PERCEPTION

Today

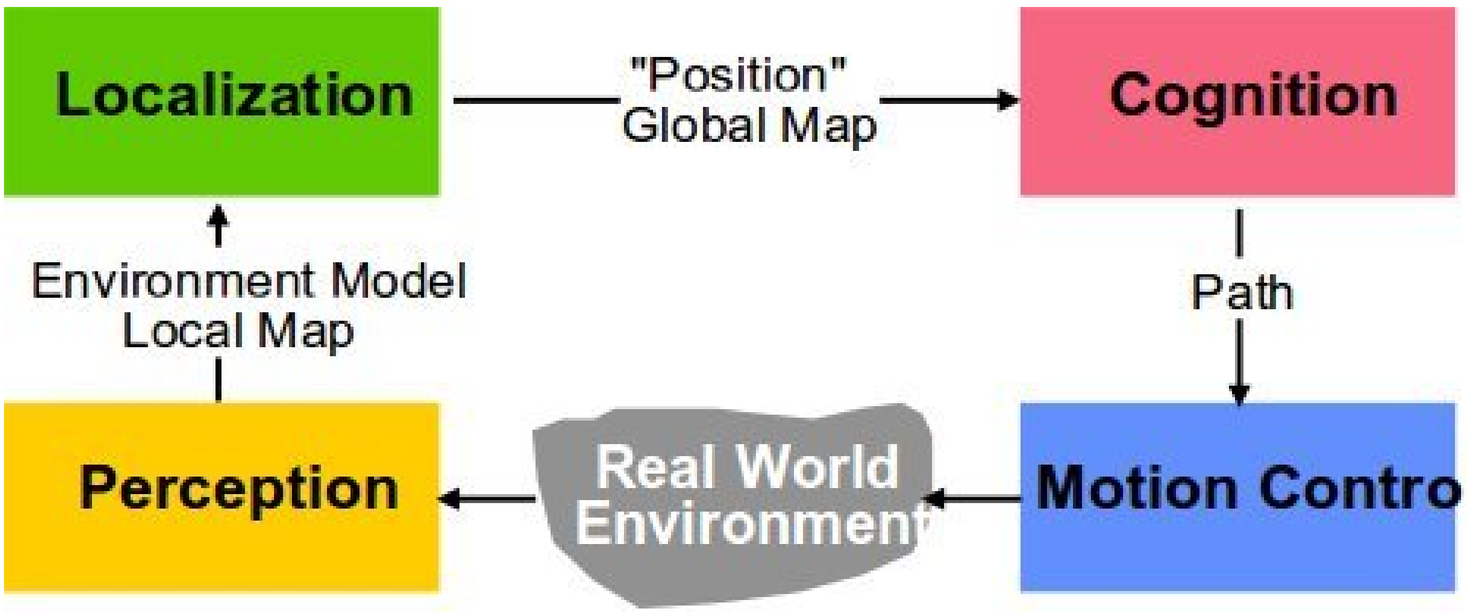

- Last time we looked at motion, on the bottom right hand side.
- This time we'll move round to the bottom left.

Perception

- So far we have programmed robots more or less "blind"
  - Almost no information from the environment.

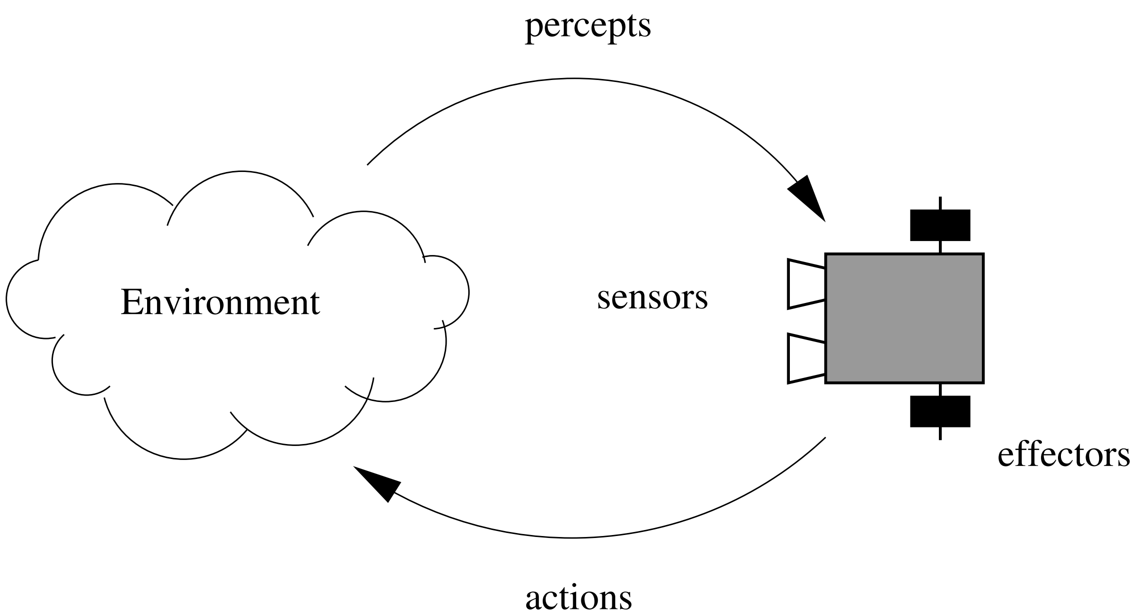

- Perception is all about what can be sensed and what we can do with that sensing.

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 for the environment.
- Active sensors
  - Emit their own energy and measure the reaction
  - Better performance, but some influence on envrionment

Helpmate

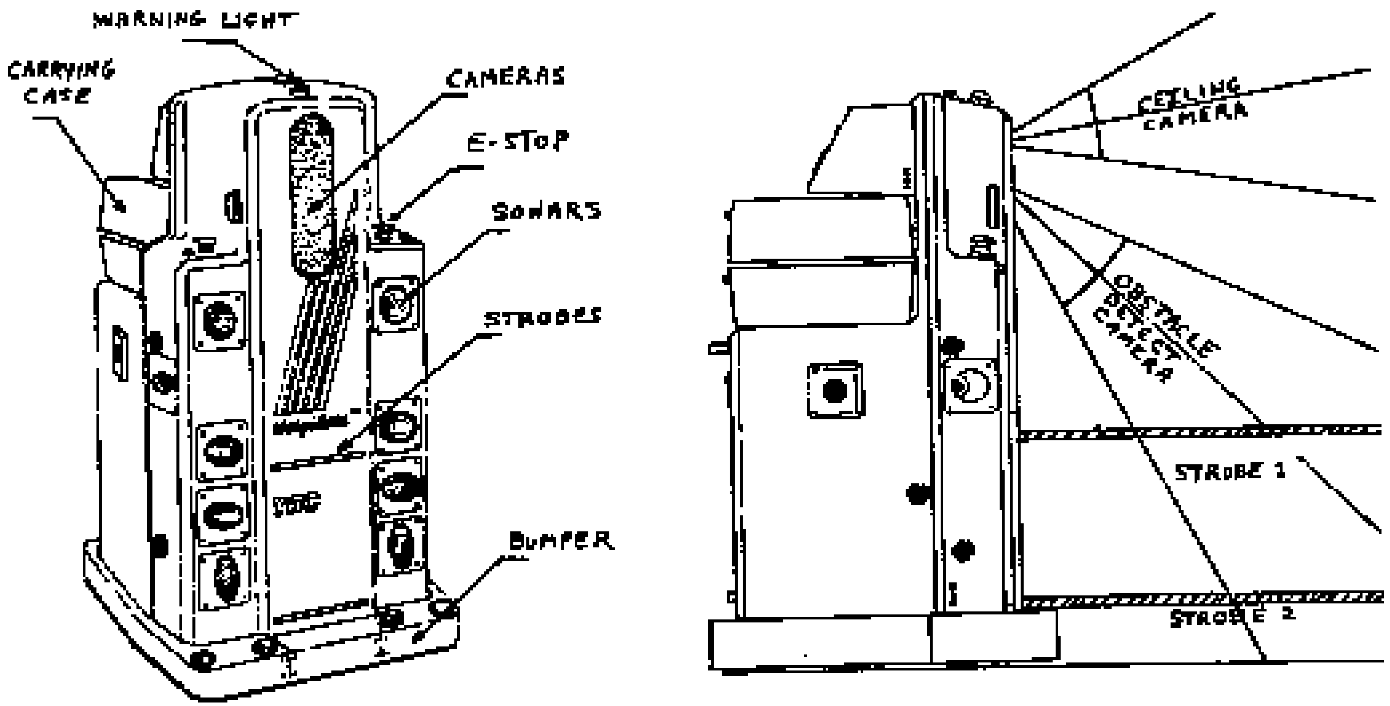

- Note that the robot has a number of different sensors.

B14 and B21

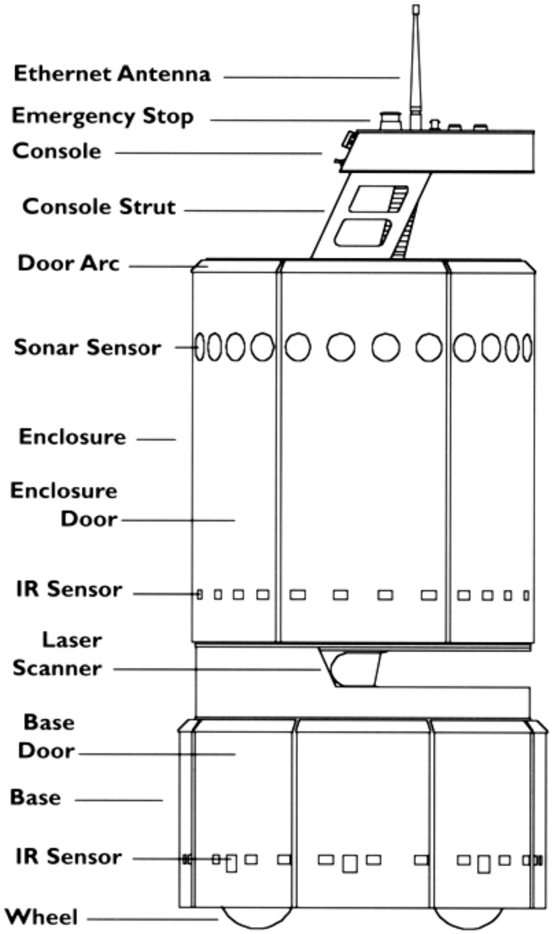

Ethernet Antenna
Emergency Stop
Console
Console Strut
Door Arc
Sonar Sensor
Enclosure
Enclosure Door
IR Sensor
Laser Scanner
Base Door
Base
IR Sensor
Wheel

RWI

Xavier

- Built at CMU.

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

- Also includes a 4-wheel synchrodrive.

BibaBot

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

- BlueBotics SA, Switzerland

General classification (1)

General classification (typical use)	Sensor Sensor System	PC or EC	A or P
Tactile sensors (detection of physical contact or closeness; security switches)	Contact switches, bumpers	EC	P
	Optical barriers	EC	A
	Noncontact proximity sensors	EC	A
Wheel/motor sensors (wheel/motor speed and position)	Brush encoders	PC	P
	Potentiometers	PC	P
	Synchros, resolvers	PC	A
	Optical encoders	PC	A
	Magnetic encoders	PC	A
	Inductive encoders	PC	A
	Capacitive encoders	PC	A
Heading sensors (orientation of the robot in relation to a fixed reference frame)	Compass	EC	P
	Gyroscopes	PC	P
	Inclinometers	EC	A/P

A, active; P, passive; P/A, passive/active; PC, proprioceptive; EC, exteroceptive.

General classification (2)

General classification (typical use)	Sensor Sensor System	PC or EC	A or P
Ground-based beacons (localization in a fixed reference frame)	GPS Active optical or RF beacons Active ultrasonic beacons Reflective beacons	EC EC EC EC	A A A A
Active ranging (reflectivity, time-of-flight, and geometric triangulation)	Reflectivity sensors Ultrasonic sensor Laser rangefinder Optical triangulation (1D) Structured light (2D)	EC EC EC EC EC	A A A A A
Motion/speed sensors (speed relative to fixed or moving objects)	Doppler radar Doppler sound	EC EC	A A
Vision-based sensors (visual ranging, whole-image analysis, segmentation, object recognition)	CCD/CMOS camera(s) Visual ranging packages Object tracking packages	EC	P

Sensor characteristics (1)

- Range
  - The upper limit that a sensor can measure.
- Dynamic range
  - Ratio between lower and upper limits, usually in decibels (dB) which is a measure of power.
  - Power measurement from 1 Milliwatt to 20 Watts

$$-10.\boldsymbol{log}\left[\frac{0.001}{20}\right] = 43\boldsymbol{dB}$$

  - Voltage measurement from 1 Millivolt to 20 Volt

$$-20.\boldsymbol{log}\left[\frac{0.001}{20}\right] = 86\boldsymbol{dB}$$

20 instead of 10 because square of voltage is equal to power

Sensor characteristics (2)

- 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))

Sensor performance (1)

- Sensitivity
  - Ratio of output change to input change
  - In real world environment, the sensor has very often high sensitivity to other environmental changes, e.g. illumination.
- Cross-sensitivity
  - Sensitivity to environmental parameters that are orthogonal to the target parameters
- Error / Accuracy
  - Difference between the sensors output and the true value

$$accuracy = 1 - \frac{|m - v|}{v}$$

where m = measured value and v = true value.

Sensor performance (2)

- Systematic error $\rightarrow$ deterministic errors
  - Caused by factors that can (in theory) be modeled $\rightarrow$ prediction

    e.g. distortion caused by the optics of a camera.
- Random error $\rightarrow$ non-deterministic
  - No prediction possible
  - However, they can be described probabilistically

    e.g. error in wheel odometry.
- Precision
  - Reproducibility of sensor results

$$\frac{range}{\sigma}$$

Bumpers

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

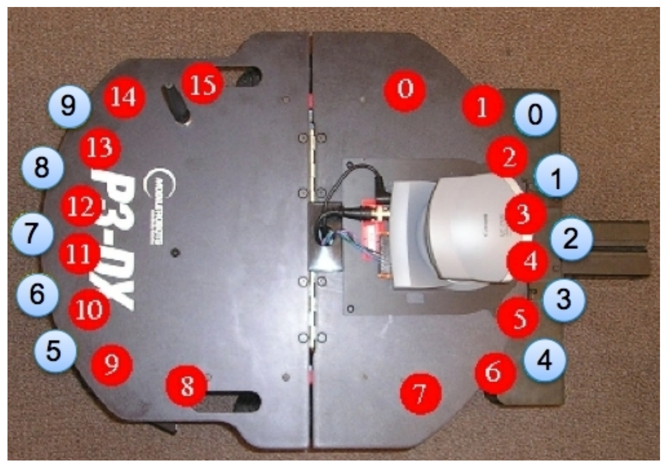

- While the Create has just two bumpers, a robot can have many.

Wheel encoders

- Measure position or speed of the wheels or steering.
- Wheel movements can be integrated to get an estimate of the robot's position
  - Odometry (this is what the position proxy is doing)
- Optical encoders are proprioceptive sensors
  - Position estimate is only useful for short movements.
- Typical resolutions: 2000 increments per revolution.

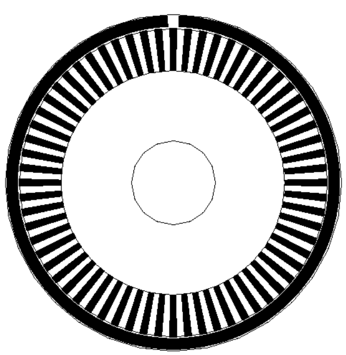

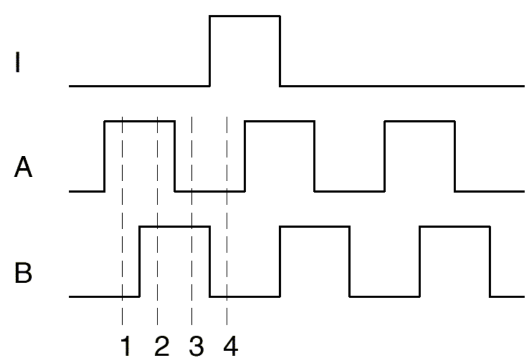

State	Ch A	Ch B
S_1	High	Low
S_2	High	High
S_3	Low	High
S_4	Low	Low

Heading sensors

- Heading sensors can be:
  - proprioceptive (gyroscope, inclinometer); or
  - exteroceptive (compass).
- Used to determine the robots orientation and/or inclination.
- Allow, together with an appropriate velocity information, to integrate the movement to an position estimate.
  - A bit more sophisticated than just using odometry.

Compass

- Used since before 2000 B.C.
  - Chinese suspended a piece of naturally magnetite from a silk thread and used it to guide a chariot over land.
- Magnetic field on earth
  - Absolute measure for orientation.
- Large variety of solutions to measure the earth magnetic field
  - Mechanical magnetic compass
  - Direct measure of the magnetic field (Hall-effect, magnetoresistive sensors)

- Major drawback
  - Weakness of the earth field
  - Easily disturbed by magnetic objects or other sources
    Not feasible for indoor environments in general.

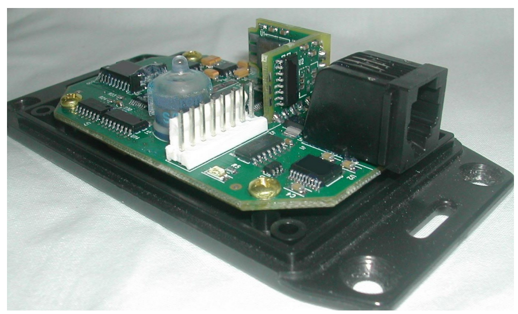

- Modern devices can give 3D orientation relative to Earth's magnetic field and integrate to give distance.

Gyroscope

- Heading sensors, that keep the orientation to a fixed frame
  - Provide an absolute measure for the heading of a mobile system.
  - Unlike a compass doesn't measure the outside world.
- Two categories, mechanical and optical gyroscopes
- Mechanical Gyroscopes
  - Standard gyro
  - Rated gyro
- Optical Gyroscopes
  - Rated gyro

- Concept: inertial properties of a fast spinning rotor
  - *gyroscopic precession*
- Angular momentum associated with a spinning wheel keeps the axis of the gyroscope inertially stable.
- No torque can be transmitted from the outer pivot to the wheel axis
- Spinning axis will therefore be space-stable
- Quality: 0.1 degrees in 6 hours

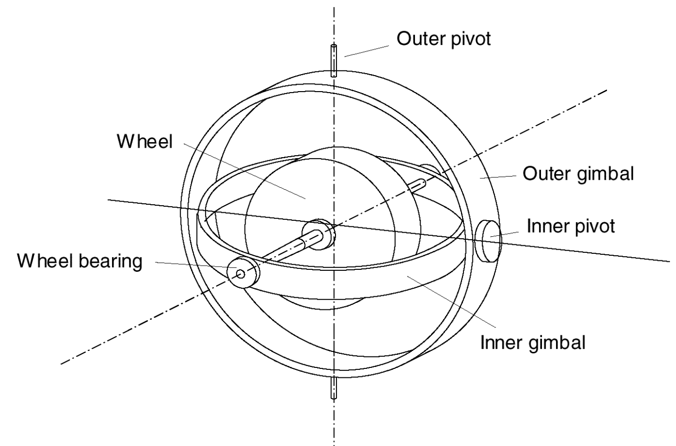

- In *rate* gyros, gimbals are held by torsional springs.
  - Measuring force gives angular velocity.

Optical gyroscopes

- Use two monochromatic light (or laser) beams from the same source.
- One beam travels clockwise in a cylinder around a fiber, the other counterclockwise.
- The beam traveling in direction of rotation:
- Slightly shorter path $\rightarrow$ shows a higher frequency
- Difference in frequency Δf of the two beams is proportional to the angular velocity Ω of the cylinder/fiber.

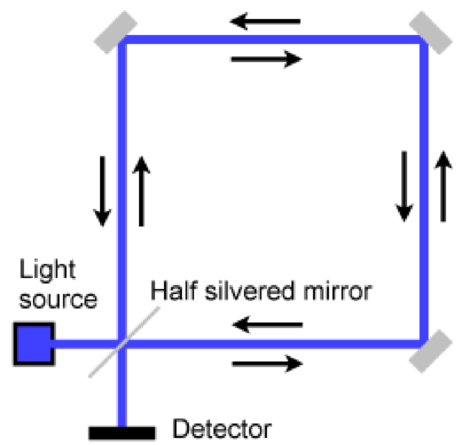

- Newest optical gyros are solid state.

Ground-based beacons

- Elegant way to solve the localization problem in mobile robotics
- Beacons are objects with a precisely known position
- Used since the humans started to travel
  - Natural beacons (landmarks) like stars, mountains or the sun
  - Artificial beacons like lighthouses

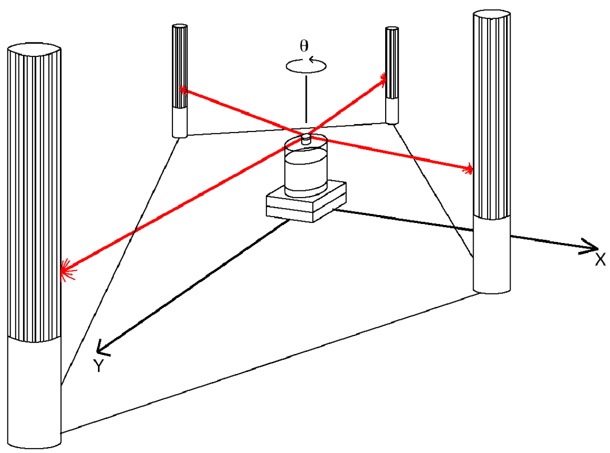

- Cost prohibits too much reliance on beacons.

GPS

- Extension of the beacon idea, developed for military use
- 24 satellites (including three spares) orbiting the earth every 12 hours at a height of 20,190 km.
- Four satellites are located in each of six planes inclined 55 degrees with respect to the plane of the earths equator
- Location of any GPS receiver is determined through a time of flight measurement
- Technical challenges:
  - Time synchronization between the individual satellites and the GPS receiver
  - Real time update of the exact location of the satellites
  - Precise measurement of the time of flight
  - Interference with other signals

GPS (2)

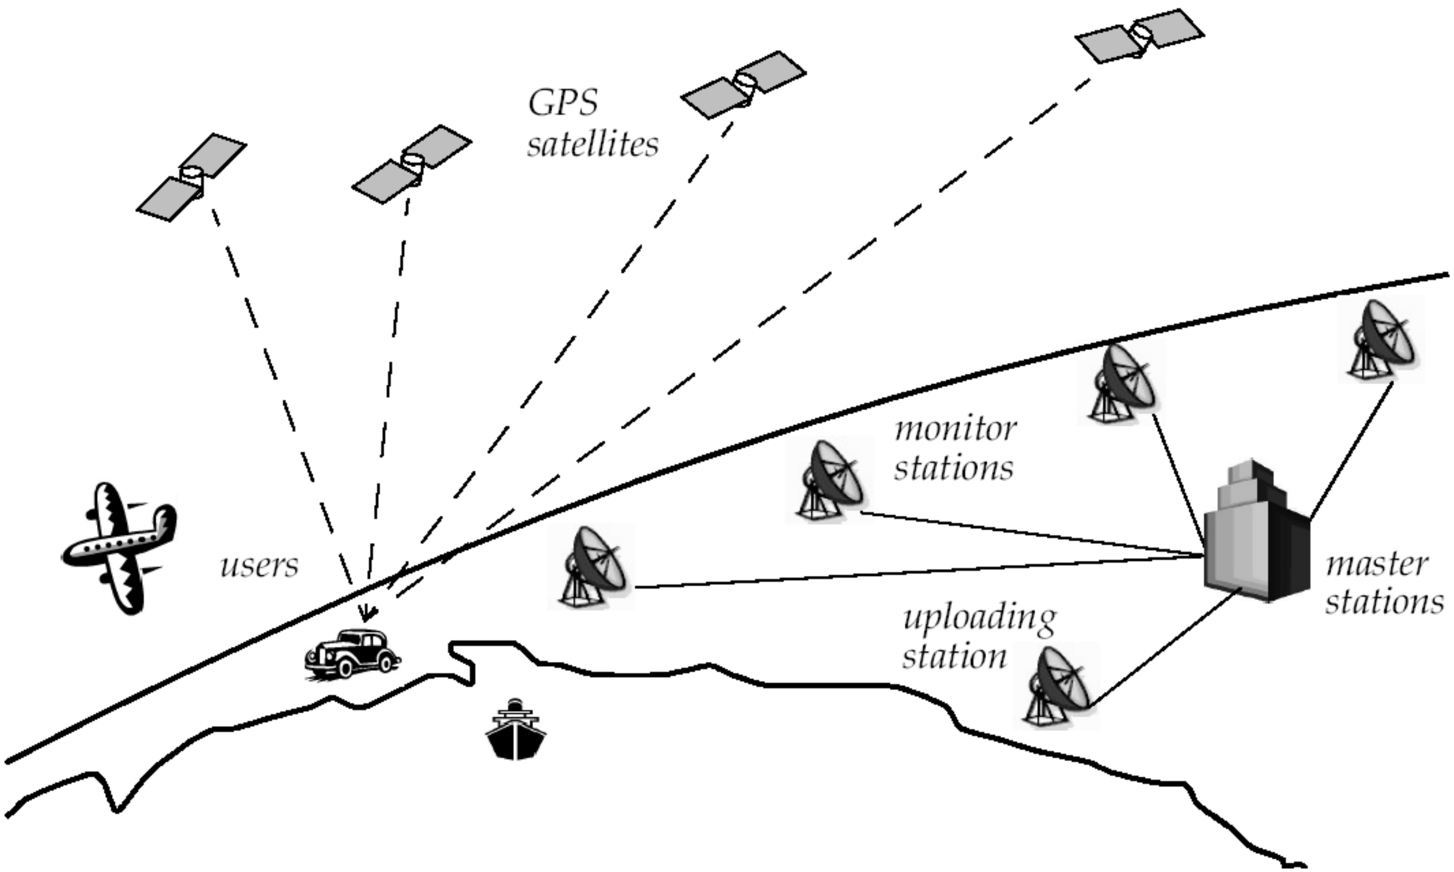

GPS
satellites
monitor
stations
users
master
stations
uploading
station

GPS(3)

- Time synchronization:
  - Atomic clocks on each satellite
  - Monitored from different ground stations
- Ultra-precision time synchronization is extremely important
  - Electromagnetic radiation propagates at light speed,
  - Roughly 0.3 m per nanosecond.
- Position accuracy proportional to precision of time measurement.
- Real time update of the exact location of the satellites:
  - Monitoring the satellites from ground stations
  - Master station analyses all the measurements and transmits the actual position to each of the satellites

Range sensors

- Large range distance measurement $\rightarrow$ 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.

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

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

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

where:

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

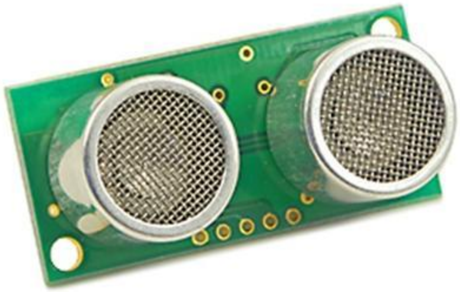

Sonar timing

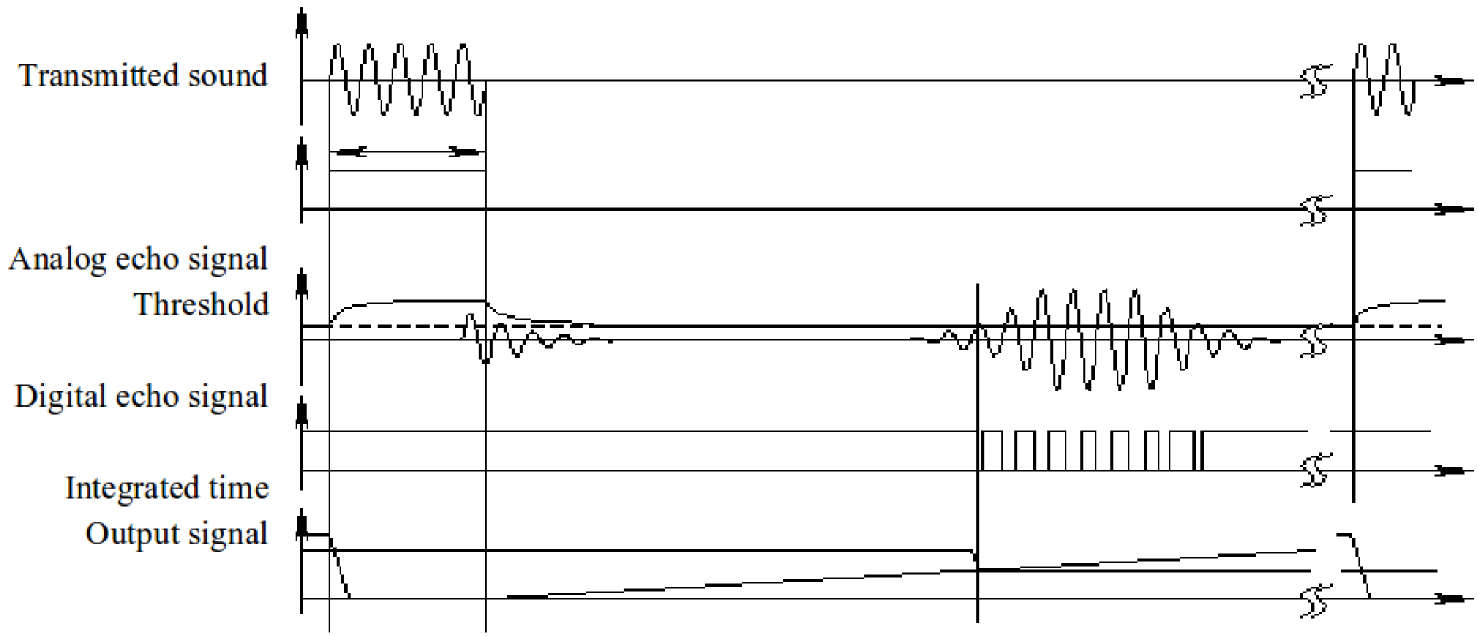

What gets measured

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

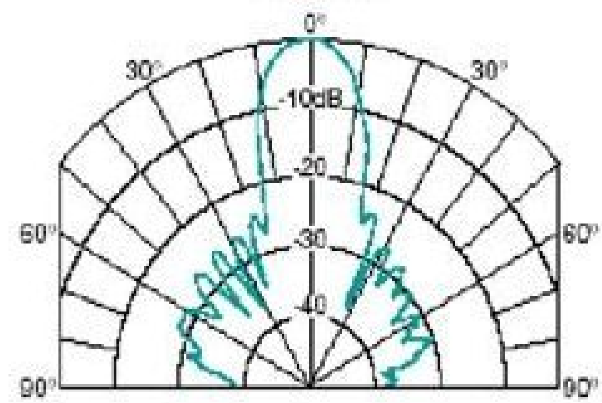

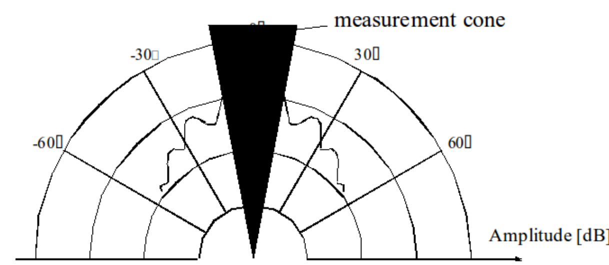

- Piezo electric transducer generates frequency: $40 - -180$ kHz

Typical sonar scan

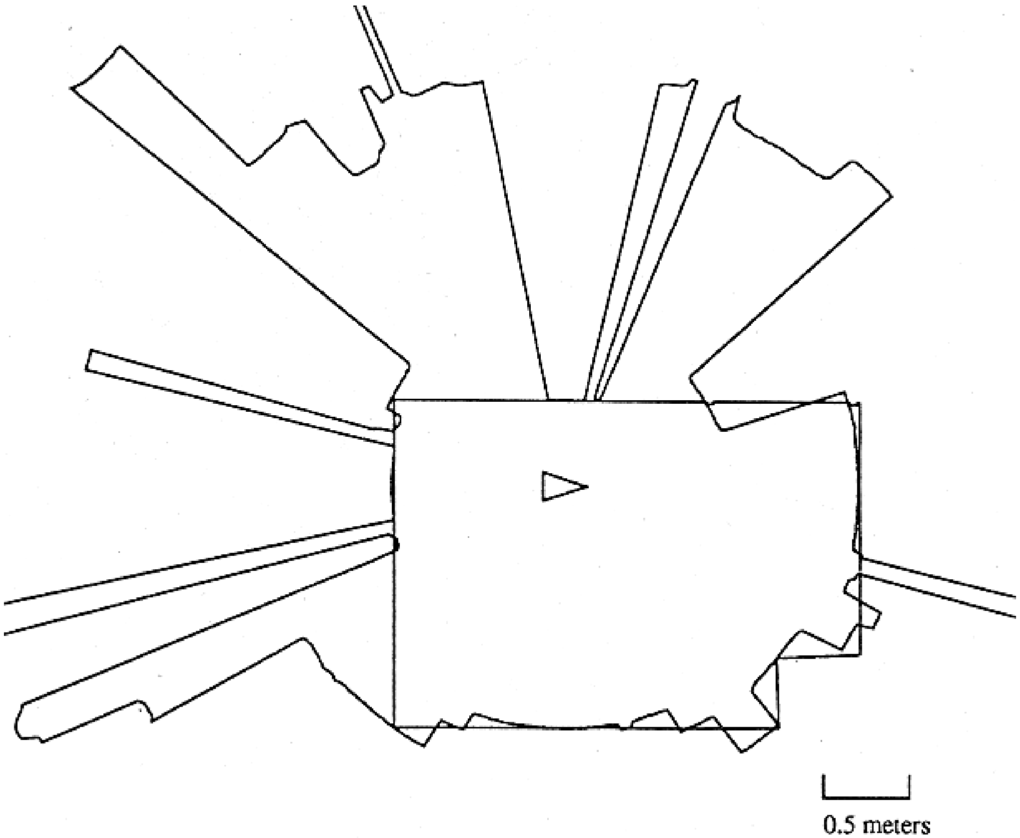

- *Specular reflection* from non-perpendicular surfaces.
- Absorption from soft surfaces.

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

Laser range finder

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

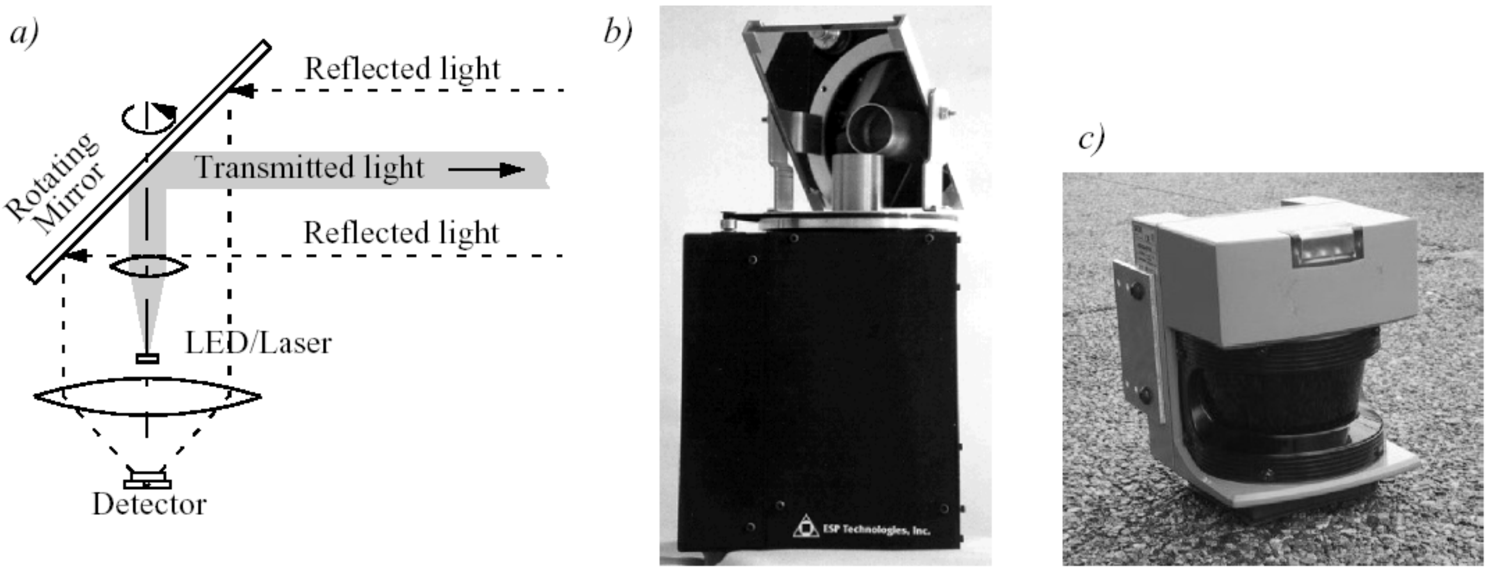

Figure 4.11
(a) Schematic drawing of laser range sensor with rotating mirror; (b) Scanning range sensor from EPS Technologies Inc.; (c) Industrial 180 degree laser range sensor from Sick Inc., Germany

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

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

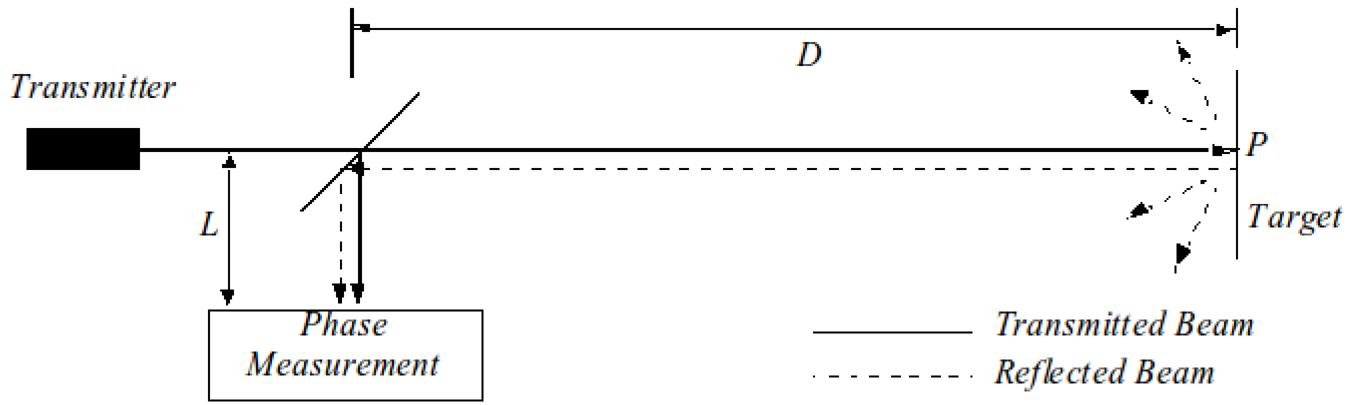

- The total distance covered by the light is:

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

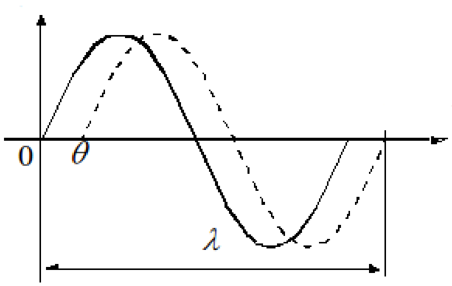

- The distance of the target is then:

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

where θ is the phase shift.

Typical laser scan

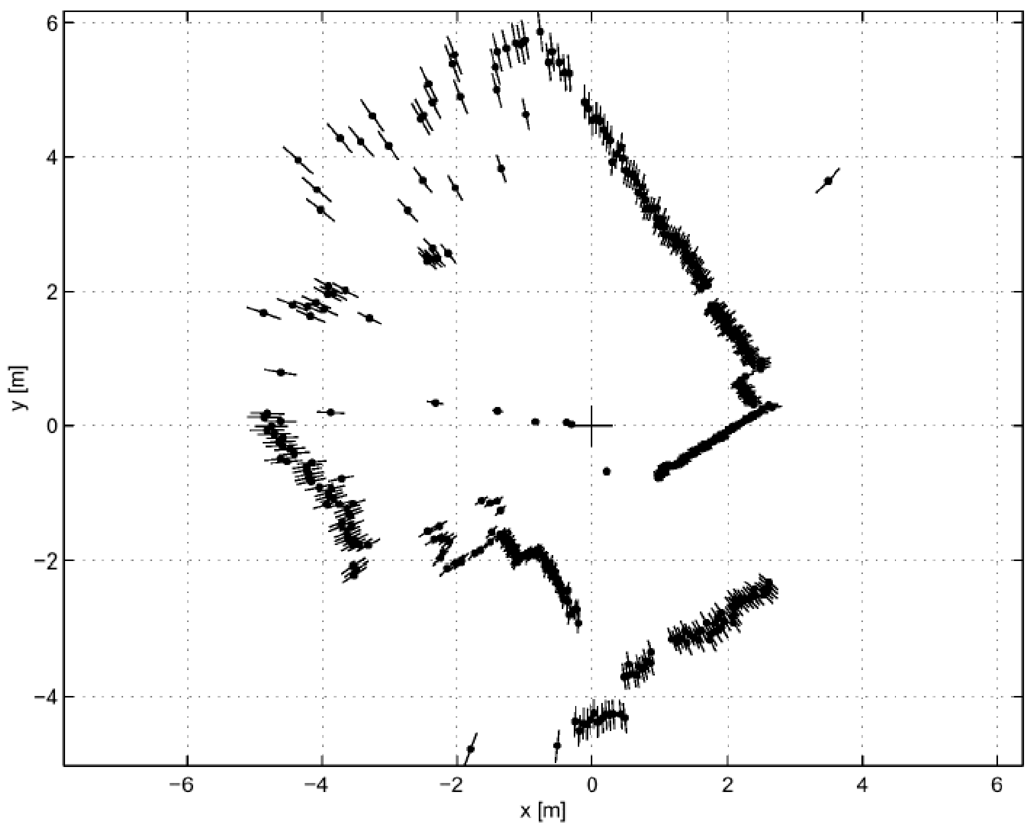

- Length of bars is an estimate of the error.

IR distance sensor

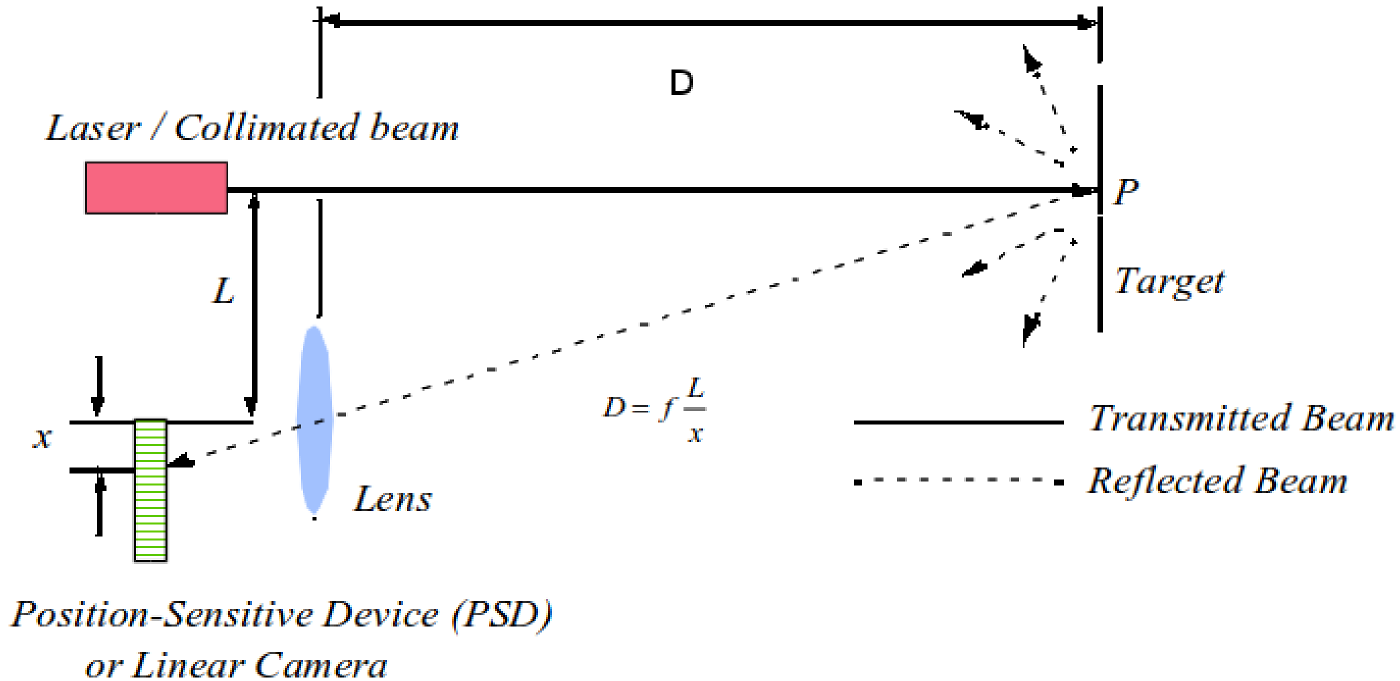

- Distance is inversely proportional to x

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

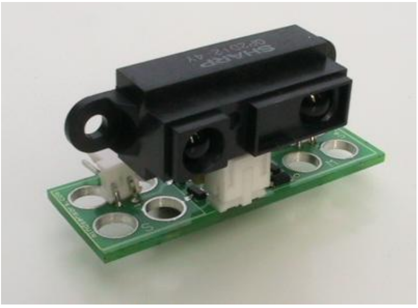

State of the art

- Hokuyu manufacture a cheap laser scanner.

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

Vision

- 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

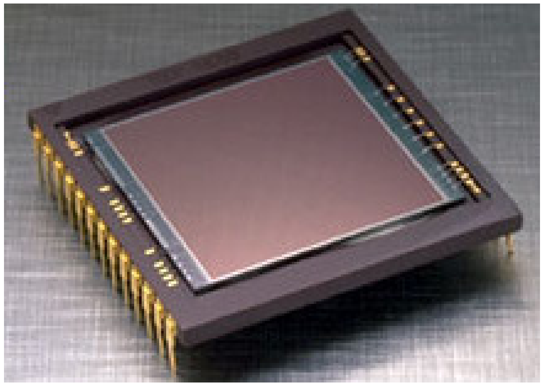

- Today, with cheap CMOS cameras, the hardware cost of adding a camera to a robot is negligible.

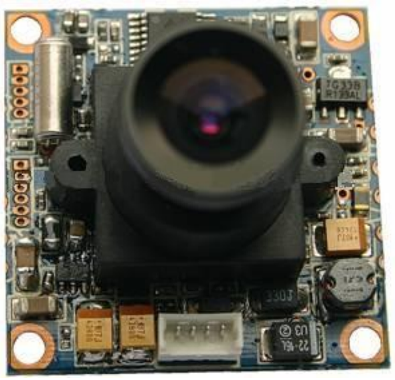

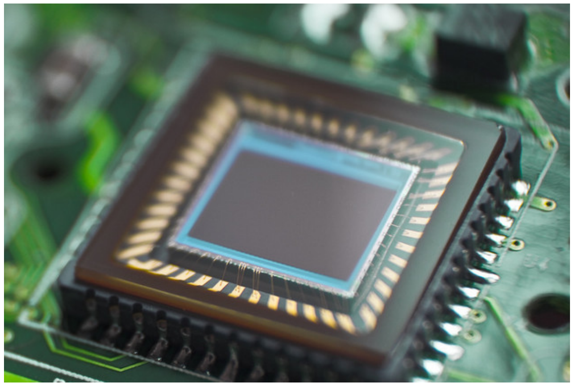

- Although vision seems to be easy for humans, it is hard for machines.

  (as always, remember how long it takes us to learn to "see").

- Reasons include:
  - variable illumination,
  - uncontrolled illumination,
  - shadows,
  - irregular objects,
  - occulsion 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.

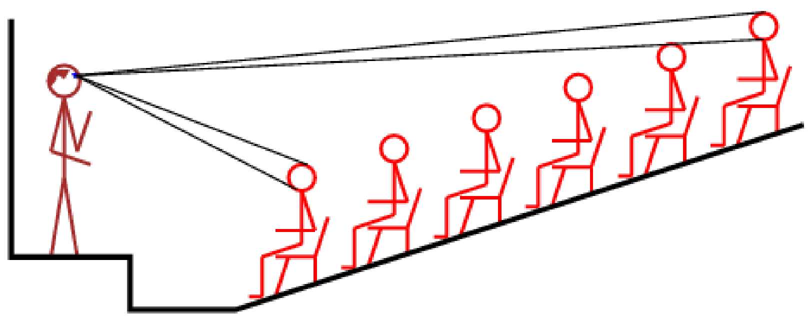

- What the camera sees in one view can be generated by many different scenes.

Problems in processing

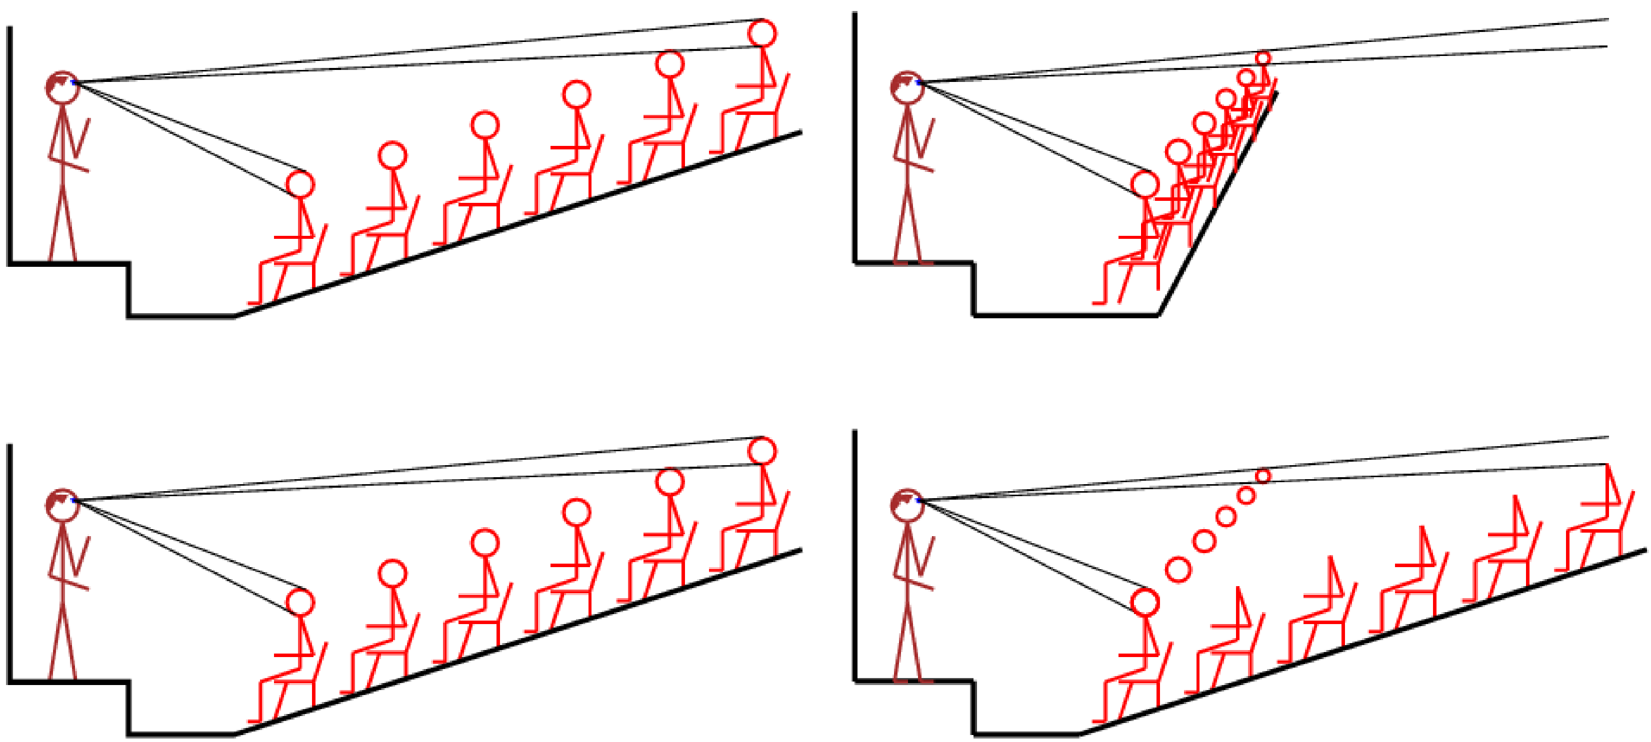

Vision techniques

- 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:

$$\begin{aligned} red &\in [30, 100] \\ blue &\in [70, 120] \\ green &\in [150, 230] \end{aligned}$$

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

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

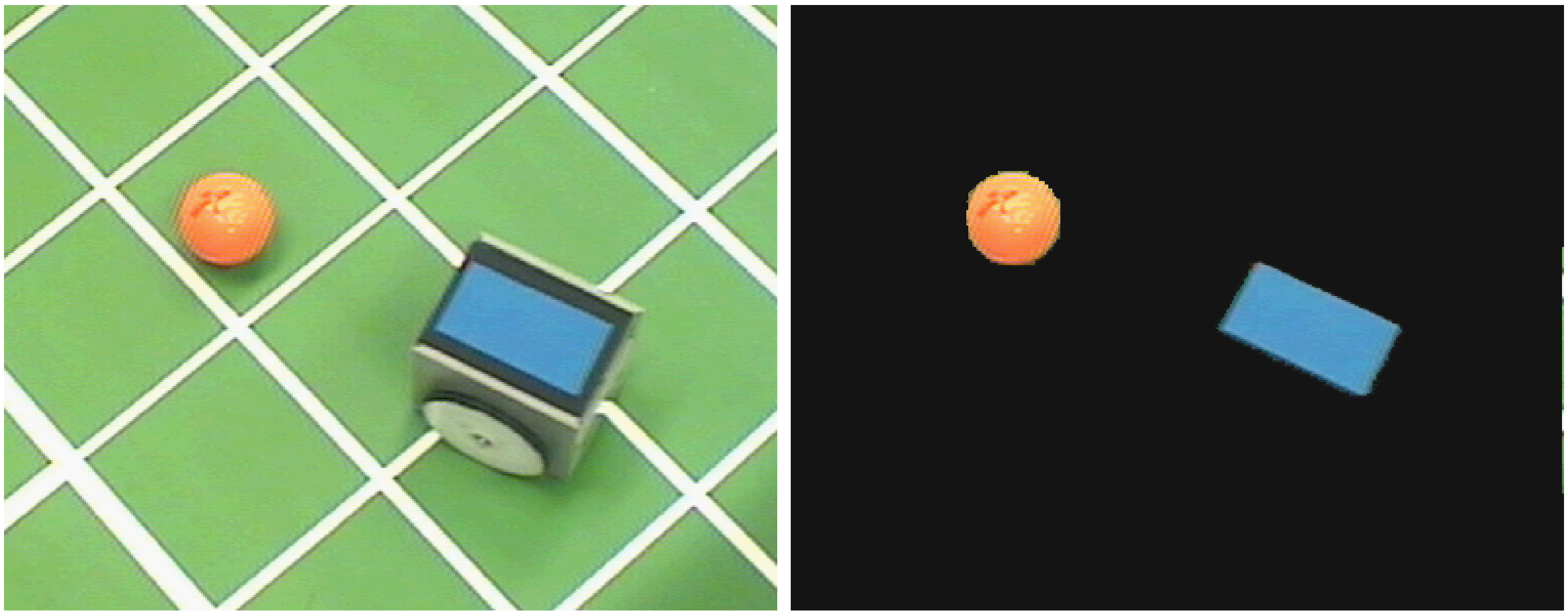

- Example: segmentation in robot soccer:

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

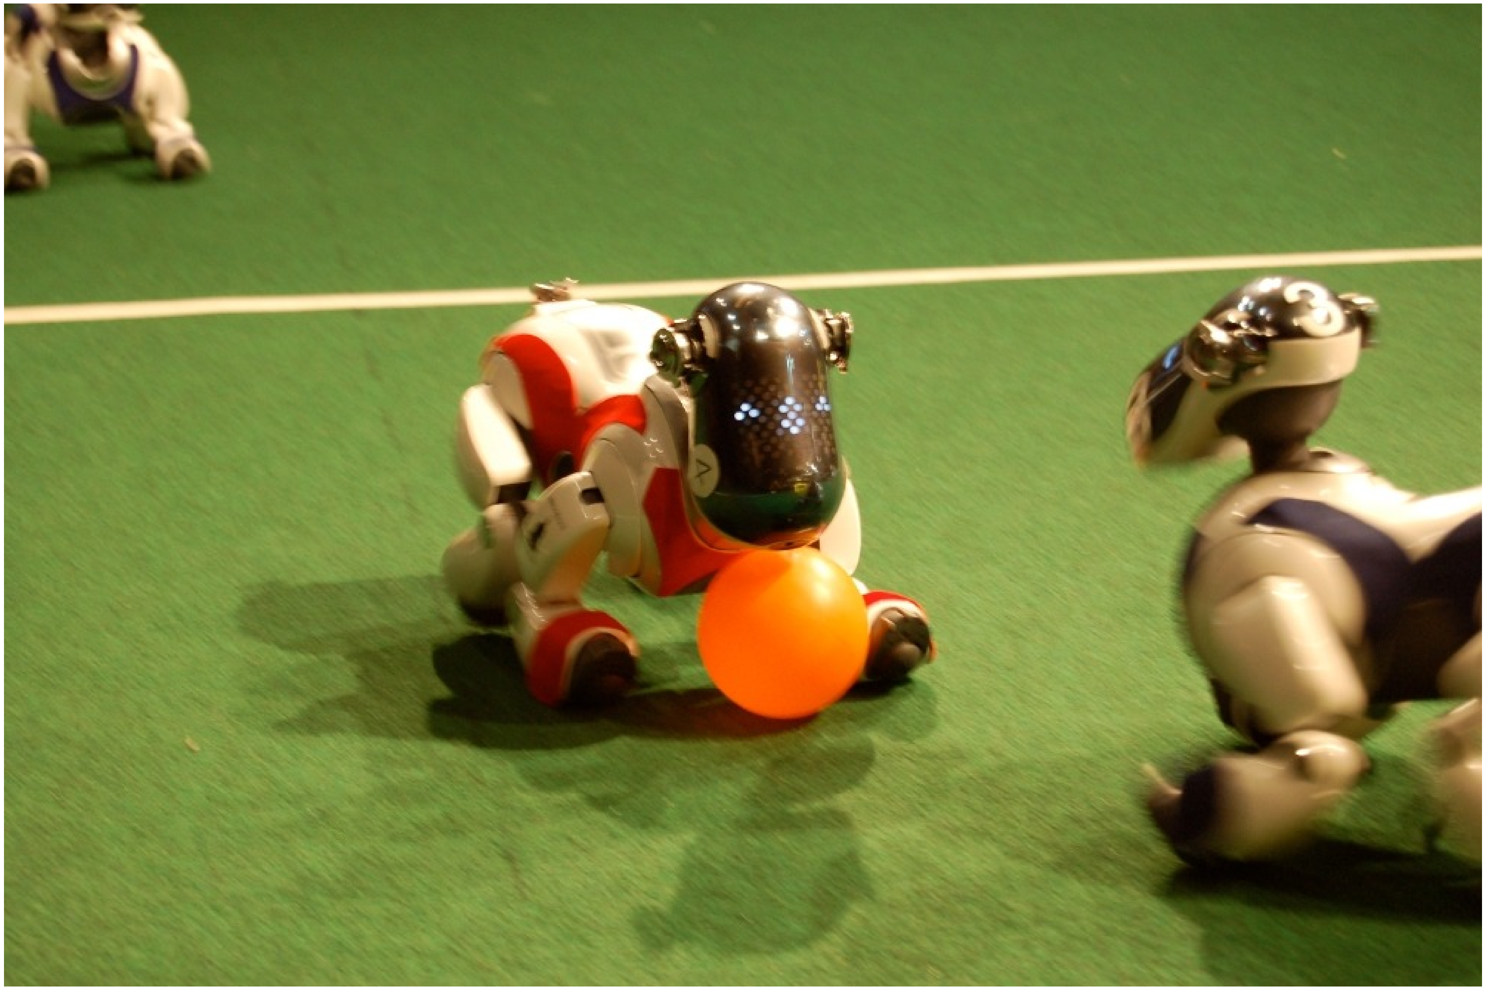

Edge detection

- 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.

• Gives us something like:


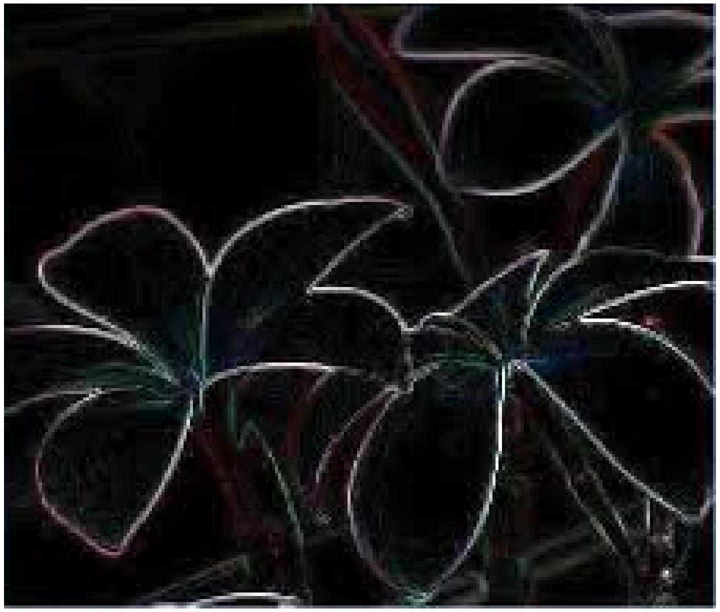

- Often edge detection gives us many mini-edges that need to be merged or removed:

- 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.
- They also lose more information (see `http://nskyc.com/`).
- Of course, all this processing is expensive, and slows down the speed of reaction of the robot.

Stereo vision

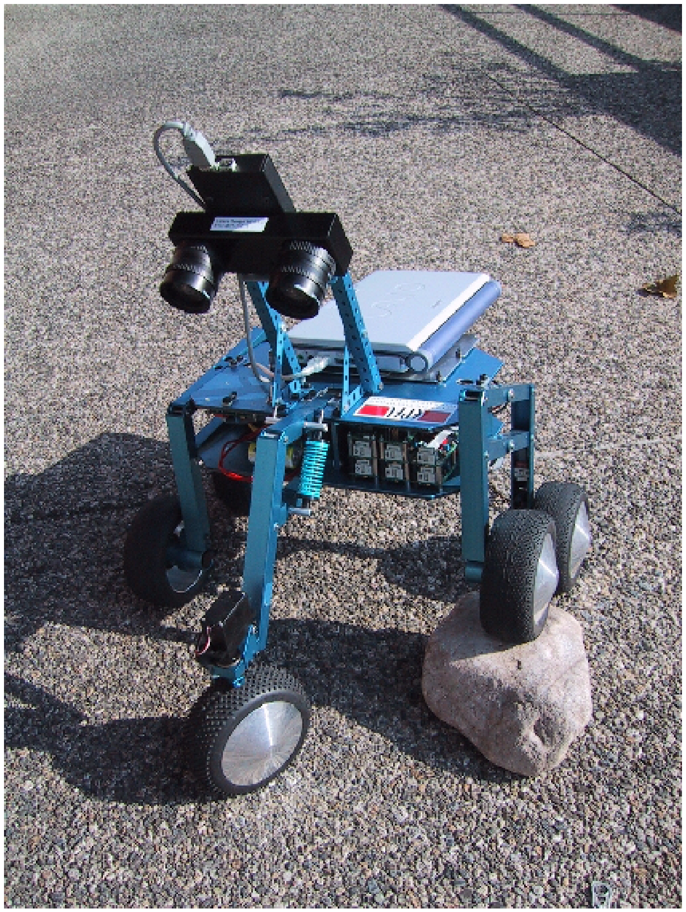

- Two cameras, spaced as widely as possible.
- Can get depth information if we can identify the common point(s) in two images.

$$
\begin{aligned}
x &= b\left(\frac{x_l + x_r}{2(x_l - x_r)}\right) \\
y &= b\left(\frac{y_l + y_r}{2(x_l - x_r)}\right) \\
z &= b\left(\frac{f}{x_l - x_r}\right)
\end{aligned}
$$

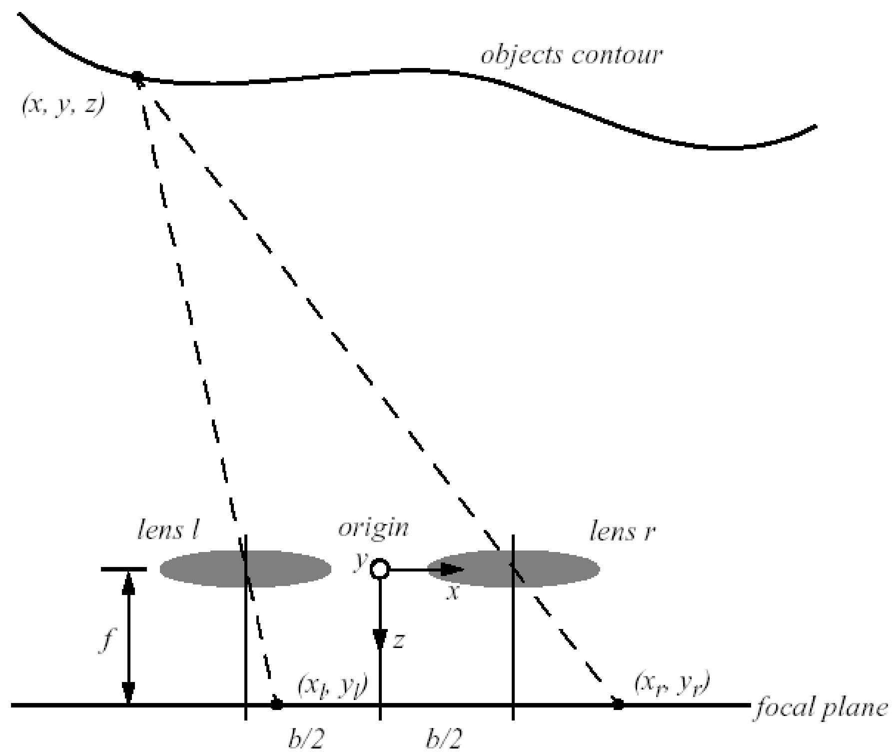

- The accuracy of the depth estimate increases with increasing baseline b.

Using stereo vision

- Also equipped with several other sensors.

The current frontier

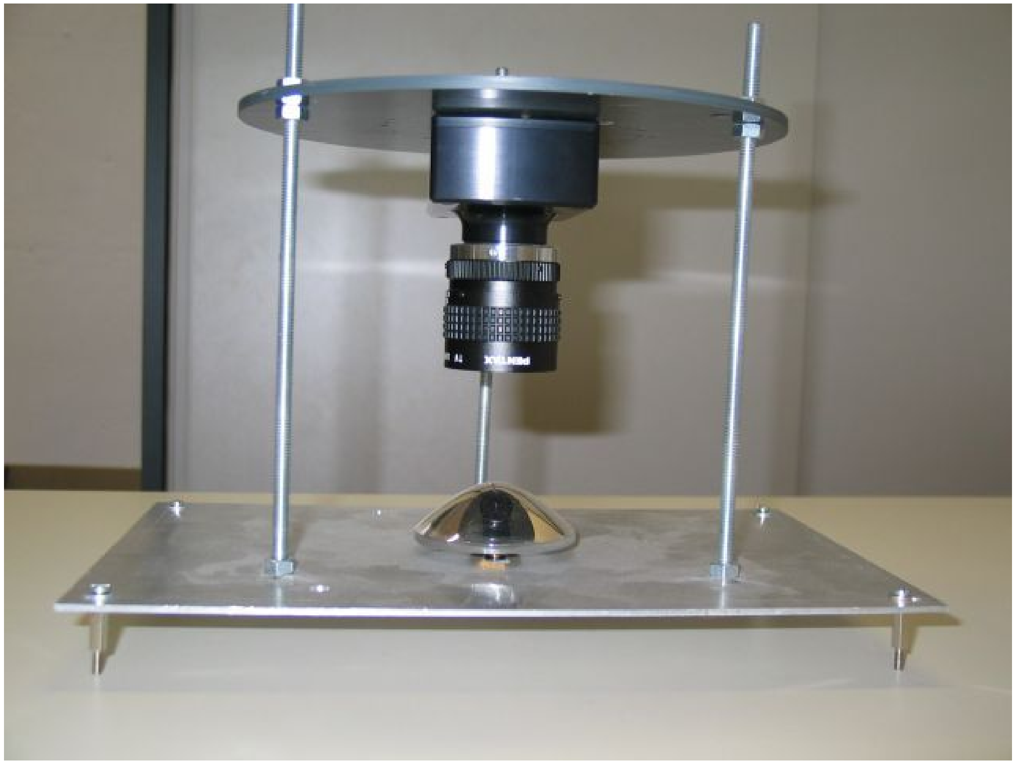

- Omnidirectional cameras allow robots to see all around them.
  - Usually mounted with the lens above the camera.

- Presents new challenges in machine vision.

Summary

- This lecture discussed *perception*, how the robot views its world.
- This means *sensors*.
- We looked at a range of different sensor types.
  - Including range sensors
- In particular we looked at cameras, and that led to a discussion of some simple techniques from computer vision.