Statistical Filtering and Control for AI and Robotics

Part I. Bayes filtering

Riccardo Muradore

UNIVERSITÀ di VERONA
Dipartimento di INFORMATICA

ALTAIR

Outline of the Course

- lesson 1 *Introduction to Probabilistic Robotics; Basics of Probability; Bayes filtering* [R.M.]
- lesson 2 *Basics of Linear methods for Regression; Kalman filtering and applications* [R.M.]
- lesson 3 *Nonparametric filters; Particle filter* [R.M.]

- lesson 4 *Planning and Control: Markov Decision Processes* [A.F.]
- lesson 5 *Exploration and information gathering* [A.F.]
- lesson 6 *Plan monitoring for robotics; Applications for mobile robots* [A.F.]

Outline of this Lesson

Motivation

Basics of probability

Bayes filtering

Motivation

50 years of robotics

Boston dynamics

DARPA challenge

Basics of probability

Why Probabilistic Robotics?

- At the core of probabilistic robotics is the idea of estimating state from sensor data. State estimation addresses the problem of estimating quantities from sensor data that are not directly observable, but that can be inferred.
- Sensors carry only partial information about those quantities, and their measurements are corrupted by noise. State estimation seeks to recover state variables from the data. Probabilistic state estimation algorithms compute belief distributions over possible world states.
- In probabilistic robotics, quantities such as sensor measurements, controls, and the states a robot and its environment might assume are all modeled as random variables.
- Probabilistic inference is the process of calculating these laws for random variables that are derived from other random variables, such as those modeling sensor data.

Reference

This lecture is based on the following book

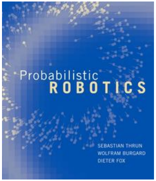

Sebastian Thrun, Wolfram Burgard and Dieter Fox, “Probabilistic Robotics”, MIT Press, 2005

Several pictures from this book have been copied and pasted here

Discrete random variables

Let X be a Discrete random variable, i.e.

$$X \in \mathcal{X} := \{x_1, \dots, x_N\}, \qquad N \text{ is countable}$$

$p(X = x) = p(x)$ probability than X takes the value $x \in \mathcal{X}$

$p(\cdot)$ is called probability mass function, $p(\cdot) \geq 0$

Law of total probability

$$\sum_{x \in \mathcal{X}} p(x) = 1$$

Let X be a Continuous random variable, i.e. X takes on an uncountably infinite number of possible outcomes (support S)

$$P(a < X < b) = \int_a^b p(x)dx, \qquad (a, b) \subset S$$

$p(\cdot)$ is called probability density function (PDF)

Definition (PDF)

The probability density function of a continuous random variable X with support S is an integrable function $p(x)$ such that

1. $p(x)$ is positive everywhere in the support S; $p(x) > 0, \forall x \in S$
2. $p(x)$ satisfies the Law of total probability

$$\int_S p(x)dx = 1$$

3. the probability that $X \in A$, where $A \subseteq S$, is given by $P(X \in A) = \int_A p(x)dx$

Notation

We will ofter refer to the *probability mass function* and to the *probability density function* as probability.

Let X and Y be two random variables, the joint distribution is

$$p(x, y) = p(X = x \text{ and } Y = y)$$

X and Y are independent if

$$p(x, y) = p(X = x)p(Y = y) = p(x)p(y)$$
$$p_{xy}(x, y) = p_x(X = x)p_y(Y = y) = p_x(x)p_y(y)$$

Conditional probability: probability that X has value x conditioned on the fact that Y value is y

$$p(x|y) = p(X = x|Y = y)$$

If $p(y) > 0$, the conditional probability of x given y is

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

if X and Y are independent

$$p(x|y) = p(x)$$

Theorem of Total probability

Discrete random variables

$$p(x) = \sum_{y \in \mathcal{Y}} p(x|y)p(y)$$

Continuous random variables

$$p(x) = \int_{S_y} p(x|y)p(y)dy$$

Discrete random variables

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} \stackrel{(*)}{=} \frac{p(y|x)p(x)}{\sum_{x' \in \mathcal{X}} p(y|x')p(x')} \stackrel{(**)}{=} \eta p(y|x)p(x)$$

Continuous random variables

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} \stackrel{(*)}{=} \frac{p(y|x)p(x)}{\int_{S_x} p(y|x')p(x')dx'} \stackrel{(**)}{=} \eta p(y|x)p(x)$$

(*) $=$ theorem of total probability
(**) $= \eta$ is the normalization symbol

Let us focus on the continuous r.v.

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

- x is the quantity we need to infer from the data y
- $p(x)$ is the prior probability (or a priori probability), i.e. it is the knowledge about x we have before using the information in y
- $p(y)$ is the probability of the measurements y (e.g. how the sensor works)
- $p(x|y)$ is the posterior probability
- $p(y|x)$ is the "inverse" probability. It describes how the x causes the measurement y

Remark 1. if y is independent of x (i.e. if y carries no information about x) we end up with

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

Remark 2. It is possible to condition the Bayes rule on $Z = z$

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

Conditional independence

Let x and y be two independent r.v., we know that

$$p(x, y) = p(x)p(y)$$

What is the meaning of?

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

$\rightarrow$ x and y are conditionally independent on another r.v. $Z = z$.

$\boxed{\text{the r.v. } y \text{ carries no information about the r.v. } x \text{ if } z \text{ is known}}$

Conditional independence

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

is equivalent to

$$\begin{aligned} p(x|z) &= p(x|y, z) \\ p(y|z) &= p(y|x, z) \end{aligned}$$

Pay attention!

Conditional independence does not imply independence

$$p(x, y|z) = p(x|z)p(y|z) \quad \not\Rightarrow \quad p(x, y) = p(x)p(y)$$

Independence does not imply conditional independence

$$p(x, y) = p(x)p(y) \quad \not\Rightarrow \quad p(x, y|z) = p(x|z)p(y|z)$$

Mean and Variance

Let X be a discrete r.v., the expectation (or expected value, or mean) is

$$\mathbb{E}[X] := \sum_{x \in \mathcal{X}} xp(x)$$

The conditional mean of X assuming $\mathcal{M}$ is given by

$$\mathbb{E}[X|\mathcal{M}] := \sum_{x \in \mathcal{X}} xp(x|\mathcal{M})$$

Let X be a continuous r.v., the expectation (or expected value, or mean) is

$$\mathbb{E}[X] := \int_{S_x} xp(x)dx$$

The conditional mean of X assuming $\mathcal{M}$ is given by

$$\mathbb{E}[X|\mathcal{M}] := \int_{S_x} xp(x|\mathcal{M})dx$$

If $\mathcal{M} = \{Y = y\}$ then

$$\mathbb{E}[X|y] := \int_{S_x} xp(x|y)dx$$

Theorem

Given the r.v. X and a function $g(\cdot)$, the mean of the random variable $Y = g(X)$ is

$$\mathbb{E}[Y] = \int_{S_x} g(x)p(x)dx$$

Theorem (Linearity)

$\mathbb{E}[a_1 g_1(X) + \ldots + a_N g_N(X)] = a_1\mathbb{E}[g_1(X)] + \ldots + a_N\mathbb{E}[g_N(X)]$

($\mathbb{E}[aX + b] = a\mathbb{E}[X] + b$)

Let X be a discrete r.v. with mean $\mu = \mathbb{E}[X]$, the variance σ^2 is

$$\sigma^2 := \mathbb{E}[(X-\mu)^2] = \sum_{x \in \mathcal{X}} (x-\mu)^2 P(x)$$

Let X be a continuous r.v. with mean $\mu = \mathbb{E}[X]$, the variance σ^2 is

$$\sigma^2 := \mathbb{E}[(X-\mu)^2] = \int_{S_x} (x-\mu)^2 p(x) dx$$

The following relationship holds

$$\sigma^2 = \mathbb{E}[(X-\mu)^2] = \mathbb{E}[X^2] - \mathbb{E}^2[X]$$

σ is called standard deviation

Let X and Y be two r.v. with mean $\mu_x = \mathbb{E}[X]$ and $\mu_y = \mathbb{E}[Y]$, respectively. The covariance of X and Y is by definition the number

$$\Sigma_{xy} = \mathbb{E}[(X - \mu_x)(Y - \mu_y)].$$

The following relationship holds

$$\Sigma_{xy} = \mathbb{E}[(X - \mu_x)(Y - \mu_y)] = \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$$

The correlation coefficient r is the ratio

$$r_{xy} = \frac{\Sigma_{xy}}{\sigma_x \sigma_y}$$

with $|r_{xy}| \leq 1$

Remark. the r.v. X, Y and $X - \mathbb{E}[X]$, $Y - \mathbb{E}[Y]$ have the same covariance and correlation coefficient

Uncorrelation and Orthogonality

Definition

Two r.v. X, Y are uncorrelated if their covariance is zero, i.e.

$$\Sigma_{xy} = 0 \quad \Leftrightarrow \quad r_{xy} = 0 \quad \Leftrightarrow \quad \mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$$

Definition

Two r.v. X, Y are orthogonal ($X \perp Y$) if

$$\mathbb{E}[XY] = 0$$

Uncorrelation and Orthogonality

Observations:

- if X and Y are uncorrelated, then $X - \mu_x$ and $Y - \mu_y$ are orthogonal

$$X - \mu_x \perp Y - \mu_y,$$

- if X and Y are uncorrelated and $\mu_x = 0$ and $\mu_y = 0$, then $X \perp Y$,
- if X and Y are independent, then they are uncorrelated (the converse is false),
- if X and Y are Gaussian and uncorrelated, then they are independent,
- if X and Y are uncorrelated with mean μ_x, μ_y and variance σ_x^2, σ_y^2, then the mean and the variance of the r.v. $Z = X + Y$ are

$$\begin{aligned} \mu_z &= \mu_x + \mu_y \\ \sigma_z^2 &= \sigma_x^2 + \sigma_y^2 \end{aligned}$$

Conditional Mean and Variance

We already introduced the conditional mean of the r.v. X assuming $Y = y$

$$\mu_{x|y} = \mathbb{E}[X|y] = \int_{S_x} xp(x|y)dx$$

We can also define the conditional variance of the r.v. X assuming $Y = y$

$$\sigma^2_{x|y} = \mathbb{E}[(X - \mu_{x|y})^2|y] = \int_{S_x} (x - \mu_{x|y})^2 p(x|y)dx$$

Observations:

- $\mathbb{E}[g(X, Y)|y] = \int_{S_x} g(x, y)p(x|y)dx \qquad = \mathbb{E}[g(X, y)|y]$
- $\mathbb{E}\left[\mathbb{E}[X|y]\right] = \mathbb{E}[X]$

Conditional Mean and Variance

Is there any difference between $\mathbb{E}[X|y]$ and $\mathbb{E}[X|Y]$?

YES!!!

$\varphi(y) = \mathbb{E}[X|y]$ is a function of y

whereas

$\varphi(Y) = \mathbb{E}[X|Y]$ is a random variable

Observations:

- $\mathbb{E}\left[\mathbb{E}[X|Y]\right] = \mathbb{E}[X]$
- $\mathbb{E}\left[\mathbb{E}[g(X, Y)|Y]\right] = \mathbb{E}[g(X, Y)]$

Mobile robots

We mainly focus on mobile robots.

- robot acquires information about the surrounding environment by analyzing the data/measurements collected by its on-board sensors (cameras, laser scanners, bumpers, odometers, GPS)

$$\boxed{\text{data} \longrightarrow \text{elaboration} \longrightarrow \text{information}}$$

  Elaboration means to estimate things that are not directly measured and/or to clean noisy measurements

- robot interacts with the environment through its actuators. The robot 'changes' the environment that must be estimated continuously

---

We are drowning in information and starving for knowledge.

–Rutherford D. Roger

Examples of Unmanned Vehicles

UAV, UGV, etc






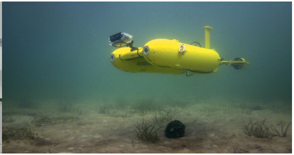

Mobile robots: Sensors

Stereo Camera, RGBD Camera, Infrared sensor, Laser scanner, Odometer, GPS, Force sensor

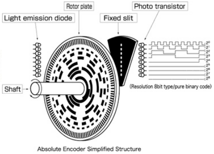

Mobile robots

Sensors

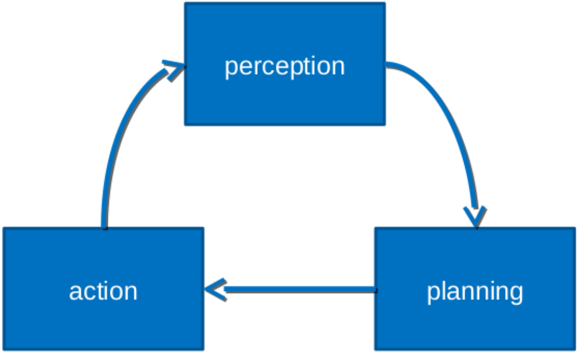

Actuators
(motors in mobile robots
and manipulators)

Learning
Data elaboration (e.g. filtering)
Make decisions
Control architecture

In this course we call state the collection of all information that describes the robot AND the environment. We indicate the state with x, or x_t if it is important to highlight the time.

- dynamic state: all the states that change over the time.
  E.g. velocity of the robot, position of a moving obstacle, status of some object within the scene (open door vs closed door), ...
- static state: all the states that do not change over the time.
  E.g. position of a door/wall, static obstacles, dimension of the robot

Important examples:

- POSE: actual position and orientation of the robot
- LANDMARKS: stationary features of the environment that can be identified and used to construct a map and then to plan collision-free trajectories to reach a target (pre-defined or decided at run-time)

The state x_t is complete if it contains all the knowledge needed to predict the future. Past measurements $(z_{t-1}, z_{t-2}, \ldots)$, past states $(x_{t-1}, x_{t-2}, \ldots)$, past commands $(u_{t-1}, u_{t-2}, \ldots)$ do not carry additional information to predict the future more accurately.

The state at time t can be seen as the value of a process X_t

Definition (Markov process)

A Markov process is a stochastic process whose past has no influence on the future if its present is specified.

A complete state means that the process X_t is Markovian.

The goal of probabilistic robotics is to provide tools to compute "the best" estimation of x_t given the available measurements till time t

Even though X_t is defined for $t \in \mathbb{R}$ our information update has a discrete nature due to the way we collect measurements from the sensors. $\rightarrow$ the motion of the mobile robot is continuous (it is describe by differential equations) whereas, for example, the odometer gives a new measurement every T_s seconds.

From now on the state x_t will be updated with $t \in \mathbb{Z}$, i.e. $x_t, x_{t-1}, \ldots, x_{t-N}$. X_t is a discrete time Markov process. However, some element of the state can take value in $\mathbb{R}$.

Properties of Markov processes

- $p(x_n | x_{n-1}, x_{n-2}, \ldots, x_1) = p(x_n | x_{n-1})$
- $\mathbb{E}[X_n | X_{n-1}, X_{n-2}, \ldots, X_1] = \mathbb{E}[X_n | X_{n-1}]$

Notation

- Environment Measurement data z_t provides information about the actual state of the environment. E.g. camera images, laser scanner measurements.

$$
\begin{aligned}
z_t &: \text{ measurement data at time } t \\
z_{t_1:t_2} = \{z_{t_1}, z_{t_1+1}, \ldots, z_{t_2}\} &: \text{ measurement data from time } t_1 \text{ to } t_2 \geq t_1
\end{aligned}
$$

- Control data u_t carries information about the change of state in the environment. E.g. robot velocity

$$
\begin{aligned}
u_t &: \text{ control data at time } t\text{, i.e.} \\
& \text{ change of state in the interval } (t-1, t] \\
u_{t_1:t_2} = \{u_{t_1}, u_{t_1+1}, \ldots, u_{t_2}\} &: \text{ control data from time } t_1 \text{ to } t_2 \geq t_1
\end{aligned}
$$

If x_t is complete (i.e. the process X_t is a Markov process) the following equalities hold

- state transition probability: how past states, past measurements, and past and actual commands change the actual state

  $$p(x_t | x_{0:t-1}, z_{1:t-1}, u_{1:t}) = p(x_t | x_{t-1}, u_t)$$

  ($\rightarrow$ process equation)

- measurement probability: how past states, past measurements, and past and actual commands influence the actual measurement

  $$p(z_t | x_{0:t}, z_{1:t-1}, u_{1:t}) = p(z_t | x_t)$$

  ($\rightarrow$ measurement equation)

These equalities are examples of conditional independence:
the state x_t is sufficient to predict the (potentially noisy) measurement z_t.
Knowledge of any other variable, such as past measurements,
controls or even past states, is irrelevant if x_t is complete

Complete state ⇔ Markov process

Conditional independence is the main source
of tractability of probabilistic robotics algorithms

It is enough to store the last value of the state

The state transition probability and the measurement probability describe the dynamical stochastic system of the robot and the environment

Dynamic Bayes network

Dynamic Bayes network (DBN) or Hidden Markov Model (HMM) of our dynamical stochastic system

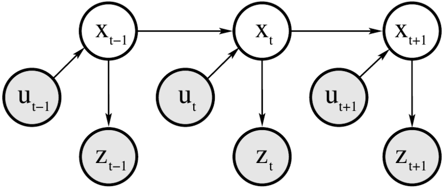

state transition probability $p(x_t|x_{t-1}, u_t)$

measurement probability $p(z_t|x_t)$

A key concept in probabilistic robotics is that of belief:

- x_t is the true state of the environment at time t
- $bel(x_t)$ is the robot knowledge about the state of the environment at time t based on past and actual measurement $z_{1:t}$, and past and actual commands $u_{1:t}$

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

  $p(x_t | z_{1:t}, u_{1:t})$ is the posterior probability

- $\overline{bel}(x_t)$ is the robot knowledge about the state of the environment at time t based on past measurement $z_{1:t-1}$, and past and actual commands $u_{1:t}$

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

  $p(x_t | z_{1:t-1}, u_{1:t})$ is the probability before incorporating z_t (i.e. prior probability)

Bayes filtering

Bayer recursive algorithm

Algorithm 1 $bel(x_t)$ =BayerFilter$(bel(x_{t-1}), u_t, z_t)$

1: **forall** $x_t \in S_x$ **do**
2: $\quad \overline{bel}(x_t) = \int_{S_x} p(x_t|u_t, x_{t-1}) bel(x_{t-1}) dx_{t-1}$
3: $\quad bel(x_t) = \eta p(z_t|x_t) \overline{bel}(x_t)$
4: **end**
5: **return** $bel(x_t)$

the algorithm updates recursively the belief distribution $bel(x_t)$ by integrating the actual control data u_t (step 2) and the new measurement z_t (step 3) starting from the previous belief distribution $bel(x_{t-1})$

Algorithm 2 $bel(x_t)$ =BayerFilter($bel(x_{t-1}), u_t, z_t$)

1: **forall** $x_t \in S_x$ **do**
2: $\overline{bel}(x_t) = \int_{S_x} p(x_t|u_t, x_{t-1}) bel(x_{t-1}) dx_{t-1}$
3: $bel(x_t) = \eta p(z_t|x_t)\overline{bel}(x_t)$
4: **end**
5: **return** $bel(x_t)$

- $\overline{bel}(x_t)$ is the prediction in this two-step statistical filtering (computed using the old $bel(x_{t-1})$ and the current controls u_t)

Bayer recursive algorithm

Algorithm 3 $bel(x_t)$ =BayerFilter($bel(x_{t-1}), u_t, z_t$)

1: **forall** $x_t \in S_x$ **do**
2: $\quad \overline{bel}(x_t) = \int_{S_x} p(x_t|u_t, x_{t-1}) bel(x_{t-1}) dx_{t-1}$
3: $\quad bel(x_t) = \eta p(z_t|x_t) \overline{bel}(x_t)$
4: **end**
5: **return** $bel(x_t)$

- $bel(x_t)$ is the measurement update in this two-step statistical filtering (computed integrating $\overline{bel}(x_t)$ and the new measurements z_t)

In the book "Probabilistic Robotics", the authors report the following Assumptions:

- the state x_t is complete
- u_t are chosen at random

before deriving the Bayer filter

The second assumption does not hold when we want to decide u_t to fulfill some well defined goal.

We will see that it is enough to ask that u_t is a function of past measurements/states

Prediction step

$$
\begin{aligned}
\overline{bel}(x_t) &= p(x_t|z_{1:t-1}, u_{1:t}) \\
&\overset{\text{th total prob.}}{=} \int_{S_x} p(x_t|x_{t-1}, z_{1:t-1}, u_{1:t})p(x_{t-1}|z_{1:t-1}, u_{1:t})dx_{t-1} \\
&\overset{\text{Markov}}{=} \int_{S_x} p(x_t|x_{t-1}, u_t)p(x_{t-1}|z_{1:t-1}, u_{1:t})dx_{t-1} \\
&\overset{u_t \text{ random}}{=} \int_{S_x} p(x_t|x_{t-1}, u_t)p(x_{t-1}|z_{1:t-1}, u_{1:t-1})dx_{t-1} \\
&= \int_{S_x} p(x_t|x_{t-1}, u_t)bel(x_{t-1})dx_{t-1} \qquad \boxed{\text{line 2: prediction}}
\end{aligned}
$$

Measurement update

$$
\begin{aligned}
bel(x_t) &= p(x_t|z_{1:t}, u_{1:t}) \\
&\overset{\text{Bayes rule}}{=} \frac{p(x_t|z_{1:t-1}, u_{1:t})\,p(z_t|x_t, z_{1:t-1}, u_{1:t})}{p(z_t|z_{1:t-1}, u_{1:t})} \\
&= \eta\, p(x_t|z_{1:t-1}, u_{1:t})\,p(z_t|x_t, z_{1:t-1}, u_{1:t}) \\
&\overset{\text{Markov}}{=} \eta\, p(x_t|z_{1:t-1}, u_{1:t})\,p(z_t|x_t) \\
&= \eta\, p(z_t|x_t)\,\overline{bel}(x_t)
\end{aligned}
$$

line 3: measurement update

Open questions

- How strong the Markov assumption is (i.e. x_t complete)?

  [we have always to deal with noisy measurements and uncertain model]

- When can the a priori and posteriori probabilities ($\overline{bel}(x_t)$ and $bel(x_t)$) be really computed?

- What about the computational efficiency?