Fakultas Ilmu Komputer, Universitas Indonesia

Probabilistic Robotics

Probabilities and Bayes Rule

Axel Grossmann

15 August 2007

- Probabilities
- Bayes Rule
- Example

Probabilistic Robotics

Key ideas

- Explicit representation of uncertainty using the calculus of probability theory
- Perception as state estimation
- Action selection as utility optimisation

Axioms of Probability Theory

$\Pr(A)$ denotes probability that proposition A is true.

1. $0 \leq \Pr(A) \leq 1$
2. $\Pr(\mathit{True}) = 1$
   $\Pr(\mathit{False}) = 0$
3. $\Pr(A \vee B) = \Pr(A) + \Pr(B) - \Pr(A \wedge B)$

A Closer Look At Axiom 3

$$\Pr(A \vee B) = \Pr(A) + \Pr(B) - \Pr(A \wedge B)$$

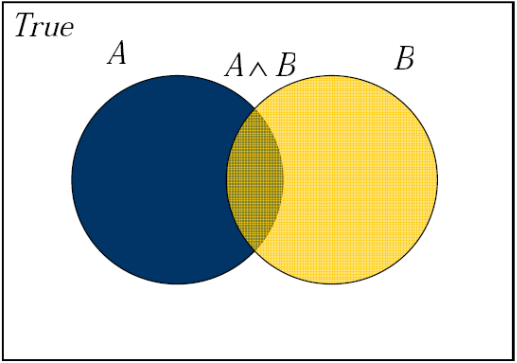

Using the Axioms

$$\begin{aligned}
\Pr(A \vee \neg A) &= \Pr(A) + \Pr(\neg A) - \Pr(A \wedge \neg A) \\
\Pr(\mathit{True}) &= \Pr(A) + \Pr(\neg A) - \Pr(\mathit{False}) \\
1 &= \Pr(A) + \Pr(\neg A) - 0 \\
\Pr(\neg A) &= 1 - \Pr(A)
\end{aligned}$$

Discrete Random Variables

X denotes a random variable.

X can take on a countable number of values in $\{x_1, x_2, \ldots, x_n\}$.

$P(X = x_i)$, or $P(x_i)$, is the probability that the random variable X takes on value x_i.

$P(\cdot)$ is called the probability mass function.
A *probability mass function* (abbreviated PMF) is a function that gives the probability that a discrete random variable is exactly equal to some value.

Example: $P(Room) = \langle 0.7, 0.2, 0.08, 0.02 \rangle$

Continuous Random Variables

X takes on values on the continuum.

$p(X = x)$, or $p(x)$, is a *probability density function* (abbreviated PDF) with

$$\Pr(x \in (a, b)) = \int_a^b p(x)dx$$

$$\int_{-\infty}^{\infty} p(x)dx = 1$$

Example:

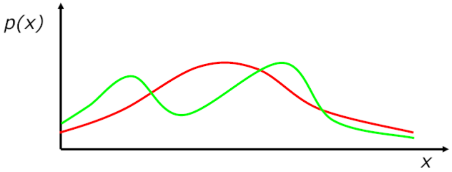

An unimodal and a multimodal case.

Joint and Conditional Probability

$P(X = x \wedge Y = y) = P(x, y)$

If X and Y are independent then

$$P(x, y) = P(x)\ P(y)$$

$P(x|y)$ is the probability of x given y such as

$$P(x|y) = \frac{P(x, y)}{P(y)}$$
$$P(x, y) = P(x|y)\ P(y)$$

If X and Y are independent then

$$P(x|y) = P(x)$$

Law of Total Probability, Marginals

Discrete case	Continuous case
$\sum_x P(x) = 1$	$\int p(x)dx = 1$
$P(x) = \sum_y P(x, y)$	$p(x) = \int p(x, y)dy$
$P(x) = \sum_y P(x\|y)\ P(y)$	$p(x) = \int p(x\|y)\ p(y)dy$

Bayes Theorem (1)

Definition of conditional probability:

$$P(x|y) = \frac{P(x,y)}{P(y)}$$

$$P(y|x) = \frac{P(x,y)}{P(x)}$$

Rearranging and combining these two equations, we find

$$P(x|y)\ P(y) = P(x,y) = P(y|x)\ P(x)$$

Dividing both sides by $P(y)$, we obtain

$$P(x|y) = \frac{P(y|x)\ P(x)}{P(y)}$$

Bayes Theorem (2)

$$P(x|y) = \frac{P(y|x)\ P(x)}{P(y)} \qquad \text{posterior} = \frac{\text{likelihood prior}}{\text{evidence}}$$

$P(x)$ is the *prior probability* or *marginal probability* of x. It is "prior" in the sense that it does not take into account any information about y.

$P(x|y)$ is the conditional probability of x, given y. It is also called the *posterior probability* because it is derived from or depends upon the specified value of y.

$P(y|x)$ is the conditional probability of y given x.

$P(y)$ is the *prior* or *marginal probability* of y, and acts as a *normalising constant*. It is also referred to as *evidence*.

Normalisation

$$P(x|y) = \frac{P(y|x)\ P(x)}{P(y)} = \eta\ P(y|x)\ P(x)$$

$$\eta = P(y)^{-1} = \frac{1}{\sum_x P(y|x)\ P(x)}$$

Algorithm:

- $\forall x : \mathrm{aux}_{x|y} = P(y|x)\ P(x)$
- $\eta = \frac{1}{\sum_x \mathrm{aux}_{x|y}}$
- $\forall x : P(x|y) = \eta\ \mathrm{aux}_{x|y}$

Conditioning

Law of total probability:

$$P(x) = \int P(x, z)dz$$

$$P(x) = \int P(x|z)P(z)dz$$

$$P(x|y) = \int P(x|y, z)P(z|y)dz$$

Bayes Theorem With Background Knowledge

$$P(x|y,z) = \frac{P(y|x,z)\ P(x|z)}{P(y|z)}$$

Conditional Independence

Given x and y are independent:

$$P(x, y|z) = P(x|z)\ P(y|z)$$

equivalent to

$$P(x|z) = P(x|z, y)$$

and

$$P(y|z) = P(y|z, x)$$

Simple Example of State Estimation

Suppose a robot obtains a measurement z.

What is $P(open|z)$?

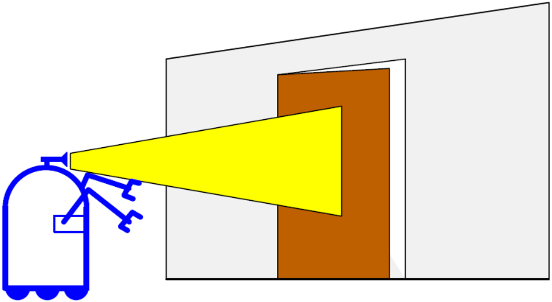

Causal Versus Diagnostic Reasoning

$P(open|z)$ is diagnostic.

$P(z|open)$ is causal.

Often causal knowledge is easier to obtain.

Bayes rule allows us to use causal knowledge:

$$P(open|z) = \frac{P(z|open)\ P(open)}{P(z)}$$

$P(z|open)$ ... count frequencies

Example

$$P(z|open) = 0.6$$
$$P(z|\neg open) = 0.3$$
$$P(open) = P(\neg open) = 0.5$$

$$P(open|z) = \frac{P(z|open)\ P(open)}{P(z|open)\ P(open) + P(z|\neg open)\ P(\neg open)}$$

$$P(open|z) = \frac{0.6 \cdot 0.5}{0.6 \cdot 0.5 + 0.3 \cdot 0.5} = \frac{2}{3} = 0.67$$

z raises the probability that the door is open.

Combining Evidence

Suppose our robot obtains a another observation z_2.

How can we integrate this new information?

More generally, how can we estimate $P(x|z_1, \ldots, z_n)$?

Recursive Bayesian Updating

$$P(x|z_1,\ldots,z_n) = \frac{P(z_n|x,z_1,\ldots,z_{n-1})\ P(x|z_1,\ldots,z_{n-1})}{P(z_n|z_1,\ldots,z_{n-1})}$$

Markov assumption: z_n is independent of $z_1,\ldots,z_{n-1}$ if we know x.

$$P(x|z_1,\ldots,z_n) = \frac{P(z_n|x)\ P(x|z_1,\ldots,z_{n-1})}{P(z_n|z_1,\ldots,z_{n-1})}$$

$$= \eta\ P(z_n|x)\ P(x|z_1,\ldots,z_{n-1})$$

$$= \eta_{1\ldots n}[\prod_{i=1\ldots n} P(z_i|x)]P(x)$$

Example: Second Measurement

$$P(z_2|open) = 0.5$$
$$P(z_2|\neg open) = 0.6$$
$$P(open|z_1) = \frac{2}{3}$$

$$P(open|z_2, z_1) = \frac{P(z_2|open)\ P(open|z_1)}{P(z_2|open)\ P(open|z_1) + P(z_2|\neg open)\ P(\neg open|z_1)}$$

$$= \frac{\frac{1}{2} \cdot \frac{2}{3}}{\frac{1}{2} \cdot \frac{2}{3} + \frac{3}{5} \cdot \frac{1}{3}} = \frac{5}{8} = 0.625$$

z_2 lowers the probability that the door is open.

A Typical Pitfall

Two possible locations x_1 and x_2.

$P(x_1) = 0.99$

$P(z|x_2) = 0.09$ $\qquad$ $P(z|x_1) = 0.07$