Visual landmark learning

Giovanni Bianco
Computer Science Serv.
University of Verona
Via S. Francesco 8
I-37129 Verona - Italy
bianco@chiostro.univr.it

Alexander Zelinsky
Dept. of Syst. Eng. - RSISE
Australian National University
Canberra, ACT 0200 - Australia
Alex.Zelinsky@anu.edu.au

Miriam Lehrer
Dept. of Neurobiology
University of Zurich
Winterthurerstr. 190
CH-8057 Zurich - Switzerland
Miriam@zool.unizh.ch

Abstract

Biology often offers valuable example of systems both for learning and for controlling motion. Work in robotics has often been inspired by these findings in diverse ways. Though, the fundamental aspects that involve visual landmark learning and motion control mechanisms have almost exclusively been approached heuristically rather than examining the underlying principles. In this paper we introduce theoretical tools that might explain how the visual learning works and why the motion is attracted by the pre-learnt goal position. Basically, the theoretical tools emerge from the navigation vector field produced by the visual behaviors. Both the learning process and the navigation scheme influence the motion field. We apply classical mathematical and dynamic control to analyze the efficiency of our method.

1 Introduction

Animals, including insects, are proficient in navigating and, in general, several biological schemas for solving navigational tasks seem to be promising for robotics applications. Basically biological systems involve two fundamental mechanisms: visual learning and visual motion control. There are, of course, diverse ways of implementing such mechanisms among animals. Nevertheless, insects provide basic methods that are extremely valuable for robotics. To this extent, the use of landmarks for piloting and for goal identification is fundamental to orientation, although other mechanisms, such as the optomotor response, path integration and skylight navigation are clearly involved [19]. However, whereas the optomotor response and the path integration mechanism require no more than an appropriately hard-wired neural network, the use of the skylight compass and of landmarks requires, in addition, a learning process that must take place in every new situation.

On the other hand, referring to motion planning, it is widely agreed that landmark-guided navigation in insects is based on storing the image of the landmark(s) as some kind of a *snapshot* or a *template* [1, 4]. On its next journey, the insect navigator strives to accomplish a match between the currently viewed image of the landmark(s) or of the goal and the snapshot stored in memory. Many examples of the above mentioned abilities transposed to real robots can be found in the review [18]. From these case studies the implementation of biological mechanisms appear to follow a pragmatic approach where the final behavior is the main objective. The learning and control aspects are of utmost importance but a thorough study on the theoretical principles has only recently been performed [13, 3]

Basically, the core of the theory is represented by the navigation vector field, whose study provide two main results:

- the *visual potential function* generating the field represents the driving principle to perform visual guidance. When proven to be a *Lyapunov* compliant function, we can state the navigation system exhibits convergence to the goal.
- The *conservativeness* of the navigation vector field deals with the concept of repeatability of the trials and provides key information to perform landmark learning.

In this paper we address, in particular, principles involved in landmark learning. Details about the navigation system and the nature of the potential function can be found in references [2] and [3] respectively.

The organization of the paper is as follows. In Section 2 aspects both related to findings on biological learning and to biological navigation will be introduced. In Section 3 the theoretical principles specifically involved with visual learning and visual navigation are detailed. Final remarks conclude the paper.

2 Biological findings

Over many decades, studies of the visual performance of bees have exploited the fact that bees keep returning to a profitable feeding site once found, even when it is an artificial food source established by an experimenter.

2.1 Landmark learning

As soon as the bee encounters a novel place, she turns by 180 degrees to inspect the place and performs the initial phase of training, termed the *Turn-Back-and-Look (TBL)* phase [10]. A similar behavior was also observed in other insects thus categorizing this phase a typical behavior of an insect when a new visual learning phase is needed [20, 21].

In references [10, 11] and [14] the details and results on the visual parameters learned by TBL are introduced. Basically, findings show that TBL performed on departure serves primarily for acquiring depth information by exploiting image motion, whereas color, shape and size of landmarks are mainly acquired on arrival.

Attempts to understand in detail the geometric significance of learning flights have only recently been made. Essentially, the flights are invariant in certain dynamic and geometric structures thus allowing the insects to artificially produce visual cues in specific areas of the eyes [22]. Perhaps, the main reason is that the precision for the homing mostly depends upon the proximity of chosen landmarks to the goal [5]. In fact, those flights need to be repeated whenever some changes in the goal position occur [12].

2.2 Landmark guidance

Landmarks guidance in insects is retinoptically driven and animals tend to reduce the discrepancies between the stored view and the actual one by a matching procedure (reviews in [6] and [19]). The survey work presented in [18] addresses biological navigating behaviors from a robotics point of view.

Referring to landmark guidance in bees, the seminal work is presented in [4]. The authors show how bees learn landmarks by storing an unprocessed two dimensional snapshot of the panorama. The model matches landmarks in the stored snapshot with landmarks in the actual image. If this match is performed far from the goal every matched pair could differ both in angular size and compass bearing. These differences drive a bee toward the right position. Template matching has also recently been noticed in ants [9].

The model introduced in [4] has some shortcomings and interesting extensions have been addressed in recent works [7, 3] by introducing omnidirectional vision and color images.

3 The motion field

According to what has been previously expressed, starting with local visual information, a vector needs to be computed by the agent which will be used it to perform the next movement. In our case, the computation of the navigation vector is based on information involving the chosen landmarks. How to get navigation information from landmarks is briefly introduced here for completeness and details can be found in [2, 3].

Basically, once landmarks have been learned, they can provide two kind of information to perform motion:

- their actual size, compared to the size learned at the goal site, reports how far/close the agent is to the goal position
- their actual orientation, compared to the orientation learned at the goal site, speaks about the actual left/right shift of the agent.

This kind of data come from each individual landmark and we need to fuse them in order to get the overall navigation vector. Intriguingly, the fusion procedure has strong biological bases as detailed in [18].

To formalize aspects related to the motion field generated in the environment, we call $\mathbf{p}$ the vector representing the robot's Cartesian position $[x\ y]$ in a world reference $\mathbf{W}$. We also define step k the discrete time k of *robot dynamic state*.

Let $\vec{V}(\mathbf{p}(k)) = [V_x(\mathbf{p}(k))\ \ V_y(\mathbf{p}(k))]$ be the output of the motion strategy at a given step k, i.e. the robot movement at step k. If the robot operates in *position mode*, i.e. at each step it updates its Cartesian position, then

$$\mathbf{p}(k+1) = \mathbf{p}(k) + \vec{V}(\mathbf{p}(k)) \tag{1}$$

where $\mathbf{p}(k)$ represents the coordinates of robot at step k, and $\mathbf{p}(k+1)$ represents the new position at step

$k+1$. The goal position is defined as an equilibrium point $\mathbf{p}^*$ for the system.

The computation over the whole environment of vector $\vec{V}$ defines a vector field $\mathbf{V}$. Let us consider a partial set of equivalent statements about a generic vector field $\mathbf{V}$ [15].

- any oriented simple closed curve $\mathbf{c}$: $\oint_{\mathbf{c}} \mathbf{V} \cdot d\mathbf{s} = 0$
- $\mathbf{V}$ is the gradient of some function U: $\mathbf{V} = \nabla U$

The former is related to the concept of *conservativeness* of the field. The latter is concerned with the existence of a *potential function* generating an unique field. From a different point of view, conservativeness is a measure of the quality of landmark learning, whereas the existence of a Lyapunov potential function indicates the robot's capability to reach the goal. The following Section addresses the former aspect. Details of the other aspect can be found in [2, 3].

The robot *Nomad200* was used to accomplish the tests. It includes the *Fujitsu Tracking Card (TRV)* which performs real-time tracking of full color templates at a NTSC frame rate (30Hz).

4 Principles for landmark learning

A landmark must be reliable for accomplishing a task as detailed in Section 2.1. Landmarks that appear to be appropriate for human beings are not necessarily appropriate for other agents (animals, insects or artificial beings) because of the completely different sensor apparatus and matching systems [17]. Therefore we need to state the meaning of landmark reliability in advance for the system in use before to solve the problem of selecting landmarks.

For our system, a template is a region of the grabbed image identified by two parameters m_x and m_y representing the sizes along X and Y axes. The size ranges from 1 to 8, i.e. from *small* (2^1 pixels wide) to *large* (2^8 pixels wide) templates. The TRV can simultaneously track many templates. For each template the card performs a match in a sub-area of the actual video frame adopting the block matching method [8]. This introduces the concept of *correlation* between the template being used and the actual video frame. The sub-area is composed of 16×16 positions in the frame usually taken around the origin (o_x, o_y) of the template (its upper-left corner). The whole set of computed correlation measures forms the *correlation matrix*. Examples of correlation matrices are reported in figure 1.

We can take advantage of the matrices to compute a measure that states upon the reliability of the template under study [16]. As reported in [2, 13] we calculate a figure r, ranging from 0 to 1, which states how deep is the global minimum of the matrix in relation to its neighborhood. Therefore, we define reliable landmarks as *templates which are uniquely identifiable* in their neighborhoods: the greater r the more uniquely identifiable the landmark in its sub-area.

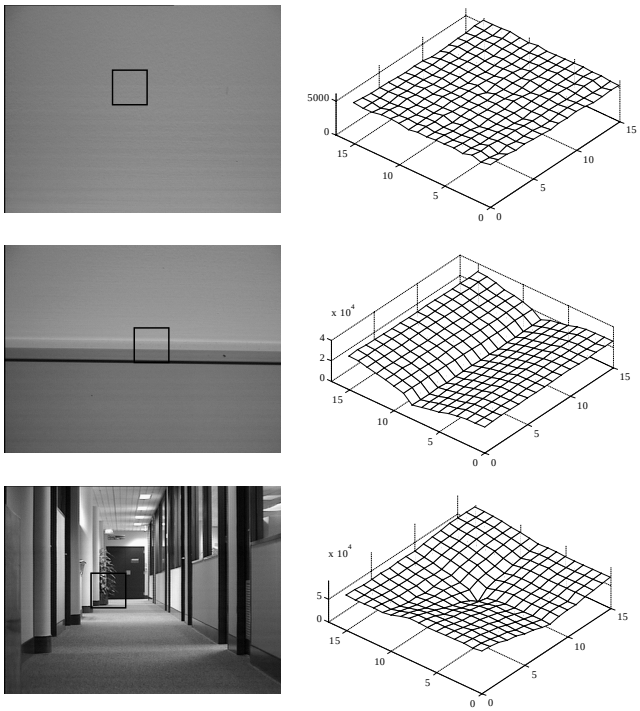

Figure 1: Examples of correlation matrices. These are computed within the local sub-area of the templates (square box in the pictures).

Once that the measure for the reliability of a landmark has been stated, the next step consists of searching the whole panorama for landmarks. There are several degrees of freedom in searching for the best landmarks within a video frame [2], but some simplifications can be introduced: only square templates are used, and the position of a landmark is searched for by maximizing the following:

$$(o_x^*, o_y^*) = \arg \max_{(o_x, o_y) \in \text{grid}} r_l(o_x, o_y) \qquad (2)$$

where $r_l(o_x, o_y)$ identifies the reliability factor for a landmark l whose origin is located in (o_x, o_y) representing a generic place on the grid. The position (o_x^*, o_y^*) represents the cells with the highest r. In or-

der to assure that different landmarks occupy different positions, previously chosen coordinates are not considered. In figure 2, examples of landmarks chosen have been reported. When different sizes are considered, different sets of landmarks are extracted.

The landmarks which have been *statically* chosen are used for navigation tasks. This is done by testing the landmarks to verify that they represent good guides for navigation tasks.

TBL helps to verify landmarks by testing whether during the motion the statically chosen landmarks are robustly identifiable. Through a series of stereotyped movements small perturbations (local lighting conditions, changes in camera heading, different perspectives and so on) can influence the reliability of the statically chosen landmarks. Figure 3 shows the robot moving away from the goal while the camera (arrows pointing to the top) is continously pointing towards the goal.

These perturbations to images naturally occur in typical robot journeys thus allowing us to state that the TBL phase represents a *testing framework* for landmarks. In other words, the robot tries to *learn* which landmarks are suitable for use in real navigation tasks by simulating the conditions the robot will encounter along the paths. At the end of the TBL process only those landmarks whose reliability r_l is above a certain threshold ϵ are suitable to be used in navigation tasks.

The reliability factor r_l for landmark l is continuously computed during the TBL phase through the following:

$$r_l = \frac{\sum_{i=1}^{TBL} r_l^i}{TBL} \tag{3}$$

where TBL is the total number of steps exploited till that time, and r_l^i is the reliability of landmark l calculated at time i. In the tests, at the end of the phase, TBL usually consists of 400 steps (it takes about 13 seconds to be performed). The set of landmarks is tracked along the whole TBL phase and r_l is continuously monitored for each landmark (details in [13]).

4.1 The quality of learning

There are strong connections between the learning phase and navigation actions. The conservativeness of the motion field bridges these two aspects.

A vector field $\mathbf{V}$ is said to be *conservative* when the integral computed on any closed path is zero. Conversely, if the field is not conservative then diverse potential functions can be associated with the field. This translates into *non-repeatability* of robot navigation trails in [13].

Figure 2: Different choices of landmarks for different landmark sizes. Landmarks are box-shaped.

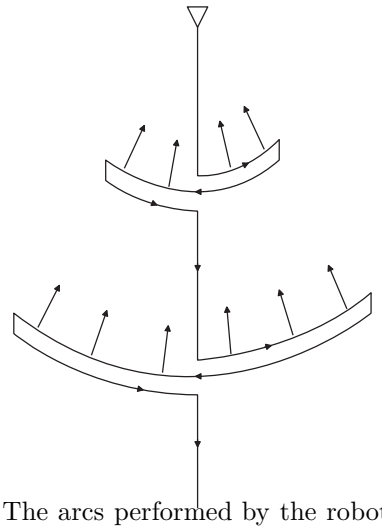

Figure 3: The arcs performed by the robot to implement the TBL phase

If the vector field is defined on a connected set in the environments, then the null *circuitation* property is equivalent to [15]:

$$\frac{\partial V_x(x,y)}{\partial y} = \frac{\partial V_y(x,y)}{\partial x} \tag{4}$$

We can measure how this equation differs from the theoretical null value as follows:

$$\frac{\partial V_x(x,y)}{\partial y} - \frac{\partial V_y(x,y)}{\partial x} \tag{5}$$

The property expressed by Equation 5 is referred to as *degree of conservativeness*. The degree of conservativeness of the vector field computed with a threshold set to 0 and landmarks sized 6 is shown in figure 4. Only small regions of the whole area roughly satisfy the constraint.

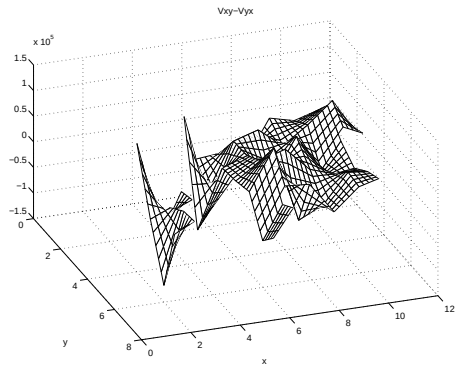

Figure 4: Conservativeness of a vector field computed with a TBL threshold of 0 and landmarks sized 6

A small change in the threshold for TBL can dramatically change the situation. In figure 5 the degree of conservativeness for each point is plotted.

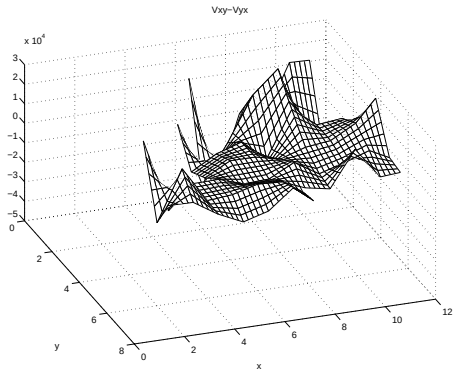

Figure 5: Conservativeness of a vector field computed with a TBL threshold of 0.1 and landmarks sized 6

A key consideration is concerned with the scale along Z: it is about one order of magnitude less than the one reported in figure 4. A trend toward a conservative field is thus becoming evident.

The situation obtained with a threshold of TBL set to 0.2 has been reported in figure 6. A large area of the environment has a degree of conservativeness that roughly equals 0.

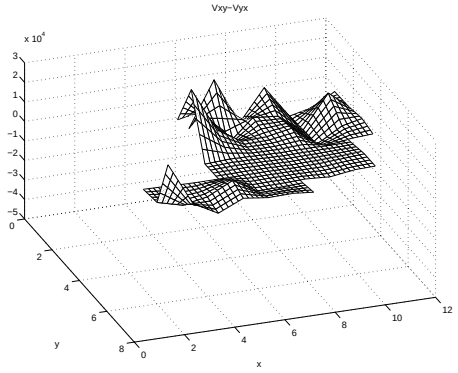

Figure 6: Conservativeness of a vector field computed with a TBL threshold of 0.2 and landmarks sized 6

Similar considerations can be expressed dealing with a different landmark size. For example, figure 7 shows the case where the TBL threshold is 0.2 and landmarks have a size of 4. The *template* of the graph is the same as before. Therefore, with a good choice of threshold the field becomes conservative regardless of the size of the landmarks.

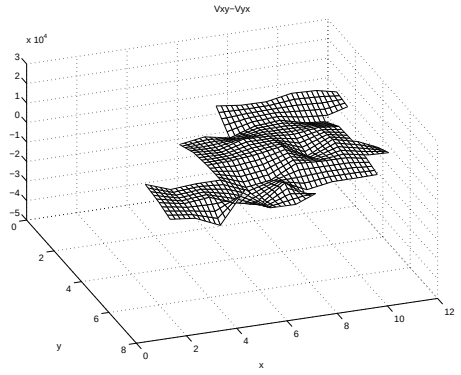

Figure 7: Conservativeness of a vector field computed with a TBL threshold of 0.2 and landmarks sized 4

5 Conclusions

Landmarks learning for robots can take inspiration from Biology but it needs to be well formalized for its efficient implementation in artificial agents. First, a definition for landmark reliability must be stated. Second, a measure that can assess about the quality of the learning phase needs to be introduced.

In this paper, we have shown how both these aspects can be efficiently addressed. Particularly, we have shown how the learning phase affects the navigation motion field. Further improvements to this study can be achieved by the use of omni directional visual sensors.

References

[1] A. Anderson. A model for landmark learning in the honey-bee. *Journal of Comp. Phys.*, (114):335–355, 1977.

[2] G. Bianco and A. Zelinsky. Biologically-inspired visual landmark navigation for mobile robots. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Kyongju (Korea), October 17-21 1999.

[3] G. Bianco and A. Zelinsky. Dealing with robustness in mobile robot guidance while operating with visual strategies. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, San Francisco (CA), April 24-28 2000.

[4] B. Cartwright and T. Collett. Landmark learning in bees. *Journal of Comp. Phys.*, A(151):521–543, 1983.

[5] K. Cheng, T. Collett, A. Pickhard, and R. Wehner. The use of visual landmarks by honeybees: Bees weight landmarks according to their distance from the goal. *Journal of Comp. Phys.*, A(161):469–475, 1987.

[6] T. Collett. Landmark learning and guidance in insects. *Phil. Trans. R. Soc. London*, B(337):295–303, 1992.

[7] J. Hong, X. Tan, B. Pinette, R. Weiss, and E. Riseman. Image-based homing. In *Proceeding of the IEEE Int. Conf. of Robotics and Automation*, pages 620–625, Sacramento, April 1991.

[8] H. Inoue, T. Tachikawa, and M. Inaba. Robot Vision System with a Correlation Chip for Real-time Tracking, Optical Flow and Depth Map Generation. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, pages 1621–1626, Nice, France, May 12–14 1992.

[9] S. Judd and T. Collett. Multiple stored views and landmarks guidance in ants. *Nature*, 392:710–714, 1998.

[10] M. Lehrer. Bees which turn back and look. *Naturwis*, (78):274–276, 1991.

[11] M. Lehrer. Why do bees turn back and look? *Journal of Comp. Physiology*, A(172):549–563, 1993.

[12] M. Lehrer. Honeybees' visual spatial orientation at the feeding site. In M. Lehrer, editor, *Orientation and communications in arthropods*, pages 115–144. Birkhauser verlag, Basel/Switzerland, 1997.

[13] M. Lehrer and G. Bianco. The turn-back-and-look behaviour: Bee versus robot. *Biological Cybernetics*, 2000.

[14] M. Lehrer and T. Collett. Approaching and departing bees learn different cues to the distance of a landmark. *Journal of Comp. Phys.*, A(175):171–177, 1994.

[15] J. Mardsen and A. Tromba. *Vector calculus*. W.H. Freeman and Company, 1996.

[16] T. Mori, Y. Matsumoto, T. Shibata, M. Inaba, and H. Inoue. Trackable attention point generation based on classification of correlation value distribution. In *JSME Annual Conf. on Robotics and Mechatronics (ROBOMEC 95)*, pages 1076–1079, Kavasaki (Japan), 1995.

[17] S. Thrun. A bayesian approach to landmark discovery and active perception in mobile robot navigation. Technical report, School of Computer Science Carnegie Mellon University, 1996.

[18] O. Trullier, S. Wiener, A. Berthoz, and J. Meyer. Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51:483–544, 1997.

[19] R. Wehner. Arthropods. In F. Papi, editor, *Animal Homing*, pages 45–144. Chapman and Hall, London, 1992.

[20] J. Zeil. Orientation flights of solitary wasps 1: Description of flights. *Journal of Comp. Phys.*, A(172):189–205, 1993.

[21] J. Zeil. Orientation flights of solitary wasps 2: similarities between orientation and return flights and the use of motion parallax. *Journal of Comp. Phys.*, A(172):207–222, 1993.

[22] J. Zeil, A. Kelber, and R. Voss. Structure and function of learning flights in bees and wasps. *Journal of Experimental Biology*, 199:245–252, 1996.