Statistical Techniques in Robotics (16-831, F10) Lecture #02 (Wednesday, August 29)

Introduction to Filtering

Lecturer: Drew Bagnell *Scribe: Bryan Hood and Juan Pablo Mendoza*[1]

1 Probability

1.1 Definitions

A Random Variable[2] is a measurable function from the probability space to a measurable space, known as the state space. Usually, random variables are denoted with a capital letter (e.g. X, Y, A ...). For example, a random variable might be the outcome of a coin flip, which takes one of two possible values: heads or tails. For a fair coin, $P(X = head) = P(X = tail) = 0.5$.

In the discrete random variable case, such as a coin flip, the probability mass function (PMF) is a function that assigns a probability value for each specific value of a random variable. Values of the PMF must sum up to 1.

In the case of a continuous random variable, the probability density function (PDF) is the function that represents the relative likelihood of the random variable to take a given value. The probability for the random variable to fall within a particular region is given by the integral of this variable's density over the region: $P(a \leq X \leq b) = \int_a^b p(x)dx$.

The PDF must always integrate to 1: $\int p(x)dx = 1$

Note 1. *For a continuous random variable, the probability that the random variable takes a specific value x is zero, or infinitesimally small. In other words: $P(X = x) = 0$.*

1.2 Probability Axioms

The axioms of probability are:

- $0 <= P(x) <= 1$
- $P(\texttt{true}) = 1$
- $P(\texttt{false}) = 0$
- $P(X \vee Y) = P(X) + P(Y) - P(X \wedge Y)$, (see Figure 1)

From these axioms only we can *prove* several other statements, such as $P(\neg X) = 1 - P(X)$.

The Gaussian PDF is probably the most commonly used PDF. It is parameterized by two values: the mean (μ) and the variance (σ^2). It is defined as:

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}\right) \tag{1}$$

[2] It is not random and it is not a variable; it is easier to think of a random variable as an *unknown quantity.*

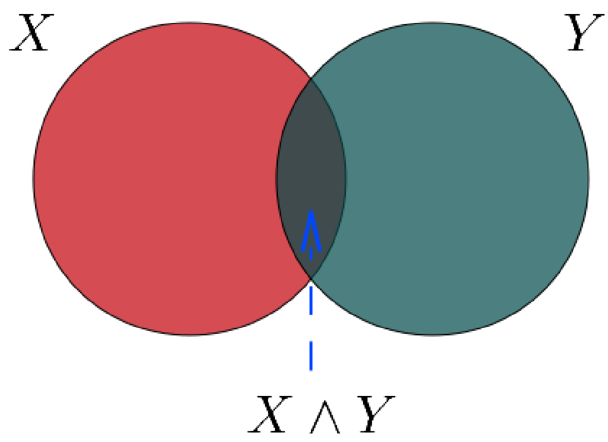

Figure 1: The probability of the event $X\ or\ Y$ is the sum of the probability of X and the probability of Y subtracted by the probability of X and Y. We subtract the union of X and Y to prevent overcounting.

This definition can be generalized for a multidimensional random variable $X \in \mathbb{R}^n$, by replacing the mean with a mean vector and replacing the variance with a symmetric positive semidefinite matrix, called the covariance matrix (Σ). This is called a multivariate Gaussian distribution and can be written as:

$$p(x_1, \ldots, x_n) = \frac{1}{(2\pi)^{n/2}|\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x}-\mu)^T\Sigma^{-1}(\mathbf{x}-\mu)\right), \tag{2}$$

1.3 Expectation and Variance

Given a random variable X, that can take on value x_1 with probability p_1 and x_2 with probablity p_2 and so forth through value x_n and probability p_n, the expectation of X is defined as

$$\mu = E[X] = \sum_{i=1}^{n} x_i p_i \tag{3}$$

The expectation of X is an indicator of the mean or first moment of the random variable.

The variance of a random variable is a measurement of the spread of the values that the random variable can assume and is expressed as

$$Var(X) = E[(X-\mu)^2] \tag{4}$$

$$= E[X^2] - (E[X])^2 \tag{5}$$

1.4 Law of large numbers

The law of large numbers states that, given a set of independent and identically distributed samples from a distribution, the sample average evolves towards the expected value, as the number of samples approaches infinity. For example, as a robot takes more sonar measurements, the average of the samples would converge to the expected value of the sonar measurement. For a random variable X let the sample average be $\bar{x}_n = 1/n(x_1 + \cdots + x_n)$. Then, $\lim_{n\to\infty} P(|\bar{x}_n - \mu| > \epsilon) = 0$.

1.5 Bayes' Rule

The conditional probability of x given y, represented as $P(X = x|Y = y)$ and $P(x|y)$ in shorthand, is the probability of x given that y has occurred. The conditional probability is the probability of the union of events x and y divided by the probability of event y.

$$P(x|y) = P(x, y)/P(y) \tag{6}$$

Bayes' rule provides a way to calculate the posterior ($P(x|y)$) given the likelihood ($P(y|x)$), using the prior ($P(x)$) and the evidence ($P(y)$):

$$P(x|y) = \frac{P(y|x)P(x)}{P(y)} = \frac{P(y|x)P(x)}{\sum_{x'} P(y|x')P(x')} \tag{7}$$

The evidence is also known as the normalization factor and is used to ensure that the posterior distribution integrates to 1. In many cases, η, or $1/Z$ is used to denote the normalizing constant: $P(x|y) = \eta P(y|x)P(x)$.

1.6 Indicator Function

An indicator function is a function defined on a set X that indicates the membership of an element in a subset A.

$$\mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A. \end{cases} \tag{8}$$

2 State Estimation

State estimation is the problem of estimating a system's non directly observable *state* from its observable outputs; the state of the system is a vector of values that fully describe the robot and its world [3].

2.1 Filtering

In filtering, we want to obtain an estimate of state x_t (s_t is also commonly used) over time. For example, in a localization problem x_t is the pose of the robot. In mapping, the state is the map of the world, and in SLAM, the state is both.

Inputs to a general filtering problem are:

Data: Time-indexed set of observations $z_{1:t}$ (also used are $y_{1:t}$ and $o_{1:t}$) and control actions $u_{1:t}$ (also $a_{1:t}$). Generally, the robot observes the environment and performs a control action at every time step. Data can be denoted as: $d = \{z_{1:t}, u_{1:t}\}$.

[3] In reality, the state never truly captures every aspect of the world

Initial probability: The probability distribution of the initial state $p(x_0)$.

Motion model: Also called action model, transition probabilities, Markov kernel, plant and transition kernel; It is a model that relates the probability of the current state to the previous state and the action taken: $p(x_t|x_{t-1}, u_t)$

Sensor model: Also called observation model, measurement model, emmision probabilities; It is a model of how measurements are generated at each state: $p(z_t|x_t)$.

See Figure 2 to see the evolution of control, states and observations.

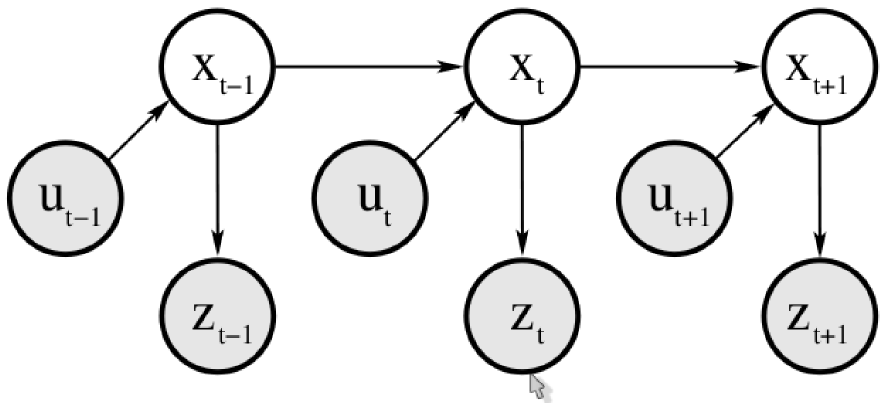

Figure 2: Evolution of controls, states, and observations in a Markovian processs (see Section 2.5). Pay no attention to the pointer to z_t.

2.2 Belief distribution

A belief is a reflection of the robot's internal knowledge about the state and it is represented through conditional probability distributions.

$$bel(x_t) = p(x_t|z_{1:t}, u_{1:t}), \tag{9}$$

which is the posterior probability over state variables conditioned on data (observations and actions).

2.3 The Bayes Filter

Bayes filter is a general algorithm to compute belief from observations and control data. A discrete Bayes filter algorithm is shown in Algorithm 1.

Algorithm 1 Discrete_Bayes_Filter $(Bel(x), d)$

$\eta = 0$
if d is a *perceptual* data item z **then**
  for all x **do**
    $Bel'(x) = P(z|x)Bel(x)$
    $\eta = \eta + Bel'(x)$
  end for
  for all x **do**
    $Bel'(x) = \eta^{-1}Bel'(x)$
  end for
else if d is an *action* data item u **then**
  for all x **do**
    $Bel'(x) = \sum_{x'} P(x|u, x')Bel(x')$
  end for
end if
return $Bel'(x)$

2.4 Derivation of the Bayes filter

Below is the mathematical derivation of the Bayes filter:

$$
\begin{aligned}
Bel(x_t) &= P(x_t|u_1, z_1, \ldots, u_t, z_t) \\
&= \eta P(z_t|x_t, u_1, z_1, \ldots, z_{t-1}, u_t)P(x_t|u_1, z_1, \ldots, z_{t-1}, u_t) && \text{Bayes rule} \\
&= \eta P(z_t|x_t)P(x_t|z_1, u_1, \ldots, z_{t-1}, u_t, x_{t-1}) && \text{Markov (see section 2.5)} \\
&= \eta P(z_t|x_t) \int P(x_t|z_1, u_1, \ldots, z_{t-1}, u_t, x_{t-1})P(x_{t-1}|z_1, u_1 \ldots, z_{t-1}, u_t)dx_{t-1} && \text{Total probability} \\
&= \eta P(z_t|x_t) \int P(x_t|u_t, x_{t-1})P(x_{t-1}|z_1, u_1, \ldots, z_{t-1}, u_{t-1})dx_{t-1} && \text{Markov (removed } u_t\text{)} \\
&= \eta P(z_t|x_t) \int P(x_t|u_t, x_{t-1})Bel(x_{t-1})dx_{t-1} && \text{Def. of } Bel(x_t)
\end{aligned}
$$

2.5 The Markov Assumption

The Markov assumption states that the probability distribution of future states of the robot depends only upon the present state, and not on past ones. While the Markovian world assumption may not be realistic in many situations[4].

For an observation model, the Markov assumption can be expressed as:

$$P(z_t|x_0, x_1, u_1, z_1, \ldots, x_t, u_t, z_{t-1}) = P(z_t|x_t) \tag{10}$$

[4]For example, if a robot's state consists only of its position in the world, but not the time of the day, there could be correlations between subsequent observations that violate the Markovian assumption (e.g., video frames taken during the day will be different than those taken during the night).

And for an action model

$$P(x_{t+1}|x_0, x_1, u_1, z_1, \ldots, x_t, u_t, z_{t-1}, u_{t+1}) = P(x_{t+1}|x_t, u_{t+1}) \tag{11}$$

3 Questions

1) Q: Do measurements always make $p(x|z)$ more concentrated? A: On average, measurements always increase the concentration or certainty of the state, but a specific measurement could decrease the certainty.

2) Q: Do actions always increase uncertainty? A: Generally, yes, but in some cases they reduce uncertainty. For example, if a robot knows there is a wall in front of it somewhere, but does not know its own location, it can drive forward for a long period of time and be certain that it is up against the wall.