1 Particle Filtering

1.1 Particle Filtering Summary

In particle filtering, the value of a particle is one of the possible values that the state variable, X, can take on.

As time goes on we consistently want to know "what is the probability of the state variable X at the *current time*, given what we *currently know about the evidence*". As we move forward in time, and collect more evidence, our specific probability query keeps changing, but in general, we are just updating our belief about the distribution of X. We can approximate this belief distribution, $B(X)$, using our particles at any given time.

1. **Initialize** We start by initializing a list of particles, where each particle has a value randomly selected from the prior probability $P(X_1)$. Set the timestep to $t \leftarrow 1$.

   Our initial estimate of the distribtion of X is:

   $$P(X_1 = x) \approx B'(X_1 = x) = \frac{\text{num particles with value } x}{\text{num particles}}$$

2. **Update** As we observe new information $O_t = o_t$, we would like to update our estimate of the distribution of X.

   Each particle has a weight equal to the likelihood of the observation given it's parent in the Bayes net:

   $$\text{weight}(p_i) \leftarrow P(O_t = o_t \mid X_t = p_i)$$

   If we combine this likelihood weighting with our current belief distribution, we get the updated approximation:

   $$B(X_t = x)_{\text{unnormalized}} = \text{weight}(x) \cdot B'(X_t = x)$$

   where

   $$B'(X_t = x) = \frac{\text{num particles with value } x}{\text{num particles}}$$

   This is an approximation of the forward algorithm update step:

   $$P(X_t = x, O_t = o \mid o_{1:t-1}) = P(O_t = o_t \mid X_t = x)P(X_t = x \mid o_{1:t-1})$$

   Note that in the particle filtering algorithm, we seek to approximate the following:

   $$B'(X_t) \approx P(X_t \mid o_{1:t-1})$$

   $$B(X_t) \approx P(X_t \mid o_{1:t})$$

3. **Normalize** Normalize the values in the belief table to make sure the probabilities sum to one:

   $$B(X_t) \leftarrow \frac{B(X_t)_{\text{unnormalized}}}{\sum\limits_{x \in X} B(X_t = x)_{\text{unnormalized}}}$$

   This is an approximation of the normalization part of the forward algorithm update step:

   $$P(X_t = x \mid o_t) = \frac{P(X_t = x, o_t \mid o_{1:t-1})}{\sum_{x \in X_t} P(X_t = x, o_t \mid o_{1:t-1})}$$

4. **Resample** Next, sample a brand new list of particles from $B(X_t)$.

5. **Predict** Change each particle p_i, by randomly sampling a new value from $P(X_{t+1} \mid X_t = p_i)$.
   This is equivalent to the forward algorithm predict step:

$$P(X_{t+1} \mid o_{1:t}) = \sum_{x \in X_t} P(X_{t+1} \mid P(X_t = x)P(X_t = x \mid o_{1:t})$$

   Repeat at **Update** for $t \leftarrow t + 1$ until t reaches desired timestep.

1.2 Particle Filters Problem

Consider the following HMM:

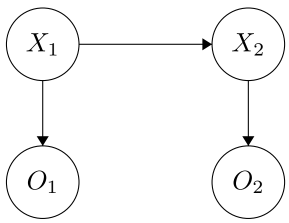

X_t	X_{t+1}	$P(X_{t+1} \mid X_t)$
0	0	0.4
0	1	0.6
1	0	0.8
1	1	0.2

X_t	O_t	$P(O_t \mid X_t)$
0	a	0.9
0	b	0.1
1	a	0.5
1	b	0.5

Suppose that we observe $O_1 = a$ and $O_2 = b$.
Using the forward algorithm, compute the probability distribution $P(X_2 \mid O_1 = a,\ O_2 = b)$ We start with two particles representing our distribution for X_1.
$p_1 : X_1 = 0$
$p_2 : X_1 = 1$
Use the following random numbers to run particle filtering:

$$[0.22, 0.05, 0.33, 0.20, 0.84, 0.54, 0.79, 0.66, 0.14, 0.96]$$

1. **Observe (Update)**: Compute the weight of the two particles after evidence $O_1 = a$.

$$\text{weight}(p_1) = P(O_1 = a \mid X_1 = 0) = 0.9$$
$$\text{weight}(p_2) = P(O_1 = a \mid X_1 = 1) = 0.5$$

2. **Resample**: Using the random numbers, resample p_1 and p_2 based on the weights.

   First we calculate the beliefs:

$$B(X_1 = 0)_{\text{unnormalized}} = 0.9 \cdot \frac{1}{2} = 0.45$$

$$B(X_1 = 1)_{\text{unnormalized}} = 0.5 \cdot \frac{1}{2} = 0.25$$

$$B(X_1 = 0) = \frac{B(X_1 = 0)_{\text{unnormalized}}}{B(X_1 = 0)_{\text{unnormalized}} + B(X_1 = 1)_{\text{unnormalized}}} = \frac{9}{14} = 0.643$$

$$B(X_1 = 1) = \frac{B(X_1 = 1)_{\text{unnormalized}}}{B(X_1 = 0)_{\text{unnormalized}} + B(X_1 = 1)_{\text{unnormalized}}} = \frac{5}{14} = 0.357$$

   We now sample from the weighted distribution we found above. After normalizing the weights, we find that p_1 maps to range $[0, 0.643)$, and p_2 maps to range $[0.643, 1)$. Using the first two random samples, we find:

$$p_1 = \text{sample}(B(X_1),\ 0.22) = 0$$
$$p_2 = \text{sample}(B(X_1),\ 0.05) = 0$$

3. **Elapse Time (Predict)**: Now let's compute the elapse time particle update. Sample p_1 and p_2 from applying the time update.

$$p_1 = \text{sample}(P(X_2 \mid X_1 = p_1),\ 0.33) = \text{sample}(P(X_2 \mid X_1 = 1),\ 0.33) = 0$$
$$p_2 = \text{sample}(P(X_2 \mid X_1 = p_2),\ 0.20) = \text{sample}(P(X_2 \mid X_1 = 0),\ 0.20) = 0$$

4. **Observe (Update)**: Compute the weight of the two particles after evidence $O_2 = b$.

$$\text{weight}(p_1) = P(O_2 = b \mid X_2 = 0) = 0.1$$
$$\text{weight}(p_2) = P(O_2 = b \mid X_2 = 0) = 0.1$$

5. **Resample**: Using the random numbers, resample p_1 and p_2 based on the weights.

Because our only two particles have $X = 0$, we have $B(X_2 = 0) = 1$ and $B(X_2 = 1) = 0$, so resampling will still leave us with two particles with $X = 0$.

$$p_1 = 0$$

$$p_2 = 0$$

6. What is our estimated distribution for $P(X_2 \mid O_1 = a,\ O_2 = b)$?

$$P(X_2 = 0 \mid O_1 = a,\ O_2 = b) = 2/2 = 1$$
$$P(X_2 = 1 \mid O_1 = a,\ O_2 = b) = 0/2 = 0$$

2 Game Theory

2.1 Terminology

- Forms of Game
  - What is the normal form of a game? What does it represent and how?

    Approximate the game. Both agents act together once. Payoff represented by payoff matrix.
  - What is the extensive form of a game? What does it represent and how?

    Represent agents, each agent's move, agents' knowledge when making every move, and the payoffs (tree form, listing all possibilities, can consider as players taking turns to act with each action having a different probability.
- What's the difference between pure vs mixed strategy? How do we calculate the utility of a strategy?

  Pure: deterministic.
  Mixed: stochastic. To calculate utilities, we take the weighted average of the utilities with weights being joint probability of different actions combinations of all players.
- What are zero-sum games?

  Each cell of the table sums up to 0. (Loss for one, gain for the other, etc)
- What's the difference between strictly dominant strategy and weakly dominant strategy?

  Dominant strategy: No matter what's the opponent's strategy, for player i the dominant strategy is always better than all other strategies. For strictly dominant the utility of such strategy is always the highest (i.e., ">" than all other strategies' utilities), where weakly dominant strategies have utilities "≥" other strategies' utilities.
- What is "Tragedy of the Commons"?

  Individual players act against the common good.
- What is the Social Welfare?

  The sum of utilities of all the players. Social Welfare is not always maximized. (e.g., Prisoner's dilemma)
- What is Nash Equilibrium?

  NE: No player has any tendency of unilaterally changing its strategy. (i.e., utility of current strategy is greater than other strategies for a player i).
- Finding Pure Strategy Nash Equilibrium(PSNE)
  - Naive looping through
  - Elimination of rows/cols other than strictly dominating strategies.
  - Elmination of strictly dominated strategy.
- Find Mixed Strategy Nash Equilibrium(MSNE)
  - The utilities of the weighted actions are equal.(Solve system of equations)

2.2 A toy example

Below is a toy example for a game between player 1 and player 2. The first number in each cell represents the utility of player 1 while the second number represents the utility of player 2.

		Player 2	
		A	B
Player 1	A	5,6	1,2
	B	4,4	3,3

- Is there any pure strategy Nash equilibrium in this game? If so, what is it?

  Yes, there are 1 PSNE in this game, (A,A) as choices of Player 1 and Player 2 respectively. Since if Player 1's choosing A, Player 2 has no incentive to deviate from A to B since 6>2, and if Player 2 chooses A, Player 1 has no incentive to deviate from A to B since 5>4. No player has the incentive to deviate from their strategy. So (A,A) is a Nash Equilibrium.

- Is there any mixed strategy Nash Equilibrium in this game? If so, what is it?

  If there's a Mixed Strategy Nash Equilibrium, we need to make sure that no matter what the other players' strategies are, no player has the tendency to deviate from their current mixed strategy. i.e., the total utility of the person's mixed strategy remains the same given others choosing different strategies Therefore, let u_1 denote the utility of player 1's mixed strategy and u_2 denotes utility of player 2's mixed strategy:

$$
\begin{aligned}
u_1(\text{P1 chooses A}) &= u_1(\text{P1 chooses B}) \\
\Rightarrow 5q + 1(1-q) &= 4q + 3(1-q) \\
\Rightarrow q &= 2/3 \\
u_2(\text{P2 chooses A}) &= u_2(\text{P2 chooses B}) \\
\Rightarrow 6p + 4(1-p) &= 2p + 3(1-p) \\
\Rightarrow p &= -1/3
\end{aligned}
$$

  It doesn't make any sense for a mixed strategy to be choosing an action with negative probability. When there exists a dominant strategy for a player, the probability would become negative, and observe that in this case, action A is player 2's dominant strategy. So this specific search space does not allow both players to have mixed strategies. In this case, if both players play optimally, player 2 would always choose the dominant strategy A, forcing player 1 to choose action A. Therefore, we result in our Pure Strategy Nash Equilibrium, which indicates that there's no Mixed Strategy Nash Equilibrium in this case.(However, it's still possible for both MSNE and PSNE to exist in some other cases.)

- Which strategies should both players take to maximize social welfare?

  The strategy (A,A) would maximize the social welfare since the total utility is $5 + 6 = 11$ which is the maximum among all cells.