

Robotics

Lecture 4: Probabilistic Robotics

See course website


<http://www.doc.ic.ac.uk/~ajd/Robotics/> for up to date information.

Andrew Davison
Department of Computing
Imperial College London

Review: Sensors Practical from Lecture 3

- Bump sensor (returns yes/no state): use to detect collision and trigger avoidance action.
- Sonar sensor (returns depth value in cm): can be used for smooth *servoing* behaviour with proportional gain.
- Both are examples of negative feedback.

Review: Wall Following with Sonar



- Use sideways-looking sonar to measure distance z to wall.
- Use velocity control and a while loop at for instance 20Hz.
- With the goal of maintaining a desired distance d , set difference between left and right wheel velocities proportional to difference between z and d :


$$v_R - v_L = K_p(z - d)$$

Symmetric behaviour can therefore be achieved using a constant offset v_C :


$$v_R = v_C + \frac{1}{2}K_p(z - d)$$

$$v_L = v_C - \frac{1}{2}K_p(z - d)$$

Review: Wall Following with Sonar




- Problem if angle between robot's direction and wall gets too large because sonar doesn't measure the perpendicular distance.
- Solutions: ring of sonar sensors would be most straightforward. Clever combination of measurements from different times?
- A simple way to a better is to mount the sonar forward of the wheels. This couples rotation and the distance from the wall:



Probabilistic Localisation


Over the next two weeks we will aim beyond local reactive behaviours towards reliable long-term navigation, via *localisation* relative to a map.



Probabilistic Robotics


- Problem: simple sensing/action procedures can be locally effective but are limited in complicated problems in the real-world; we need longer-term representations and consistent scene models.
- 'Classical AI' approaches to long-term estimation based on logical reasoning about true/false statements fall down when presented with real-world data.
- Why?
 - Advanced sensors don't lend themselves to straightforward analysis like bump and light sensors.
 - All information which a robot receives is uncertain.
- A probabilistic approach acknowledges uncertainty and uses models to abstract useful information from data.
- Our goal is an incrementally updated probabilistic estimate of the position of the robot relative to the map.

Uncertainty in Robotics



- Every robot action is uncertain.
- Every sensor measurement is uncertain.
- When we combine actions and measurements and want to estimate the state of a robot, the state estimate will be uncertain.
- Usually, we will start with some uncertain estimate of the state of a robot, and then take some action and receive new information from sensors. We need to update the uncertain state estimate in response to this.

Probabilistic Inference



- What is my state and that of the world around me?
- Prior knowledge is combined with new measurements; most generally modelled as a Bayesian Network.
- A series of weighted combinations of old and new information.
- **Sensor fusion**: the general process of combining data from many different sources into useful estimates.
- This composite state estimate can then be used to decide on the robot's next action.

Bayesian Probabilistic Inference


- ‘Bayesian’ has come to be known as a certain view of the meaning of probability theory as a **measure of subjective belief**.
- Probabilities describe our state of knowledge — nothing to do with randomness in the world.
- Bayes’ Rule relating probabilities of discrete statements:

$$\begin{aligned} P(\mathbf{XZ}) &= P(\mathbf{Z|X})P(\mathbf{X}) = P(\mathbf{X|Z})P(\mathbf{Z}) \\ \Rightarrow P(\mathbf{X|Z}) &= \frac{P(\mathbf{Z|X})P(\mathbf{X})}{P(\mathbf{Z})} \end{aligned}$$

- Here $P(\mathbf{X})$ is the **prior**; $P(\mathbf{Z|X})$ the **likelihood**; $P(\mathbf{X|Z})$ the **posterior**; $P(\mathbf{Z})$ sometimes called **marginal likelihood**.
- We use Bayes’s Rule to incrementally **digest** new information from sensors about a robot’s state. Straightforward use for discrete inference where \mathbf{X} and \mathbf{Z} each have values which are one of several labels such as the identity of the room a robot is in.

Probability Distributions: Discrete and Continuous

- Discrete probabilistic inference generalises to large numbers of possible states as we make the bin size smaller and smaller.
- A continuous Probability Density Function $p(x)$ is the limiting case as the widths of bins in a discrete histogram tend to zero.




- The probability that a continuous parameter lies in the range a to b is given by the area under the curve:

$$P_{a \rightarrow b} = \int_a^b p(x) dx$$

- But generic high resolution representation of probability density is very expensive in terms of memory and computation.


Probability Representations: Gaussians



$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

- Explicit Gaussian (or normal) distributions are often used to represent the uncertainty in sensor measurements very well.
- Wide Gaussian *prior* multiplied by *likelihood* curve to produce a *posterior* which is tighter than either. The product of two Gaussians is always another Gaussian.


Probability Representations: Particles



- Here a probability distribution is represented by a finite set of *weighted samples* of the state $\{\mathbf{x}_i, w_i\}$, where $\sum_i w_i = 1$.
- Big advantages are simplicity and the ability to represent and shape of distribution, including *multi-modal* distributions (with more than one peak) in ambiguous situations.
- Disadvantages are a poor ability to represent detailed shape of distribution when number of particles is low. If we increase the number of particles to improve this, the computational cost can be very high.

Probabilistic Localisation


- The robot has a map of its environment in advance.
- The only uncertain thing is the position of the robot.



- The robot stores and updates a *probability distribution* representing its uncertain position estimate.

Monte Carlo Localisation (Particle Filter)

- Cloud of particles represent uncertain robot state: more particles in a region = more probability that the robot is there.



(Dieter Fox *et al.* 1999, using sonar. See animated gif at <http://www.doc.ic.ac.uk/~ajd/Robotics/RoboticsResources/montecarlolocalization.gif> .)

The Particle Distribution

- A particle is a point estimate \mathbf{x}_i of the state (position) of the robot with a weight w_i .

$$\mathbf{x}_i = \begin{pmatrix} x_i \\ y_i \\ \theta_i \end{pmatrix}$$

- The full particle set is:

$$\{\mathbf{x}_i, w_i\} ,$$


for $i = 1$ to N . A typical value might be $N = 100$.

- All weights should add up to 1. If so, the distribution is said to be *normalised*:

$$\sum_{i=1}^N w_i = 1 .$$

Displaying a Particle Set

- We can visualise the particle set by plotting the x and y coordinates as a set of dots; more difficult to visualise the θ angular distribution (perhaps with arrows?) — but we can get the main idea just from the linear components.




Steps in Particle Filtering

These steps are repeated every time the robot moves a little and makes measurements:

1. Motion Prediction based on Proprioceptive Sensors.
2. Measurement Update based on Outward-Looking Sensors.
3. Normalisation.
4. Resampling.

Motion Prediction



- We know that uncertainty grows during blind motion.
- So when the robot makes a movement, the particle distribution needs to shift its mean position but also spread out.
- We achieve this by passing the state part of each particle through a function which has a deterministic component and a random component.

Motion Prediction

- During a straight-line period of motion of distance D :

$$\begin{pmatrix} x_{new} \\ y_{new} \\ \theta_{new} \end{pmatrix} = \begin{pmatrix} x + (D + e) \cos \theta \\ y + (D + e) \sin \theta \\ \theta + f \end{pmatrix}$$

- During a pure rotation of angle α :

$$\begin{pmatrix} x_{new} \\ y_{new} \\ \theta_{new} \end{pmatrix} = \begin{pmatrix} x \\ y \\ \theta + \alpha + g \end{pmatrix}$$

- Here e , f and g are zero mean *noise* terms — i.e. random numbers typically with a Gaussian distribution. We generate a different samples for each particle, which causes the particles to spread out.

Motion Prediction

- We need to find the appropriate standard deviations of the Gaussian distributions from which to sample e , f and g . Make some initial estimates; then adjust by looking at the particle spread over an extended motion and match the distribution to experiments (see practical sheet).
- If your robot is to move through variable step sizes, the values of these standard deviations should change with the distance move. The correct way to do this is to scale the **variance** proportional to the linear or angular distance moved — variance is additive, so for instance two 1m steps will cause the same spread as one 2m step.

Measurement Updates

- A measurement update consists of applying Bayes Rule to each particle; remember:

$$P(\mathbf{X}|\mathbf{Z}) = \frac{P(\mathbf{Z}|\mathbf{X})P(\mathbf{X})}{P(\mathbf{Z})}$$


- So when we achieve a measurement z , we update the weight of each particle as follows:

$$w_{i(new)} = P(z|\mathbf{x}_i) \times w_i ,$$

remembering that the denominator in Bayes' rule is a constant factor which we do not need to calculate because it will later be removed by normalisation.

- $P(z|\mathbf{x}_i)$ is the *likelihood* of particle i ; the probability of getting measurement z given that it represents the true state.

Likelihood Function



- The form of a likelihood function comes from a probabilistic model of the outward-looking sensor.
- Having calibrated a sensor and understood the uncertainty in its measurements we can build a probabilistic measurement model for how it works. This will be a probability distribution (specifically a likelihood function) of the form:

$$P(z|\mathbf{x}_i)$$

Such a distribution will often have a Gaussian shape.


Inferring an Estimate and Position-Based Navigation

- At any point, our uncertain estimate of the location of the robot is represented by the whole particle set.
- If we need to, we can make a point estimate of the current position and orientation of the robot by taking the mean of all of the particles:

$$\bar{\mathbf{x}} = \sum_{i=1}^N w_i \mathbf{x}_i .$$

- i.e. the means of the x , y and θ components are all calculated individually and stored in $\bar{\mathbf{x}}$.
- The robot could use this for instance to plan A \rightarrow B waypoint navigation.

Position-Based Path Planning



Assuming that a robot has *localisation*, and knows where it is relative to a fixed coordinate frame, then position-based path planning enables it to move in a precise way along a sequence of pre-defined waypoints. Paths of various curved shapes could be planned, aiming to optimise criteria such as overall time or power usage. Here we will consider the specific, simple case where we assume that:

- Our robot's movements are composed by straight-line segments separated by turns on the spot.
- The robot aims to minimise total distance travelled, so it always turns immediately to face the next waypoint and drives straight towards it.

Position-Based Path Planning

In one step of path planning, assume that the robot's current pose is (x, y, θ) and the next waypoint to travel to is at (W_x, W_y) .

- It must first rotate to point towards the waypoint. The vector direction it must point in is:

$$\begin{pmatrix} d_x \\ d_y \end{pmatrix} = \begin{pmatrix} W_x - x \\ W_y - y \end{pmatrix}$$

The absolute angular orientation α the robot must drive in is therefore given by:


$$\alpha = \tan^{-1} \frac{d_y}{d_x}$$

Care must be taken to make sure that α is in the correct quadrant of $-\pi < \alpha \leq \pi$. A standard \tan^{-1} function will return a value in the range $-\pi/2 < \alpha \leq \pi/2$. This can be also achieved directly with an `atan2(dy, dx)` in most languages.

Position-Based Path Planning

- The angle the robot must rotate through is therefore $\beta = \alpha - \theta$. If the robot is to move as efficiently as possible, care should be taken to shift this angle by adding or subtracting 2π so make sure that $-\pi < \beta \leq \pi$.
- The robot should then drive forward in a straight line through distance $d = \sqrt{d_x^2 + d_y^2}$.

Practical 4: Probabilistic Motion



- Today's practical is on probabilistic robot motion, and position-based motion planning.
- **Everyone should read the practical sheet fully!**
- This is an **ASSESSED** practical: we will assess your achievement next week in the live practical session. Your whole group should be there to demonstrate and discuss your robot via screenshare.
- A reminder that no reports or code need to be submitted. Assessment is via demonstration and discussion.