Example of Bayes Rule

Suppose we are trying to estimate if a pedestrian will cross the street or wait. We have some processed sensor measurement z. What is $P(cross \mid z)$?

Given: $P(z \mid cross) = 0.6$
$P(z \mid wait) = 0.3$

can you think of how we might get this / what this means?

this is called an uninformative prior!

$$P(cross) = P(wait) = 0.5$$

$$P(cross \mid z) = \frac{P(z \mid cross)\,P(cross)}{\underbrace{P(z)}_{\text{how to get this?}}}$$

$$P(z) = P(z \mid cross)\,P(cross) + P(z \mid wait)\,P(wait)$$

$$P(cross \mid z) = \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5}$$

$$= 2/3 = 0.67$$

Example of Recursive Bayesian Filtering

Again, let's look at whether or not a pedestrian will cross. We will denote cross as c and wait as w. Let's suppose that the vehicle can stop (s) or move forward (m). Initially, assume an uninformed prior:

$$bel(x_0 = c) = bel(x_0 = w) = 0.5$$

Suppose our perception system outputs two measurements:

detect cross (dc)
detect wait (dw)

We are given the following sensor model:

$$p(\underset{\text{dc or dw}}{z_t} \mid \underset{\text{c or w}}{x_t}) \rightarrow \underbrace{p(dc \mid c)}_{\text{the probability of detecting cross given that the person is crossing}} = 0.6$$

this model can be represented as a table:

		x_t: c	x_t: w
z_t	dc	.6	.2
z_t	dw	.4	.8

$$\left.\begin{aligned} p(dw \mid c) &= 0.4 \\ p(dc \mid w) &= 0.2 \\ p(dw \mid w) &= 0.8 \end{aligned}\right\} \text{what do these probabilities represent?}$$

now consider the impact of the vehicle.

$p(x_t \mid u_t, x_{t-1})$

$\hookrightarrow$ s or m $\rightarrow$ these actions can provide more information about the env state

intuitively, what does this mean? can you list out the probabilities as we gave for the sensor model?

we can visualize this model as a graph:

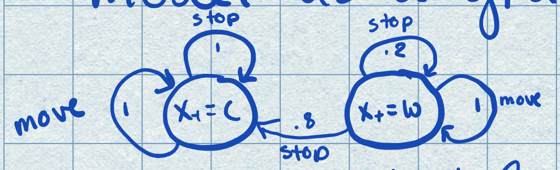

... or as tables:

u_t = move (m)

x_t \ x_{t-1}	c	w
c	1	0
w	0	1

u_t = stop (s)

x_t \ x_{t-1}	c	w
c	1	.8
w	0	.2

at t=1, stop and see what happens:

$$\overline{bel}(x_1) = \sum_{x_0} p(x_1 \mid u_1, x_0)\, bel(x_0)$$

$$= p(x_1 \mid u_1 = s, x_0 = c)\, bel(x_0 = c) \quad \leftarrow \text{case where } x_0 \text{ is cross}$$

$$+\, p(x_1 \mid u_1 = s, x_0 = w)\, bel(x_0 = w) \quad \leftarrow \text{wait case}$$

note that this is fixed

if $x_1 = c$, then:

init at .5

$$\overline{bel}(x_1 = c) = p(x_1 = c \mid s, c)\, bel(c) + p(x_1 = c \mid s, w)\, bel(w)$$

$$= 1 \cdot 0.5 + 0.8 \cdot 0.5$$

$$= 0.9$$

if $x_1 = w$, then:

$$\overline{bel}(x_1 = w) = p(w \mid s, c)\, bel(c) + p(w \mid s, w)\, bel(w)$$

$$= 0 \cdot 0.5 + 0.2 \cdot 0.5 = 0.1$$

recall: $bel(x_1) = \eta\, p(z_1 | x_1)\, \overline{bel}(x_1)$

suppose z_1 = detect cross (dc)

note fixed!

1) $x_1 = c$

$bel(x_1 = c) = \eta\, p(z_1 = dc | x_1 = c)\, \overline{bel}(x_1 = c)$

$= \eta \cdot 0.6 \cdot 0.9$

$= \eta \cdot 0.54 \approx 0.964$

2) $x_1 = w$

$bel(x_1 = w) = \eta\, p(z_1 = dc | x_1 = w)\, \overline{bel}(x_1 = w)$

$= \eta \cdot 0.2 \cdot 0.1$

$= \eta \cdot 0.02 \approx 0.036$

3) normalizer: $\eta = \dfrac{1}{0.54 + 0.02} = 1.786$

now substitute in!

at this point, we are pretty confident that the person is crossing.

(but are we confident enough?)

let's do one more round of updates.

suppose we move the vehicle ($u_2 = m$) and get a new measurement: ($z_2 = dc$).

$$\overline{bel}(x_2=c) = p(c|m,c)\,\overbrace{bel(c)}^{t=1} + p(c|m,w)\,bel(w)$$
$$= 1\cdot 0.964 + 0\cdot 0.036$$
$$= 0.964$$

$$\overline{bel}(x_2=w) = p(w|m,c)\,bel(c) + p(w|m,w)\,bel(w)$$
$$= 0\cdot 0.964 + 1\cdot 0.036$$
$$= 0.036$$

→ why didn't these values change?

$$bel(x_2=c) = \eta\, p(dc|c)\,\overline{bel}(c)$$
$$= \eta\cdot 0.6\cdot 0.964 = 0.988$$
$$bel(x_2=w) = \eta\, p(dc|w)\,\overline{bel}(w)$$
$$= \eta\cdot 0.2\cdot 0.036 = 0.012$$
$$\eta = \frac{1}{0.5856} = 1.708$$

food for thought:

1. how confident should we be in our estimation before we trust it?
2. if we received a measurement $z_2 = dw$, how would things change?
3. what hidden assumptions are we making?