

Robotics and Autonomous Systems

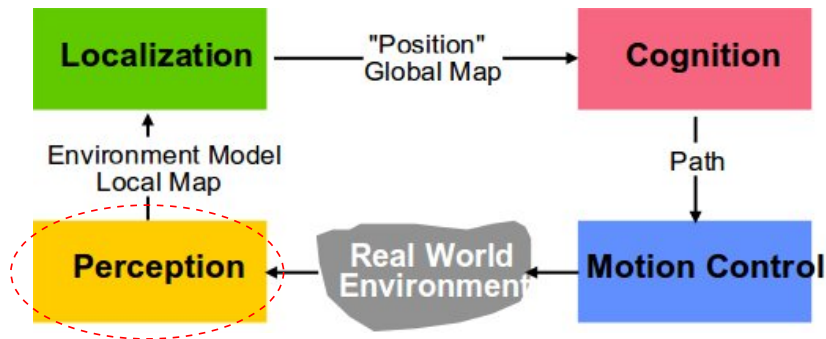
Lecture 7: Perception

Terry Payne

Department of Computer Science
University of Liverpool



UNIVERSITY OF
LIVERPOOL



- We'll finish talking about perception.

Classification of sensors

- Proprioceptive sensors
- Exteroceptive sensors
- Passive sensors
- Active sensors

Classification of sensors



- Proprioceptive sensors
 - Measure values internally to the system (robot)
(motor speed, wheel load, heading of the robot, battery status).
- Exteroceptive sensors
 - Information from the robots environment
(distances to objects, intensity of the ambient light, unique features).
- Passive sensors
 - Energy coming from the environment.
- Active sensors
 - Emit their own energy and measure the reaction.
 - Better performance, but some influence on environment.

Classification of sensors

- Proprioceptive sensors
 - Measure values internally to the system (robot), (motor speed, wheel load, heading of the robot, battery status).
- Exteroceptive sensors
 - Information from the robots environment (distances to objects, intensity of the ambient light, unique features).
- Passive sensors
 - Energy coming from the environment.
- Active sensors
 - Emit their own energy and measure the reaction.
 - Better performance, but some influence on environment.


Classification of sensors

- Proprioceptive sensors
 - Measure values internally to the system (robot),
(motor speed, wheel load, heading of the robot, battery status)
- Exteroceptive sensors
 - Information from the robots environment
(distances to objects, intensity of the ambient light, unique features.)
- Passive sensors
 - Energy coming from the environment.
- Active sensors
 - Emit their own energy and measure the reaction
 - Better performance, but some influence on environment



- Note that the robot has a number of different sensors.

B14 and B21



- Built at CMU.



Sensors include bump panels, a Denning sonar ring, a Nomadics laser light stripper, and twin cameras mounted on a Directed Perception pan/tilt head for stereo vision.

- Also includes a 4-wheel synchrodrive.

- Omnidirectional and pan/tilt camera.
- Sonar
- Wheel encoders
- Laser range finder
- Bumpers




- BlueBotics SA, Switzerland

Sensor characteristics

- Range
 - The upper limit that a sensor can measure.
- Linearity
 - Variation of output signal as function of the input signal.
 - Less important when signal is post-processed.
- Bandwidth or Frequency
 - The speed with which a sensor can provide readings
 - Usually an upper limit. Depends on sensor and the sample rate.
 - Lower limit is also possible, e.g. acceleration sensor.
- Resolution
 - Minimum difference between two values. Usually the lower limit of dynamic range.
 - For digital sensors it is usually the A/D resolution. (e.g. $5V / 255$ (8 bit))

Bumpers

- You should be a bit familiar with the bumper on our NXT robot by now.
 - Each bumper says when it has hit something.
- Bumpers are just contact switches — indicate when they are pressed.



- While our NXT robot has just two bumpers, a robot can have many.

Range sensors

- Large range distance measurement → called range sensors.
- Range information is the key element for localization and environment modeling.
- Ultrasonic sensors, infra-red sensors and laser range sensors make use of propagation speed of sound or electromagnetic waves respectively.
- The distance traveled by a sound or electromagnetic wave is given by

$$d = c.t$$

- Where:
 - d = distance traveled (usually round-trip)
 - c = speed of wave propagation
 - t = time of flight.

Ultrasound (sonar) sensor

- Transmit a packet of (ultrasonic) pressure waves
- Distance d of the echoing object can be calculated based on the propagation speed of sound c and the time of flight t .


$$d = \frac{c \cdot t}{2}$$

- The speed of sound c (340 m/s) in air is given by:


$$c = \sqrt{\gamma \cdot R \cdot T}$$

where:

- γ : ratio of specific heats
- R : gas constant
- T : temperature in degree Kelvin





Sonar timing




What gets measured

- Sound beam propagates in a cone.
- Opening angles around 20 to 40 degrees
- Detects regions of constant depth on segments of an arc



- Piezo electric transducer generates frequency: 40 – 180 kHz

Typical sonar scan



- **Specular reflection** from non-perpendicular surfaces.
- Absorption from soft surfaces.

- Note that in places the result is far from accurate.

- To use the sonar you first need to link the right library:

```
import lejos.nxt.UltrasonicSensor;
```

- Then you need to create an instance of the sensor:

```
UltrasonicSensor us = new  
UltrasonicSensor(SensorPort.S1);
```

(make sure you use the right port for your robot).

- Then you can read the distance of the nearest object in front of the robot:

```
us.getRange();
```

- You will likely want to use the sonar in your first assignment, so get familiar with it.

Laser range finder

- A laser range-finder uses light rather than sound.





Figure 4.11

(a) Schematic drawing of laser range sensor with rotating mirror; (b) Scanning range sensor from EPS Technologies Inc.; (c) Industrial 180 degree laser range sensor from Sick Inc., Germany

- The rotating mirror allows the laser to take many measurements.

Laser range finder

- For any wave, speed is related to frequency and wavelength by:
 $c = f \cdot \lambda$




- The total distance covered by the light is:

$$\text{distance} = L + 2D = L + \frac{\theta}{2\pi} \lambda$$


- The distance of the target is then:

$$D = \frac{\lambda}{4\pi} \theta$$




where θ is the phase shift.

Typical laser scan



- Length of bars is an estimate of the error.

IR distance sensor



*Position-Sensitive Device (PSD)
or Linear Camera*

- Distance is inversely proportional to x

$$D = f \frac{L}{x}$$





State of the art

- Hokuyo manufacture a cheap laser scanner.




- The Kinect has made accurate range-finder data much cheaper to acquire.



- Vision is the sense that humans rely upon most.
- It provides the means to gather lots of data very quickly.
- Attractive for use on robots which are often data-poor.
- However, presents a new challenge
 - How can we extract data from an image
 - Or from a sequence of images.



Cameras



- Today, with cheap CMOS cameras, the hardware cost of adding a camera to a robot is negligible.



- Although vision seems to be easy for humans, it is hard for machines.
(as always, remember how long it takes us to learn to “see”).
- Reasons include:
 - variable illumination,
 - uncontrolled illumination,
 - shadows,
 - irregular objects,
 - occultion of objects,
 - noisy sensors,
 - ...
- Typically these problems are worse outside.

- The lens produces a **perspective projection** of the scene.
- The 3-d scene becomes a 2-d image:


$$I(x, y, t)$$

x and y are the co-ordinates of the array, t is time.

- The image is just an array.
- Well, typically 3 arrays — each with one entry per pixel in the image.
 - Why?
- These must be processed to extract the information that we need.





Problems in processing

- The projection from 3D to 2D introduces massive ambiguity.



- What the camera sees in one view can be generated by many different scenes.

Problems in processing



- We will look briefly at a couple of basic computer vision techniques.
- These don't come close to solving the general vision problem.
 - Nobody has come close to solving that.
- However, they give us some ways to extract data that can help our robots in some domains.
- Where we know what to expect, we can look for it.

Color segmentation

- An image is a two dimensional array of pixels.
- Each pixel is a set of three values:

$\langle red, green, blue \rangle$

typically with a value between 0 and 255 (8 bit).

- (Well, most computer vision uses something other than RGB, but the principle is the same.)
- Define a color you want to recognise as a box in RGB space:

$$red \in [30, 100]$$


$$blue \in [70, 120]$$

$$green \in [150, 230]$$

- Label each pixel 1 if it falls in the box, 0 if it falls outside the box.

Color segmentation


- Result is a set of “blobs” which, if you calibrated correctly, identify objects of interest.



- Example: segmentation in robot soccer.

Color segmentation

- Result is a set of “blobs” which, if you calibrated correctly, identify objects of interest.




- Example: segmentation in robot soccer.

What you can do with a segmented image

- Object identification:
 - “I see an orange blob” means “I see the ball”.
- Object tracking:
 - Keep the orange blob in the center of the frame.
- Limited navigation:
 - Walk towards the orange blob.
 - When you get to the orange blob, kick it towards the blue blob.
- Localization.
 - Measure angles of blue and yellow blobs to me.
 - If I know the location of the blobs, I can tell where I am.

Works well enough for some applications



Edge detection

- We often want to identify edges.
- We can then use the edges to identify shapes that we are looking for in an image.
- What we do has a mathematical interpretation in terms of **convolution**, but there's also a simple way to think about this.
- Edges are all about changes in color (in a color image) or intensity (in a black and white image).
- So identifying pixels that are on an edge is relatively easy.
 - We look for sudden changes in R, G and B in a color image or the single value in a b/w image.


Edge detection

- Gives us something like:




Edge detection

- Often edge detection gives us many mini-edges that need to be merged or removed:



- Pre-processing the image can also help.
- For example, noise can be removed by **smoothing** the image.
 - Averaging across the pixel values.
- For example we might replace the value of every pixel by the average of the values of the 8 pixels around it.
- The larger the area we average over, the more robust the results are against noise.
- Of course, all this processing is expensive, and slows down the speed of reaction of the robot.

Stereo vision




- Two cameras, spaced as widely as possible.
- Can get depth information if we can identify the common point(s) in two images.

Stereo vision

$$x = b \left(\frac{x_l + x_r}{2(x_l - x_r)} \right)$$

$$y = b \left(\frac{y_l + y_r}{2(x_l - x_r)} \right)$$

$$z = b \left(\frac{f}{x_l - x_r} \right)$$




- The accuracy of the depth estimate increases with increasing baseline b .

Using stereo vision



- Also equipped with several other sensors.

The current frontier




- Omnidirectional cameras allow robots to see all around them.
 - Usually mounted with the lens above the camera.

The current frontier




- Presents new challenges in machine vision.

The current frontier



The current frontier



Summary

- This lecture finished our look at sensors and perception.
- We spent most of our time looking at:
 - Range sensors
 - Cameras and image data.
- These are probably the most widely used sensors in robotics today.