ELSEVIER

Robotics and Autonomous Systems 30 (2000) 39–64

Robotics and Autonomous Systems

www.elsevier.com/locate/robot

A mobile robot employing insect strategies for navigation

Dimitrios Lambrinos [a], Ralf Möller[a,b,*], Thomas Labhart [b], Rolf Pfeifer [a], Rüdiger Wehner [b]

[a] *AI Lab, Department of Computer Science, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland*

[b] *Department of Zoology, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland*

Abstract

The ability to navigate in a complex environment is crucial for both animals and robots. Many animals use a combination of different strategies to return to significant locations in their environment. For example, the desert ant *Cataglyphis* is able to explore its desert habitat for hundreds of meters while foraging and return back to its nest precisely and on a straight line. The three main strategies that *Cataglyphis* is using to accomplish this task are *path integration*, *visual piloting* and *systematic search*. In this study, we use a synthetic methodology to gain additional insights into the navigation behavior of *Cataglyphis*. Inspired by the insect's navigation system we have developed mechanisms for path integration and visual piloting that were successfully employed on the mobile robot *Sahabot 2*. On the one hand, the results obtained from these experiments provide support for the underlying biological models. On the other hand, by taking the parsimonious navigation strategies of insects as a guideline, computationally cheap navigation methods for mobile robots are derived from the insights gained in the experiments.

Keywords: Insect navigation; Robot navigation; Polarization vision; Path integration; Visual landmark navigation

1. Introduction

In recent years the idea of "learning from nature" is rapidly spreading through a number of scientific communities: computer science (artificial intelligence, artificial evolution, artificial life), engineering (bionics), and robotics (biorobotics). The main goal is to exploit the impressive results achieved by the blind but potent designer "Evolution". Among the most awesome capabilities exhibited by natural systems are the navigational skills of insects. Despite their diminutive brains, many insects accomplish impressive navigation tasks. Desert ants *Cataglyphis*, for example, make foraging excursions that take them up to 200 m away from their nest. On finding a suitable prey, they return home unfailingly and in a straight line [47] (see Fig. 1).

Cataglyphis cannot use pheromones to retrace its trail in order to return back to its nest, since the pheromones evaporate in a few seconds because of the high ground temperatures. More than two decades of field work (for a review see [50]) have revealed many details about the behavioral repertoire and the underlying mechanisms that *Cataglyphis* employs when homing. The three main strategies used are *path integration*, *visual piloting*, and *systematic search* [52]. Whereas path integration based on compass information gained from the polarization pattern of the sky is the primary navigation strategy of the ants, geocentered information based on landmarks is also used in order to finally pinpoint the nest.

Although there is a large number of behavioral data about the navigation behavior of *Cataglyphis*,

* Corresponding author. Tel: +41-1-63-65722; fax: +41-1-63-56809.

E-mail address: moeller@ifi.unizh.ch (R. Möller).

0921-8890/00/$ – see front matter
PII: S0921-8890(99)00064-0

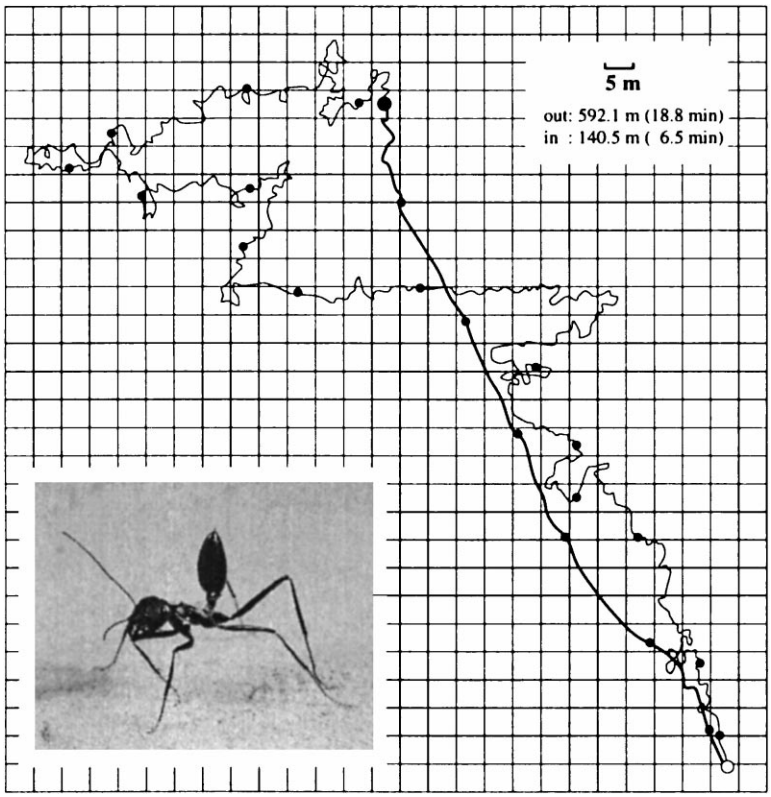

Fig. 1. A typical foraging trip of the Saharan ant *Cataglyphis* (inset). Starting at the nest (open circle), the ant searches for food on a random course (thin line) until it finds a prey (position marked with the large filled circle). The food is carried back to the nest on an almost straight course (thick line). Adapted from [56].

and some mechanisms of peripheral signal processing have been unraveled, it is still largely unknown how the navigation system is implemented in the insect's brain. In this paper we use the *autonomous agents* approach (see e.g. [37]) to gain additional insights into the navigation behavior of insects. The goal of this approach is to develop an understanding of natural systems by building a robot that mimics some aspects of their sensory and nervous system and their behavior.

This "synthetic" methodology has a number of advantages. Computer simulations of models are a first step of synthetic modeling. While it is often the case that models of biological agents are only described verbally or outlined implicitly, computer simulations require an explicit, algorithmic model, which helps to avoid pitfalls in terms of unwarranted assumptions or glossing over details. Especially the behavior of feedback systems is difficult to predict without simulations, and moving agents receive a rich and complex feedback on their actions from the environment. However, the value of computer simulations is limited by the fact that properties of the environment are usually difficult to reproduce in simulations. Wrong assumptions about these properties may severely misguide the development of models. The necessary step from a simulation to the real world is done by constructing artificial agents (mobile robots) and exposing them to the same environment that also the biological agents experience. Moreover, in contrast to animal experiments, the observed behavior of an artificial agent can be linked to its sensory inputs and its internal state. The advantages of proceeding in this way are illustrated in our recent studies, in which an autonomous agent navigated using directional information from skylight [29]. A similar line of research

is also pursued by other groups. Autonomous agents have, for example, been used to study the visuomotor system of the housefly [17], visual odometry in bees [9,38,39], cricket phonotaxis [31,41–45], six-legged locomotion of insects [14,15], and lobster chemotaxis [13,21].

In this study, (1) we give a brief overview of a polarized-light compass, and how it was employed in a path-integration system, (2) we describe an implementation of a biological model of visual landmark navigation using a panoramic visual system, (3) we present and discuss data on the navigational performance of the mobile robot *Sahabot 2*, and (4) we propose a new, more parsimonious model for visual landmark navigation.

2. Path integration

Path integration — the continuous update of a home vector by integrating all angles steered and all distances covered — is a navigation method widely employed by both animals and mobile robots. To use this navigation mechanism, both distance information, and even more important, directional information must be available. In the simplest method of path integration used in robotics, distance and directional information are derived from wheel encoders. There are several reasons for the wide application of this simple path-integration method in robotics. First, for short distances path integration using wheel encoders can provide relatively accurate position estimation, second, it is computationally cheap, and third, it can be done in real time. However, path integration with wheel encoders is prone to cumulative errors. Especially accumulation of orientation errors will cause large position errors that increase significantly as a function of the distance traveled. Several approaches for dealing with these errors have been proposed (see [5] for an overview). The most common approaches use either specialized heading sensors, like gyroscopes and magnetic compasses, or dedicated methods for reducing odometry errors.

An example of this kind of approach is work done in projects related to NASA's Mars missions. Path integration on the *Sojourner* Mars rover was performed with wheel encoders and a solid state turn rate sensor. Because of the errors introduced in the path-integration system, the position of the rover had to be updated daily by observing the rover from the lander, sending the images to Earth, detecting the rover in the images, and sending the rover position and heading back to the rover via the lander. The performance of the rover's path-integration system was evaluated prior to the mission [32]. For a distance of 10 m, standard deviations of 125 and 24 cm were predicted for the lateral and the forward errors, respectively. When the rover had to follow trajectories where turning or driving over rocks was necessary, these errors were much larger. The main conclusion from these experiments was that reaching a target position even in a distance of less than 10 m from the lander would require external update of the position of the rover. Future missions involve scenarios where the rover will have to cover greater distances. This will require solving the navigation problems and granting greater autonomy to the rover. There are plans to use a sun compass for obtaining the orientation of the rover, but this has not been tested yet.

Central-place foragers such as bees and ants, which primarily employ path integration to return to important places in their environment, are known to gain the compass direction from celestial cues, mainly from the polarization pattern of the blue sky (reviewed in [49]). In the following we describe a technical polarized-light compass system derived from the corresponding system in insects and its application in path-integration experiments.

2.1. Polarization vision in natural agents

Whereas in unpolarized light the e-vector oscillation occurs equally in all planes, in plane-polarized light there is a bias in favor of a certain oscillation plane (direction of polarization). Polarization in the sky occurs because the sunlight is scattered by atmospheric molecules. The degree of polarization is greatest for light scattered at an angle of 90° to the sunlight rays. The directions of polarization (e-vectors) observed in the sky form a regular pattern across the entire celestial hemisphere (Fig. 2). The pattern of polarization has a mirror symmetry with respect to the plane defined by the solar and the anti-solar meridian (SM and ASM in Fig. 2). Due to the daily westward movement of the sun across the sky (by some 15° per hour),

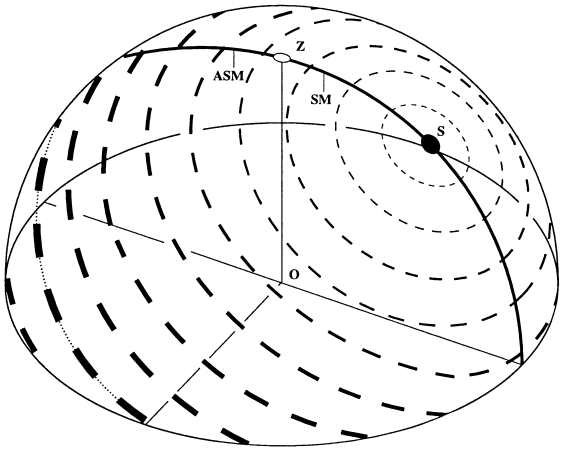

Fig. 2. 3-D representation of the pattern of polarization in the sky as experienced by an observer in point O. Orientation and width of the bars depict the direction and degree of polarization, respectively. A prominent property of the pattern is a symmetry line running through sun (S) and zenith (Z), called "solar meridian" (SM) on the side of the sun and "anti-solar meridian" (ASM) on the opposite side. Adapted from [46].

the symmetry plane, and with it the whole e-vector pattern, rotates about the zenith. The pattern retains two important characteristics over the day: its mirror symmetry, and the property that along the symmetry line the e-vectors are always perpendicular to the solar meridian.

Although the polarization pattern of the sky is invisible to humans, many insects exploit it by using the e-vector directions as a reference for compass orientation. Insects have compound eyes consisting of many discrete eyelets, the so-called ommatidia. Polarization vision is mediated by only a small group of specialized, upward-directed ommatidia situated at the dorsal rim of the eye (POL-area). Each ommatidium contains two sets of polarization-sensitive photoreceptors, which are tuned to orthogonal e-vectors (crossed-analyzer arrangement); for a review see [49].

In insects, polarization-sensitive neurons (POL-neurons) have been found in two different neural centers, the medulla of the visual lobes (crickets: [27], ants: [28]) and the central body (locusts: [40]). The medullar neurons receive input from photoreceptors of the POL-area. The activity of these neurons is a sinusoidal function of e-vector orientation with an excitatory and an inhibitory part and with the maxima and minima separated by 90°, indicating that they receive antagonistic input from two polarization-sensitive channels with orthogonal e-vector tuning orientations. It is probable that the two channels are represented by the crossed-analyzer configuration of the photoreceptor cells within each ommatidium of the POL-area (Fig. 3). The crossed-analyzer configuration has the advantage that it enhances e-vector contrast sensitivity and that it makes the e-vector response insensitive to fluctuations of light intensity. There are three types of cricket POL-neurons that are tuned to different e-vector orientations, i.e., approximately $10^\circ, 60^\circ$, and 130° with respect to the body axis.

2.2. Polarization vision in an artificial agent

Recent studies have shown that a polarized-light compass implemented in analog hardware is superior for extracting compass information compared to traditional methods based on proprioception [29]. The basic components of the POL-compass are polarization-opponent units (POL-OP units) that are functionally similar to the POL-neurons found in insects. Each POL-OP unit consists of a pair of *polarized-light sensors* (POL-sensors) followed by a log-ratio amplifier (see Fig. 4). The POL-sensors are photodiodes with a linear polarizer and a blue transmitting filter on top. In each POL-OP unit the polarizing axis of one POL-sensor was adjusted 90° to the polarizing axis of the other sensor, thus mimicking the crossed-analyzer configuration in the POL-area of insect eyes. The signals of each pair of POL-sensors were fed into a log-ratio amplifier. Three pairs of POL-sensors were mounted on our mobile robot *Sahabot 2* (see Fig. 5) and adjusted such that the polarizing axis of the positive channel was $0^\circ, 60^\circ$, and 120° with respect to the robot's body axis. The visual fields of the POL-OP units were about 60° centered around the zenith. The output of a POL-sensor is described by the following equation:

$$s(\phi) = K I (1 + d \cos(2\phi - 2\phi_{\max})). \quad (1)$$

I is the total intensity $I = I_{\max} + I_{\min}$, with $I_{\max}, I_{\min}$ being the maximum and the minimum intensity, respectively. d is the degree of polarization, ϕ is the current orientation with respect to the solar meridian, $\phi_{\max}$ is the value of ϕ that maximizes s (tuned e-vector

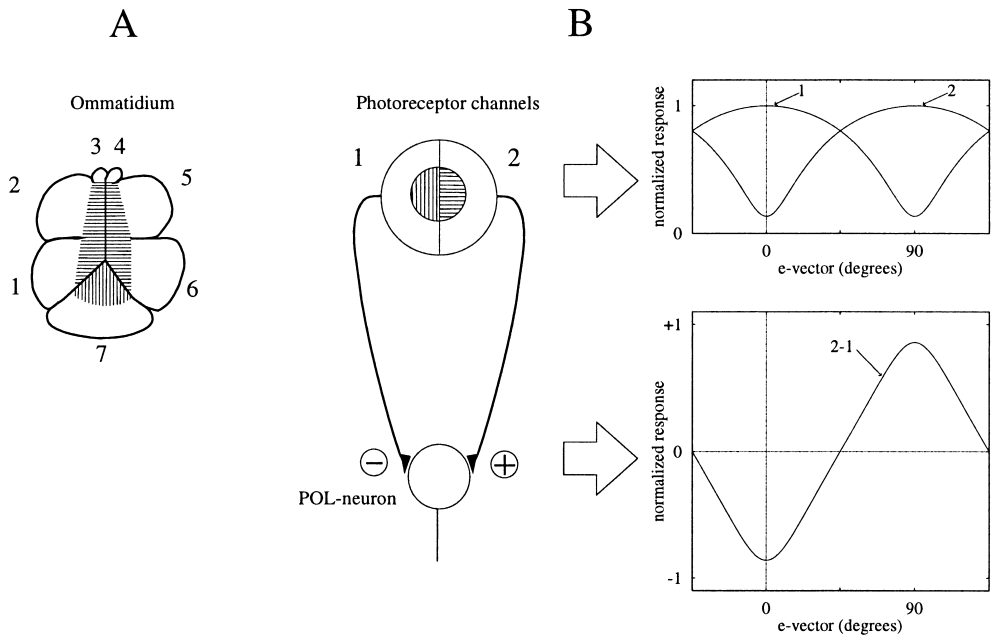

Fig. 3. Processing of polarized-light information in the insect nervous system (after [27]). (A) Cross-section through an ommatidium of the polarization-sensitive part of the cricket compound eye. Receptors are grouped in two channels with tuning to orthogonal directions of polarization (hatched areas). (B) Principle of operation of a polarization-opponent neuron (POL-neuron) in the insect optic lobe. The POL-neuron receives antagonistic input from the two receptor channels. Its response function represents a difference function (2-1) of the receptor response functions (1, 2).

direction), and K is a constant [4]. The outputs of the POL-OP units are described by

$$p_1(\phi) = \log\left(\frac{1 + d\cos(2\phi)}{1 - d\cos(2\phi)}\right),$$
$$p_2(\phi) = \log\left(\frac{1 + d\cos(2\phi - \frac{2}{3}\pi)}{1 - d\cos(2\phi - \frac{2}{3}\pi)}\right), \quad (2)$$
$$p_3(\phi) = \log\left(\frac{1 + d\cos(2\phi - \frac{4}{3}\pi)}{1 - d\cos(2\phi - \frac{4}{3}\pi)}\right),$$

where $p_1(\phi)$, $p_2(\phi)$, and $p_3(\phi)$ are the outputs of the POL-OP units tuned to 0°, 60°, and 120°, respectively.

2.3. Methods for extracting compass information

There are two groups of models for using the POL-OP responses to derive compass direction: *scanning models* and *simultaneous models* [29]. In *scanning models*, the agent has to first find the solar meridian, and use it as a reference direction (0°) for its proprioceptive system. For doing that it has to actively scan the sky by rotating around its vertical body axis. When the output signal of one POL-OP unit (or a suitable combination of multiple POL-OP units) reaches its maximum, the robot is known to be aligned with the solar meridian. After having found the reference direction, it uses proprioceptive information to find its heading direction.

In contrast, with a *simultaneous model*, the heading direction can be determined continuously and no scanning movements are necessary during the journey. In the previously performed experiments [29], the compass direction was obtained by comparing the current output values of the POL-OP units with a *lookup table* that associates the output values of the POL-OP units with the corresponding orientation of the robot. The lookup table was recorded before each experiment by a single 360° rotation of the robot. This study presents a simultaneous model which does not require a lookup table, but uses an *analytical* procedure to derive compass information from the values of the POL-OP units.

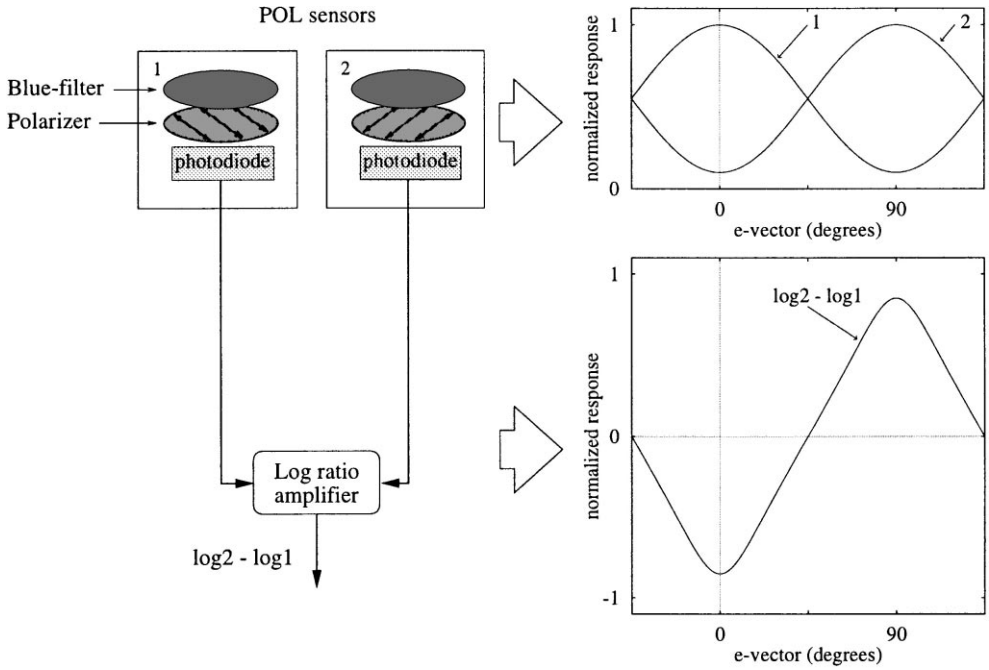

Fig. 4. Diagrammatic description of a polarization-opponent unit (POL-OP unit). A POL-OP unit consists of a pair of POL-sensors and a log-ratio amplifier. The log-ratio amplifier receives input from the two POL-sensors and delivers the difference of their logarithmized signals. The e-vector responses of the POL-sensors (1, 2) follow a $\cos^2$-function.

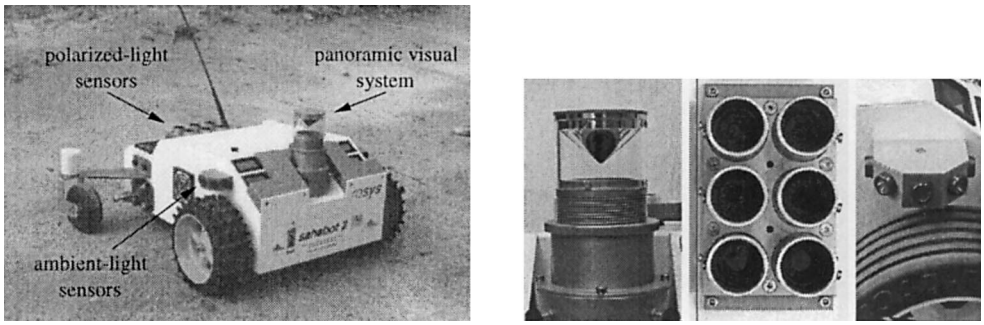

Fig. 5. Left: The mobile robot *Sahabot 2*. Right: Robot sensors. From left to right: the panoramic visual system, polarized-light sensors, ambient-light sensors.

Although the output signals of the POL-OP units are independent of the light intensity due to the cross-analyzer configuration, they still depend on the degree of polarization (see Eqs. (2)). The amplitude of the POL-OP signals is proportional to the degree of polarization, which changes during the day because of the changing elevation of the sun (see Fig. 6) and due to clouds. Fig. 7 (left) shows the outputs of the POL-OP units during a rotation of 360° at two different times of the day.

One way to eliminate the dependence on the degree of polarization is to use a *scanning model* where only the maxima of the signals are evaluated. In a *simultaneous model*, however, the change in the polarization pattern during the day has to be taken into account. This can be done by either regularly

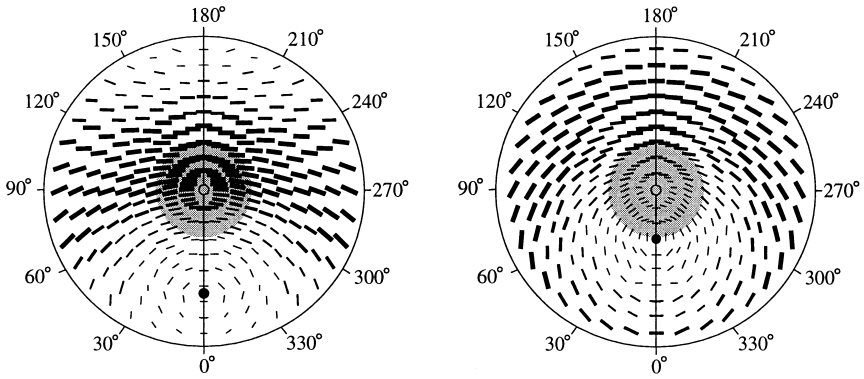

Fig. 6. 2-D representation of the pattern of polarization in the sky (for a 3-D representation see Fig. 2) for two different elevations of the sun (filled circle). The e-vector directions and the degree of polarization are indicated by the orientation and width of the black bars. The shaded area indicates the visual field of the POL-OP units. Left: Sun elevation of 25°. Right: Sun elevation of 60°. Adapted from [49].

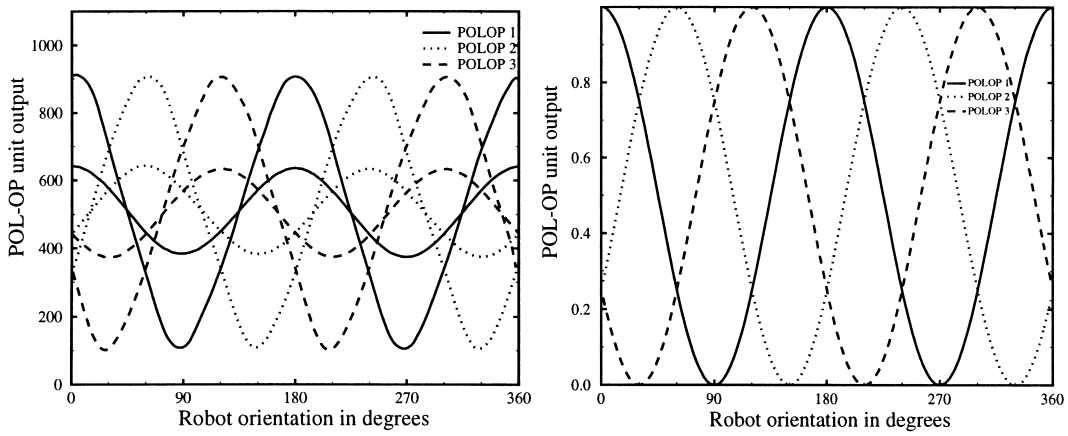

Fig. 7. Left: The outputs of the POL-OP units during a rotation of 360° at two different times of the day (i.e., with different solar elevation), late in the afternoon (large amplitude) and early in the afternoon (small amplitude). Right: The outputs of the POL-OP units shown in the left after normalization. The outputs of the corresponding units in the two different runs are completely overlapping in the figure.

updating the lookup table or by normalizing the outputs of the POL-OP units in the following way. First, the POL-OP signals are delogarithmized by applying a sigmoid function:

$$\frac{1}{10^{p(\phi)}+1} = \overline{p}(\phi). \tag{3}$$

Eqs. (2) then become:

$$1 - 2\overline{p}_1(\phi) = d\cos(2\phi), \tag{4}$$

$$1 - 2\overline{p}_2(\phi) = d\cos\left(2\phi - \frac{2\pi}{3}\right), \tag{5}$$

$$1 - 2\overline{p}_3(\phi) = d\cos\left(2\phi - \frac{4\pi}{3}\right). \tag{6}$$

Two of the above equations can be used to derive the values of d and ϕ directly. For example from Eq. (4) we find:

$$d = \frac{1 - 2\overline{p}_1(\phi)}{\cos(2\phi)} \tag{7}$$

and by substituting the value of d in Eq. (5) we get:

$$\tan(2\phi) = \frac{3 - 2\overline{p}_1(\phi) - 4\overline{p}_2(\phi)}{\sqrt{3}(1 - 2\overline{p}_1(\phi))}. \tag{8}$$

From the above we find:

$$\phi = \frac{1}{2}\arctan\left(\frac{\overline{p}_1(\phi) + 2\overline{p}_2(\phi) - \frac{3}{2}}{\sqrt{3}(\overline{p}_1(\phi) - \frac{1}{2})}\right). \tag{9}$$

Since $p_1(\phi)$, $p_2(\phi)$, $p_3(\phi)$ are known (sensor readings), and $\overline{p}_1(\phi)$, $\overline{p}_2(\phi)$, $\overline{p}_3(\phi)$ are derived from these values, we can determine the values of ϕ and d. Note that for this step only *two* POL-OP units are necessary (p_1 and p_2 in the above equations).

Since the functions in Eqs. (4)–(6) have a period of π, there will be two values of ϕ (ϕ and $\phi+\pi$) satisfying Eqs. (4)–(6) for each triplet $\overline{p}_1(\phi)$, $\overline{p}_2(\phi)$, $\overline{p}_3(\phi)$. This ambiguity cannot be resolved by referring to the POL-OP unit signals. A different sensory modality has to be involved. Disambiguation between the two candidate orientation values is done by employing a set of ambient-light sensors on the robot. Eight ambient-light sensors are arranged in two half-circles covering a visual field of 180° each. They are mounted in the front part of the robot (see Fig. 5). The values from the ambient-light sensors are used to obtain a rough estimate of the robot's heading with respect to the sun. If the robot is oriented towards the sun then the ambient-light sensor with the visual field enclosing the solar meridian will have a stronger response than the other sensors. This estimate is used to disambiguate between the two candidate orientations ($\phi, \phi+\pi$) obtained from the POL-OP units. In ants, spectral information from the sky is sufficient to serve as a compass cue (without any polarized-light information), and can, in addition, be used to solve the ambiguity problem mentioned above [50].

We can now transform the current POL-OP readings to signals that are independent of the degree of polarization. This is done by substituting the value of d from Eq. (7) in Eqs. (4)–(6) instead of using Eq. (9). To optimize precision, the value of ϕ can be calculated by selecting the equation of the POL-OP unit for which the signal curve is steepest at that point, i.e., POL-OP unit 1 for 30°–60°, POL-OP unit 2 for 60°–90° and POL-OP unit 3 for 90°–120°.

Fig. 7 (right) shows the results of applying such a normalization process to the outputs of the POL-OP units shown in Fig. 7 (left). The outputs of the POL-OP units during a rotation of 360° were recorded at two different times of the day, late and early in the afternoon. The signals become independent of the degree of polarization after the normalization process and are practically identical. Thus, they can be used in a simultaneous model to calculate the orientation of the robot with respect to the solar azimuth at any time of the day.

2.4. Ephemeris function

As shown above, it is possible to derive the orientation of the agent with respect to the solar azimuth by using the polarized-light compass. To use the POL-compass over an extended period of time, however, requires compensating for the daily movement of the sun. This task is complicated by the fact that the rate of change of the sun azimuth is not constant. The sun azimuth changes slowly in the morning and faster around noon. Moreover, the ephemeris function, i.e., the function that describes the change of sun azimuth over time, depends on the season and the geographical latitude.

There is behavioral evidence that insects are equipped with such an ephemeris function [30]. Moreover, recent studies in both bees [16] and ants [53] indicate that the ephemeris function might be innate and that insects can refine its precision by experience. In our study we extended the model of the POL-compass described above by including the ephemeris function valid for the location and season of our experiments.

2.5. Path-integration mechanism

The directional information obtained from the POL-compass was used in a path-integration mechanism to keep an estimate of the robot's position over time (for the ants' performances see [22,34]). The position of the robot was calculated as follows:

$$x(t+\Delta t) = x(t) + \cos(\theta(t))v(t)\Delta t, \tag{10}$$

$$y(t+\Delta t) = y(t) + \sin(\theta(t))v(t)\Delta t, \tag{11}$$

where $x(t+\Delta t)$, $y(t+\Delta t)$, $x(t)$ and $y(t)$ are the x and y coordinates of the robot at time $t+\Delta t$ and t, respectively, with Δt denoting the time step. The velocity of the robot $v(t)$ was estimated from the wheel encoders of the robot. The wheel encoders of the *Sahabot 2* are mounted on the axes of the two motors that drive the two front wheels (see Fig. 5) and give 6300 pulses per wheel revolution, which corresponds to 13 pulses per mm of distance traveled. $\theta(t)$ is the estimated orientation of the robot at time t, obtained either from the polarized-light compass or, alternatively, from the difference between the accumulated wheel encoder values of the left and right wheel.

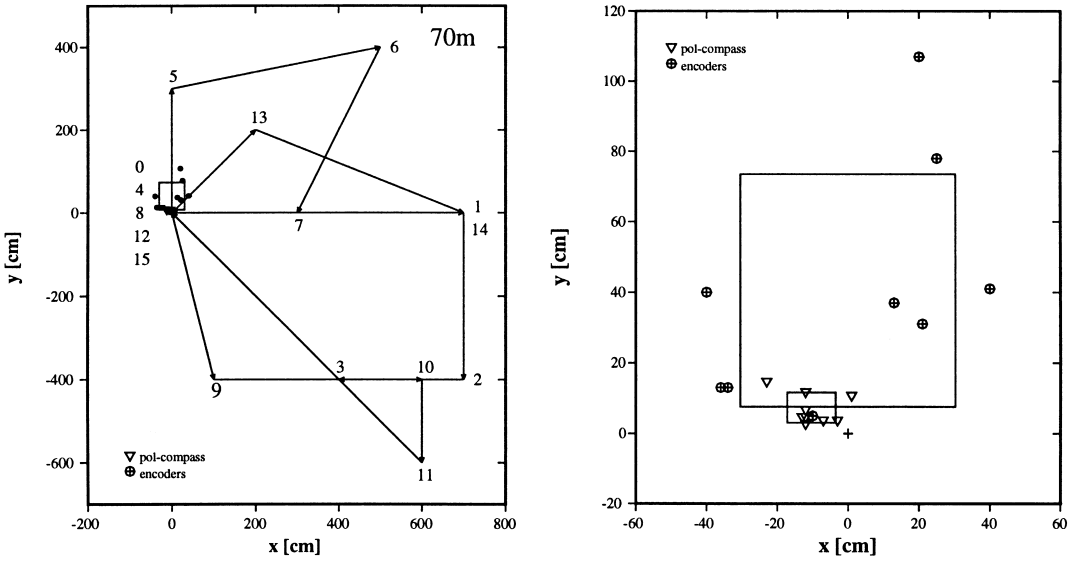

Fig. 8. Results of the experiments with the 15-segment trajectory for both POL-compass and proprioceptive system. Left: The trajectory is drawn starting from the origin with the segment number printed at the end of each segment. Note that the trajectory passes three times through the origin. Right: Zoom into the region around the starting position. The final positions of the robot during the experiments are drawn as circles (wheel-encoder experiments) or triangles (POL-compass experiments). Mean and standard deviation of the data in the direction of the axes are indicated by the centers and surroundings of the boxes for the two cases.

In the experiments described below, the performance of the path-integration system obtaining directional information ($\theta(t)$) from the POL-compass (POL-compass system) is compared with a system using proprioception to determine the orientation (proprioceptive system). Both systems were running in parallel providing independent estimates of the robot's position over time, but in each experiment only one of the systems was used to control the position of the robot. In both cases, the distance traveled (the term: $v(t)\Delta t$ in Eq. (11)) was estimated by referring to the wheel encoders.

2.6. Experiments

The path-integration experiments were performed in Zurich (47.38°N, 8.55°E) on a square covered with a soft synthetic material. Before each experiment, the robot was placed at the starting position and aligned with the geographic north. The robot was programmed to follow a trajectory consisting of a number of straight line segments of different orientations and different lengths by using one of the two path-integration systems, namely the POL-compass system or the proprioceptive system. The last segment of each trajectory led back to the starting position. Two different pre-programmed trajectories were used, one consisting of 15 segments with a total length of 70 m (Fig. 8) and the other of 30 randomly generated segments with a total length of 255 m (Fig. 9).

During the experiments, the data from the sensors, i.e., the outputs of the POL-OP units, the light sensors and the wheel encoders were logged to the robot's on-board PC. At the end of each experiment, the final position of the robot was measured with a meter stick. The distance between the final position of the robot and the starting position was used as an overall error measurement. To make a fair comparison, the wheel encoders of the robot were calibrated carefully using the square-path method described in [5].

2.7. Results

The data on overall robot performance in the path-integration experiments with the 15-segment trajectory are given in Table 1 and will be summarized below. The robot performed well with both the proprioceptive and the POL-compass system. The largest

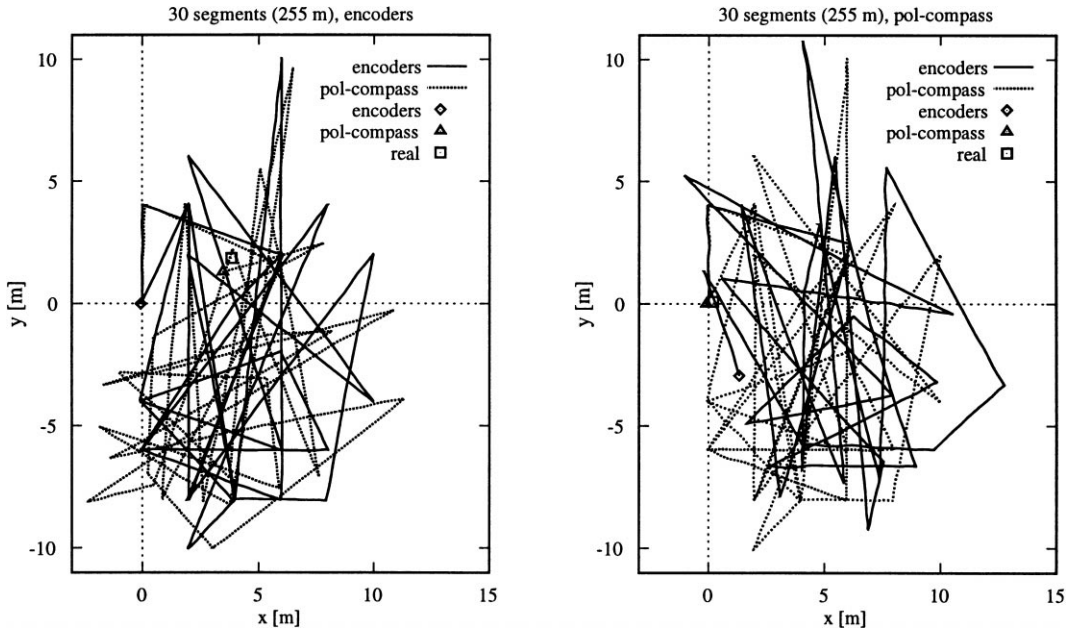

Fig. 9. Trajectories estimated by both the POL-compass system and the proprioceptive system during an experiment with a 30-segment trajectory. The final positions of the robot as estimated by both systems and the real final position are also drawn. Left: The robot was controlled by the proprioceptive system. Right: The robot was controlled by the POL-compass system.

Table 1
Overall performance of the robot with the two path-integration systems in the experiment with the 15-segment trajectory. Average error and standard deviations indicate deviations from the target position at the end of the experiments. *n* indicates the number of experiments

	POL-compass	Proprioception
Average error (cm)	13.5	51.9
Standard deviation (cm)	6.3	28.8
Min (cm)	5	11
Max (cm)	27	109
n	9	9

error recorded in these experiments was just about 1.1 m. The average errors were 13.5 and 51.9 cm for the POL-compass system and the proprioceptive system, respectively. The POL-compass system was clearly more precise than the proprioceptive system (U-test, $p = 0.0023$). Both the mean deviation from the target position and the standard deviation of the error were smaller. Fig. 8 visualizes the trajectory and the final positions of the robot.

During the experiments, the POL-compass system and the proprioceptive system produced different estimates of the position of the robot. Although the difference between these estimates was initially zero, it increased significantly in the course of a run. This can be seen clearly in two other experiments where the robot had to follow a trajectory consisting of 30 randomly generated segments. Fig. 9 shows the trajectories as estimated from the POL-compass and the proprioceptive system. In the two experiments shown, the robot was controlled by either the proprioceptive (Fig. 9, left) or the POL-compass (Fig. 9, right) system. As can be seen, the final positions of the robot as estimated with the POL-compass system (depicted by the triangles) are close to the real position of the robot (depicted by the squares). The deviation was 67 and 29 cm when controlled by the proprioceptive or the POL-compass system, respectively. The final position of the robot as estimated from the proprioceptive system is about 4.3 m away from the real position when controlled by the proprioceptive and 3.3 m when controlled by the POL-compass system. (Fig. 10) shows the deviation between the two estimates over time for the experiment depicted in Fig. 9 (left). The deviation increases over time reaching a value of 60° at the end of the 30th segment (after 12 min).

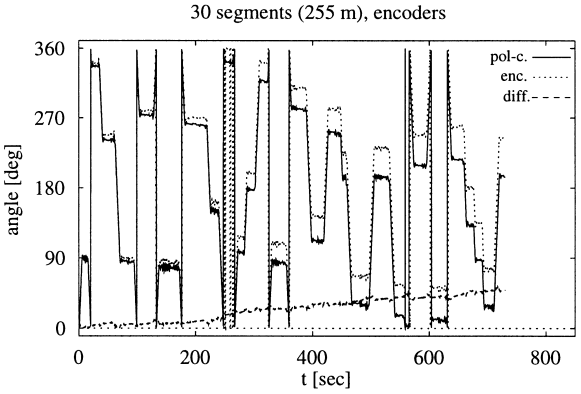

Fig. 10. Deviation between the orientation estimates by the POL-compass (solid line) and by the proprioceptive system (dotted line). At the beginning of the experiments both estimates are the same but the difference (dashed line) increases over time.

2.8. Discussion

The construction of the polarized-light compass for the *Sahabot 2* was guided by the experiences gained with its predecessor employed on the robot *Sahabot* [29]. While in the first version the polarizing axis of each POL-sensor could be adjusted independently, the polarizers are now rigidly mounted in openings of an aluminum plate, thereby excluding slackness of the adjustment gears as an error source, and reducing size and weight of the device. We retained the tuning to three e-vector directions with 60°-intervals since this guarantees a fairly constant resolution for all e-vector directions. A discussion of correspondences and differences between the POL-neurons of insects and the POL-OP units of the POL-compass can be found in [29]: Although the POL-OP units have basically the same properties as POL-neurons, there are simplifications concerning the number of sensors, the orientation of the visual fields, and the tuning directions.

Compass direction is derived from the outputs of the three POL-OP units with a novel *simultaneous model*. In contrast to a *scanning model*, no scanning movements are necessary during the journey to determine the orientation of the robot. Moreover, deviations from the set course can immediately be compensated. A previously tested *simultaneous model* [29] was based on the construction of a lookup table prior to the experiment. Changes in the degree of polarization along with changing elevation of the sun restrict the applicability of a lookup table to short-term excursions. In the *simultaneous model* presented in this study, the degree of polarization is eliminated by analytically combining the signals of two POL-OP units. This allows to determine the orientation of the robot with respect to the solar–antisolar meridian over extended periods of time.

In addition to the changes in the degree of polarization, in long-term runs errors in the orientation with respect to geographic north are introduced by the westward movement of the sun. To eliminate these errors, the novel *simultaneous model* was complemented with an ephemeris function describing the azimuthal position of the sun in dependence of the time of day. The geographical latitude, a parameter of the ephemeris function, was defined in the control program according to the location of the experiments.

In the experiments performed to evaluate the precision of the POL-compass system in the first version of the *Sahabot* [29], the robot had only to maintain a target course direction over a certain distance. The corresponding experiments performed in this study employ the POL-compass in a path-integration system that enables the robot to follow arbitrary trajectories, thus getting closer to the actual foraging behavior of the ants.

Although the distances covered in the excursions of ants and in the robot experiments are in the same range [47], it is difficult to compare the homing precision of these agents, since both their size and their method of propulsion are completely different. Therefore we merely summarize the results from animal experiments presented in [55] and contrast them with the results of the robot experiments. In the animal experiments, the end of the straight homeward run in a test field is marked by a sharp turn that indicates the position where the home vector becomes zero. One group of ants was approaching a known food source on a straight outbound run, covering a total distance of approximately 40 m. The second group of ants was foraging along tortuous routes before they accidently found a piece of food and returned home. From the foraging times given it can be concluded that the ants of the second group covered a mean distance of approximately 180 m. The distance of the turn point from the fictive nest position was on average 1.3 m for ants in the first group and 2.2 m for ants in the second group (for details see [55]). In the robot experiments,

the mean error amounted to 13 cm for a path length of approximately 70 m when using the POL-compass. For the two 30-segment trajectories with a length of 255 m estimated by the POL-compass system, errors of 67 and 29 cm were recorded.

The precision of the POL-compass system can be evaluated by a comparison of a path-integration system which derives orientation from the POL-compass with another system where orientation is computed from a carefully calibrated proprioceptive system (wheel encoders). Results obtained in experiments where the robot had to travel along a 15-segment trajectory with a total length of 70 m demonstrate that the standard deviation of the final position error is by a factor of approximately 5 smaller in the POL-compass system (Table 1). The systematic deviation of the mean position that was observed for the POL-compass system is probably caused by a slight inclination from the horizontal at the outer parts of the test field. This causes a shift of the visual field of the POL-OP units away from the zenith where the pattern exhibits a mirror symmetry (see Fig. 6). Integration over the visual field which is not centered on the solar–antisolar meridian will therefore result in a total e-vector which is not perpendicular to this meridian.

The use of an inertial navigation system for path integration would be another alternative to the above-mentioned methods. The advantage of inertial navigation systems is that they are self-contained and can provide very fast and dynamic measurements. The main disadvantage of such a system is that it drifts over time, a fact that also makes a direct comparison with the POL-compass system difficult. A sophisticated and relatively cheap inertial navigation system was developed by Barshan and Durrant-Whyte [1,2] based on combinations of different solid-state gyros. They reported a typical error of 12° after 5 minutes (commercial, high-quality and expensive inertial navigation systems have a typical drift of about 2 km after 1 hour).

While the design of the POL-OP units closely corresponds to the e-vector detection system of insects, all components of the path-integration system — the core of the *simultaneous model* (Section 2.3), the ephemeris function (Section 2.4), and the path-integration itself (Section 2.5) — are implemented in an analytical way without directly taking into account the specific processing capabilities of a neural substrate. It is, for example, known that ants solve the path-integration problem not by performing true vector summation, but by employing an approximation, resulting in systematic navigational errors under certain circumstances [34]. A neural model of the path-integration process was proposed in [22], but neurophysiological or neuroanatomical data are not available so far. Besides this, the experiments with path integration were conducted in order to evaluate the precision of the POL-compass system, which would have been more complicated by using the approximative procedures that animals resort to or by implementing plausible neural models. Taking the analytical models — which have been demonstrated to give the desired results — as a starting point, corresponding neural models could be derived.

3. Visual piloting

While path integration employing a skylight compass is the primary strategy that *Cataglyphis* ants are using to return to the vicinity of their nest, errors introduced by the path-integration process would result in a wrong estimate of the nest position [55]. Since the nest entrance is an inconspicuous hole in the desert ground (see Fig. 11) which is invisible to the insect even from a small distance, alternative strategies have to be employed in order to finally locate the entrance.

3.1. Visual piloting in natural agents

In the absence of visual landmarks, *Cataglyphis* will start a systematic search at the position where the nest is expected after having reset its path-integration system [51]. However, when landmark information is available, both bees [6] and ants [52] will exploit it, and they will relocate the target position directly with a remarkable precision (see Fig. 11).

A number of experiments performed with bees [6] and ants [52,54] have unveiled important properties of the insect's landmark navigation system. The main conclusion from these experiments is that the animal stores a rather unprocessed visual snapshot of the scene around the goal position. By matching this snapshot to the current retinal image, the insect can derive the direction it has to move in order to relocate the target position where the snapshot was taken. There is evidence from experiments in which several parame-

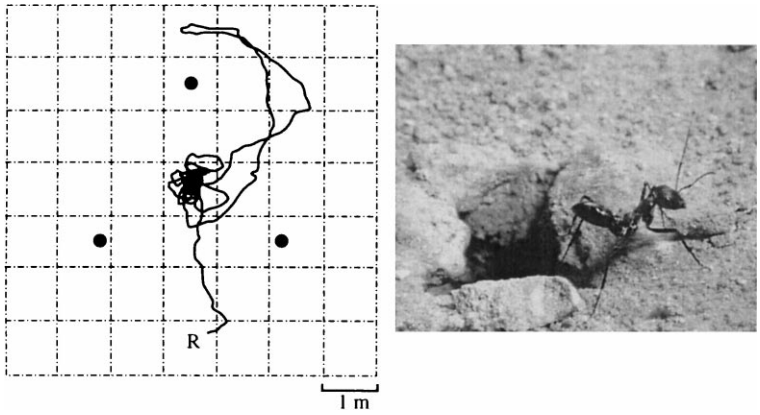

Fig. 11. Landmark navigation experiment with *Cataglyphis*. During the training period, the nest entrance was marked by three cylindrical landmarks (represented by filled circles). For the experiment, an ant that just returned from a foraging trip to the nest entrance (small hole shown in the photograph, right) was captured and displaced to a test field with an identical landmark configuration. Since the path-integration system was reset to zero, the ant could only rely on landmark information for locating the nest. It nevertheless searched very precisely at the center of the array, i.e., at the fictive location of the nest. Adapted from [52]; photograph by R. Wehner.

ters of the landmark configuration have been manipulated that the snapshot matching strategies exploit differences in the bearing, apparent width, and apparent height of the landmarks between snapshot and current retinal image. The models presented in the following only consider apparent width as one aspect of the apparent size of a landmark.

A model for the matching of snapshot and current view was proposed in [6], which reproduces many aspects of the search behavior of bees. Fig. 12 visualizes the matching process of this model. A snapshot (inner ring) is taken at the nest position marked with a cross and thick arrow in Fig. 12 (the arrow indicates the compass orientation of the snapshot, see Section 3.5). It consists of a number of filled and open sectors corresponding to landmarks and gaps between landmarks, respectively. Snapshot and current view (middle ring) match completely at the nest position. When the agent is displaced to a different position in the environment, both the azimuthal positions and apparent sizes of the landmarks on the retina change. As a result, snapshot and current view are not matched anymore. A home vector, pointing approximately to the direction of the nest position, can be derived by pairing each sector in the snapshot with the closest sector of the same type (open or filled sector corresponding to a landmark or a gap, respectively) in the current view. Note that for this pairing process snapshot and current view have to be aligned in the same compass direction (see Section 3.5). Each pairing generates two unit vectors that are attached to the center of the matched sector in the snapshot (in Fig. 12 the vectors are attached to the outer ring for clarity): a tangential vector pointing from the snapshot sector towards the paired sector in the current view, and a radial vector, which points centrifugally, if the apparent size of the current view sector is smaller than the size of its counterpart in the snapshot, and vice versa. The home vector (arrow originating from the center of the rings) is derived by summing all unit vectors. For the weighting of tangential and radial contributions a ratio of 1 : 3 was suggested [7].

3.2. Visual piloting in artificial agents

Traditional approaches to visual robot navigation are usually based on determining the robot position from geometrical maps (reviews can be found in [3,25,26]). Typically, the computational complexity required for these methods by far exceeds the limits that can be assumed for insects. Recent changes in the robotics paradigm towards “cheaper” methods replace position-based navigation by methods which regard “homing” as the basic navigation ability (for an overview see [19]). In particular, homing methods are used in conjunction with topological maps [18,26]. There are two different approaches to image-based

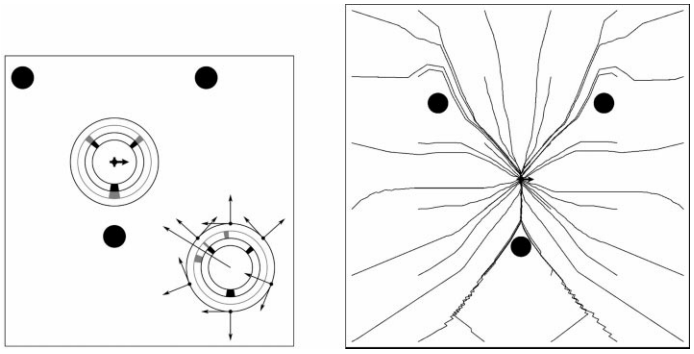

Fig. 12. Left: Diagrammatic description of the original *snapshot model* proposed in [6] for a configuration of three landmarks. The snapshot position is marked by a cross with a thick arrow (indicating the compass orientation associated with the snapshot). The matching process is visualized for two different positions, the snapshot position (center) and another position in a distance from the snapshot position (bottom right). The inner ring in the two diagrams depicts the snapshot (black), the second ring the current view (grey). Vectors contributing to the home vector (originating from the center) are attached to the outer ring. For clarity of the diagram, the weighting of the contributing tangential and radial vectors was chosen as 1 : 1, while a ratio of 1 : 3 was used in [6]. Right: Trajectories following the home vectors generated by the model. Each trajectory starts at a point on a grid.

homing, called "local" and "associative" homing [23]. *Local* homing methods gradually align a current view with a snapshot stored at the target position by deriving a direction of movement from the differences between the two images [19,23]. *Associative* homing methods associate stored scenes with movement vectors pointing towards the goal [8,20,35]. A disadvantage of the latter is that they require some additional means like path integration to determine the home direction which can be associated with a scene.

The *snapshot model* presented in the previous section belongs to the class of local homing methods. Hong et al. [23] describe another local method which is based on feature correlation. According to certain criteria, "characteristic points" are selected from each view and paired with characteristic points in the other view by computing a measure of similarity between the two windows around the characteristic points. In the snapshot model, on the contrary, only a small set of features — preferably those that are not affected by spherical distortions — have to be matched in the two images. In the original form of the model, these features are the two types of image regions; vertical contrast edges could be another choice [33]. A comparison of the complexity of the two methods is not feasible at the moment, since the environment used by Hong et al. was more complicated than the simple setup used in the experiments described below, and it is not known so far, how much effort has to be spent for the feature extraction in more complex environments. The matching process itself, though, is considerably simplified in the snapshot model, since a feature is matched to the closest feature of the same type according to their position in the image, instead of having to deal with image windows for that purpose. Therefore it was also feasible to derive a parsimonious neural equivalent of the snapshot model that provides additional support for the model from the neurobiological perspective [33], whereas it is not clear, if the same parsimony could also be achieved for a neural circuit implementing the method of Hong et al.

In order to employ the snapshot model on a mobile robot, the required sectorized representation has to be extracted from a camera image, which will be described in Section 3.4. Moreover, the snapshot model itself has to be modified so that it can be used together with the robot control strategy, as discussed in the following section.

3.3. Proportional vector model

In the original model presented in [6], all contributing vectors have unit length. This model does not take into account the magnitude of the differences between snapshot and current view. For each pair of sectors there are two unit vectors generated, irrespective of their differences in azimuthal position or size (see Fig. 12).

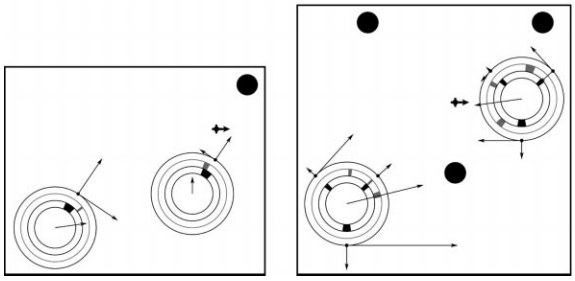

Fig. 13. The snapshot model with proportional vector contributions. Radial vectors have a length proportional to the difference in size between the paired sectors. The length of the tangential vectors is proportional to the difference in bearings of the paired sectors. For clarity only those vector contributions are shown that result from the pairing of landmark sectors in the two views, whereas gap sectors are not considered.

Having some information about the distance to the target position is important for mobile agents. When the agent is close to the target position it will have to slow down in order not to overshoot the target, while when it is far way it can speed up to reach the target area faster. In addition, distance information can be used to define the end of an experiment, i.e., to stop the robot when it is close to the target position. Here we propose an extension to the original snapshot model that incorporates this information into the model.

The unit vectors can be replaced by vectors with a length proportional to the difference in bearing, and apparent size, of the paired sectors. The resulting model, called *Proportional Vector model* (PV model), has in common with the original model that the matching process is performed the same way. Again each sector in the snapshot is paired with the closest sector of the same type in the current view resulting in two vectors. It differs from the original *snapshot model* in that the size of the radial vectors and the tangential vectors is proportional to the difference in size and bearing between the paired sectors, respectively (see Fig. 13). When the difference in azimuthal position between the paired sectors is large, the corresponding tangential vector will also be large, resulting in a larger contribution to the final home vector. The same applies for the radial vectors, i.e., when the difference in size between the paired sectors is large, then the contribution of this pairing to the final home vector will also be large. For contributions resulting from differences in bearing, a weighting of the vector length with the angular difference has been suggested before (see discussion in [19]).

This modification enables the use of the *length* of the home vector as a measure of discrepancy between current view and snapshot, called “disparity” in this study. This is not possible with the unit vector model, since the disparity is not decreased while approaching the snapshot position; it will only become zero if a perfect match between snapshot and current view is achieved (see Fig. 14). Using proportional contributions, the robot runs can be stopped when the disparity falls below a threshold. For the weighting of tangential and radial contributions, the same ratio of $1:3$ as in the original *snapshot model* was used.

3.4. Visual system and image processing

The camera system mounted in the front part of the robot (see Fig. 5) consists of a digital CCD

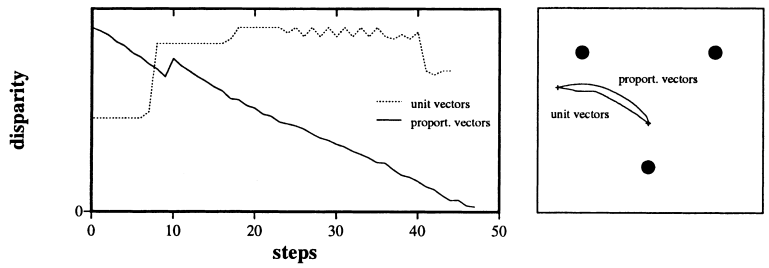

Fig. 14. Left: Time course of the disparity (length of the home vector) in a simulation of the original snapshot model using unit vector contributions, and of a model where vector contributions have a length proportional to differences in bearing and apparent size of sectors. Note that disparity is a relative measure that can be scaled arbitrarily. Within one simulation step, the home vector is computed and a movement with constant length in this direction is executed. Right: Corresponding trajectories running towards the snapshot position in the center.

camera and a conically shaped mirror in the vertical optical axis of the camera. The conical mirror was made of polished brass that went through a chrome-plating process to improve the optical quality.

With the help of the conical mirror, a panoramic, 360° view is obtained; a similar imaging technique was used in [9,19,39,57] (for a detailed description see [10]). When the axis of the cone coincides with the optical axis of the camera, horizontal slices of the environment appear as concentric circles in the image. Special adjustment screws located at the base of the camera module were used for tuning the optical axis of the camera. In order to see the horizon, the opening angle of the cone was determined by considering the visual field angle of the CCD camera. In the experiments described here, the opening angle of the cone was chosen so that the visual field extends ±10° around the horizon. The whole camera module was mounted as low as possible on the robot within the construction constraints, which brings the mirror to a height of approximately 27 cm above the ground. In order to reduce the total light intensity, a neutral density filter was mounted between camera and mirror. An additional infrared filter was necessary to prevent the influence of thermal radiation on the camera image.

Fig. 15 illustrates the image processing steps from the camera image to a horizonal view with black and white sectors as required by the matching mechanism described in Section 3.3. The same processing steps were applied to both snapshot and current views.

In a first step, the camera image (obtained from a situation similar to the one shown in Fig. 16) is transformed into an azimuth-elevation representation. A mean grey value is determined from the transformed image and used to adjust the brightness parameter of the camera in a way that the mean grey value of this image is kept constant over time. This is indispensable for the subsequent thresholding operation separating dark and bright regions. An area enclosing the horizon (horizonal area in Fig. 15) is then extracted from the thresholded image. The final one-dimensional horizontal view is obtained by counting the number of black pixels within the part of each pixel column contained in the horizonal area and applying a threshold of 50%.

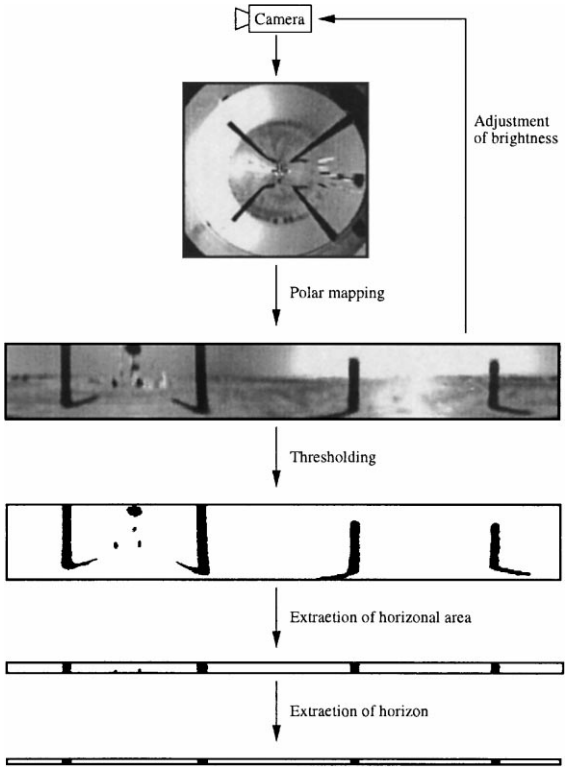

Fig. 15. Image processing for landmark navigation. The image from the 360° camera (160 × 120 pixel) is transformed into a polar view (351 × 56). After applying a threshold to each pixel, a horizonal area is extracted (351 × 20). The segmented horizon (351 × 1) is derived from the horizonal area. A pixel in the segmented horizon will be black, if more than 50% of the pixels in the corresponding column are black. The object between the two landmarks on the left is equipment in the vicinity, which was removed for the experiments.

3.5. View alignment

An important prerequisite of the matching procedure between snapshot and current view is the alignment of both views with respect to an external compass reference. Without proper alignment, the number of false matches between sectors — sectors in the two views that are paired, but do not correspond to the same landmark or gap — will increase significantly, resulting in an erroneous home vector. Bees might accomplish this alignment by maintaining a constant orientation of the body during learning and searching for a goal [12]. For some species of *Cataglyphis*, rotation on the spot interrupting the

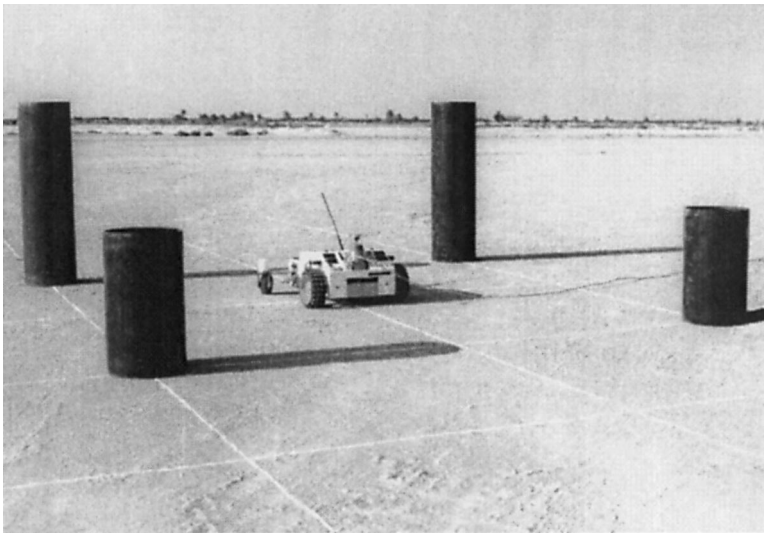

Fig. 16. Example of a landmark array used for the navigation experiments. The grid visible on the desert ground was used for the alignment of landmarks and robot, and for the registration of the final robot position.

forward movement has been observed [49], but it is not clear, if this behavior could also serve for the alignment of the views. An *internal* rotation of either snapshot or current view could be excluded for ants by experiments on intra- and inter-ocular transfer of information [52]. But even if the assumption of retinotopically fixed views holds and body rotation can be excluded, the alignment could be accomplished by storing a set of snapshots taken in different orientations at the same location and activating them selectively according to the current compass bearing [6]. Ants could derive the compass information required for one of the above-mentioned alignment methods from the polarization pattern. In the robot experiments described here, we used the POL-compass to internally rotate the current view so that it is aligned with the snapshot.

3.6. Experiments

Robot experiments were performed in the same area where the experiments with *Cataglyphis* have been performed over many years, namely in southern Tunisia near the village Maharès (34.58°N, 10.50°E) in August–September 1997. The experimental field was a sandy salt-pan flat shared with other researchers performing experiments on *Cataglyphis* navigation. The experiments were performed from 6:00 AM to 11:00 AM and from 3:00 PM to 7:30 PM in order to avoid direct sunlight on the POL-sensors.

An example of a landmark configuration used in the robot experiments is shown in Fig. 16. The grid painted on the desert ground (40×40 m) is used to align the landmarks and the robot as well as to record the robot trajectories. Before each experiment, the robot was placed at the starting position — which corresponds to the "nest position" and is identical to the target position — and aligned with one of the axes of the grid. The experiments reported here were performed with an array of three landmarks arranged in an equilateral triangle whose sides were 3 m long. The landmarks were black cylinders with a height of 80 cm and a diameter of 30 cm. The starting position was situated on a symmetry line of the triangle in a distance of 1 m from one of the landmarks.

At the beginning of each experiment, a snapshot was taken at the starting position, processed as described in Section 3.4, rotated to a standard orientation using the angle obtained from the POL-compass (see Section 3.5), and stored in the form of black and white sectors. After taking the snapshot, the robot had to follow a certain direction (80° or 220°E) for some distance (2 or 4 m) corresponding to a short foraging journey of an ant (see Fig. 17, left). For this purpose, the robot could

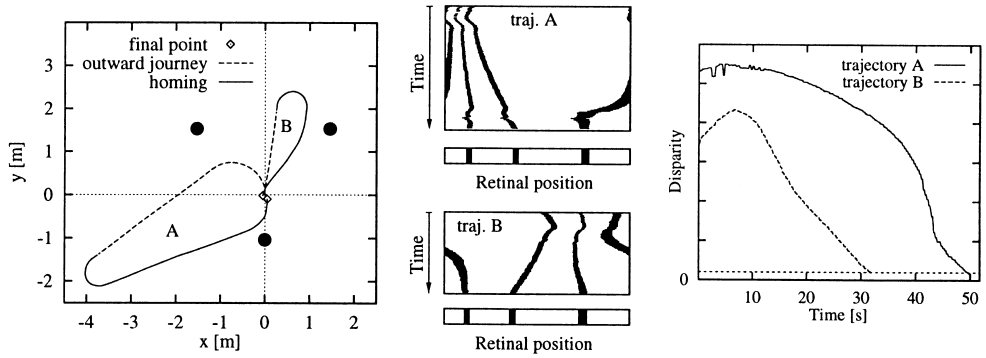

Fig. 17. Landmark navigation experiment with the robot. Left: Two typical trajectories of the robot in an array of three landmarks (filled circles). From the end of the preprogrammed outward journey (dashed) towards the target position (0,0) the robot is guided by landmark navigation (solid line). Center: Transitions of the rotated current view over time towards the snapshots (images underneath each trace) for the two trajectories. Right: Time course of disparity between snapshot and current view for the two trajectories. The dashed line depicts the threshold used for terminating the runs.

also have been displaced manually. At the end of the outward journey, control was handed over to the visual homing algorithm, which performed the extraction of a sectorized horizontal view from the camera image, the alignment of this view with the compass direction, and the computation of the home vector by matching the aligned view to the snapshot using the *PV model*. The resulting home vector was used to set the direction of movement. As soon as the disparity between snapshot and current view became lower than a threshold, the experiment was stopped (see Fig. 17, right).

3.7. Results

Eight runs (including trajectories A and B in Fig. 17) with different angles and distances of the outward journey were performed with this landmark configuration. The final parts of the eight trajectories are shown in Fig. 18 (left). The deviation between final and initial position of the camera axis was between 7 and 20 cm. This deviation does not reflect the ultimate precision of the homing mechanism, though, since the final positions are locations where the disparity reached the termination threshold, but not locations with zero disparity. In addition to the final positions, those points on the trajectories are marked, where the disparity reaches 1.5 times the threshold. For comparison, the right part of Fig. 18 visualizes the disparity for points in the vicinity of the snapshot location obtained in a simulation. Final positions as well as the points where 1.5× threshold was passed are in accordance with the disparity plot, since they can be thought to be located on roughly oval bands around the snapshot position, with all final positions (except one) on the inner band. Since the band of the final positions and the band of the intermediate (1.5 × threshold) positions are not overlapping, noise induced to the disparity from the sensor signal seems to be in a range below the termination threshold. Therefore it can be assumed that a reduction of the termination threshold will bring the final position closer to the snapshot location. This will require a specialized control algorithm for the final phase of the approach that will allow the robot to locate the target position more precisely.

3.8. Average Landmark Vector model

A common characteristic of the *snapshot model* and the *Proportional Vector model* presented in the previous sections is that they presuppose that there is an advantage in having a snapshot image stored in memory. However, in this section we will show this not to be the case. We propose a new model for visual landmark navigation in insects, called the *Average Landmark Vector* model (ALV model).

A description of the ALV model is given in Fig. 19. In the ALV model, each visual landmark feature (in this case sector centers) is associated with a unit vector pointing from the current position of the agent towards the landmark feature ($\overline{lan}_1^{\mathrm{cur}}$ and $\overline{lan}_2^{\mathrm{cur}}$ in Fig. 19). This vector is called a *landmark vector*. When the agent is at the target position (nest position) all

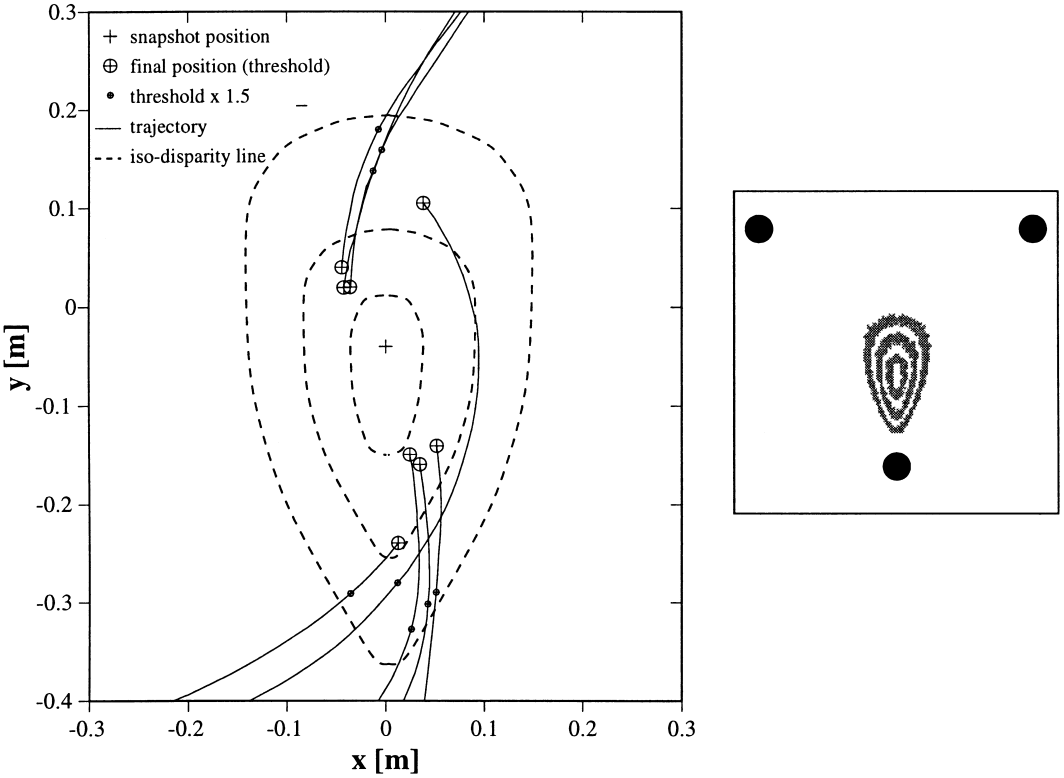

Fig. 18. Left: Final parts of 8 homing trajectories obtained in the landmark-navigation experiments (for the configuration in Fig. 17, including trajectories A and B shown there). The trajectories and final positions relate to the position of the camera axis. Circles (⊕) indicate the positions where the robot stopped as it reached the termination threshold. Small dots are placed on the trajectories at points, where the disparity dropped below a value of 1.5 times the termination threshold. Dashed lines depict iso-disparity curves. Right: Disparity in the vicinity of the snapshot location. Grey regions are restricted by equidistant iso-disparity curves; disparity decreases when approaching the snapshot position.

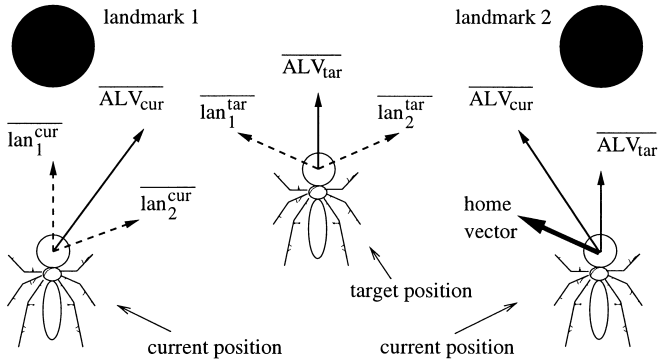

Fig. 19. Diagrammatic description of the ALV model. Each landmark feature i in the visual field — in this version of the model sector centers — is associated with a unit vector, called landmark vector ($\overline{lan}_i^{\text{cur}}$, $\overline{lan}_i^{\text{tar}}$). All landmark vectors are averaged to produce a single Average Landmark (AL) vector. At every point, the home vector — computed as the difference between the current AL vector and the vector at the target position ($\overline{ALV}_{\text{cur}} - \overline{ALV}_{\text{tar}}$) — gives the approximate direction to the target position.

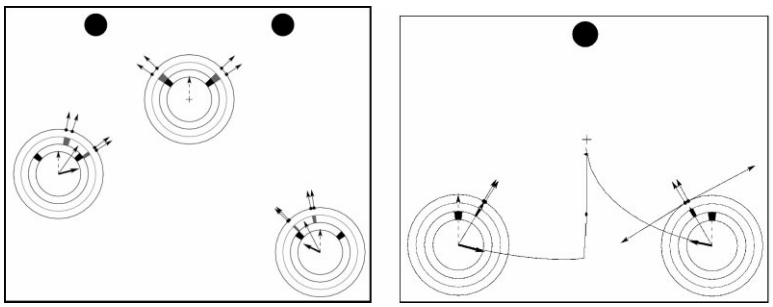

Fig. 20. Left: A version of the ALV model using edges as landmark features. Snapshot and current view are depicted as in Fig. 12. Vectors attached to the outer ring are landmark vectors contributing to the dashed Average Landmark vector (AL vector) computed for the snapshot position. The thin solid vector is the current AL vector, and the thick solid vector is the home vector. Right: To give more weight to size information contained in edges, vectors perpendicular to the landmark vectors can be added (shown on the outer circle). This is beneficial in environments with only a few landmarks as in this case. For two symmetrical positions, the home vectors were calculated without (left position), and with increased size weighting (right position). In the latter case, the agent moved in a more direct way towards the target position as shown by the corresponding trajectories.

landmark vectors are averaged to produce the *Average Landmark Vector* of the target position ($\overline{ALV_{\text{tar}}}$). This vector is stored. The same process is repeated again when the agent is in a different position, and the resulting AL vector ($\overline{ALV_{\text{cur}}}$) is "compared" to the stored AL vector at the target position by simply subtracting them ($\overline{ALV_{\text{cur}}} - \overline{ALV_{\text{tar}}}$). The "home vector" resulting from this subtraction gives the approximate direction to the target position. Note that also the ALV model requires compass information to align the vectors to a geocentered coordinate system.

The new model has an interesting property. It requires considerably less computational resources than any of the snapshot models that we saw until now: in terms of memory, all the agent has to "remember" is the AL vector in the target position (two values instead of an image); in terms of computation, determining the home vector is done by subtracting two AL vectors (vector subtraction instead of image matching procedure).

In the version of the ALV model presented in Fig. 19, sector centers are used as landmark features. As in the snapshot model, this requires determining the center of the landmark by considering its width. This is not necessary when edges are chosen as landmark features. In this case each landmark will generate two landmark vectors corresponding to the two edges that separate the landmark from the background. Fig. 20 (left) demonstrates this in an environment with two landmarks.

An emergent property of the ALV model with edges is that it implicitly incorporates information about the apparent size of landmarks. When a landmark is large, or the agent is very close to it, then the edges associated with the landmark will be further apart, and the corresponding AL vector (the average of the two edge vectors) will be very small. More specifically, the size of the AL vector is $2\cos(\frac{1}{2}\phi)$, where ϕ is the difference in bearing between the two edges. One way to increase the influence of the apparent landmark size on the final home vector is by artificially moving the edges apart. This can be done by adding a vector perpendicular to each edge vector as can be seen in Fig. 20 (right). The direction of this tangential vector depends on the polarity of an edge: it always points away from the landmark sector. Increasing the weighting of size information can be beneficial when the environment contains only a few landmarks, since in these cases it results in trajectories that lead more directly to the target. All models discussed in this paper including the ALV model have been shown in simulations to work reliably even in the presence of a large number of landmarks which are mutually occluding each other, as can be seen in Fig. 22. Occlusion also affects the apparent size of a landmark — regardless if computed explicitly by counting the number of pixels as in the snapshot model or implicitly using perpendicular vectors as suggested for the ALV model — which seems to be tolerated by these methods.

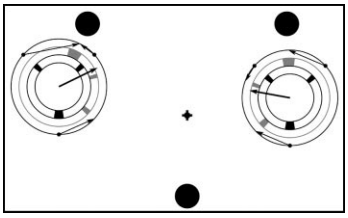

Fig. 21. The *difference vector model*, a version of the snapshot model, where the contributing vectors are computed as difference between unit vectors pointing towards the current view sectors and the snapshot sectors. Note that as in the snapshot model mismatches can occur, as visible in the diagram on the left. Views and vectors are depicted as described in Fig. 12. Only differences in bearing and only landmark sectors have been considered in this example.

Another important property of the ALV model is that under certain assumptions it generates home vectors that are identical to the home vectors generated by a version of the snapshot model, the so-called *difference vector model* (see Fig. 21). The difference vector model differs from the proportional vector model (in a version where only contributions from differences in bearing are considered) only in that the tangential vectors of the proportional vector model are replaced by "secant" vectors. These vectors are obtained by computing the difference between a unit vector pointing towards the current view sector and the corresponding sector in the snapshot. With the contributing vectors of the PV model they share the property that their length is a monotonic function of the angular difference between the matched sectors. The pairing of the sectors in the two views itself is performed in the same way as in the snapshot model.

Leaving out the increased influence of the landmark size, the Average Landmark vector is computed by simply summing (or averaging) all landmark vectors. This is described by

$$\overline{ALV_{\mathrm{tar}}} = \sum_{i=1}^{n} \overline{lan_i^{\mathrm{tar}}}, \qquad \overline{ALV_{\mathrm{cur}}} = \sum_{i=1}^{n} \overline{lan_i^{\mathrm{cur}}}, \tag{12}$$

where $\overline{ALV_{\mathrm{tar}}}, \overline{ALV_{\mathrm{cur}}}$ are the AL vectors at the target and current position, respectively, and $\overline{lan_i^{\mathrm{tar}}}, \overline{lan_i^{\mathrm{cur}}}$ are the landmark vectors at the target and current view positions, respectively. The home vector $\bar{h}$ is then derived by simply subtracting the two AL vectors:

$$\bar{h} = \overline{ALV_{\mathrm{cur}}} - \overline{ALV_{\mathrm{tar}}} = \sum_{i=1}^{n} \overline{lan_i^{\mathrm{cur}}} - \sum_{i=1}^{n} \overline{lan_i^{\mathrm{tar}}}$$
$$= \sum_{i=1}^{n} (\overline{lan_i^{\mathrm{cur}}} - \overline{lan_i^{\mathrm{tar}}}). \tag{13}$$

The same home vector $\bar{h}$ expressed by the last term is produced by difference vector model under the assumption that all landmarks are matched correctly, i.e., that $\overline{lan_i^{\mathrm{cur}}}$ and $\overline{lan_i^{\mathrm{tar}}}$ correspond to the same landmark i, which is guaranteed in the vicinity of the target location. This has an important implication. One interpretation of the above derivation is that the ALV model implicitly always establishes the correct matching. This is not the case with the other models, where depending on the position of the landmarks wrong matches can occur, as demonstrated in Fig. 21 for the difference vector model.

As was shown above, the ALV model is underpinned by its similarity with a variation of the snapshot model. In addition, the ALV model can be related to O'Keefes "centroid" model [36]. The centroid is the geometric center of mass of the landmark cues in the environment and can be computed as the average of difference vectors between the location of each landmark and the current position of the animal ("cue vectors"), provided that not only the bearings of the landmarks are known, but also their distance from the current position. The difference of the centroid vector computed for the current and the target location points from the current location exactly to the target. The ALV model can be derived from the centroid model by replacing the cue vectors with unit vectors that have the same orientation. This modification — which is comparable with the "equal distance assumption" underlying the homing scheme of Franz et al. [19] — is motivated by the fact that distance information is difficult to infer from the visual input. The tradeoff for this simplification is that the home vector has to be recomputed in the process of approaching the target location, whereas a single computation is sufficient in the centroid model.

Since the ALV model is closely related to the snapshot model and the extraction of the landmark cues from the images would be identical to the one described in Section 3.4, robot experiments using the ALV model in a similar setup would not provide additional insights at this point. We are currently inves-

tigating the application of the ALV model to indoor navigation in an office environment.

3.9. Discussion

Panoramic vision systems of the type that is used on the *Sahabot 2* — a camera with vertical optical axis facing a convex mirror — are a relatively simple technical solution to obtain a 360° view, compared to the use of multiple or rotating cameras. In the present study, this vision system emulates the full or almost full panoramic vision of insects [46,48], at least near the horizonal plane. Of course, the rectangular CCD sensor does not reproduce the distribution of ommatidia across the insect eye. In the eyes of *Cataglyphis bicolor*, maximum resolution (3°) is reached in those parts that look at the horizon, whereas the resolution is smaller in the upper and lower half of the eye (7°) [46]. In the technical vision system, resolution increases with increasing elevation of the pixel, but since only the horizonal portion is extracted from the polar view (see Fig. 15), this difference can be neglected. In the ant's eyes, the 3° resolution is approximately constant along the horizon, which is also the case for the technical counterpart with its 2° resolution. The resolution is limiting for the application range of the visual homing mechanism: as visible in Fig. 17, in a distance of approximately 4 m the landmarks occasionally disappear from the sectorized image.

It would have been interesting to use in our robot experiments those landmarks that are relevant to the ants in their natural habitat. However, whereas the ant's eyes are about 0.5 cm above ground, the mirror of the robot's visual system is at 27 cm height. Objects like small shrubs or stones that stick above the horizon of the ants and are therefore visible as a skyline against the bright sky as background may be far below the horizon for the robot. From its higher perspective, the frequent shrubs in the vicinity of the test field form a single band around the horizon which can not be separated into distinct landmarks. Therefore in the experiments black cylinders (as visible in Fig. 16) — of the same type also used in the ant experiments — served as landmarks.

In this environment and with the artificial landmarks, simple image processing steps are sufficient to link the camera image with the representation required for the *snapshot model*. For natural landmarks and environments without the advantage of a free horizon, more elaborate image processing routines will be necessary. The application of a functionally similar homing method to indoor navigation is currently under investigation.

The *snapshot model* was developed in order to reproduce the search behavior of bees [6]. Experiments have shown that the basic assumption — a snapshot image or information derived from this image is stored in the target position — also holds for ants [52,54]. However, details about the form of this snapshot (image, vector) and the specific method ants use to derive a home vector by matching snapshot and current view are not known so far. A systematic investigation would be necessary to find out if the matching procedure suggested in [6] or variations of it like the *proportional vector model* (Section 3.3), the *difference vector model* (Section 3.8), or the *ALV model* (Section 3.8) correctly predict the search behavior of ants in modified landmark setups. For this purpose, also the influence of the apparent height of the landmarks has to be considered, which has been neglected in our computer simulations and robot experiments up to now.

While it is difficult to compare the precision of the path-integration system of robot and ants (see Section 2.8), it is legitimate to compare the precision of the visual homing methods, since the visual processing is not affected by differences in size or propulsion. The precision achieved in the robot experiments — with the given threshold, the final distance to the snapshot location ranged between 7 and 20 cm (see Section 3.7) — is in a range where an ant should be able to see or smell the nest entrance. Demonstrating that this precision can be achieved in real-world experiments provides strong support for the snapshot thesis and the specific matching procedure.

The setups used in the robot experiments consisted of 2–4 landmarks. Computer simulations revealed that despite the parsimony of all models (and especially of the *ALV model*) even complex landmark situations can be mastered. Fig. 22 shows how three of the models (original *snapshot model*, *PV model*, and *ALV model*) perform in a complex environment with 27 landmarks of different sizes. All models perform surprisingly well, generating trajectories that end at the target position. The expectation was that because of the large number of landmarks many false matches would occur

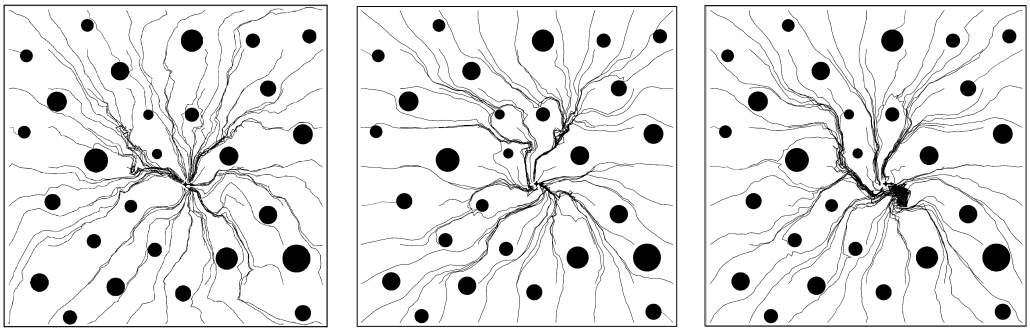

Fig. 22. The performance of the homing models in a complex environment with 27 landmarks. Left: *Snapshot model*. Center: *PV model*. Right: *ALV model*.

that cause local attractor points. Moreover, landmarks cover each other or fuse to larger landmarks from the perspective of the agent. Apparently, the home vectors produced by the models overall point to the target position, although differences in the trajectories are apparent. A detailed analysis of the differences between the behavior of the models would require an examination of a huge number of possible variations of the matching process which is beyond the scope of this paper. Note that under certain conditions some of these models exhibit an inherent obstacle-avoidance behavior without a dedicated obstacle-avoidance module.

4. Summary and conclusions

In the work presented here, we use an autonomous agent to study the navigation capabilities of insects, in particular of desert ants. We have developed a robot equipped with polarization vision and with a panoramic, 360° visual system. Three types of visual sensors employed on the robot — polarization-opponent units, ambient-light sensors, and a panoramic camera system — correspond to different parts and functions of the insect eye. The robot was used to investigate the path-integration capabilities of insects which are based on the polarized-light pattern of the sky and to test models of visual landmark navigation.

Sensors and early processing stages of the polarization-vision system of insects were reproduced in a polarized-light compass (POL-compass). We have analytically derived a *simultaneous* model for extracting compass information directly from the signals of the polarization-opponent units forming the POL-compass. In order to evaluate the precision of the POL-compass, a path-integration system employing a POL-compass was compared with a system that used proprioception only. Especially for long-range navigation, the POL-compass system proved superior to the proprioceptive system. These results are discussed in Section 2.8.

The robot was also used to test the *snapshot model* of visual landmark navigation. According to this model, an insect records a panoramic *snapshot* image of the surroundings at the target position. When it has to return to this location, it compares this snapshot with its current retinal image in order to determine a home vector. The high precision achieved in visual homing experiments with the robot demonstrates that the strategies assumed to underlie insect navigation actually work in a real-world environment. Moreover, these parsimonious strategies can also provide a guideline for robotics research. Particularly attractive for robotics applications is the novel, computationally cheap visual-homing method, the *Average Landmark Vector model*, which was developed in the course of this study. Section 3.9 provides a discussion of the results.

A fusion of the two navigation methods in a two-phase strategy — path-integration for long distances, visual homing in the vicinity of the target location — should be possible with small effort: As soon as the home vector obtained from the path-integration system is reduced to zero length, navigation switches to visual homing. As demonstrated, the path-integration system that derives orientation from the POL-compass guides the robot back to

the vicinity of the starting location with a precision of better than 1 m even after excursions of several hundred meters length. Visual piloting, on the other hand, has been shown to work in a range of up to 4 m with the given landmark setup, camera resolution, and image processing steps. The overlap of the two working ranges is indicating that a two-phase strategy can be successful. Recently, new experimental results concerning the interplay of path integration and landmark navigation have been presented [11]. Moreover, extensions of the visual homing method may increase its working range. The use of multiple snapshots was suggested in [7] and currently received support from ant experiments [24].

This case study illustrates the power of biorobotics that arises from the close relationship between engineering and biology: On the one hand, insights for innovative engineering designs can be found by analyzing the mechanisms employed by biological agents. On the other hand, biological hypotheses can be confirmed using real-world artifacts rather than simulations only, and new biological hypotheses and ideas for new animal experiments can be generated. Future work will concentrate on adapting the parsimonious strategies of insects for landmark navigation in indoor environments.

Acknowledgements

This work is supported by the Swiss National Science Foundation (grants 20-47257.96 and 2000-053915.98 to R. Pfeifer and 31-43317.95 to R. Wehner), the Swiss Federal Office for Education and Science (VIRGO TMR network), and the Human Frontier Science Program. The *Sahabot 2* was constructed in cooperation with Rosys AG, Hombrechtikon. Financial support by Rosys AG is kindly acknowledged. The authors thank Christian Gfeller (Rosys) for the fruitful cooperation. Many thanks to Marinus Maris and Hiroshi Kobayashi for their contributions in the construction of the robot.

References

[1] B. Barshan, H.F. Durrant-Whyte, Orientation estimate for mobile robots using gyroscopic information, in: Proceedings of the International Conference on Intelligent Robots and Systems (IROS 94), 1994, pp. 1867–1874.

[2] B. Barshan, H.F. Durrant-Whyte, Inertial navigation systems for mobile robots, Transactions on Robotics and Automation 11 (1995) 328–342.

[3] S.L. Bartlett, A. Hampapur, M.J. Huber, D. Kortenkamp, S. Moezzi, Vision for mobile robots, in: J. Sanz (Ed.), Advances in Image Processing and Machine Vision, Springer, Berlin, 1995, pp. 1–117.

[4] G.D. Bernard, R. Wehner, Functional similarities between polarization vision and color vision, Vision Research 17 (1977) 1019–1028.

[5] J. Borenstein, H.R. Everett, L. Feng, "Where am I?" Sensors and methods for mobile robot positioning, Technical Report, University of Michigan, 1996.

[6] B.A. Cartwright, T.S. Collett, Landmark learning in bees, Journal of Comparative Physiology 151 (1983) 521–543.

[7] B.A. Cartwright, T.S. Collett, Landmark maps for honeybees, Biological Cybernetics 57 (1987) 85–93.

[8] R. Cassinis, D. Grana, A. Rizzi, Using colour information in an omnidirectional perception system for autonomous robot localization, in: Proceedings of the First Euromicro Workshop on Advanced Mobile Robots (EUROBOT '96), Los Alamitos, CA, 9–11 October 1996, IEEE Computer Society Press, Silver Spring, MD, 1996.

[9] J.S. Chahl, M.V. Srinivasan, Visual computation of egomotion using an image interpolation technique, Biological Cybernetics 74 (1996) 405–411.

[10] J.S. Chahl, M.V. Srinivasan, Reflective surfaces for panoramic imaging, Applied Optics 36 (31) (1997) 8275–8285.

[11] M. Collett, T.S. Collett, S. Bisch, R. Wehner, Local and global vectors in desert ant navigation, Nature 394 (1998) 269–272.

[12] T.S. Collett, J. Baron, Biological compasses and the coordinate frame of landmark memories in honeybees, Nature 368 (1994) 137–140.

[13] T. Consi, J. Atema, J. Cho, C. Chryssostomidis, AUV guidance with chemical signals, in: Proceedings of the IEEE Symposium on Autonomous Underwater Vehicle Technology, Cambridge, MA, 19–20 July 1994, pp. 3–9.

[14] H. Cruse, Ch. Bartling, G. Cymbalyuk, J. Dean, M. Dreifert, A modular artificial neural net for controlling a six-legged walking system, Biological Cybernetics 72 (1995) 421–430.

[15] H. Cruse, Ch. Bartling, J. Dean, T. Kindermann, J. Schmitz, M. Schumm, H. Wagner, Coordination in a six-legged walking system. Simple solutions to complex problems by exploitation of physical properties, in: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, From Animals to Animats (SAB96), Cape Cod, MA, 1996, pp. 84–93.

[16] F.C. Dyer, J.A. Dickinson, Development of sun compensation by honey bees: How partially experienced bees estimate the sun's course, Proceedings of the National Academy of Science, USA 91 (1994) 4471–4474.

[17] N. Franceschini, J.M. Pichon, C. Blanes, From insect vision to robot vision, Philosophical Transactions of the Royal Society of London 337 (1992) 283–294.

[18] M.O. Franz, B. Schölkopf, H.A. Mallot, H.H. Bülthoff, Learning view graphs for robot navigation, Autonomous Robots 5 (1998) 111–125.

[19] M.O. Franz, B. Schölkopf, H.A. Mallot, H.H. Bülthoff, Where did I take that snapshot? Scene-based homing by image matching, Biological Cybernetics 79 (1998) 191–202.
[20] P. Gaussier, C. Joulain, S. Zrehen, J.P. Banquet, A. Revel, Visual navigation in an open environment without map, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-97) 1997, pp. 545–550.
[21] F. Grasso, T. Consi, D. Mountain, J. Atema, Locating odor sources in turbulence with a lobster inspired robot, in: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, From Animals to Animats (SAB96), Cape Cod, MA, 1996, pp. 104–112.
[22] G. Hartmann, R. Wehner, The ant's path integration system: A neural architecture, Biological Cybernetics 73 (6) (1995) 483–493.
[23] J. Hong, X. Tan, B. Pinette, R. Weiss, E.M. Riseman, Image-based homing, IEEE Control Systems (February 1992) 104–112.
[24] S.P.D. Judd, T.S. Collett, Multiple stored views and landmark guidance in ants, Nature 392 (1998) 710–714.
[25] D. Kortenkamp, Cognitive maps for mobile robots: A representation for mapping and navigation, PhD thesis, The University of Michigan, 1993.
[26] B. J. Kuipers, Y. T. Byun, A robust, qualitative method for robot spatial learning, in: Proceedings of the 7th National Conference on Artificial Intelligence (AAAI-88), St. Paul, MN, Morgan Kaufmann, Los Altos, CA, 1988, pp. 774–779.
[27] T. Labhart, Polarization-opponent interneurons in the insect visual system, Nature 331 (1988) 435–437.
[28] T. Labhart, Unpublished results, 1998.
[29] D. Lambrinos, M. Maris, H. Kobayashi, T. Labhart, R. Pfeifer, R. Wehner, An autonomous agent navigating with a polarized light compass, Adaptive Behavior 6 (1) (1997) 131–161.
[30] M. Lindauer, Angeborene und erlernte Komponenten in der Sonnenorientierung der Bienen, Zeitschrift für vergleichende Physiologie 42 (1959) 43–62.
[31] H. Lund, B. Webb, J. Hallam, A robot attracted to the cricket species gryllus bimaculatus, in: Proceedings of the Fourth European Conference on Artificial Life, 1997, pp. 131–140.
[32] L. Matthies, E. Gat, R. Harrison, Mars microrover navigation: Performance evaluation and enhancement, Technical Report 95-0096, Jet Propulsion Laboratory, NASA, Pasadena, CA, 1995.
[33] R. Möller, M. Maris, D. Lambrinos, A neural model of landmark navigation in insects, Neurocomputing 26–27 (1999) 801–808.
[34] M. Müller, R. Wehner, Path integration in desert ants, *Cataglyphis fortis*, Proceedings of the National Academy of Science Neurobiology 85 (1988) 5287–5290.
[35] R.C. Nelson, Visual homing using an associative memory, in: Proceedings of the Image Understanding Workshop, Morgan Kaufmann, Palo Alto, CA, May 1989, pp. 245–262.
[36] J. O'Keefe, The hippocampal cognitive map and navigational strategies, in: J. Paillard (Ed.), Brain and Space, Oxford University Press, Oxford, 1991, Chapter 16, pp. 273–295.
[37] R. Pfeifer, Building "fungus eaters": Design principles of autonomous agents, in: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, From Animals to Animats (SAB96), Cape Cod, MA, 1996, pp. 3–12.
[38] M.V. Srinivasan, An image-interpolation technique for the computation of optic flow and egomotion, Biological Cybernetics 71 (1994) 401–416.
[39] M.V. Srinivasan, J.S. Chahl, S.W. Zhang, Robot navigation by visual dead-reckoning: Inspiration from insects, International Journal of Pattern Recognition and Artificial Intelligence 11 (1) (1997) 35–47.
[40] H. Vitzthum, Der Zentralkomplex der Heuschrecke *Schistocerca gregaria*: Ein mögliches Zentrum des Polarisationssehsystems, PhD thesis, Universität Regensburg, 1997.
[41] B. Webb, Robotic experiments in cricket phonotaxis, in: Proceedings of the Third International Conference on Simulation of Adaptive Behavior, From Animals to Animats 3 (SAB94), MIT Press, Cambridge, MA, 1994, pp. 45–54.
[42] B. Webb, Biological orientation systems for mobile robots, in: Proceedings of the International Conference on Recent Advances in Mechatronics, 1995.
[43] B. Webb, Using robots to model animals: A cricket test, Robotics and Autonomous Systems 16 (2–4) (1995) 117–134.
[44] B. Webb, A cricket robot, Scientific American 275 (6) (1996) 94–99.
[45] B. Webb, J. Hallam, How to attract females: Further robotic experiments in cricket phonotaxis, in: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, From Animals to Animats 4 (SAB96), Cape Cod, MA, MIT Press, Cambridge, MA, 1996, pp. 75–83.
[46] R. Wehner, Himmelsnavigation bei Insekten. Neurophysiologie und Verhalten, Neujahrbl. Naturforsch. Ges. Zürich 184 (1982) 1–132.
[47] R. Wehner, Spatial organization of foraging behavior in individually searching desert ants, *Cataglyphis* (Sahara desert) and *Ocymyrmex* (Namib desert), in: J.M. Pasteels, J.L. Deneubourg (Eds.), From Individual to Collective Behavior in Social Insects, Birkhäuser, Basel, 1987, pp. 15–42.
[48] R. Wehner, Arthropods, in: F. Papi (Ed.), Animal Homing, Chapman and Hall Animal Behaviour Series, Chapman and Hall, London, 1992, Chapter 3, pp. 45–127.
[49] R. Wehner, The polarization-vision project: Championing organismic biology, in: K. Schildberger, N. Elsner (Eds.), Neural Basis of Behavioural Adaptations, Gustav Fischer Verlag, Stuttgart, 1994, pp. 103–143.
[50] R. Wehner, The ant's celestial compass system: Spectral and polarization channels, Orientation and Communication in Arthropods, Birkhäuser, Basel, 1997, pp. 145–185.
[51] R. Wehner, B. Lanfranconi, What do the ants know about the rotation of the sky?, Nature 293 (1981) 731–773.
[52] R. Wehner, B. Michel, P. Antonsen, Visual navigation in insects: Coupling of egocentric and geocentric information, Journal of Experimental Biology 199 (1996) 129–140.
[53] R. Wehner, M. Müller, How do ants acquire their celestial ephemeris function? Naturwissenschaften 331–333 (1993) 4471–4474.

[54] R. Wehner, F. Räber, Visual spatial memory in desert ants, *Cataglyphis bicolor* (Hymenoptera: Formicidae), Experientia 35 (1979) 1569–1571.
[55] R. Wehner, S. Wehner, Path integration in desert ants. Approaching a long-standing puzzle in insect navigation, Monitore Zool. Ital. (N.S.) 20 (1986) 309–331.
[56] R. Wehner, S. Wehner, Insect navigation: Use of maps or Ariadne's thread? Ethology, Ecology and Evolution 2 (1990) 27–48.
[57] Y. Yagi, M. Yachida, Real-time generation of environment map and obstacle avoidance using omnidirectional image sensor with conic mirror, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, May 1991, pp. 160–165.

Dimitrios Lambrinos received a B.S. and M.S. in Mathematics from the University of Ioannina, Greece, and an M.S. degree in Computer Science from the University of Sheffield. He received his Ph.D. with distinction in Natural Sciences at the AI Lab of the University of Zurich where he is also a Research Associate. His research interests include robot navigation, visually guide behavior and learning. He is currently studying biological models for insect navigation by using robots.

Ralf Möller studied Electrical Engineering and Computer Science and gained his Ph.D. in Engineering from the Technical University of Ilmenau, Germany, in 1996. He is currently a Postdoctoral Researcher at the AI Lab, Department of Computer Science, and at the Department of Zoology, University of Zurich. His research interests include visual robot navigation, biomimetic robots, neuromorphic systems, neuroinformatics, and parallel computation.

Thomas Labhart is a Lecturer at the Institute of Zoology of the University of Zurich. He has gained his Ph.D. in the Natural Sciences from the University of Zurich with a behavioral study on bee vision. During his two postdoctoral years at different universities in the USA, including CalTech and Yale he studied the electrophysiology of arthropoed visual systems. Working on the electrophysiological, histological and synthetic level his present research focuses on the neural mechanisms underlying the polarized-light compass of insects.

Rolf Pfeifer received his degree in Physics and Mathematics from the Swiss Federal Institute of Technology (ETH) in Zurich, Switzerland. After a number of years in the computer industry as a systems engineer, he worked at the Psychology Department of the University of Zurich in the area of simulation of cognitive processes. After his Ph.D. in Computer Science (ETH) he spent three years in the USA as a postdoc at Carnegic-Mellon University and Yale University in the areas of artificial intelligence and cognitive science. Since 1987 he has been a Professor of Computer Science at the University of Zurich. In 1990/91 he held the SWIFT AI Chair at the Free University of Brussels. He is Head of the Artificial Intelligence Laboratory at the Computer Science Department of the University of Zurich which has a strong focus on interdisciplinary research projects.

Rüdiger Wehner is Professor of Neurobiology and Behavior and Head of the Institute of Zoology of the University of Zurich, Switzerland. He obtained his Ph.D. at the University of Frankfurt, Germany, and did his postdoctoral research at both the University of Zurich and Yale University. He was A.D. White Professor at Cornell University, A. Forbes Professor at MBL, Woods Hole, Massachusetts, and received several named lectureships in the United States and Great Britain. He is permanent fellow of the Institute for Advanced Study in Berlin and foreign member of the American Philosophical Society. His research focuses on the neurobiology of insect navigation, especially on navigation in Saharan desert ants, *Cataglyphis*.