

Q1. Partying Particle #No Filter(ing)

Algorithm 1.1 Particle Filtering

```

1: procedure PARTICLE FILTERING( $T, N$ )           ▷  $T$ : number of time steps,  $N$ : number of sampled particles
2:    $x \leftarrow$  sample  $N$  particles from initial state distribution  $P(X_0)$            ▷ Initialize
3:   for  $t \leftarrow 0$  to  $T - 1$  do           ▷  $X_t$ : hidden state,  $E_t$ : observed evidence
4:      $x_i \leftarrow$  sample particle from  $P(X_{t+1}|X_t = x_i)$  for  $i = 1, \dots, N$        ▷ Time Elapse Update
5:      $w_i \leftarrow P(E_{t+1}|X_{t+1} = x_i)$  for  $i = 1, \dots, N$            ▷ Evidence Update
6:      $x \leftarrow$  resample  $N$  particles according to weights  $w$            ▷ Particle Resampling
7:   end for
8:   return  $x$ 
9: end procedure
    
```

Algorithm 1.1 outlines the particle filtering algorithm discussed in lecture. The variable x represents a list of N particles, while w is a list of N weights for those particles.

(a) Here, we consider the unweighted particles in x as approximating a distribution.

(i) After executing line 4, which distribution do the particles x represent?

- $P(X_{t+1}|E_{1:t})$
 $P(X_{1:t+1}|E_{1:t})$
 $P(X_{t+1}|E_{1:t+1})$
 None

(ii) After executing line 6, which distribution do the particles x represent?

- $P(X_{t+1}|X_t, E_{1:t+1})$
 $P(X_{1:t+1}|E_{1:t+1})$
 $P(X_{t+1}|E_{1:t+1})$
 None

(b) The particle filtering algorithm should return a sample-based approximation to the true posterior distribution $P(X_T | E_{1:T})$. The algorithm is **consistent** if and only if the approximation converges to the true distribution as $N \rightarrow \infty$. In this question, we present several modifications to Algorithm 1.1. For each modification, indicate if the algorithm is still **consistent** or **not consistent**, and if it is **consistent**, indicate whether you expect it to be **more accurate** in general in terms of its estimate of $P(X_T | E_{1:T})$ (i.e., you would expect the estimated distribution to be closer to the true one) or **less accurate**. Assume unlimited computational resources and arbitrary precision arithmetic.

(i) We modify line 6 to sample 1 or $2N - 1$ particles with equal probability $p = 0.5$ for each time step (as opposed to a fixed number of particles N). You can assume that $P(E_{t+1}|X_{t+1}) > 0$ for all observations and states. This algorithm is:

- Consistent and More Accurate
 Not Consistent
 Consistent and Less Accurate

(ii) Replace lines 4–6 as follows:

4': Compute a tabular representation of $P(X_t = s | E_{1:t})$ based on the proportion of particles in state s .

5': Use the forward algorithm to calculate $P(X_{t+1} | E_{1:t+1})$ exactly from the tabular representation.

6': Set x to be a sample of N particles from $P(X_{t+1} | E_{1:t+1})$.

This algorithm is:

- Consistent and More Accurate
 Not consistent
 Consistent and Less Accurate

(iii) At the start of the algorithm, we initialize each entry in w to 1s. Keep line 4, but replace lines 5 and 6 with the following multiplicative update:

5': For $i = 1, \dots, N$ do

6' $w_i \leftarrow w_i * P(E_{t+1}|X_{t+1} = x_i)$.

Finally, **only** at the end of the T iterations, we resample x according to the cumulative weights w just like in line 6, producing a list of particle positions. This algorithm is:

- Consistent and More Accurate
- Consistent and Less Accurate

Not Consistent

(c) Suppose that instead of particle filtering we run the following algorithm on the Bayes net with T time steps corresponding to the HMM:


1. Fix all the evidence variables $E_{1:T}$ and initialize each X_t to a random value x_t .
2. For $i = 1, \dots, N$ do
 - Choose a variable X_t uniformly at random from X_1, \dots, X_T .
 - Resample X_t according to the distribution $P(X_t | X_{t-1} = x_{t-1}, X_{t+1} = x_{t+1}, E_t = e_t)$.
 - Record the value of X_T as a sample.

Finally, estimate $P(X_T = s | E_{1:T} = e_{1:T})$ by the proportion of samples with $X_T = s$. This algorithm is:

Consistent

Not Consistent

Q2. Most Likely Estimates in HMMs




The Viterbi algorithm finds the most probable sequence of hidden states $X_{1:T}$, given a sequence of observations $e_{1:T}$. Throughout this question you may assume there are no ties. Recall that for the canonical HMM structure, the Viterbi algorithm performs the following **dynamic programming**¹ computations:

$$m_t[x_t] = P(e_t|x_t) \max_{x_{t-1}} P(x_t|x_{t-1})m_{t-1}[x_{t-1}]$$

(a) For the HMM structure above, which of the following probabilities are maximized by the sequence of states returned by the Viterbi algorithm? Mark **all** the correct option(s).

- $P(X_{1:T})$
 $P(X_T|e_T)$
 $P(X_{1:T}|e_{1:T})$
 $P(X_{1:T}, e_{1:T})$
 $P(X_1)P(e_1|X_1) \prod_{t=2}^T P(e_t|X_t)P(X_t|X_{t-1})$
 $P(X_1) \prod_{t=2}^T P(X_t|X_{t-1})$
 None of the above

(b) Consider an HMM structure like the one in part (a) above. Say for all time steps t , the state X_t can take on one of the three values $\{A, B, C\}$. Then, we can represent the state transitions through the following directed graph, also called a *Trellis Diagram*.



We wish to formulate the most probable sequence of hidden state query as a **graph search problem**.

Note in the diagram above, dummy nodes S and G have been added to represent the *start* state and the *goal* state respectively. Further, the transition from the starting node S to the first state X_1 occurs at time step $t = 0$; transition from X_T (the last HMM state) to the goal state G occurs at time step $t = T$.

(The questions for this section are on the following page)

¹ If you're not familiar with dynamic programming, it is essentially a recursive relation in which the current value is defined as a function of previously computed values. In this case, the value at time t is defined as a function of the values at time $t - 1$

Definition : Let $w_{Y \rightarrow Z}^t$, be the cost of the edge for the transition from state Y at time t to state Z at time $t + 1$. For example, $w_{A \rightarrow B}^1$ is the cost of the edge for transition from state A at time 1 to state B at time 2.

(i) For which **one** of the following values for the weights $w_{Y \rightarrow Z}^t$, $1 \leq t < T$, would the minimum cost path be exactly the same as most likely sequence of states computed by the Viterbi algorithm.

- | | |
|--|--|
| <input type="radio"/> $w_{Y \rightarrow Z}^t = -P(X_{t+1} = Z X_t = Y)$ | <input type="radio"/> $w_{Y \rightarrow Z}^t = -P(e_{t+1} X_{t+1} = Z)P(X_{t+1} = Z X_t = Y)$ |
| <input type="radio"/> $w_{Y \rightarrow Z}^t = -\log(P(X_{t+1} = Z X_t = Y))$ | <input type="radio"/> $w_{Y \rightarrow Z}^t = -\log(P(e_{t+1} X_{t+1} = Z)P(X_{t+1} = Z X_t = Y))$ |
| <input type="radio"/> $w_{Y \rightarrow Z}^t = \frac{1}{P(X_{t+1} = Z X_t = Y)}$ | <input type="radio"/> $w_{Y \rightarrow Z}^t = \frac{1}{P(e_{t+1} X_{t+1} = Z)P(X_{t+1} = Z X_t = Y)}$ |

(ii) The initial probability distribution of the state at time $t = 1$ is given $P(X_1 = Y)$, $Y \in \{A, B, C\}$.

Which **one** of the following should be the value of $w_{S \rightarrow Y}^0$, $Y \in \{A, B, C\}$ — these are the cost on the edges connecting S to the states at time $t = 1$?

- | | |
|--|--|
| <input type="radio"/> $w_{S \rightarrow Y}^0 = -P(X_1 = Y)$ | <input type="radio"/> $w_{S \rightarrow Y}^0 = -P(e_1 X_1 = Y)P(X_1 = Y)$ |
| <input type="radio"/> $w_{S \rightarrow Y}^0 = -\log(P(X_1 = Y))$ | <input type="radio"/> $w_{S \rightarrow Y}^0 = -\log(P(e_1 X_1 = Y)P(X_1 = Y))$ |
| <input type="radio"/> $w_{S \rightarrow Y}^0 = \frac{1}{P(X_1 = Y)}$ | <input type="radio"/> $w_{S \rightarrow Y}^0 = \frac{1}{P(e_1 X_1 = Y)P(X_1 = Y)}$ |
| <input type="radio"/> $w_{S \rightarrow Y}^0 = \alpha$, $\alpha \in \mathbb{R}$: (some constant) | |

(iii) Which **one** of the following should be the value of $w_{Y \rightarrow G}^T$, $Y \in \{A, B, C\}$ — these are the cost on the edges connecting the states at last time step $t = T$ to the goal state G ?


- | | |
|--|--|
| <input type="radio"/> $w_{Y \rightarrow G}^T = -P(X_T = Y)$ | <input type="radio"/> $w_{Y \rightarrow G}^T = -P(e_T X_T = Y)P(X_T = Y)$ |
| <input type="radio"/> $w_{Y \rightarrow G}^T = -\log(P(X_T = Y))$ | <input type="radio"/> $w_{Y \rightarrow G}^T = -\log(P(e_T X_T = Y)P(X_T = Y))$ |
| <input type="radio"/> $w_{Y \rightarrow G}^T = \frac{1}{P(X_T = Y)}$ | <input type="radio"/> $w_{Y \rightarrow G}^T = \frac{1}{P(e_T X_T = Y)P(X_T = Y)}$ |
| <input type="radio"/> $w_{Y \rightarrow G}^T = \alpha$, $\alpha \in \mathbb{R}$: (some constant) | |

(c) We consider extending the Viterbi algorithm for finding the most likely sequence of states in modified HMMs.

For your convenience, the computations performed by the Viterbi algorithm for the canonical HMM structure, like in part (a), are repeated below:

$$m_t[x_t] = P(e_t|x_t) \max_{x_{t-1}} P(x_t|x_{t-1})m_{t-1}[x_{t-1}]$$


(i) Consider the HMM below with additional variables U_t . The HMM can be interpreted as : The state X_t at time t is caused due to some action U_t and previous state X_{t-1} . The state X_t emits an observation E_t . Both U_t and X_t are unobserved.



We want to find the most likely sequence of states $X_{1:T}$ and actions $U_{1:T}$, given the sequence of observations $e_{1:T}$. Write a dynamic programming update for $t > 1$ **analogous** to the one for the canonical HMM structure.

$m_t[x_t, u_t] =$ _____

(ii) Consider the HMM below with two emission variables at each time step F_t and E_t . E_t is observed while X_t and F_t are unobserved.




We want to find the most likely sequence of states $X_{1:T}$ and the unobserved emissions $F_{1:T}$, given the sequence of observations $e_{1:T}$. Write a dynamic programming update for $t > 1$ **analogous** to the one for the canonical HMM structure.

$m_t[x_t, f_t] =$ _____

Q3. Hidden Markov Models

Consider the following Hidden Markov Model.



X_1	$\Pr(X_1)$
0	0.3
1	0.7

X_t	X_{t+1}	$\Pr(X_{t+1} X_t)$
0	0	0.4
0	1	0.6
1	0	0.8
1	1	0.2

X_t	O_t	$\Pr(O_t X_t)$
0	A	0.9
0	B	0.1
1	A	0.5
1	B	0.5

Suppose that $O_1 = A$ and $O_2 = B$ is observed.

- (a) Use the Forward algorithm to compute the probability distribution $\Pr(X_2, O_1 = A, O_2 = B)$. Show your work. You do not need to evaluate arithmetic expressions involving only numbers.

- (b) Use the Viterbi algorithm to compute the maximum probability sequence X_1, X_2 . Show your work.

For the next two questions, use the specified sequence of random numbers $\{a_i\}$ generated independently and uniformly at random from $[0, 1)$ to perform sampling. Specifically, to obtain a sample from a distribution over a variable $Y \in \{0, 1\}$ using the random number a_i , pick $Y = 0$ if $a_i < \Pr(Y = 0)$, and pick $Y = 1$ if $a_i \geq \Pr(Y = 0)$. Similarly, to obtain a sample from a distribution over a variable $Z \in \{A, B\}$ using the random number a_i , pick $Z = A$ if $a_i < \Pr(Z = A)$, and pick $Z = B$ if $a_i \geq \Pr(Z = A)$. Use the random numbers $\{a_i\}$ in order starting from a_1 , using a new random number each time a sample needs to be obtained.

- (c) Use likelihood-weighted sampling to obtain 2 samples from the distribution $\Pr(X_1, X_2 | O_1 = A, O_2 = B)$, and then use these samples to estimate $E[\sqrt{X_1 + 3X_2} | O_1 = A, O_2 = B]$.

a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8	a_9	a_{10}
0.134	0.847	0.764	0.255	0.495	0.449	0.652	0.789	0.094	0.028

- (d) [*true or false*] In general, particle filtering using a single particle is equivalent to rejection sampling in the case that there is no evidence. Explain your answer.

- (e) [*true or false*] Performing particle filtering twice, each time with 50 particles, is equivalent to performing particle filtering once with 100 particles. Explain your answer.

- (f) [*true or false*] Variable elimination is generally more accurate than the Forward algorithm. Explain your answer.