UNIVERSITÉ DE TECHNOLOGIE DE BELFORT-MONTBÉLIARD

SLAM

RO51 - Introduction to Mobile Robotics

Zhi Yan

April 4, 2023

https://yzrobot.github.io/

www.utbm.fr

What?

- SLAM stands for Simultaneous Localization And Mapping.
- A fundamental problem in mobile robotics, which studies the problem of constructing a map of an unknown environment while simultaneously keeping track of a robot's location within it.
  - Goal: build a map (because we don't have it).
  - Problem: robot needs (always) know where I am (need a map).
  - => A chicken-and-egg problem!
- An operational difference from exploration (c.f. next lecture):
  - SLAM: the movement of the robot is usually teleoperated by a human.
  - Exploration: the robot moves autonomously, which can be regarded as SLAM plus decision-making problem (i.e. where to go).

Why?

- Taking humans as an example:
  - When you first came to UTBM campus, you didn't know the location of Building A, B, etc.
  - So you wandered around, and in the process you discovered and remembered the appearance (features) of the campus.
    => You get a very abstract map in your mind.
  - When you step into the UTBM campus again and need to go to Building A, you are no longer confused, you will choose the best route (usually the shortest) to reach your destination (to avoid being late).
- So for a mobile robot, it also needs a map of the environment in which it will work (e.g. robot vacuum).

How?

- Theoretical and conceptual level (can be considered a solved problem): Bayesian formulation (i.e. probabilistic robotics)
- Practical and computational level (still an open problem): extended Kalman filter (EKF-SLAM), Rao-Blackwellized particle filters (FastSLAM), and more.

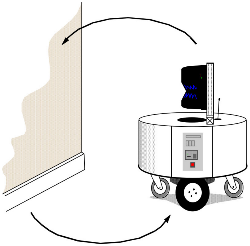

A brief chronicle before the formal discussion

- Before 1986: the stone-age of mobile robotics, most were rule-based approaches.
- 1986: Genesis discussion on probabilistic methods at the International Conference on Robotics and Automation (ICRA).
- 1987-1988: Key papers established a statistical basis for describing relationships between landmarks and manipulating geometric uncertainty.
- 1995: SLAM acronym coined at the International Symposium of Robotic Research (ISRR).
- 1995-1999: Convergence proofs and first demonstrations of real systems.
- 2000-present: Wide interest in SLAM started.

Preliminaries

- The essential SLAM problem:
  - A simultaneous estimate of both robot and landmark locations is required.
  - The true locations are never known or measured directly.
  - Observations are made between true robot and landmark locations.

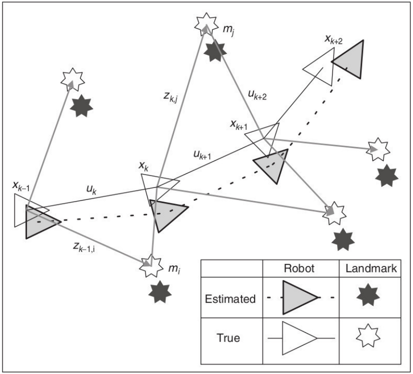

Why landmarks matter?

Why probabilistic methods rule?

- Remember that there will always be errors in sensor readings?
- If a deterministic approach is used, the robot will eventually get lost due to the accumulation of errors!

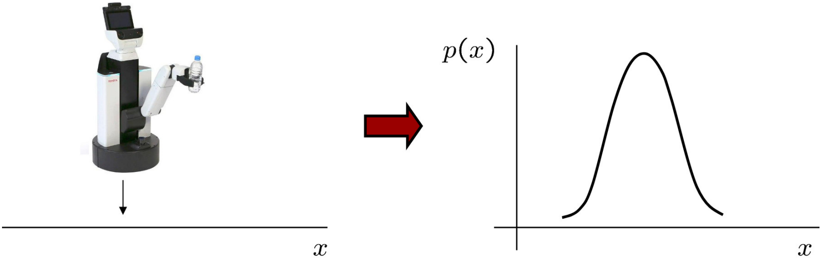

"The robot is exactly here"

"The robot is somewhere here"

Probabilistic robotics

- Face the inevitable uncertainty in mobile robotics, in particular the robot's motions (action) and observations (perception).
- Use the probability theory to explicitly represent the uncertainty:
    - perception = state estimation
    - action = utility optimization

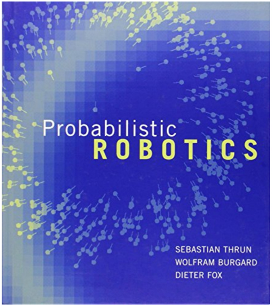

Sebastian Thrun

Wolfram Burgard

Dieter Fox

Probability theory

Axioms

- $P(A)$ denotes probability that proposition A is true:
  1. $0 \leq P(A) \leq 1$
  2. $P(True) = 1, P(False) = 0$
  3. $P(A \cup B) = P(A) + P(B) - P(A \cap B)$
- A closer look at Axiom 3:

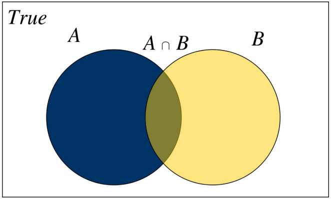

- Using the axioms:
  - $P(A \cup \neg A) = P(A) + P(\neg A) - P(A \cap \neg A)$
  - $P(True) = P(A) + P(\neg A) - P(False)$
  - $1 = P(A) + P(\neg A) - 0$
  - $P(\neg A) = 1 - P(A)$

Probability theory

Random variable

- **Discrete random variable**:
  - X denotes a random variable and 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.
  - $P(X)$ is called probability mass function (pmf).
  - For example, $P(Room) =< 0.7, 0.2, 0.08, 0.02 >$
- **Continuous random variable**:
  - X takes on values in the continuum.
  - $p(X = x)$ or $p(x)$ is a probability density function (PDF), and $p(x \in [a, b]) = \int_a^b p(x)dx$.
  - For example, 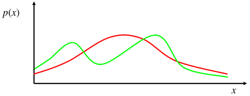

- **Probability sums up to 1:**
  - Discrete case: $\sum_x P(X) = 1$
  - Continuous case: $\int p(x)dx = 1$

Probability theory

Joint probability

- The probability of two events co-occurring, denoted as $P(A \cap B)$, or even simpler, $P(A, B)$
- If A and B are independent (i.e. not influence each other), then $P(A \cap B) = P(A) \times P(B)$
  - The probability of clouds in the sky has an impact on the probability of rain that day. => dependent events
  - The probability of getting a head on the first coin toss does not have an impact on the probability of getting heads on the second coin toss. => independent events
- Time for questions:
  - Q1: What is the probability of drawing a red card from a deck of cards?
  - Q2: What is the joint probability of drawing a number 6 card that is red?

Probability theory

Joint probability

- For Q2, we first confirm that the question is correct, as the events "6" and "red" are independent.
- Then two ways of thinking:
  - A deck of cards has two red 6 (hearts and diamonds), thus $P(6 \cap red) = 2/52 = 1/26$
  - The probability of drawing a 6 $= 4/52$, while the probability of drawing a red card $= 26/52$, thus $P(6 \cap red) = P(6) \times P(red) = 4/52 \times 26/52 = 1/26$

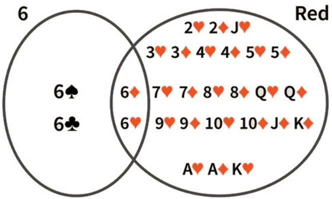

Probability theory

Conditional probability

- The probability that A occurs under the conditions that B occurs, denoted as $P(A, \textit{given } B)$, or even simpler, $P(A \mid B)$.
- Q3: What is the probability of drawing a 6 from a deck of cards, given that you drew a red one?
  - $P(6 \mid red) = 2/26 = 1/13$, since there are two 6 out of 26 red cards.
- From Q3 to Q2:
  - Difference: "given that you drew a red one", meaning $P(red)$.
  - Connection: $P(6 \mid red) \times P(red)$ actually give us $P(6 \cap red)$.
- Therefore,

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}$$

- What about if A and B are independent?

Probability theory

Marginalization (given $P(A, B)$)

- **Discrete case**:
  - $P(A) = \sum_B P(A, B) = \sum_B P(A \mid B)P(B)$
  - Sum all the values that can take B.
  - Probability related only to a single random variable.
  - Dimensionality reduction is actually performed.
- **Continuous case**: $P(A) = \int_{-\infty}^{\infty} P(A, X = x)dx$

Probability theory

Law of total probability

- **Discrete case**:
  - $P(A) = \sum_n P(A \cap B_n)$ or alternatively,
    $P(A) = \sum_n P(A \mid B_n)P(B_n)$ where,
    $\{B_n : n = 1, 2, 3, \dots\}$ and must be measurable.
  - In case of conditional probability:
    $P(A \mid C) = \sum_n P(A \mid C \cap B_n)P(B_n \mid C)$, and in case of B_n and C being independent:
    $P(A \mid C) = \sum_n P(A \mid C \cap B_n)P(B_n)$
  - The above formulation in human language: *Given an event A, with known conditional probabilities given any of the B_n events, each with a known probability itself, what is the total probability that A will happen?*
- **Continuous case**:
  - $P(A) = \int_{-\infty}^{\infty} P(A \mid X = x) f_X(x) dx$

Probability theory

Bayes' theorem (for Thomas Bayes)

- Let's worship first.

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)} = \frac{\textbf{likelihood} \cdot \textbf{prior}}{\textbf{evidence}}$$

Probability theory

Bayes' theorem (for Thomas Bayes)

- A type of inference in inferential statistics: A fourth probability is derived from the three known probabilities. => Very important and useful in mobile robotics.

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}$$

- A and B are different events and $P(B) \neq 0$.
- $P(A \mid B)$: the probability A occurring given that B is true (also called the posterior probability of A given B).
- $P(B \mid A)$: the probability of event B occurring given that A is true (also be interpreted as the likelihood of A given a fixed B as $P(B \mid A) = L(A \mid B)$, L for *Likelihood*).
- $P(A)$ and $P(B)$: the probabilities of observing A and B respectively without any given conditions (also known as the marginal probability or prior probability).

Probability theory

Bayes' theorem (for Thomas Bayes)

- Note the difference between probability and likelihood:
  - **Probability**: predict the outcome of the next observation given some parameters.
  - **Likelihood**: estimate the parameters of the relevant event when the results obtained by some observations are known, that is, to guess the relevant parameters after an event has been observed.

Probability theory

Bayes' theorem (for Thomas Bayes)

- Derivation of Bayesian Theory from Conditional Probabilities:
  - According to the definition of conditional probability, the probability that event A occurs under the condition that event B occurs is:

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}, \text{ if } P(B) \neq 0$$

  - where $P(A \cap B)$ is the probability of both A and B being true. Similarly,

$$P(B \mid A) = \frac{P(A \cap B)}{P(A)}, \text{ if } P(A) \neq 0$$

  - Solving for $P(A \cap B)$ and substituting into the above expression for $P(A \mid B)$ yields Bayes' theorem:

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}, \text{ if } P(B) \neq 0$$

Why the Bayesian approach?

A minimalist example

- Suppose a robot obtains measurement z.
- What is $P(open \mid z)$?

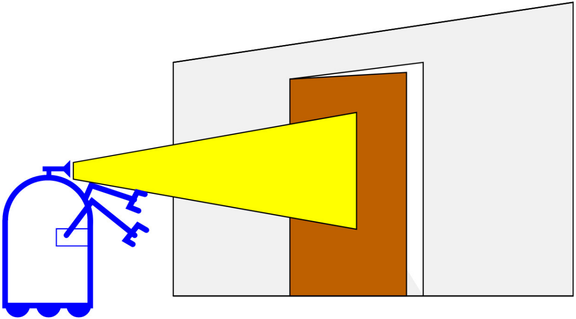

Why the Bayesian approach?

A minimalist example

- According to factory testing, the sensor is 60% accurate for "open" while 30% for "not open". => likelihood obtained through multiple tests.
- Deterministic (rule-based) approach:
  - We just trust whatever the sensor tells you, turning a blind eye to the existence of sensing errors.
  - The errors don't just disappear out of thin air, but accumulate.
- Probabilistic approach:
  - $P(z \mid open) = 0.6$, $P(z \mid \neg open) = 0.3$
  - $P(open) = 0.5$, since "open" is a binary event. => priori possessed by humans.
  - One last value left:

$$\begin{aligned} P(z) &= P(z \cap open) + P(z \cap \neg open) \\ &= P(z \mid open)P(open) + P(z \mid \neg open)P(\neg open) \end{aligned}$$

=> evidence obtained by applying the Law of total probability.

Why the Bayesian approach?

A minimalist example

- So, according to Bayes' theorem:

$$
\begin{aligned}
P(open \mid z) &= \frac{P(z \mid open)P(open)}{P(z \mid open)P(open) + P(z \mid \neg open)P(\neg open)} \\
&= \frac{0.6 \times 0.5}{0.6 \times 0.5 + 0.3 \times 0.5} \\
&= 0.67
\end{aligned}
$$

- Therefore, z raises the probability that the door is open. In other words, even if the sensor tells you that the door is open, the probability that the door is open is not 100%.

Preliminaries

Consider a mobile robot moving through an environment taking relative observations of a number of unknown landmarks using a sensor located on the robot.

- At a time instant k, the following quantities are defined:
  - $\boldsymbol{x}_k$: the state vector describing the location and orientation of the robot.
  - $\boldsymbol{u}_k$: the control vector, applied at time $k-1$ to drive the robot to a state $\boldsymbol{x}_k$ at time k.

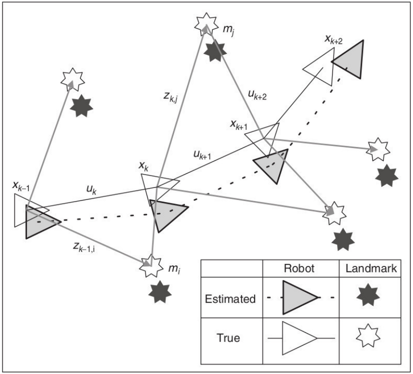

Preliminaries

- $\boldsymbol{m}_i$: a vector describing the location of the ith landmark whose true location is assumed time invariant.
- $\boldsymbol{z}_{ik}$: an observation taken from the robot of the location of the ith landmark at time k. When there are multiple landmark observations at any one time or when the specific landmark is not relevant to the discussion, the observation will be written simply as $\boldsymbol{z}_k$.

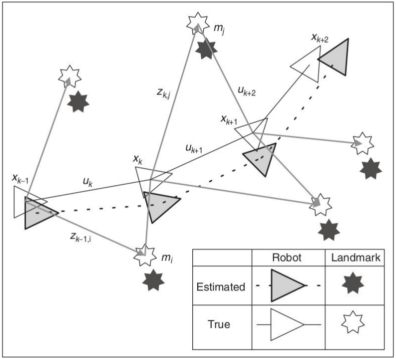

Preliminaries

In addition, the following sets are also defined:

- $\boldsymbol{X}_{0:k} = \{\boldsymbol{x}_0, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_k\} = \{\boldsymbol{X}_{0:k-1}, \boldsymbol{x}_k\}$: the history of robot locations.
- $\boldsymbol{U}_{0:k} = \{\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_k\} = \{\boldsymbol{U}_{0:k-1}, \boldsymbol{u}_k\}$: the history of control inputs.
- $\boldsymbol{m} = \{\boldsymbol{m}_1, \boldsymbol{m}_2, \ldots, \boldsymbol{m}_n\}$: the set of all landmarks (i.e. the map).
- $\boldsymbol{Z}_{0:k} = \{\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_k\} = \{\boldsymbol{Z}_{0:k-1}, \boldsymbol{z}_k\}$: the set of all landmark observations.

Probabilistic SLAM

- SLAM in probabilistic form:

$$P(\boldsymbol{x}_k, \boldsymbol{m} \mid \boldsymbol{Z}_{0:k}, \boldsymbol{U}_{0:k}, \boldsymbol{x}_0)$$

=> This probability distribution describes the joint posterior density of the landmark locations and vehicle state (at time k) given the recorded observations and control inputs up to and including time k together with the initial state of the robot.

- In general, a recursive solution to the SLAM problem is desirable, starting with:

$$P(\boldsymbol{x}_{k-1}, \boldsymbol{m} \mid \boldsymbol{Z}_{0:k-1}, \boldsymbol{U}_{0:k-1})$$

at time $k-1$, and so on.

- An observation model and a motion model are needed to describe the effect of the observation and control input and respectively.

Probabilistic SLAM

- The **observation model** describes the probability of making an observation $\boldsymbol{z}_k$ when the robot location and landmark locations are known and is generally described in the form:

$$P(\boldsymbol{z}_k \mid \boldsymbol{x}_k, \boldsymbol{m})$$

It is reasonable to assume that once the robot location and map are defined, observations are conditionally independent given the map and the current robot state.

- The **motion model** for the robot can be described in terms of a probability distribution on state transitions in the form:

$$P(\boldsymbol{x}_k \mid \boldsymbol{x}_{k-1}, \boldsymbol{u}_k)$$

The next state $\boldsymbol{x}_k$ depends only on the immediately preceding state $\boldsymbol{x}_{k-1}$ and the applied control $\boldsymbol{u}_k$ and is independent of both the observations and the map (i.e. a Markov process).

Probabilistic SLAM

The SLAM algorithm is now implemented in a standard two-step recursive (sequential) prediction (time update, with the motion model) correction (measurement update, with the observation model) form:

- **Prediction**:

$$P(\boldsymbol{x}_k, \boldsymbol{m} \mid \boldsymbol{Z}_{0:k-1}, \boldsymbol{U}_{0:k}, \boldsymbol{x}_0) = \int P(\boldsymbol{x}_k \mid \boldsymbol{x}_{k-1}, \boldsymbol{u}_k) P(\boldsymbol{x}_{k-1}, \boldsymbol{m} \mid \boldsymbol{Z}_{0:k-1}, \boldsymbol{U}_{0:k-1}, \boldsymbol{x}_0) d\boldsymbol{x}_{k-1}$$

- **Correction**:

$$P(\boldsymbol{x}_k, \boldsymbol{m} \mid \boldsymbol{Z}_{0:k}, \boldsymbol{U}_{0:k}, \boldsymbol{x}_0) = \frac{P(\boldsymbol{z}_k \mid \boldsymbol{x}_k, \boldsymbol{m}) P(\boldsymbol{x}_k, \boldsymbol{m} \mid \boldsymbol{Z}_{0:k-1}, \boldsymbol{U}_{0:k}, \boldsymbol{x}_0)}{P(\boldsymbol{z}_k \mid \boldsymbol{Z}_{0:k-1}, \boldsymbol{U}_{0:k})}$$

Solutions to the SLAM Problem

- EKF-SLAM, FastSLAM, and more.

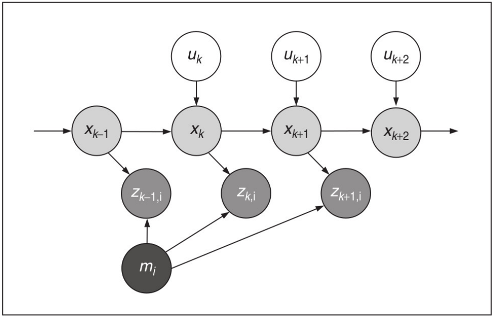

A graphical model of the SLAM algorithm. If the history of pose states are known exactly then, since the observations are conditionally independent, the map states are also independent. For FastSLAM, each particle defines a different robot trajectory hypothesis.

Relaxation before the end

SLAM Applications

SLAM is central to a range of indoor, outdoor, in-air and underwater applications for both manned and autonomous vehicles. For example,

- at home: vacuum cleaner, lawn mower => video
- air: surveillance with unmanned air vehicles => video
- underwater: reef monitoring => video
- underground: exploration of mines => video
- space: terrain mapping for localization => video

Summary

- SLAM, what, why and how.
- Probabilistic SLAM.
- Further reading on Bayesian theory: normalization, rule with background knowledge, conditional independence, etc.
- The practice of SLAM will be covered in TD and TP.

The end

Thank you for your attention!

Any questions?