

Particle Filtering

The Hidden Markov Model analog to Bayes' net sampling is called **particle filtering**, and involves simulating the motion of a set of particles through a state graph to approximate the probability (belief) distribution of the random variable in question.

Instead of storing a full probability table mapping each state to its belief probability, we'll instead store a list of n particles, where each particle is in one of the d possible states in the domain of our time-dependent random variable.

Once we've sampled an initial list of particles, the simulation takes on a similar form to the forward algorithm, with a time elapse update followed by an observation update at each timestep:

- *Prediction update* - Update the value of each particle according to the transition model. For a particle in state W_t , sample the updated value from the probability distribution given by $Pr(W_{t+1}|w_t)$. Note the similarity of the prediction update to prior sampling with Bayes' nets, since the frequency of particles in any given state reflects the transition probabilities.
- *Observation update* - During the observation update for particle filtering, we use the sensor model $Pr(O_t|W_t)$ to weight each particle according to the probability dictated by the observed evidence and the particle's state. Specifically, for a particle in state w_t with sensor reading o_t , assign a weight of $Pr(o_t|w_t)$. The algorithm for the observation update is as follows:
 1. Calculate the weights of all particles as described above.
 2. Calculate the total weight for each state.
 3. If the sum of all weights across all states is 0, reinitialize all particles.
 4. Else, normalize the distribution of total weights over states and resample your list of particles from this distribution.

Note the similarity of the observation update to likelihood weighting, where we again downweight samples based on our evidence.

Utilities

Rational agents must follow the **principle of maximum utility** - they must always select the action that maximizes their expected utility. However, obeying this principle only benefits agents that have **rational preferences**.

- If an agent prefers receiving a prize A to receiving a prize B , this is written $A \succ B$
- If an agent is indifferent between receiving A or B , this is written as $A \sim B$

- A **lottery** is a situation with different prizes resulting with different probabilities. To denote lottery where A is received with probability p and B is received with probability $(1 - p)$, we write

$$L = [p, A; (1 - p), B]$$

In order for a set of preferences to be rational, they must follow the five **Axioms of Rationality**:

- **Orderability:** $(A \succ B) \vee (B \succ A) \vee (A \sim B)$
A rational agent must either prefer one of A or B , or be indifferent between the two.
- **Transitivity:** $(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$
If a rational agent prefers A to B and B to C , then it prefers A to C .
- **Continuity:** $A \succ B \succ C \Rightarrow \exists p [p, A; (1 - p), C] \sim B$
If a rational agent prefers A to B but B to C , then it's possible to construct a lottery L between A and C such that the agent is indifferent between L and B with appropriate selection of p .
- **Substitutability:** $A \sim B \Rightarrow [p, A; (1 - p), C] \sim [p, B; (1 - p), C]$
A rational agent indifferent between two prizes A and B is also indifferent between any two lotteries which only differ in substitutions of A for B or B for A .
- **Monotonicity:** $A \succ B \Rightarrow (p \geq q \Leftrightarrow [p, A; (1 - p), B] \succeq [q, A; (1 - q), B])$
If a rational agent prefers A over B , then given a choice between lotteries involving only A and B , the agent prefers the lottery assigning the highest probability to A .

If all five axioms are satisfied by an agent, then it's guaranteed that the agent's behavior is describable as a maximization of expected utility.

Example: Consider the following lottery:

$$L = [0.5, \$0; 0.5, \$1000]$$

This represents a lottery where you receive \$1000 with probability 0.5 and \$0 with probability 0.5. Now consider three agents A_1 , A_2 , and A_3 which have utility functions $U_1(\$x) = x$, $U_2(\$x) = \sqrt{x}$, and $U_3(\$x) = x^2$ respectively. If each of the three agents were faced with a choice between participating in the lottery and receiving a flat payment of \$500, which would they choose? The respective utilities for each agent of participating in the lottery and accepting the flat payment are listed in the following table:

Agent	Lottery	Flat Payment
1	500	500
2	15.81	22.36
3	500000	250000

These utility values for the lotteries were calculated as follows:

$$U_1(L) = U_1([0.5, \$0; 0.5, \$1000]) = 0.5 \cdot U_1(\$1000) + 0.5 \cdot U_1(\$0) = 0.5 \cdot 1000 + 0.5 \cdot 0 = \boxed{500}$$


$$U_2(L) = U_2([0.5, \$0; 0.5, \$1000]) = 0.5 \cdot U_2(\$1000) + 0.5 \cdot U_2(\$0) = 0.5 \cdot \sqrt{1000} + 0.5 \cdot \sqrt{0} = \boxed{15.81}$$

$$U_3(L) = U_3([0.5, \$0; 0.5, \$1000]) = 0.5 \cdot U_3(\$1000) + 0.5 \cdot U_3(\$0) = 0.5 \cdot 1000^2 + 0.5 \cdot 0^2 = \boxed{500000}$$

With these results, we can see that agent A_1 is indifferent between participating in the lottery and receiving the flat payment (the utilities for both cases are identical). Such an agent is known as **risk-neutral**. Similarly, agent A_2 prefers the flat payment to the lottery and is known as **risk-averse** and agent A_3 prefers the lottery to the flat payment and is known as **risk-seeking**.

1 Particle Filtering

Let's use Particle Filtering to estimate the distribution of $P(W_2|O_1 = a, O_2 = b)$. Here's the HMM again. O_1 and O_2 are supposed to be shaded.



W_1	$P(W_1)$
0	0.3
1	0.7

W_t	W_{t+1}	$P(W_{t+1} W_t)$
0	0	0.4
0	1	0.6
1	0	0.8
1	1	0.2

W_t	O_t	$P(O_t W_t)$
0	a	0.9
0	b	0.1
1	a	0.5
1	b	0.5

We start with two particles representing our distribution for W_1 .

$P_1 : W_1 = 0$

$P_2 : W_1 = 1$

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]

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

(b) **Resample:** Using the random numbers, resample P_1 and P_2 based on the weights.

(c) **Predict:** Sample P_1 and P_2 from applying the time update.

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

(e) **Resample:** Using the random numbers, resample P_1 and P_2 based on the weights.

(f) What is our estimated distribution for $P(W_2|O_1 = a, O_2 = b)$?

2 Utilities

1. Consider a utility function of $U(x) = 2x$. What is the utility for each of the following outcomes?

(a) 3

(b) $L(\frac{2}{3}, 3; \frac{1}{3}, 6)$

(c) -2

(d) $L(0.5, 2; 0.5, L(0.5, 4; 0.5, 6))$

2. Consider a utility function of $U(x) = x^2$. What is the utility for each of the following outcomes?

(a) 3

(b) $L(\frac{2}{3}, 3; \frac{1}{3}, 6)$

(c) -2

(d) $L(0.5, 2; 0.5, L(0.5, 4; 0.5, 6))$

3. What is the expected monetary value (EMV) of the lottery $L(\frac{2}{3}, \$3; \frac{1}{3}, \$6)$?

4. For each of the following types of utility function, state how the utility of the lottery $U(L)$ compares to the utility of the amount of money equal to the EMV of the lottery, $U(EMV(L))$. Write $<$, $>$, $=$, or $?$ for can't tell.

(a) U is an arbitrary function.
 $U(L) \text{ ___ } U(EMV(L))$

(b) U is monotonically increasing and its rate of increase is increasing (its second derivative is positive).
 $U(L) \text{ ___ } U(EMV(L))$

(c) U is monotonically increasing and linear (its second derivative is zero).
 $U(L) \text{ ___ } U(EMV(L))$

(d) U is monotonically increasing and its rate of increase is decreasing (its second derivative is negative).
 $U(L) \text{ ___ } U(EMV(L))$