- Friday, February 2, 1996
  Everything That You Would Like To Know About Colour Vision But Were (previously) Afraid to Ask!

  1 Here are the minutes for the YORKVIS meeting, the next meeting will be on March 1st and will be hosted by Peter Kaiser. He has decided to try a new style in which he would like to be asked questions BEFORE he gives the talk. He would like to talk about colour vision and he would like to take as his starting point EVERYTHING THAT YOU WOULD LIKE TO KNOW ABOUT COLOUR VISION BUT WERE (previously) afraid to ask!! Please email him your questions ASAP: pkaiser@yorku.ca

  2 Michael Jenkin gave the first-Friday-in-Feb YORKVIS discussion group. He talked about ROBOTS and how they sense the world with an emphasis on the use of stereo-vision in robotics. This is my summary of his talk, please forgive the shorthand of using 'you' when I mean 'a robot'....

  3 Mobile robots have lots of applications: security, going into dangerous places, doing jobs that only over-paid people are available to do (eg. delivering food in hospitals, lecturing at universities -- no, no, only kidding), bomb disposal, exploring hostile environments (eg, volcanoes at the north and south Poles: the Dante project), or just fun. Robots of many different designs, sizes and function have in common that they need to move and sense the environment. Their sensing is partly to enable them to move and navigate successfully. They can move in a variety of ways: using conventional wheels, swivel wheels, legs... and can sense using a range of sensors.

  4 One of the most taxing computational tasks is maintaining an estimate of the robot's pose that is getting the robot to know where it is and where it is pointing (posed). It is not possible to navigate perfectly using only INTERNAL sensors (eg. measuring wheel rotation) because unpredictable errors will always accumulate between where your measures say you are and where you are really. For example, the slippage of wheels on the ground will eventually produce a mismatch between wheel circumference times the number of rotations made and the ACTUAL distance covered. Also things change in the environment, either accidentally or maliciously, that need to be taken into account. Therefore sensors of the external world are REQUIRED to calibrate the position of a robot relative to the outside world.

  5 It is not NECESSARY always for a robot to have an internal map of space to navigate. Although that is one solution. Alternatives include using a invisible railway lines in the form of a hidden magnetic strip. But it is hardly 'intelligent' behaviour to run along railway tracks. To act intelligently you need sensors to assess environmental information.

  6 Types of sensors:

  ULTRASOUND: transmitter/receiver pairs (like on bats). measuring time for echo to return. not good for close distances, or far. vulnerable to cross-talk from other transmitters, vulnerable to confusion by surfaces that do not reflect the waves in a direct way.

  INFRARED: transmitter/receiver pairs (like on TV remote controls) looks for reflected infrared light, good for detecting there is a close object present, doesn't tell you much about it. Easily flooded out by other light or confused by absorbent surfaces.

  RADAR: dead expensive, but has some use

  LASAR: transmitter/receiver pairs. using very accurate measures of 'time-of-flight' of the emitted light can tell you about distance. Takes about 3 secs to work out each sample. Potentially dangerous. you can get faster systems. also slower less accurate systems

  SCANNING LASAR: transmitter/receiver pairs. Looking for the emitted spot (like the infrared). Can triangulate a detected, reflected spot and get accurate information about the distance and location of the reflecting surface. Vulnerable to misalignment.

  In short all these sensors are limited in some or several ways. Combinations of these and others can be useful, but an additional candidate, based on how humans do the job is:

  STEREOVISION: using two or more cameras and looking at the DIFFERENCES in the images the two cameras return. Calculations are the tricky part. In order to find DIFFERENCES, you need first to find which points in one image match with which points in the other image so you can measure their differences. This is known as the correspondence problem (both in this mathematical form and in the biological task of extracting information from having two eyes). A correlation of gaussian windows was briefly alluded to as a possible way of solving this task. It takes ages for this (and related algorithms) to do the job and only works very well over a restricted range of distances. But that is if you want to extract all the 3D info in the whole image, that is about the whole world. Generally the robot (or the human, for that matter) is only concerned about a restricted part of the visual world. For example a robot might be very concerned about the floor.

  7 In the images of two cameras (eyes) there will be points in the world whose images will fall on the same parts (corresponding points) of the detector surfaces in the two cameras (eyes). The position of these points define what is called a horopter. The horopter has a definite shape and

is basically a circle going through the two eyes and the fixation point (the Vieth Muller circle) and also vertical line, also going through the fixation point. An algorithm that only looks at the horopter can be more efficient.

8 By having cameras that can move in all three dimensions (up/down, left/right, torsion), it is possible to adjust the orientation of the horopter. Thus it can be adjusted to match a surface that you are interested in eg. a vertical wall (the default setting) or the floor (by tilting and torting). Then you can just look to see what is 'on the horopter' ie. on the surface of interest. Although the not-the-circle part of the horopter is only a straight line, it is possible to extend the algorithm to look at a plane orthogonal to the plane of the cameras that includes the horopter. This is done by warping space (or the internal representation of it, at least).

9 Using these sorts of tricks to narrow down what is being asked about 3D space by robots makes stereovision a potentially important sensor in their array of possible sensors.

10 See you first Friday in March!!

Michael Jenkin
York University