Sampling Based Filters

Lecturer: Drew Bagnell *Scribe:Haley Dalzell, Matt Barnes, De-An Huang* [1]

1 Monty Hall Problem

The Monty Hall Problem is a probability puzzle originally posed by Steve Selvin. It is loosely based on the game show Deal or No Deal and is named after its host, Monty Hall. The problem is as follows.

> *Suppose you're on a game show, and you're given the choice of three doors: Behind one door is a car; behind the others, goats. You pick a door, say No. 1, and the host, who knows what's behind the doors, opens another door, say No. 3, which has a goat. He then says to you, "Do you want to pick door No. 2?" Is it to your advantage to switch your choice? (Whitaker, 1990, as quoted by vos Savant 1990a)*

The problem is of particular interest because most people find the solution counter-intuitive. The solution also changes if the defined behavior of the host changes.

1.1 Canonical Solution

In order to solve this problem, some assumptions must be made about the host's behavior:

- The host will never select the door with the car behind it.
- The host will never select the same door you have selected.
- If you initially select the door with the car behind it, the host is equally likely to open either remaining door.

Because the car is equally likely to be behind any of the three doors, it does not matter which door is selected first. Without loss of generality, say we select door 1. The host then opens a door with a goat behind it. Say this happens to be door 2.

Our goal is to compute the probability that the door we selected has a car behind it, given that we selected door 1 and that we've observed that door 2 does not have a car behind it. In other words, we want to compute:

$$P(Door_1 = 1 | Select_1, Observe_2 = 0) \tag{1}$$

[1]Some content adapted from previous scribes: M Taylor and M Shomin

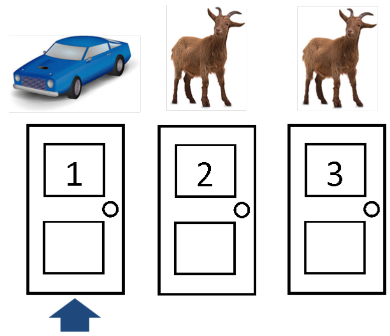

Figure 1: We select door 1, it contains the car

where

$$Door_n = \begin{cases} 1 & \text{if door } n \text{ contains car} \\ 0 & \text{if door } n \text{ contains a goat} \end{cases}$$

$$Observe_n = \begin{cases} 1 & \text{if Monty opens door } n \text{ and we observe a car} \\ 0 & \text{if Monty opens door } n \text{ and we observe a goat} \end{cases}$$

Using Bayes' Rule on the above probability, we will compute the following for each door:

$$P(Observe_2 = 0|Select_1, Door_n = 1) * P(Door_n = 1|Select_1) \tag{2}$$

The right hand quantity is easy to compute. Since the door containing the car is independent of the door,

$$P(Door_n = 1|Select_1) = \frac{1}{3} \text{ for } n \in 1, 2, 3 \tag{3}$$

Door 1:

$$P(Observe_2 = 0|Select_1, Door_1 = 1) * P(Door_1 = 1|Select_1) \tag{4}$$

In this case, we want to know the probability that the host shows us a goat behind door 2, given that we've selected door 1 and the car is behind door 1. Since the remaining two doors contain goats, the host is equally likely to pick either door. Thus, the probability that he opens door two and we observe that it does not contain a car is 1/2.

$$P(Observe_2 = 0|Select_1, Door_1 = 1) * P(Door_1 = 1|Select_1) = \frac{1}{2} * \frac{1}{3} = \frac{1}{6} \tag{5}$$

Door 2:

$$P(Observe_2 = 0|Select_1, Door_2 = 1) * P(Door_2 = 1|Select_1) \tag{6}$$

In this case, we want to know the probability that the host shows us a goat behind door 2, given we've selected door 1 and the car is behind door 2. Because the car is behind door 2, the host cannot show us a goat behind door 2, so the probability that we observe that there is a goat behind door 2 is 0.

$$P(Observe_2 = 0|Select_1, Door_2 = 1) * P(Door_2 = 1|Select_1) = 0 * \frac{1}{3} = 0 \tag{7}$$

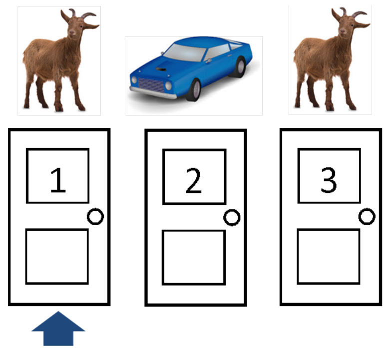

Figure 2: We select door 1, door 2 contains the car

Door 3:

$$P(Observe_2 = 0|Select_1, Door_3 = 1) * P(Door_3 = 1|Select_1) \tag{8}$$

In this case, we want to know the probability that the host shows us a goat behind door 2, given that we've selected door 1 and the car is behind door 3. Because the host cannot select our door (door 1) and he cannot select the door containing the car (door 3), he must pick door 2. Thus, the probability that we observe a goat behind door 2 is 1.

$$P(Observe_2 = 0|Select_1, Door_2 = 1) * P(Door_2 = 1|Select_1) = 1 * \frac{1}{3} = \frac{1}{3} \tag{9}$$

Now we must normalize the probabilities over all the doors so they sum to 1:

$$P(Door_1 = 1|Select_1, Observe_2 = 0) = \frac{\frac{1}{6}}{\frac{1}{6} + 0 + \frac{1}{3}} = \frac{1}{3} \tag{10}$$

$$P(Door_2 = 1|Select_1, Observe_2 = 0) = \frac{0}{\frac{1}{6} + 0 + \frac{1}{3}} = 0 \tag{11}$$

$$P(Door_3 = 1|Select_1, Observe_2 = 0) = \frac{\frac{1}{3}}{\frac{1}{6} + 0 + \frac{1}{3}} = \frac{2}{3} \tag{12}$$

We see that the probability that door 1 contains a car remains 1/3, but the probability that door 3 contains a car is now 2/3. Thus, we should take the host's offer to switch doors.

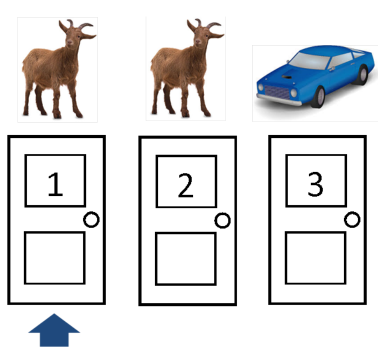

Figure 3: We select door 1, door 3 contains the car

1.2 100 Doors Variation

The intuitive reasoning behind why one should switch doors in the three door example becomes much more obvious when more doors are introduced.

Say now we have 100 doors, with one containing a car, and the other 99 containing goats. We select one door, and the host opens 98 other doors containing goats. The probability that our first guess was correct is very small (1/100). By the same logic used in the 3 door example, the probability that host intentionally avoided the remaining closed door because it contained a car is 99/100. Thus the obvious choice is to switch doors when the host offers.

1.3 Alternate Host Behavior

If the behavior of the host is altered, the solution can change. Say now that the host is able to select the door with the car behind it.

Door 1: Like the original version, the remaining two doors do not contain a car, and the host is equally likely to pick either one so the probability that he opens door two and we observe that it does not contain a car is 1/2.

$$P(Observe_2 = 0|Select_1, Door_1 = 1) * P(Door_1 = 1|Select_1) = \frac{1}{2} * \frac{1}{3} = \frac{1}{6} \tag{13}$$

Door 2: Again, because the car is behind door 2, it is impossible for us to observe a goat behind door 2.

$$P(Observe_2 = 0|Select_1, Door_2 = 1) * P(Door_2 = 1|Select_1) = 0 * \frac{1}{3} = 0 \tag{14}$$

Door 3: Because the host is now equally likely to open either remaining door, the probability that we observe that there is a goat behind door 2 is changed to 1/2

$$P(Observe_2 = 0|Select_1, Door_2 = 1) * P(Door_2 = 1|Select_1) = \frac{1}{2} * \frac{1}{3} = \frac{1}{6} \tag{15}$$

Now we must normalize the probabilities so that they sum to 1:

$$P(Door_1 = 1|Select_1, Observe_2 = 0) = \frac{\frac{1}{6}}{\frac{1}{6} + 0 + \frac{1}{6}} = \frac{1}{2} \tag{16}$$

$$P(Door_2 = 1|Select_1, Observe_2 = 0) = \frac{0}{\frac{1}{6} + 0 + \frac{1}{6}} = 0 \tag{17}$$

$$P(Door_3 = 1|Select_1, Observe_2 = 0) = \frac{\frac{1}{3}}{\frac{1}{6} + 0 + \frac{1}{6}} = \frac{1}{2} \tag{18}$$

Thus, because of the host's altered behavior, the car is equally likely to be behind either the door you picked or the remaining door. Switching does not help or harm you.

2 Monte Carlo Method

As discussed in previous lectures, the belief of a robot, for example, is defined as:

$$Bel(x_t) := p(x_t | u_{1:t}, z_{1:t}) \tag{19}$$

where x_t is the robot's position, u_t is the control input, and z_t is the observation, all at time t. If we wish to use the belief of x_t in any useful manner we are going to use the expected value. The expected value of any function f is given by:

$$\mathbb{E}[f] = \sum_x f(x)p(x), \tag{20}$$

where $p(x)$ is the distribution of x. Rarely, if ever, is the distribution known in closed form. However, if you can generate samples of x_i from the distribution $p(x)$ the expectation can be estimated as:

$$\mathbb{E}[f] \approx \frac{1}{N} \sum_{i=1}^{N} f(x_i). \tag{21}$$

This sampling based approach falls into a general class of algorithms called Monte Carlo Methods. This method of using random sampling to compute results have been reinvented/rediscovered many times, with the name being given by physicists at Los Alamos Scientifc Laboratory in reference to the Monte Carlo Casino in Monaco.

2.1 Example 1: Which room am I in?

We have a map with rooms numbered 1, 2, 3 as shown in Fig. 4.

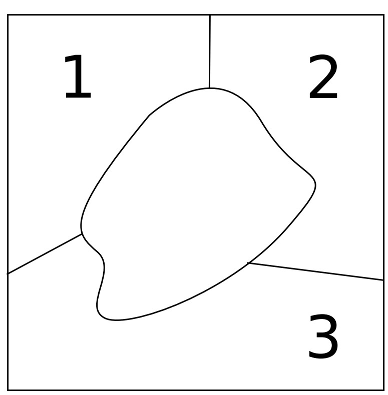

Figure 4: Rooms in which robot may be located

We want probability that the robot is in room 1. But how would we write this as an expectation? We define rooms as subsets of continuous space, and invent a function $f()$ which is 1 if $x_t \in room_1$, 0 otherwise. This is essentially the indicator function of the room: $\mathbb{1}_{room_1}$. Then the probability the robot is in the room is also the expectation of the indicator function:

$$\mathbb{E}[\mathbb{1}_{room_1}(x)] = \sum_x p(x)\mathbb{1}_{room_1}(x) = \sum_{x \in room_1} p(x) = Pr(robot\ is\ in\ room\ 1)$$

In our case, this is the sum of the probabilities of being in a particular room. In general, the expectation of an indicator function is the probability of the indicated expression.

In many cases, we don't know $p(x)$ in closed form, but we can generate marginal and independent samples x_i from $\mathbb{X}$. Supposing this is true, by law of large numbers (LLN), we can do

$$\frac{1}{N}\sum x_i \xrightarrow[N\to\infty]{} \mathbb{E}[x] \tag{22}$$

or more strongly

$$\frac{1}{N}\sum f(x_i) \xrightarrow[N\to\infty]{} \mathbb{E}[f(x)] \tag{23}$$

2.2 Example 2: Area Estimation

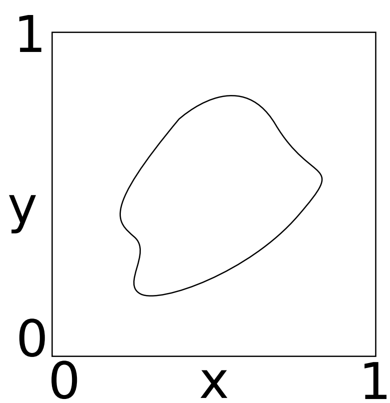

Figure 5: Area we wish to calculate

Suppose we have a shape we wish to calculate the area of, as depicted by the amoeba shape in Fig. 5. Conceptually we can find the desired area by randomly sampling the area and taking the proportion of samples that fell inside the area to the total number of samples and multiplying by the area of the box.

More formally:

- Let $X = [x\ y]^T$ be our uncertain quantity (Random Variable).
- Assume we have an indicator function $\mathbb{1}_{amoeba}(X_i)$ which takes the value 1 if X_i lies in the area and 0 otherwise.
- We can sample uniformly from $\mathbb{X}$

The percentage of the total area the desired area covers is the expected value of the indicator function.

$$\mathbb{E}[Y] \approx \frac{1}{N} \sum_{i=1}^{N} g(X_i) \tag{24}$$

Then area of the shape is

$$A_{amoeba} \approx \mathbb{E}[Y] A_{total} \tag{25}$$

2.3 Return to Robots

In the case of robots we would like to compute the expected value of a function f knowing previous location x_{t-1} and measurement z_t.

$$\mathbb{E}_{p(x_t|z_t, x_{t-1})}[f(x_t)] \tag{26}$$

We will likely never know a closed form solution for this distribution. So, we pursue the previously described sample based Monte Carlo Method. This would require samples drawn from the distribution

$$p(x_t|z_t, x_{t-1}), \tag{27}$$

which can be rewritten as

$$\frac{1}{\mathcal{Z}} \, p(z_t|x_t) p(x_t|x_{t-1}). \tag{28}$$

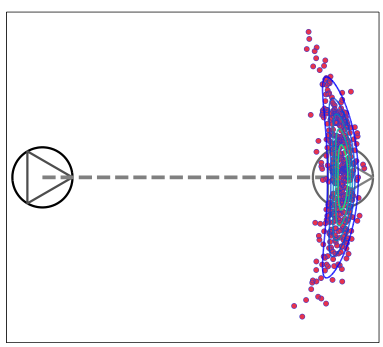

Figure 6: The motion model: Posterior distributions of the robot's pose upon executing the motion command illustrated by the solid line. The darker a location, the more likely it is. This plot has been projected into 2-D. The original density is three-dimensional, taking the robot's heading direction θ into account.

The distribution is the product of the measurement model, $p(z_t|x_t)$, and the robot motion model, $p(x_t|x_{t-1})$. It is not obvious how to sample from the product of these distributions. It is, however, straightforward to sample from the robot motion model, a classic example of which is shown in Fig. 6.

3 Importance sampling (IS)

Although we are unable to sample from the required distribution, $p(x)$, we are able to evaluate the function at desired x_i. Thus, we can sample from a known distribution $q(x)$ and evaluate $p(x)$ to correctly reweigh the samples. This trick is referred to as Importance Sampling. Manipulating the expectation over $p(x)$ yields the following approximation:

$$\begin{aligned}
\mathbb{E}_{p(x)}[f(x)] &= \sum p(x) f(x) && \text{(29a)} \\
&= \sum p(x) f(x) \frac{q(x)}{q(x)} && \text{(29b)} \\
&= \sum q(x) \frac{p(x)}{q(x)} f(x) && \text{(29c)} \\
&= \mathbb{E}_{q(x)}[\frac{p(x)}{q(x)} f(x)] && \text{(29d)} \\
&\approx \frac{1}{N} \sum_{i=1}^{N} \frac{p(x_i)}{q(x_i)} f(x_i). && \text{(29e)}
\end{aligned}$$

where $\frac{p(x)}{q(x)}$ is called the importance weight. For a given $p(x)$, when we are undersampling an area (i.e. $q(x)$ is small), it weights it stronger; likewise, oversampled areas are weighted more weakly. Consider the following (somewhat contrived) pair of functions:

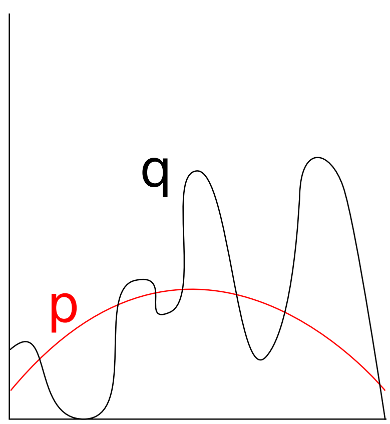

Figure 7: Comparison of distribution we wish to sample from, p, and distribution we can sample from, q.

Notice that there are places where p is nonzero and q is zero, which means we have regions to sample that we aren't sampling at all. This problematic issue will be discussed in future lectures. Our estimate of the expectation of p becomes:

$$\frac{1}{N} \sum_{i=1}^{N} \frac{p(x_i)}{q(x_i)} f(x_i) \xrightarrow[N \to \infty]{} \mathbb{E}_p[f(x)] \tag{30}$$

3.1 Example

Applying the importance sampling to Section 2.3, we can set $q = p(x_t|x_{t-1})$ as the motion model, and $p = \frac{1}{\mathcal{Z}}p(z_t|x_t)p(x_t|x_{t-1})$.

$$\begin{aligned}
&\text{Draw } N \text{ samples}\\
&\quad x^i \sim p(x_t|x_{t-1})\\
&\quad w^i = \frac{p}{q} = p(z_t|x_t)/\mathcal{Z}\\
&\text{Provide estimate } \frac{1}{N}\sum w^i f(x_i)
\end{aligned}$$

Notice that we need the normalizer $\mathcal{Z}$, discussed in the next lecture.