

Biomimetic robot navigation [☆]

Matthias O. Franz ^{*}, Hanspeter A. Mallot

Max-Planck-Institut für biologische Kybernetik, Spemannstr. 38, 72076 Tübingen, Germany

Abstract

In the past decade, a large number of robots has been built that explicitly implement biological navigation behaviours. We review these biomimetic approaches using a framework that allows for a common description of biological and technical navigation behaviour. The review shows that biomimetic systems make significant contributions to two fields of research: First, they provide a real world test of models of biological navigation behaviour; second, they make new navigation mechanisms available for technical applications, most notably in the field of indoor robot navigation. While simpler insect navigation behaviours have been implemented quite successfully, the more complicated way-finding capabilities of vertebrates still pose a challenge to current systems. ©2000 Elsevier Science B.V. All rights reserved.

Keywords: Robot navigation; Spatial behaviour; Cognitive map; Topological map; Landmark

1. Introduction

In its original sense, the term navigation applies to the process of directing a ship to its destination.¹ This process consists of three repeating steps [69]: (a) the navigator determines the ship's position on a chart as accurately as possible; (b) on the chart, he relates his position to the destination, reference points and possible hazards; (c) based on this information, he sets the new course of the vessel.

[☆] This paper builds on a tutorial given at the Fifth International Conference on Simulation of Adaptive Behavior (SAB'98). We thank Tony Vladusich for helpful comments on the paper. Financial support was provided by the Max-Planck-Gesellschaft and the Human Frontier Science Program.

^{*} Corresponding author. Current address: DaimlerChrysler AG, Research and Technology, P.O. Box 2360, D-89013, Ulm, Germany. Tel.: +49-731-505-2120, fax: +49-731-505-4105. E-mail address: matthias.franz@daimlerchrysler.com (M.O. Franz).

¹ From Latin: *navis*, ship; *agere*, to drive.

The nautical practice of navigation has entered almost unchanged into the domain of robotics. For instance, Levitt and Lawton [37] define navigation as a process answering the following three questions: (a) “Where am I?”; (b) “Where are other places with respect to me?”; (c) “How do I get to other places from here?”. Applied to robot navigation, this means that the robot's sensory inputs are used to update a single global representation of the environment, from which motor actions are derived by an elaborate inference procedure. This view of navigation has not only been adopted in standard robotics textbooks (e.g., [44]), but also forms the basis for many robot navigation systems (see, e.g., the list in [33]).

None of these systems has yet reached the flexibility and navigation performance of bees or ants, let alone migrating birds or fish. This has motivated robotics researchers to look for biological navigation mechanisms that can be implemented on an autonomous mobile robot. The notion of navigation described above turned out to be of little help in biomimetic naviga-

tion systems. Indeed, ethological research in the past decades has shown that many animals (including humans) are able to navigate without answering all of Levitt and Lawton's questions, or even without answering any of them. "Where am I?" is not necessarily the first question to ask. For navigating animals, the most important question is "How do I reach the goal?"; this does not always require knowledge of the starting position. A necessary prerequisite for a biomimetic approach therefore is a broader idea of navigation, since otherwise most biological mechanisms are excluded from consideration at the outset.

In the present paper, we provide a framework for describing navigation phenomena which builds on recent discussions in the biological literature. Within this framework, we give an overview over the efforts that have been undertaken so far in the field of biomimetic robot navigation. We consider an approach as biomimetic if the authors try to implement a mechanism described in the biological literature, and explicitly refer to the biological inspiration of their approach. Out of the large number of biomimetic approaches to navigation, we have chosen a subgroup for this review, namely those that (i) were implemented on a real mobile robot, and (ii) mimic actually observed navigation behaviour of animals or humans. To keep our review focused, we did not include biomimetic systems based on more general biological ideas that are not specific to navigation, such as robots using artificial neural networks or biological learning mechanisms.

In the next section, we try to give a useful definition of navigation and discuss its relation to other forms of spatial behaviour. We present a hierarchical classification of navigation phenomena which is our guideline for discussing the biomimetic robot navigation systems in Sections 3 and 4. We conclude our review with a discussion of the technical and biological relevance of biomimetic approaches to navigation.

2. Navigation

2.1. Basic concepts

Since the classical notion of navigation captures only part of the observed biological navigation

phenomena, we adopt a slightly modified version of Gallistel's [19] definition²

"Navigation is the process of determining and maintaining a course or trajectory to a goal location."

The minimal capabilities for navigation are thus to move in space, and to determine whether or not the goal has been found. To that end, the sensory features which identify the goal have to be stored in some form of long-term memory. In contrast to the definition of Levitt and Lawton [37], this notion of navigation does not imply that the current location must be recognized, nor that a map-like representation must be used to find the goal.

The reference to a goal location distinguishes navigation from other forms of spatial behaviour such as exploration, foraging, obstacle avoidance, body orientation or course stabilization. In particular, "taxis" is not always a navigation mechanism. Originally, it refers to an active body orientation into a direction with respect to a stimulus field [31]. This implies that taxis is certainly involved in most navigation behaviours (together with obstacle avoidance and course stabilization), but taxis does not necessarily include the essential capabilities of locomotion and goal recognition.

The forms of spatial behaviour captured by the above definition fall into two fundamentally different groups: local navigation and way-finding. *Local navigation* requires the recognition of only one location, namely the goal. The agent chooses its actions on the basis of current sensory or internal information, without the need of representing any objects or places outside the current sensory horizon [76]. Local navigation methods have also been called "tactics" [82] or local control strategies [33]. *Way-finding* involves the recognition of several places, and the representation of relations between places which may be outside the current range of perception [59]. Way-finding relies on local navigation skills to move from one place to another, but it allows the agent to find places that could not be found by local navigation alone.

The distinction between local navigation and way-finding is not necessarily associated with the

² Gallistel restricts this definition in the following sentences to navigation based on self-localization and map-like representations which, again, is too narrow for our purpose.

Table 1
The navigation hierarchy^a

	Behavioural prerequisite	Navigation competence
Search	Goal recognition	Finding the goal without active goal orientation
Direction-following	Align course with local direction	Finding the goal from one direction
Aiming	Keep goal in front	Finding a salient goal from a catchment area
Guidance	Attain spatial relation to the surrounding objects	Finding a goal defined by its relation to the surroundings
Recognition-triggered response	Association sensory pattern–action	Following fixed routes
Topological navigation	Route integration, route planning	Flexible concatenation of route segments
Survey navigation	Embedding into a common reference frame	Finding paths over novel terrain

^aNavigation behaviours are classified according to navigation competences that can be tested experimentally. The upper half of the table contains local navigation behaviours, the lower half way-finding behaviours.

scale of the environment in which the agent moves: By simply following a locally measured compass direction, for instance, one can find places that are thousands of kilometers away, whereas finding a particular room in a campus building already requires considerable way-finding skills.

The large range of existing biological and technical navigation behaviour calls for a further classification of the extant phenomena. There is a number of classification schemes in the literature, each based on different criteria [10,31,56,58,76,82]. Here, we use a hierarchy of navigation competences that can be tested experimentally. This allows for a common classification of both biological and technical navigation behaviour which is especially suitable for the purpose of this text.

2.2. The navigation hierarchy

The navigation hierarchy is based on the classification scheme of Trullier et al. [76] which we have modified and extended. In this hierarchy, navigation behaviour is classified according to the complexity of the task that can be performed (cf. Table 1). Thus, each level is characterized by a certain navigation competence which can be tested experimentally. Local navigation behaviours are divided into four levels: *search*, *direction-following*, *aiming* and *guidance*; way-finding behaviours into three levels: *recognition-triggered response*, *topological* and *survey navigation*. An agent at a given level has all the capabilities of the lower levels of its respective group, but a way-finding agent does not necessarily have all

local navigation skills. In computer science, a similar hierarchy has been presented by Kuipers [32].

2.2.1. Search

An agent navigating by search alone shows no active orientation towards the goal (Fig. 1(a)). The goal can only be found by chance if the agent hits it while moving around. This simplest form of navigation requires only the basic competences of locomotion and goal detection, without the need of any type of spatial representation. Search requires a large amount of time compared to other navigation methods. Since the direction towards the goal needs not to be known, search can serve as a backup strategy when the agent cannot find its goal.

2.2.2. Direction-following and path integration

For this type of navigation behaviour, the agent must be able to align its course with a locally available direction to find the goal. The goal itself need not to be perceivable during approach. An example of this behaviour would be a ship setting its course along a fixed compass direction leading to the goal. Direction information may be extracted either from *allothetic* (based on an external reference) sources such as the magnetic field, celestial cues or odour trails, or from *idiothetic* (based on an internal reference) sources such as an inertial compass or proprioceptive signals. Whereas direction following is more effective than search, it allows the agent to find the goal only when it moves on the trail defined by the direction (Fig. 1(b)). If the agent is displaced from the trail, it will miss the goal. Thus, this method is not very tolerant with respect




Fig. 1. Local navigation behaviours: (a) A searching agent shows no active goal orientation. (b) Direction-following allows an agent to find the goal along a trail. (c) Aiming needs a salient sensory cue at the goal. (d) Guidance enables an agent to find a goal defined by the spatial relationship to the surrounding objects.

to inaccurate directional information and alignment. Deviations from the trail may accumulate so that the goal is missed.

If, in addition, the distance to the goal is known, direction-following becomes more effective, since the agent can switch to another strategy, e.g., search, after the goal has been passed undetected. Again, there are various allo- and idiothetic sources for distance information such as step number, energy consumption, optic flow, etc. A strategy for acquiring both the direction and the distance to the goal is called *path integration* in biology, or *odometry* in robotics: On the outward journey, the agent continuously integrates every change of direction and the covered distance, thus maintaining a continuous estimate of the current egocentric position of the starting point. This computation does not require a global representation of all places visited on the outward journey. It can be done by using only locally available information on the current movement of the animal. An agent using path integration is not confined to a fixed trail with respect to the ground, it is able to return to the start position from any point of the trajectory, as long as the integration process is not interrupted.

2.2.3. Aiming

An agent aiming at a goal has to orient its body axis such that the goal is in front of it. The goal must be associated with some salient cue, be it olfactory, auditory, or other, which is always perceivable during approach (Fig. 1(c)). In the visual domain, this cue is often called a *beacon* [36]. A ship, for instance, can navigate by aiming at a widely visible lighthouse which functions as a beacon. In contrast to direction-following, the goal can be approached from

various directions without the danger of cumulative error. The area in which the salient cue can be perceived defines a “catchment area” around the goal. However, not every location can serve as a goal since it must be marked by a salient cue.

2.2.4. Guidance

When the goal is not marked by a salient cue, the agent can be guided by the spatial configuration of the surrounding objects. *Guidance* is a process whereby “a certain egocentric relationship” with respect to “a particular landmark or object” is maintained [54]. Moving so as to attain this spatial relationship to the configuration of the surrounding objects leads the agent to the location where the relationship was memorized [76]. Thus, the spatial information required is not only a single direction or location, but the spatial relationship between the current location, the goal, and the currently perceptible environment (Fig. 1(d)). Guidance is still a local navigation method since it requires only the processing of current sensory or internal information, without the need of representing anything outside the current sensory horizon. A nautical example of guidance would be a ship that tries to reach a fixed position between several islands.

2.2.5. Recognition-triggered response

All previous navigation methods are local since they lead to a single location with the help of locally available information. *Recognition-triggered responses* connect two locations by means of a local navigation method. They may be formally denoted by a pair $\langle \text{starting location, local navigation method} \rangle$ and thus involve the recognition not only of the goal,




Fig. 2. Way-finding behaviours: (a) A route consisting of recognition-triggered responses cannot be adapted to new obstacles such as a closed door (S start, G goal). (b) Topological navigation allows a flexible concatenation of routes. Route integration occurs at places where routes overlap (circles with double rim). (c) Survey navigation enables an agent to find shortcuts over previously unvisited terrain (dotted line).

but also of a starting location. The recognition of the starting location triggers the activation of a local navigation method leading to the goal. In this context, a location is defined as a certain sensory situation in which a particular local navigation method is selected. Thus, an association between a sensory pattern defining the start and an action has to be learned. There is no planning of a sequence of subsequent movements, only the selection of the very next action. Thus, the agent responds in an inflexible manner to the current situation.

Recognition-triggered responses can serve as an elementary navigation step for building *routes*. Routes are sequences of recognition-triggered responses, in which the attainment of the goal of one step triggers the start of the next step (cf. Fig. 2(a)). The local navigation method can be different in each step according to the local environment. A route may connect locations that cannot be reached by local navigation alone. Still there is no planning involved, as knowledge is limited to the next action to perform. If one route segment is blocked, e.g., by an obstacle, the agent has to resort to a search strategy until it reaches a known place again.

2.2.6. Topological navigation

An agent using recognition-triggered responses is confined to using always the same sequences of locations. Routes are generated independently of each other and each goal needs its own route. Navigation is more adaptive if the spatial representation is goal-independent, i.e., if the same representation can be used for multiple goals. To this end, the agent must have the basic competence of detecting whether

two routes pass through the same place. Two possibly different sensory configurations associated with the different routes leading through the same place have to be merged by *route integration*. A collection of integrated routes thus becomes a topological representation of the environment. This can be expressed mathematically as a graph, where vertices represent places and edges represent a local navigation method connecting two vertices.

Any vertex can become the start or the goal of a route, so that, in the case of obstacles, alternative intersecting routes may be found (cf. Fig. 2(b)). The fact that alternative routes may lead to one goal requires *planning* abilities which generate routes from the graph. Planning and route integration are the capabilities required for *topological navigation*. The resulting routes are concatenations of sub-sequences from previously visited routes. As a consequence, an agent relying on topological navigation cannot generate novel routes over unvisited terrain.

2.2.7. Survey navigation

Whereas for topological navigation different routes have to be integrated locally, *survey navigation* requires the *embedding* of all known places and of their spatial relations into a common frame of reference. In this process, the spatial representation must be manipulated and accessible as a whole, so that the spatial relationship between any two of the represented places can be inferred. In contrast, topological navigation needs only the spatial relations between connected places. An agent using survey navigation is able to find novel paths over unknown terrain, since the embedding of the current location into the common frame

of reference allows the agent to infer its spatial relation to the known places. Examples include finding of shortcuts in unknown terrain between unconnected routes (cf. Fig. 2(c)), or detours around obstacles in unknown terrain.

2.3. Biomimetic robot navigation

Examples of local navigation mechanisms, recognition-triggered responses and topological navigation can be found throughout the animal kingdom, whereas survey navigation may be limited to vertebrates. Generally, each level of the navigation hierarchy requires new skills on top of the lower level skills. This could also indicate the direction taken during evolution, since new behavioural capabilities are usually built on simpler pre-existing mechanisms. The hierarchy of competences and their underlying mechanisms in a biomimetic robot should thus reflect an “evolutionary scaling” as discussed by Mallot [41].

The approach of classical robotics to navigation is most reminiscent to survey navigation since spatial knowledge is represented in a common global map. This contrasts with the above considerations in which survey navigation is the very last stage of the evolutionary development. Biomimetic approaches are therefore constructed in a bottom-up manner: higher navigation abilities are used on top of simple, but reliable mechanisms. Sometimes these simpler mechanisms turn out to be sufficient for a given task, so that the higher levels need not to be implemented. The distinction between biomimetic and purely technical approaches, however, is not always clearly cut. Often technical solutions correspond very closely to those found by nature. In these cases, biomimetic approaches differ from the purely technical ones only by their explicit reference to a biological example, not necessarily by the type of mechanism implemented.

Biomimetic approaches can be found on all levels of the navigation hierarchy, with the exception of search and survey navigation. Perhaps due to its low efficiency, search has not received much attention from biomimetic research. Biomimetic survey navigation, on the other hand, requires all lower level skills, many of which are still poorly understood. In the following review, we therefore consider only the local mechanisms *direction-following*, *aiming* and *guidance*, and

the way-finding mechanisms *recognition-triggered response* and *topological navigation*.

As we stated in the beginning, we included only those biomimetic systems in this review that implemented a biological navigation behaviour on a real mobile robot, not just in simulations. Readers interested in theoretical or simulation models are referred to the comprehensive review of Trullier et al. [76].

3. Local navigation

3.1. Direction-following

Almost every commercially available mobile robot comes equipped with idiotactic odometry and active proximity sensors which allow for simple wall and corridor following behaviours. Thus, it is not surprising that direction following is a very common navigation mechanism in robotics. The contribution of biomimetic approaches consists mainly of two behaviours: (i) reactive trail following mechanisms with passive sensors; (ii) allothetic path integration.

3.1.1. Trail following

Sharpe and Webb [68]. The robot of Sharpe and Webb is unique in that it navigates with the help of chemical cues. It is inspired by the trail following behaviour of many ant species which lay chemical trails along the ground. The ants walk in a sinusoidal path through the “vapour tunnel” created by the evaporating chemicals (Fig. 1(c)). The mechanism assumed to underlie this behaviour is called “osmotropotaxis”: The ant detects the concentration of the chemical with both antennae. The concentration difference between left and right determines the turning tendency towards the higher concentration [23].

In Sharpe and Webb’s robot, the antennae function was replaced by two chemical sensors which detected alcohol vapours. The sensors were mounted on “antennae” extending forwards and sideways at a span of 29 cm. The osmotropotactic mechanism was implemented by a simple neural controller. A stronger sensor signal on one side of the robot accelerated the wheel speed on the opposite side which in effect caused the robot to head towards the greater concentration of vapour. The overall speed depended on the concentration measured by both sensors so that

the robot slowed down when it received only faint odour signals. Sharpe and Webb were primarily interested in the question of whether the neural controller could reproduce the observed ant behaviour under real world conditions. In their experiments, they laid alcohol trails of various lengths and concentrations on the laboratory floor which the robot had to follow. Sharpe and Webb found a close qualitative correspondence between robot and ant behaviour, e.g., a similar dependence of trail following performance on body axis orientation, speed and antennae span.

Coombs and Roberts [14]. Trails can also be defined by walls or corridors. Instead of using ultrasonic or infrared sensors, flying insects [71] are able to centre their flight path in a corridor by balancing the image motion in their two eyes (cf. Fig. 3). Since the optic flow during translation depends on the distance of the objects being passed, balancing the image flow results in balancing the object distances and thus in centring. In addition, honey bees have been shown to regulate flight speed by trying to keep the overall image motion as constant as possible [72]. When forced to fly down a tunnel, bees slow down at narrow passages and accelerate in wider sections. In this way, the bee adapts its speed to the current situation since tight passages require highly controlled flight manoeuvres.

The centering behaviour of bees inspired Coombs and Roberts to build “bee-bot”, a robot equipped with a wide-angle camera with a 115° field of view. In




Fig. 3. A robot finds the midline between obstacles by balancing the optic flow on both sides. Since the optic flow during translation is inversely proportional to the distance of the passed objects, balancing the image flow results in balancing the distances.

the right and left third of the visual field, the vertical optical flow was computed using a gradient method, while the central third was ignored. The maximal flow value on each peripheral field indicated the nearest obstacles on the right and on the left. By balancing the maximal flow on both sides, bee-bot centred its course between the nearest objects. While correcting its course, bee-bot's camera actively counter-rotated to prevent the rotatory flow from contaminating the flow field. When the difference between gaze direction and heading exceeded a threshold, the camera was returned to the heading in a fast saccade-like rotation. At the same time, bee-bot controlled its forward speed to keep the flow in a measurable range. Coombs and Roberts tested bee-bot in a laboratory environment demonstrating its ability to centre its course between obstacles.

Santos-Victor et al. [66]. Instead of actively compensating for rotation, Santos-Victor et al. reduced rotatory effects by using a steering algorithm that always kept the rotary flow component at a sufficiently low level. Their robot, called “Robee”, was equipped with two cameras pointing at opposite lateral directions (“divergent stereo”). In both cameras, the horizontal flow was measured using a standard gradient scheme and averaged over the image to increase robustness. The robot steered into the direction of the smaller average flow while adapting its forward speed such that the overall flow remained constant. Thus, Robee centred between the average object distances, not between the nearest objects. In addition, the robot switched from centring to wall following behaviour when the flow on one side vanished while passing an untextured wall. Robee could follow straight and curved corridors with non-uniform texture. An approach similar to the one of Santos-Victor et al. was taken by Duchon and Warren [15] who instead used a single camera pointing into the forward direction. Centring was achieved by comparing the average flow in both halves of the visual field.

Weber et al. [78]. The distance of the corridor walls can be computed from their image velocity, if the insect or robot knows its own speed [49]. This idea was used in the robot of Weber et al. which balanced flow-derived distances instead of image flow. Weber et al. used a V-shaped mirror mounted above the camera

to obtain the lateral views. The image motion was computed with the help of an image interpolation technique that delivers not only the image velocity of the walls, but also their angle with respect to the driving direction [70]. Weber et al. removed the rotatory flow component by subtracting the current rotation obtained from the robot's wheel encoders. From the corrected flow, the robot computed the range of the lateral walls using its forward speed, again obtained from its odometry. The robot set its course such that it followed the midline defined by the currently measured wall orientation and distance. Weber et al. demonstrated in a number of experiments that the robot was able to successfully traverse differently formed corridors with textured walls. Due to the odometry-based removal of the rotatory flow, the robot did not need to restrict its turning velocity, nor did it need to compensate for rotation by active gaze control.

3.1.2. Path integration

Lambrinos et al. [35]. Idiothetic path integration based on wheel encoder signals is very common in robotics, but is subject to error accumulation as no external reference is used. Allothetic compasses can help to reduce error accumulation, and are therefore, often used by animals. Celestial cues such as the sun or the stars can serve as a compass, if the time of day is known. One way to estimate the sun's position is to analyse the polarization pattern of the sky that arises from the scattering of the sunlight in the atmosphere. This is particularly useful when the sun itself is obscured by clouds or large objects such as trees or buildings. Many insects analyse the celestial polarization pattern for compass orientation with the help of polarization-sensitive photoreceptors (for review, see [79]).

The "Sahabot" of Lambrinos et al. used such a polarized light compass for path integration, together with distance estimates from its wheel encoder signals. The compass consisted of three polarization analysers oriented along different polarization planes which mimicked the spatial layout and neural processing of the insect eye. This sensor arrangement was used to estimate the predominant polarization plane of the scattered light at the zenith. The polarization plane at the zenith is perpendicular to the sun's position and thus allows for the determination of the solar meridian. An

additional suite of ambient light sensors indicated on which half of the solar meridian the sun was currently positioned. From the geographical position and the time of day and year, the compass direction could be computed from the estimated position of the sun.

Lambrinos et al. tested three different compass models that were proposed for the Saharan desert ant *Cataglyphis* and compared their performances. Interestingly, the experiments were conducted in the same natural habitat as that of the desert ant. The path integration accuracy of Sahabot was surprisingly high, in a range comparable to that of real desert ants which are well known for their highly developed path integration abilities.

Chahl and Srinivasan [8]. Visual input allows insects not only to estimate their global orientation from visual input, but also the distance travelled. As a possible mechanism, it has been proposed that distance estimates are obtained by integrating the optic flow over time [63,72]. This inspired Chahl and Srinivasan to investigate whether a mobile robot is able to navigate by path integration using only vision as sensory input.

Chahl and Srinivasan's robot carried two cameras, one at the front and one at the rear of the robot. Each of them pointed up to a conical mirror that allowed to capture a panoramic image of the environment (see Fig. 4, [9] for a detailed description). This conical mirror camera has become quite popular in biomimetic vision systems since it provides a two-dimensional model of the almost omnidirectional insect eye which is relatively easy to construct [17,18,28,47].

To measure translation, the robot recorded two one-dimensional reference images of the horizon from both cameras before moving. After a forward movement, the new image of the rear camera was compared to an interpolated image generated from the reference images assuming a linear transformation between them. The distance travelled was then derived from the interpolation parameter and the known distance of the two cameras. A similar procedure was applied to estimate rotation. The setup allowed the robot either only to translate or to rotate, not to do both at the same time. The experiments of Chahl and Srinivasan demonstrated that visual path integration based on image motion is indeed possible. Compared to conventional path integration based on wheel encoders, visual path integration did not prove superior.

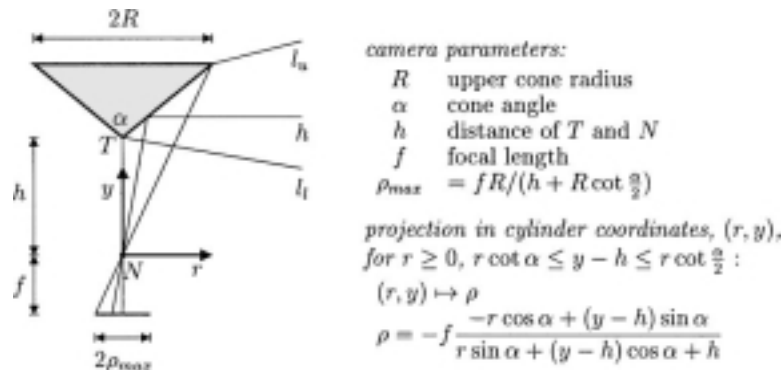


Fig. 4. Imaging geometry for a conical mirror camera. N , camera nodal point; T , tip of conical mirror; l_u , upper limiting ray; h , horizontal ray; l_l , lower limiting ray. The conical mirror camera allows for capturing omnidirectional images without rotating a camera. A ring-shaped visual field between l_l and l_u is mapped to a circular disk with radius ρ_{\max} in the image plane. The visual field contains the horizon if $\alpha > 90^\circ$ and $R > -h \cos \alpha \tan \frac{1}{2}\alpha$.

While not subject to wheel slippage, it suffered from non-uniform contrasts and occlusion effects of passing objects during translation. However, the results were in an acceptable range indicating that visual path integration is a suitable alternative for cases where wheel encoders are not available such as, for instance, in flying agents.

Weber et al. [78]. The issue of estimating the distance travelled from image motion was also taken up by Weber et al. (cf. Section 3.1.1). Using the same robot as in their corridor following experiments, Weber et al. integrated the lateral flow on both sides over time. In contrast to Chahl and Srinivasan [8], they could not rely on reference images with a known spatial distance. Therefore, their method did not yield absolute values of the distance travelled. Weber et al. did not take into account rotatory movements as the rotatory flow component was removed using odometry. A series of corridor traversals at different speeds showed that the integrated flow values were indeed very consistent as long as the same corridor was travelled. This was not the case for varying environments, since the optic flow depends on the environment-specific distance distribution.

3.2. Aiming

Aiming is among the navigation behaviours most easily implemented on a mobile robot. In his influential book *Vehicles*, Braitenberg [3] describes the basic




Fig. 5. Approaching a goal marked by an object using fixation behaviour [28]. The image flow created by passing the object elicits a turning response towards the object by decreasing the wheel speed on the side of the object, and increasing the wheel speed on the opposite side.

architecture (Fig. 5): Two sensors and two motors on each side of the robot with either cross-wise and excitatory (Vehicle 2), or uncrossed and inhibitory (Vehicle 3) connections lead to a simple aiming behaviour by turning the robot towards the stronger stimulated side during forward motion. These vehicles find point-like stimuli such as a bright light or an odour source by balancing the inputs of the two sensors during approach, i.e., aiming is achieved using a tropotactic mechanism. Here, we describe two biomimetic type 3 vehicles which were designed to test computational models of insect aiming behaviour. The first vehicle finds sound sources with a specific frequency characteristic, the second finds conspicuous objects showing strong visual contrasts.

Webb [77]. Female crickets are able to find a conspecific male by turning and approaching the sounds produced by the male (phonotaxis). The turning response is elicited when the cricket detects a phase difference between its ears on the two front legs, but only if the sound frequency is in the specific range of the cricket species. The detection of phase differences is achieved by the anatomy of the auditory apparatus: sound travels to the cricket's tympana either externally or via an internal air-filled tracheal connection. Upon arrival at the tympanum, the sound signals from the two pathways have different phases. The phase difference leads to different degrees of amplification or cancellation of the incoming signals in both ears, thus allowing the cricket to detect the direction of the sound source. Webb argued that this mechanism might also be responsible for the frequency selectivity of the cricket response since the detection mechanism works only in a very limited frequency range.

To support this hypothesis, Webb devised a controller which electronically mimicked the relative phase discrimination mechanism and the different internal travel times of the incoming sound signals. The controller was set on a mobile robot with two laterally displaced microphones to test whether it could produce a cricket-like phonotaxis. In this way, Webb's robotic simulation could use the complex sound structure of a real environment which is very hard to model in a computer simulation. In a number of experiments, Webb's robot performed very similar to the real cricket. It found an artificial sound source under a variety of conditions, even when the source was located behind an obstacle. Thus, the experiments confirmed that the same mechanism could be responsible for both phase discrimination and frequency selectivity. In more recent work, the capabilities of Webb's system were extended so that the robot was able to find real crickets instead of an artificial sound source [39].

Huber and Bülthoff [28]. Many image-based approaches in robotics rely on sophisticated aiming mechanisms that require the segmentation of salient objects from the background. Aiming towards conspicuous objects may be realized using much simpler mechanisms. Flies, for instance, orient towards black stripes in a homogeneous environment (fixation behaviour, [57]) and approach them. This could be due

to the fly's optomotor response: The fly compensates for disturbances causing a large rotatory image motion by rotating into the direction of the image motion, thus reducing the disturbance. When an isolated object passes from front to back during flight, the fly compensates for the optic flow created by the object by counter-rotating until the object is brought in front of the insect where it creates almost no image motion (cf. Figs. 4 and 5). Thus, both a compensatory optomotor response and an aiming behaviour might be implemented in a common sensomotor controller [22].

Huber and Bülthoff tested the viability of this hypothesis on a mobile robot equipped with a conical mirror camera. The robot's image processing system was a relatively detailed one-dimensional model of the insect motion processing pathway. The motion signals were analyzed by an array of Reichardt motion detectors [24] the outputs of which were integrated over each hemisphere. The integrated flow from both hemispheres was used to control the robot's driving direction in the same way as in Braitenberg's type 3 vehicle (cf. Fig. 5). The system of Huber and Bülthoff resembled those designed for flow-based trail following (cf. Section 3.1.1), but was used for aiming instead.

Similar to the insect experiments of Poggio and Reichardt [57], Huber and Bülthoff conducted their robot experiments in a circular arena with black stripes as visual stimuli. The robot produced both fixation and compensatory optomotor behaviour, depending on whether a single isolated stripe or a wide-field pattern was presented. Thus, the plausibility of a common controller for both behaviours could be confirmed in a real world test. The resulting aiming mechanism allowed the robot to find salient objects without requiring sophisticated segmentation techniques.

3.3. Guidance

Bees and ants are able to use visual guidance (scene-based homing) to find a location which is only defined by an array of locally visible landmarks (for review, see [11]). The experimental evidence suggests that these insects store a relatively unprocessed snapshot of the surrounding panorama as seen from the goal. Cartwright and Collett [7] developed a computational model that could find a goal by matching the snapshot with the current view (cf. Fig. 6). In their model, they assumed that the views are omnidirec-




Fig. 6. Scene-based homing with a panoramic snapshot taken at the goal: In the current view, the image positions of the surrounding landmarks are displaced with respect to the snapshot. The goal can be found by moving into a direction that diminishes the image displacements (after [7]).

tional and aligned with an external reference direction. Landmarks appeared as black regions in the views, the background being white. The matching between views was done by searching for black regions in the current view near the image position where it was seen in the snapshot. From the difference between the image positions in the snapshot and the current view, a movement direction was computed to reduce the perceived difference. The computed movement directions for all visible landmarks were summed up to give the overall flight direction of the model bee. Computer simulations showed that the model could indeed account for the observed search behaviour of honeybees.

This simple form of visual guidance has inspired several robot implementations since no complex scene representations have to be handled to find an inconspicuous goal. The biomimetic approaches introduced below mainly differ in the way they establish correspondences between views; this is one of the principal problems when transferring the idealized model of Cartwright and Collett [7] to a real world robot application. Similar approaches that do not explicitly refer to biological behaviours have been presented by Hong et al. [27] and Neven and Schöner [51].

Röfer [61]. Röfer used a robot with a synchronous drive that kept it at a nearly constant orientation during operation. A rotating photoreceptor allowed for recording one-dimensional omnidirectional grey value signatures. These views were assumed to be roughly aligned to a global reference direction because of the

constant orientation of the robot platform. This made the system susceptible to cumulative error since a slow orientation change due to wheel slippage is hard to avoid during operation. In contrast to the model of Cartwright and Collett, Röfer could not rely on a segmentation of the image into landmarks and background. Instead, correspondences between the images were learnt by a one-dimensional Kohonen network. The nodes of the network were initialized with the grey values of the snapshot and had to converge to the grey values of the current view by changing their image position. The movement directions were computed from the different positions of the nodes in the snapshot and the current view, and summed up as in the model of Cartwright and Collett. The robot could only move step-wise since it had to take a new scan of the photosensor between each movement. In a later implementation [62], Röfer improved this scheme in several respects: Instead of the rotating photoreceptor, he used an omnidirectional sensor similar to the conical mirror camera which allowed for continuous robot motion. In addition, he generalized his correspondence algorithm to non-aligned and coloured views to improve its robustness.

Franz et al. [18]. Franz et al. analyzed the computational foundations of snapshot-based guidance and provided proofs on convergence and error properties. Based on the mathematical analysis, they developed an alternative model to that of Cartwright and Collett. Instead of computing correspondences locally as in the other approaches, Franz et al. used a set of templates for the entire correspondence field derived from simplifying assumptions. Each correspondence template generated a novel view from the stored snapshot which was compared to the current view. The template leading to the best reconstruction of the current view was taken as a base for the computation of the goal direction. As in Röfer's approach, the algorithm of Franz et al. worked with unsegmented grey value images, but the views did not need to be aligned with a reference direction. Robust performance was shown in a number of experiments in a realistic low contrast environment using a miniature robot with a conical mirror camera. Part of this robustness can be attributed to the global image matching technique which largely restricts the search space for correspondences. However, the limited set of correspondence templates and

the underlying assumptions mostly restricted the area from which the goal could be found (catchment area) to the open space around the snapshot.

Möller et al. [47]. Möller et al. directly implemented the model of Cartwright and Collett [7] on the Sahabot 2, the successor of Sahabot (see Section 3.1.2) which also carried a conical mirror camera. The experiments were conducted on a flat plane in the Sahara desert with four black cylinders as landmarks, similar to the ant experiments of Wehner et al. [80]. In contrast to the approaches described above, the omnidirectional image was segmented into the black landmarks and the bright background, and the snapshots were aligned at a compass direction obtained from the polarized light compass of Sahabot. Correspondences were found by pairing black image regions, as in the original model. The experiments showed that Sahabot 2 could reach the goal with high accuracy, even from outside the area defined by the immediately surrounding landmarks. This demonstrated that the original model of Cartwright and Collett can be transferred to a real world application, provided that the environment allows for a reliable segmentation of the landmarks from the background. While this might be feasible in the Sahara desert, it is not clear how the scheme of Möller et al. performs in environments without highly distinctive landmarks, and how much of its performance can be attributed to the use of additional compass information which was not used by the other approaches.

4. Way-finding

4.1. Recognition-triggered response

Insects can associate movement decisions with visual landmarks. Ants, for instance, may learn to always pass a landmark on the right side. This association persists, even when the order of the landmarks or their positions relative to the nest are changed [12]. Bees are able to learn routes, i.e., connected chains of recognition-triggered responses. Collett et al. [13] observed that bees, after associating a sequence of landmarks with movement decisions, showed inconsistent behaviour when the order of landmarks was changed. When a navigation task requires only stereo-

typed route following, vertebrates also appear to use recognition-triggered responses [21,36].

Recognition-triggered responses have been used in a number of biomimetic navigation systems which differ in the way they recognize a location and in the local navigation method associated with it. They fall into two groups: The first type uses a collection of unconnected associations of compass directions with recognizable sensory situations. The memorized directions all lead to a single goal location so that this behaviour lies in the middle between local navigation and way-finding. For an implementation of this idea into a neural associative memory, see Barto and Sutton [2].

The second type connects the responses to routes, i.e., ordered sequences of recognition-triggered responses. All of these systems only require routes as building blocks for topological navigation skills. We will discuss them in the next section.

Nelson [50]. Nelson presented a robot capable of a view recognition-triggered response together with a theoretical analysis of the associative memory used in his approach. Instead of a mobile robot, Nelson used a small camera mounted on a robot arm which could move at a fixed height over a miniature city called “Tinytown”. Places were recognized from the bird’s eye view of the camera. Each camera image of a location was divided into a 5×5 grid of adjacent subwindows in which the predominant edge direction was determined. In the training phase, Nelson’s system learned associations between local edge orientation patterns with directions towards the goal in the town centre at 120 preselected places covering Tinytown. Homing was accomplished by comparing the currently perceived pattern against all patterns stored in the associative memory, and moving into the specified direction of the winning pattern. Nelson’s system was able to find the goal from almost all places in Tinytown and thus demonstrated the feasibility of recognition-triggered responses in robot applications. However, his system relied on some idealizations: pictures were taken in exactly the same camera orientation, the system had to be trained at all possible locations and could not select places autonomously.

Gaussier and Zrehen [20]. Similar to a theoretical study by Krakauer [30], Gaussier and Zrehen pro-

posed recognition-triggered responses instead of guidance as an alternative explanation of the scene-based homing behaviour of insects (cf. Section 3.3). Their mobile robot learned associations between compass directions and landmark configurations. The landmark configurations were extracted from panoramic images obtained from a rotating camera. A neural preprocessing stage determined so-called “focal points” (corners or junctions) in the panoramic image. Around each focal point, a local landmark view was extracted and its compass direction determined. In addition, the system stored the camera movements leading from one focal point to the next. Thus, a place was characterized by a sequence of local landmark views and bearings connected by camera movements. In contrast to Nelson’s [50] approach, the recognition-triggered responses were learned during an autonomous exploration phase in which Gaussier and Zreben’s robot made repeated excursions from the goal in various directions. During each excursion, a group of neural networks learned the landmark configuration at a certain distance from the goal, and associated it with a vector pointing into the goal direction. While homing, the vector belonging to the most similar landmark configuration was activated which led to a zig-zag trajectory along the decision boundary between neighbouring landmark configurations (cf. Fig. 7). The system of Gaussier and Zreben could find its goal from any position inside an office room. In addition, Gaussier and Zreben showed that three scene-vector-associations




Fig. 7. Finding a goal using recognition-triggered responses [20]. During exploration, the robot learns associations between the currently perceived scene and a compass direction. During homing, the vector belonging to the most similar scene is activated.

are sufficient for reliable homing. Thus, their system had to learn only a few recognition-triggered responses as opposed to the large number required by Nelson’s [50] approach.

Recce and Harris [60]. A class of neurons in the rat hippocampus, the so-called “place cells”, have been found to consistently discriminate between different parts of an environment, i.e., they fire only in one part of the environment, but not in others [53]. It is generally believed that place cells play an important role in navigation, and a large number of theoretical hippocampus models have been presented (e.g., [4,6,46,48,54,65,73]). The system of Recce and Harris is an implementation of Marr’s hippocampus model [42] on a mobile robot. In this model, the hippocampus was viewed as an autoassociative memory which stores a scene representation consisting of the bearings and distances of the surrounding landmarks and of a goal location. Place recognition was achieved by feeding the current scene into the autoassociator which activated the stored scene memory if it matched the current scene. In the model of Recce and Harris, this corresponded to the activation of a place cell. The stored direction and distance of the goal were activated together with the scene memory and could be used to directly drive towards the goal. In this respect, the approach of Recce and Harris was very similar to Nelson’s which also required the storing of scene representations all over the environment. In contrast to the other approaches, the landmark bearings were not extracted from the visual input, but from omnidirectional sonar scans which also yielded the landmark distances. The robot did not need compass information or a constant orientation since it could “mentally” rotate the scene representation to find the maximal activation of the autoassociator. Recce and Harris compared the activation patterns of the simulated place cells with the experimental evidence and found a good qualitative correspondence. In addition, the robot could find the goal from all locations in the test environment which confirmed that Marr’s model could be used for a real world navigation task.

Burgess et al. [5]. Burgess et al. describe a robot implementation of an earlier neurophysiological model of the rat hippocampus [6]. Some place cells can be shown to fire at a relatively fixed distance from a

Table 2
Biomimetic topological navigation systems

Author(s)	Place recognition	Local navigation	Route integration
Matarić [43]	Context of preceding	Wall and corridor following	Context-based place recognition
Kortenkamp and Weymouth [29]	Gateway type and stereoscopic features	Wall and corridor following	Human trainer
Bachelder and Waxman [1]	Configuration of landmark views	Direction following	Place recognition
Mallot et al. [40]	Binary views	Corridor following	View recognition
Franz et al. [17]	Panoramic views	Visual guidance	View similarity and verification
Owen and Nehmzow [55]	Sonar signature	Direction following	Rehearsal procedure

wall [52]. This property inspired the place recognition mechanism of the robot of Burgess et al. which visually estimated the distances to the surrounding walls. Compared to the very general environments of Recce and Harris [60], their approach was designed for the rather specialized case of a rectangular arena with a visual marker on the northern wall. The wall distance was derived from the perceived height of the edge between wall and floor in the visual field. Thus, each place in the arena was characterized by a specific combination of wall distances.

During a first exploration phase, the robot rotated on the spot at all locations of the arena, to face all walls and to estimate their distance. The robot's orientation with respect to a geocentric reference direction was derived from path integration which was periodically reset by using the visual marker on the northern wall. A competitive learning mechanism selected a number of place cells to represent the specific wall distances for each place. In a second learning phase, the robot associated "goal cells" with the place cells representing four reference locations from which the direction towards the goal was known. Thus, the goal cells played a role similar to Gaussier and Zrehen's [20] recognition-triggered responses. However, Burgess et al. computed the goal direction from the relative activity of all goal cells instead of selecting the most active one. This should result in smooth trajectories instead of the zig-zag course of Gaussier and Zrehen's robot. The system showed good agreement with the neurophysiological data and thus demonstrated the validity of the hippocampus model in a simplified, yet real world situation.

4.2. Topological navigation

One of the most influential biomimetic ideas in robotics is that of topological navigation. It provides

an alternative to the computationally expensive map paradigm discussed in the introduction, without necessarily restricting its performance. Therefore, many recent approaches are topological, or construct a topological representation from a global map (e.g., [74]). In contrast, biological systems seem to construct topological representations by integrating routes in a bottom-up manner [38]. This ability has been observed in many animals, ranging from honeybees [16] to humans [21]. The systems described below follow this bottom-up approach using different local navigation, place recognition and route integration strategies (cf. Table 2).

Matarić [43]. Experiments on rat navigation motivated Matarić to build a robot that concatenated different wall and midline following behaviours to routes and graphs. The robot was equipped with a ring of ultrasonic sensors and a compass. It executed a single behaviour as long as the sensory conditions remained qualitatively the same, e.g., as long as the robot's ultrasonic sensors detected a corridor leading into the south. A new behaviour was triggered by arriving at a distinctive place showing a qualitative change of the immediate environment such as a corner or a dead end. In contrast to the other approaches, the recognition of these places was only determined by their context, i.e., by the sequence of actions preceding the current action. Thus, the only information stored in the graph representation were actions, not place descriptions. In the graph, each action was labelled by the type of the local environment and by the navigation behaviour used to traverse it (e.g., follow left wall heading south). These features allowed the robot to acquire routes autonomously by simply following the walls of the experimental room. Routes were integrated as soon as the robot encountered a previously visited local environment requiring the same local nav-

igation behaviour. This shows a limitation of Mataric's approach: The number of uniquely identifiable combinations of wall type and compass direction is small. In general environments, the problem of *perceptual aliasing* arises, i.e., distinct locations appear identical to the robot's sensors [55]. Arriving at a new situation identical to an already stored one, Mataric's robot would perform a false route integration. Thus, the navigable environment remained confined to rooms with a uniquely identifiable wall configuration.

Kortenkamp and Weymouth [29]. Kortenkamp and Weymouth tried to reduce the problem of perceptual aliasing by enriching the sensory information defining a place. Their approach was based on concepts derived from a general theory of human cognitive mapping that also involved topological navigation [10]. As in Mataric's approach, Kortenkamp and Weymouth's robot was designed to navigate in indoor environments. A place was recognized by its associated "gateway", i.e., an opening in the lateral walls such as, a door. The gateways were detected and classified with the help of a sonar ring around the robot. Kortenkamp and Weymouth found 24 gateway types in a typical indoor environment. The detected gateway type was combined with stereovision data to achieve a rich sensory place characterization. Between gateways, the robot navigated by wall or corridor following. While the richer place representation allowed Kortenkamp and Weymouth's system to recognize a much greater number of places than Mataric's, this alone did not exclude the possibility of perceptual aliasing. However, they avoided this problem by employing a human supervisor during exploration of a new environment. At each gateway, the supervisor provided the identity of the place so that route integration was basically done by the trainer.

Bachelder and Waxman [1]. Bachelder and Waxman designed a neurocomputational architecture to emulate the place cells in the rat hippocampus. In contrast to the recognition-triggered response models of Recce and Harris [60] and Burgess et al. [5] (cf. Section 4.1), they assume an involvement of the hippocampus in topological navigation. Place cells could encode a topological representation since they are highly interconnected via modifiable synapses. In Bachelder and Waxman's system, place cells formed the vertices of a

graph interconnected by movement decisions. In contrast to the other recognition-triggered response models, places were defined only by landmark views, not by landmark distances.

The first layer of Bachelder and Waxman's architecture partitioned a room into regions defined by a specific configuration of landmark views. The landmarks were differently shaped objects covered by an arrangement of lamps to simplify the image processing. While following a path prescribed by a supervisor, Bachelder and Waxman's robot recorded at each place the characteristic configuration of lamps. This was done by rotating a camera on the robot and relating the detected landmarks to a compass direction. Regions with similar lamp configurations were associated with the same place cell. The regions described by the activation of a single place cell correspond to the place concept of the other topological approaches. Movements leading from one region to another were learnt by the second layer, a heteroassociative network. This network, however, required several hours of off-line training on the data gathered by the robot, and was not subsequently tested on the robot. As they were mainly interested in neurocomputational questions, Bachelder and Waxman did not address the problems of topological navigation such as exploration, path planning and perceptual aliasing.

Mallot et al. [40]. The system of Mallot et al. is based on the view graph theory of navigation developed by Schölkopf and Mallot [67]. According to this theory, biological navigation behaviour in maze-like environments can be explained if one assumes an underlying topological representation. This so-called "view graph" consists of local views as vertices and their spatial relationships (adjacencies or movement directions) as edges.

Mallot et al. used a miniature robot to explore hexagonal mazes. The robot was equipped with "two pixel vision", i.e., two infrared sensors looking downward to the textured floor. The local views were binary patterns on the floor which were recognized while a maze junction was approached. Between junctions, the robot travelled by means of corridor following using infrared proximity sensors. In contrast to the other approaches, Mallot et al. did not integrate the local views into a common place representation. The view graph was learned by a neural architecture that asso-

ciated sequences of views with movement decisions, similar to the one used by Bachelder and Waxman [1]. The robot explored a maze of 12 places and 24 “views” in about an hour. Afterwards, the robot was able to find the shortest path between arbitrary pairs of views. Perceptual aliasing did not occur in the specific setup of Mallot et al. since every view was unique in the maze. However, their network facilitated view recognition by taking into account context information from previously visited views which decreased the probability of perceptual aliasing.

Franz et al. [17]. The robot of Franz et al. was designed to demonstrate that view graphs can also be learned in open environments without a maze-like structure. Similar to Nelson [50], the experiments were conducted in a miniature town, but the robot moved on the ground between the toy houses. The only sensory input used by the robot were one-dimensional grey value signatures obtained from an omnidirectional conical mirror camera. Places were recognized by comparing the current view to stored omnidirectional snapshots. Between snapshots, the robot navigated with the help of the visual guidance strategy described in Section 3.3 [18]. Routes consisted of chains of snapshots taken during autonomous exploration of the environment. In contrast to Mallot et al. [40], only the adjacencies of the snapshots were stored, not the movement decisions connecting them. Whenever view similarities to another route were found, the system tried to navigate to the most similar snapshot. If successful, both routes were connected. New snapshots were taken only when the current view was classified as sufficiently distinct from all previously taken snapshots. This feature prevented the system from including ambiguous information into the view graph, but limited the navigable range of the system to regions with non-ambiguous visual input.

During exploration, the robot tried to move into the least explored regions of its view graph, i.e., regions with a smaller number of snapshots and less connections between them. This accelerated the exploration process which still took more than an hour. The robot recorded between 20 and 50 snapshots that typically covered two thirds of the test arena. In contrast to Mallot et al., Franz et al. did not use contextual information to disambiguate visual information. However, the view information used by the robot proved to be

sufficiently rich to prevent perceptual aliasing in their test arena.

Owen and Nehmzow [55]. The topological navigation system of Owen and Nehmzow recognized places from omnidirectional sonar readings. Their robot travelled between places by following a recorded vector (distance and compass direction) pointing from one place to another. The robot could not explore autonomously, but had to be guided along a path by a trainer. Similar to the system of Franz et al. [17], the robot recorded a new place during exploration whenever the sonar readings were sufficiently distinct from the familiar places. When the robot encountered a sensory situation which was similar to a familiar place in the graph, it automatically tested the place identity by verifying its graph context. To that end, the robot travelled to all recorded neighbours of the place (rehearsal procedure, cf. [33]). When it failed to find the recorded neighbours, a new place node was included in the graph. Although the ultrasonic signatures provide much less information than the views used by Franz et al. [17], the rehearsal procedure allowed Owen and Nehmzow’s system to map areas with ambiguous sensory input while keeping the risk of perceptual aliasing low.

5. Discussion

The above described robot systems show that, mostly within the last five years, biomimetic approaches have been developed for most types of biological navigation, with the exception of search and survey navigation. However, the field is still at the beginning of its development. This is documented by the fact that all of the reviewed approaches aim at testing biological models or mechanisms rather than at finding an optimal technical solution to a given problem.

Researchers in this field are motivated by two basic interests: the first (biological) interest is to subject biological hypotheses to a rigorous test in the real world; the second (bionic) interest is to find new mechanisms that might be utilized in future navigation systems. In most cases, both interests cannot be pursued in the same system: Either, realistic modelling of animal behaviour restricts the technical application to very spe-

cialized cases (such as the cricket robot of Webb [77]), or the investigated mechanism is so highly abstracted that predictions of actual animal behaviour are difficult (as, e.g., in Mataric's [43] topologically navigating robot).

There are good reasons to test biological theories of navigation on real mobile robots. Webb [77] argues that, even in the most elaborate simulations, some potentially important effect may be incorrectly modelled or overlooked. A computational step left out for convenience may prove to be crucial in the real world. This cannot happen in a robot test, as all system components are evaluated simultaneously. Thus, robot implementations provide a more rigorous test for behavioural models than computer simulations. In addition, realistic computer models of visual or auditory input are extremely complex. Often a robot implementation can be done with less effort.

From the bionic point of view, biomimetic research has yielded a number of interesting new navigation mechanisms. In particular, corridor-following based on optic flow, path integration with a polarization compass, snapshot-based guidance and recognition-triggered responses are now well understood. In terms of robustness and performance, their robotic implementations proved to be comparable or superior to existing, purely technical approaches. The biomimetic topological navigation systems, however, are still outperformed by current technical systems. This is due to the fact that the described topological systems either cannot explore their environment autonomously, or do not use sufficiently sophisticated place recognition techniques that would allow for mapping a larger environment.

Interestingly, most of the navigation mechanisms presented here are inspired by insect behaviour. The main reason might be the fact that, due to their limited information processing capacity, insects have to rely on computationally inexpensive, yet robust minimalistic mechanisms. This makes them more accessible to biological experiments, and simplifies their robot implementation. The complicated navigation behaviour of vertebrates is still less understood so that most corresponding biomimetic approaches have to remain on a very abstract level. An exception are the robots using models of the rat hippocampus that can draw on a large number of neurophysiological investigations. But even here, the possible interpreta-

tions are diverse. Bachelder and Waxman [1], together with many computational models, assume that the hippocampus is involved in topological navigation behaviour, while Burgess et al. [5] and Recce and Harris [60] propose a role in recognition-triggered responses. McNaughton et al. [45,65], finally, argue that the network of hippocampal place cells is primarily a path integration system. Landmarks provide additional information used to prevent error accumulation.

5.1. Cognitive maps

A major distinction of navigation competences which has been elaborated by O'Keefe and Nadel [54] is the one between stereotyped and flexible behaviours. Mechanisms like path integration, aiming, or even route behaviour allow an agent to go to one goal. If the agent wants to go to a different goal, it would have to learn everything anew, even if parts of the required route had been used already for the previous goal. In contrast, topological and survey navigation are goal-independent memories of space that can be used for many different routes as well as for the planning of routes. Learning of such goal-independent memories is "latent", i.e., it is driven by curiosity, not reinforcement and may occur long before the agent knows which goal might become interesting. Goal-independent ("declarative") memories of space are called cognitive maps (in the sense of O'Keefe and Nadel [54]) whereas route-knowledge is a form of procedural memory. The flexibility of cognitive maps is not free, however. While procedural memories of space generate an action command (such as "go left!") as the output of the memory, cognitive maps have to be "consulted" by an additional device. In themselves, they can only provide information like "if you want to go to place B, you should turn left here".

In their original introduction of the term *cognitive map*, Tolman et al. [75] stressed the idea that cognitive maps can be used as a route planning stage, which is well in line with the idea of a declarative memory of space. Consequently, Tolman et al. suggested the use of shortcut behaviour as evidence in support of cognitive maps (see also [19]). This is problematic, however, since shortcuts can also be found by path integration or aiming to a distant landmark; in order to conclude the involvement of a cognitive map from

shortcut behaviour, such local mechanisms would have to be ruled out first.

5.2. Relationship to other terminologies

In the literature, a number of terminologies for navigation behaviour have been introduced. The approach presented here is closest to the navigation hierarchy of Trullier et al. [76] in which we included the additional levels *search* and *direction-following* at the lower end. The main difference, however, is that in our terminology, behaviour is categorized only on the basis of competences, rather than on the basis of information structure and content. Our intention hereby is to make the categories accessible to behavioural experiments and benchmarks so that both biological and robot behaviours can be related in the same framework.

Some authors insist that recognition-triggered responses must be directly associated with a place, independent of the viewing angle or the orientation of the agent [58,76]. This is an unnecessary restriction as it requires the integration of different sensory configurations into a common place representation. In fact, many way-finding skills may be accomplished without the need for this integration, as is demonstrated by the system of Mallot et al. [40]. Their robot learned a topological representation of a maze without integrating local views into a common place representation.

Although survey navigation is the highest level in our hierarchy, this does not necessarily imply that an agent first has to acquire a topological representation, as suggested by some authors [34,58]. There is evidence that humans can immediately acquire survey knowledge during exploration of a novel environment [64]. Thus, the proposed hierarchy of spatial information processing does not necessarily result in a hierarchical organization of knowledge acquisition. Similarly, metric information cannot be associated with a particular level of the navigation hierarchy: it may arise from very simple mechanisms, such as path integration, as well as from the inference of spatial relationships in survey navigation.

In general, it seems sensible to distinguish between different types of navigation hierarchies that are to some degree independent: the logical hierarchy

describing the complexity of navigation behaviour and underlying mechanisms (e.g., [32]); the evolutionary hierarchy found in the animal kingdom (e.g., [76]); the ontogenetic sequence of the development of behavioural competences in the child; and finally, the acquisition hierarchy during exploration of an environment.

5.3. Future directions

Our analysis has shown that two types of navigation behaviours are still awaiting a biomimetic robot implementation: search at the lower end of the hierarchy, and survey navigation at the higher end. As we argued at the beginning, an effective search mechanism could be very useful as a backup strategy when other navigation mechanisms fail (a common case in current systems). Moreover, there are some interesting biological examples of search: Desert ants [81] and desert isopods [25,26], for instance, use an effective systematic search strategy when they miss their burrow after a foraging trip. A robot implementation of these search strategies could be a realistic test of the theoretical considerations discussed by Wehner and Srinivasan [81] and Hoffmann [25,26].

Before a biomimetic survey navigation system can be built, the hurdle of topological navigation has to be taken. Although the principle problems are well understood, there is currently no system that autonomously builds graphs from routes in a larger scale environment. In our opinion, reliable topological navigation could be the challenge of the near future. Once this is achieved, a number of interesting questions concerning survey navigation can be addressed; for instance, how the topological representation should be embedded in a common reference frame, or how much metric information is necessary for this behaviour. A biomimetic robot capable of survey navigation would finally close the loop that began with abandoning the traditional map-based navigation systems.

Although many biological examples remain to be explored, the number of well-studied biological mechanisms suitable for robot implementation is relatively small. Therefore, the current development cannot go on infinitely. Future biomimetic systems will have to be designed in close collaboration with on-going biological research. This will allow technical results to

directly influence the direction of empirical research, thus providing a novel, synthetic approach to biology.

References

- [1] I.A. Bachelder, A.M. Waxman, A view-based neurocomputational system for relational map-making and navigation in visual environments, *Robotics and Autonomous Systems* 16 (1995) 267–289.
- [2] A.G. Barto, R.S. Sutton, Landmark learning: An illustration of associative search, *Biological Cybernetics* 42 (1981) 1–8.
- [3] V. Braitenberg, *Vehicles*, MIT Press, Cambridge, MA, 1984.
- [4] M.A. Brown, P.E. Sharp, Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens, *Hippocampus* 5 (1995) 171–188.
- [5] N. Burgess, J.G. Donnett, K.J. Jeffery, J. O’Keefe, Robotic and neuronal simulation of the hippocampus and rat navigation, *Philosophical Transactions of the Royal Society of London B* 352 (1997) 1535–1543.
- [6] N. Burgess, M. Recce, J. O’Keefe, A model of hippocampal function, *Neural Networks* 7 (1994) 1065–1081.
- [7] B.A. Cartwright, T.S. Collett, Landmark learning in bees, *Journal of Computational Physiology A* 151 (1983) 521–543.
- [8] J.S. Chahl, M.V. Srinivasan, Visual computation of egomotion using an image interpolation technique, *Biological Cybernetics* 74 (1996) 405–411.
- [9] J.S. Chahl, M.V. Srinivasan, Reflective surfaces for panoramic imaging, *Applied Optics* 36 (31) (1997) 8275–8285.
- [10] E. Chown, S. Kaplan, D. Kortenkamp, Prototypes, location, and associative networks (PLAN): Towards a unified theory of cognitive mapping, *Cognitive Science* 19 (1995) 1–51.
- [11] T.S. Collett, Landmark learning and guidance in insects, *Philosophical Transactions of the Royal Society of London B* 337 (1992) 295–303.
- [12] T.S. Collett, E. Dillmann, A. Giger, R. Wehner, Visual landmarks and route following in desert ants, *Journal of Computational Physiology A* 158 (1986) 835–851.
- [13] T.S. Collett, S.N. Fry, R. Wehner, Sequence learning by honeybees, *Journal of Computational Physiology A* 172 (1993) 693–706.
- [14] D. Coombs and K. Roberts, Centering behavior using peripheral vision, in: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, IEEE Computer Society Press, Los Alamitos, CA, 1993, pp. 440–451.
- [15] A.P. Duchon, W.H. Warren, Robot navigation from a Gibsonian viewpoint, in: *Proceedings of IEEE Conference on Systems, Man and Cybernetics*, IEEE Computer Society Press, Los Alamitos, CA, 1994, pp. 2272–2277.
- [16] F.C. Dyer, Bees acquire route-based memories but not cognitive maps in a familiar landscape, *Animal Behaviour* 41 (1991) 239–246.
- [17] M.O. Franz, B. Schölkopf, H.A. Mallot, H.H. Bülthoff, Learning view graphs for robot navigation, *Autonomous Robots* 5 (1998) 111–125.
- [18] 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.
- [19] C.R. Gallistel, *The Organization of Learning*, MIT Press, Cambridge, MA, 1990.
- [20] P. Gaussier, S. Zrehen, Perac: A neural architecture to control artificial animals, *Robotics and Autonomous Systems* 16 (1995) 291–320.
- [21] S. Gillner, H.A. Mallot, Navigation and acquisition of spatial knowledge in a virtual maze, *Journal of Cognitive Neuroscience* 10(4) (1998) 445–463.
- [22] K.G. Götz, The optomotor equilibrium of the *Drosophila* navigation system, *Journal of Computational Physiology* 99 (1975) 187–210.
- [23] W. Hangartner, Spezifität und Inaktivierung des Spurpheromons von *Lasius fuliginosus* Latr. und Orientierung der Arbeiterinnen im Duftfeld, *Zeitschrift für Vergleichende Physiologie* 57(2) (1967) 103–136.
- [24] B. Hassenstein, W. Reichardt, Systemtheoretische Analyse der Zeit-Reihenfolgen und Vorzeichenauswertung bei der Bewegungsperzeption des Rüsselkäfers *Chlorophanus*, *Zeitschrift für Naturforschung* 11b (1956) 513–524.
- [25] G. Hoffmann, The random elements in the systematic search behavior of the desert isopod, *Hemilepistus reaumuri*, *Behavioral Ecology and Sociobiology* 13 (1983) 81–92.
- [26] G. Hoffmann, The search behavior of the desert isopod *Hemilepistus reaumuri* as compared with a systematic search, *Behavioral Ecology and Sociobiology* 13 (1983) 93–106.
- [27] J. Hong, X. Tan, B. Pinette, R. Weiss, E.M. Riseman, Image-based homing, in: *Proceedings of IEEE International Conference on Robotics and Automation*, Sacramento, CA, 1991, pp. 620–625.
- [28] S.A. Huber, H.H. Bülthoff, Simulation and robot implementation of visual orientation behaviors of flies, in: R. Pfeifer, B. Blumberg, J.A. Meyer, S.W. Wilson (Eds.), *From Animals to Animats 5*, Proceedings of SAB’98, MIT Press, Cambridge, MA, 1998, pp. 77–85.
- [29] D. Kortenkamp, T. Weymouth, Topological mapping for mobile robots using a combination of sonar and vision sensing, in: *Proceedings of AAAI-94*, Seattle, WA, 1994.
- [30] D.C. Krakauer, Simple connectionist models of spatial memory in bees, *Journal of Theoretical Biology* 172 (1995) 149–160.
- [31] A. Kühn, *Die Orientierung der Tiere im Raum*, Gustav Fischer, Jena, 1919.
- [32] B. Kuipers, A hierarchy of qualitative representations for space, in: *Working Papers of the Tenth International Workshop on Qualitative Reasoning about Physical Systems (QR-96)*, AAAI Press, Menlo Park, CA, 1996.
- [33] B.J. Kuipers, Y. Byun, A robust, qualitative method for robot spatial learning, in: *Proceedings of AAAI-88*, St. Paul/Minneapolis, MN, 1988, pp. 774–779.
- [34] B.J. Kuipers, Y. Byun, A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations, *Robotics and Autonomous Systems* 9 (1991) 47–63.

- [35] D. Lambrinos, M. Maris, H. Kobayashi, T. Labhart, R. Pfeifer, R. Wehner, An autonomous agent navigating with a polarized light compass, *Adaptive Behavior* 6 (1997) 131–161.
- [36] B. Leonard, B.L. McNaughton, Rat: Conceptual, behavioural and neurophysiological perspectives, in: R.P. Kesner, D.S. Olton (Eds.), *Neurobiology of Comparative Cognition*, Lawrence Erlbaum, New York, 1995, Chapter 13.
- [37] T.S. Levitt, D.T. Lawton, Qualitative navigation for mobile robots, *Artificial Intelligence* 44 (1990) 305–360.
- [38] I. Lieblch, M.A. Arbib, Multiple representations of space underlying behavior, *Behavioral and Brain Sciences* 5 (1982) 627–659.
- [39] H.H. Lund, B. Webb, A robot attracted to the cricket *Gryllus bimaculatus*, in: P. Husbands, I. Harvey (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, MIT Press, Cambridge, MA, 1997.
- [40] H. Mallot, H. Bülthoff, P. Georg, B. Schölkopf, K. Yasuhara, View-based cognitive map learning by an autonomous robot, in: F. Fogelman-Soulié, P. Gallinari (Eds.), *Proceedings of ICANN'95—International Conference on Artificial Neural Networks*, Vol. II, EC2, Nanterre, France, 1995, pp. 381–386.
- [41] H.A. Mallot, Behavior-oriented approaches to cognition: Theoretical perspectives, *Theory in Biosciences* 116 (1997) 192–220.
- [42] D. Marr, Simple memory: A theory of archicortex, *Philosophical Transactions of the Royal Society of London B* 176 (1971) 23–81.
- [43] M.J. Mataric, Navigating with a rat brain: A neurobiologically-inspired model for robot spatial representation, in: J.A. Meyer, S.W. Wilson (Eds.), *From Animals to Animats*, MIT Press, Cambridge, MA, 1991.
- [44] P.J. McKerrow, *Introduction to Robotics*, Addison-Wesley, New York, 1991.
- [45] B.L. McNaughton, C.A. Barnes, J.L. Gerrard, K. Gothard, M.W. Jung, J.J. Knierim, H. Kudrimoti, Y. Qin, W.E. Skaggs, M. Suster, K.L. Weaver, Deciphering the hippocampal polyglot: The hippocampus as a path integration system, *Journal of Experimental Biology* 199 (1996) 173–185.
- [46] B.L. McNaughton, L. Nadel, Hebb-Marr networks and the neurobiological representation of action in space, in: M.A. Gluck, D.E. Rumelhart (Eds.), *Neuroscience and Connectionist Theory*, Lawrence Erlbaum, Hillsdale, NJ, 1990, pp. 1–64.
- [47] R. Möller, D. Lambrinos, R. Pfeifer, T. Labhart, R. Wehner, Modeling ant navigation with an autonomous agent, in: R. Pfeifer, B. Blumberg, J.A. Meyer, S.W. Wilson (Eds.), *From Animals to Animats 5*, Proceedings of SAB'98, MIT Press, Cambridge, MA, 1998, pp. 185–194.
- [48] R.U. Muller, M. Stead, J. Pach, The hippocampus as a cognitive graph, *Journal of General Physiology* 107 (1996) 663–694.
- [49] K. Nakayama, J.M. Loomis, Optical velocity patterns, velocity-sensitive neurons, and space perception: a hypothesis, *Perception* 3 (1974) 63–80.
- [50] R.C. Nelson, Visual homing using an associative memory, *Biological Cybernetics* 65 (1991) 281–291.
- [51] H. Neven, G. Schöner, Dynamics parametrically controlled by image correlations organize robot navigation, *Biological Cybernetics* 75 (1996) 293–307.
- [52] J. O'Keefe, N. Burgess, Geometric determinants of the place fields of hippocampal neurones, *Nature* 381 (1996) 425–428.
- [53] J. O'Keefe, J. Dostrovsky, The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat, *Brain Research* 34 (1971) 171–175.
- [54] J. O'Keefe, L. Nadel, *The Hippocampus as a Cognitive Map*, Clarendon Press, Oxford, UK, 1978.
- [55] C. Owen, U. Nehmzow, Landmark-based navigation for a mobile robot, in: R. Pfeifer, B. Blumberg, J.-A. Meyer, S.W. Wilson (Eds.), *From Animals to Animats 5*, Proceedings of SAB'98, MIT Press, Cambridge, MA, 1998, pp. 240–245.
- [56] F. Papi, *Animal Homing*, Chapman & Hall, London, 1992.
- [57] T. Poggio, W. Reichardt, A theory of the pattern induced flight orientation of the fly *Musca domestica*, *Kybernetik* 12 (1973) 185–203.
- [58] B. Poucet, Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms, *Psychological Review* 100 (1993) 163–182.
- [59] T.J. Prescott, Spatial representation for navigation in animats, *Adaptive Behaviour* 4 (1996) 85–123.
- [60] M. Recce, K.D. Harris, Memory of places: A navigational model in support of Marr's theory of hippocampal function, *Hippocampus* 6 (1996) 735–748.
- [61] T. Röfer, Controlling a robot with image-based homing, in: B. Krieg-Brückner, C. Herwig (Eds.), *Kognitive Robotik*, ZKW-Bericht 3/95, Center for Cognitive Sciences, Bremen, 1995.
- [62] T. Röfer, Controlling a wheelchair with image-based homing, *AISB Workshop on Spatial Reasoning in Mobile Robots and Animals*, Technical Report, UMCS-97-4-1, Department of Computer Science, Manchester University, 1997.
- [63] B. Ronacher, R. Wehner, Desert ants *Cataglyphis fortis* use self-induced optic flow to measure distances travelled, *Journal of Computational Physiology A* 177 (1995) 21–27.
- [64] R. Rothkegel, K.F. Wender, S. Schumacher, Judging spatial relations from memory, in: C. Freksa, C. Habel, K.F. Wender (Eds.), *Spatial Cognition*, Lecture Notes in Artificial Intelligence, Vol. 1404, Springer, Berlin, 1998, pp. 79–105.
- [65] A. Samsonovich, B.L. McNaughton, Path integration and cognitive mapping in a continuous attractor neural network model, *Journal of Neuroscience* 17 (1997) 5900–5920.
- [66] J. Santos-Victor, G. Sandini, F. Curotto, S. Garibaldi, Divergent stereo for robot navigation: A step forward to a robotic bee, *International Journal of Computer Vision* 14 (1995) 159–177.
- [67] B. Schölkopf, H.A. Mallot, View-based cognitive mapping and path planning, *Adaptive Behavior* 3 (1995) 311–348.
- [68] T. Sharpe, B. Webb, Simulated and situated models of chemical trail following in ants, in: R. Pfeifer, B. Blumberg, J.-A. Meyer, S.W. Wilson (Eds.), *From Animals to Animats 5*, Proceedings of SAB'98, MIT Press, Cambridge, MA, 1998, pp. 195–204.
- [69] H.H. Shufeldt, G.D. Dunlap, *Piloting and Dead Reckoning*, Naval Institute Press, Annapolis, MD, 1991.

- [70] M.V. Srinivasan, An image-interpolation technique for the computation of optic flow and egomotion, *Biological Cybernetics* 71 (1994) 401–415.
- [71] M.V. Srinivasan, M. Lehrer, W.H. Kirchner, S.W. Zhang, Range perception through apparent image speed in freely-flying honeybees, *Visual Neuroscience* 6 (1991) 519–535.
- [72] M.V. Srinivasan, S.W. Zhang, M. Lehrer, T.S. Collett, Honeybee navigation en route to the goal: Visual flight control and odometry, *Journal of Experimental Biology* 199 (1996) 237–244.
- [73] R.J. Sutherland, J.W. Rudy, Configural association theory: The role of the hippocampal formation in learning, memory, and amnesia, *Psychobiology* 17 (1989) 129–144.
- [74] S. Thrun, Learning metric-topological maps for indoor mobile robot navigation, *Artificial Intelligence* 99 (1997) 21–71.
- [75] E.C. Tolman, B.F. Ritchie, D. Kalish, Studies in spatial learning. I. Orientation and the short-cut, *Journal of Experimental Psychology* 36 (1946) 13–24.
- [76] O. Trullier, S.I. Wiener, A. Berthoz, J.-A. Meyer, Biologically based artificial navigation systems: Review and prospects, *Progress in Neurobiology* 51 (1997) 483–544.
- [77] B. Webb, Using robots to model animals: A cricket test, *Robotics and Autonomous Systems* 16 (1995) 117–134.
- [78] K. Weber, S. Venkatesh, M.V. Srinivasan, Insect inspired behaviours for the autonomous control of mobile robots, in: M.V. Srinivasan, S. Venkatesh (Eds.), *From Living Eyes to Seeing Machines*, Oxford University Press, Oxford, 1997, pp. 226–248.
- [79] R. Wehner, The polarization vision project: Championing organismic biology, in: K. Schildberger, N. Elsner (Eds.), *Neural Basis of Behavioural Adaptation*, G. Fischer, Stuttgart, 1994, pp. 103–143.
- [80] R. Wehner, B. Michel, P. Antonsen, Visual navigation in insects: Coupling of egocentric and geocentric information, *Journal of Experimental Biology* 199 (1996) 129–140.
- [81] R. Wehner, M.V. Srinivasan, Searching behaviour of desert ants, genus *Cataglyphis* (Formicidae, Hymenoptera), *Journal of Computational Physiology* 142 (1981) 315–338.

- [82] S. Werner, B. Krieg-Brückner, H.A. Mallot, C. Freksa, Spatial cognition: The role of landmark, route, and survey knowledge in human and robot navigation, in: W. Brauer (Ed.), *Informatik'97. Informatik als Innovationsmotor*, 27. Jahrestagung Gesellschaft für Informatik, Springer, Berlin, 1997.



Matthias O. Franz graduated with an M.Sc. in Atmospheric Sciences from SUNY at Stony Brook, NY, in 1994. In 1995, he received a Diploma in Physics from the Eberhard-Karls-Universität, Tübingen, Germany, where he also finished his Ph.D. in Computer Science in 1998. His thesis research on minimalistic visual navigation strategies was done at the Max-Planck-Institut für biologische Kybernetik, Tübingen. He is currently working as a researcher at DaimlerChrysler Research & Technology at Ulm, Germany. His scientific interests include optic flow, autonomous navigation, and natural image statistics.



Hanspeter A. Mallot received his Ph.D. in 1986 from the University of Mainz, Germany, where he studied Biology and Mathematics. From 1986 to 1987 he was a Postdoctoral Fellow at the Center for Biological Information Processing of the Massachusetts Institute of Technology. In 1993, he received his Habilitation for “Biological Information Processing” from the Ruhr-Universität-Bochum. Since 1993, he is a Research Scientist at the Max-Planck-Institute for Biological Cybernetics in Tübingen, Germany. Hanspeter A. Mallot is editorial board member of the journals *Biological Cybernetics*, *Neural Networks*, and *Kognitionswissenschaft*. From 1997–1999, he served as President of the European Neural Network Society, ENNS.