t)$
- $P(X_t, Z_t | e_{1:t}) \propto P(X_t, Z_t | e_{1:t-1})P(e_t | X_t)$
- $P(X_t, Z_t | e_{1:t}) \propto P(X_t, Z_t | e_{1:t-1}) \sum_{X_t} P(e_t | X_t)$

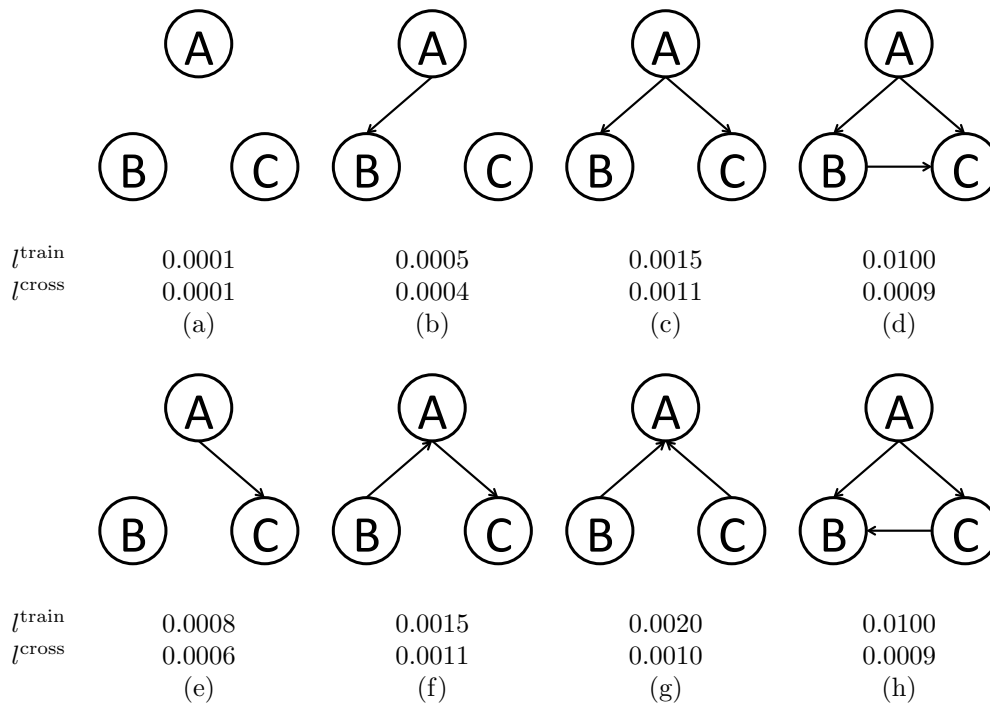
(ii) [4 pts]

- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}, z_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t | x_{t-1}, z_{t-1})P(Z_t | e_{t-1})$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}, z_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(Z_t | e_{t-1})P(X_t | x_{t-1}, Z_t)$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}, z_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t, Z_t | x_{t-1}, e_{t-1})$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t | x_{t-1}, z_{t-1})P(Z_t | e_{t-1})$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(Z_t | e_{t-1})P(X_t | x_{t-1}, Z_t)$
- $P(X_t, Z_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1}, z_{t-1} | e_{1:t-1})P(X_t, Z_t | x_{t-1}, e_{t-1})$

- $P(X_t, Z_t | e_{1:t}) \propto P(X_t, Z_t | e_{1:t-1})P(e_t | X_t, Z_t)$
- $P(X_t, Z_t | e_{1:t}) \propto \sum_{X_t} P(X_t, Z_t | e_{1:t-1})P(e_t | X_t, Z_t)$
- $P(X_t, Z_t | e_{1:t}) \propto \sum_{Z_t} P(X_t, Z_t | e_{1:t-1})P(e_t | X_t, Z_t)$
- $P(X_t, Z_t | e_{1:t}) \propto P(X_t, Z_t | e_{1:t-1})P(e_t | X_t)P(e_t | Z_t)$
- $P(X_t, Z_t | e_{1:t}) \propto P(X_t, Z_t | e_{1:t-1})P(e_t | X_t)$
- $P(X_t, Z_t | e_{1:t}) \propto P(X_t, Z_t | e_{1:t-1}) \sum_{X_t} P(e_t | X_t)$

Q9. [9 pts] Learning a Bayes' Net Structure

You want to learn a Bayes' net over the random variables A, B, C . You decide you want to learn not only the Bayes' net parameters, but also the structure from the data. You are willing to consider the 8 structures shown below. First you use your training data to perform maximum likelihood estimation of the parameters of each of the Bayes' nets. Then for each of the learned Bayes' nets, you evaluate the likelihood of the training data (l^{train}), and the likelihood of your cross-validation data (l^{cross}). Both likelihoods are shown below each structure.



- (a) [3 pts] Which Bayes' net structure will (on expectation) perform best on test-data? (If there is a tie, list all Bayes' nets that are tied for the top spot.) Justify your answer.
- (b) [3 pts] Two pairs of the learned Bayes' nets have identical likelihoods. Explain why this is the case.
- (c) [3 pts] For every two structures S_1 and S_2 , where S_2 can be obtained from S_1 by adding one or more edges, l^{train} is higher for S_2 than for S_1 . Explain why this is the case.

THIS PAGE IS INTENTIONALLY LEFT BLANK