

Robotics

Lecture 3: Sensors

See course website

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

Andrew Davison
Department of Computing
Imperial College London

Sensors: Proprioceptive and Outward-Looking

- Sensors are either *proprioceptive* (literally self-sensing) or *exteroceptive* (outward-looking).
- Proprioceptive sensors (such as motor encoders or internal force sensors) will improve a robot's sense of its own internal state and motion.
- But without outward-looking sensors a mobile robot is moving blindly. It needs them to, for example:
 - Localise without drift with respect to a map.
 - Recognise places and objects it has seen before.
 - Map out free space and avoid obstacles.
 - Interact with objects and people.
 - In general, be *aware* of its environment.

Sensor Measurements: Proprioceptive

- Sensors gather numerical readings or *measurements*. In the case of proprioceptive sensors, the value of the measurement \mathbf{z}_p will depend on (be a function of) just the state of the robot \mathbf{x} :

$$\mathbf{z}_p = \mathbf{z}_p(\mathbf{x}) .$$

- More generally, a proprioceptive measurement might depend not just on the current state but also previous states in the robot's history or the current rate of change of state. e.g. wheel odometry will report a reading depending on the difference between the current and previous state. A gyro which is part of an Inertial Measurement Unit (IMU) will report a reading depending on the current *rate* of rotation.

Sensor Measurements: Outward-Looking

- A measurement from an outward-looking sensor will depend both on the state of the robot \mathbf{x} and the state of the world around it \mathbf{y} :

$$\mathbf{z}_o = \mathbf{z}_o(\mathbf{x}, \mathbf{y}) .$$

- The state of the world might be parameterised in many ways; e.g. a list of the geometric coordinates of walls or landmarks; and may either be uncertain or perfectly known.

Single and Multiple Value Sensors

- Touch, light and sonar sensors each return a *single value* within a given range.
- Sensors such as a camera or laser range-finder return an *array* of values. This can be achieved by scanning a single sensing element (as in a laser range-finder) or by having an array of sensing elements (such as the pixels of a camera's CCD chip).

Touch Sensors



- Binary on/off state — no processing required.
- Switch open — no current flows.
- Switch closed — current flows (hit).

Light Sensors



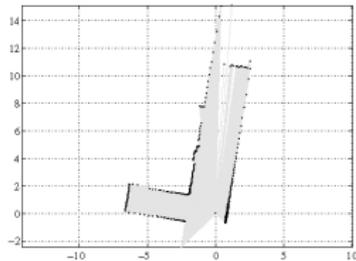
- Detect intensity of passive light incident from a single forward direction, with some range of angular sensitivity.
- Multiple sensors pointing in different directions can guide steering behaviours.
- The Lego sensors also have a mode where they emit their own light, which will reflect off close targets and can be used for following a line on the floor or quite effective short-range obstacle avoidance.

Sonar (Ultrasonic) Sensors



- Measures depth (distance) by emitting an ultrasonic pulse and timing the interval until echo returns. Sonar beam typically has an angular width of 10 to 20 degrees.
- Fairly accurate depth measurement (centimetre) in one direction but can give noisy measurements in the presence of complicated shapes. Maximum range a few metres.
- Robots sometimes have a ring of sonar sensors for obstacle detection.
- Especially important underwater where it is the only serious option beyond very short ranges.

External Sensing: Laser Range-Finder



- Like a sonar, measures depth using an active signal. Commercial Ladar sensors return an array of depth measurements from a scanning beam.
- Very accurate measurement depth (sub-millimetre) and works on most types of surface.
- Normally scans in a 2D plane but 3D versions are also available
- Rather bulky (and expensive) for small robots

External Sensing: Vision

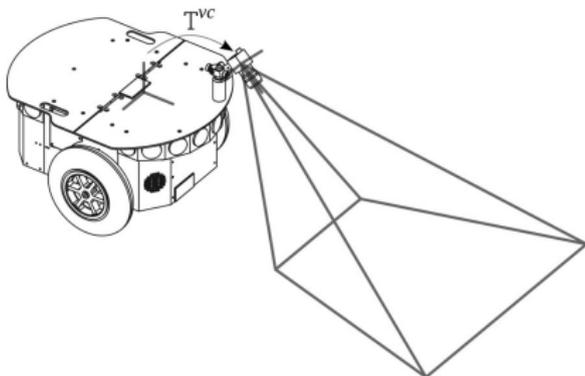


- The generalisation of a light sensor. A camera measures passive light intensity in many directions simultaneously by directing incident light onto a light sensitive chip.
- Returns a large, rectangular array of measurements.
- A single camera is a projective sensor which measures light intensity, rather than any direct information about geometry. From one image, we do not know if objects are “small and close” or “large and far away” unless we have some other information.

<https://www.youtube.com/watch?v=MMiKyfd6hA0>

Using a Camera as a Planar Sensor

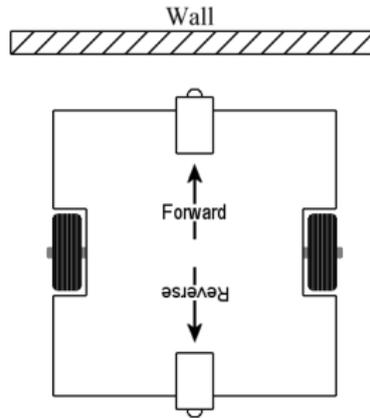
- We can do better if we have some extra scene knowledge.
- A useful case is if we know that the camera is observing the *ground plane*. Then every pixel in the image captured corresponds to a specific point on the ground plane.



The relationship between points on the ground plane and pixels in the camera image depends on:

- The position of the robot on the ground plane.
- The position and orientation of the camera relative to the robot.
- The *intrinsic calibration parameters* of the camera (focal length and principal point).

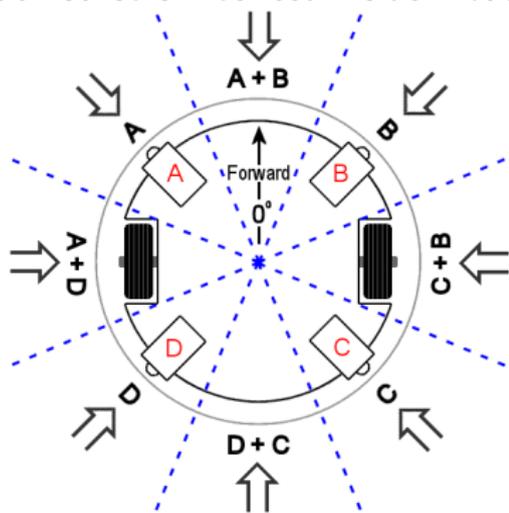
Touch Sensors for Bump Detection



- Sensor detects when an obstacle has been hit (last line of defence).
- Demands immediate reaction — evasive manoeuvre, or stop forward motion at least.

Multiple Touch Sensors — Where was I hit?

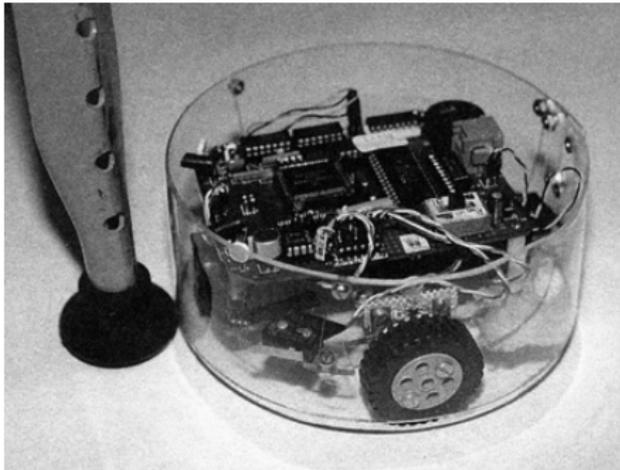
Touch sensors mounted inside 'floating skirt' around circular robot:



A	B	C	D	Sector	Centre
1	1	0	0	337.5° to 22.5°	0°
0	1	0	0	25.5° to 67.5°	45°
0	1	1	0	67.5° to 112.5°	90°
0	0	1	0	112.5° to 157.5°	135°
0	0	1	1	157.5° to 202.5°	180°
0	0	0	1	202.5° to 247.5°	225°
1	0	0	1	247.5° to 292.5°	270°
1	0	0	0	292.5° to 337.5°	315°

- Four sensors give the ability to measure eight bump directions.

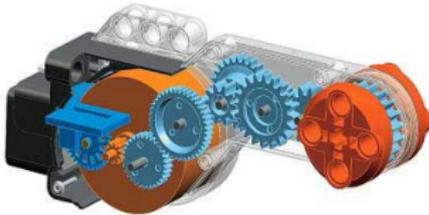
Strategies After a Collision



- Try to move around object: reverse, and try to go around (e.g. turn to a fixed angle from hit and proceed).
- Random bounce: rotate through random angle and head off straight again until next collision.

Servoing

- Servoing is a robot control technique where control parameters (such as the desired speed of a motor) are coupled directly to a sensor reading and updated regularly in a *negative feedback loop*. It is also sometimes known as *closed loop control*.
- Servoing needs high frequency update of the sensor/control cycle or motion may oscillate.
- The way that a motor controls its motion using encoder feedback is one example of servoing and negative feedback. This concept can also be used with outward-looking sensors.

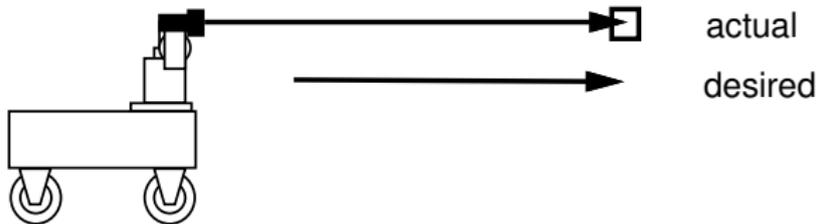


Proportional Control using an External Sensor: Servoing

- In servoing, a control demand is set which over time aims to bring the current value of a sensor reading into agreement with a desired value.
- Proportional control: set demand proportional to negative error (difference between desired sensor value and actual sensor value):
e.g. set velocity proportional to error:

$$v = -k_p(z_{desired} - z_{actual}) ,$$

where k_p is the proportional gain constant. (Note that since this is a different control loop, k_p will not have the same value as in last week's motor tuning but will need to be individually adjusted through trial and error.)



- Proportional control is a special case of more general *PID Control* (Proportional, Integral, Differential).

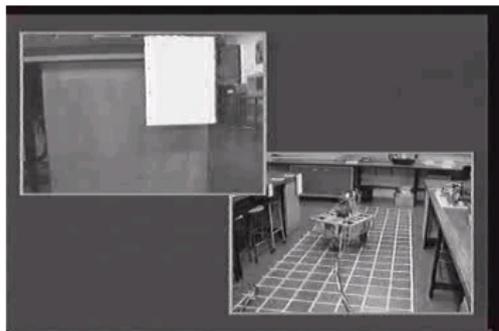
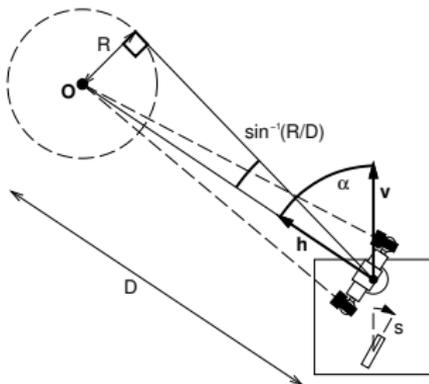
Visual Servoing to Control Steering

- For a robot with a tricycle or car-type wheel configuration.
- Simple steering law which will guide robot to collide with target:

$$s = k_p \alpha$$

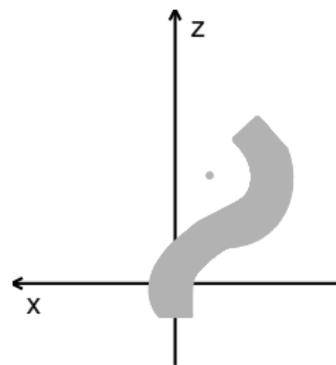
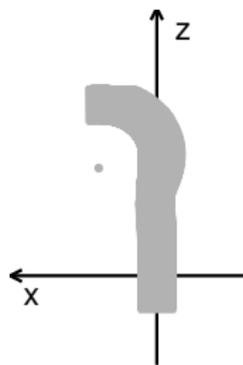
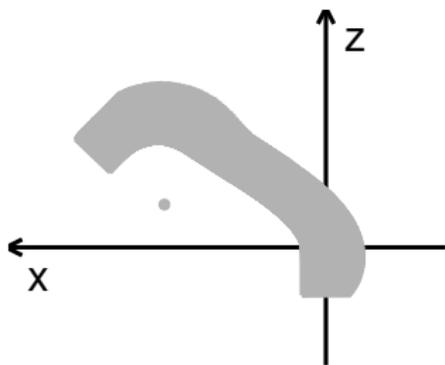
- Steering law which will guide robot to avoid obstacle at a safe radius: subtract offset:

$$s = k_p \left(\alpha - \sin^{-1} \frac{R}{D} \right)$$

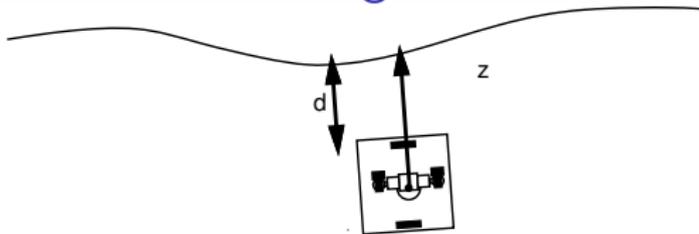


- <https://youtu.be/Ce2FrSBKwB8>

Visual Servoing Trajectories



Wall Following with Sonar



- Use sideways-looking sonar to measure distance z to wall.
- Use velocity control and a 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 :

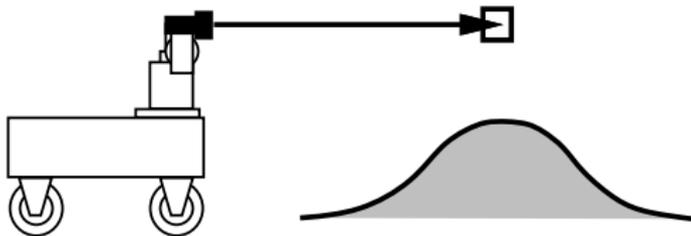
$$v_R - v_L = K_p(z - d)$$

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

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

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

Probabilistic Sensor Modelling



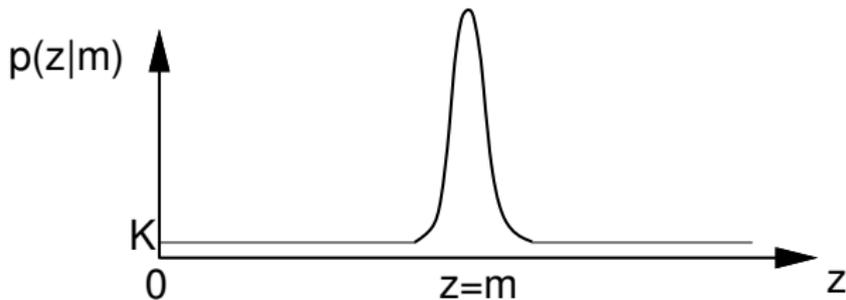
- Like robot motion, robot sensing is also fundamentally *uncertain*. Real sensors do not report the exact truth of the quantities they are measuring but a perturbed version.
- Having characterized (modelled; calibrated?) a sensor and understood the uncertainty in its measurements we can build a probabilistic measurement model for how it works. This will be a probability distribution (specifically a *likelihood function*) of the form:

$$p(\mathbf{z}_o | \mathbf{x}, \mathbf{y}) .$$

Such a distribution will often have a Gaussian shape.

Robust Likelihood for Sonar Measurements

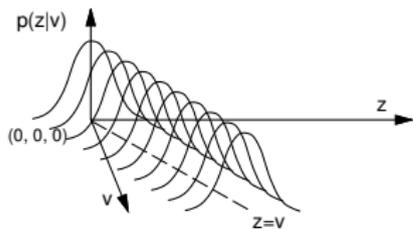
- We will return to this in later lectures on probabilistic robotics, but a suitable likelihood function for a sonar sensor is as follows: it says 'what is the probability of obtaining sensor measurement z given that the ground truth value I expect is m '
- This distribution has a narrow Gaussian band around the expected value, plus a constant additive band representing a fixed percentage of 'garbage' measurements.



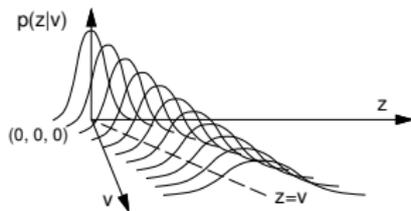
$$p(z|m) \propto e^{-\frac{(z-m)^2}{2\sigma_s^2}} + K$$

Likelihood Functions

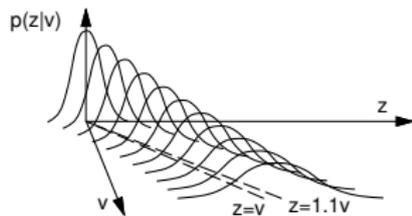
- A likelihood function fully describes a sensor's performance; in particular the fact that the uncertainty in its performance will often be variable.
- $p(z|v)$ is a function of both measurement variables z and ground truth v and can be plotted as a probability surface. e.g. for a depth sensor:



Constant Uncertainty



Growing



Systematic Error (Biased)

Sensor Filtering/Pre-processing/Abstraction

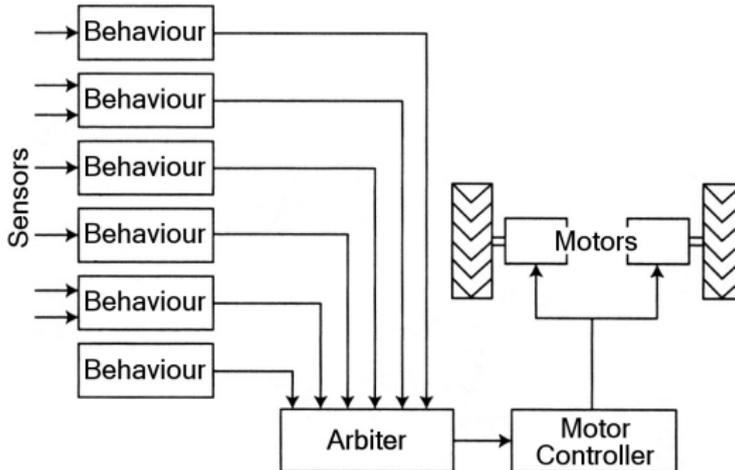
- Most generally, we can design algorithms (usually probabilistic) which will use sensor measurement in a **direct** way. The algorithm itself (e.g. Monte Carlo Localisation) will model and take account of the uncertainty in a sensor's performance.
- For simpler algorithms such as servoing/wall following we will often need to do some pre-processing of raw sensor readings to turn them into a more useful form. Common type are:
 1. Temporal filtering: e.g. smoothing or finding the median of the last few measurements of a sensor which reports a single reading such as a sonar. This is good at reducing the effect of occasional large **outliers**, which are measurements with an unusually large amount of error.
 2. Geometric fitting (can be called feature detection) to data from a sensor which reports an array of measurements, such as a laser range finder or scanning sonar, where we might fit geometric shapes such as straight lines or corners to the measurements and output the parameters of those shapes rather than the raw measurements.

Combining: Sensing/Action Loops

- No modelling and planning! Consider each local 'servo'-like sensing-action loop as a **behaviour**.

Sense → Act

- The challenge is in combining many behaviours into useful overall activity: see TR Programs, Subsumption, Braitenberg vehicles.



Combining Sensors: World Model Approach

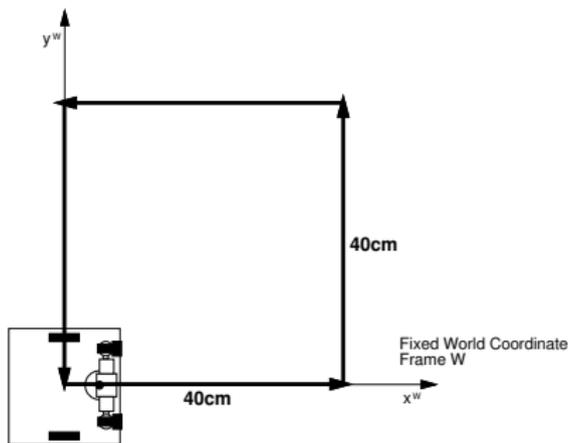
- Capture data; store and manipulate it using symbolic representations.
- *Plan* a sequence of actions to achieve a given goal
- Execute plan.
- If the world changes during execution, stop and re-plan.
- Powerful, but computationally expensive and complicated!



Shakey: one of the first mobile robots.

- Probabilistic state inference and planning is the modern version of this able to cope with uncertainty in sensors.

Practical 1: Accurate Robot Motion



- Today's practical is on accurate robot motion. How well is it really possible to estimate robot motion from wheel odometry?
- **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.