VISUAL HOMING IN ANALOG HARDWARE

Ralf Möller
Department of Computer Science, Artificial Intelligence Lab
and Department of Zoology, Neurobiology Lab
University of Zurich, Winterthurerstr. 190, CH-8057 Zurich, Switzerland

Insects of several species rely on visual landmarks for returning to important locations in their environment. The "average landmark vector model" is a parsimonious model which reproduces some aspects of the visual homing behavior of bees and ants. To gain insights in the structure and complexity of the neural apparatus that might underly the navigational capabilities of these animals, the average landmark vector model was implemented in analog hardware and used to control a mobile robot. The experiments demonstrate that the apparently complex task of visual homing might be realized by simple and mostly peripheral neural circuits in insect brains.

1 Introduction

The ability to return to important locations in their environment (like feeding sites or the nest) is essential particularly for social insects. Experiments with honey bees [1] and desert ants [2, 3] demonstrated that visual information about the landmarks in the environment of a target location is stored in the insect's brain and subsequently used to find back to that location. The "snapshot hypothesis" asserts that it is a relatively unprocessed *panoramic image* of the surroundings of the target location which is memorized [4, 2, 1]. How a home direction can be derived by establishing correspondences between this stored image — the "snapshot" — and the image perceived in another location was specified in the "snapshot model" and tested in computer simulations [1]. However, a mathematical simplification of the snapshot model revealed that the visual information identifying the target location can be compressed into a single *two-component vector* — the "average landmark (AL) vector" — without affecting the homing behavior [5, 6]. Moreover, in the "average landmark vector (ALV) model" determining the home direction does not require an image matching process but only a subtraction of the AL vectors of current and target location. Being parsimonious and underpinned by ethology, the ALV model is the basis of the present study.

Visual homing abilities have been investigated in numerous behavioral experiments, but up to now, nothing is known about the neural apparatus mediating these abilities in insect brains. On the contrary, the neural circuitry for visual motion detection is well studied, especially in the fly brain [7]. The reason for this discrepancy may be that while in the case of motion detection it is clear what to look for by anatomical or electrophysiological methods, there is a lack of functional hypotheses that can guide neurobiological investigations in the case of visual homing. It may therefore be worthwhile to start the quest for the neural apparatus underlying visual homing by constructing biologically plausible neural models. In the approach presented here, it was attempted to ensure plausibility by imposing the following constraints on the model:

(1) The model is derived from behavioral data: As shown elsewhere [6], the ALV model is a mathematical simplification of the snapshot model which reproduces the searching behavior of bees [1].

(2) The medium of the model shares basic computational principles with biological neural networks: The ALV model is implemented in analog hardware, where, as in neural systems, the computation is analog, asynchronous, and inherently parallel, and where the same operations are difficult or easy to realize. Analog circuits as the medium for biological models have been used before, especially for studying the elementary motion detectors of the fly [8, 9]. In contrast to the "neuromorphic engineering" approach [10, 11], this project uses discrete analog components instead of sub-threshold aVLSI circuits.

(3) The model is tested in the real world and not only in computer simulations: The analog circuit im-

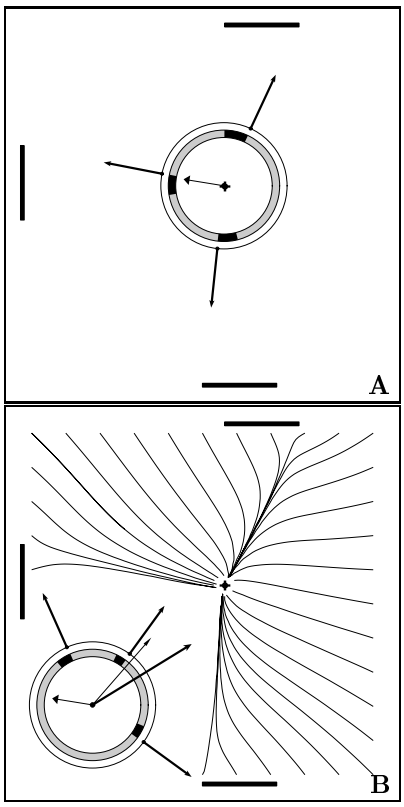

Figure 1. Average landmark vector model. **A:** In the target location (cross), the agent computes the AL vector (thin vector, wide head) by averaging all landmark vectors (thick vectors), which point towards the landmark cue, in this case the black-white edge of each landmark (black bars). The inner ring visualizes the image perceived by the agent. **B:** In another location (in the center of the ring diagram), the robot computes another AL vector (thin vector, small head). The difference between this AL vector and the stored AL vector (thin vector, wide head) gives the home vector (thick vector originating from the center). In addition, some trajectories obtained by following the continuously updated home vector are shown.

plementing the ALV model controls a mobile robot, and the homing precision of the robot was systematically investigated in a series of experiments.

This paper reviews the ALV model of visual homing (section 2), describes the analog circuit implementing this model and the robot hardware (section 3), presents results of homing experiments with the robot (section 4), and discusses correspondences between the analog implementation and neural architectures (section 5).

2 Average Landmark Vector Model

Figure 1 describes the homing mechanism of the ALV model. In the simple form presented here, the model assumes that the agent perceives a one-dimensional, horizonal panorama, which is shown in the inner ring. In each location, the agent can determine the corresponding AL vector by averaging unit-length "landmark" vectors pointing from the location of the agent towards a specific landmark cue, in this case towards black-white edges (clockwise). The AL vector obtained for the target location is stored (A). The home vector can then be determined by computing the difference between the AL vector of the current location and the stored AL vector (B). By following the continuously updated home vector, the agent will approach the target location. Note that in the figure all vectors relate to a fixed world coordinate system, or, put differently, a constant orientation of the agent is presumed. If the latter can not be guaranteed, the ALV model requires an external compass reference to align the AL vectors to the same coordinate system.

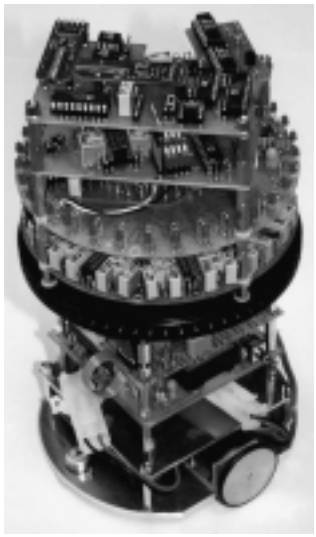

Figure 2. Robot (height 22 cm, diameter 12.8 cm, weight 1.3 kg). The black ring in the center contains the 32 photo diodes, the boards above implement the ALV model, the boards below belong to the motor control.

3 Circuit and Robot Hardware

3.1 Overview

Figure 3 gives an overview of the analog circuit implementing the ALV model. The implementation is based on discrete analog components (operational amplifiers, multipliers, analog switches). Visual input comes from a horizontal array of 32 photo diodes. The signals of the photo diodes are amplified. Edges of one polarity are detected by combining the signals of two neighboring pixel elements. Lateral inhibition between neighboring edge filters ensures that

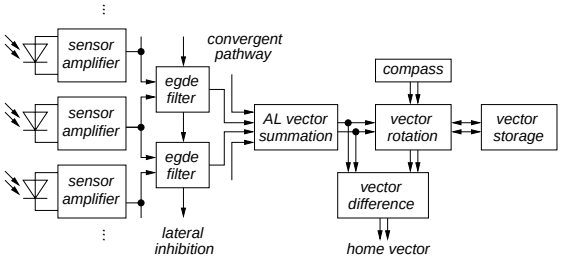

Figure 3. Overview of the analog implementation of the ALV model.

exactly one pixel per edge will become active. From the binary edge-filtered image the AL vector for the current location is determined. Using an electronic compass, the AL vector of the target location is rotated to world coordinates and stored when a switch is pressed. In the homing process, the stored vector is transformed back to robot coordinates and subtracted from the AL vector of the current location, which gives the home vector. The components of the home vector directly determine the speed of the two motors.

3.2 Sensor Array

The 32 Si photo diodes are mounted in a horizontal circular array (diameter: 12.8 cm, height above ground: 12 cm) with an inter-sensor angle of $\beta = 11.25°$; see figure 2. An aperture in front of each diode restricts the opening angle to approximately $\alpha = 8°$. A thread inside the diode's mounting hole reduces the influence of light from outside the desired opening angle.

3.3 Sensor Amplifiers and Edge Filters

The signals of the photo diodes are amplified using a standard circuit with one operational amplifier per diode. Amplification and offset of each amplifier are calibrated in a way that the amplified signals of all diodes are roughly identical when facing a white (0.4 V) or a black surface (0 V). The outputs of two neighboring amplifiers are compared in order to extract edges of one polarity that serve as landmark cue; see schematic in figure 4. One of the amplifier signals is directly fed to the comparator, the other is slightly attenuated and shifted with the threshold voltage $V_{ref} = 2$ V.

The diode connecting the output of one edge-filter with the negative input of one of the neighboring edge filters implements "lateral inhibition". This mechanism guarantees that the edge detection can operate

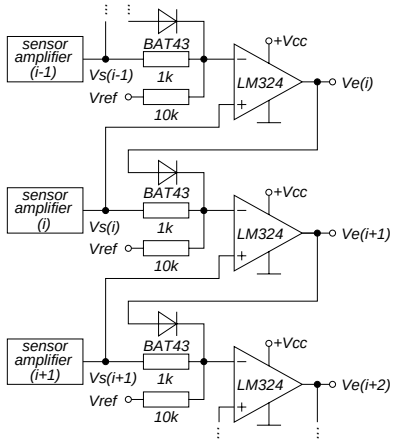

Figure 4. Edge-filter circuit.

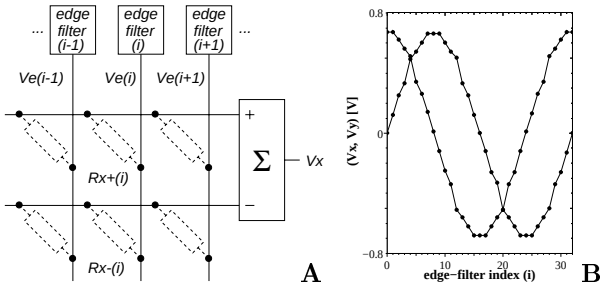

Figure 5. **A:** Circuit for the computation of the x-component of the AL vector; the circuit for the y-component is identical except for the values of the resistor array (dashed symbols). Landmark vectors are encoded in the resistors $R_{x+}(i)$ or $R_{x-}(i)$; values used are ∞, $390k$, $196k$, $150k$, $100k$, $91k$, $82k$, $75k$, $75k$ for the first quarter of a sine function. **B:** Voltages (V_x, V_y) produced at the output of the circuit in A, when only the edge filter with index i is activated.

with low thresholds, but nevertheless responds with only one active edge pixel to a sharp visual edge, as required by the ALV model. In cases where two neighboring edge detectors would be activated because one sensor is integrating over the edge, lateral inhibition switches off one of the edge filters.

3.4 AL Vector Computation

Because of the binary output of the edge filters, each landmark vector in robot coordinates can be encoded by just two resistors with appropriate values that connect to the input of adder circuits for the x- and y-component of the AL vector (figure 5, A). Since the landmark vectors have positive and negative components, the resistors are either connected with the positive or negative input line of the adder circuit. A landmark vector is considered in the summation

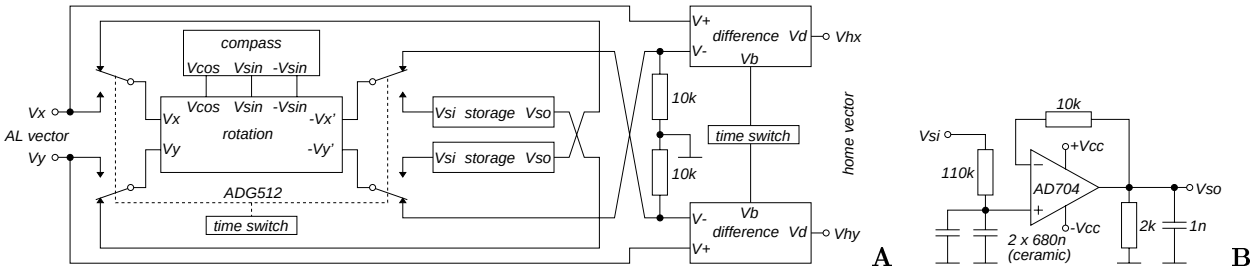

Figure 6. **A:** Circuit for rotation and storage of AL vectors and computation of the home vector. The four switches are shown in the position used for homing, in the other position the AL vector of the target location is transformed to world coordinates and stored. The blocks marked "difference" compute the difference $V_d = V_+ - V_-$. The time switch at V_b sets the home vector to zero and thereby stops the robot. **B:** Analog storage for one vector component.

if and only if the corresponding edge pixel is active. Figure 5 (B) shows the outputs V_x and V_y if the edge filters are activated individually. For simplicity, the analog circuit computes the *sum* of all landmark vectors, and not the average; this simplification does not affect the behavior as long as the number of landmarks is constant.

3.5 Rotation, Storage, Home Vector Computation

The circuit shown in figure 6 (A) rotates the AL vector of the target location to world coordinates, stores it, transforms the stored vector to robot coordinates and subtracts the result from the AL vector of the current location. The electronic compass used for this alignment procedure comprises two fluxgate magnetic field sensors. Since the strength of the magnetic field is not constant indoors, the compass vector is normalized to constant length. Multiplication of the AL vector with a rotation matrix derived from the compass vector is implemented with four four-quadrant multipliers (figure 7). The same circuit can be used for both the transformation from robot to world coordinates and from world to robot coordinates by transposing the matrix; transposition can be achieved by exchanging the components in the input and output vector (double 'X' crossover in figure 6, A). The transformed AL vector of the target location is stored in the circuit depicted in figure 6 (B). Analog switches with low leakage currents (< 100 pA) and amplifiers with low input bias currents (< 150 pA) have been used to achieve sufficiently long storage times in the range of 15 min (voltage change approx. 50 μV/sec, vector components go up to 3 V at $V_{cc} = \pm$ 5 V).

3.6 Motor Unit

The components of the home vector can directly control the speed of the two motors that are independently driving the two wheels of the robot. If the vector (1,1) of the robot coordinate system is pointing towards the front of the robot and the value on the coordinate axis pointing towards the left side is determining the speed of the right wheel and vice versa, the robot will turn around if facing away from the home vector direction and afterwards follow the vector. The motor unit uses open-loop control with pulse-width modulation and dead-band compensation.

4 Robot Experiments

Experiments were performed in an 1 m $\times$ 1 m arena with white walls (30 cm high) and floor. Black pieces of paper (21 cm$\times$29 cm) attached upright to the walls served as landmarks. Light came from the ceiling lamps of the room. Figure 8 shows the AL vector voltages for one and for three landmarks. The two voltages have been measured (using multimeters with computer interface) while the robot was placed at 64 locations on a grid and aligned with the world coordinate system.

For a visualization of the home vectors (figure 9), the robot was rotated by 60° with respect to the world coordinates and placed at the target location. An AL vector was computed and stored while a switch on the robot was pressed. Then the robot was aligned with the world coordinates (rotation 0°) and positioned at 64 grid points. The orientation was changed in order to demonstrate the operation of the rotation circuit. Note that the home vectors become shorter in the vicinity of the target location which will automatically slow down the robot.

In the experiment shown in figure 10 (A), the robot was first placed at the target location facing the upper wall, where the AL vector was computed and stored. Then it was moved in the same orientation to 11 different starting points. After the time of a motor blocking elapsed (V_b in figure 6, A), the robot started

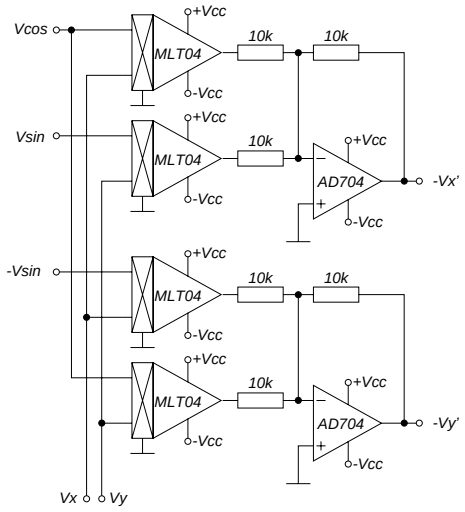

Figure 7. Circuit for vector rotation.

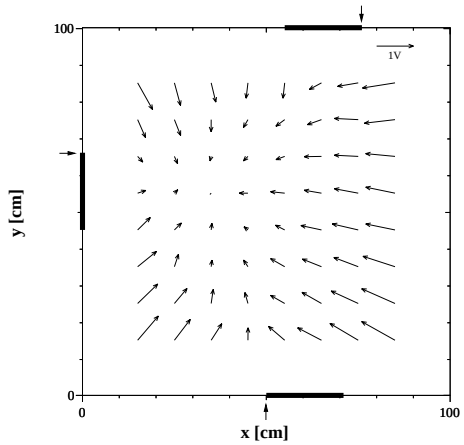

Figure 8. AL vectors for a setup with three landmarks. Black landmarks are depicted by bars. Detectable edges are marked with arrows. The vector scaling is shown in the upper right corner.

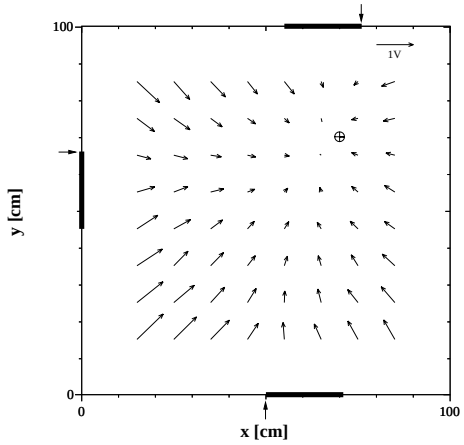

Figure 9. Home vectors for the target location marked by a cross-circle in a setup with three landmarks. See figure 8 for a description.

to move. A pen in the center of the robot drew the trajectories on paper covering the floor; the trajectories were digitized from a photo of the paper. Note the change of direction in the upper four trajectories, which are resulting from the type of motor control used (see section 3.6). In the first part of the curved trace in the bottom left the robot probably did not detect the landmark on the left or on the bottom wall, since they appear under a small angle. All trajectories end in the vicinity of the target location. The mean final distance from the target location in this case is 6.4 cm. This deviation can be fully explained by the fact that the small visual resolution will result in relatively large regions around the target location where the view (and thus the AL vector) does not change. This region can be constructed as shown in figure 10 (B): The sectors of the active edge pixels are attached to the detectable edges of the landmark; their cross section gives the iso-view region. The size of the iso-view region for this setup is in the same range as the deviation observed in the experiment.

In order to systematically investigate the precision of homing, 99 trajectories have been registered with different target locations for the setup with three landmarks, and 28 trajectories in a setup with an additional fourth landmark on the right wall. The average deviation of the final points of the trajectories from the target point was 68 mm for three landmarks. As in the experiment in figure 10, this deviation can be fully explained by the limited visual acuity. For four landmarks, the robot deviates on average by 48 mm. This significant improvement is due to the fact that an additional landmark will usually reduce the size of the iso-view region.

5 Discussion

As was stated above, the precision of homing observed in the robot experiments could be shown to be close to the optimum achievable with the given visual acuity, thus the third constraint imposed on the model — real-world suitability — is at least partly satisfied. Limiting for the validity of the real-world experiments is the simplified laboratory setup with its strong visual contrasts and the simple mechanism of feature detection which relies on these properties of the environment. The analog circuit may define the “lower

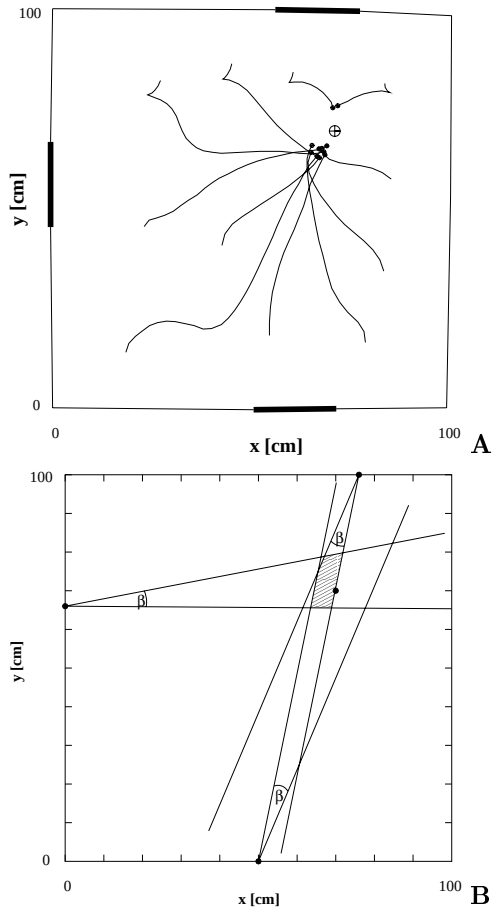

Figure 10. **A:** Trajectories of the robot while approaching the target location (cross-circle). Small circles are the positions where the robot movement stopped. **B:** Iso-view region (hatched) around the target location (dot) for the given sensor resolution. Dots on the margin depict detectable edges.

limit of complexity" that can be assumed to underly the visual homing abilities of insects. With only 91 operational amplifiers, 12 multipliers (eight of them implementing the normalization of the compass vector), and four analog switches, the circuit is surprisingly simple for the apparently complex navigation task it fulfills. An interesting observation is that the majority, namely 64 of the 91 operational amplifiers are part of the retinotopically organized feature detection circuit, although the latter is extremely simple. Retinotopical organization and small-field neurons like the local feature detectors in the model are typical for the peripheral parts of the optical nervous system of insects (lamina and medulla, respectively) [12, 13], where, according to this model, most of the functionality of visual homing seems to be concentrated.

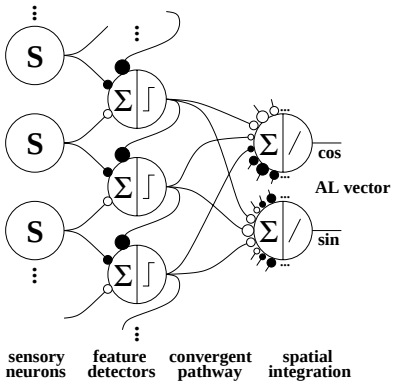

Figure 11. Equivalent neural model for feature detection and AL vector computation. Large circles depict neurons, with spatial summation at the input and a step function or a linear function at the output. Small circles symbolize synapses, with black synapses being inhibitory and white synapses excitatory. The size of the synapse symbols corresponds to their synaptic efficacy.

Figure 11 shows that the retinotopically organized part (compare figure 4) and the convergent pathway (compare figure 5, A) can be translated into an equivalent neural model by a 1:1 projection of operational amplifiers onto graded-response model neurons. Feature-detecting neurons with steep output functions receive antagonistic input from two neighboring sensory neurons, and strong inhibitory input from neighboring detectors (lateral inhibition). The outputs of all feature-detecting neurons converge on (in the simplest case) two neurons with linear output function. In the spatial integration, the output of each feature detector is weighted according to its field of view relative to the robot-bound coordinate system. For this specific encoding it may be sufficient to assume a certain spatial arrangement of the cells and transmission losses depending on the length of the neural processes [14]. "Large-field" neurons like the ones mediating the spatial integration can be found in higher stages of the visual system of insects (medulla, lobula plate) [12].

For the analog rotation circuit with its precision multipliers it is not very likely to find direct correspondences in a biological system. In general, the question if and how insects relate landmark information to some external reference is still largely unresolved [15, 3]. Biologically plausible candidate mechanisms for an internal rotation of image or vector information are self-stabilizing neural circuits, which have been suggested for the path-integration system of insects [16].

6 Conclusion

The model of insect visual homing presented here is constrained by behavioral data, the medium of the model, and real-world suitability. The ALV model — a mathematical simplification of the snapshot model that reproduces the homing behavior of bees — was implemented in analog hardware and successfully tested on a mobile robot. Parts of the analog circuit can be translated into an equivalent neural network model with less than 100 neurons. This model may define the lower limit of complexity of the neural apparatus underlying the visual homing abilities of insects. The structure of the network indicates that the largest part of the functionality of visual homing might be implemented in the peripheral parts of the visual system of insects. Future work will focus on eliminating the limitations of the feature detection circuit and the small visual resolution.

Acknowledgments

The project is supported by the Swiss National Science Foundation (grants 2000-053915.98 and 5002-044889 to Rolf Pfeifer), the Swiss Federal Office for Education and Science (VIRGO TMR network), the Human Frontier Science Program, and a personal grant from the Kommission zur Förderung des akademischen Nachwuchses der Universität Zürich. Many thanks to Bernhard Schmid and colleagues, Koh Hosoda, Hiroshi Kobayashi, Dimitrios Lambrinos, Marinus Maris, and Peter Paschke.

References

[1] B. A. Cartwright and T. S. Collett. Landmark learning in bees. *J. Comp. Physiol. A*, 151:521–543, 1983.

[2] R. Wehner and F. Räber. Visual spatial memory in desert ants, *Cataglyphis bicolor* (Hymenoptera: Formicidae). *Experientia*, 35:1569–1571, 1979.

[3] R. Wehner, B. Michel, and P. Antonsen. Visual navigation in insects: coupling of egocentric and geocentric information. *J. Exp. Biol.*, 199:129–140, 1996.

[4] T. S. Collett and M. F. Land. Visual spatial memory in a hoverfly. *J. Comp. Physiol.*, 100:59–84, 1975.

[5] D. Lambrinos, R. Möller, R. Pfeifer, and R. Wehner. Landmark navigation without snapshots: the average landmark vector model. In N. Elsner and R. Wehner, editors, *Proc. Neurobiol. Conf. Göttingen*, page 30a, Georg Thieme Verlag, Stuttgart, 1998.

[6] D. Lambrinos, R. Möller, T. Labhart, R. Pfeifer, and R. Wehner. A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems, special issue: Biomimetic Robots* (to appear)

[7] M. Egelhaaf and A. Borst. A look into the cockpit of the fly: Visual orientation, algorithms, and identified neurons. *J. Neurosci.*, 13:4563–4574, 1993.

[8] N. Franceschini, J. M. Pichon, and C. Blanes. From insect vision to robot vision. *Phil. Trans. R. Soc. Lond. B*, 337:283–294, 1992.

[9] R. R. Harrison and C. Koch. An analog VLSI model of the fly elementary motion detector. In M.I. Jordan, M.J. Kearns, and S.A. Solla, editors, *Advances in Neural Information Processing Systems*, pages 880–886, MIT Press, Cambrigde, MA, 1998.

[10] C. Mead. *Analog VLSI and Neural Systems.* Addison-Wesley, Reading, MA, 1989.

[11] R. Douglas, M. Mahowald, and C. Mead. Neuromorphic analogue VLSI. *Ann. Rev. Neurosci.*, 18:255–281, 1995.

[12] N. J. Strausfeld and D. R. Nässel. Neuroarchitectures serving compound eyes. In H. Autrum, editor, *Handbook of Sensory Physiology, VII/6B*, pages 1–132. Springer, Berlin, 1981.

[13] S. R. Shaw. Early visual processing in insects. *J. Exp. Biol.*, 112:225–251, 1984.

[14] R. Möller, M. Maris, and D. Lambrinos. A neural model of landmark navigation in insects. *Neurocomputing*, 26:801–808, 1999.

[15] T. S. Collett and J. Baron. Biological compasses and the coordinate frame of landmark memories in honeybees. *Nature*, 368:137–140, March 10, 1994.

[16] G. Hartmann and R. Wehner. The ant's path integration system: a neural architecture. *Biol. Cybern.*, 73:483–497, 1995.