

Robot Navigation by Light

Ulrich Nehmzow and Brendan McGonigle

Robotics Laboratory
Laboratory for Cognitive Neuroscience, Psychology Department
Edinburgh University
Appleton Tower
Edinburgh EH8 9LE
ulrich@castle.ed.ac.uk, B.O.McGonigle@festival.ed.ac.uk

Abstract

The ability to return to a base location is crucial for moving animals; so it is for mobile robots. Insects like bees and ants achieve this task by a combination of different mechanisms: coarsegrained dead-reckoning, using a global reference frame such as the sun, takes them back close to their nest. Once near the desired location, either a search behaviour (ants) or a navigation by local landmarks (bees) is initiated and takes the animal home.

We present a navigation system for a mobile robot that shows similarities to this solution. Using a distant light source (the sun, for example), the robot is able to return back to base within about 10% of the total distance travelled. The robot then switches to a finegrained local navigation, exploiting the structure of its environment and returns to its base by means of local landmarks.

1 Introduction

Animals move around their niche, and as a consequence of such locomotion create a number of problems. The necessary competences can be arranged in a logical hierarchy ([McGonigle 91]). Initially, the most immediate obligation of a moving system is not to collide with rigid objects. These competences are achieved at a very early stage in the ontogeny of all species including human primates.

Such competences are not, however, designed to solve a further problem which arises from locomotion: that of relocating in a principled way, to a region of space once occupied by the agent and out of range of its sensors. For animals, particular locations within their niche may be important sites for food. Finding these places directly, based on locative memory, for instance, materially reduces the costs of foraging as compared with one controlled by some random search process.

The same problem exists in the case of an artificial agent. It is one thing to avoid objects in the work space; another to recover a privileged location which has some vital importance (such as a fuel source). It is this problem which we address specifically in this paper.

As in the case of biological systems we need first to consider the niche and the operating conditions under which some form of navigation system can be implemented. In open daylight conditions, for example, rather than in a subterranean passageway with physical constraints on movement, a moving agent is enabled to detect which aspects of the array change continuously with motion from those aspects which do not. Distant cues, for instance, remain relatively stable with locomotion and afford a relatively invariant reference system within which the agent can reliably orientate. Under daylight conditions many species use some aspects of sunlight as the basis for such orientation. Honey bees *apis mellifera* and desert ants *cataglyphis bicolor*, for example, can navigate as long as there is some polarised light available ([Gould & Gould 88, Waterman 89] and [Wehner & Räber 79]). Other species can use the position of the sun (for example starlings). Even rodents, well adapted for subterranean route learning, can use distal visual clues as a basis for determining their general orientation in space.

A stable frame of reference is, of course, vital to help relocate to a position in space outside the immediate sensor range of the agent. Whether it is ever sufficient, however, in itself is unlikely, especially over large distances. Gallistel ([Gallistel 90, p39ff]) reports the average human error of navigation as from 5-10 percent of the total distance travelled, and a similar error has been observed in desert ants *cataglyphis bicolor* ([Wehner & Srinivasan 81]).

This suggests that the solution must be ultimately based on a shift in the granularity of the processes involved: One part is to provide a relatively coarsegrained schema enabling the subject to return to the neighbourhood of the sought after location; another is to then search that neighbourhood for other sources of information¹.

¹Putting a key in a door lock is a common place example for this. Finding the door with respect to a general reference frame involves a much more global referencing system than is required for the precision movements necessary to put the key accurately in the lock once the appropriate door has been located.




Figure 1: The Edinburgh R2.

This then is the basis for our approach here. The particular task looked at here is that of exploration and subsequent return to the base. Leaving from an arbitrary starting position (which is defined as the base) the robot moves in its environment, avoiding obstacles it encounters. Upon command by the experimenter² the robot returns to its base by means of dead reckoning, based on phototaxis.

In order to be able to achieve this task, a uniform light gradient is required. Such a light gradient often exists in laboratory environments (window front at one side of the room), if it doesn't, it can be created by artificial light. We also have conducted experiments outdoors, which show that sunlight is suitable for the type of navigation presented here.

2 Experiments with the Differential Light Compass

2.1 The Edinburgh R2

The robot we used for the experiments described here is shown in figure 1.

²It is conceivable that the decision to return is taken by the robot itself, for example upon having grasped an object.

It is a mobile robot of roughly quadratic shape (25cm x 25cm), called the "Edinburgh R2". The robot is controlled by a distributed control architecture, using six 68HC11 microcontrollers (one master and five slaves, used to control sensors and actuators). It has a gripper with two degrees of freedom. Eight infrared proximity sensors are mounted around the robot at a height of 25 cm. These detect white objects as far away as 40 cm, objects coated with reflective tape can be sensed from distances of up to 150 cm. Additionally, the robot has six light sensors (five around the perimeter, one pointing upwards), as well as two break beam sensors and one colour sensor in the gripper. The light sensors are based on cadmium sulphide light dependent resistors; by means of a voltage divider a light-dependent voltage is obtained and fed into the analog input of the 68HC11 microcontroller of the light-sensor slave.

Mounted at the front of the robot at a height of 5 cm are five tactile sensor pads, two additional tactile sensor pads are fitted inside the gripper. Through wheel encoders the robot can also sense the velocity of each of its two wheels. The speed of the robot can be controlled by software and goes from -40 cms^{-1} to $+40 \text{ cms}^{-1}$.

2.2 Dead Reckoning

If current heading with respect to some reference direction and current speed (or distance travelled) are known, the current position of an agent can be computed as follows (see also figure 2):

$$dx = \cos(\alpha_i)\bar{v}\Delta t, \quad dy = \sin(\alpha_i)\bar{v}\Delta t \quad (1)$$

$$x(t+1) = x(t) + dx, \quad y(t+1) = y(t) + dy. \quad (2)$$

α_i is the current heading of the robot, Δt is the time interval between measurements and \bar{v} is the average speed of left and right motor, which is assumed to be constant during Δt .

The current heading can either be *measured* by referring to a reference direction (through a magnetic sense, or, as in our case, through light sensors), or it can be *estimated*. Both methods introduce errors. Estimating angles soon becomes very unreliable as the robot performs several turns in a row (due to the accumulation of error). Using an external reference, on the other hand, removes the accumulative error in angle measurement completely and leaves only a measurement error. It is therefore more robust.

On returning home, direction and distance are given by

$$\tan(\gamma) = \frac{y}{x}, \\ l = \sqrt{x^2 + y^2}.$$




Figure 2: Dead reckoning, not using reference landmarks.

2.2.1 The Differential Light Compass in the Edinburgh R2

The dead reckoning mechanism used in the Edinburgh R2 follows the basic principle outlined before: Position estimation is performed by dead reckoning, using a distant light source (our laboratory window) as a reference direction. In order to reduce computation, a constant speed is assumed. This assumption is reasonable because the robot's speed is controlled by independent PD controllers.

The Edinburgh R2 is equipped with six light sensors. Sensors zero to four are mounted on the robot as shown in figure 3 (sensor five points upward and can be used to determine the ambient light level or detect ceiling lights. It is not used in the experiments described here).

Rather than measuring current heading and applying equation 1, we compute dx and dy directly:

$$dx = \frac{S1 - S4}{|S1 - S4| + \left| \frac{S2+S3}{2} - S0 \right|}, \quad dy = \frac{\frac{S2+S3}{2} - S0}{|S1 - S4| + \left| \frac{S2+S3}{2} - S0 \right|} \quad (3)$$

where $S0 \dots S4$ denote the signal values of light sensors zero to four, dx the heading in x-direction and dy the heading in y-direction (scalars). This provides more accurate results at lower computational costs, as no trigonometric computations at all have to be performed. Because dx and dy are dependent upon the difference in light detected by opposite light sensors, we call this mechanism *differential light compass*.

During the outward journey (during which the robot avoids obstacles using its infrared sensors), the current position of the robot is estimated in discrete time steps according to equation 2.

To return back to base, the robot turns on the spot until both dx and dy




Figure 3: Arrangement of light sensors on the Edinburgh R2.

(see equation 3) are such that the robot is heading towards the starting location. The robot then starts moving, again avoiding obstacles on the way, and again updating its current position according to equation 2. Once x and y -coordinate lie below a predefined threshold, the robot stops. Control is then handed over to a finegrain navigation system that homes by means of a passive beacon which marks the starting location.

3 Experimental Results

3.1 Experiments

Figure 4 shows typical outward and return paths of the robot (no beacon at base). The points marked 'S' denote the starting location of the robot, at points 'R' the robot was instructed to return to the starting location.

The top half of figure 4 shows paths where an obstacle is encountered on the outward journey, but not on the return trip. The bottom half shows the more complicated case where the direct return path is blocked by an obstacle (the indicated obstacle is not present during the outward journey). The final course followed by the robot is the result of two competing behaviours, that of obstacle avoidance and that of homing.

The positions shown in the diagrams denote the center of the robot. The actual robot (and therefore its sensors) extends 25 cm to either side of the marked points. This means that the starting location was within sensor range from all return positions, i.e. a beacon placed there would have been detected.

The experiment shown in figure 5 investigates the accuracy of navigation for two cases: a) the current heading of the robot is estimated, and equation 1 is used (return positions for this case are marked by letters), and b) the current heading of the robot is measured by using the light sensors, and equation 3 is




Figure 4: Typical paths of the robot.

applied (the return positions of the robot for this case are marked by numbers).

The figure shows that all four return positions (using the DLC) were within sensor range of the starting location, whereas for the runs without DLC they were not. This is due to the accumulated error when each turn angle of the robot is only estimated.

3.2 Measuring the Navigational Displacement

Experiments with desert ants show that the typical navigational error is about 10% ([Wehner & Srinivasan 81]). Similarly, dead reckoning navigation in humans is subject to an error in the range of 5 to 10% of the total distance travelled ([Gallistel 90]).

In order to estimate the precision at which the robot would be able to navigate according to equation 3, the robot was made to follow a path as shown in figure 6, of about 570 cm length.

The return positions in eleven separate runs are shown in table 1. Also shown is the displacement in per cent of total distance travelled.

To interpret table 1, it should be borne in mind that the sensors of the robot are 12 to 25 cm away from the center, and have a range of another 150 cm for passive, reflective beacons. In other words: the beacon was visible from all return positions in the eleven runs of the experiment.

This data can be used to estimate the maximum trip length over which the

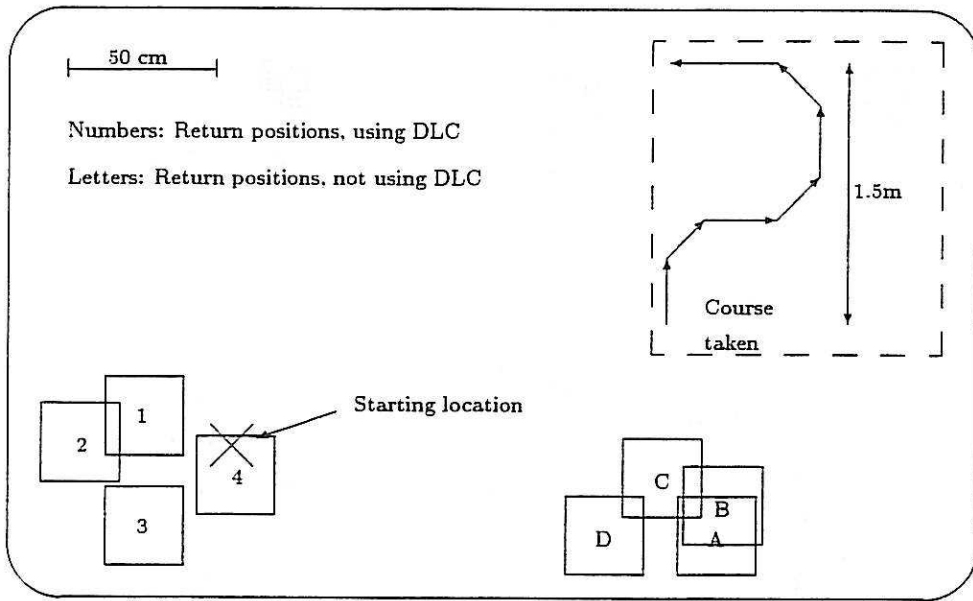


Figure 5: Navigation with and without DLC.




Figure 6: Course steered to measure navigational displacement.

x [cm]	70	0	20	-30	-10	0	-5	-25	-25	-25	-30	-5,5
y [cm]	30	60	35	0	25	75	-10	-10	0	60	-50	19,5
d [cm]	76	60	40	30	27	75	11	27	25	65	58	44,9
Error [%]	13	10,5	7	5,3	4,7	13,2	1,9	4,7	4,4	11,4	10,2	7,9

Table 1: Navigational displacement, using the DLC and dead reckoning.

navigational system would still work. Given that the beacon is visible from distances of up to 150 cm, and a navigational displacement of about 8% occurs, the maximum travel distance that still allows homing must not exceed 18 m; however, due to physical constraints the travel distances we used seldomly exceeded 6 m.

3.3 Local Fluctuations in Light Gradient

Local fluctuations in light gradient occur for instance if the robot moves into shadows, or approaches a window off-center. As long as a gradient exists at all this does not affect navigation, provided the robot returns along the exact same path it used during the outward journey. However, as this is impossible, local fluctuations in light gradient do introduce an error. In our experiments we have found that navigation based solely on the differential light compass is indeed dependent on a constant light gradient, a requirement that is not always met indoors. For trip lengths of about four to six sensor ranges (4-6 m) we have found, however, that if the starting location is marked by a passive beacon, the robot reliably returns to the starting location, despite the fact that local fluctuations in light gradient are encountered. A reliable navigational competence therefore requires the use of both local landmarks and reference landmarks. This, in fact, can be perceived in desert ants, who use dead reckoning until they are near the nest, and then use local navigation to find the actual entrance ([Gallistel 90, p.61], [Wehner & Srinivasan 81]). The same is true for gerbils, who will even run past their (displaced) nest if their dead reckoning navigation indicates they are nowhere near the nest ([Gallistel 90]).

4 Discussion

We report our first attempts to solve an inevitable problem arising out of locomotion: to return to a previous location without benefit of direct sensing of the location in the first instance. This solution we based on a two-fold task decomposition: a global reference frame first enables the robot to return close to its starting location; it then navigates by means of local landmarks (a passive beacon in our case). In our particular case, the final search was made easy by having only one object available. However, provided the differential light compass can deliver the agent to the search area within which direct sensing is a plausible strategy, then there is no reason why we cannot make the object selection criteria more exclusive to cater for a multiple object scenario. Given the colour sensors on the grippers, we can readily fine tune the system to select only landmarks of a particular colour. In this way, other distractor objects can be avoided, even within the crucial search space.

It is also clear that there are many other ways of implementing solutions to the navigation problem we pose. Location recognition whilst following canonical paths is one way ([Nehmzow 92]). Route learning is another powerful yet low level solution which many animals use³. There is no reason in principle why a

³Particularly those which occupy subterranean niches, such as rodents in burrows.

succession of landmarks could not be learnt by a robot in an invariant chain in which each component is the trigger to search for its successor. In this way, an extensive range of exploratory behaviours could be controlled avoiding the cumulative error problems inherent in dead reckoning solutions. Indeed, as humans do, we could and should incorporate both systems (dead reckoning *and* route following) in the same robot!

Acknowledgements

The work reported here is funded by the Medical Research Council, grant number G9110835.

References

- [Gallistel 90] C.R. Gallistel, 1990. *The Organisation of Learning*. MIT Press, Cambridge Mass. and London, England.
- [Gould & Gould 88] James L. Gould and Carol Grant Gould, 1988. *The Honey Bee*. Scientific American Library 1988, New York.
- [McGonigle 91] Brendan McGonigle, 1991. *Incrementing Intelligent Systems by Design*. In: Jean-Arcady Meyer and Stewart Wilson (eds.), 1991. *From Animals to Animats*, Proc. 1st Intern. Conf. on Simulation of Adaptive Behaviour. MIT Press, Cambridge Mass. and London, England.
- [Nehmzow 92] Ulrich Nehmzow, 1992. *Experiments in Competence Acquisition for Autonomous Mobile Robots*. Ph.D. Thesis, Edinburgh University, 1992.
- [Waterman 89] Talbot H. Waterman, 1989. *Animal Navigation*. Scientific American Library, New York.
- [Wehner & Srinivasan 81] R. Wehner and M.V. Srinivasan, 1981. Searching behaviour of desert ants, genus *cataglyphis* (formicidae, hymenoptera). *Journal of Comparative Physiology* 1981, 142, 315-338, after [Gallistel 90].
- [Wehner & R aber 79] R. Wehner and F. R aber, 1979. Visual spatial memory in desert ants *Cataglyphis bicolor*. *Experientia* 35 (1979)