An Investigation into Obstacle Avoidance as an 'Emergent' Behaviour from two Different Perspectives

Verena V. Hafner, Hanspeter Kunz, and Rolf Pfeifer
Artificial Intelligence Lab,
Department of Information Technology,
University of Zurich, Switzerland,
{*vhafner, hkunz, pfeifer*}*@ifi.unizh.ch*

Abstract

In this paper, we investigate in which ways a particular behaviour can be achieved with simple mechanisms without being explicitly encoded. It is suggested that enrichening the overall behaviour of an agent in such a way simplifies the creation of intelligent agents. We focus on obstacle avoidance as one particular kind of such a behaviour by an agent. We demonstrate that obstacle avoidance can be achieved on both a strategic level, caused by the intrinsic structure of the environment, and on a morphological level, caused by the properties of the agent's body. The underlying principles are studied on the example of local visual homing and an agent without sensory-motor coupling.

1 Introduction

Whereas in nature many types of behaviours interact smoothly with each other, it is still a great challenge for robotics and artificial intelligence to create even a limited number of behaviours for an artificial agent. Since many behaviours have common subsets (e.g. search and exploration), and cannot be easily separated, it is obvious that pursuing a purely modular approach cannot be the whole answer.

In this paper, we investigate in which ways a particular behaviour – i.e. obstacle avoidance – can be achieved with simple mechanisms without being explicitly encoded. We will refer to such a behaviour as 'emergent' behaviour. It is suggested that enrichening the overall behaviour of an agent in such a way simplifies the creation of intelligent agents.

Obstacle avoidance is not a recent invention but has evolved very early after the first living creatures started to move. It can be achieved by a number of different sensors. The most common is vision, others are echolocation or somatosensory cues. In the area of robotics, obstacle avoidance is considered as one of the most basic behaviours, and was already present among the earliest autonomous mobile robots.

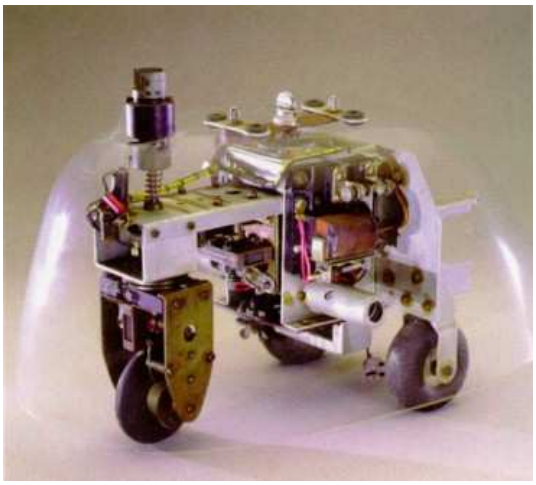

Figure 1: Second generation turtle (tortoise) designed by Grey Walter around 1945. Its main two behaviours were obstacle avoidance and light seeking. [picture taken from the Grey Walter Online Archive (Holland, 1996)]

At the end of the 1940s, the neurophysiologist Grey Walter built some of the first autonomous robots 'Elsie' and 'Elmer' (see Figure 1). Their mechanics consisted of two propulsion wheels and one steering wheel; as sensors, he used one light and one touch sensor. The robots were completely analogue, and could perform two main actions: obstacle avoidance, by retreating on contact, and light following. By placing a lamp on each turtle's shell, a kind of 'social behaviour' emerged from the interaction of the two turtles.

A similar approach was taken by Valentino Braitenberg. His vehicles (Braitenberg, 1984) (Figure 2), although being thought experiments, inspired many roboticists to design robots in a parsimonious way and exploit the environment to create interesting behaviours, among them obstacle avoidance.

In the approaches by Grey Walter and Valentino Brait-

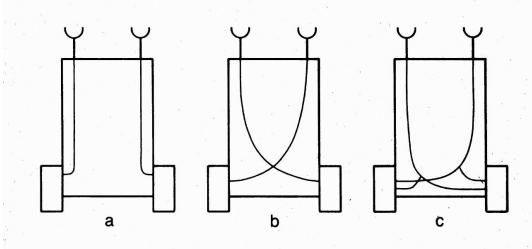

Figure 2: Schematics of different vehicles of type 2 from V. Braitenberg's thought experiments.

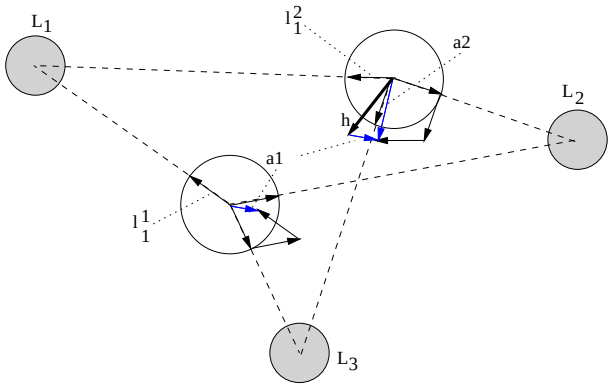

Figure 3: The ALV model applied to two snapshots within a virtual environment.

enberg, obstacle avoidance behaviour is achieved through sensory-motor coupling, involving sensors, actuators, and simple controllers. The more recent robots capable of doing obstacle avoidance we know of are also built along these lines, most of them having this behaviour explicitly encoded. From an engineering viewpoint and for obvious safety reasons to avoid damage to robots and property, this is a reasonable approach.

We will introduce two examples where emergent obstacle avoidance is achieved and study their underlying principles. The first can be seen on a more strategic level, represented by an agent which is able to perform a local visual homing strategy. The second is an agent without sensory-motor coupling with focus on the influence of morphology. We will question the usual assumption that obstacle avoidance is amongst the most basic behaviours of a robot and therefore has to be implemented before any other behaviour. We will also question the usual way of implementing obstacle avoidance using dedicated sensors and a controller.

The paper is divided into the following sections: In section 2, visual homing, an important local navigation strategy used by many animals, is introduced, and the emergence of obstacle avoidance behaviour is discussed within this context (perspective I). Section 3 presents obstacle avoidance in the context of a robot without sensory-motor coupling (perspective II). Section 4 discusses the impact of these two different approaches coming from different directions.

2 Perspective I: Visual Homing and Obstacle Avoidance

In this section, we investigate the emergent obstacle avoidance properties of two visual homing methods using snapshots, the snapshot model (Cartwright and Collett, 1987) and the Average Landmark Vector (ALV) model (Lambrinos et al., 2000). Visual homing is an instance of the group of local navigation strategies (for a review see (Franz and Mallot, 2000)). The underlying principle of snapshot visual homing is the following: One-dimensional omni-directional snapshots along the horizon are taken at two different positions. Those are usually the home position and the current position. The snapshots are aligned towards a common global orientation. The visual homing models allow to infer the vector of displacement for the two positions of the snapshots. Note that a large subset of the landmarks has to be visible from both snapshot positions.

2.1 The ALV Model

We will concentrate on the ALV model since it is more parsimonious than the snapshot model and produces almost the same behavioural results. Another advantage of considering the ALV model is that it does not need to perform a grouping of landmarks between the two snapshots as does the snapshot model. The ALV model can explain some of the navigation behaviour observed in insects. Under certain circumstances, it is identical to the difference vector model, a version of the snapshot model where the contributing vectors are calculated as the difference between unit vectors pointing towards the current view sectors and the snapshot sectors.

The ALV model calculates the homing vector h by subtracting the AL vector at the target position from the AL vector at the current position:

$$h = a_c - a_t,$$

where $a_t = \sum_{i=1}^{n} l_i^t$ and $a_c = \sum_{i=1}^{n} l_i^c$ with l_i^c and l_i^t being the landmark vectors in unit length (compare Figure 3). For simplicity, the AL vectors are expressed as the sum (not the average) of the landmark vectors.

The ALV model has been successfully implemented on a mobile robot built completely in analogue hardware (Möller, 2000) using two capacitors in order to store the AL vector at the home position. It has also been shown by Hafner and Möller (Hafner and Möller, 2001) that using the two snapshots as inputs and the homing vector as desired target to a feed-forward neural network, the ALV model can be learned.

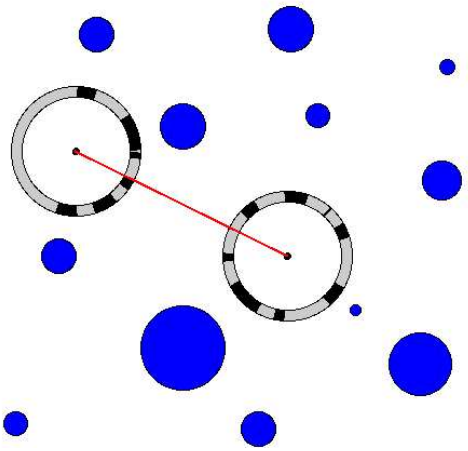

Figure 4: Schematics of an agent within a virtual environment at two different positions. The two-dimensional omni-directional visual field of the agent is represented by a ring, containing the projections of the landmarks.

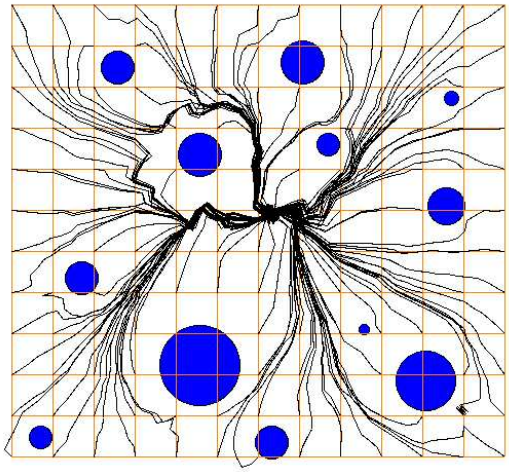

Figure 5: Trajectories from different starting positions on a grid towards a goal position where a snapshot had been taken. The homing strategy used here is the average landmark vector (ALV) model, using the edges of the cylinders as landmarks.

2.2 Image Preprocessing

Both the snapshot model and the ALV model reduce the image to a one-dimensional binary array, where each landmark is represented by one pixel at the position pointing towards the center of the landmark (or alternatively two pixels pointing towards the left and right edge of the landmark). In a natural environment, it is difficult to separate the landmarks from the background. Experiments of visual homing have been performed with normalised, lowpass filtered grayscale images, showing promising results (Hafner and Möller, 2001). However, with this method, the notion of a landmark is different.

2.3 Visual Homing Simulations

In Figure 4, the schematics of an agent within a virtual environment at two different positions are displayed. The two-dimensional omni-directional visual field of the agent is represented by a ring, containing the projections of the landmarks. The robot trajectories during visual homing using the ALV model from different starting positions on a grid towards the goal position can be seen in Figure 5. The home snapshot has been taken at a position near the center of the virtual world. The homing algorithm takes a new snapshot at the current position at each time step, calculates the homing vector (see section 2.1), and moves a small step in this direction.

2.4 Emergent Obstacle Avoidance in the ALV Model

As can be seen in Figure 5, the homing trajectories move around the obstacles. This behaviour is not explicitely encoded in the ALV model. In previous publications on visual homing and the ALV model, the obstacle-avoidance behaviour has been mentioned, however, it has not been analysed in detail: "Note that under certain conditions some of these models exhibit obstacle-avoidance behavior without a dedicated obstacle-avoidance module" (Lambrinos et al., 2000).

An interesting aspect shows up if we consider the learned visual homing model (Hafner and Möller, 2001), which resembles the original ALV model very closely. The neural network is trained with a set S of snapshot pairs $(s_i, s_j)^P$ and vectors v_h^P, which directly point from position i^P to position j^P, regardless of whether there are any obstacles in between the two snapshot positions or not. The resulting learned model, however, will most often avoid these obstacles.

The reason for this strange behaviour can be explained by some geometric properties of the environment. In Figure 6, a scenario for an agent homing in an environment with three landmarks is plotted from a bird's perspective at two different time steps. The home position is positioned somewhere between the landmarks, the current position of the agent is below the landmarks. The region not visible from the current position of the agent (occluded by a landmark) is shaded in grey. In Figure 6 top, the agent is moving straight towards the goal, however, the direct way is blocked by a landmark. As soon as this landmark is occluding another landmark (Figure 6 bottom), the landmark vector l_i^c which was formerly pointing from the current position of the agent in direction of the occluded landmark L_i, is now subtracted from the home vector without occlusions. For $L_m, \ldots, L_n$ being the occluded landmarks at the current position with $m < n$, we get a new

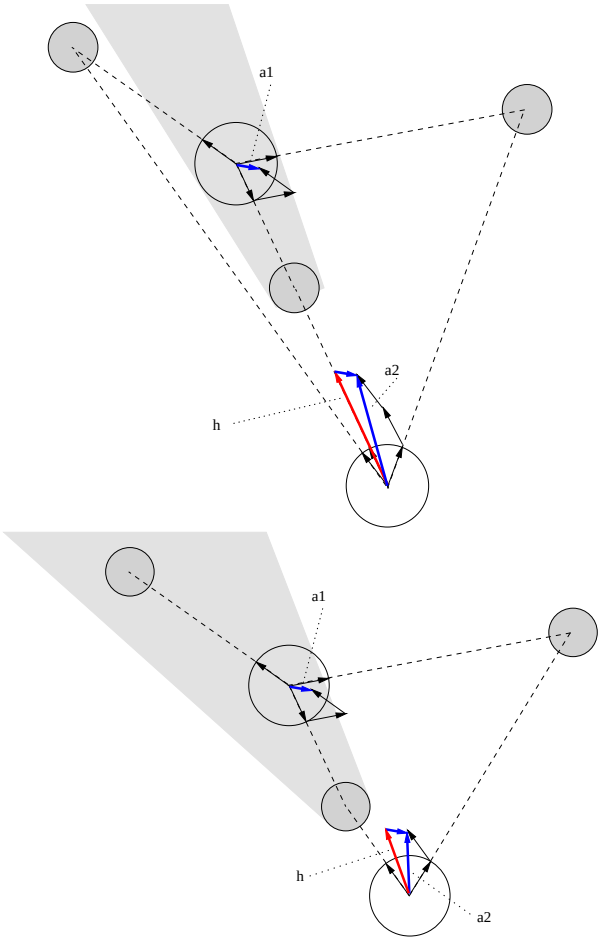

Figure 6: Top: Homing vector for none occluded landmarks. Bottom: Homing vector for one occluded landmark. The agent is repulsed from the landmark which causes the occlusion.

homing vector h':

$$h' = a'_c - a_t \quad (1)$$

$$= \sum_{i=1}^{m} l_i^c - \sum_{i=1}^{n} l_i^t \quad (2)$$

$$= h - \sum_{i=m+1}^{n} l_i^c \quad (3)$$

On average, the sum of the vectors pointing from the agent towards the occluded landmarks is pointing straight in the center of the occluding landmark (see Figure 7). As a consequence, the agent is repulsed from the landmark in front of it, allowing for any sidewards movement to take over, resulting in trajectories as can be found in Figure 5.

We have shown an example of a situation, where a navigation strategy, which is usually considered to be on a higher level than basic obstacle avoidance behaviour, results in exactly this behaviour without having it explicitly encoded.

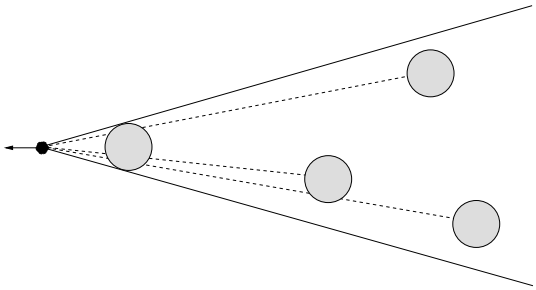

Figure 7: Repulsive vector component in the ALV model resulting from occlusions.

3 Perspective II: Obstacle Avoidance without Sensory-Motor Coupling

In this part a different route is taken. Instead of using dedicated sensors to detect obstacles and a controller to move the robot away from it, the embodiment of the robot itself should act as information gathering and processing device as well as actuator. The desired behaviour should be directly generated by mechanical robot-environment interaction, exploiting the physics of both environment and robot embodiment. Thus the robot will not avoid collisions but by exploiting the interaction with the obstacle turns away, i.e. avoids it. In this way, the need for sensors, actuators and control could be reduced to a minimum.

A similar approach was the design of a suitable morphology in the context of an artificial compound eye (Lichtensteiger, 2000) for obstacle avoidance. The morphology has been used to simplify control but neither as information gathering device nor as actuator.

3.1 Prerequisites

As a starting point for our investigations we were using a holonomic platform (see next section). A holonomic robot moving on the floor, as in our case, has three independently actuated degrees of freedom that allow for movements in arbitrary directions.

In order to use morphology as a sensing device, the robot has to physically interact with the environment and record this interaction by being changed; that is, we need a flexible robot morphology. To use this flexibility as a control device, the same change has to go along with a change in actuation, which will lead to a change in robot behaviour. Change in morphology quite naturally results in change in spatial actuator configuration, which can be exploited to achieve what we call meccano-motor coupling, which in turn can be used to achieve the desired behaviour, as in our case obstacle avoidance.

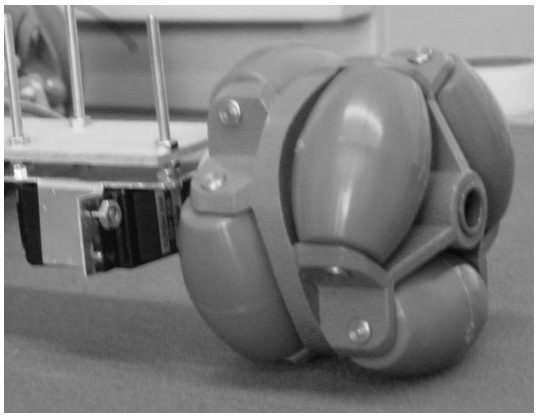

Figure 8: Stanford omnidirectional wheels used in the present design. Two rows of free-rotating rolls are mounted on the actual wheel to allow arbitrary movements.

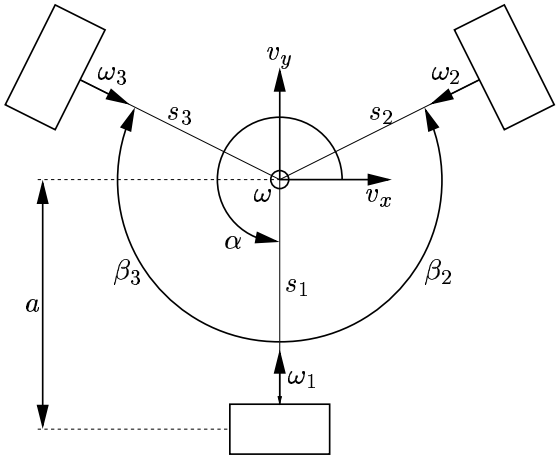

Figure 9: Schematic of a PPRK-type holonomic robot. Refer to text for further information.

3.2 Holonomy

A robot is holonomic if its movement is not restricted by any non-integrable kinematic constraint. For a robot moving in two dimensions (i.e. on a flat surface), holonomy is the ability to rotate freely and to move in any direction independently of orientation. A differential-wheel driven robot is a typical example of a non-holonomic robot, as it cannot move sideways; ordinary wheels do not allow for a movement parallel to the turning axis.

A holonomic driving system involves wheels (or legs) that allow for movements along all three degrees of freedom independently, e.g. driven castor wheels, ball wheels, or so-called omni-directional wheels, which were used in the present design (see Figure 8). Using these wheels, a holonomic robot platform is easily built, as for example the PPRK (Palm Pilot Robot Kit (Reshko et al., 2000)).

The omni-directional wheels can freely roll sideways due to the rollers, but can only control the speed in the direction where the wheel is pointing. The inverse kinematics are easily found (see (Reshko et al., 2000) or (Saner, 2002)),

$$\begin{bmatrix}\omega_1\\ \omega_2\\ \omega_3\end{bmatrix} = \underbrace{-\frac{1}{r}\begin{bmatrix}1 & 0 & a\\ \cos\beta_2 & \sin\beta_2 & a\\ \cos\beta_3 & \sin\beta_3 & a\end{bmatrix}}_{=:\mathbf{A}}\begin{bmatrix}v_x\\ v_y\\ \omega\end{bmatrix} \tag{4}$$

which determines the wheel velocities ω_i for given linear $\mathbf{v} := (v_x, v_y)^T$ and angular speed ω of the robot. As the matrix $\mathbf{A}$ is invertible (for appropriate angles β_2 and β_3, respectively[1]) we can also calculate the robot movement, given the wheel speeds. r denotes the diameter of the wheels[2], for the other notation check Figure 9.

3.3 Morphology

Quite a few holonomic robots have been built, most of them using 4-wheel driving-systems. Nevertheless, the base for our further investigations was the 3-wheeled PPRK.

The general configuration of the PPRK was kept while making the morphology flexible. The most simple way is a star-like robot with three movable segments s_1, s_2, s_3 to which the wheels are attached (see Figure 9). This platform has five degrees of freedom: position $\mathbf{x} = (x, y)$ of center of mass, orientation α of the robot (i.e., of the reference segment s_1) plus relative orientations of the two remaining segments (β_2 and β_3).

If all these five degrees of freedom were actuated, we would have a redundant system, i.e. for given linear $\mathbf{v}$ and angular speed ω of the robot there are several solutions for wheel speeds ω_i and segment angles β_j. In fact, redundancy in this case allows the robot to move holonomically while concurrently performing arbitrary movements of the segments.

In what follows, we describe how this redundancy can be exploited to construct a meccano-motor coupling.

3.4 Dynamics

If we fix the speed of two wheels in a way that the robot performs a forward movement,

$$\omega_3 = -\omega_2 = \text{const.} \tag{5}$$

we still have three degrees of freedom. By changing the segment angles and the speed of the third wheel appropriately, it is possible to control the robot. The interesting difference to the PPRK is that the robot movement is not controlled by wheel speeds alone, but also by morphology and morphological change.

To investigate the relationship between robot morphology and movement, a series of simulations was performed. Two wheels were set to constant speeds (see equation (5)), the

[1] For the PPRK we have $\beta_2 = \beta_3 = \frac{2\pi}{3}$

[2] For the PPRK we have: $r = 0.02\,m$, $a = 0.14\,m$

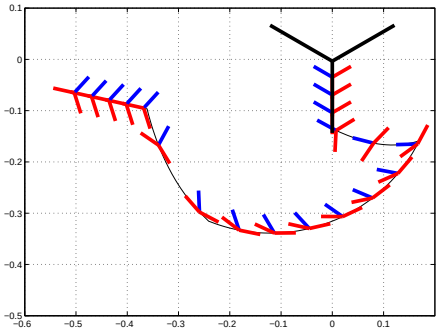

A

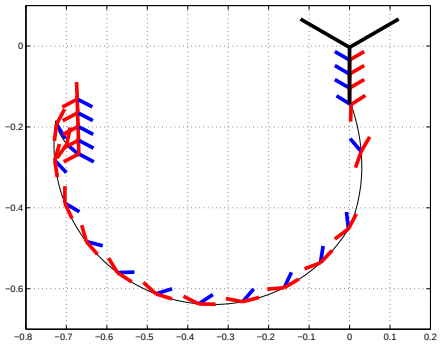

B

Figure 10: Resulting trajectories rotating the front segment (panel A) and the left segment (panel B) by $\frac{\pi}{4}$.

third wheel's speed was set to zero

$$\omega_1 = 0\,rad/s \tag{6}$$

while segment angles β_1, β_2 were varied (see Figure 10). In all simulations the robot started near the top right corner moving forward (i.e. in the direction of the reference segment s_1) for two seconds. The segment angles were changed in several ways during one second leading to a new form of the robot that was maintained for five seconds. During these five seconds the robot performed a turn to the right. Then, during another second, the original form of the robot was restored, by moving the segments back to their original angles. For the remaining two seconds the robot moved straight forward.

Depicted are the simulations involving rotation of the front segment s_1 (see Figure 10A) and of the left segment s_2 (see Figure 10B). In both cases the robot performed a right turn. Rotating the front and the left segment simultaneously resulted in a trajectory that was a combination of the results obtained when rotating only one of the segments (data not shown). Although the turn radius was of similar magnitude for all three simulations, transients (the time when the segment angles are changing) were different. It was also possible to generate left turns by rotating the right segment s_3 instead of the left one.

The conclusion is that the segment angles significantly influence the behaviour of the robot. Furthermore, it is actually feasible to control the robot by changing its morphology appropriately. These simulations illustrating basic characteristics of the morphology-movement coupling were the starting point for the design of the obstacle avoiding robot.

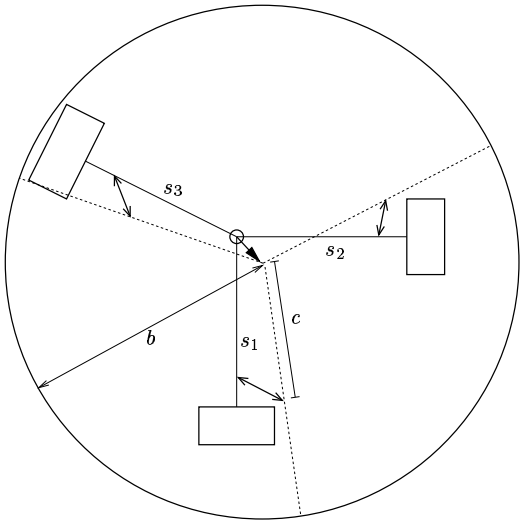

Figure 11: Interaction of the robot with the environment is mediated by a circular bumper (large circle), connected elastically to the three segments (illustrated by double-arrow lines). Bumping into an obstacle results in a displacement of the bumper relative to its equilibrium position (illustrated by single-arrow line). This leads to a rotation of the segments, i.e. a morphological change, that influences the robot trajectory. Equilibrium position of the segments (relative to the bumper) are depicted by dashed lines.

3.5 Obstacle Avoidance

So far, change in morphology was artificially generated. To obtain obstacle avoidance behaviour, this same change needs to result from the interaction of the robot with the environment, i.e. with the obstacles. Several possibilities were investigated, among them direct interaction of the segments with the environment, which did not yield satisfying results.

Best results have been obtained using a rigid circular bumper mounted on top of the robot's segments (see Figure 11). The bumper was connected by three elastic springs to the segments of the robot. This was done such that the equilibrium position of the segments corresponded to the PPRK configuration[3]. The bumper's function is to center the segments when the robot is not touching an obstacle. Furthermore, as we will see later, the bumper avoids direct contact of the segments with the obstacles, in our case walls.

Again, the wheel speeds were set in a way, that the robot performed a forward movement

$$\omega_1 = 0\,rad/s, \quad \omega_2 = -3\,rad/s, \quad \omega_3 = 3\,rad/s \tag{7}$$

[3] $\beta_2 = \beta_3 = \frac{2\pi}{3}$

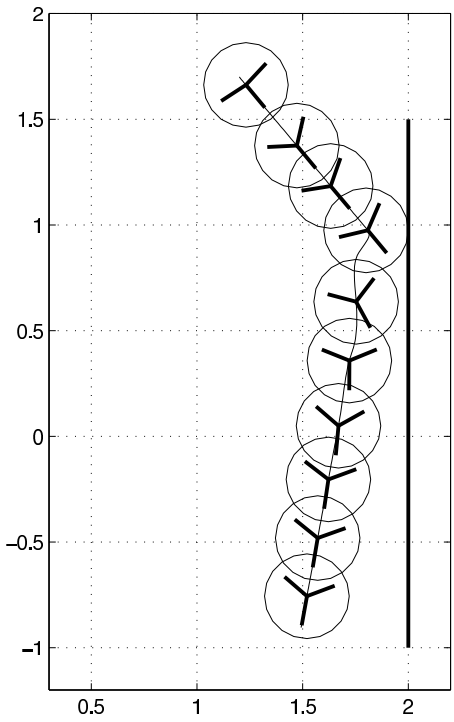

Figure 12: Obstacle avoidance. The robot approaches a wall under an angle of $40°$, avoiding it by touching it (the robot is depicted every $4\,s$, the line illustrates the trajectory of the robot's center of mass).

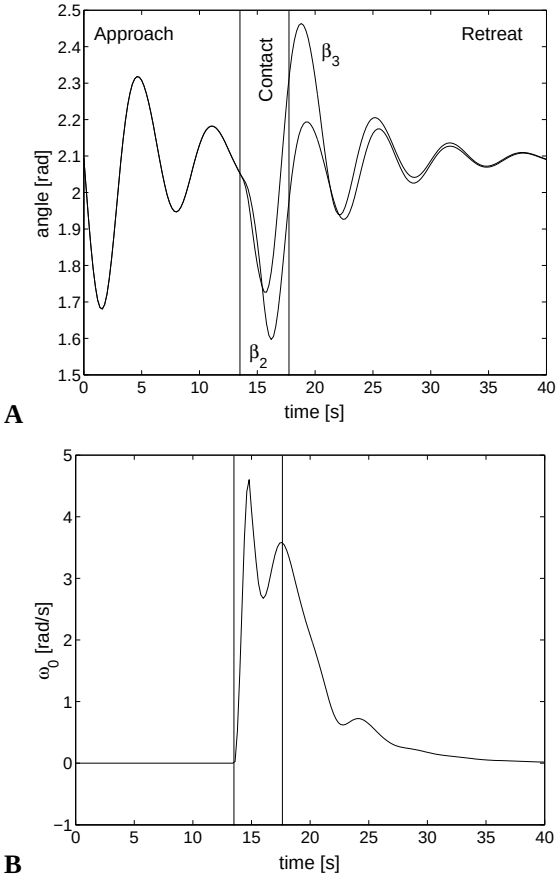

Figure 13: Segment angles β_1 and β_2 (panel A) and speed ω_0 of the front wheel (panel B). Before contact with the wall ($t \approx 14\,s$) segment angles are oscillating symmetrically ($\beta_1 = \beta_2$) as the robot accelerates. Front wheel speed is zero. On contact, the segments $s_{1,2}$ rotate forward, therefore $\beta_{1,2}$ decrease. At the same time the front wheel starts to spin as the robot rotates. After wall contact ($t \approx 18$) segment angles and front wheel speed approach their default values.

in the direction of the segment s_1[4]. In contrast to the simulations described above, where the wheel speeds were forced to constant values, we assumed a weak motor control (i.e. small gain values), leading to a "soft" or compliant behaviour of the wheels.

In this setup simulations of the robot approaching a wall and interacting with it were performed. The simulations can be divided into three phases: the approach phase, where the robot accelerates and moves towards the wall; the contact phase, where the bumper touches the wall; and the retreat phase, where the robot moves away from the wall.

In the **approach phase** the robot starts from zero velocity. The acceleration of the wheels is accompanied by a rowing movement of the segments s_2 and s_3 indicated by (synchronous) oscillation of the segment angles β_2 and β_3 (see Figure 13). The oscillations are generated by the elastic connections of the segments to the bumper and the deviation from their respective equilibrium positions by the force generated by the accelerating wheels. The oscillations are damped and will die out eventually. The rowing movement of the segments is also evident from Figure 12.

On **contact** ($t \approx 14\,s$), the bumper's movement is restricted by the wall. The segments s_2 and s_3 bend forwards driven by the wheels. The resulting forces between the obstacle and the bumper are fed back by the springs onto the segments and induce a change in the segment angles (see Figures 12 and 13A). As the wheels continue to drive, the segments get pulled away further from their equilibrium positions. This leads to increasing forces on the segments. The strong force component perpendicular to the front segment evokes a rotation of the front wheel (see Figure 13B), and therefore the front segment moves away from the wall. The robot temporarily loses wall contact ($t \approx 15\,s$, the front wheel speed decreases) but immediately bumps into the wall a second time ($t \approx 16\,s$). The front wheel speed increases again, and the robot turns away from the wall and finally looses contact ($t \approx 18\,s$).

Because of the weak motor control, the front wheel continues its rotation for a few seconds in the **retreat phase**, resulting in a further rotational movement of the robot. The segments s_2 and s_3 express a similar rowing movement as in

[4] Additional parameters: $a = 0.14\,m$, $b = 0.2\,m$, $c = 0.1\,m$ (see Figure 11)

the starting phase.

Although the details of the dynamics during the contact phase are rather complex, the following basic features seem to account for the obstacle avoidance behaviour. The first is the redundancy in the (holonomic) driving system, which allows for a high flexibility of movement.

The second is the elastic connections between the bumper and the segments. As we have seen, this plasticity allows for morphological change which influences the robot behaviour. Here it is important to note that the change in morphology is associated with an increased energy of the system (stored in the springs). This energy surplus is not only used, in combination with the morphological change, to turn the robot, but actually to push it away from the wall. Without this elastic energy the turning of the robot would stop as soon as it is loosing contact with the wall. Consequently the robot would slide along the wall, never actually turning away from it.

Besides the flexible morphology another form of compliancy is needed, namely the weak motor control, which is most important for the front wheel. This allows the rotation of the front wheel in the first place thus enhancing the rotation of the robot. As the front wheel keeps its speed for some time, even after the robot lost contact to the wall, the compliancy in the wheels also contributes to the ability of the robot to turn away from the wall.

The obstacle avoidance by meccano-motor-coupling therefore exploits the redundancy of the driving system and the flexibility of the morphology as well as the compliancy in the motors.

4 Conclusions

There are two common unquestioned assumptions about obstacle avoidance: The first is, that obstacle avoidance is considered a very basic behaviour which has to be implemented as a low-level behaviour before any other behaviour is considered. The second assumption is, that distance or collision sensors are a necessary prerequisite for obstacle avoidance.

We have shown that these assumptions have to be reconsidered. Obstacle avoidance can be an 'emergent' behaviour resulting from both an abstract, strategic behaviour like visual homing, or purely from properties of the morphology of the agent and its interaction with the environment. Furthermore, there is no need for dedicated sensors. Obstacle avoidance can be achieved using very abstract sensors like visual snapshot memories or the whole body of the agent itself.

There are other interesting examples of emergent obstacle avoidance, like for example light following robots. If there is an obstacle between the robot and the light source, this obstacle generates a shadow in front of it. The robot avoids the shadow since it is seeking the light and therefore does not collide with the obstacle (Reimann and Mansour, 2000). The emergence of obstacle avoidance is task-specific, as in the example of visual homing. The design of the meccano-motor coupling is also task specific. It would be interesting to see how the two perspectives generalise to arbitrary tasks, and whether they can be put in more general, theoretical frameworks.

Acknowledgements

Many thanks go to Christian Saner, who did excellent work in implementing the simulation of the holonomic robot described in this paper. This research has been supported in part by grant #20-61372.00 of the Swiss National Science Foundation (VVH).

References

Braitenberg, V. (1984). *Vehicles. Experiments in Synthetic Psychology*. MIT Press, Cambridge.

Cartwright, B. A. and Collett, T. S. (1987). Landmark Maps for Honeybees. *Biological Cybernetics*, 57:85–93.

Franz, M. O. and Mallot, H. A. (2000). Biomimetic robot navigation. *Robotics and Autonomous Systems*, 30:133–153.

Hafner, V. V. and Möller, R. (2001). Learning of visual navigation strategies. In Quoy, M., Gaussier, P., and Wyatt, J., (Eds.), *Proceedings of the European Workshop on Learning Robots (EWLR-9)*, pages 47–56, Prague.

Holland, O. (1996). The Grey Walter Online Archive. http://idle.uwe.ac.uk/IAS/gwonline.html, University of the West of England, Bristol.

Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., and Wehner, R. (2000). A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems, special issue on Biomimetic Robots*, 30:39–64.

Lichtensteiger, L. (2000). Towards optimal sensor morphology for specific tasks: Evolution of an artificial compound eye for estimating time to contact. In McKee, G. T. and Schenker, P. S., (Eds.), *Sensor Fusion and Decentralized Control in Robotic Systems III*, volume 4196 of *Proceedings of SPIE*, pages 138–146.

Möller, R. (2000). Insect visual homing strategies in a robot with analog processing. *Biological Cybernetics, special issue: Navigation in Biological and Artificial Systems*, 83:231–243.

Reimann, S. and Mansour, A. (2000). Orientation by weighted randomness. *Journal of Artificial Life and Robotics*, 4:119–123.

Reshko, G., Mason, M. T., and Nourbaksh, I. R. (2000). Rapid prototyping of small robots. Technical report, Carnegie Mellon University.

Saner, C. (2002). Design eines holonomen Roboters mit veränderbarer Morphologie. Master's thesis, University of Zurich, AILab.