

Modelling Navigational Knowledge by Route Graphs

Steffen Werner¹, Bernd Krieg-Brückner², and Theo Herrmann³

¹ Institute of Psychology, University of Göttingen, Gosslerstr. 14, D-37073 Göttingen
swerner@uni-goettingen.de, www.uni-goettingen.de/~sppraum

² Bremen Institute of Safe Systems, University of Bremen, PBox 330440, D-28334 Bremen
bkb@Informatik.Uni-Bremen.DE, www.uni-bremen.de/~sppraum

³ Institute of Psychology, University of Mannheim, Schloss–E0, D-68131 Mannheim
fprg@rumms.uni-mannheim.de

Abstract. Navigation has always been an interdisciplinary topic of research, because mobile agents of different types are inevitably faced with similar navigational problems. Therefore, human navigation can readily be compared to navigation in other biological organisms or in artificial mobile agents like autonomous robots. One such navigational strategy, route-based navigation, in which an agent moves from one location to another by following a particular route, is the focus of this paper. Drawing on the research from cognitive psychology and linguistics, biology, and robotics, we present a simple, abstract formalism to express the key concepts of route-based navigation in a common scientific language. Starting with the distinction of places and route segments, we develop the notion of a route graph, which can serve as the basis for complex navigational knowledge. Implications and constraints of the model are discussed along the way, together with examples of different instantiations of parts of the model in different mobile agents. By providing this common conceptual framework, we hope to advance the interdisciplinary discussion of spatial navigation.

1 Introduction

Our daily life is filled with activities travelling from one location to another. Every morning, we walk from the bedroom to the bathroom, from there to the kitchen, etc. Later, we drive from home to work or take the bus into town. We stroll through the city we live in or drive to places far away to visit friends. All these activities share one important component: navigation in a spatial environment.

The navigational abilities of humans, and of other organisms, always have been an interdisciplinary topic of research. This is due to the fact that problems of navigation, in their general form, similarly apply to very different mobile agents in very different environments. The desert ant, foraging for food under the desert sun, has to return home after it has found a dead insect (Wehner & Menzel, 1990) as does the hamster

after pouching food (Etienne, Maurer, Georgakopoulos, & Griffin, 1999). Pigeons, like other birds, take extensive flights away from their nest but still return home safely, while migrating birds fly for thousands of miles to reach their winter quarters (Wiltschkow & Wiltschkow, 1999).

Human navigation ranges from our daily way to work to global travel. Unlike most other organisms, humans have also developed navigational tools to complement or enhance their navigational abilities. Maps, compasses, sextants, and global positioning devices (GPS) help to locate oneself in an environment or to plan one's travel. By using language, navigational information can be communicated to others, e.g. by giving route directions. Using remotely controlled objects, the space in which humans have to navigate has been greatly extended. Navigation today might entail directing minuscule medical instruments within the human body or controlling a remote vehicle on a different planet.

Artificial autonomous mobile agents, such as robots, also have to deal with similar navigational problems. Although relying on a host of different sensory and motor mechanisms, the abstract tasks of planning a route and executing the right actions along the way are nearly the same (e.g. Werner *et al.*, 1997). It is thus not surprising that in this field, disciplines such as biology, psychology, linguistics, artificial intelligence, and geography jointly contribute to our understanding of the basic processes underlying spatial navigation.

In this paper, we wish to outline a simple model to describe how mobile agents can acquire spatial knowledge to successfully navigate in an environment. Our approach is aimed at bringing together the basic ideas and models from different disciplines and providing a common platform to express concepts relevant for navigation and related types of spatial knowledge. Our main interest lies in so-called *route-based navigation* because it has proven to be an important navigational method for humans and other organisms, and in addition constitutes a common method employed in robot navigation. Of course this approach is not new, and many aspects of the proposed model can be found in a number of predecessors. Before we will describe the model, a brief overview of similar models from different disciplines will therefore set the stage.

2 Route-Based Navigational Knowledge for Different Agents

The number of particular navigational strategies employed by different agents is quite large. However, only a few basic kinds of navigational strategies can be distinguished at an abstract level (e.g., Allen, 1999). Agents can follow a marked trail, as in the case of Theseus following Ariadne's thread to the exit of the labyrinth, or use a gradient ascent or descent strategy, for example when moving in the direction increasing the intensity of a specific smell or following the magnetic field of the earth (Wiltschkow & Wiltschkow, 1999). Alternatively, agents can compute their position in relation to a starting point or goal location by integrating all their movements in space as in path-integration (Loomis *et al.*, 1999). In this case, the direction and distance of the home-base or the location of the goal are continually updated. Some other agent might be able to compute its position in terms of a global co-ordinate system, allowing the

computation of the direction and distance of any goal specified in the same reference system. Other navigational means might entail systematic search strategies, such as a grid search to find a shipwreck in the ocean or meandering to locate one's nest for desert ants (Wehner *et al.*, 1996).

In this paper we will mainly deal with route-based navigation, in which an agent moves from one location to another by following a particular route, i.e. a sequence of different places and means to get from one place to the next (an exact terminology will be developed later). Other strategies, such as the ones mentioned above, will be discussed only as they relate to route-based navigation. This, of course, does not imply that these other strategies are in any way inferior to route-based methods as general navigational strategies. However, route-based navigation plays a dominant role in human and robot navigation (and probably in animal navigation, too) and is thus the main focus of this paper.

2.1 Route-Based Navigation in Humans

Routes are a concept commonly encountered when dealing with human spatial navigation. When asking another person for directions to a goal, the usual response consists of a route description, sometimes accompanied by a sketch-map (Habel, 1988; Tversky & Lee, 1998). When planning a trip from one point to another, a number of intermediate points are usually identified and the total trip is broken down into a number of different segments, the endpoint of one segment also serving as the start of the next. The routes we travel are often physically defined in our environment, e.g. by the flow of a river when travelling by boat, or by the highway system, determining the sequence of cities and exits when driving in a particular direction.

At least since the classic paper by Siegel and White (1975), drawing heavily on earlier work by Shemyakin (1962), the distinction of landmark knowledge, route knowledge, and survey knowledge has played a prominent role in cognitive psychology. In their view, route knowledge is an integral part in the acquisition of spatial information. Route knowledge consists of a sequence of decisions, which are triggered by the perception of landmarks identifying a particular location (for the definition of landmarks see for example Lynch, 1960)¹. Simple learning paradigms, such as paired-associate learning, can account for route learning. Learning is thus “organised around the nodes of the decision system, the landmarks” (p. 29), while learning between landmarks is incidental and largely irrelevant. With more experience, however, information about different routes is integrated into a network-like structure and distance information is added. This integration of spatial knowledge ultimately leads to survey knowledge, which contains configurational information, such as regions or the geometrical shape of the layout of landmarks.

The model of Siegel and White (1975) has been criticised for its rigid developmental sequence of spatial knowledge acquisition, both for adult spatial cognition and

¹ We will later distinguish different uses of landmarks, as global landmarks or local routemarks, see Section 3.1.2.

children's development of spatial knowledge (e.g., McDonald & Pellegrino, 1993; Montello, 1998). However, the concept of route knowledge has stimulated a number of research questions.

One way to investigate route knowledge is to conduct linguistic analyses of route descriptions. In one study, for example, Habel (1988) asked 73 participants to describe a route through the town of Trier. As his analyses show, more than two-thirds of the descriptions used landmarks together with a directional change at that landmark to describe the route. Results from Tversky & Lee (1998) suggest that similar elements are used when giving a verbal route description and sketching a map depicting the route. This indicates the important role of landmark and directional-change information for route navigation.

Route descriptions necessarily consist of sequences of verbal statements, e.g., about landmarks and direction changes. However, route knowledge seems to be bound to a sequential order of places or objects even outside the realm of verbal descriptions. As Herrmann *et al.* (1995) were able to show, retrieval of information about objects along a route is easier in the direction of the route than in the opposite direction. If an object is briefly presented as a cue it is thus easier to recognise the name of the subsequent object than the preceding object on the route. This so-called *route direction effect* does not occur if the objects are just *displayed* in the same sequence as when travelling the path or when the optical flow is reversed while learning the route (Herrmann *et al.*, 1995; Schweizer *et al.*, 1998). These results indicate that spatial relations are encoded in a direction-specific way, and that the relation between two objects or places A and B is not necessarily the same as the relation between B and A.

In light of the distinction between route and survey knowledge, the question of distance estimation between two points on a route has also been addressed in a number of publications (Rothkegel, Wender, & Schumacher, 1998; Sadalla, Staplin, & Burroughs, 1979; Thorndyke & Hayes-Roth, 1982). Two different distance estimations are usually compared: route-distance vs. Euclidean distance estimation. In their study, Rothkegel *et al.* analysed the time it took participants to estimate the route-distance between different objects. Their results suggest that route-distance estimation times increase with increasing number of intervening objects along the route, while the length of the path itself only marginally determines the estimation time. They interpret their results as indicating that only route-distances between neighbouring objects (or places) are explicitly stored in memory. In contrast, route-distances encompassing more than two objects are mentally computed by adding the different legs from one object to another along the route.

In contrast, Rothkegel *et al.* were able to show in a different experiment that the required time for estimates of the Euclidean distance between two objects learned on a route increases with increasing Euclidean distance. In this case, the estimation time was largely independent of the route-distance between the two objects. Rothkegel *et al.* see this potential difference between the estimation times for route and Euclidean distance as one way to distinguish between route and survey knowledge. Route-distances are computed by integrating the length of individual legs along the route and are thus affected by the number of intervening objects. Survey knowledge should

be largely independent of the route(s) experienced within an area and Euclidean distance estimations should therefore not be affected by the number of objects between the target pair. Assuming a mental scan process to estimate Euclidean distances in survey knowledge, an increase in the required time with increasing Euclidean distances fits well with this notion of survey knowledge.

In summary, the concept of route knowledge for human navigation has led to a significant body of research. Routes are commonly seen as a sequence of places which act as decision points for further travel. Route knowledge is direction-specific. Route knowledge in human memory also seems to consist of local information, e.g. the distance between a place and a successor, whereas integrated spatial knowledge of configurations or arbitrary Euclidean distances is attributed to survey knowledge.

In the following section we will briefly outline the role of route-based navigation in animals and how it compares to research on humans.

2.2 Route-Based Navigation in Animals

In the vast literature on animal navigation and spatial competences of different organisms, the role of route-based navigation has been a topic of research in many different species (for an overview see Able, 1996). Research in this area focuses on two main areas. First, the spatial behaviours of foraging animals, such as rodents or insects, enable researchers to focus on navigational abilities at time scales of minutes to hours, covering long distances with respect to the organism's body size. Second, migrating behaviour in long-distance migrating birds or sea animals, such as sea-turtles or fish, allows researchers to investigate navigational skills at a much larger scale of both time and space. In this short overview we will mainly focus on the first part, excluding migration for reasons of brevity.

As in the discussion on human spatial cognition, the issue whether navigation relies on a spatial cognitive map or whether it is based on different principles has also been addressed in the animal literature. A prominent model for the development of such a spatial cognitive map, explicitly representing the spatial relations between a large number of locations within the same reference system, has been put forward by Poucet (1993). The basic units in his hierarchical model are place representations. A place is defined by a collection of inter-linked views that an animal experiences through movement, especially rotations, within a small part of space. This set of views allows an animal to recognise the same place from different vantage points, e.g., when approaching a place from different directions. Connections between two places are encoded in a local reference system in terms of the direction and distance from one place to another. These multiple local reference systems make it impossible to compare directional information from one place directly with directional information from another place. The set of places within an animal's environment, which are connected by this vector information, constitutes a local chart. It is important to note that not all connections between places have to be available, nor does knowledge of the direction from A to B imply that the reverse direction from B to A is known. Besides this multiple-point reference system connecting some of the places, a second,

so-called topological network yields information about the connectivity and order of the places. In Poucet's model both kinds of information, vectorial and topological, work together to enable an organism to find novel routes and shortcuts. The final step in the model consists of computing an overall reference direction. This global, environment-centred reference direction is assumed to correspond to the preferred axes of movement. It allows an organism to integrate the directional information from different local reference systems with respect to a common reference system, enabling the computation of previously unknown spatial relations and thus enhancing an organism's navigational precision and flexibility.

The model of Poucet (1993) captures a number of points relevant to route-based navigation. First, the role of places for the integration of navigational knowledge is emphasized. Second, the model clearly distinguishes between a network of connected places with multiple location-dependent reference systems and topological information on the one hand, and spatial information integrated within a common, global reference system on the other. Route-following behaviour and the finding of new routes is assigned to the lower level of integration of spatial information. Third, the gradual construction of an increasing number of connections between places, which takes place through active spatial exploration by the animal, can result in highly integrated spatial cognitive maps.

One goal of Poucet's (1993) model is the analysis of how integrated spatial knowledge might develop out of locally available information, mainly focussing on mammalian spatial navigation. However, for many species it can be demonstrated that no such maps are necessary to explain their navigational skills. As Wehner *et al.* (1996) point out, an animal's brain did not necessarily "evolve ... to reconstruct a full representation of the three-dimensional world, but to find particular solutions to particular problems within that world" (p. 138). They argue against a cognitive map metaphor and propose instead that a close analysis of the computational strategies used by a particular brain in a particular situation might prove more useful for our understanding of spatial behaviour.

In their elaborate model of spatial navigation in the desert ant (*Cataglyphis*), Wehner and colleagues were able to identify a number of important components of the ant's navigational techniques (Wehner & Menzel, 1990; Wehner *et al.*, 1996). The most striking ability of the desert ant rests in its capacity to use *path-integration* to integrate all directional changes and distances travelled into a homing vector, representing the exact direction and distance of the nest entrance (from which its travel originated). To avoid cumulation of directional errors, the ant is equipped with a simple but efficient means of computing a global reference direction from an innate skylight compass. This enables the ant to correctly identify the direction it is travelling at each point in time. The importance of the homing vector becomes evident when the ant is transferred to a different environment before it is able to return to the nest. In this case, the ant ignores the repositioning and travels exactly to the location where its nest would have been if it were still in the old environment, clearly indicating that the homing vector drives the ant's spatial behaviour. More precisely, the behaviour is governed by two vectors: the homing vector and the target vector, both re-adjusted by path-integration at any moment in time. For a honeybee, the target vector may have

been communicated prior to the foraging flight by a scout sister using the honeybee language (von Frisch, 1967).

Besides the dominant role of path-integration for *Cataglyphis*-navigation, the use of *landmark information* has also been demonstrated. One piece of evidence rests on the fact that individual desert ants show consistent trajectories from the nests to the feeding site, presumably by using landmark information on the way to guide their course. In addition, if an ant has reached the nest entrance, thus resetting its homing vector to zero, and is subsequently passively transported and released at the former feeding point, it returns to the nest along the same route as before (Wehner *et al.*, 1996). This indicates that features of the environment have been learned and can be used by the ant to find its way back home. There is also evidence that insects follow routes, e.g. a honeybee systematically flying from one food source to the next following the same route as on the previous flight.

One mechanism by which insects, such as desert ants or bees, can use landmark information has been proposed by Cartwright and Collett (1983). In their snapshot model, a 360° view of the environment is encoded by the animal at a particular location, e.g., the nest entrance or a salient part of a route. Using simple techniques, this information can then be used to steer towards the original position from which the snapshot was taken whenever the animal is close to it. By physically turning the landmark configuration it can be demonstrated that these snapshots are encoded with respect to a particular direction. If the configuration is turned, the animals are less likely to recognise it and thus don't use it for navigation (Judd, Dale, & Collett, 1999; Wehner *et al.*, 1996). Honeybees, for example, make a few circular flights around their nest before leaving to learn this landmark configuration; they do so in a well-defined fixed orientation, the same as when returning, to avoid the influence of rotation on the 360° view. Interestingly, a landmark strategy is not always used by the animal. The desert ant, for example, only reacts to the visual landmark configuration around the nest entrance when its homing vector indicates that it should be close to it. Expectancies, therefore, seem to play an important role in landmark usage. This simple mechanism to identify a previously visited location can also be coupled with additional navigational information. However, as Wehner *et al.* (1996) show, such route following is less accurate and slower when the homing vector is missing than for the original route home including the homing vector.

Similar processes like those proposed above have also been investigated in humans (Gillner & Mallot, 1998). Gillner and Mallot's view-graph model assumes that particular views are represented in a network of views together with the action required to get from one view to another (see also section 2.3).

In conclusion, research on route knowledge in non-human organisms has identified a number of detailed mechanisms how landmarks can be used to guide an organism's travel. Although there are differences in the importance and exact role of route knowledge for each species, the main topics are the same. Landmarks or places have to be recognised, expectancies for certain places can be built up, and navigational decisions have to be triggered at the correct point of the route. In addition, research on detailed navigational mechanisms in different animals, in particular non-vertebrates,

provides an excellent means to explain complex human spatial behaviour by comparing it to simple mechanisms that can be found in other species.

In the last section of this chapter we will now focus on the navigational strategies employed in artificial mobile agents, e.g. autonomous robots.

2.3 Route-Based Navigation in Robots

The set of navigational strategies found in artificial moving agents mirrors the complexity of navigational strategies employed by biological organisms. In many instances the close resemblance between biological and artificial navigation is not a mere coincidence. Artificial navigation often mimics evolutionary proven strategies in an attempt to build robust technologies. In some instances, artificial agents are even specifically designed to test biological models of navigation, e.g. by Möller *et al.* (1998).

There are, however, a number of important differences between biological and artificial agents. First, modern technology makes a wide range of very accurate sensors possible, such as laser range finders, radar and ultrasonic sensors, global positioning systems, etc., providing information about the environment or the position of an agent. Unlike biological organisms, who might possess a few highly developed sensory systems for particular sources of information (e.g., the innate compass for the desert ant), the designer of a robot system is free to use an arbitrary combination of different sensors. In addition, robots sometimes are equipped from the start with a precise representation of their spatial environment (e.g., a map). In many cases, the environment is also specifically designed to match the robots navigational or sensory abilities, e.g. by placing markers at important points of a route or using well demarcated paths along which the robot travels. The basic tasks of spatial navigation, however, remain the same for both robots and biological organisms. According to Trullier *et al.* (1997), four different, broad categories to classify natural and artificial navigation can be distinguished: guidance, place recognition – triggered response, topological navigation, and metrical navigation.

Guidance is mainly concerned with directly leading an agent by external cues – either by following a particular gradient or moving to match the current sensory image with a stored image of the target or of the surroundings. In all these cases, the agent tries to locally maximise a predefined criterion without knowledge of spatial relations in the environment or about its own position. A robot could, for example, try to follow a wall of a corridor by keeping a constant distance to it (Krieg-Brückner *et al.*, 1998). For vehicle control, the guidance system could be designed to steer the vehicle parallel to the side of the road. A different robot could try to increase the match of its current visual image of the surroundings with a target image. This latter approach is directly modelled after the snapshot model of insect navigation, as discussed above (Röfer, 1998; Franz *et al.*, 1997).

Place Recognition – Triggered Response. For place recognition based strategies, complex spatial behaviours are triggered at distinct points in space. Once the correct place is recognised, the associated action (e.g., movement in a particular direction or guided behaviour) will lead to complex trajectories. The main problem of this strategy obviously consists in the correct identification of a place. One approach uses the perceptual input (the current state of the sensors of a robot) as the defining characteristic of a place (e.g., Schölkopf & Mallott, 1995). Another method is to detect landmarks (Krieg-Brückner *et al.*, 1998). Of course both might lead to problems if the same sensory state or landmark can be evoked at different spatial locations (*perceptual aliasing*; Duckett & Nehmzow, 1997). A second problem is sensory noise. Even with highly accurate sensors it is unlikely that a second visit to the identical location will yield exactly the same sensory state.

Topological Navigation describes navigation based on topological networks and is thus a more flexible extension of place-triggered navigation. The basic elements of this type of network are places and the connections between these places. In this case the agent possesses knowledge how to get from one place to a second place which is connected to it. The agent's spatial knowledge is confined, however, to the topological network. New places not included in the network, or new connections between two places cannot be found by navigating on the basis of a topological network. One example for a topological network is Schölkopf and Mallott's (1995) view graph model. In this type of graph, each view (i.e. a location coupled with a viewing direction) is connected to a second view by an action. The view graph thus consists of all transitions from one view to the next, and the actions to achieve these transitions (Franz *et al.*, 1997). As before, no further spatial knowledge about the directions or distances between different places is necessary for navigation.

Metrical Navigation. Unlike the last two approaches, which divide space in a small number of distinct places and the space in between, metrical navigation does not require such a distinction in principle. The metric most frequently used is Euclidean, thus distances and angles are well defined and can be used to drive spatial navigation. Pre-existing maps, which specify the metrical relations between objects in the environment of the agent, are often supplied directly or are autonomously constructed by triangulation and integration of sensory information. A coarse version of metrical navigation can be seen in world representations such as *occupancy grids*. Unlike navigation based on topological networks, the agent can determine its position in space and its spatial relation to each other object within the same metrical space for each position. However, knowledge of its position does not, as in topological navigation, automatically trigger the correct action. Instead, the action usually has to be computed for the current situation.

The Spatial Semantic Hierarchy Model. Kuiper's (1998) Spatial Semantic Hierarchy (SSH) model nicely illustrates the interplay of the different navigational strategies listed above. At the lowest level, *sensorimotor systems* of an agent provide information about its current state, which can be used at the *control level* to reach so-called

distinctive states. A distinctive state is characterised by a local maximum of some *distinctiveness measure* which is computed from the change of the agent's sensorimotor states, but also more complex operations as detecting a landmark (Krieg-Brückner *et al.*, 1998) or recognising a kink in the robot's trajectory (Musto *et al.*, 1999; Röfer, 1999) can be employed to generate such a state. The measure thus provides the agent with a discrete set of distinctive states which can be reached by simple guidance techniques, such as gradient hill-climbing or following a particular trajectory from one distinct state to the next. At the causal level, these distinct states are treated as *places*, from which another place can be reached through a particular *action*. Combining different associations between two places through an action, the agent can build a *topological* representation, consisting of *places*, *paths*, and *regions*. At this level, route-finding problems can be solved. Finally, at a *metrical level*, the topological representation is enhanced with metrical properties, such as distance or direction between different places. This hierarchical strategy has been successfully employed in different simulated and physical robots (Kuipers & Buyn, 1991). However, it does not seem clear that the metrical level is necessarily acquired last or acquired at all. As Gutmann and Nebel (1997) have shown, knowledge at the topological level might be more important and more difficult to acquire for navigational tasks than a metrical representation (for a real-life example of the acquisition of such a topological map see Fu, Hammond & Swain, 1996).

Route based navigation can be found at different levels in the SSH model of Kuipers (1998). Starting with the causal level, different places are connected through spatial actions that the agent can perform. At the topological level, the flexibility of continuing from one place to another is greatly enhanced. The level of metrical spatial information, on the other hand, only enhances existing knowledge about spatial relations between different places, without adding significantly to route following behaviours. Eventually, however, metrical information can be used to expand topological networks by exploring new connections or defining new places.

3 The Model

In the following section, we will present a simple model describing key elements for route based navigation as part of an agents' general spatial knowledge. Drawing on the topics discussed above for human, animal, and artificial navigation, we will present basic concepts relevant for route based navigation. Our main goal consists of sketching a potential formalism expressing the key notions of route based navigation without restricting the formalism to a particular implementation, agent, domain, or discipline. Therefore, instead of focussing on a particular implementation or a specialised purpose, we will develop the general concepts step-by-step and discuss the implications of different possible assumptions at each stage, leaving open many details of the model for the reader to fill in. We will emphasise the core assumptions and the implications that they carry at each stage. This attempt of a flexible and open model is mirrored by our use of the term "agent" to describe any kind of autonomous mobile entity, whether it is a biological organism, a machine, or a virtual entity ex-

isting only in a simulation. It is our experience that the introduction of abstract concepts in this way aids the communication between disciplines, both at the empirical and theoretical level. In addition, new empirical questions will arise from a better understanding of the relevant concepts for route-navigation.

3.1 Single Routes

A **Route** is a concatenation of directed Route Segments from one Place to another. As we will see below, a Place is a tactical decision point where to continue, i.e. about the choice of the next Route Segment. This dichotomy of Places and Route Segments is consistent with most approaches taken in the literature. As an example of a simple route, i.e. a linear chain of directed Route Segments, consider the commuter train line S6 in Munich from Erding to Tutzing. Routes may also be cyclic, as for example part of the underground in London (cf. Fig. 1). Although we will sometimes refer to entities in the physical world (such as the train lines or the underground in this example), our interest lies in the knowledge that an agent has to possess to successfully navigate in space. Concepts such as a Place thus have to be interpreted in terms of an agent's knowledge, not in terms of particular physical locations, even though they often go hand in hand.




Fig. 1. Simple Route (a) and Cyclic Route (b)

A **Path along a Route** is distinct from a Route. It embodies the dynamic usage of a Route or a contiguous part of it. An example would be a ride on the train S6 from Munich Main Station to Tutzing. A Path may be periodic (even infinite) on a cyclic Route (for a detailed discussion of formalisms to denote paths and spatial trajectories see Eschenbach et al., 2000).

A **Route Segment** consists of two Places, a Source and a Target, which are connected by a Course (Places are discussed further below). As is evident from much of the research on route knowledge, the Course information is not always reversible. The connection between Source and Target is thus directed. Being able to walk from A to B does not imply any knowledge about how to get back from B to A.

In between the Source and the Target, the Course contains information that allows a subject to follow the Route Segment. This may be a particular trajectory, e.g. by re-

playing a specific motor program, or following a gradient hill-climbing strategy, such as walking towards the bakery by the increasing intensity of the scent. An ant might follow its homing vector, trying to minimise the distance to its nest entrance while integrating its path. A robot, on the other hand, might use a simple following strategy, e.g. “wall following” (Krieg-Brückner, 1998; Krieg-Brückner *et al.*, 1998). It is important to note that using the abstract notion of a Course allows the model to be adapted to very different navigational strategies.

A Route Segment contains the necessary information to make a tactical decision at the Source how to enter this particular Route Segment as a continuation in a Path. This Entry to the Route Segment is non-trivial. Depending on the scenario, the specification of the Entry may differ substantially. For a human walking through a system of passages in a town, this may be the decision or action “next right” (cf. Fig. 2a), “second right” (Fig. 2b) or “right when facing the city hall” (Fig. 2c). When steering a boat, it might entail changing one’s course to 135° SE (Fig. 2d).




Fig. 2. Entries and Exits of Route Segments

In a similar way, the Exit of the Route Segment contains the necessary information to arrive at the target in a well-defined way. In the above examples, taking the “next path to the right” assumes that “next right” is unambiguously specified. However, this does imply that the traveller does not walk into an intersection backwards or sideways, but instead is walking forwards to avoid a conflict between a body-defined “left” and “right” and the “left” and “right” determined by the travel direction. Another example is the final Route Segment when docking a boat. One usually wants to end up parallel to the dock to avoid damage.

In some scenarios, some of the attributes may be superfluous: in a labyrinth, for example, Course and Exit may be empty, only the tactical decision when entering a Route Segment is necessary. Similarly, some information may be intentionally redundant. In some cases, either Entry or Exit, or both, may be used to model a well-defined position or bearing: when navigating in space, a (new) direction may be part of the Exit (e.g. the harbour entrance) or the Entry to a Route Segment (e.g. the compass bearing for a new course); in this model we prefer to associate the new direction to the Entry when it models a tactical decision related to the next rather than the previous segment.

Kinds of Route Segments. As is evident from the range of possibilities above, the particular information associated with Entry, Course or Exit depends on the context and grain of detail, i.e. the particular scenario one is trying to model. Let us call this the Kind of the Route Segment. Examples for common Kinds of Routes would be CommuterTrainLine, ShipRoute, FootPassage, CityRoad, Highway or Labyrinth, distinguished by different human task-related behaviour or navigational tactics involved (Krieg-Brückner, 1998). For other animals, common Kinds of Routes might entail UndergroundPassage, TreeBranches, or MigrationRoute. Route Kinds have to be consistent with Place Kinds, of course (compare also *Transfers* below); in general, the information associated with each component of a Route Segment has to be consistent with its Kind.

Places. Route Segments are linked by Places, e.g. concatenated to a simple Route: the Target of the preceding Route Segment has to be the same as the Source of the next. When constructing Route Segments incrementally, this is quite straightforward as we will continue from the Place we arrived at. For more complex Route Graphs, this is not so obvious, see *Place Integration* below.

Of course, the notion of a Place has different implications depending on the scenario at hand. In rats, for example, the concept of place representations has been used to describe the response pattern of single neurons which fire whenever the rat finds itself in the same position in its environment (O’Keefe & Nadel, 1978). However, this location is probably not absolutely defined, but instead in relation to environmental cues (such as distal landmarks or the geometry of the surroundings, etc.). If the rat knows its position with regard to these cues, however, the neuron’s response pattern can also be driven by egomotion, with no further environmental cues. This implies that a place is defined as a physical location with respect to an allocentric reference system, which in turn is updated by environmental features. Similarly, navigational charts for boating often refer to the relative position of salient landmarks to define the place of bearing changes.

In contrast, bees and ants do not only respond to the relative position of landmarks, but also to the absolute orientation of the configuration (Collet & Baron, 1994). Even more so, a hiker using a global positioning device will also rely on the global coordinates to determine his or her location. Two Places can therefore be identical in one situation and different in another, depending on the definition of a Place.

In another scenario, the dancing instructions “two steps forward, side-step, turn” might be interpreted as a Route across the dance floor, where each Place is defined solely in terms of the preceding Route Segment. In this example, there is no reference to the environment at all.

Reference System and Position. Quite importantly, each Place comes with its own local Reference System (RefSystem). Uniform global reference systems will be introduced in Section 3.4. The term reference system is meant to be very general. For example, in case of the desert ant it may refer to absolute compass bearings, derived from the direction of polarised light. It may also be a constellation of *landmarks* that can be seen in the distance (e.g. church squares, towers, celestial bodies such as the

sun), or a particular collection of *routemarks* which characterise the Place in a passage scenario (e.g. the city hall, the ice-cream parlour)². Probably the most basic RefSystem for a Place is the direction of travel in which one approaches that Place – which is often implied when giving directions.

Relative to its RefSystem, a Place is characterised by a RefPosition, i.e. a well-defined reference *position* in it and, in general, a *bearing* (cf. the crosses plus arrowheads in Fig. 3). As examples consider “in front of the city hall, facing the main entrance”, “on Königsplatz, leaving Luisenstrasse, coming from the direction of the main station”, or “in the centre of the market place, bearing North”. Other ways to indicate a RefPosition include the snapshot model (cf. Section 2.2), whereby a RefPosition is defined as the position where the image stored for that position matches the current percept, or simply by referring to a salient point.

The local RefSystem is not universally defined for each Place. Depending on the kind of RefSystem, Places may have different Kinds, consistent with the Kinds of Route Segments above. Thus the *same physical location* might have multiple RefSystems associated with it for different Routes; in our model, we distinguish these as *different Places* unless they are integrated to the same RefSystem (see Place Integration). For example, when using the direction of entry as a local RefSystem, a Place is automatically accompanied by a different RefSystem whenever it is approached from a different direction. This means, of course, that the interpretation of instructions such as “next left” depends on the preceding Route Segment.

In Figure 3, a few examples of Places and their corresponding RefPositions and RefSystems are depicted. As the two Figures 3b (“second right”) and 3c (“right when facing the city hall”) show, the same physical location might have different RefPositions and RefSystems even if the trajectory is very similar.




Fig. 3. Places with Individual Reference Systems

² The “same” routemark may be used for different purposes: to help in an exit approach to decide about arrival (e.g. “until the entrance to the harbour is reached”) and to define a proper exit position and bearing, to help a choice of continuation route and entry to it, or to help characterise the reference system of a place as part of a collection of routemarks.

3.2 Route Graphs

Union of Routes into Route Graphs. The representation of Places and Route Segments corresponds to the mathematical notion of a directed graph with a set of nodes (Places) and edges (pairs of Source and Target in Route Segments), both enriched with their particular attributes (there may be several edges between a given pair of nodes, in both directions). The union of different Routes into Route Graphs therefore corresponds directly to the union of the corresponding sets of nodes and sets of edges, respectively. This union of different Routes into a Route Graph is often non-trivial. In this section we will deal with the most prevalent problems that might occur when trying to integrate different Routes into a coherent whole.

In Fig. 4, two previously separate Routes are joined into one Route Graph, sharing a common sequence of Places and Route Segments (i.e. a sub-Route). As an example consider the commuter train lines S6 from Erding to Tutzing and S8 from Munich Airport to Pasing³. Of course, it is essential for the union of these two Routes that the identity of the Places and Route Segments of the two sub-Routes is noticed. In the case of two subway-stations with the same names, this might be fairly obvious. However, in the general case, the recognition of one Place as a known Place which was visited before might be much more demanding.




Fig. 4. Union of two separate Routes (a) into one Route Graph (b)

Place Integration. In the scenario of Fig. 3a and 3b, for example, two positions for two Routes are marked. In the current model, they thus signify two separate Places and are therefore two unrelated parts of two different Routes. The two Places might share a common RefSystem (e.g. a collection of route marks) but each has its own RefPosition. The way the union of Routes into a Route Graph was introduced above, these two Routes would be treated as separate, with no relationship between them. What is needed is a way to integrate Places which only differ in their respective RefSystem or RefPosition – and this is indeed a non-trivial activity.

³ It is a choice of modelling, whether the routes S6 and S8 are kept separate (and drawn side-by-side in different colours, as e.g. in the diagram of the Munich commuter train system), or whether the same track is regarded as an identical sub-route.

To integrate two Places, one thus has to come up with a common RefSystem and to agree on a common RefPosition. Choosing the same RefPosition “in front of the city hall” in Figures 3c and 3d, we are able to integrate the two Places.

Depending on the scenario, Place integration may be easy or require complicated updating of Route Segment attributes to conform with the new Place. Exits and Entries of the various Route Segments associated with a Place all refer to the same RefPosition in its RefSystem. For each Place that is affected by an integration, each Exit of every Route Segment leading towards it and each Entry of every Route Segment leading away from it has to be re-computed. In the example above, the Exit of the left-hand Route Segment in Fig. 3a and the Entry to the segment leading downwards have to be re-computed (as is depicted in Fig. 3d) relative to the new RefPosition.

Choice of Route Continuations in a Graph. In general, a common RefPosition and RefSystem is necessary to achieve the increased functionality of a Route Graph: after Place integration, we are able to continue with *either* of the two Route continuations.

Our own experience reveals that we often have a hard time with Place integration, having to “look around” to recognise a Place to be the same as one that was previously reached by another Route. This is especially obvious when trying to retrace ones Path through an unfamiliar environment, such as a forest, with new passages never seen before opening up when walking back. One reason for this problem might lie in the limited field of vision for humans, which dramatically changes the visual input depending on the walking direction. Other animals, such as the desert ant, have a 360° visual field, which makes them less vulnerable to this problem.

A second problem for Place integration are false matches (undesired aliases) for a Place. In modern buildings, for example, one intersection might appear to be the same as the intersection encountered just minutes ago, although they are in different parts of the building. Whereas humans can be confused in such a situation, the desert ant can usually rely on its well-developed path integration system, not to be fooled by the similar appearance.

3.3 Layers and Transfers

Modelling Aspects by Layers. As we have mentioned above, it is often necessary to distinguish between different Kinds of Routes when modelling different scenarios. One possible way to conceive of these different Kinds of Routes can be found in separate Layers. This approach is used, for example, in constructing cartographic maps of an area, where the information for water and power lines, for the sewer system, for roads, train tracks, and other significant entities are kept separate on different maps. Each layer represents a different aspect of the environment, in analogy to aspect maps (Berendt *et al.* 1998). In our model, each Layer represents a Route Graph of a particular Kind.

Munich main station, for example, would then correspond to different Places on different Layers. When arriving by train, the corresponding RefPosition might be “at the head of the tracks, facing the timetable”. If travelling via commuter train, it might

be “underground, in the centre of the platform in the direction of Pasing”. If travelling by foot, the entire train station might correspond to a Route Graph, instead of a single Place. To achieve appropriate links in each Layer, these Places have to be distinct.




Fig. 5. Transfer Routes (half-filled arrows)

Abstraction Hierarchies and Transfers between Layers. Obviously, modern day travel requires the use of a combination of different Kinds of Route Graphs at different levels of detail. When travelling from Hamburg to Munich, for example, one will probably have to walk towards the subway, ride the subway to the train station, take a fast train from Hamburg to Munich, get off the train in Munich, walk downstairs to the commuter train, take it for a few stops, maybe changing trains, and finally walk towards one’s final destination. Understandably, the level of detail and the kind of Route Graph needed for walking towards the commuter train is different from the one dealing with the train ride. So far, the model cannot deal with transitions between different kinds of such Route Graphs since we tacitly assumed that Source and Target nodes of a Route Segment are of the same Kind; such Route Graphs shall be called *homogeneous*. To allow transitions in heterogeneous Route Graphs, need a special type of Route Segment, whose Source and Target nodes can be of different Kinds. We will call these special Route Segments “Transfers”. To allow different complexities of Transfers, they can also consist of a whole Route or Route Graph.

In Fig. 5, such Transfers are depicted by half-filled arrows. Transfers are also directed (e.g., highway ramps). They may be one-to-one, one-to-many (e.g., Exit Transfers) or many-to-one (e.g. Entry Transfers). Transfers do not only occur when switching the mode of transportation or ego-motion, though. In some cases, place integration might not be possible even within the same Kind of Route Graph or Layer. In these cases, Transfers between Layers then also serve as refinements or abstraction projections. One such example is the set of entry and exit ramps of a highway intersection. In this case, the exit ramp taken in one direction is physically different from the entry ramp on the other side of the intersection. However, we still wish to think of the highway crossing as the same Place at a higher level of abstraction. In the current model, this would be modelled by a second, more abstract Layer. Both, the entry and the exit ramp would be connected by a Transfer to a single Place at this higher level of abstraction.

Uniform Reference System. The essential property of an Overview Graph, or Overview for brevity, is a uniform RefSystem from the observers RefPosition as far as the view allows. The RefSystem is therefore not completely “global” but restricted to the sub-graph of the original graph from which the Overview is constructed. Consequently, the Overview Layer requires a re-computation of the local RefSystems of all Places included in the Overview, based on the uniform RefSystem of the observer’s location, e.g. viewpoint. The RefPositions for each Place are not necessarily affected by this transformation (whereas the bearing often is).

The variety of possible RefSystems is again large. A special case can be seen in the introduction of a Euclidean space, such as \mathfrak{R}^2 or \mathfrak{R}^3 , which allows an approximation to common cartographic maps.

In the uniform RefSystem, the construction of new relations between Places becomes easily possible: new distances or angles can be denoted in pseudo-segments. Similarly, deviations and shortcuts can now be planned. Note that the relation between the original (sub-)graph and the Overview Graph can be maintained in different Layers, with appropriate Transfers for co-ordination.

4 Closing Remarks

Our main goal in this paper was to present a common, simple formalism to express the key notions of route based navigation as they occur in different disciplines. We hope that this endeavour will help researchers focussing on different research questions pertaining to spatial navigation to communicate their findings and theories more efficiently. Of course, we were not able to present thorough overviews of the navigational literature in all different areas. Instead, we tried to present a number of exemplary research questions and results from human, animal, and robot navigation to emphasise the close resemblance of navigational problems encountered in each field.

We are well aware that the general modelling framework presented above leaves much freedom to model the same scenario in different ways. There are often trade-offs, e.g. between complexity of generating a route graph vs. using it for navigation, between modelling by place integration or explicit transfers, or between several levels of granularity. It is not immediately obvious, for example, whether several trains on the same track (e.g. a fast train and a slow train between Hamburg and Munich) correspond to the same or different routes; this is rather a matter of abstraction. Similarly, a path along the route from Hamburg to Munich with or without changing trains of the same kind might make a difference; is it the same or two different routes?

Obviously, the model does not provide answers to the particular navigational strategy humans or other organisms use. However, the explicit definition of the basic elements and the preconditions for integrating local knowledge put forward in the model might help to specify the questions to ask. From a psychological point of view, for example, many questions remain open. To what extent are separate routes learned before they are integrated using place integration? What are the mechanisms that humans use to identify a place? How and when are the adjustments made to the local reference systems so that different routes can be combined to a route graph? What

kind of information is used to specify the course of a route segment, after a particular route segment has been entered? How is survey knowledge related to overviews constructed from sub-graphs with a common reference system? These questions and more, we feel, will benefit from a common language of route based navigation – thus enabling an interdisciplinary research effort.

Acknowledgements

This paper is the result of an interdisciplinary dialog within the context of the priority program on “Spatial Cognition” funded by the German Science Foundation. We gratefully acknowledge many fruitful discussions with our colleagues. We especially thank Thomas Röfer for his comments on the section on robot navigation, which is in part based on his thesis work, and Carola Eschenbach, Ute Schmidt, and an anonymous reviewer for their detailed comments on an earlier version of this paper.

References

- Able, K. P. (1996). Large-scale navigation. *The Journal of Exp. Biology*, 199, 1-2.
- Allen, G.L. (1999). Spatial abilities, cognitive maps, and wayfinding: Bases for individual differences in spatial cognition and behavior. In R. Golledge (ed.), *Wayfinding behavior* (46-80). Baltimore: Johns Hopkins.
- Berendt, B., Barkowsky, T., Freksa, C., & Kelter, S. (1998). Spatial representation with aspect maps. In C. Freksa, C. Habel, & K. F. Wender (Eds.), *Spatial cognition - An interdisciplinary approach to representing and processing spatial knowledge*. Berlin: Springer.
- Cartwright, B.A. & Collet, T.S. (1983). Landmark Learning in Bees. In *Journal of Comparative Physiology, A* 151, 521-543.
- Collet, T.S. & Baron, J. (1994). Biological compasses and the coordinate frame of landmark memories in honeybees. *Nature*, 368, 137-140.
- Duckett, T. & Nehmzow, U. (1997). Experiments in evidence based localisation for a mobile robot. In *Spatial Reasoning in Mobile Robots and Animals* (25-34). Manchester University. AISB-97 Workshop.
- Etienne, A.S., Maurer, R., Georgakopoulos, J. & Griffin, A. (1999). Dead reckoning (path integration), landmarks, and representation of space in a comparative perspective. In R. Golledge (ed.), *Wayfinding behavior* (197-228). Baltimore: Johns Hopkins.
- Eschenbach, C., Tschander, L., Habel, C., & Kulik, L. (2000). The specifications of paths. *This issue*.
- Franz, M. O., Schölkopf, B., Georg, P., Mallot, H. A., and Bühlhoff, H. H. (1997). Learning view graphs for robot navigation. In W. L. Johnson (ed.), *Proc. 1st Int. Conf. on Autonomous Agents* (138-147). New York. ACM Press.
- Fu, D.D., Hammond, K.J., & Swain, M.J. (1996). Navigation for everyday life. *Technical Report 96-03. Dept. of Computer Science, University of Chicago*.
- Gillner, S. & Mallot, H.A. (1998) Navigation and acquisition of spatial knowledge in a virtual maze. *Journal of Cognitive Neuroscience*, 10, 445-463.
- Golledge, R.G. (1999). Human wayfinding and cognitive maps. In R. Golledge (Ed.), *Wayfinding behavior* (5-45). Baltimore: Johns Hopkins.

- Gutmann, J.-S. & Nebel, B. (1997). Navigation mobiler Roboter mit Laserscans. In P. Levi, T. Bräunl, and N. Oswald (eds.), *Autonome Mobile Systeme* (36-47). Informatik aktuell, Berlin, Heidelberg New York. Springer.
- Habel, C. (1988). Prozedurale Aspekte der Wegplanung und Wegbeschreibung. In H. Schnelle & G. Rickheit (eds.), *Sprache in Mensch und Computer* (107-133). Opladen: Westdeutscher Verlag.
- Herrmann, Th., Buhl, H.M. & Schweizer, K. (1995). Zur blickpunktbezogenen Wissensrepräsentation: der Richtungseffekt. *Zeitschrift für Psychologie*, 203, 1-23.
- Judd, S.P.D., Dale, K., & Collett, T.S. (1999). On the fine-structure of view-based navigation in insects. In R. Golledge (ed.), *Wayfinding behavior* (229-258). Baltimore: Johns Hopkins.
- Krieg-Brückner, B. (1998). A Taxonomy of Spatial Knowledge for Navigation. In: Schmid, U., Wysotzki, F. (Eds.). *Qualitative and Quantitative Approaches to Spatial Inference and the Analysis of Movements*. Technical Report, 98-2, Technische Universität Berlin, Computer Science Department.
- Krieg-Brückner, B., Röfer, T., Carmesin, H.-O., Müller, R. (1998). A Taxonomy of Spatial Knowledge for Navigation and its Application to the Bremen Autonomous Wheelchair. In Freksa, Ch., Habel, Ch., Wender, K. F. (Eds.), *Spatial Cognition. Lecture Notes in Artificial Intelligence 1404* (373-397). Springer.
- Kuipers, B. (1998). A hierarchy of qualitative representations for space. In Freksa, Ch., Habel, Ch., Wender, K. F. (Eds.), *Spatial Cognition. Lecture Notes in Artificial Intelligence 1404*. Berlin: Springer.
- Kuipers, B. J. & Byun, Y.-T. (1991). A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8, 47-63.
- Loomis, J.M, Klatzky, R.L., Golledge, R.G., & Philbeck, J.W. (1999). Human navigation by path integration. In R. Golledge (ed.), *Wayfinding behavior* (125-151). Baltimore: Johns Hopkins.
- Lynch, K. (1960). *The Image of the City*. Cambridge: MIT-Press.
- McDonald, T.P. & Pellegrino, J.W. (1993). Psychological perspectives on spatial cognition. In T. Gärling & R.G. Golledge (Eds.), *Behavior and Environment: Psychological and geographical approaches*, p. 47-82. Amsterdam: Elsevier.
- Möller, R., Lambrinos, D., Pfeifer, R., Wehner, R., & Labhart, T. (1998). Modeling Ant Navigation with an Autonomous Agent. In: *From Animals to Animats, Proc. Fifth International Conference of The Society for Adaptive Behavior* (185-194). MIT Press.
- Montello, D.R. (1998). A new framework for understanding the acquisition of spatial knowledge in large-scale environments. In M.Egenhofer & R.G.Golledge (Eds.), *Spatial and temporal reasoning in Geographic Information Systems* (143-154). Oxford University Press.
- Musto, A., Stein, K., Eisenkolb, A., Röfer, T. (1999). Qualitative and Quantitative Representations of Locomotion and their Application in Robot Navigation. In: *Proc. of the 16th International Joint Conference on Artificial Intelligence* (1067-1073). Morgan Kaufman Publishers, Inc. San Francisco, CA.
- Nigro, G. & Neisser, U. (1983). Point of view in personal memories. *Cognitive Psychology*, 15, 467-482.
- O'Keefe, J. & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.
- Poucet, B. (1993). Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review*, 100, 163-182.
- Röfer, T. (1998). *Panoramic Image Processing and Route Navigation*. PhD thesis. BISS Monographs 7. Shaker-Verlag.
- Röfer, T. (1999). Route Navigation Using Motion Analysis. In: Freksa, C., Mark, D. M. (Eds.), *Spatial Information Theory, Proc. COSIT'99. Lecture Notes in Computer Science, 1661*, 21-36. Berlin: Springer.
- Rothkegel, R., Wender, K.F., & Schumacher, S. (1998). Judging spatial relations from memory. Freksa, Ch., Habel, Ch., Wender, K. F. (Eds.): *Spatial Cognition. Lecture Notes in Artificial Intelligence 1404*. Berlin: Springer. (79-106).

- Saade, C. & Werner, S. (in press). Flexibilität mentaler Repräsentationen räumlicher Information in Abhängigkeit von der Erwerbsperspektive. *Zeitschrift für Exp. Psychologie*
- Sadalla, E.K., Staplin, L.J., & Burroughs, W.J. (1979). Retrieval processes in distance cognition. *Environment and Behavior*, *12*, 167-182.
- Schölkopf, B. and Mallot, H. A. (1995). View-based cognitive mapping and planning. In *Adaptive Behavior* *3* (311-348).
- Schweizer, K., Herrmann, T., Janzen, G., & Katz, S. (1998). The route direction effect and its constraints. In C. Freksa, C. Habel & K.F. Wender (eds.), *Spatial Cognition. Lecture Notes in Artificial Intelligence 1404*, 19-38. Berlin: Springer.
- Shemyakin, F.N. (1962). Orientation in space. In B.G. Anan'yev et al. (Hrsg.), *Psychological Science in the U.S.S.R.*: (Bd. 1, S. 186-255). Washington, D.C.: U.S.Department of commerce, Office of Technical Services.
- Siegel, A.W. & White, S.H. (1975). The development of spatial representations of large-scale environments. In H.W. Reese (Ed.), *Advances in Child Development and Behavior* (vol.10, S. 9-55). New York: Academic.
- Thorndyke, P.W. & Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, *14*, 560-589.
- Trullier, O., Wiener, S. I., Bertholz, A., and Meyer, J.-A. (1997). Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, *51*, 483-544.
- Tversky, B. & Lee, P.U. (1998). How space structures language. In C. Freksa, C. Habel & K. F. Wender (eds.), *Spatial Cognition* (157-175). Berlin: Springer.
- v. Frisch, K. (1967). *The dance language and orientation of bees*. Oxford University Press, London.
- Wehner, R. & Menzel, R. (1990). Do insects have cognitive maps? *Ann. Rev. Neuroscience*. *13*, 403-413.
- Wehner, R., Michel, B., & Antonsen, P. (1996). Visual navigation in insects: Coupling of egocentric and geocentric information. *The Journal of Experimental Biology*, *199*, 129-140.
- Werner, S., Krieg-Brückner, B., Mallot, H.A., Schweizer, K., & Freksa, C. (1997). Spatial cognition: The role of landmark, route, and survey knowledge in human and robot navigation. In M. Jarke, K. Pasedach, & K. Pohl (eds.), *Informatik '97* (41-50). Berlin: Springer.
- Wiltschko, R. & Wiltschko, W. (1999). Compass orientation as a basic element in avian orientation and navigation. In R. Golledge (ed.), *Wayfinding behavior* (259-293). Baltimore: Johns Hopkins.