

Agent-Environment Interaction in Visual Homing

Verena V. Hafner*

Artificial Intelligence Lab,
Department of Information Technology,
University of Zurich, Switzerland,
vhafner@ifi.unizh.ch

Abstract. This study illustrates how obstacle avoidance can emerge from a visual homing strategy, caused by the intrinsic geometric structure of the environment. An example is shown where an agent performs visual homing in a virtual environment with several obstacles which also serve as visual landmarks. The agent has omnidirectional vision similar to many prey animals. The applied visual homing strategy is the Average Landmark Vector (ALV) model by Lambrinos et al.[1]. When observing the homing trajectories of the agent, it can be seen that it performs obstacle avoidance without having this behaviour explicitly encoded. It will be shown that the dynamic feedback the agent gets from its environment is crucial for this kind of behaviour.

1 Introduction

Obstacle avoidance in animals evolved soon after the first living creatures started to move. Even bacteria such as *E.coli* perform attractant or repellent behaviour. By doing so, they can move towards certain objects (or chemicals in this case), avoid others, and consequently perform obstacle avoidance. In artificial agents, obstacle avoidance is considered as one of the most basic behaviours, and was already present among the earliest autonomous mobile robots, such as the turtles ‘Elsie’ and ‘Elmer’ built by Grey Walter [2] (see Figure 1). For moving, they had two propulsion wheels and one steering wheel; as sensors, one light and one touch sensor were used. The robots’ control system was completely analogue, and they could perform two main actions: obstacle avoidance, by retreating on contact, and light following. These simple behaviours in combination with an interaction with the environment led to other kinds of behaviour which are perceived as very complex and even intentional by outside observers. For example, by placing a lamp on each turtle’s shell, a kind of ‘social behaviour’ emerged from the interaction of the two turtles.

Since the pioneering work of Grey Walter, increasing numbers of robots have been built in order to understand adaptive behaviour in the real world. The new interest in *embodied artificial intelligence (EAI)* builds on this early work, and recognised the importance of interaction with the real world. In this chapter, I will present an example of agent behaviour that profits from the interaction with the environment in a way that has

* Now at Sony CSL Paris, France, *hafner@csl.sony.fr*




Fig. 1. Second generation turtle designed by Grey Walter. Its main two behaviours are obstacle avoidance and light seeking (picture taken from the Grey Walter Online Archive [3]).

not been anticipated by the designer. The example shows an agent navigating back to its home location by means of visual homing. What makes this experiment particularly interesting is the fact that this agent gets obstacle avoidance “for free” through the interaction with the environment and intrinsic geometric properties.

2 An Example: Navigation and Emergent Obstacle Avoidance

In this section, I introduce the emergence of obstacle avoidance induced by a behaviour which is usually considered as being on a higher level than pure obstacle avoidance: navigation¹. I investigate the emergent obstacle avoidance properties of a visual homing method using snapshots called the Average Landmark Vector (ALV) model [1]. The underlying principle of snapshot visual homing is the following: Omni-directional one-dimensional snapshots along the horizon are taken at two different positions. Those are usually the home position and the current position of the agent. The snapshots are aligned toward a common global orientation. The visual homing model allows to infer the vector of displacement for the two positions of the snapshots. Note that this homing strategy is local and therefore a large subset of the landmarks has to be visible from both snapshot positions.

The ALV model can explain certain aspects of the navigation behaviour observed in insects. It calculates the homing vector h by subtracting the AL vector at the target position from the AL vector at the current position:

$$h = a_c - a_t,$$

¹ This experiment has been described first in the context of emergence together with the example of a holonomic robot in [4].

where $a_t = \sum_{i=1}^n l_i^t$ and $a_c = \sum_{i=1}^n l_i^c$ are the sums of the landmark vectors at the target and the current position respectively. The landmark vectors l_i^c and l_i^t have unit length since distance information is not used. For simplicity, the AL vectors are expressed as the sum (not the average) of the landmark vectors. The ALV model reduces the image to a one-dimensional binary array, where each landmark is represented by one pixel at the position pointing toward the centre of the landmark (or alternatively two pixels pointing toward the left and right edge of the landmark). The ALV model has been successfully implemented on a mobile robot built completely in analogue hardware [5] using two capacitors in order to store the AL vector at the home position. In a natural environment, it is difficult to separate the landmarks from the background. An alternative version of the ALV model with continuous image processing has been introduced [6]. It works on normalised, low-pass-filtered greyscale images, where a vector pointing towards the centre of mass is used rather than an AL vector.

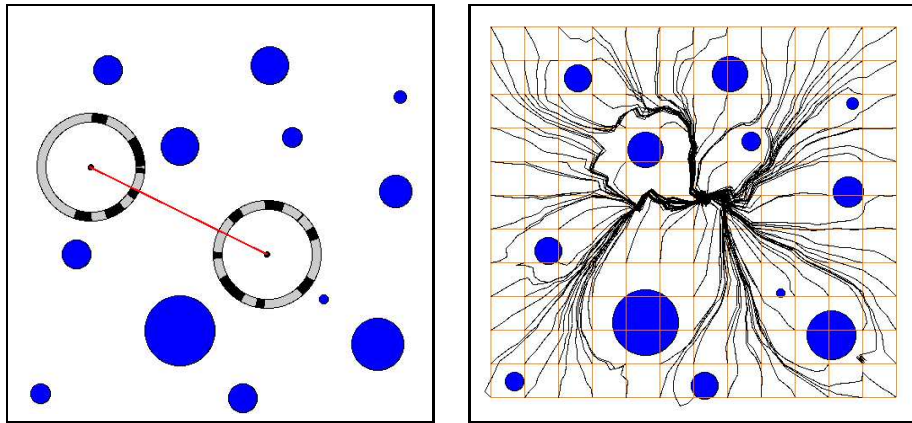


Fig. 2. Left: Schematics of an agent within a virtual environment at two different positions. The omni-directional one-dimensional visual field of the agent is represented by a ring, containing the projections of the landmarks. These rings do not represent the size of the agent. Right: Trajectories from different starting positions on a grid toward a goal position at the centre 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.

In Figure 2 left, the schematics of an agent within a virtual environment with sparsely distributed cylindrical obstacles at two different positions are displayed. The omni-directional one-dimensional 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 toward the goal position can be seen in Figure 2 right. The home snapshot has been taken at a position near the centre of the virtual world. The homing algorithm takes a new snapshot at the current position at each time step, calculates the homing vector, and moves a small step in this direction.

As can be seen in Figure 2 right, the homing trajectories move around the obstacles. This behaviour is not explicitly encoded in the ALV model.

It has been shown by Hafner[6] and Hafner and Möller[7] 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 in a self-supervised manner. An interesting aspect shows up if we consider the learned visual homing model, 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

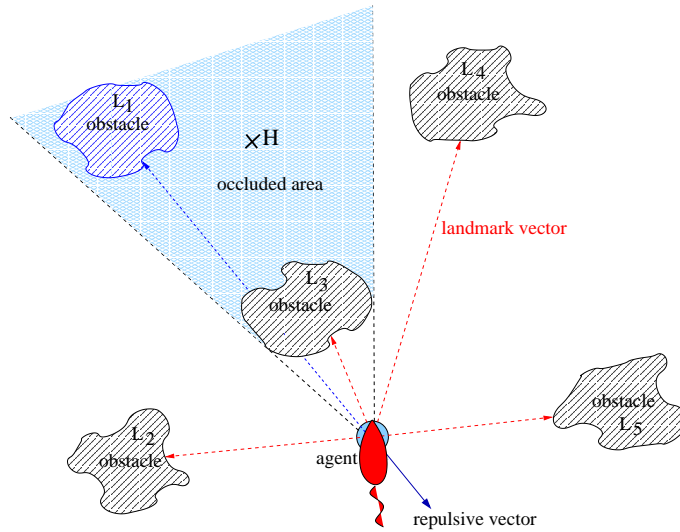


Fig. 3. Schematics of an agent homing in a complex environment with several landmarks which can at the same time be obstacles. The nearest obstacle is occluding other objects therefore producing a repulsive vector away from the obstacle. \mathbf{H} indicates the home position.

reason for this strange behaviour can be explained by some geometric properties of the environment.

In Figure 3, a scenario for an agent homing in an environment with several landmarks is plotted from a bird's eye view. The agent is moving straight towards the goal, however, the direct way is blocked by an obstacle. As soon as this obstacle is occluding another landmark, 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, \dots, L_n being the occluded landmarks at the current position with $m < n$, we get a new homing vector h' from the difference of

the two average landmark vectors a'_c and a_t :

$$h' = a'_c - a_t = \sum_{i=1}^m l_i^c - \sum_{i=1}^n l_i^t = h - \sum_{i=m+1}^n l_i^c$$

On average, the sum of the vectors pointing from the agent toward the occluded landmarks is pointing straight in the centre of the occluding landmark. As a consequence, the agent is repulsed from the landmark in front of it, allowing for any sideways movement to take over, resulting in trajectories as can be found in Figure 2 right. This repulsion provides the same results as described in [8], where a force field method results in steering away from a surface. The problem of a dead spot, where the agent is repulsed exactly in the direction negative to its movement vector is avoided by noise in real robots. The agent is also automatically more strongly repulsed from close obstacles than from others further away, since the close ones are occluding a higher number of other landmarks.

3 Discussion

I have shown a situation, where a navigation strategy, which is usually considered to be on a higher level of complexity than basic obstacle avoidance behaviour, results in exactly this behaviour without having it explicitly encoded. We call this behaviour *emergent*, since it results from the agent-environment interaction, is not pre-programmed, and cannot be separated into a sub-module independent from the homing behaviour. This emergence of a simple behaviour (obstacle avoidance) from a more complex behaviour (visual homing) is unusual. In nature, evolution clearly does not come up with more complex behaviour before the most basic skills crucial for survival have been developed. In robotics, obstacle avoidance is usually one of the first behaviours to be implemented on a mobile robot in order to avoid damage to the property, the robot and others.

The effect of emergent obstacle avoidance during visual homing can be observed both in the real world using a mobile robot and in simulation, as long as the dynamics of vision and egomotion are simulated correctly. In both simulation and the real world, the dynamical interaction between agent and environment are important. A homing vector that is derived at only one position either points to a position away from the goal, or points directly towards the goal without taking any obstacles in between into account. Only through the dynamics arising from a constantly updated visual input caused by the movements of the agent within its environment can trajectories be produced which avoid obstacles and lead to the goal. The robot does not influence its environment substantially, but it influences its own sensory input through its movements in space.

What makes this example particularly interesting for *embodied AI* is its focus on the interaction with the environment. A navigating agent with sensors, motors and a brain² *has* to be embodied. It also shows that this concept cannot be restricted to mobile

² Note that having a brain is not necessarily a requirement for *EAI*, but for navigation. Animals evolved brains because they had to move.

robots, but should also apply to simulated agents: If we call an agent that exploits the interaction with its environment *embodied*, both mobile robots and simulated agents that fulfil this requirement should be called *embodied*.

The importance of real world interaction for the study of intelligence and intelligent behaviour in humans and other animals has been recognised a long time ago. This interpretation applied to the study or the design of artificial agents has recently been termed *embodied AI*. The aspects of the real world environment can vary. If we restrict *EAI* to aspects of the environment as we (humans) perceive them, we exclude many of the non-humanoid artificial agents. If we do not restrict it at all, then the concept of *EAI* becomes extremely vague, including even abstract algorithmic software such as a sorting algorithm which interacts with the memory stack. One problem with the *EAI* approach lies in the vagueness of the central concept of *embodiment*. This vagueness has hindered precise communication between researchers in the field so far.

As is clear from the example in section 2, there is no such thing as a dis-embodied navigating agent, and going one step further, there is no dis-embodied agent at all. Every agent, even a pure software agent, is defined by its interactions with the environment. One has to be careful not to restrict the definition of environment in this context too much. In conclusion, any interaction between an agent and an environment in their most open definition can be interesting for the study of intelligence.

References

1. Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., Wehner, R.: A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, special issue on Biomimetic Robots **30** (2000) 39–64
2. Walter, W.G.: An imitation of life. *Scientific American* **182** (1950) 42–45
3. Holland, O.: The Grey Walter Online Archive.
http://www.ias.uwe.ac.uk/extra_pages/gwonline.html, University of the West of England, Bristol (1996)
4. Hafner, V.V., Kunz, H., Pfeifer, R.: An Investigation into Obstacle Avoidance as an ‘Emergent’ Behaviour from two Different Perspectives. *Proceedings of the EPSRC/BBSRC International Workshop on Biologically-Inspired Robotics: The Legacy of W. Grey Walter* (2002) 166–173
5. Möller, R.: Insect visual homing strategies in a robot with analog processing. *Biological Cybernetics*, special issue: Navigation in Biological and Artificial Systems **83** (2000) 231–243
6. Hafner, V.V.: Adaptive Homing - Robotic Exploration Tours. *Adaptive Behavior* **9** (2001) 131–141
7. Hafner, V.V., Möller, R.: Learning of visual navigation strategies. In: *Proceedings of the European Workshop on Learning Robots (EWLR-9)*, Prague (2001) 47–56
8. Reynolds, C.W.: Not bumping into things. Notes for the SIGGRAPH 88 course *Developments in Physically-Based Modeling* (1988)