CS 188 Spring 2026 — Introduction to Artificial Intelligence — **Exam Prep 8 Solutions**

Q1. 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.

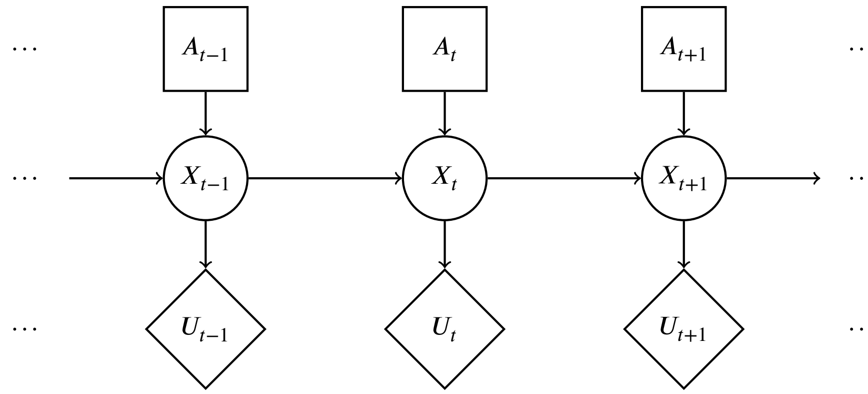

(a) The belief at time t is defined as $B_t(X_t) = p(X_t|a_{1:t})$. The forward algorithm update has the following form:

$$B_t(X_t) = \underline{\quad \textbf{(i)} \quad}\ \underline{\quad \textbf{(ii)} \quad}\ B_{t-1}(x_{t-1}).$$

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

(i)	○ $\max_{x_{t-1}}$	● $\sum_{x_{t-1}}$	○ $\max_{x_t}$	○ $\sum_{x_t}$	○ 1
(ii)	○ $p(X_t\|x_{t-1})$	○ $p(X_t\|x_{t-1})p(X_t\|a_t)$	○ $p(X_t)$	● $p(X_t\|x_{t-1}, a_t)$	○ 1

○ None of the above combinations is correct

$$\begin{aligned} B_t(X_t) &= p(X_t|a_{1:t}) \\ &= \sum_{x_{t-1}} p(X_t|x_{t-1}, a_t)p(x_{t-1}|a_{1:t-1}) \\ &= \sum_{x_{t-1}} p(X_t|x_{t-1}, a_t)B_{t-1}(x_{t-1}) \end{aligned}$$

(b) Pacman would like to take actions $A_{1:T}$ that maximizes the expected sum of utilities, which has the following form:

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

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

(i)	● $\max_{a_{1:T}}$	○ $\max_{a_T}$	○ $\sum_{a_{1:T}}$	○ $\sum_{a_T}$	○ 1
(ii)	○ $\max_t$	○ $\prod_{t=1}^{T}$	● $\sum_{t=1}^{T}$	○ $\min_t$	○ 1
(iii)	○ $\sum_{x_t, a_t}$	● $\sum_{x_t}$	○ $\sum_{a_t}$	○ $\sum_{x_T}$	○ 1
(iv)	○ $p(x_t\|x_{t-1}, a_t)$	○ $p(x_t)$	● $B_t(x_t)$	○ $B_T(x_T)$	○ 1
(v)	○ U_T	○ $\frac{1}{U_t}$	○ $\frac{1}{U_T}$	● U_t	○ 1

○ None of the above combinations is correct

$$\text{MEU}_{1:T} = \max_{a_{1:T}} \sum_{t=1}^{T} \sum_{x_t} B_t(x_t) U_t(x_t)$$

(c) 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 $p(x_t|x_{t-1}, a_t)$ is not deterministic. **With respect to the utility** U_t, mark all that can be True:

- [x] $\text{VPI}(X_{t-1}|X_{t-2}) > 0$
- [ ] $\text{VPI}(X_{t-2}|X_{t-1}) > 0$
- [x] $\text{VPI}(X_{t-1}|X_{t-2}) = 0$
- [x] $\text{VPI}(X_{t-2}|X_{t-1}) = 0$
- [ ] None of the above

It is always possible that VPI = 0. Can guarantee $\text{VPI}(E|e)$ is not greater than 0 if E is independent of parents(U) given e.

(d) 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 $B_T(X_T)$ is? If different methods give an equivalently accurate estimate, mark them as the same number.

	Most accurate			Least accurate
Exact inference	● 1	○ 2	○ 3	○ 4
Particle filtering with no resampling	○ 1	● 2	○ 3	○ 4
Particle filtering with resampling before every time elapse	○ 1	○ 2	○ 3	● 4
Particle filtering with resampling before every other time elapse	○ 1	○ 2	● 3	○ 4

Exact inference will always be more accurate than using a particle filter. When comparing the particle filter resampling approaches, notice that because there are no observations, each particle will have weight 1. Therefore resampling when particle weights are 1 could lead to particles being lost and hence prove bad.

Q2. Sampling in Bayes Nets

We would like to analyze people's ice cream eating habits on sunny and rainy days. Suppose we consider the weather, along with a person's ice-cream eating, over the span of two days. We'll have four random variables: W_1 and W_2 stand for the weather on days 1 and 2, which can either be rainy R or sunny S, and the variables I_1 and I_2 represent whether or not the person ate ice cream on days 1 and 2, and take values T (for truly eating ice cream) or F. We can model this as the following Bayes Net with these probabilities.

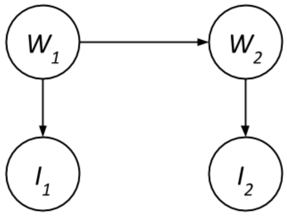

W_1	$P(W_1)$
S	0.6
R	0.4

W_1	W_2	$P(W_2\|W_1)$
S	S	0.7
S	R	0.3
R	S	0.5
R	R	0.5

W	I	$P(I\|W)$
S	T	0.9
S	F	0.1
R	T	0.2
R	F	0.8

Suppose we produce the following samples of (W_1, I_1, W_2, I_2) from the ice-cream model:

~~R, F, R, F~~ ~~R, F, R, F~~ ~~S, F, S, T~~ ~~S, T, S, T~~ S, T, R, F
~~R, F, R, T~~ ~~S, T, S, T~~ ~~S, T, S, T~~ S, T, R, F ~~R, F, S, T~~

(a) Using these samples, what is our estimate of $P(W_2 = \mathtt{R})$? $5/10 = 0.5$

(b) Cross off samples above which are rejected by rejection sampling if we're trying to estimate $P(W_2|I_1 = \mathtt{T}, I_2 = \mathtt{F})$

Rejection sampling seems to be wasting a lot of effort, so we decide to switch to likelihood weighting. Assume we generate the following six samples for (W_1, I_1, W_2, I_2), given the evidence $I_1 = \mathtt{T}$ and $I_2 = \mathtt{F}$:

(S, T, R, F) (R, T, R, F) (S, T, R, F) (S, T, S, F) (S, T, S, F) (R, T, S, F)

(c) Calculate the weight of each sample.

In this case, the evidence is $I_1 = \mathtt{T}, I_2 = \mathtt{F}$.
The weight of the first sample is thus $w = \Pr(I_1 = \mathtt{T}|W_1 = \mathtt{S}) \cdot \Pr(I_2 = \mathtt{F}|W_2 = \mathtt{R}) = 0.9 \cdot 0.8 = 0.72$
Similarly, we can calculate the weights of the other samples as follows:

(W_1, I_1, W_2, I_2)	w
S, T, R, F	0.72
R, T, R, F	0.16
S, T, R, F	0.72
S, T, S, F	0.09
S, T, S, F	0.09
R, T, S, F	0.02

(d) Estimate $P(W_2|I_1 = \mathtt{T}, I_2 = \mathtt{F})$ using our likelihood weights from the previous part.
To compute the probability estimate, we normalize the weights and find

$$\hat{P}(W_2 = \mathtt{R}|I_1 = \mathtt{T}, I_2 = \mathtt{F}) = \frac{0.72 + 0.16 + 0.72}{0.72 + 0.16 + 0.72 + 0.09 + 0.09 + 0.02} = 0.889$$
$$\hat{P}(W_2 = \mathtt{S}|I_1 = \mathtt{T}, I_2 = \mathtt{F}) = 1 - 0.889 = 0.111.$$

Q3: Pedestrian Tracking & Variable Elimination

An autonomous vehicle is tracking a pedestrian. To predict their movement, the vehicle's AI uses a DBN with three variables per time step t:

- I_t: The pedestrian's unobserved intention (e.g., cross street, stay on sidewalk). Domain size $|I| = 2$.
- P_t: The pedestrian's true physical position. Domain size $|P| = 50$.
- C_t: The vehicle's camera observation. Domain size $|C| = 10$.

The network structure is defined by the following edges for all t:

- $I_{t-1} \rightarrow I_t$
- $P_{t-1} \rightarrow P_t$
- $I_t \rightarrow P_t$
- $P_t \rightarrow C_t$

1. **Transition Model:** Write the factored representation of the transition model $P(I_t, P_t \mid I_{t-1}, P_{t-1})$ exploiting the conditional independencies of this specific DBN.

   Looking at the edges coming into the time t nodes from time $t-1$ and within time t:
   $P(I_t, P_t \mid I_{t-1}, P_{t-1}) = P(I_t \mid I_{t-1})P(P_t \mid P_{t-1}, I_t)$

2. **Markov Assumptions:** Using the fundamental conditional independencies of this Bayesian Network, simplify the following expression as much as possible: $P(P_3 \mid P_1, P_2, I_1, I_2, I_3)$. Briefly justify your answer.

   Answer: $P(P_3 \mid P_2, I_3)$
   By the local semantics of Bayesian Networks, a node is conditionally independent of all its non-descendants given its parents. The direct parents of P_3 are P_2 and I_3. Since P_1, I_1, and I_2 are non-descendants of P_3, they can be dropped from the conditioning set once the parents are given.

3. **Exact Inference (Variable Elimination):** We want to find the exact distribution of the pedestrian's intention at time 2, $P(I_2 \mid C_1, C_2)$. To do this, we run Variable Elimination on the unrolled network for $t = 1, 2$.

   The initial factors are:

   $$P(I_1), P(P_1 \mid I_1), P(C_1 \mid P_1), P(I_2 \mid I_1), P(P_2 \mid P_1, I_2), P(C_2 \mid P_2)$$

   Assume we eliminate the hidden variables in the following order: I_1, P_1, P_2.

   When we eliminate P_1, we must join several factors and then sum out P_1 to create a new factor f'. **What is the total number of entries (size) in this newly generated factor f'?** Show your work.

   Size: 100
   First, look at the state of the factors right before eliminating P_1. We have already eliminated I_1, which generated a new factor $f_1(P_1, I_2)$. (Because we joined $P(I_1)$, $P(P_1 \mid I_1)$, and $P(I_2 \mid I_1)$ and summed out I_1).

   Crucially, because we are querying $P(I_2 \mid C_1, C_2)$, the variables C_1 and C_2 are observed evidence. This means they are instantiated to specific observed values (e.g., $C_1 = c_1$) before Variable Elimination begins, removing them as free variables in the factor domains.

Now, to eliminate P_1, we collect all remaining factors that contain P_1:

- $f_1(P_1, I_2)$ (from eliminating I_1)
- $P(c_1 \mid P_1)$ (instantiated evidence)
- $P(P_2 \mid P_1, I_2)$

We join these to form an intermediate factor $f_{join}(P_1, I_2, P_2)$.
We then sum out P_1 to create the new factor $f'(I_2, P_2)$.
The size of f' is strictly the product of the domain sizes of its uninstantiated variables:
Size $= |I| \times |P| = 2 \times 50 = 100$.