

Biologically-based Artificial Navigation Systems: Review and prospects

Olivier Trullier (1,2,3), Sidney I. Wiener (3),
Alain Berthoz (3), and Jean-Arcady Meyer (2)

(1) To whom correspondence should be addressed.

Electronic mail address: trullier@wotan.ens.fr

(2) AnimatLab

Département de Biologie

Ecole Normale Supérieure

46, rue d'Ulm

75230 Paris Cedex 05

FRANCE

(3) CNRS-Collège de France

Laboratoire de Physiologie de la Perception et de l'Action

11, place Marcelin Berthelot

75005 Paris

FRANCE

October 1996

Telephone: 01 44 32 36 33 / 01 43 29 28 75

Fax: 01 44 32 39 01

Abstract

Diverse theories of animal navigation aim at explaining how to determine and maintain a course from one place to another in the environment, although each presents a particular perspective with its own terminologies. These vocabularies sometimes overlap, but unfortunately with different meanings. This paper attempts to precisely define the existing concepts and terminologies, so as to comprehensively describe the different theories and models within the same unifying framework.

We present navigation strategies within a 4 level hierarchical framework based upon levels of complexity of required processing (Guidance, Place recognition-triggered Response, Topological navigation, Metric navigation). This classification is based upon what information is perceived, represented and processed. It contrasts with common distinctions based upon availability of certain sensors or cues and rather stresses the information structure and content of central processors. We then review computational models of animal navigation or animats. These are introduced along with the underlying conceptual basis in biological data drawn from behavioral and physiological experiments, with emphasis on theories of “spatial cognitive maps”.

The goal is to aid in deriving algorithms based upon insights into these processes, algorithms that can be useful both for psychobiologists and roboticists. The main observation is however that despite the fact that all reviewed models claim to have biological inspiration and that some of them explicitly use “Cognitive Map”-like mechanisms, they correspond to different levels of our proposed hierarchy and that none of them exhibits the main capabilities of real “Cognitive Maps” — in Tolman’s sense — that is, a robust capacity for detour and shortcut behaviors.

Keywords: Computational Models, Navigation, Cognitive Maps, Hippocampus, Spatial Representations.

Contents

1	Introduction	1
2	The four-level hierarchy of navigation strategies: definitions	2
2.1	Navigation	2
2.2	Body alignment and target approaching	4
2.3	Guidance	4
2.4	Place recognition-triggered Response	5
2.5	Topological navigation	7
2.6	Metric navigation	8
2.7	Summary	10
3	Navigation in animals and animats	12
3.1	Guidance	12
3.1.1	Behavioral evidence	12
3.1.1.1	Beacon learning	12
3.1.1.2	Landmark-array learning	14
3.1.1.3	Changing aspects of landmark arrays	14
3.1.1.4	Internal map-like representation or guidance	15
3.1.2	Computational models	16
3.1.2.1	Cartwright and Collett (1983)	17
3.1.2.2	Benhamou et al. (1994)	18
3.1.2.3	Wilkie and Palfrey (1987)	20
3.1.2.4	Kuipers and Byun (1991)	21
3.1.2.5	Mataric (1990)	22
3.2	Place recognition-triggered Response	23
3.2.1	Behavioral and physiological evidence	23
3.2.1.1	Behavioral evidence: rats in mazes	23
3.2.1.2	Electrophysiological evidence: Place cells	24
3.2.1.3	Head-direction coding	28
3.2.1.4	Coherence between position and orientation representations	29
3.2.2	Computational models	29
3.2.2.1	Zipser (1986)	30
3.2.2.2	Cartwright and Collett (1987)	32
3.2.2.3	Gaussier and Zrehen (1994)	33
3.2.2.4	Burgess, Recce and O'Keefe (1994)	36
3.2.2.5	Blum and Abbott (1996)	38
3.2.2.6	Brown and Sharp (1995)	40
3.2.2.7	Krakauer (1995)	42
3.3	Topological navigation	43
3.3.1	Behavioral experiments	43
3.3.2	Electrophysiological evidence	45
3.3.2.1	Connections between place cells	45
3.3.2.2	Whole body movement representation	46
3.3.3	Computational models	46

3.3.3.1	Muller, Kubie and Saypoff (1991)	47
3.3.3.2	Kuipers and Byun (1991)	48
3.3.3.3	Mataric (1990)	49
3.3.3.4	Schmajuk and Thieme (1992)	50
3.3.3.5	Schölkopf and Mallot (1994)	53
3.3.3.6	Penna and Wu (1993)	55
3.3.3.7	Bachelder and Waxman (1994)	58
3.4	Metric navigation	61
3.4.1	Behavioral experiments	61
3.4.1.1	The capability to make detours requires distance measurements	61
3.4.1.2	Shortcuts	63
3.4.2	Computational models	64
3.4.2.1	Wan, Touretzky and Redish (1994)	64
3.4.2.2	Worden (1992)	68
3.4.2.3	Prescott (1994)	69
4	Discussion	71
4.1	Summary	71
4.2	Principal suggested directions of future research	72
4.2.1	The inputs to the navigation system	72
4.2.2	The internal representation of space	74
4.2.3	Learned features of the internal spatial representation	74
4.2.4	Computing the motor output	76
4.2.5	Multiple environments	78
4.2.6	Conclusions	78

1 Introduction

Animals spend much of their time moving from one place to another, within or between different environments (habitats). Survival of the individual is contingent upon adaptive skills to find, learn and return to specific places (and often at specific times) — such as the nest or the feeding site — quickly and safely. These skills are summarized by the word *navigation*.

Navigation abilities are also very important for mobile robots. But current robots have a limited set of possible behaviors and although they are able to *move* efficiently through a cluttered environment and to avoid obstacles, their *navigation* systems tend to be specialized and brittle [78]. Navigation is thus one of the most elaborate tasks that current mobile robots try to accomplish.

Animals, on the other hand, can be very good at navigating. The paths they take may be sub-optimal in a mathematical sense, but they are rapidly selected, flexible and the resulting navigation behavior is very adaptive [156]. Biological processes thus seem to be a promising source of ideas for building more efficient robots and roboticists have begun to elaborate “biomimetic” navigation systems. However, the biological data inspiring their work is sometimes outdated, misquoted, or misinterpreted.

In contrast to the usual classification of navigation behaviors according to the availability and use of certain types of sensors and cues (e.g. [1]), this paper proposes a typology of navigation strategies and subsequently organizes the “biologically inspired” computational models into a 4-level hierarchy, while presenting the sometimes controversial biological evidence supporting the hypotheses and assumptions used in each model. The 4 levels are, in order of complexity, Guidance (the animat¹ can move in relation to what it perceives), Place recognition-triggered Response (it can orient itself relative to specific places), Topological navigation (it can move along known paths), and Metric navigation (it can move in relation to an “overview” of the whole environment). This paper attempts to exhaustively review the literature concerning *computational models* of animal navigation that explicitly refer to biological findings and that have been either simulated or implemented on a mobile robot. Thus papers from the wide literature of robotics that don’t have explicit links to biology are beyond the scope of this paper. Descriptions of formal theories of animal navigation, such as the ones by Thinus-Blanc [145] or by Poucet [110], are also beyond the scope of this paper. Because the models that will be reviewed herein have been devised in reference to specific biological findings, and because these findings cannot be easily generalized, the paper will explicitly mention whether each model concerns bees, gerbils, rats or if it intends to represent some “universal” biological principle.

Finally, the main characteristics of these different computational models will be compared in a summary. Some future research approaches are suggested, both for biologists (behavioral and electrophysiological experiments) and for engineers (conceiving and building animats). This should lead to improved understanding of animal behavior, as well as better autonomy and adaptability of current mobile robots.

¹An animat [77, 80, 79, 29, 66] is an artificial organism — either a simulated animal or an animal-like robot — the structure and functionalities of which are based substantially on mechanisms observed in real animals. We will henceforth use the word “animat” to denote the embodiment of any biologically-based artificial navigation system.

2 The four-level hierarchy of navigation strategies: definitions

2.1 Navigation

Navigating seems to be a rather specific task, in the sense that its requirements can be well defined in mathematical (geometrical) terms and that its potential solutions seem to be straightforward. It however involves many distinct sensory inputs and computational processes. Elementary decisions like *turn left*, or *run*, or *stop*, are made on the basis of thousands of incoming signals [156]. These signals come from sensors, but also from memory and expectations derived from previous actions. The global mechanism by which multisensory, internal and external information is fused, stored, and used, is poorly, if at all, understood.

According to Gallistel [42] in *The Organization of Learning* (p.35):

“Navigation is the process of determining and maintaining a course or trajectory from one place to another. Processes for estimating one’s position with respect to the known world are fundamental to it. The known world is composed of the surfaces whose locations relative to one another are represented on a map.”

This definition contains two crucial hypotheses, one being that the world the animal lives and moves in is *represented* within the brain, and the other that this representation can be called a *map*. Waterman [156] also assumes that animal navigation requires the use of a map, either innate or learned.

But what exactly is this “map in the head” (to quote Kuipers [57])? How is it implemented at the neural level? What is the observable behavior during navigation? Studies from ethology, psychology, artificial intelligence, neuroscience and robotics provide different tentative answers. Taken together as a whole, they are diverse and often conflicting.

It is thus necessary to first define what navigation is and what the functions of a navigation system are. For instance, according to Levitt and Lawton [62], navigation is defined by the following three questions: (a) “Where am I?”, (b) “Where are other places relative to me?”, and (c) “How do I get to other places from here?” Underlying the question (a) is the problem of recognizing and identifying the particular place the animal or the animat is in. The answer to the question is not necessarily a specific *position* in relation to a coordinate frame, which might be more related to the second question, although a large portion of the robotics literature deals with this type of answer (see [40]) for a review). Indeed, as far as the place is recognized as a *situation* experienced in past history, a movement selection procedure can be applied and lead to the desired goal place, without the need to answer any of the three questions. A more general formulation of the first question should therefore be: “What are the identifiable characteristics of this place?” Question (b) amounts to processing a global spatial representation — for instance a spatial relation network where every other position has a known relation to the current place, but not necessarily to one another.² This contrasts with the processing of a small-scale information relevant to the goal at hand. Indeed, if the system relies on a global, unitary spatial representation, navigation is probably slowed by the extra amount of time devoted to processing irrelevant pieces of information. However, if the system relies on a collection of

²However, the spatial relation between two places other than the current one is at least implicitly coded since it can be recovered from vector subtraction.

local spatial representations at different scales, question (b) should be: “where are the other *relevant* places relative to me?” These first two questions build up the necessary framework from which the appropriate action can be derived, that is, the object of question (c). Note that question (c) also implicates all the possible places in the global spatial representation, and not only the current goal. Furthermore, it involves a notion of “planning”, that is, asking for the whole trajectory from the current location to the considered goal location.

Navigation according to Poucet [110] is also characterized by the questions of self-localization and planning of spatial actions. The animat, according to Poucet, computes its current position and the goal position within a map-like spatial representation, and plans the sequence of moves from the former to the latter in terms of intermediate positions in the same “map”.

For the present application, these definitions of navigation are too restrictive, since navigation is possible without processing a global spatial representation and without planning. It is thus useful to distinguish between several types of navigation.

First let us distinguish between *local navigation* skill and *way-finding* skill [112]. Local navigation is the process of moving around in the immediate environment, that is, an environment in which only objects within the animat’s range of perception are useful. Thus, there is no need for any type of internal representation of objects and places outside the immediate environment. The animat chooses its action only on the basis of current sensory information (for instance, the overall direction to the perceived goal). Way-finding is the process of moving in a *large-scale environment*, that is, an environment in which there are relevant cues out of the range of the perception, and in particular where the goal is not in the immediate environment.

Way-finding can occur with or without planning, and can lead to a trajectory made only of portions of known paths, or to a trajectory containing new paths. As we will show, the former type of trajectories can be deduced from a *topological*³ representation of the environment, whereas the latter type of trajectories can only be planned from a *metric* representation of the environment.

Consequently, we suggest that there are at least four general types of navigation, which permit successively more complex behaviors: guidance — limited to the local environment, place recognition-triggered response — limited to way-finding without planning, topological navigation — limited to using known paths, and metric navigation.⁴

Each navigation type requires some minimum spatial information from the environment, not necessarily giving rise to a “Cognitive Map”. Tolman [148] introduced the concept of “Cognitive Maps” as a way to interpret experimental findings that, in path selection behaviors, rats did not merely respond reflexively to cue stimuli, but instead used some form of internal spatial representation. Since this seminal paper, many studies concerning navigation have persisted in using the term “Cognitive Map” rather loosely as a metaphor to describe whatever mechanism was used by animals to navigate. But no physical evidence of maps-in-the-head has yet been brought forward.

³Here by “topological relationships between places”, we mean the fact that they are “connected” in a broad sense. Two places are connected when they are adjacent to each other, or when there’s a “known” path leading from one to the other that doesn’t traverse a third known place.

⁴On the basis of the models available in 1991, O’Keefe [98] defines two alternative navigational strategies and groups the models into either the class of “orientation hypotheses” or the class of “guidances”. The former consists of associating the current sensory inputs with motor outputs and predicting the next sensory inputs. This relates to our topological navigation strategy. The latter consists of different mechanisms by which a goal could be attained and relates to our guidance or place recognition-triggered response navigation strategies. Our paper also includes more recent models, some of which fall into the topological or metric navigation strategies.

In the following sections (2.2 to 2.6), we characterize each of the four types of navigation according to selected criteria. These are: (i) what spatial information is used? (ii) how is this information used to guide the movement selection process? and (iii) what kind of situations can be dealt with? This will provide a framework to assess the advantages and limitations of the different computational models in the literature.

Moreover, because navigation is directed toward a goal, this paper will be focused on the situations where an animat first *explores* its environment and *learns* an internal spatial representation, and subsequently uses this representation to guide its movements to return to or to avoid already encountered goals. Certain skills will be assumed, for example the ability to recognize that the goal has been attained, even though this is a basic requirement for navigation. Likewise, although we will consider the problem of dealing with multiple goals, the actual decision process for selecting which goal is desirable at any moment will not be considered here.

2.2 Body alignment and target approaching

Navigation would not be possible without the basic ability of approaching (or moving away from) a perceived object. However, implementing such abilities should not be underestimated, despite Morris' claim that "it requires no more than the operation of a motor 'taxis' system" [86]. Animals have to learn how to coordinate their movements, first in their prehensive space (visuomotor spaces such as eye-hand or eye-head-hand-mouth coordination), then in their locomotor space. As for robots, either the appropriate capacities must be built in, or some learning abilities must be implemented. A simple relevant mechanism could be the feedback loop, where, for instance, reducing the distance to the goal is equivalent to following a gradient of increasing intensity of sensory cues (visual, olfactory) specific to the goal.

When such an ability is acquired, we can assume that whenever the goal object is detectable (a "proximal-cue" situation [86]) the animat is first able to orient only on the basis of sensory information in the direction of the goal (with gaze-shifts and head- or body-turns). This is "**body alignment**". The animat must then be able to move toward the goal. This is "**target approaching**" or "**homing**". In psychological terminology, this is a typical Stimulus-Response (S-R) type of behavior.

When the goal is not perceivable, the animat cannot use this mechanism. However, it may learn that a **landmark** — a salient⁵ spatial cue in the environment — is close to the goal. Such a landmark is called a **beacon** [60], and approaching the beacon leads the animat to the hidden goal. The goal and the beacon are functionally equivalent for guiding approach movements. The animat can use a target approaching (taxis) strategy with the beacon as the principal sensory cue, thus serving as a surrogate goal.

2.3 Guidance

When the goal is hidden and there is no available beacon, the animat can be guided by the spatial distribution of landmarks, that is, a **landmark-configuration** or a landmark-array. At the goal, the animat *memorizes* the spatial relationship between itself and the landmark-configuration. Later on, it will attempt to move so as to replicate this view.

⁵Saliency has no universal definition. In fact, little is known about the types of landmarks used by animals [9] though they can be characterized as proximal/distal, global/local landmarks. Usually, they are tall (the vertical dimension is more important than the other dimensions [39]), and stable [10].

Maintaining “a certain egocentric relationship” with respect to a “particular landmark or object”⁶ is what O’Keefe and Nadel [101] call **guidance** (p.82). For instance, the animat memorizes its distance to a wall and moves so as to maintain this distance. This results in a wall-following behavior.

Moving so as to *attain* a certain spatial relationship between the animat and the landmark-configuration leads the animat to the location where this spatial relationship was memorized. For instance, a hidden goal location can be defined by its distances to three distinct landmarks and the animat successively approaches or moves away from each landmark so as to match its current distances to the memorized ones. The most critical sensory information for this type of navigation was referred to as a **snapshot** by Cartwright and Collett [23]. Collett calls such a process a **landmark guidance** [31]. However, the mechanism is more general, not necessarily involving landmarks but also other types of sensory cues.

Moreover, the memorization of a specific spatial relationship with respect to a landmark-configuration doesn’t necessarily require high-level information such as the identities of landmarks, their positions or the distances to them. Indeed, the animat can memorize the corresponding raw sensory information — how the landmarks are perceived, without any actual recognition. For instance, in the case of wall-following, the knowledge of the distance to a wall can be replaced by the memorization of ultrasonic range sensors’ values. Likewise, the knowledge of the distances to landmarks can be replaced by the memorization of the retinal image of these landmarks at the goal location (which will be called a snapshot). In other words, the selection of the appropriate action that guides navigation doesn’t require the processing of an internal spatial representation. The memorized sensory state has no spatial “meaning”, and thus, is not a “representation” of some spatial feature in the environment. The appropriate movement is selected from the discrepancy between the current state of the sensors and their required state.

We will henceforth say that the animat uses a **guidance** type of navigation strategy whenever it is sufficient for the animat to optimize some sensor-related criterion, based on memorized sensory information, to reach its goal. In contrast, the higher levels of navigation strategies, as described below, will require some kind of “spatial representation”.

2.4 Place recognition-triggered Response

Target approaching and guidance strategies concern local navigation since the goal, a beacon, or a specific landmark-configuration needs to be in the perceptible range of the animat at any moment. The following navigation strategies concern large-scale environments and entail the notion of “place”. In contrast with guidance where “places” are points, we define a place in a large-scale environment as a set of contiguous locations that are equivalent with regard to action selection. In other words, the animat selects the same locomotor action from every location within a given place. A place can also be defined as the set of locations from which a set of landmarks or a landmark-configuration is perceived in an identical or very similar fashion. If the set of locations are sensorily identical (from the animat’s point of view), then the same movement will be selected. These two definitions are thus operationally equivalent.

⁶An *egocentric* relationship is a spatial relationship defined relative to the animat’s body and the animat’s body axes (left/right, front/back, up/down) . The environmental features that these refer to thus change as the animat moves. An *allocentric* relationship — also called *exocentric* — is an object’s spatial relationship defined relative to a reference frame attached to the environment, that is, relative to other stationary objects. It is thus independent of the animat’s position and orientation.

Note that the identification of a place must be independent of the viewing angle, that is, the orientation of the animat. This ability to set the equivalence between different views from the same place is also an interesting question but is beyond the scope of this paper.

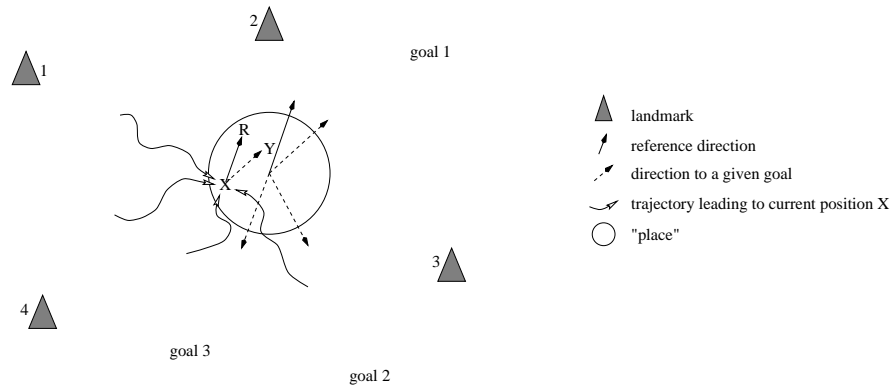


Figure 1: The animat must recognize its position **and** its orientation within a “place” (enclosed by circle), before selecting the next direction of movement. The landmark-configuration can give positional information — which place the animat is in regardless of the exact location X within this place. The landmark-configuration can also define a local and arbitrarily selected reference direction (for example R), which gives directional information. The subsequent movement is selected from the relationship between the reference direction R and the direction to the currently pursued goal (Y for goal 1). Scaling is disproportionate — goals and landmarks are actually more distant.

In this type of navigation, the animat’s strategy in reaching a known but not currently visible goal is threefold (Fig. 1): (a) recognizing the place the animat is currently situated in; (b) orienting itself within this place; (c) selecting in which direction to move so as to reach its current goal. There is no planning of a sequence of subsequent movements, only the selection of the very next action. Thus, the animat responds in an inflexible manner to the situation. We call this **place recognition-triggered response**.

On the basis of previous exploration, the current place can be associated with the memorized direction to the goal from this place. Selecting this direction (part c of the strategy) can be done in two different ways. This direction can be defined by a snapshot. Then, the animat will have to rotate so as to match the current retinal image to the snapshot that defines the direction to the goal. This guidance strategy will re-orient the animat toward the current goal. On the other hand, the memorized direction to the goal can be defined with respect to a local and arbitrarily selected reference direction. This reference direction can be determined by the configuration of spatial cues. Then, the animat will have to rotate by a certain angle given with respect to this reference direction.

Such strategies of inflexible responses based upon local cue-configurations will lead the animat to the next place, where the same strategies can be applied. With the knowledge of what movement to make in each place, the animat is then able to reach the goal by a series of successive trajectories.

Note that at this level, the animat identifies places on the basis of local cues and has a series of fixed responses. But it has no internal representation of the relations between the current place and other places in the environment. The animat still cannot, from the information available at this level, *plan* its path, that is, represent the complete trajectory from the current place to a distant goal. For each step, knowledge is limited to the very next action to perform. If performing the selected action leads to a “wrong” place, because of obstacles for instance, the animat will not be able to reach the goal unless it can wander a little, recognize another place, and use the same local response strategies there (Fig. 2).

The navigation system uses a mapping from place and goal to directed actions, a type of “action model” [6, 78, 118]. We will call such mappings “place-goal-action associations”.

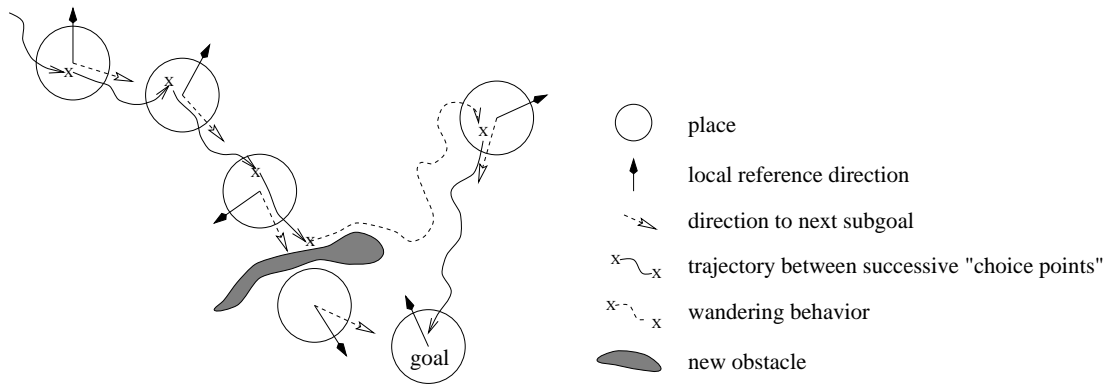


Figure 2: The place recognition-triggered response strategy enables the animat to move from place to place if it knows the direction to follow from each place. When it gets lost because of an obstacle, for instance, it has to wander around until it gets to a known place again.

2.5 Topological navigation

One way for the animat to reach distant places without getting lost — for instance because of drift errors or obstacles — while following the place recognition-triggered response navigation strategy, is to anticipate subsequent incoming stimuli, that is to predict in what place it will arrive next. This is a Stimulus-Response-Stimulus (S-R-S) type of association, and a sequence of such associations will be called a **route** [101]. Being able to predict the next stimulus from the current stimulus and the next action is also called having a “world model” [41, 78, 118, 139].

However, routes are independent of one another and each route leads to a unique goal. They are inflexible because there is no taking into account of the fact that two different routes may pass through the same places. In this sense, the representation of a route is a simple extension of what is used by the place recognition-triggered response navigation strategy, that is, place-goal-action-place associations instead of place-goal-action associations. Navigation would be more adaptive if the spatial representation was goal-independent, that is, if the same spatial representation could be used for multiple goals. Such a result can be obtained by merging place-action-place associations derived from the collection of routes (thus dropping the goal representation) into a **topological** representation of the environment. Any place can now become the origin or the goal of a path and in the case of obstacles, alternate intersecting paths can be taken.

In this type of navigation, the animat can follow routes or sub-sequences of routes. At any rate, the animat goes *only* through already visited (sub)paths. We call this **topological navigation**.

A topological representation can be expressed in mathematical terms as a graph, where *nodes* represent places and *edges* represent *adjacency*, or direct connectivity. Then, two nodes are linked if there is a previously visited direct path which leads from one corresponding place to the other corresponding place, without going through a third intermediate known place.

In this type of navigation, the animat’s strategy in reaching a known but not currently visible goal is four-fold (Fig. 3): (a) recognizing the place the animat is currently situated in; (b) locating the corresponding node in the topological graph; (c) searching for the sequence of nodes (places) to

go through to reach the goal; (d) the resulting route is a *concatenation* of segments of experienced routes, and the animat can follow it even if it has never followed this particular sequence before.

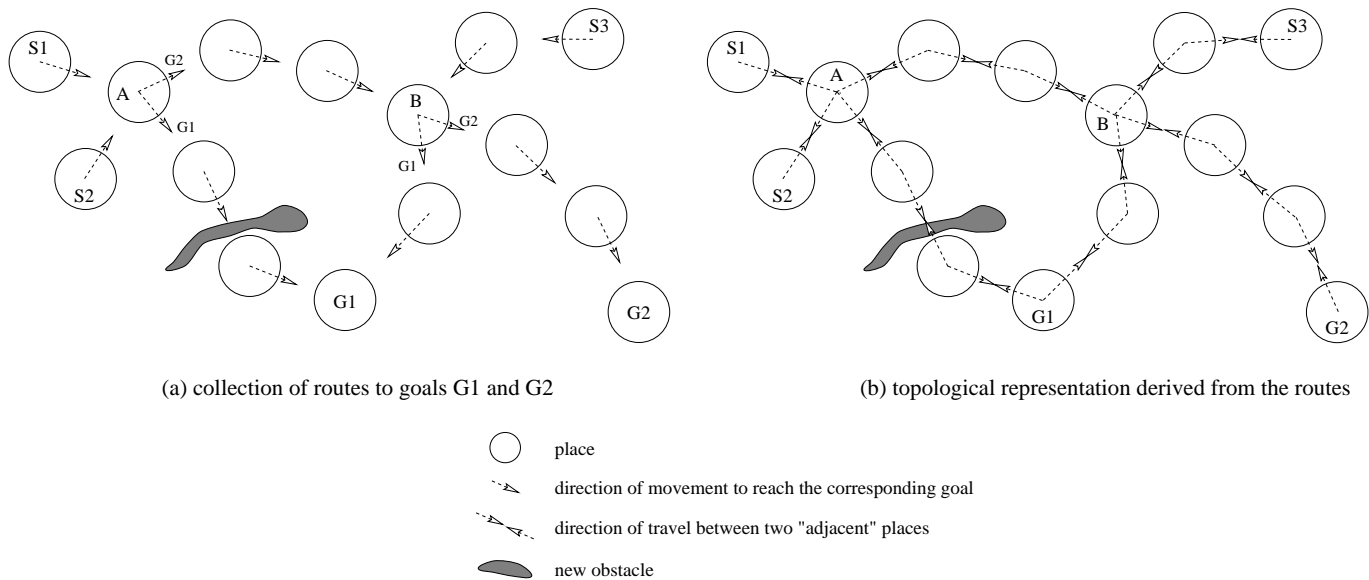


Figure 3: (a) With the place recognition-triggered response strategy there can be an ensemble of intersecting routes. The animat is able to go from S1 to G1, from S2 to G2, and from S3 to G1. However, if there is a new obstacle on the way from S1 to G1, as on this figure, the animat is lost because the route from S1 to G1 is unique (see also Fig. 2). (b) In contrast, if the animat merges its representations of routes into a topological representation, the animat can go back to place A, take the sub-route between places A and B, and take the sub-route from place B to the goal G1. The resulting path is the concatenation of three sub-sequences, derived from three different routes.

At this level, we now have a representation of some of the spatial (topological) relationships between places. The main result is that the animat has *planned* the whole sequence of places to be visited. Note that in order to follow this sequence, the animat can generate a series of commands to the place recognition-triggered response strategy by specifying every next place (node) as a subgoal.

2.6 Metric navigation

The three previous levels of navigation strategies enable successful movements to goal locations although not necessarily through optimal (i.e. shortest) paths. Indeed, the corresponding information cannot be used to compute or plan *novel* trajectories, nor detours in front of unpredicted obstacles, nor shortcuts. Here “detours” and “shortcuts” mean *metric* detours and *metric* shortcuts, and two prototypic examples illustrate behaviors that the previous navigation strategies are incapable of (Fig. 4).

As a first example, consider that an animat, starting from place A, is suddenly confronted with a new wall at place B but has planned to continue the other side, namely from place D to E (Figure 4a). It has to skirt around the obstacle, for instance by following the wall until it can go behind it (place C) and then following it back on the other side. The exact location of place D on the other side of the wall from place B cannot be recognized by visual inspection because the wall obstructs the relevant landmarks. The animat needs to know the distance along the wall — the length of the *detour* (BCD). There is a need for path integration.⁷ Also, we do not consider here the case where

⁷Path integration is an animat’s ability to estimate its current position relative to a known starting position by

landmarks identifying place E can be perceived from place C, in which case the animat doesn't need to complete its detour but can take a shortcut — using a guidance strategy — from C to E. As a second example, consider that an animat, starting from place A, needs to go to the other side of a forest, to a place C it knows but cannot perceive from where it is, namely place B (Figure 4b). The animat also knows a much longer path from place B to place C, around the forest. From the knowledge of this longer path, it tries to estimate the correct heading through the forest to take a *shortcut*. The animat needs to know the relative orientation between places and once in the forest, doesn't have access to landmarks of place B nor of place C.

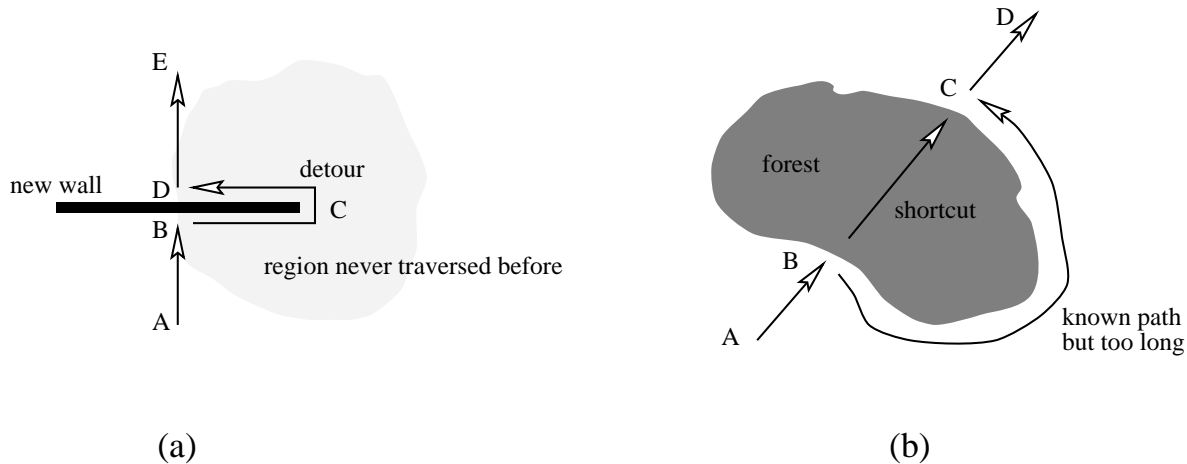


Figure 4: (a) Metric detour and (b) metric shortcut behaviors. In both cases, the animat takes a path never experienced before, without being able to use familiar landmarks (the new wall is assumed to be tall and the forest is assumed to be dense). Note that in (a), the animat could, in principle, go directly from C to E. This would be an illustration of a metric shortcut.

Thus the introduction of **metric** information — distances and angles between places — in addition to the topological information of the previous level, is a necessary requirement for generating novel paths that are metric detours and metric shortcuts. The question remains as to how this information is manipulated in order to yield the necessary instructions to follow the paths. From a computational point of view, a distance and an angle with respect to a reference frame can be represented by a vector and the navigation system should be able to perform vector subtraction to devise a novel path \vec{BC} from known paths \vec{AB} and \vec{AC} . The selection of the movement is thus the result of some deductive reasoning.

From a functional point of view, the animat could use an internalized metric map of the environment — although whether animals use such a map is still controversial — by examining it from above, that is, exploiting what is called a “survey” strategy.

We will henceforth say that the animat exploits a metric navigation strategy whenever it successfully navigates from one place to another by eventually going through new places, *regardless* of previously followed routes, that is, whenever it realizes metric detours or metric shortcuts.

computing the displacement through the integration of speed and direction, that is, through the exclusive use of information gained *en route* [83]. This process is also referred to as ‘dead-reckoning’.

2.7 Summary

Although Gallistel [42] defines the requirements for navigation as the extraction, processing and storage of the “geometric relations among points, lines and surfaces that define the macroscopic shape of the animal’s behavioral space”, it is clear from the above considerations that the four types of navigation strategies do not really require the extraction, processing and storage of all the spatial information available in the environment.

Table 1 summarizes the 4 levels of navigation strategies according to three criteria: (i) the information structure and content, (ii) the movement selection procedure, and (iii) the behavioral repertoire of the navigation strategy. The first level of navigation strategy involves sensory-based information (comparison of sensory information memorized at the goal location with current and ongoing inputs). The second level requires information about landmarks, such as their identities and their spatial configuration in an egocentric reference frame, as well as ability to determine a local directional reference. But there is no information about the spatial interrelations between different places. The third and fourth levels, on the other hand, code some type of spatial interrelations between places, that is *topological* or *metric* information.

The four levels of navigation strategies interact and lower level strategies could very well *override* higher level strategies, which might be more efficient since lower level strategies are simpler than higher level ones. For instance, a novel path computed by a metric navigation strategy may go through an unpredicted but recognizable place. Once in this known place, the animat can choose to follow a known route, as computed by the topological navigation strategy, for instance considering it to be safer than continuing the unknown path. While following this route, the animat may suddenly perceive the goal, which was earlier hidden from the route, and consequently use the guidance strategy instead of continuing on the previously determined route. This latter behavior could be considered as calculating metric shortcut from a planned route to the goal through a novel trajectory, but it is — in our terminology — a simple reflexively opportunistic behavior, resulting from the interactions of different navigation strategies. Other examples of such interactions will be given in the next sections.

	name	stored spatial information	procedure	characteristics
0	Target approaching	None.	Taxis	Basic requirement for navigation.
1	Guidance	Identity of the landmark-configuration. Raw state of the sensory inputs at goal location.	Minimize the mismatch between the perceived configuration and the memorized configuration (approach). Self-localize by recognizing the current place as an already experienced place. Orient relative to it. Move in the goal-associated direction.	Local navigation. Only when direct perception is available.
2	Place recognition-triggered Response	Landmark-configurations defining places. A local directional reference frame for each. The direction of movement that leads to the goal from each place.	Search for the sequence of places linked by experienced routes from the current place to the goal.	Way-finding. Stimulus-Response type of behavior.
3	Topological navigation	A set of landmark-configurations linked by topological relationships.	Plan a trajectory which will be followed by lower level strategies. The resulting path is not necessarily a previously taken one.	Way-finding. Stimulus-Response-Stimulus type of behavior. Topological detours (path selection).
4	Metric navigation	A set of landmark-configurations linked by metric relationships.		Way-finding. Metric detours, metric shortcuts, novelty.

Table 1: A hierarchy of navigation strategies

3 Navigation in animals and animats

This section presents some of the behavioral, anatomical and physiological evidence supporting the existence of the proposed four types of navigation strategies in animals, and the computational models that correspond to each type.

3.1 Guidance

3.1.1 Behavioral evidence

Target approaching, and its primitive form, taxis, is observed in virtually all animals that are capable of locomotion. There is a wide literature describing the different patterns of behavior and studying how animals acquire and express this ability (see [49]). In fact, animals are able to approach proximal cues very early (e.g. rats only 17-days-old [124]) and the underlying mechanisms seem more innate than learned. Thus, target approaching is frequently assumed as present and virtually all models include this without explicitly modeling the mechanism of taxis.

Although guidance can account for many observed behaviors, there is no anatomical or physiological evidence for a discrete module in the nervous system dedicated to such a function, that is, evidence for each of the components described in Sections 2.2 and 2.3. However, there is much behavioral evidence. Collett [31] reviews in detail behavioral evidence for “landmark guidance” in insects. He describes the properties of the snapshots, that is, the memorized environmental scenes, that enable an insect to return to specific locations, and then summarizes hypotheses about the environmental features that might be included in these snapshots. Here, we will focus on three main issues for these local navigation strategies: (i) what type of objects are selected as landmarks or beacons, (ii) how landmarks may be grouped together to guide navigation, and (iii) how perceived and memorized configurations are compared.

3.1.1.1 Beacon learning

Collett et al. [33] studied how gerbils use salient objects in the environment as spatial cues to locate hidden food. They trained the animals to retrieve food in a cue-controlled environment where the spatial relation between the food location and a given landmark was held constant, but this pair was translated from trial to trial, in the absence of the animal. They also showed that gerbils are able to discriminate between a relevant landmark — the object that determines the food location in a constant way — and an irrelevant landmark — an other object that is moved within the environment from trial to trial, independently of the food location. Thus the gerbils learned to use the relevant beacon efficiently.

The landmark used in these experiments was a vertically-oriented cylinder — a radially symmetrical cue. Thus, if the food was always 5cm from it, this location was ambiguous since the cylinder had the same appearance from all angles. The gerbils must have had additional directional information to determine the reward site.⁸ Consequently, beacon homing is not sufficient to explain how gerbils retrieved the food in this single-landmark case.

Biegler and Morris [10] also showed that rats are able to learn to discriminate between relevant and irrelevant landmarks. However, they further showed that the quality of beacon learning depends

⁸Collett et al. acknowledge this without being able to suggest how gerbils got this information in these experiments. This illustrates the difficulty of devising an environment where all cues are controlled.

on what type of environmental features remains invariant from trial to trial. They trained rats to find food at a specific distance and a specific direction from a landmark (filled circle in Figure 5). Directional information was provided by the fact that the rectangular arena was composed of 3 black walls and one white wall. The latter was a salient cue under experimental control. The environment contained another landmark which bore no constant relation to the goal location (unfilled circle). Two groups of rats were trained in two different contexts. For the first group, both landmarks remained at fixed locations, that is, the spatial configuration was stable from trial to trial (Figure 5a, group “fixed”). The second group experienced a variable configuration, where the two landmarks were moved independently from each other from session to session (Figure 5b, group “varied”).




Figure 5: The two training contexts in Biegler and Morris’ experiments. In each case, the environment is a rectangular arena with three sides painted in black and one side in white (shaded side of the rectangle in the overhead views), providing a strong directional reference. (a) group “fixed”: there are two landmarks in the environment, both fixed relative to the arena across the trials. During training, the food (F) was located to the “south” of one of the landmark (filled circle). In probe trials, the animal searched at the exact same location (S). (b) group “varied”: there is a relevant landmark (filled circle) which defines the hidden food location (F). the food is always hidden at the same distance, “south” of the relevant landmark. There is also an irrelevant landmark (unfilled circle). Both are moved within the environment from trial to trial. In probe trials, the animal searched in the vicinity of the relevant landmark (shaded area labeled S), instead of at the exact location defined by the landmark.

In probe trials, the maze was not baited. Rats from the “fixed” group searched only at the correct location. Thus, they went to the beacon, and used directional information to search at the reward site. Rats from the “varied” group also went to the relevant landmark but instead, their search included a circle all around the vicinity of the beacon.⁹ This result shows that beacon learning is possible — rats from the “varied” group were able to discriminate between the two landmarks and searched around the landmark that was the best predictor of food availability — but also that “stability” aided spatial learning — rats from the “varied” group didn’t learn the constant spatial relationship between the beacon and the food location (namely “opposite the white wall”) but only the “vicinity property”.

The approach to a beacon is usually under continuous visual control, which corresponds to the definition of target approaching given in Section 2.2. However, if the goal is some distance away from a single given landmark, the sudden displacement of the landmark is ignored by gerbils, suggesting

⁹This result is in contrast with Collett et al. experiments on gerbils. Collett et al. reported that the animals search *at* the correct location, and not *around* the relevant landmark. However, they didn’t give details as to whether the relevant landmark remained fixed relative to the environment — in which case the landmark-configuration is intermediate between “fixed” and “varied” — or not — in which case the experiment is the same as group “varied”.

that the landmark is no more considered as a beacon and that other orienting mechanisms are used by the animal [30].

3.1.1.2 Landmark-array learning

Collett et al. [33] showed that gerbils are able to learn a food location defined entirely by the configuration formed by a pair of distinguishable landmarks, both in terms of distance and in terms of direction (Figure 6). Thus this is a straightforward example of the use of landmark configuration for guidance. This is interesting since it shows that these animals are capable of ignoring the position of the landmark-reward array relative to walls in the experimental chamber. Cue such as the shape of the room, possible cues on the walls and the floor, lighting were all ignored by the gerbils that learned this navigation task.

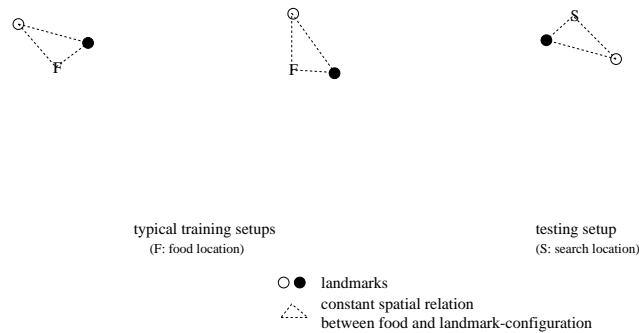


Figure 6: Gerbils are able to learn the food location entirely defined by two distinct landmarks. The local distance and angle relations between the two landmarks and the food, that is, the configuration, were kept constant but the whole array was displaced and rotated randomly during training. After (Collett et al., 1986)

3.1.1.3 Changing aspects of landmark arrays

An interesting question then is how animals use perception (or the knowledge) of a landmark-configuration to guide their trajectory to a goal. Is there a geometrical transformation from perception to an internal spatial representation or are the sensory features directly exploited to trigger trajectories in a reflexive manner? The results in the corresponding literature are mixed according to the particular species, tasks and environments.

Cartwright and Collett [23] tested landmark learning in bees and found that “bees learn no more than the apparent size and bearing of the landmark as seen from the food source.”¹⁰ Their first experiment consisted of training the bee to feed at a specific location in a cue-controlled environment containing but a single radially symmetrical landmark. Subsequently, the size of the landmark was changed; the bee then renewed its search at a location where the visual angle subtended by the landmark — and hence its apparent size — was the same as in the training condition. The search was also at the same relative bearing with respect to an unspecified reference direction. The experiment was repeated with landmark arrays, which the bees also learned to use to localize the food. Then

¹⁰they also observed the same tendency in gerbils but the apparent size of the landmark was not the only cue used by gerbils [33].

in another manipulation, the sizes of the landmarks were changed, but since their absolute positions were maintained fixed, this created a conflict (Figure 7). The bees then went to a location that was the best compromise between the apparent sizes of the landmarks and the relative orientations of the landmarks.

Thus, in these experiments, the bees learned locations on the basis of the visual scene and not on the basis of absolute geometric relations between landmarks. As a consequence, navigation in these animals seems to be based more on attempting to reproduce direct sensory impressions than on a geometrical reconstruction of the external world.

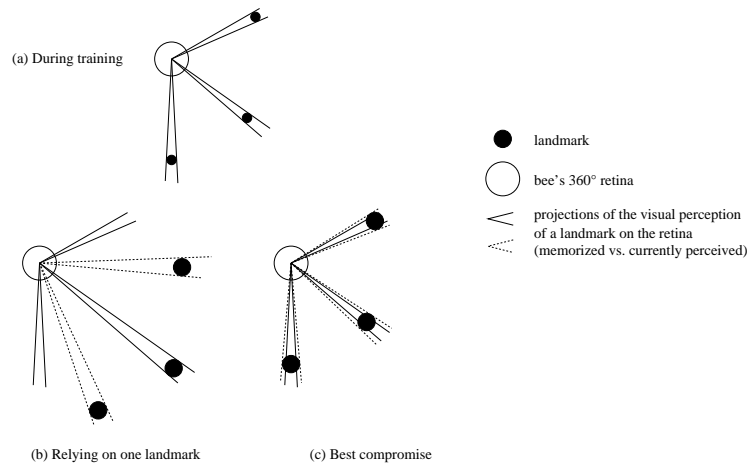


Figure 7: Bees locate the food site relative to a landmark array by matching the apparent sizes and relative bearings of the landmarks. In overhead views: (a) during training, the scene observed at the food location (the visual angles subtended by the landmarks, in continuous lines) is memorized. If the landmarks are enlarged but remain in the same positions, (b) the bees don't search food at a place providing a view perfectly matching only one of the remembered landmarks, but (c) at a place providing the best compromise between the apparent sizes and relative bearings of all memorized landmarks. (Dashed lines represent the current visual scene and continuous lines represent the memorized snapshot.) After (Cartwright and Collett, 1983)

3.1.1.4 Internal map-like representation or guidance

There is controversy as to whether insects navigate by using internal map-like representations. Wehner [159, 160] found that insects use specialized sensory processing systems, exploiting relatively simple mechanisms to navigate within their environments. For instance, ants have an array of photoreceptors at the top of their heads, each responding to different orientations of the polarization streak of sunlight in the sky (which is invisible to humans). Thus they have a built-in compass. This permits the ants to orient without the need for a matching process between sensory inputs and an internal representation. Wehner called these specialized processing systems “matched filters”. They are *ad hoc* processes but their existence cautions us against invoking internal map-like representations to explain insect navigation.

Gould [50], on the other hand, reported findings from honeybees that he interpreted as supporting the existence of an internal map-like representation. Honeybees trained at a feeding site A were captured upon their departure from the hive and released at another site B (see Figure 8a). Four possibilities existed: (a) the bees would be completely disoriented and would depart at random directions, (b) the bees would return to the hive, (c) the bees would attempt to fly to site A and follow a stored compass direction, that is, they would depart from B in the direction corresponding to the hive-to-site A direction, as if they hadn't been displaced, and (d) the bees would recognize

site B (which means they had already been there) and orient directly toward A. Gould observed that bees flew in a straight line from B to A. Since he assumed that they had never traveled along this path, he deduced that bees were able to reorient themselves on the basis of a spatial representation that included the relationship between site A and site B.

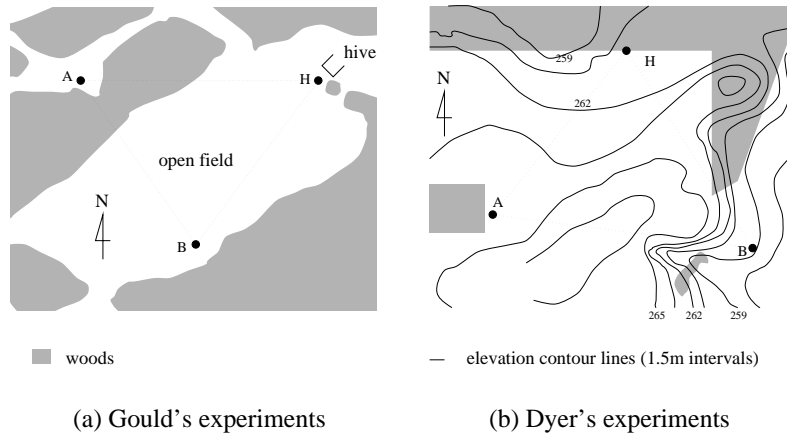


Figure 8: Setup of behavioral experiments conducted on honeybees by (a) Gould (1986) and (b) Dyer (1991). Bees from the hive (H) were trained to feed at site A. They were then captured as they were departing from the hive toward site A and passively transported in darkness to site B. They were subsequently released from site B. Depending on the terrain, some bees flew away in the same compass direction as if from the hive, some flew back to the hive, some flew directly to site A, and some flew at random directions. After (Gould, 1986) and (Dyer, 1991)

Although several authors have failed to replicate these results, Dyer [36] succeeded partially. He used two feeding sites, A and B (see Figure 8b). From site A, the bees had an unobstructed view of distant landmarks, some of which might characterize site B, whereas from site B, located at a bottom of a ravine, landmarks defining site A were hidden. He observed that bees trained at B and released at A could fly directly to B. In contrast, most bees trained at A and released at B flew in the compass direction they would have taken, had they not been displaced (solution c). A minority headed toward the hive (solution b). His conclusion was that honeybees “do not have the ‘mental maps’ posited by Gould, or any other mechanism to compute novel short cuts between familiar sites that are not in view of each other” [36, p.245]. In fact, the success of bees released at A to fly directly to B could be explained by a simple sensory *guidance* mechanism involving distant landmarks. Similar interpretations can be given to Gould’s results.

Thus, these results suggest that, even in the absence of specialized senses, it is not necessary to invoke an internal spatial representation in the form of a “mental map”. The intermediate solution is a mechanism that directly processes sensory information to guide navigation, for instance a matching process between remembered landmark-configuration and currently perceived landmark-configuration (coupled with the capacity to recognize distant cues even if only previously experienced from proximal locations).

3.1.2 Computational models

Here we present four navigation systems that fall in the category of guidance strategies, the first of which is a computational model of navigational behavior in bees and was devised by behavioral scientists. The second is a computational model of navigational behavior in rats in the Morris water-maze task. The last two are taken from more complex models devised by AI researchers and

roboticists, and will be returned to in later chapters. They also illustrate how navigation strategies at different levels of the hierarchy may interact.

3.1.2.1 Cartwright and Collett (1983)

Cartwright and Collett [23] propose several models with the common concept that bees do not use “maps” but move so as to reduce the discrepancy between the current retinal image and the memorized “snapshot” viewed at the goal site.

Two types of information are available from the retinal image: the orientation of each landmark relative to the current head-direction (egocentric orientation) and the landmark’s apparent size (which is a monotonic function of its distance from the bee). The models differ slightly in the way these two types of information are processed but the general principle consists of the following rules: (i) if the apparent size of a landmark is smaller than the remembered size, the model bee tries to move toward this landmark and if the apparent size is bigger than remembered, it tries to move away from it; (ii) if the landmark is to the right (or left) of the corresponding remembered orientation, the model bee will rotate to the left (or right). The two movements — forward/backward and left/right — are independently specified, and computed for each landmark profile in the snapshot. They are represented by vectors, radial for forward/backward movements and tangential for left/right movements (Figure 9). The model bee executes the movement that is the sum of the individual vectors.

However, Cartwright and Collett do not assume that individual landmarks can be identified. Instead, the retinal image and the memorized snapshot are each divided into light and dark areas. Each area on the retinal image is then paired with the closest area of the same type on the snapshot (Figure 9 only shows pairing of dark areas). The model bee maintains a constant orientation in flight, as do real bees [32]. Consequently, the two images are constantly aligned. The computation described above for each landmark is performed on each corresponding pair of areas, on the retina and the memorized snapshot.

One of the biggest problems with this model stems from the problem of matching the correct light and dark areas between the current retinal image and the memorized snapshot, since the landmarks are not distinguishable. The order of landmarks in the retinal image doesn’t necessarily correspond to the order in the snapshot, so that areas corresponding to inappropriate landmarks could be paired, yielding incorrect direction vectors. The pairing process gets even more difficult as the number of landmarks grows. This limitation on the pairing process leads to local minima, that is, positions other than the goal location where individual non-zero vectors cancel out, and thus wrongly advise the model bee not to move. Cartwright and Collett confirm that the model bee can be “trapped by partial matches”.

More recent behavioral observations show that although real bees use this guidance strategy, it is only employed as a second phase of navigation. Indeed, guidance by image matching seems to be limited to the vicinity of the goal [31]. The first phase of navigation then consists in approaching this vicinity by other means (e.g. beacon homing as in the experiments presented by Collett and Baron [32]) and then aligning the body in a specific compass orientation, corresponding to the orientation taken when the goal location was learned. As the second phase begins, the bee is sufficiently close to the goal that the retinal image and the snapshot are approximately aligned. The discrepancy is thus small and only needs minor adjustment.

Another way to avoid mispairing is to use sparse images [31]. For instance, a snapshot would consist of only the configuration of nearby landmarks, or inversely, only the configuration of distant




Figure 9: The snapshot model of Cartwright and Collett. The model bee is represented by two concentric circles, the outer circle represents its current retinal image and the inner one represents its memorized snapshot. (a) At the food location, the images of the landmarks (filled circles) project onto the retina and the snapshot, as shown. To return to the feeding site (b), the model bee tries to match its retinal image (dark areas on the outer circle) and the remembered snapshot (dark areas on the inner circle). Each dark area on the snapshot is paired to the closest dark area on the retinal image (dashed lines). Each comparison yields a pair of vectors (unfilled arrows) indicating the required adjustments: a radial vector (indicating required forward/backward motion) for the size adjustment and a tangential vector (indicating required left/right rotation) for the orientation adjustment. The resulting movement, which globally reduces the discrepancy between the snapshots, is computed (c) from the summed vector (filled arrow).

After (Cartwright and Collett, 1983)

landmarks. The mechanism by which such a filtering from the retinal image would be done is unclear (the use of motion parallax or limited depth of focal planes would be two possibilities).

3.1.2.2 Benhamou et al. (1994)

Benhamou, Bovet and Poucet [7, 8] extend the previous model by changing the way movement toward the goal is selected. In contrast with Cartwright and Collett’s model, where the movement selection was based on the local discrepancy measure (discrepancy between currently perceived landmark configuration and the landmark configuration perceived at the goal), this animat estimates the discrepancies that would be found at nearby locations, compares these estimates, and moves toward the location of least discrepancy. The estimations are based on learned landmark configuration changes. This model can be considered as using an internal spatial representation.

Figure 10 shows the simulated environment in which the model was tested. The animat’s movements are constrained to 6 directions, so that in practice, it moves from place to place (hexagons). At each place, the animat perceives all the landmarks in the environment, in a 360° view, and measures their apparent sizes and head-referred bearings. A local reference direction is then computed as the angular mean of the head-referred bearings of the landmarks, weighted by their apparent sizes. This computation yields a reference direction independent of the current head-orientation of the animat. Benhamou et al. then define a “panorama” at each place, i.e. a 360° view of the environment.¹¹ This panorama is stored in terms of landmark identities, the visual angles they subtend and their orientations with respect to the local reference direction.

Exploration of the environment consists in learning how the views of landmarks change as the ani-

¹¹This is almost equivalent to the snapshots of Cartwright and Collett’s model. Here, however, landmarks can be identified, and the panorama is not memorized as simply a retinal image.

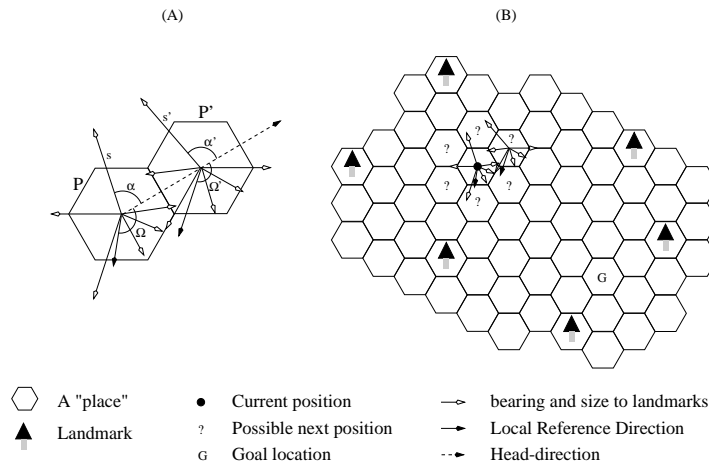


Figure 10: The simulated environment to test the model proposed by Benhamou et al. (1994) (a) Each place (hexagons) has its own “local reference direction” defined as the vectorial angular mean (solid arrow) of the bearings (directions of open arrows) of the visible landmarks, weighted by their respective angular sizes (lengths of open arrows). During exploration, the animat learns that when moving from place P to place P’, the positions of landmarks change, for example from head-referred bearings α to head-referred bearings α' . In addition, apparent sizes change from s to s' , and that the local reference direction changes from head-referred bearing Ω to head-referred bearing Ω' . (b) Consequently, when the animat is at the center of place P, it estimates the discrepancies (question marks) between the landmark-configurations, as seen from adjacent places, and the snapshot memorized at the goal. The animat moves to the adjacent place where the estimated discrepancy is the lowest.

mat moves, and thus, how the panoramas change. This will consequently allow the animat to “mentally move” to adjacent places, predicting the landmark configuration at these places and compute the discrepancy between the landmark configuration perceived there and the panorama memorized at the goal, and choose the appropriate move to reduce the discrepancy. First, the change in panoramas is learned in egocentric references, that is, how the identified landmarks of the configuration change in their head-referred (thus egocentric) bearings and the change in their apparent sizes (α and s in the figure). However, measures of the egocentric bearings cannot be compared to one another for the computation of discrepancies. Bearings with respect to a global reference direction D are needed. D is arbitrarily defined at the beginning of the exploration and related to the local reference directions at each place, by exploiting the fact that during the displacement from place P to place P’, the animat maintains a constant head-direction (dashed line linking the 2 points), and thus, $\Omega'_{referred\ to\ D} = \Omega_{referred\ to\ D} + (\Omega'_{head-referred} - \Omega_{head-referred})$. Thus, the change in panoramas is learned first in terms of egocentrically referred changes and secondly in its allocentric form. The allocentric representation is derived from changes in the local reference directions with respect to the global reference direction.

Benhamou et al. use associative memories to learn the dynamics of panorama changes. These memories are implemented in the form of two-layer feedforward networks. There is one such network for each landmark. Each network receives as inputs the head-referred bearing (α) and the apparent size (s) of the corresponding landmark at the current place P. It outputs the change as a head-referred bearing ($\alpha - \alpha'$) that corresponds to a move of the animat from an adjacent place P’ to the current place P, and the apparent size (s') of the landmark at place P’. Another two-layer feedforward network computes the overall reference direction from the bearings of all the landmarks with respect to the local reference direction. Note however that this computation doesn’t require the identification of the current place, nor its recognition. In fact, it is completely independent of where the animat actually is.

Although all the necessary information is acquired and stored by the animat, there is no actual explicit use of “where the animat is.” Indeed, each place is uniquely characterized by a panorama. Comparing the current panorama and memorized panoramas would enable the recognition of the place. However, this model takes each landmark into account separately in order to estimate the panoramas at neighboring places. Navigation is then performed on the basis of discrepancy measures, as in the model of Cartwright and Collett [23]. Thus, this model belongs to the guidance strategy.

However, the ability to predict panoramas at neighboring places could in principle be extended to places further away. The animat could “simulate” moving through the environment from place to place. If such a computational ability was added to this model, the animat could perform path planning and topological navigation, without requiring the storage of any additional information.

3.1.2.3 Wilkie and Palfrey (1987)

Wilkie and Palfrey [164] propose a “simple perceptual memory-matching model” that is able, in simulations, to find the platform in the Morris water-maze task [86, 87].

As shown in Figure 11a, the behavioral experiment consists of placing a rat into a circular pool of water made opaque by the addition of milk or paint. A small platform lies hidden just below the water surface. Rats try to escape from the water by finding and climbing onto the platform. At first they succeed by chance, but on subsequent trials, they rapidly learn to take direct paths to the hidden platform from different starting locations. There is no sensory guidance since the rats presumably cannot directly detect the position of the platform. Motor response strategies can be ruled out by starting the rat from different initial positions. Spatial memory is tested by removing the platform on probe trials and observing if the rats search in the vicinity of its former position, which they do. This indicates that navigation is guided by configurations of landmark cues in the environment outside the pool.




Figure 11: The Morris water maze. (a) behavioral experiment: a rat is able to learn the association between the distribution of visual cues (black objects) and the position of the immersed platform (dashed cylinder). The animal is then able to reach the platform by a direct trajectory from any starting point in the pool (arrows). (b) (overhead view) Wilkie and Palfrey propose a model where the animat successively approaches or moves away from each landmark so as to match their perceived distances to memorized values. The animat first moves toward landmark 1 because the observed distance ($R1'$) is greater than the memorized distance ($R1$) and then moves away from landmark 2 because $R2'$ is smaller than $R2$.

Hippocampal lesions in rats cause impairments in this task. Since O’Keefe and Nadel proposed that the hippocampus served as a “cognitive map” [101], Morris interpreted his observations as supporting the idea that the rats are using an internal map-like representation to guide their movements to the hidden platform.

Here, Wilkie and Palfrey propose that instead of using an internal map-like representation, the rat moves so as to match its perceived distance from each of the respective cues with memorized

distances learned when the rat first reached the platform. This is similar to Cartwright and Collett’s model. However, this animat takes only the distance and not the angular information into account. Moreover, individual landmarks must be identifiable.

Wilkie and Palfrey’s animat moves toward or away from landmarks according to the difference between the memorized distance from the goal location (platform) to the landmark and the current distance from the animat to the landmark. Thus, for instance, the animat on the Figure 11b approaches landmark 1 and then moves away from landmark 2 and repeats the process until it reaches the hidden platform.

This model does not replicate rat behavior perfectly since it produces markedly jagged zigzag paths. The important point here is that a rather simple mechanism can guide an animat to a memorized goal location on the basis of multiple distant environmental cues, without using configurational information nor a topological or metric internal representation.

3.1.2.4 Kuipers and Byun (1991)

The navigation system of the robot NX, devised by Kuipers and Byun [58], maintains a spatial representation of the environment in terms of “places” and is further discussed in Section 3.3.3.2. Each such place is defined as the attraction basin of some “distinctiveness” measure — a function that depends on the different sensor readings of the animat. In other words, it is defined as a region of the environment from which a specific location — defined as a local maximum of the “distinctiveness” measure — can be reached by a simple hill-climbing procedure. The animat’s movements result from a gradient ascent process.

For instance, a T-junction in an indoor environment can be recognized as that location where all three walls are equidistant (Figure 12). The distinctiveness measure “chosen” by NX for the region around this location will be differences of the range sensor readings in all directions around the animat. The animat will have reached this location when all its range sensor readings are the same.



Figure 12: How the robot NX reaches and recognizes a T-section in a maze. The T-section is characterized by location T, which is defined as the location from where the three sonar readings are equal.

One interesting aspect of this application is that the animat has several different sensors and eventually learns which sensors are relevant to specific places. Pierce and Kuipers [107] also deal with the sensor/feature selection problem for goal-oriented behavior.

A place can be defined as the set of locations from which the same distinctiveness measure is selected by the animat. Consequently, the local goal for the animat at any of these locations is to reach the peak of this measure. The animat will recognize the place only when arriving at this peak. Thus, this sensor-related navigation is a guidance strategy that is limited only to distinctive locations, but that can be useful in place recognition systems (cf. Section 3.3.3.2).

3.1.2.5 Mataric (1990)

The navigation system proposed by Mataric [68] also maintains a representation of the environment in term of places (This will be further discussed in Section 3.3.3.3). Similar to the definition of a location adopted by Kuipers and Byun, Mataric defines a landmark as “a feature or location which is robustly and reliably detectable by the sensors. Consequently, a landmark is an extreme point in sensor space.”¹² Kuipers and Byun use instantaneous sensory patterns for their place definition and consequently, an extreme point in sensor space corresponds to a precise location in the environment. In contrast, Mataric chooses a dynamic matching process, thus avoiding many inaccuracies coming from the sensor position, orientation, measurement and other noise. A “landmark” is therefore recognized as the animat is moving. Consequently, an extreme point in sensor space corresponds to an extended region in the environment.

Mataric’s animat moves within an indoor environment. It is equipped with ultrasonic range sensors and an internal compass. The animat moves in a stereotyped manner, in environmentally-constrained movements that consist of keeping the animat at a constant distance with walls. This is a wall-following behavior which can be considered a *guidance* process since the animat maintains “a certain egocentric relationship” relative to the wall. As a consequence, the animat’s sensor readings remain constant over time. The animat uses the compass measurement and some of the sonar signals to characterize right walls, left walls, corridors, and irregular boundaries. For instance, the animat will register a “North-right-wall” whenever it detects a wall on the right as it is moving North. In this context, a place corresponds to a stable reading of the animat’s sensors over a period of time, and the guidance process ensures that the same “landmarks” are perceived similarly each time the animat passes by. As was the case for Kuipers and Byun’s animat, a guidance strategy is used to define places. The resulting place representation will be used by higher-level navigation strategies (cf. Section 3.3.3.3).

¹²The current location of the animat can be described by the animat’s (x,y) position and orientation in a coordinate frame attached to the physical environment. To each location also correspond specific sensor readings. Thus, the current location can be described by the values of all the animat’s sensor readings. The corresponding coordinate frame is in sensor space, an n-dimensional space. The apparent increase in complexity from a 3-dimensional space to an n-dimensional space is compensated for by the fact that the animat has direct access to its sensor readings whereas it would have to compute its position and orientation from these sensor readings, a difficult inverse problem.

3.2 Place recognition-triggered Response

In the previous sections, we described navigation where the animals merely responded to specific stimuli, by approaching (or avoiding) them or by moving in a specific fashion in relation to them. Here we will deal with the situation where the animals also “know” where they are, that is, recognize the place they are in and consequently choose the next direction of movement. This second strategy was defined as **Place recognition-triggered Response**.

3.2.1 Behavioral and physiological evidence

First, how is the environment partitioned into distinct and discriminable places? What criteria are used to characterize individual places? Behavioral and electrophysiological experiments, described below, yield some answers — at least for animals such as gerbils, hamsters or rats. This has been investigated in behavioral experiments that test for place recognition in mazes. The type of “place” can be characterized by egocentric or allocentric references, and by local or distal cues, depending upon experimental context. Electrophysiological experiments suggest that some hippocampal cells are selectively active when the rat occupies specific “places”. Other neurons also code for other information relevant for navigation, such as head-orientation, locomotion, or turning.

3.2.1.1 Behavioral evidence: rats in mazes

When rats are trained to find a food reward in a T-maze (Figure 13a), how do they remember which of the two arms to choose on subsequent trials? Contradictory results led to a long-running controversy between place learning and response learning theories: do animals plan their movements on the basis of internal representations of cue configurations or, alternatively, are they only learning to make locomotor responses to specific stimuli?

Of course, rats are capable of learning both strategies, depending upon training conditions or more specifically, on “what is invariant throughout the trials” [117]. Indeed, the different motor or cognitive responses observed in the T-maze experiments can be grouped into three types (Figure 13), each corresponding to the environmental or behavioral feature that has been most successful for predicting the food reward location. The same experimental setup (Fig. 13a) is used but the features that are varied from trial to trial are different. (1) The two arms of the T-maze are physically distinguishable, for instance a tactile cue may be present in only one arm. If food is always in the arm with the tactile cue, irrespective of other sensory or motor cues, the “goal place” will be characterized by the presence of this cue, which serves as a beacon (Fig. 13b). (2) If the ensemble of environmental landmarks, and the maze itself, are moved from trial to trial, but food is always in the arm to the right of the T-maze, then the “goal place” will be characterized by an egocentric motor-response, that is, a right turn of the animal in the T-maze (Fig. 13c). (3) If the whole maze is rotated and intramaze cues are moved about from trial to trial, but the reward location doesn’t vary with respect to the extramaze landmarks, then the “goal place” will be characterized by its position with respect to stable extramaze landmarks (Fig. 13d).

We can explain the first type of associations between environmental features and reward by a guidance navigation strategy, or more precisely, by a “beacon homing” strategy, since the animal approaches an intramaze cue. The second type of associations can correspond to a navigation strategy based on an egocentrically-defined motor response — which can be thought of as a simple reflex — “turn right at the choice-point of the T-maze”. The third type of associations can however be explained by two different navigation strategies: (1) guidance navigation strategy, where the




Figure 13: The different behavioral responses of a rat trained to search for food in a T-maze (a). The “goal place”, that is, the arm containing the food, can be characterized by intramaze cues (b), an egocentric motor-response (c), or by extramaze landmarks (d). The different characterizations depend upon the training procedure, where some environmental features remain stable with respect to the reward location across trials, whereas other features are displaced from trial to trial. Depending on what association the rat makes, it will go where the cue is (b), execute a right turn (c), or go to the arm in a position defined by the configuration of extramaze landmarks but not by intramaze cues (d). After (Leonard and McNaughton, 1990)

landmark-configuration is used continuously and where turning left corresponds to reducing the discrepancy between a memorized snapshot and the current retinal image; or (2) place recognition-triggered response strategy, where the landmark-configuration need be observed only once, at the start, and turning left corresponds to part of the sequence of movements planned at the start, without the need for further sensory updating.

This ambiguity then leads us to question the necessity for invoking the place recognition-triggered response type of navigation strategy. It seems difficult to test, behaviorally or electrophysiologically, the difference between guidance and place recognition-triggered response. However, the models that are described in this section specifically implement the latter type of navigation strategy.

3.2.1.2 Electrophysiological evidence: Place cells

Definitions

Discovered by O’Keefe and Dostrovsky [100] in the rat hippocampus in 1971, “place cells” brought evidence that at the neural level, representations of the environment are divided into small regions or places and that this could provide a neural substrate with which the rat could “recognize” where it is at every moment. O’Keefe [96] gives two definitions of “place cells”:

“cells whose firing rate or pattern consistently discriminates between different parts of an environment.”

or

“cells whose firing rate or pattern varies as a function of the animal’s location in an environment but can not be shown to be dependent on a single specific sensory input.”

The first definition, combined with the fact that the place fields — defined as the restricted region of the environment where the corresponding place cell is active — are distributed all about the maze [99], leads to the idea that the place fields partition the environment and that place cells,

as an ensemble, represent the environment (they “discriminate”). Indeed, there is a mathematical mapping between where the rat is and which place cell fires.¹³ Moreover, The second definition considers how these place cells are driven, that is, how “place recognition” is done for each place cell, individually, stressing the fact that place cell activity is the result of some fusion of processed multimodal information, or even the result of a highly “abstract” neurocomputation, instead of being like sensory cells, responding to single or relatively simple qualities.

What drives place cell activity?

Electrophysiological recording experiments in cue-controlled environments reveal how landmarks play a role in the location-selective firing of place cells. Muller et al. [89] (Figure 14a) showed that the position of a given place field is determined by its spatial relationship to a prominent landmark. Place cells were recorded while the rat was foraging after food pellets within a homogeneous cylindrical environment with only one salient landmark (a white cue-card on a black wall). The cue-card was then rotated to a different position on the cylinder wall in the absence of the animal. When the rat was subsequently put back in this environment, the place fields of the same cells shifted in register with the cue-card.

Place cell activity doesn't depend only on visual information, since it is sustained in the absence of the controlled cues or in darkness, as evidenced by O'Keefe and Speakman [103] and by Quirk et al. [113]. It has been suggested that dead-reckoning — the updating of the estimate of the animal's position by its own movement information — played a major role in maintaining the spatial representation in the absence of visual or auditory cues. Moreover, the influence of inertial information — which can play a major role in dead-reckoning — on hippocampal cell activity has been independently evidenced by recordings during passive motions of restrained rats [45]. However, dead-reckoning is known to be subject to integration errors and the position estimation quickly drifts [162]. Thus, the animal needs other means to update the estimate of its position. Tactile information (e.g. borders of the maze) may be one such means. Indeed, Hill and Best [53] showed that hippocampal cells of blind and deaf rats also exhibited spatial selectivity. However, after rotation of the radial arm maze in the absence of the animal, most place fields rotated in register. This suggests that in intact animals, extramaze cues are the principal basis for spatial selectivity, whereas blind and deaf animals only have access to intramaze cues. Hill and Best also suggested that the use of vestibular inputs could explain that the place fields of the few remaining neurons stayed fixed with respect to the experimental room after arena rotation.

Although visual, tactile, kinesthetic or vestibular influences can be shown, sensory information doesn't directly drive the place cells. Context and memory also seems to play a role, as suggested by Sharp et al. [134], who observed that place fields — in some cells but not others — were related to one cue-card and not to another identical cue-card placed in symmetric position within the same environment (Figure 14b). These place fields also depend on the point of entry of the animal into the environment in previous trials. In the figure, the place cell is “associated” to the cue-card at the left when entering the environment.

Position representation by the ensemble of place cells

¹³Although in open-field environments, place cell activity depends mainly on the animal's position [88], it also depends, in radial-arm mazes, on the animal's direction of movement [71]. In other words, a place cell can fire when the rat is running outward on a given arm but remain silent when the rat is running inward. Conversely, another place cell can fire when the rat is at this same location but only when the rat is running inward.

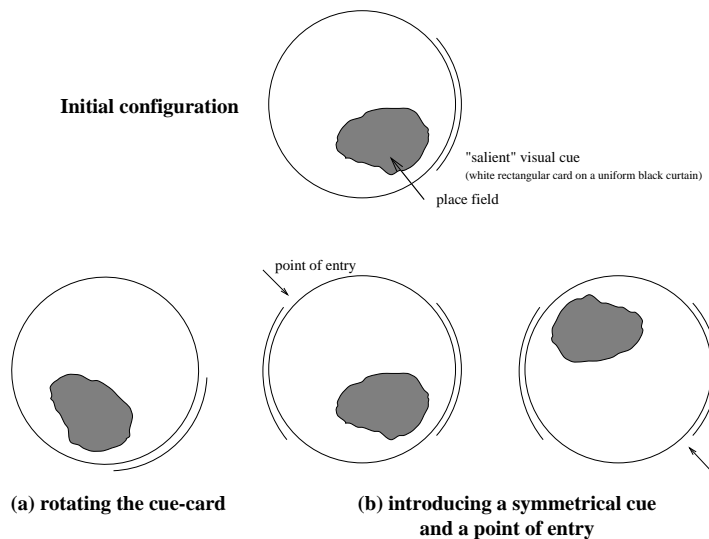


Figure 14: Manipulations of the apparatus influence place cell activity. Once the place field of a hippocampal cell is determined (Initial configuration), the environmental setup is manipulated in the absence of the animal, and the animal subsequently put back in it and the *same* place cell recorded. (a) Simply rotating the cue-card leads to an equivalent rotation of the place field (schematized on the basis of data from Muller et al. (1987)). (b) Introducing a second identical cue-card leads to a place field defined by one of the two cue-cards and the actual point of entry. (schematized on the basis of data from Sharp et al. (1990))

Place cells encode information concerning places in the environment in an abstract way, independent of (actual) sensory views of the environment. But how precise is the place recognition? If only a single place cell is considered, then its firing could be taken to indicate that the animal is or is not within a restricted region, which is a “low precision” positioning. Wilson and McNaughton [165] showed that even relatively small ensembles of place cells (about 100 neurons) could predict the animal’s location. Such a coding implies that the brain computes the barycenter (weighted average) of the (previously determined) place fields’ centers, weighted by the corresponding place cells’ activities. They found that the estimation of the animal’s position was all the better as they used more place cells to compute the population coding. This computation assumed that each place cell corresponds to a position in a Cartesian coordinate frame, described by the (x,y) coordinates of the place field’s center with respect to a reference frame chosen by the experimenter. It remains unclear how this spatial position coding by the place cell population might be used by the rat’s brain to guide navigation.

Place representation is independent of goal locations

Place cell activity is goal independent [137], that is, the place fields don’t change if the goal is moved within the cue-controlled environment. (However Breese et al. [15] reported that selectively delivering rewards to different locations in a cue-controlled environment led hippocampal cells to shift their place fields to the location where reward was available.) Goal locations are also not preferentially represented in the hippocampus and there is no evidence, to date, of a neural representation of future intended positions of the animal. In other words, it is not clear how the signals coming out of the hippocampus can be related to the goal approach and the selection of the appropriate direction or distance of movement.

Note that this independence of place representation and goal representation (if the latter exists) is consistent with the fact that rats can learn a reward location even if the reward is never consumed,

an example of “latent learning” [148].

Possible learning mechanisms at the neural level

Long term synaptic potentiation (LTP) — a persistent increase in synaptic transmission efficacy that can be induced by brief trains of synaptic stimulation — was first evidenced in hippocampal cells [17]. This synaptic mechanism produces modifications that resemble “Hebb’s postulate of learning” [51]:

“When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A’s efficiency, as one of the cells firing B, is increased.”

Thus, it has been suggested that this process is important for adaptive plasticity and memory functions of the hippocampus. Then, associative memories could be stored as a specific distribution of modifiable synaptic strengths. McNaughton et al. [72] presented evidence supporting the idea that spatial information is stored in hippocampal synapses — at least temporarily — during the formation of “cognitive maps”.

Most computational models simulate the learning process at different levels of their spatial representation. In this section, most models (excepting Blum and Abbott’s) learn the associations between landmark configurations and place cells. In other words, the place fields are defined during exploration, by learning the correspondence between places and landmark configurations. In models of higher levels of navigation strategies and in Blum and Abbott’s model (1996), learning concerns the coding of the fact that place fields are adjacent (topological representation) or the coding of the distances between place fields (metric representation).

Place cell activity and phase coding

Lastly, another property of hippocampal place cells, discovered by O’Keefe and Recce [102] (and confirmed by Skaggs et al. [136]), is the relationship between place cell firing and the phase of the theta rhythm, a sinusoidal EEG oscillation between 4Hz and 10Hz.¹⁴ As the rat runs through a place field, the corresponding place cell fires successively at “late”, “middle” and “early” phases, with respect to the theta cycles (Figure 15). O’Keefe and Recce proposed that the brain might utilize the phase of firing of a place cell to detect the position of the rat within the place field. This property will be crucial in the functioning of the model proposed by Burgess et al. [19], discussed below.

Spatial selectivity of cells outside the hippocampus proper

Non-hippocampal cells with activities selective for the animal’s position also exist but are not called place cells because of low resolution. In the rat, Quirk et al. [114] reported that cells in the entorhinal cortex, the major relay for sensory inputs to the hippocampus, have a spatial selectivity but are more “sensory bound” than place cells: entorhinal cells fire similarly in different cue-controlled environments (e.g. circular versus rectangular enclosures) at a fixed orientation to

¹⁴The theta rhythm (review: [81]) occurs essentially when the rat is performing “voluntary” movements including locomotor behaviors, such as walking, running, rearing and jumping [154]. However, the theta rhythm has also been shown to be triggered by passive displacements [46, 47]. The hippocampal EEG exhibits large amplitude irregular activity (LIA) during other “automatic” behaviors, such as eating and scratching, but also when the rat is in a state of immobile vigilance [20].




Figure 15: Phase shift in place cell firing during the rat's movement, as discovered by O'Keefe and Recce (1993). The horizontal axis represents both space (top, rat's movement) and time (bottom, theta cycle and spikes). When the rat is at A (it runs from left to right in this overhead view), the place cell corresponding to place field PF1 fires **late** in the theta cycle (spikes shown as vertical bars). In B, this cell fires at the **middle** phase and when in C, at the **early** phase of the theta cycle. C also corresponds to the point of entry into place field PF2 with respect to the heading direction. The corresponding place cell fires **late** in the theta cycle. (schematized on the basis of data from O'Keefe and Recce (1993) and from Skaggs et al. (1996))

landmark cues common to the two environments. In contrast, the hippocampal neurons re-map each environment in an independent manner, suggesting that the hippocampal neurons make an abstract sensory-independent representation of environments. Likewise, Sharp and Green [133] reported neurons of the rat subiculum, a major output of the hippocampus, which have rather low spatial selectivity. These observations suggest that these cells might encode another type of information with spatial components, or that population coding is used to analyze these signals. For these cells with low resolution or mixed spatial correlates, we will call "activity fields" the restricted region of the environment where the corresponding cell is active, similarly to the definition of place fields.

Electrophysiological evidence for a neural place recognition system has not yet been found in other species. For instance, despite many attempts, there has been no report of place cells in the hippocampus of monkeys, although space-related neurons [105] and view-selective cells have been reported [122, 123].

3.2.1.3 Head-direction coding

In addition to a positional representation, behavioral experiments have shown that some animals (e.g. rats and hamsters) have a sense of direction [84, 39]. A possible neural basis for this is provided by "head-direction cells" in rats. These neurons discharge selectively when the animal is oriented in a specific direction, regardless of its position in the environment. Such neurons have been found in several different parts of the rat brain (postsubiculum [142, 143], posterior cortex [27, 26], thalamus [141], and striatum [161]).

Like place cells, their preferred direction can be influenced by the rotation of the principal landmark cue in the environment [143]. Like place cells, their directional specificity and preference are disrupted if the rat is placed on the maze in darkness, but are restored when the light is subsequently turned on [85]. Moreover, their directional specificity is maintained briefly in darkness, if the light is turned




Figure 16: The position representation — shown by the distribution of place fields of 5 different hippocampal neurons (gray patches) — and the head-direction representation — shown by one reference direction (arrow) — are coherent and interrelated. Whenever one representation shifts (here the reference direction rotates by 90°), the other also shifts (the place cells coding for the maze arms code for the arms rotated by 90°). The position representation and the selection of the goal arm by the animal are also coherent: if the goal arm is indicated by the arrow in the left figure during training, the animal selects the arm indicated by the arrow in the right figure during trial, in accordance with the fact that the goal arm is assumed to be coded by the darkest place field by the internal representation. Schematized on the basis of data from Knierim et al. (1995)

off after the rat is placed on the maze [85].

Some head-direction cells (in anterodorsal nucleus) shift their directional firing preference as a function of angular head velocity [12]. It is thus postulated that the directional position of the rat's head is computed by integrating head rotations, for instance through vestibular information.

3.2.1.4 Coherence between position and orientation representations

There is a strong relation between place cells and head-direction cells [54]. When the preferred directions of the head-direction cells shift, after cue rotation for example, the hippocampal place fields also shift in register so that the two representations remain coherent (Figure 16).

Moreover, there is coherence between these representations and the goal-directed behavior in rats. O'Keefe and Speakman [103] showed that “the rat's choice of the goal arm at the end of the trial continues to show the usual spatial relationship to the [place] fields.” In other words, mistakes in the sense that the reward was not in the selected arm were not mistakes in the sense that the animal correctly remembered where the reward was with respect to the place representation.

3.2.2 Computational models

The place recognition-triggered response strategy was described in Section 2.4 as involving (a) self-localization (place recognition), (b) retrieval or determination of the local reference direction, and (c) association of a direction to the goal (or intermediate goal) from this place. Self-localization is the recognition of the place where the animal currently is. The problem of self-localization can be dealt with by mechanisms similar to those required for pattern recognition. A stored representation is compared to the sensory information available in the current location [169], that is, the current sensor readings — possibly from several modalities — are matched against stored values. The following models mainly differ in what types of sensors are used by the animals, how the sensor readings are coded and stored, and how this stored information is then processed to give a place representation. Most models (except Zipser's as well as Blum and Abbott's) rely on an internal compass to define

the reference direction. Thus, although the reference direction put forth in Section 2.4 was based on local landmark-configurations and therefore changed from place to place, the reference direction in these models is global and constant. Most models code the direction to the goal by a vector, that is, an angle relative to the reference direction. The length of this vector can correspond to an estimate of the distance to the goal. Brown and Sharp's model is the only one in this section that associates a movement to a place recognition, instead of a direction.

3.2.2.1 Zipser (1986)

In one of the first simulations of a computational model of hippocampal function, Zipser [168] starts from the geometrical proof that a point in 2-D space can be uniquely defined by its distances to 3 non-collinear landmarks. Assuming that the hippocampus receives highly processed information in the form of landmark identities and their distances with respect to the animat, he suggests a mechanism of how hippocampal place cells could be driven. In this neural network model, each cell is associated with 3 landmarks¹⁵ in the environment. Each cell discharges at a specific location defined by 3 visual angles subtended by the respective landmarks. The cell's activity is actually a graded function of the degree of match between the visual angles subtended by the 3 landmarks at the current location and those at the location associated with the cell.

Zipser [169] then proposes a goal-centered representation where the simulated cells guide approaches to a single goal from anywhere in the environment (Figure 17). However, instead of using place cells as in the previous model (Zipser, 1985), this animat uses "view cells", that is, cells whose activity depends on the animat's position *and* on its orientation. Each view cell corresponds to a place cell with an associated local reference direction.¹⁶ Similarly to place fields, the restricted region of the environment in which a view cell is active is called a view-field. The animat's current position is indicated by the activity profile of the view cell ensemble. In contrast, hippocampal place cell activity does not depend on the rat's orientation in open-field environments — such as the one simulated for Zipser's animat — although this holds true in radial-arm mazes.

The population of view-cells can be considered to represent the environment. This gradually occurs during exploration of different regions. When the animat starts at the goal location, none of the view-cells exist. The view-cells are then "recruited", i.e. created, at regular time intervals as the animat wanders. Meanwhile, the animat keeps track of its position relative to the goal through path integration (dead-reckoning). A "goal-cell" is also recruited when a view-cell is created and this stores the direction to the goal from the current location. Consequently, each view-cell is associated with a vector pointing to the goal location. This vector is defined with respect to the local reference direction of the view-field.

Once the spatial representation is built, the animat can return to the goal if it is disoriented or if the path integrator has been disabled, by using the information associated with each view-cell. However, there is no one-to-one mapping from locations to view-cells, since this would require too many cells. Thus, there is a need for a mechanism to infer the correct direction to the goal from locations that do not correspond to the center of a view-field. This is done by approximation, taking the information from view cells that are nearby (Figure 18). This idea of approximating a continuous point-to-point vector map by an incomplete population of discrete vectors will also be found in other models. At any location, several view-cells are simultaneously active, with activity levels proportional

¹⁵If more than 3 landmarks are available, each cell selects an arbitrary subset of 3 of them.

¹⁶However, this direction is not very precise because it is defined by the low-resolution egocentric orientations of the 3 landmarks associated with the place (each landmark is said to lie to the left, ahead, or to the right of the animat, with respect to the its current head-direction).




Figure 17: Zipser’s view-field model. Each view-cell in the model is associated with a view-field, defined by the distances from the center of the field to 3 landmarks and by an orientation. This orientation, which serves as a local reference direction, corresponds to the orientation the animat had at its first visit to the place, and is defined by where the 3 landmarks are with respect to the animat’s head-direction (left, ahead, or right). The direction to the goal from the view-fields is coded as vectors, oriented with respect to each local reference direction. The direction to the goal is interpolated from the activities and the vectors associated with all the view-cells (see text and Figure 18).

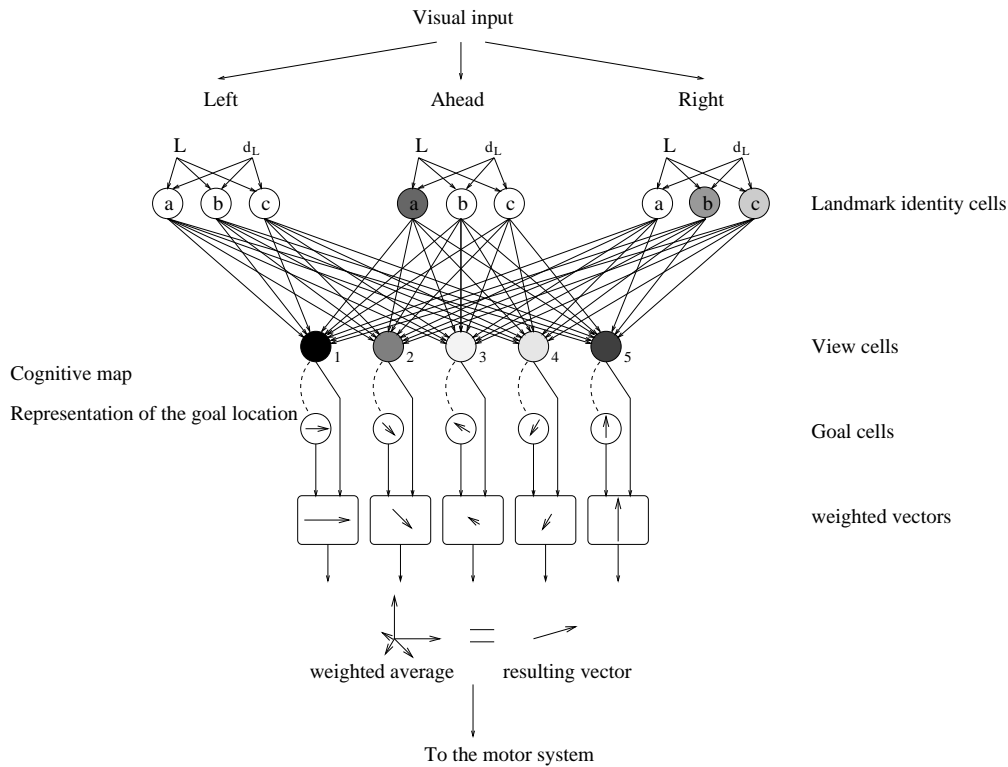


Figure 18: Zipser’s navigation system. The upper module (bordered in gray) of the neural network shows how the oriented place representation is developed by the view-cells. Each landmark (a, b or c), indicated by its identity L, is seen left, ahead or right relative to the head-direction and at a certain distance d_L . The landmark identity cells a, b, and c fire proportionally to the distances to the associated landmarks (d_L) when the associated landmarks are in their gross angular field (L is 0 or 1). For instance, in the position of the animat in Figure 17, landmark identity cell a for the “Ahead” visual field and landmark identity cells b and c for the “Right” visual field are active (shading) while all the other landmark identity cells remain silent. The combination of the landmark identity cells activates the view cells. Here, view-cells 1 and 5 fire the most and view-cell 3 fires the less (shading). The lower module represents the subsequent population coding of the direction to the goal. Each view-cell is associated with the direction to the goal, and the corresponding vector is represented by a goal cell. The vectors representing these directions are weighted by the activities of their associated view cells and the weighted average yields the estimate of the direction to the goal from the animat’s current position.

(figures after (Zipser, 1986))

to the match between the sensory information at the current location and the sensory information associated with the view-cell. Thus, one popular algorithm for the approximation process is to take a weighted average of the different vectors of the neighboring view-cells, weighted by these cells' activity levels. This is similar to simulations of the population coding in the motor cortex of monkeys [48, 21]. Another popular algorithm will be presented in the next section (Cartwright and Collett's model).

This model disagrees with the neurophysiological findings since hippocampal place representation is independent of the goal location [137]. Here, the exploration of the environment and the "map" building is subsequent to a goal specification and the spatial representation is goal-centered. New view-cells cannot be included in the spatial representation and be used to guide a goal-oriented navigation once the animat has lost track of the goal position: the animat would have to return to the goal, start the path integration process and go back (how?) to the place where it recruited the new view-field. Some of the other models presented below are inspired by Zipser's model, but try to palliate this limitation.

3.2.2.2 Cartwright and Collett (1987)

In the previously discussed model of Cartwright and Collett [23], the animat's navigation capabilities were limited to a small area around each goal, that is, the set of locations from which the model bee would be attracted to the corresponding goal. Cartwright and Collett [24] call it the "catchment area" associated with the snapshot taken at the goal. Each such area is defined a posteriori from the results of the simulations with the model bee, but it corresponds to the attraction basin of the discrepancy measure described in Section 3.1.2.1. To increase the range of the model bee's foraging area, that is to enable navigation in a large-scale environment, Cartwright and Collett [24] define an "album of snapshots, each taken at a *different* location within the terrain", and each associated with the information concerning direction and distance to the hive. Each snapshot thus corresponds to a "place" in the environment, from where the model bee has an associated "action" to perform to return to the hive — each place, once recognized, leads to a unique response.

Like in Zipser's model, during an exploration phase, the model bee stores a new snapshot whenever the discrepancy between its current retinal view and the best matching snapshot in its current album is greater than some threshold value. During this phase, the model bee "knows" the direction and the distance to the hive.

Figure 19 shows an example of a resulting vector-map, where each point in the environment is associated with a vector pointing approximately toward the hive. Note that all the locations within one place (defined by the closest location where a snapshot was made) correspond to the same vector (hollow arrows) as the one associated to the snapshot (filled arrows). This vector-map is thus an approximation of a continuous hive-centered vector-map, through a finite set of vector values. Consequently, the return to the hive follows a two-stage process: first the model bee follows the directions indicated by the closest snapshots (for instance from A to B and then from B to C) until the closest snapshot corresponds to the one defining the hive (C); then it uses guidance navigation strategy to reach the hive.

The assumption that the direction to the hive can be computed from each point where a snapshot was taken requires that the model bee has access to a reference direction, namely a compass sense, relative to which the vector gives the direction to fly. Thus, as in Cartwright and Collett's previous model (Section 3.1.2.1), there is a need for a mechanism that maintains orientation information.

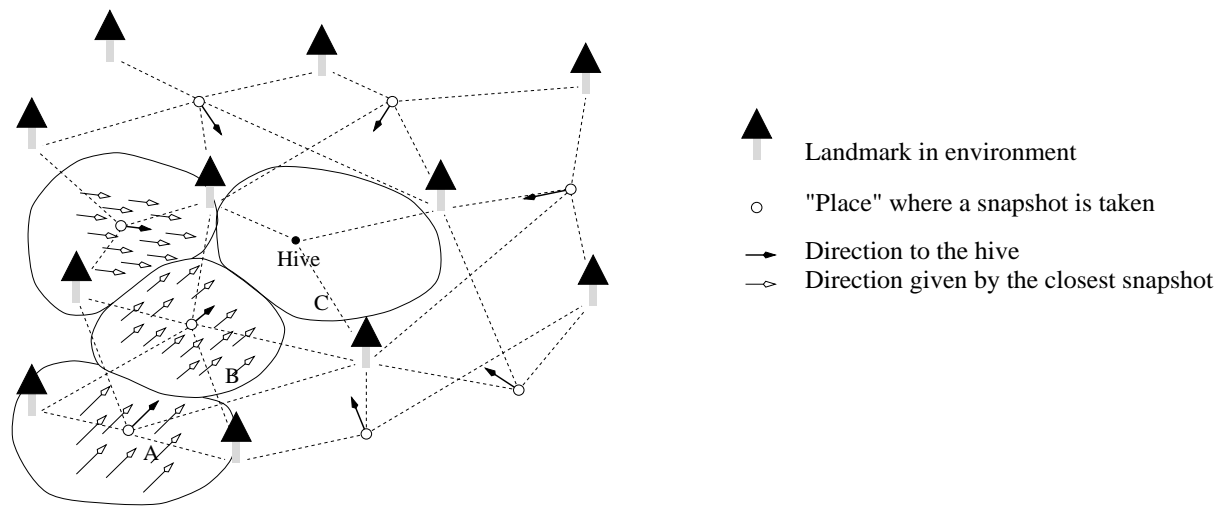


Figure 19: Cartwright and Collett’s navigation model (1987). For each place, defined as a location where a snapshot was memorized, dashed lines indicate the landmarks that are included in the snapshot. Continuous lines define the basins of attraction corresponding to the snapshots. Each snapshot is then associated to a vector (solid arrows) defining the direction and the distance to the goal with respect to a global reference direction. During navigation, the model bee compares its current retinal view(s) to all the snapshots in its repertoire. The one providing the best matching defines the current place the model bee is in. The model bee follows the direction (hollow arrows) corresponding to the previously learned vector required to return to the hive. When the model bee is in place C, which corresponds to the region containing the hive, there is no more vector information. The model bee then uses guidance (see Section 3.1.2.1) to reach the hive. Note that distance information is present but is not used here. The bee has an “album” of snapshots which define respective places and each with the associated hive-directed vector. This gives a hive-centered vector map.

3.2.2.3 Gaussier and Zrehen (1994)

While all of the previously described animats start from an a priori known goal location to build a goal-centered spatial representation, the animat proposed by Gaussier and Zrehen [44] is able to explore the environment (by random wandering) until it finds a reward and it subsequently builds a goal-centered spatial representation that ensures the ability to return to the corresponding place.

The originality of this model resides in the architecture of the navigation system, which consists of a variety of so-called PerAc¹⁷ modules. These modules are basic building blocks that implement an active perception-recognition process through feedback loops (Figure 20). The working principle of these modules is of biological inspiration, although the architecture doesn’t mimic any specific brain structure. A place is defined as a scanning sequence of landmarks. These are linked by successive movements (eye, head, or body rotations) fixating from one landmark to the next. The PerAc module associates the recognition of each individual landmark (module c) with the relative movement (modules c and d) that will bring the perception system onto the next salient landmark (the association is the result of learning). Thus, the global recognition consists of detecting and recognizing a landmark, executing the associated movement, and recognizing the next landmark in the sequence. Likewise, a landmark is defined as a sequence of salient points on the object, linked by ocular saccades.

However, the recognition is not based on the scanning sequence itself, which can vary on successive visits to the place, but rather on the set of the recognized landmarks with the associated movements. In practice, the place recognition scheme implements the notion of convergence of *object recognition*

¹⁷PerAc stands for Perception-to-Action.

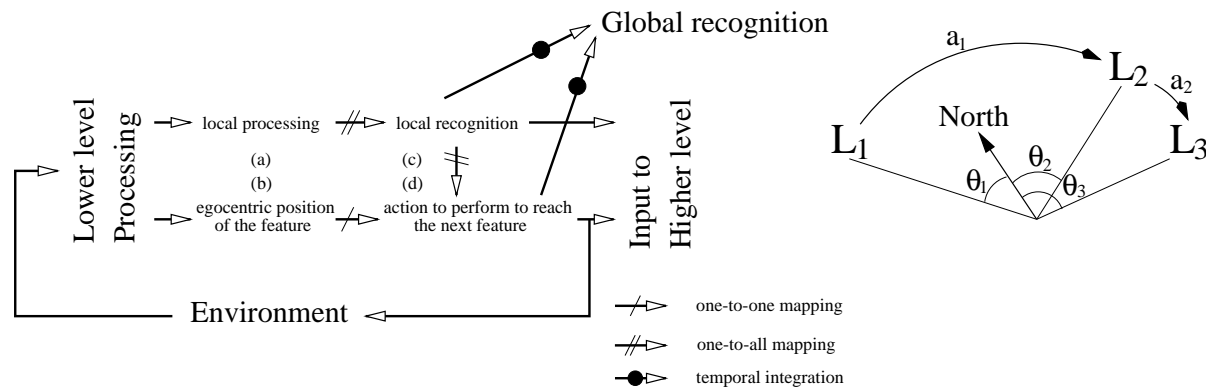


Figure 20: The PerAc module of Gaussier and Zrehen. Each small block in the module is a “neural map”, that is, a population of neurons that codes for different information: (a) complex features extracted from lower-level visual processing (contours from visual scene or landmark identities), (b) position of these features (in eye-centered or body-centered coordinates), (c) a stimulus pattern that enables the local recognition (categorization from the feature extraction. See Figure 21), or (d) action (ocular or locomotor command) associated with the current position (given by b) and with the currently recognized feature (given by c). As illustrated on the right, the extraction of features (landmarks L_1 , then L_2 and then L_3) and the knowledge of their egocentric orientations (θ_1 , θ_2 and θ_3) lead to the selection of an appropriate action (here, a_1 and a_2 are body rotations rather than ocular saccades) that triggers the system to orient toward the next feature. After (Gaussier and Zrehen, 1994)

and *spatial localization*, which Mishkin et al. [82] proposed to be the respective function of the “ventral and dorsal” streams in the primate visual system. One pathway (modules a and c) in the PerAc block codes for landmark identities (the ventral “what” pathway), another pathway (modules b and d) codes for the orientation of these landmarks in egocentric space (the dorsal “where” pathway), and a downstream module (Global recognition) fuses the two types of information through AND neurons, that is, neurons that are activated only when both inputs are simultaneously active (Figure 21).

This model thus recognizes landmarks and places through the same general mechanism. It is the first model that we have seen so far that doesn’t make the assumption that the landmarks are known and recognized a priori. Of course, there must have been some prior learning to permit characterization of the different objects that will be used as landmarks.

When a reward site is found, this triggers a reflex behavior to wander “around” it and to memorize how to approach it from different directions (Figure 22). The animat wanders in a circle about the goal, until the goal returns into its field of view (which is limited).¹⁸ The animat stops at this location, then “looks around” to memorize the landmark configuration that will define this place. Specifically, a “place cell” is recruited that will be selective for the current pattern of activity in the “global recognition” module (Figure 21). In this manner, the animat creates a population of “place cells” downstream of the PerAc modules described above. The animat further associates with this place, the direction to the visible goal with respect to absolute North as detected by its internal compass. It then heads directly to the goal (by guidance), and repeats the same wandering process until enough place cells have been recruited to cover the surroundings of the goal.

When the animat later attempts to return to the goal, it first recognizes the place it is in, represented by the most activated place cell (in other words, it looks for the closest place description matching the current visual scene). Then the animat retrieves the direction to the goal (the “re-

¹⁸It is implicitly assumed here that the animat can visually recognize the reward site from a certain distance. This model is also one of the rare ones that doesn’t have a 360° field of view.

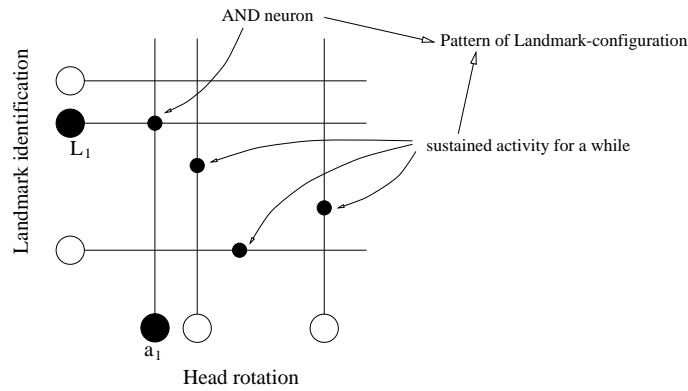


Figure 21: Gaussier and Zrehen’s global recognition module implementing the convergence of the *What* and *Where* types of information to yield a landmark-configuration pattern. The landmark identification comes from a lower level PerAc block. The simultaneous activity of a landmark identity neuron (*what* the animat is looking at) and a head rotation neuron (the angular movement made by the animat, that is, *where* the next predicted landmark is with respect to the current one) drives the activity of an AND neuron, which stays activated for a prolonged period of time. For instance, the inputs in this figure code for the recognition of landmark L_1 and for the suggested head rotation a_1 that leads to landmark L_2 , and thus correspond to the place shown on the right of Figure 20. The sequence of actions that lead from one landmark to another will yield the simultaneous activation of a subset of AND neurons, which results in a unique pattern of activation. The activation level of each neuron of the subset depends on the position of the landmark recognition within the temporal sequence. However, the pattern is recognized regardless of the activation level of each neuron. A “place cell” is associated with each characteristic pattern. After (Gaussier and Zrehen, 1994)

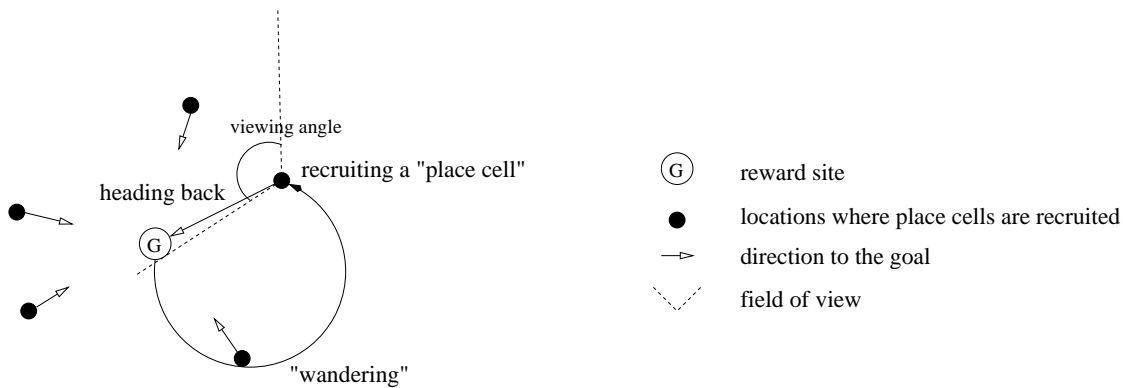


Figure 22: The acquisition of a goal-centered “action-map” (what to do at each location). When the animat encounters a reward, it wanders around the site in circles (one shown). At locations on these circles where the reward is visible again, the animat recruits a place cell and memorizes the direction to the goal from this place. Subsequently, the recognition of a place will be associated with the direction to the goal.

sponse”) associated with the place cell.

Recall that in Zipser’s model, place cells are recruited at regular time intervals and the animat has to update its estimation of the direction to the goal during the whole exploration phase. In other words, Zipser’s model requires an additional mechanism to give the directional information to the goal at any time (during exploration). In contrast, Gaussier and Zrehen’s animat exploits reflex behaviors. It relies on the fact that wandering in circles brings the goal back into view and gives the directional information as well as the triggering for place cell recruitment.

However, place cells in Zipser’s model are rather regularly distributed over the environment, like place cells in the hippocampus of the rat, whereas in this model, they are clustered around the goal.

3.2.2.4 Burgess, Recce and O’Keefe (1994)

The model of Burgess et al. [18, 19] derives most closely from physiological evidence of the hippocampal neurons. Their model incorporates to a greater degree than most the neurobiological data on the hippocampal formation.

Burgess et al. propose a 5-layer feedforward neural network (Figure 23) that largely reproduces the biological hippocampus architecture. The first of the five layers represents the highly processed information the hippocampus receives as inputs (sensory neurons) and the last layer represents neurons downstream of the hippocampus, and consists of “goal cells” whose existence is postulated by Burgess et al., but awaits experimental confirmation.




Figure 23: Burgess et al.’s feedforward network inspired by the hippocampus architecture. Cells (circles) in some layers are organized into clusters (ellipses). There are 5 clusters of 50 place cells, 10 groups of 25 subicular cells and 8 goal cells for each goal, corresponding to 8 head-direction cells. After (Burgess et al., 1994)

Each sensory neuron in this model discharges selectively when an associated landmark is visible at a certain distance from the animal.

Units in the next layer (EC) compute the product of two associated sensory neurons’ activities. Burgess et al. use a heuristic to hard-wire the connections between the sensory layer and the EC layer: the pairs of sensory neurons are chosen so that the associated landmarks lie at approximately opposite sides of the environment and so that the associated distances correspond approximately to the midpoint between the two landmarks. The result is that the activity fields of the EC neurons lie within the borders of the environment but do not overlap one another “too much”.

In rat hippocampus, O’Keefe and Recce found that place cells of area CA₁ tend to fire at a late phase with respect to the theta cycle as the rat is entering the corresponding place field, at an intermediate phase as the rat is running through the field, and at an early phase as the rat is leaving the field (Section 3.2.1.2). Burgess et al. postulate that this phase coding is transmitted downstream to the hippocampus from their EC layer. To determine when the EC neurons should fire with respect to the theta rhythm, Burgess et al. assume that the center of the EC activity fields approximately corresponds to the midpoint between the landmarks of the associated pair. The phase of firing of

the EC neurons is then computed by looking whether this point is located relative to the animat (Figure 24).




Figure 24: Firing phase computation yields α , the angle between the heading direction of the animat and the direction defined by the position of the animat and the centroid of landmarks 1 and 2; the phase will be “Late” if $|\alpha|$ is smaller than 60° , “Middle” if $|\alpha|$ is between 60° and 120° and “Early” if $|\alpha|$ is greater than 120° . After (Burgess et al., 1994)

Each of the cells in the EC layer is connected to 50% of the neurons in the next layer (PC). The synaptic weights of these connections are binary (0 or 1). Initially, only a small fraction of the weights are set to 1 and the remainder at 0. The neurons in the PC layer are organized into clusters and there is a winner-take-all process [126]) within each cluster. Consequently, only a limited number of PC neurons (one in each cluster) is active at any given time. When the animat moves, the activity in the network is propagated through from the EC layer to the PC layer. The active PC neurons “represent” the fact that the animat is in a given place that is defined by the conjunction of the activities of a given subset of EC neurons. They behave like place cells. However, the place fields that result from random synaptic weights are fragmented. Sharp [131] has previously shown that a Hebbian-type learning mechanism between sensory neurons (here the EC cells) and place cells (PC cells), coupled with the competitive mechanism, leads to a partitioning of the environment in “equivalent-size” unitary place fields. Thus, Burgess et al. introduce a learning mechanism during an exploration phase, inspired by Sharp’s work. The connections from the active EC neurons to the simultaneously active place cells are reinforced (set to 1) while the animat is randomly moving around in the environment.

A similar process takes place in the next layer (SC). The only difference between the PC and the SC layers is the size of the clusters. Clusters are smaller in the SC layer, so that each neuron competes with fewer neurons. As a result, activity fields in the SC layer are larger than in the PC layer.

This spatial representation is built up by the competitive network during a first phase of exploration. In a second phase of exploration, the animat learns the reward sites. There are 8 goal cells in the last layer of Burgess et al.’s network for each reward site in the environment (one cluster for each goal, as shown in Figure 23). Each of the 8 goal cells is associated with a direction. The 8 directions are represented by the head-direction cells and are arbitrarily labeled North, North-East, East, and so on. When the animat encounters a reward, it looks into each of the 8 directions. In each of these directions, the SC neurons that fire at a late phase with respect to the theta rhythm

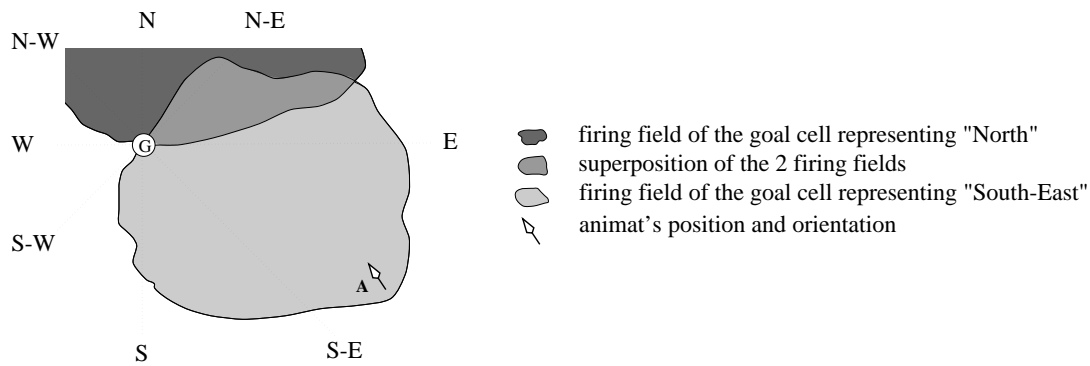


Figure 25: Examples of goal cells in Burgess et al.'s model. The cell with the field in dark gray discharges when the animat is north of the goal while the cell with the field in light gray discharges when the animat is generally south-east of the goal. Thus, the activity profile of 8 goal cells indicate the position of the animat relative to the goal.

correspond to activity fields that are *ahead* of the animat in the corresponding direction. In other words, when the animat looks North for instance, the activity fields of all the SC neurons that fire during the late phase of the theta cycle lie at the North of the reward site. Thus each goal cell draws connections from all the active SC neurons. The activity field of each goal cells then amounts to the sum of the activity fields of the active SC neurons, corresponding to the associated direction. The result of this learning process is that the goal cells partition the environment in 8 directional sectors (North, North-East, East, and so on) around the goal location. Figure 25 illustrates the activity fields of two goal cells.

Subsequently, when the animat is for example in location A in Figure 25, cells representing “South”, “South-East” and “East” of the goal location fire more than any of the 5 other goal cells. Taken as a whole, the population of goal cells indicates that the animat is globally “South-East” of the goal. This, in turn, triggers a signal to move to the “North-West” to reach the goal.

Figure 26 shows that this model can represent multiple goals. Each goal corresponds to a distinct set of goal cells but all the goal cells receive inputs from the same population of SC neurons. Obstacles can also be taken into account by considering them as negative goals.

EC neurons in this model behave like place cells. Thus, the contribution of the PC and SC layers for computations are questionable. In fact, a 3-layer feedforward version of this network that would consist of the sensory layer, the EC layer and the goal cells also enables successful navigation (Trullier et al., unpublished observations). The principle advantage of the PC layer seems to be in maintaining the architecture similar to the neurobiological model: the size of the PC neurons' firing fields look like the experimentally recorded place fields of hippocampal place cells in rats. However, the large activity fields of the SC layer are essential, so that the goal cells can be activated from locations distant from the goal.

Most importantly, this model is one of the rare ones that exhibit latent learning, that is, a goal-independent spatial representation.

3.2.2.5 Blum and Abbott (1996)

Blum and Abbott [14] propose a hippocampal-inspired model whose central feature is a goal-centered navigational map. Their model shows how an animat can solve the Morris water-maze task




Figure 26: Vector representation of the network’s output in each part of two square environments (viewed from above). The large circle is the goal location in each case. From points all over the environment (small circles) the navigation system indicates the direction to the goal (lines). The lengths of the lines correspond to the degree of activation of the goal cells, and roughly correspond to the distance to the goal. The same spatial representation (the first 4 layers of the network) is used for both goals. Two distinct populations of goal cells correspond to two distinct goals. Filled circles surrounding the environment are landmarks. (simulations by Trullier)

[86, 87] (see Section 3.1.2.3).

In this model, it is assumed that there are place cells whose place fields overlap one another. They are evenly distributed so that the whole environment is covered. It is further assumed that given two arbitrary locations in the environment, several of the place cells whose place fields are around each of these locations are synaptically connected through modifiable synapses (the weights are initially set at random).

During the first trials of the animat in the water-maze, the synaptic weights are changed through a modified Hebbian rule. When the animat reaches the hidden platform (by random movements), the model hippocampus receives a reinforcement signal that triggers the enhancement of the synaptic weights between place cells that simultaneously fired in the last part of the trajectory that led to the goal. Consequently, the activation of a place cell triggers discharges in other cells that have neighboring place fields along the path *in the direction* to the goal (Figure 27). The level of synaptic enhancement however depends on the time it took the animat to move from the place field to the goal. Paths that are more circuitous (and long) have weaker reinforcement of synaptic weights. When the environment is “sufficiently” explored, connection strengths are such that only place fields on direct paths to the goal are linked.¹⁹

Following the discovery by Muller and Kubie [90] that place cell firing can predict the future position of the animal by about 120ms, Blum and Abbott consider in their model, place cells that represent where the animat “intends to go”, beside the ones that represent where it currently is. Thus, during subsequent trials, the animat doesn’t move at random in search of the hidden platform, but moves into the place fields of the activated place cells. In other words, the current position (given as an input) activates the corresponding place cells, which in turn activate other place cells with neighboring place fields through the previously enhanced connections. Then, the model computes the difference between the current position and the intended future position as given by the activity of these latter place cells, and this output signal drives movement.

¹⁹However, it is not completely clear how a direct path to a distant goal and a circuitous path to a nearby goal are distinguished for this reinforcement mechanism.




Figure 27: Paths (bold arrows) that successfully led to the goal (black filled circle) are reinforced through the enhancement of the connections between the place cells whose place fields (dashed circles) are along the path (solid lines). The synaptic weights between these place cells are represented in gray-scale where place fields overlap. Consequently, when the animat re-enters the environment, the activation of the place cell corresponding to where the animat currently is drives the strongly connected place cell with the neighboring place field. The animat then moves toward this place field, which is on a path leading to the goal.

This model is the only one using the predictive nature of place cell firing, a property reported but once in the literature. This model is original in this respect, suggesting that the motor output can be directly computed from the hippocampal place representation. However, this model hippocampus doesn't represent goal-independent places but places on paths to the goal. This is in contradiction with the results from Speakman and O'Keefe [137] in hippocampal recording (cf. Section 3.2.1.2).

3.2.2.6 Brown and Sharp (1995)

All the preceding models in this section rely on an allocentric framework for directing movements. The direction to follow to return to the goal is defined as a vector with respect to absolute references as given by some internal registers like a compass or head-direction cells. Brown and Sharp [16] propose a model with a place representation and a head-direction representation. These control a movement coded egocentrically in the form of left and right turns.

The place representation is based on a simplified version of an earlier model by Sharp [134] (Figure 28). It consists of a 3-layer feedforward network. The first layer represents the sensory inputs to the navigation system, the second layer builds a place representation from the latter while the head-direction is given a priori. Place cells and head-direction cells project onto cells of the output layer, which is divided into left-turn and right-turn clusters. This last layer produces motor commands, consisting of a small step forward and a turn to the left or right.

Each sensory cell in the first layer is activated whenever a given landmark is within a certain distance from the animat, and at a fixed angular relationship with respect to the head-direction.

These cells project onto a layer of hippocampal place cells which is organized into a winner-take-all cluster (a mechanism also used by Burgess et al. [19]), so that only one hippocampal cell is active at any one time. Thus, each place cell discharge corresponds to a set of active cells in the sensory layer, capturing the conjunction of specific sensory features, that is, a specific landmark configuration in a particular part of the field of view of the animat (that is, with relation to the heading direction).

Each of the output layer cell clusters is organized as shown on Figure 29. There is a one-to-one mapping from each place cell to interneurons in the cluster. Each interneuron is in turn connected to all but one output cell. Each head-direction cell is connected to all of the output cells.

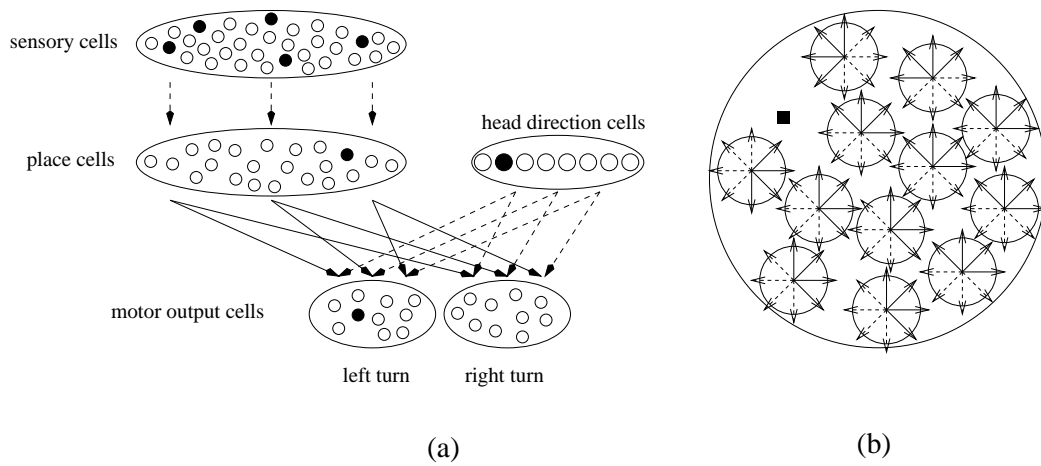


Figure 28: (a) Brown and Sharp’s neural network maps a place representation and a head-direction representation onto an egocentrically-coded movement representation. Continuous arrows are fixed connections and dashed arrows are connections with modifiable synapses. Filled circles represent activated cells while hollow circles correspond to silent cells. (b) The simulated environment (landmarks not shown). Each place (circle) and each head-direction (arrow) correspond to an unique movement toward the goal computed by the network (left-turns shown as continuous arrows and right-turns shown as dashed arrows). The goal is indicated by the black square. After (Brown and Sharp, 1995)




Figure 29: Detail of an output layer cell cluster. Each place cell activates a given inhibitory interneuron that inhibits all but one output cell. This output cell is activated by the population of head-direction cells. After (Brown and Sharp, 1995)

As a result, the place cell representation selects in each cluster one output cell that may be activated by the head-direction cells. The synaptic weights between head-direction cells and output cells determine the place-dependent mapping from direction to left-or-right motor output. The animat subsequently moves a little step forward and turns in the direction indicated by the most active of the two activated output cells.

The synaptic weights between the sensory cells and the place cells are modifiable, as are those between head-direction cells and motor output cells. The synaptic weights between place cells and the inhibitory interneurons in the motor output clusters are fixed, as are those between the inhibitory interneurons and the output cells.

The synaptic weights between the sensory cells and the place cells are modified according to a Hebbian rule. The details and the properties of this place learning process are similar to the process described in Section 3.2.2.4, although the weights here can have real values between 0 and 1, instead of being binary.

The synaptic weights between head-direction cells and motor output cells are modified according

to a reinforcement learning rule (modified Hebbian rule). Weights are changed only when the animat encounters the reward in the environment, and the changes depend upon the time-history of the synaptic activity, that is, how often and how recently the pre- and post-synaptic cells fired. This rule follows the idea that recent actions — which correspond to the simultaneous firing of a given set of head-direction cells and of a given motor output cell — most likely led to the goal location, so that the corresponding connections should be reinforced.

The network devised by Brown and Sharp implements a goal-dependent action model where there is a mapping from head-direction to movement at each place. In fact, the mapping takes the form of two functions, one for each of the possible movements (left turn or right turn). Each of the functions computes the “degree of relevance” of performing the corresponding movement on the basis of the current place and head-direction. The navigation system then selects the movement corresponding to the highest value. This is reminiscent of Q-learning in the reinforcement learning literature [157]. Q-learning is an algorithm that makes use of an utility function $Q(x,a)$, which estimates the long-term reward of performing action a when the system is in state x . When the goal is attained, the learning algorithm backpropagates the reinforcement signal to update this function (as in Dynamic Programming). This function Q is usually represented by arrays but can also be represented by neural networks [5, 64].

One difficulty with this type of learning is that if the animat encounters the reward after a long sequence of actions, the system can’t determine which of the series of actions were important. This is the problem of *delayed reward* [158]. Brown and Sharp acknowledge that “only actions which take place quite close to the goal can be rewarded.” The solution they propose is to generate “anticipatory goal responses”, a subject for future study.

3.2.2.7 Krakauer (1995)

Another computational model of animal navigation [56] is a neural network implementation of a “hierarchical snapshot model” inspired by the work of Cartwright and Collett [23, 24]. A first-level neural network receives filtered input from distal landmarks (identities, distances and bearings) and learns to have only one output channel active at a given time. This channel represents the “catchment area” [24], of one specific snapshot.

This output activates a second-level neural network that receives a filtered input from proximal landmarks (identities, distances and bearings), and is trained through backpropagation to produce a vector directed towards the goal.

This animat thus navigates according to a place recognition-triggered response strategy.

3.3 Topological navigation

3.3.1 Behavioral experiments

The navigation strategies presented up until this point could not be used for planning the whole sequence of movements from the current location to a goal because all of the spatial information available to the animal at a given moment concerned local spatial relations, that is, restricted to the current place. However, it is also important to choose between paths and to plan actions. This requires facilities for accessing information about places that are beyond the current field of perception. To start with, we will look at evidence for an internal representation of the topological layout of different places, that is, how places are connected to one another.

Thinus-Blanc [144] demonstrated that, within a certain range of metric distortions in the apparatus, hamsters were able to discriminate between two topologically different experimental setups, as shown in Figure 30. The environment was a cylinder (1m in diameter) within which was disposed a circular barrier (9cm high, 45cm in diameter during the training trials), which divided the arena into an interior region (A) and an exterior region (B). The rats were capable of jumping over the barrier. The barrier had a removable section (15cm wide). The task was to avoid a mild electric shock that was applied in one of the regions of the environment (area A or B) by moving to the other. The region where the electric shock was delivered was defined by the topological property of the region A (open or closed).

Hamsters were trained in the task and were then tested in the same environment, but with different barriers. The testing barriers varied in diameter (from 17cm to 80cm) and in height (from 1.5cm to 12cm). The removable section was always 15cm wide, except for the barrier that was 17cm in diameter, where the removable section was 7.5cm wide. Hamsters exhibited correct responses in a very few cases, for barriers whose dimensions were very close to the dimensions of the training barrier. When the barrier was closed, the performance decreased more when the height was changed than when the diameter was changed. When the barrier was open, changes in height and diameter had equivalent effects. Thus, hamsters could attend to the topological property of this experimental environment, but they were disturbed by metrical properties. The difficulty came also from the fact that hamsters have a natural tendency to move to the periphery of an arena and also that in the case of closed barriers, changes in height led to changes in locomotor movements required for moving from one area to the other.




Figure 30: Two topologically different setups used by Thinus-Blanc in her conditioning experiments. The rats were to avoid the region where there is an electric shock (-) and run to the other region (+). The correct response was determined by the topology (open/closed) of the region A. The shock placements were reversed for a second group of rats. After (Thinus-Blanc, 1978)

Cats are also influenced by metric properties of the environment in tasks where only topological properties are relevant for performance. Poucet [109] studied the processing of visuospatial information in cats required to choose between two paths to reach a visible reward (Figure 31). Two white paths were drawn on a completely black floor and these paths were bordered with transparent walls, preventing the cat from walking outside the paths but enabling it to visually inspect the whole environment. One feeding bowl was placed at the end of one path in one of the three possible arrangements and the hungry cat was released at the starting point. All three environments were topologically equivalent. All cats were able to choose the correct path in environment 1 and their performance in environment 2 was significantly above chance level. However, they tended to choose the wrong path in environment 3. In environment 1, moving directly toward the visible goal leads to the correct choice. In environment 2, the direction to the goal is straight ahead, and thus ambiguous. But cats were able to attend to the fact that the path was circuitous, a topological property of the environment. In environment 3, however, visual attraction and topological properties were contradictory. The results showed that visual attraction, a guidance strategy, tended to dominate.

This “perceptual dominance” will also be found in experiments testing the ability of animals to attend to metrical properties.

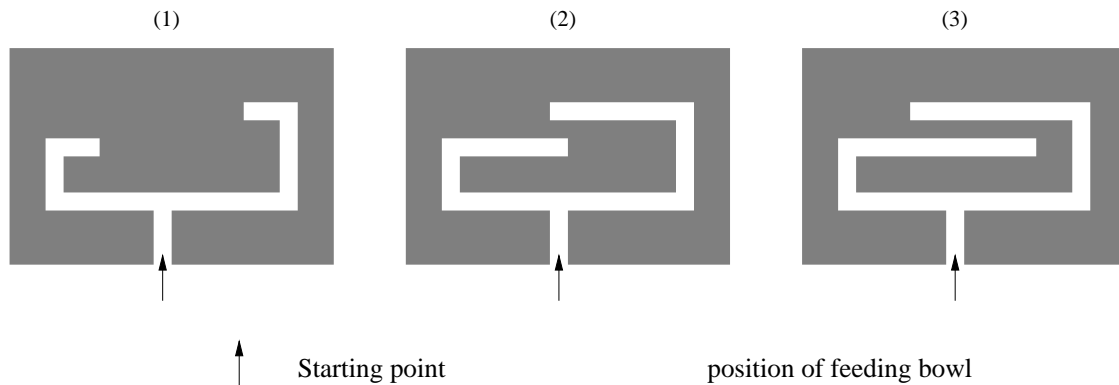


Figure 31: Experimental setups of the path-selection problem used by Poucet. In all three setups, the feeding bowl is put randomly on either path on successive trials. Here, the correct choice is to turn to the right at the choice point in all three setups. In (1), cats usually select the correct path. In (2), cats select the correct path with probability well above chance level. In (3), cats usually select the wrong path. After (Poucet, 1984)

Some of the experiments in the behavioral literature that were interpreted as indicating the ability to make detours using metric internal maps can be accounted for by a topological navigation strategy, without the need to take into account metric information. Tolman and Honzik’s detour problem (Figure 32) is such a case. In this experiment, a rat is required to select one of three paths leading to a reward. It quickly learns to use path 1, which is the shortest one. When barrier A is put to block path 1, the rat shifts, after a few trials, its choice to path 2, the second shortest path. However, when barrier B is put in place instead of barrier A, while the rat is using path 1, the rat shifts its choice to path 3, *without* trying path 2. The authors’ interpretation was that the rat has the “insight” that the barrier blocks simultaneously paths 1 and 2. It turns out that such an “insight” doesn’t necessarily invoke a map-like, metric representation of the environment²⁰ because the detour problem is here a *path selection* problem, that is choosing one complete path among a *finite* set of

²⁰However, the fact that a rat prefers path 1, which is the shortest, demonstrates its ability to compare distances, and therefore some knowledge of metric information.

possibilities. The presence of barrier B suppresses the adjacency between place p1 and place p2. Consequently, it suffices for the topological representation to be modified accordingly (by removing the edge corresponding to the connectivity between node p1 and node p2) to induce the choice of path 3, the remaining path leading to place p2.




Figure 32: Tolman and Honzik’s detour problem. The path usually taken by the rats is path 1, but when the barrier A is put in place, the rat shifts its choice to path 2 after a few trials. When barrier B is put in place instead of barrier A, and the rat tries path 1, the rat then shifts its choice to path 3, without trying path 2. After (Tolman, 1948)

3.3.2 Electrophysiological evidence

3.3.2.1 Connections between place cells

Place cells in the rat hippocampus are interconnected, for example through recurrent collateral connections in the CA₃ region. Recurrent collaterals are distributed to other CA₃ cells throughout the hippocampus [2] and the connectivity is about 4%. Therefore, any two hippocampal cells in the CA₃ region are connected together directly or via 2 or 3 intermediate steps [121].

This interconnectivity between place cells is one possible evidence for the existence of physical links between place representations. Furthermore, since these connections have modifiable synapses [17], one argument usually put forward is that they could code for some *learned* spatial relationship between the corresponding place cells, such as proximity and overlap. Thus, hippocampal cells could implement a topological graph of the environment. The place cells could be the nodes and the connections could be the edges. An implicit assumption in this case would be that one synapse is sufficient for one place cell to drive the next, which is unlikely to happen in the rat hippocampus. A more reasonable assumption is to consider each model place cell as an ensemble of place cells coding for the same region in the environment, so that many simultaneously activated cells can drive the next ensemble.

However, if two cells have overlapping place fields in one environment, they do not *necessarily* have overlapping place fields in a second environment [99, 89, 147].²¹ Thus, there doesn’t seem to be any *a priori* topological relation encoded in the anatomical connections between two place cells.

The connections between hippocampal cells could correspond to other associative connections, which could be equivalent to topological relationships within one given environment. Skaggs and McNaughton [135] reported that two cells with overlapping fields simultaneously recorded in an environment tended to fire more synchronously during the following sleep period of the animal, in the same temporal order as during the spatial behavior. One interpretation of these results is that the

²¹Eichenbaum et al. [38] however showed that the place fields of anatomically close place cells are significantly clustered — “closer in space and considerably more overlapped than chance” — although these cells were not tested in a second environment.

two place cells were connected and that spatial exploration tended to enhance the synaptic efficacy from the first place cell to the second one.

3.3.2.2 Whole body movement representation

Assuming the rat hippocampus implements a topological representation of the environment, a command to move from one place to a neighboring one is required. If a route is defined as a sequence of places to traverse, what is required is information concerning the movement that leads the animal from one place to the next. McNaughton [69] suggested that place representations are linked through movement representations, in the form of a “transition matrix” (Figure 33). The idea is: “when starting from place B (implicitly facing place D) and making a left-turn at the choice point, the animal ends up in place C.” In other words, the spatial representation incorporates simultaneously places and movements, the latter corresponding to the links between the former.

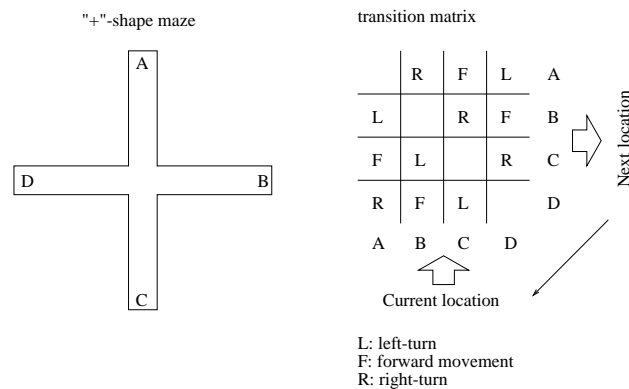


Figure 33: A transition matrix maps places to places through specific movements. After (McNaughton, 1989)

There is evidence for “movement” representations in the rat brain. McNaughton et al. [69, 74] recorded in the primary sensory-motor cortical area and the posterior parietal cortex of rats, and found cells responding to one of three broad categories of movements: left turns, right turns and forward motion. It is not known if these cells are somatosensory, driven by vestibular information or by visual information. But McNaughton et al. tested for somatic and kinesthetic responses by passively moving the body of the rat, and found no responses.

Likewise, O’Mara et al. [104] recorded neurons in the macaque hippocampus and Gavrilov et al. [45] in the rat hippocampus, that discharged selectively during passive displacements of the animal. This suggests that self-movement information derived from inertial cues can influence hippocampal cell activity.

Lastly, rat hippocampal place cells have been shown by Sharp et al. [132] and Wiener et al. [163] to be influenced by passive rotation of the animal, suggesting that hippocampal spatial representations are updated by active self-movement information.

3.3.3 Computational models

Models in this section allow computations to exploit the topology of the spatial layout of different places and will hence enable path planning.

Models described in previous sections built a place representation from a set of visible landmarks, partitioning a small-scale environment into “small” places. In contrast, although some of the models

in this section also deal with small-scale environments (Penna and Wu [106]; Bachelder and Waxman [4]), most models in this section deal with places in a large-scale environment, where there may be landmarks that are not simultaneously visible. Consequently, the models don't incorporate lower levels of information processing, that is, the animat is assumed to be able to recognize the place it is situated in (as well as adjacent places, in Schmajuk and Thieme's model [127]). The necessary information is provided to the navigation system, assuming, for instance as Levenick [61] does, that there are units representing places and these result from upstream spatial pattern recognition. The recognition of a wide range of "places" is a difficult problem, not yet resolved, though robots devised by Kuipers and Byun [58] or by Mataric [68] include place recognition systems, as defined by their guidance mode of navigation (Section 3.1).

3.3.3.1 Muller, Kubie and Saypoff (1991)

Assuming place cells are given a priori and using the two properties on hippocampal connectivity described in Section 3.2.2.5 — (i) place cells are sufficiently interconnected so that there are two interconnected place cells for any two given locations in the environment, and (ii) connections have synapses modifiable through a Hebbian rule — Muller et al. [91] propose a model of how the hippocampus might implement a topological representation of the environment. Their model also encodes some metric information. The animat exhibits topological navigation and is able to select paths on the basis of the path's length, but doesn't exhibit metric navigation because it is not able to make metric detours or metric shortcuts.

As in virtually all models, the connection weights are learned during random exploration. When the rat moves (in a straight line at a constant speed) from one place to another, the corresponding place cells fire with a temporal delay that is proportional to the distance between the place fields. The closer the place fields, the more probably the place cells will fire simultaneously, and thus the more often the (real-valued) weights of the synapses connecting these cells will get enhanced. The synaptic weights eventually become proportional to the (inverse of the) *distance* between place fields, a metric quantity. The temporal order of firing of the two place cells are not taken into account. Thus, when the animat moves from place A to place B, the synaptic weights from cells representing A to cells representing B, as well as those from cells representing B to cells representing A, are modified. However, these connections between place cells do not contribute to the firing of these place cells, in contrast with the model of Blum and Abbott. They only passively store the spatial information.

The animat needs to extensively explore the environment for the connection weights between all the place cells to be learned. However, the initial values of the weights are zero. Thus, after some exploration, non-zero connection weights signify that two corresponding place fields overlap, although they don't correspond to the exact distances. This proves to be sufficient for topological path planning, thus extensive exploration is not required.

Once learning is complete, the animat is able to use this implicit spatial representation — the "cognitive graph" as Muller et al. call it — to plan a route from the current location to a given goal location. Muller et al. propose a classical graph search algorithm, employing as yet undetermined physiological mechanisms. The goal location is defined by one place cell and the current location is defined by the place cell that is currently active. Each "neural" path, i.e. each ordered set of connected place cells, corresponds to a physical path in the environment. The most direct path in the environment is represented by the set of connected place cells for which the sum of the synaptic weights is the greatest (maximizing the sum of the inverses of the distances).

As explained by Muller et al. [92], the addition of a few ad hoc mechanisms enable the generation of topological detours and shortcuts. The addition of an obstacle, requiring a detour, is represented by the removal from the graph of the place cells occupied by the obstacle. Thus, paths that previously included these places are no longer taken into account by the graph search algorithm. The appearance of a hole in the obstacle, enabling a shortcut, requires a new phase of exploration that recreates the adequate links in the graph.

An interesting aspect of this model is that a simple Hebbian learning rule — which transforms the temporal characteristics of cell firing into a code for proximity relationship — enables the implicit building of a *metric* map. Indeed, all the necessary information to reconstruct a physical map are available in the connections between the place cells. As shown on Figure 34, the knowledge of the distances between all the place fields uniquely determine their relative positions (except the crucial left/right relation).

However, the model proposed by Muller et al. only exploits the topological information encoded and uses the metric information only to select between different paths.




Figure 34: Implicit metric spatial representation. Left: the physical *metric* layout of the place fields (dashed circles) can be reconstructed on the basis of sufficient information about distance between them. Right: these distances are implicitly encoded in the connection weights (labels on bidirectional arrows, arbitrary units) between place cells (small circles).

3.3.3.2 Kuipers and Byun (1991)

The model proposed by Kuipers and Byun [58] also constructs a topological graph of the environment during the exploration phase. Its original aspect resides in its interactions with a (lower-level) guidance navigation strategy.

As described in Section 3.1.2.4, places are defined as specific locations that the animat reaches by using a guidance strategy. The place representation is not given a priori. Likewise, paths can be defined by a guidance strategy. The animat follows a wall (“move-along-object-on-right”) or a corridor (“follow-the-midline of the corridor”), or moves toward a specific remote landmark (taxi).

The spatial representation is built as the animat selects distinctiveness measures to define new places or by control strategies to follow a new path until another place is found. It consists of a graph where nodes are places and edges are paths between places. However, in contrast to the other topological graphs described in this paper, there are many different types of information that can be associated with the nodes and the edges. These associations are implemented symbolically.

Firstly, each node stores (1) the definition of the corresponding place by its distinctiveness measure, (2) the identities of the sensors that are involved in the computation of this measure, and (3) the

values of the sensor readings at the local maximum of the measure. Likewise, each edge stores the definition of the movement control strategies that enable the animat to follow the corresponding path.

Secondly, any information that can help disambiguate different places is also stored. If the animat was equipped with appropriate sensors, it could memorize non-geometrical features in the vicinity of the place, such as the colors or the texture of the nearby objects. The animat also memorizes local metrical information, such as the distances and direction to nearby objects, their shape and apparent extent, or the directions to open space. Likewise, the animat is equipped with a compass and thus can memorize the direction of travel along a path. It can also measure the distance traveled.

The topological representation is augmented with metrical information. Thus, the system generates predictions and plans for exploration to disambiguate sensorily similar places. Figure 35 shows one example and illustrates the fact that here, exploration and path planning are simultaneously performed. When a place (P4) is found that is similar to an already known place (P1), the animat tries to match its spatial representation with its sensory information. It then plans a route (follow direction 1) that will enable it to check its hypotheses (it should find place P2). This is called a “rehearsal procedure”. Path planning is not detailed in computational terms in the paper but follows a classical graph search with the use of metric information.




Figure 35: The “rehearsal procedure” performed by Kuipers and Byun’s animat. The animat starts from place S and first finds place P1. From there, it selects direction 1 (arrow) which is defined with respect to compass. It also memorizes the existence of an open space in direction 2. It then successively finds P2 and P3. When it finds place P4, which has sensory attributes similar to place P1, it first tries to take metrical information into account (distance and direction traveled). With the hypothesis that P4 might be the same as P1, it generates a prediction about the next place to be found by selecting direction 1 again. Since it finds a new place P5 instead of P2, the system updates its spatial representation by storing place P4.

It would seem that since distance and direction measurements are available, this model is able to perform metric navigation. Indeed, it can reconstruct the metrical layout of the environment, but Kuipers and Byun use it only for the rehearsal procedure. Since the animat can only generate movements referred to environmental features, it cannot move into a completely open space and perform metric shortcuts for instance.

3.3.3.3 Mataric (1990)

Mataric’s model [68] is also based on the construction of a topological representation of an indoor environment, with a place representation derived from the guidance strategy (described Section 3.1.2.5). The basic behavior of the robot is to follow walls and the environment is decomposed into “right-northward-walls”, “left-eastward-walls”, “southward-corridors”, and so on (Figure 36).

As the animat wanders in the environment and detects new places, it recruits new nodes to represent them and links them to existing nodes through movement information, that is, comparing compass readings at both places. Each node also stores the size of the place (the distance between one end and the other end of the wall). This is a simple graph growing mechanism.

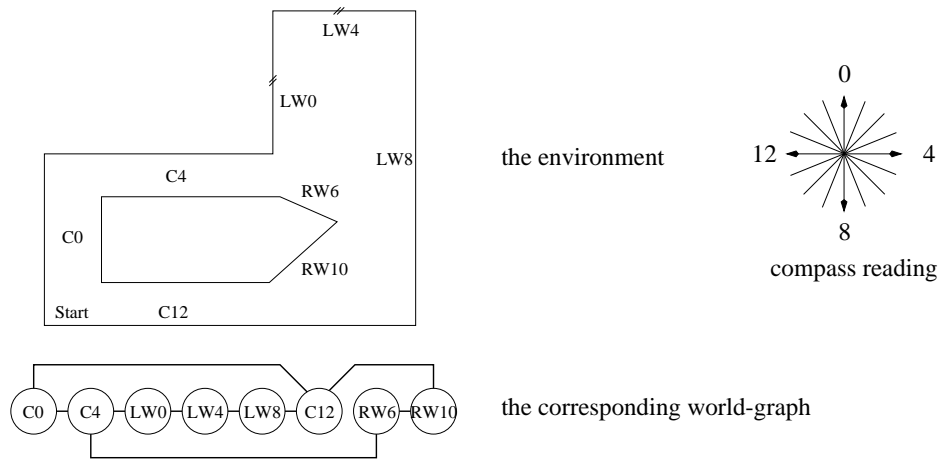


Figure 36: Example of an environment and of the associated graph in Mataric's model. *Left above:* each place in the L-shaped room is labeled according to the characteristics of the range sensor readings (RW: right wall, LW: left wall, C: corridor) and to the compass reading (*right:* from 0 to 15) when the robot moves about, using the wall-following procedure (lower-level guidance strategy). *Below left:* in the graph, each place is represented by a node and any two nodes can be linked dynamically (lines) to code the adjacency relationship. After (Mataric, 1990)

Orienting toward the goal is accomplished by using a spreading of activation through the graph, from the node representing the goal to the node representing the place the animat is currently in. Once a goal location is defined, the node representing it keeps sending a signal which is propagated from node to node through the network. The propagation speed is proportional to the size of each place. When the current place node receives the signal, the adjacent place lying on the shortest path corresponds to the adjacent node from which the signal comes from. The direction to this adjacent place is retrieved from the topological graph and the animat moves in this direction.

As in the model of Muller et al. [91, 92], the animat is not able to generate metric shortcuts through previously unvisited places, but the spatial representation includes sufficient metric information so that the animat can select between alternate paths on the basis of their lengths.

3.3.3.4 Schmajuk and Thieme (1992)

Schmajuk and Thieme [127] propose a model of animat navigation based on a neural network with specific dynamical properties of activity propagation that can represent the topological layout of places and thus enable path planning. The model is composed of two modules, one that encodes the topological representation and one that selects movements on the basis of predictions generated by the first module (Figure 37).

As in many models, place cells are defined a priori. The navigation system's inputs consist of the identity of the current place and the identities of the adjacent places (called "views"). As shown on Figure 38, the neural network has dedicated neurons for each place and each view. As the animat moves from place to place, the synaptic weights between the place nodes and the view nodes are modified through a generalized Hebbian learning rule. The weights are initially set at random values. During learning, the activation of the current place node is propagated through the

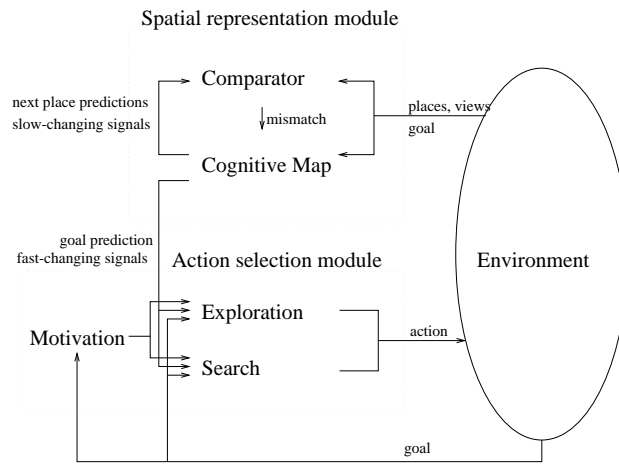


Figure 37: Functional schema of Schmajuk and Thieme’s navigation model. The first module (above) encodes the topological relationships between places by constantly comparing the internal predictions of what the animat expects to see and what it actually sees. The second module (below) uses the predictions generated by the first module to select appropriate movements to reach the goal. After (Schmajuk and Thieme, 1992)

network and place-prediction nodes are activated according to the synaptic weights. The predicted places should correspond to current views. Thus, the activities of the place-prediction nodes are compared to the activities of the view nodes. The synaptic weights are updated whenever there is a mismatch.²² When learning is completed, these synaptic weights represent the transformation from places to views and, since views correspond to adjacent places, they represent the adjacency properties between places.

The activities of the place-prediction nodes then propagate back through the recurrent connections and drive the place-prediction nodes again. Their activities now represent the predictions of places that are adjacent to the current views, that is, places that are two steps away from the current place. For instance, assume the animat is in place 1 of Figure 38. Nodes p2 and p3 will first be activated since places 2 and 3 are “viewed” from place 1. The activities of p2 and p3 feed back to drive the nodes p1, p3, and pG, and the nodes p1 and p2, respectively, through strongly weighted synapses. Thus, the existence of a path leading to the goal from place 1 can be “inferred”. However, because such recurrent signals could create interferences with local signals perceived at the current place, Schmajuk and Thieme assume that there are two rates of change of the transmitted signals, fast and slow (the dynamics of the activities of cells are governed by differential equations). The recurrent connections involve an intermediate layer of neurons (labeled *d* in the figure) that output the temporal derivatives of their inputs. Consequently, only fast-changing signals are efficiently propagated through the recurrent connections. Moreover, these derivatives are small, ensuring that the recurrent activity is always weaker than the direct activity. Repeating this looping process continues to reduce the signal, so that the activity of a place-prediction node is inversely proportional to the topological distance (number of intermediate places) between the current place and the predicted place.

Only the slow-changing signals are sufficiently large enough to modify the synaptic weights. Consequently, prediction of distant places with the weak, fast-changing signals does not lead to a mismatch detection (for instance, node pG is activated by the prediction of p2 from place 1, but there is no

²²the weight between place node P and view node V is increased whenever the animat is at P and perceives view V, that is, the corresponding nodes are simultaneously activated. The weight is decreased if the animat is at P but doesn’t perceive view V, that is, the corresponding nodes are not simultaneously activated.

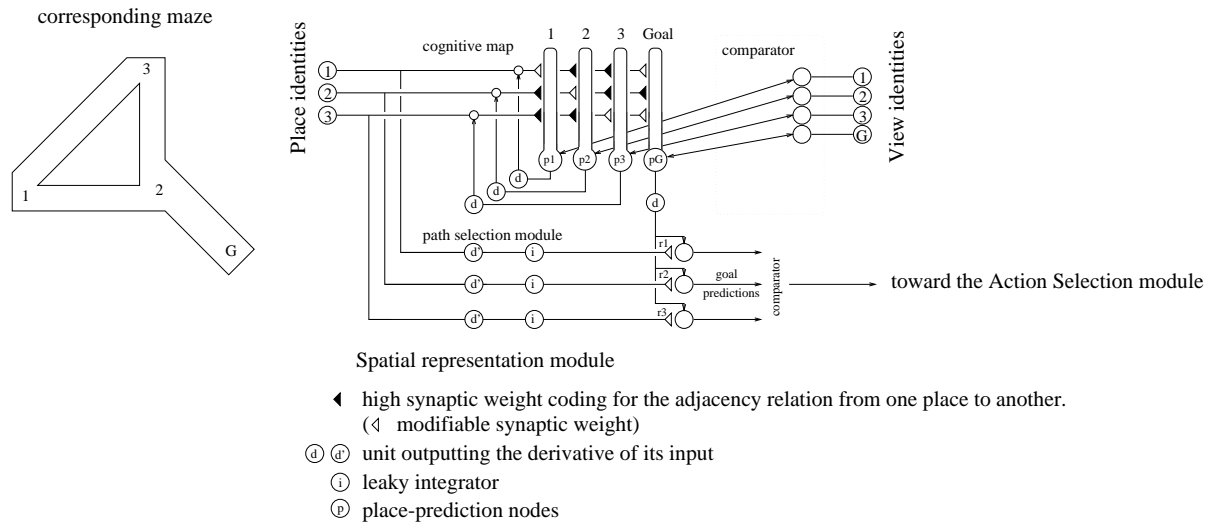


Figure 38: (Left) The maze in which the animat was tested consists of four distinct places. The animat is able to recognize each place and all of its neighboring places. (Right) The neural network implementing the spatial representation module described in Figure 37. The synaptic weights (filled and unfilled triangles) in the cognitive map are modified so that the recognition of the current place triggers high activity in the place-prediction nodes that correspond to the neighboring places. The recurrent connections enable further predictions that can then be compared by the path selection module. See text for details. After (Schmajuk and Thieme, 1992)

mismatch between pG being activated and the absence of the corresponding view of place G). Thus, “the network operates as a heteroassociative non-recurrent network for slow-changing signals and as a recurrent network for fast-changing signals.”

Before initiating a movement the animat “peeks” into each adjacent place, in sequence, generating a fast-changing signal from the current place to each adjacent place. This signal is sent simultaneously into the “cognitive map” and into the “path selection module”. Within the cognitive map module, it spreads through the network and the recurrent connections, until it eventually activates the node predicting the goal. Within the path selection module, it is transformed into a pulse by the d' node corresponding to the place peeked into. This pulse initiates a trace in the corresponding node i , a leaky integrator. The connection weights r_j are modifiable and store the signals coming from the goal-prediction node. In practice, a connection weight r_j increases whenever there is simultaneous activity on the corresponding i node and on the d node coming from the pG node. Thus, the connection weights r correspond to the topological distance between the adjacent places and the goal.

The Action Selection module compares these weights. The shortest path between the current place and the goal (in the number of intermediate places) starts at the adjacent place which generated the strongest signal at goal-prediction node pG. In other words, in order to reach the goal by the topologically shortest route, the animat should move into the place corresponding to the strongest connection weight r . If the prediction signals are too weak, which means either that the goal is far away, or that the spatial representation is incomplete, the Action Selection module generates a random movement.

Schmajuk and Thieme successfully simulated performance of Tolman and Honzik’s experiment (Figure 32) where the rat took the shortest detours (in terms of the number of places to traverse

in the model) after blockage of the normal path. It thus demonstrated their model's capacity of generating detours.

Schmajuk et al. [128] extend this model with a "route module" where each place is associated with a specific movement toward the goal. They simulated the interactions between "route-following" behavior (stereotyped sequence of place-response movements) and "path planning" behavior (as in the model here), and the effects of "lesioning" either the "route" module or the "cognitive map" module.²³

This model raises an interesting question: are there really two types of signals embedded within a single biological network, one for "building a map" and the other for "planning paths"? Schmajuk and Thieme do not explicitly refer to the hippocampus in their paper but since the hippocampus is strongly implicated in spatial representations, it would be interesting to compare this model with actual hippocampal functioning. In particular, this model predicts the existence of cells selectively active for places where the rat "intends to go" (similar to the d' nodes) and also weak activities in cells representing "adjacent" places.

3.3.3.5 Schölkopf and Mallot (1994)

Schölkopf and Mallot [129] extend the previous model of Schmajuk and Thieme by adding information about the movements required between respective places. In other words, while the output in the model of Schmajuk and Thieme consists of a place identity (movement was implicit), it consists here of the movement command that will lead the animat to the desired place.

The model was tested in a maze environment, shown in Figure 39, where places are predefined and numbered (as in the previous model) and views are also assumed to be all distinguishable and numbered. Contrary to the previous model, a place can have many corresponding views, depending on the direction from which the place is reached. For instance, from place p_4 , the animat will see view v_9 if it came from place p_2 , view v_8 if it came from place p_5 , and view v_7 if it came from place p_6 . Each view is uniquely determined by a corridor and an orientation. For instance, view v_9 corresponds to corridor p_2 - p_4 (oriented from p_2 to p_4). Thus, the animat of Schölkopf and Mallot navigates between views, instead of between places.

The animat makes one of three possible movements at each timestep, that is, turn left (by 60 degrees) and continue down the corridor, turn right (by 60 degrees) and continue down the corridor, or turn around (by 180 degrees) and continue down the corridor.

As we have seen in other models, during exploration, the animat builds a topological representation of the maze environment, a graph where nodes are the views and edges represent adjacency property between views. This follows from the association of the perception of one view and the subsequent locomotion (turning at a junction and moving down a corridor), and then the perception of the next view. For instance, v_9 is adjacent to v_2 because the animat can move from the former to the latter by a left-rotation at p_4 followed by a translation along the p_4 - p_5 corridor. This "view graph", as Schölkopf and Mallot call this topological graph, includes a metric component, namely the angle of rotation required to move from one view to an adjacent one. Thus, the spatial representation here incorporates the topological and the (local) directional structure of the environment, which is a major improvement over the previous model.

²³Note that "route" is taken here in its strict sense, as defined in Section 2.5, that is, an inflexible sequence of places linked by specific movements, and not as "route representation", that is, a flexible topological representation.




Figure 39: The maze used to simulate the navigation model of Schölkopf and Mallot. Places are defined as junction of corridors but are not used by the animat. The animat uses views, which are defined as the junction of two walls, as seen when approaching a place through the opposite corridor (e.g. v_9 is the view available when moving from place p_2 to place p_4). After (Schölkopf and Mallot, 1994)

Figure 40 shows the architecture of the artificial neural network that implements the coding of the topological spatial layout.

First, a unique pattern of activity (randomly chosen and normalized) across the input layer is associated with each possible view in the maze. Each pattern is meant to correspond to a unique visual stimulus. The input nodes are fully connected to the nodes in the “map” layer through initially random weights. There is a competitive mechanism among the map layer nodes (a variant of Kohonen’s self-organizing network [55]) and an adaptive plasticity of the synaptic weights between the input layer and the map layer, so that each map layer node ultimately becomes a view node.²⁴ Subsequently, a given pattern of input activity will drive a specific node. This node represents the corresponding view.

Second, a movement layer encodes the action of the animat: there are three nodes, representing respectively “turn left”, “turn right”, and “turn around” (or “move backward”). These nodes project onto the connections between the view nodes, so that the propagation of activity from one view node to another is gated by the activity of the movement nodes. During exploration, the synaptic weights of the local connections between view nodes are adapted through a modified Hebbian rule. They are enhanced whenever two connected view nodes are activated at consecutive timesteps and the propagation of activity through the connection is gated by the activity of the appropriate movement node.

Once exploration is completed, there are two complementary ways that a given view node can be activated: while at the end of a corridor, the visual stimulus will selectively excite the input nodes, which in turn will drive the corresponding view node; or, during locomotion (or planning), the combination of the activities in a view node and in a movement node will lead to the activation of the appropriate adjacent view node. For instance, view node v_6 can be activated either by the pattern of activity in the input layer corresponding to what the animat sees when at place p_4 and coming from place p_5 , or by the node representing view v_8 and the node representing the left-turn.

Schölkopf and Mallot assume an external “neural system which makes use of the map contained in the network.” This system corresponds to a path planning module. Schölkopf and Mallot functionally

²⁴These view nodes are somewhat similar to the view-cells of Zipser’s model [169], in that their activity depends upon the position and orientation of the animat.

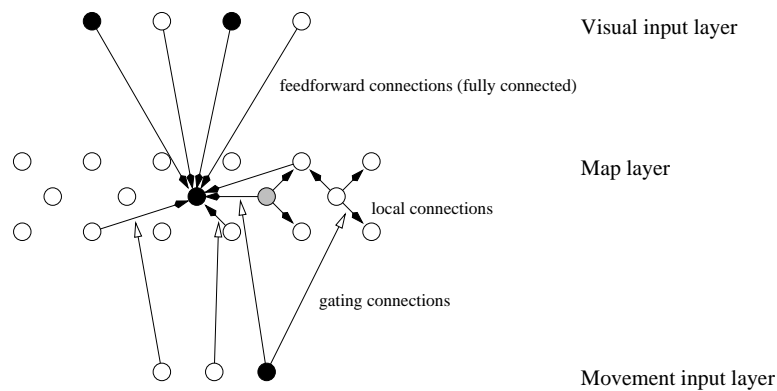


Figure 40: Architecture of Schölkopf and Mallot’s model. The current pattern of activity in the input layer (filled circles are active nodes and unfilled circles are inactive nodes) drives one of the view nodes (filled in black), which is the winner in the competition among all the nodes in the map layer. A second input comes through a local connection from the previously active view node (filled in gray), modulated by movement information (movement node filled in black). For instance, this state could represent the transition from view v_8 to view v_6 with a left-turn. After (Schölkopf and Mallot, 1994)

simulate such a system by using a path planning algorithm that resembles the classical graph search algorithm. First (i) the input pattern corresponding to the start view is activated, thus activating the starting view node; (ii) different possible movements (there is a finite discrete set) are then tried by activating the corresponding movement nodes, letting the activation spread for one timestep to see where the movement leads to; and (iii) the search then restarts from the new place until the goal is reached.

The simulation takes into account how a real robot would move between places, that is, along corridors. However, this model has been tested only in a well-structured maze. It is unclear how it would function in open-fields. In particular, since unique views may not be available in open-fields, and since each view corresponds in fact in this model to a corridor, that is, to “oriented places”, view nodes should be replaced by combinations of place cells and head-direction cells.

An interesting feature of this model is how movement information is processed by the network. The idea of associating places and movements is usually implemented in an abstract way, that is, by building transition matrices. The introduction of the gating connectivity is innovative. In the transition matrix of Figure 33, places and movements play equivalent roles. In other words, “coming from place A and turning left leads to place B” is the same as “turning left and coming from place A leads to place B.” In contrast, the principal part of the network here is the view nodes. The movement information only plays a modulatory role. An interesting question posed by this model is whether there is any physiological evidence for such connections from cells selective for the animal’s movements to place cells and head-direction cells.

3.3.3.6 Penna and Wu (1993)

Although inspired by physiological data concerning the hippocampal formation, Penna and Wu [106] consider navigation from a computational point of view. Their navigation system is based on an artificial neural network in which all the connections are built according to mathematical rules derived from geometrical considerations, instead of adapting the weights through biologically plausible learning mechanisms.

Following the work by Kuipers and Levitt [59], Penna and Wu partition the environment into small areas called “observation regions” (OR) (Figure 41). An OR is the area enclosed within a set of sectors, defined by “landmark pair boundaries” (LPB). A LPB is a line that passes through two landmarks in the environment, and is arbitrarily oriented from one of the landmarks to the other. For instance, LPB “L2-L3” in Figure 41 is oriented from L2 to L3, as shown by the arrows on the line. Consequently, a LPB divides space into a “left” half and a “right” half. A location is said to lie at the right of a LPB — joining landmark L1 and landmark L2 — when a viewer moving from L1 toward L2 sees it on the right. Thus, ORs are uniquely characterized by “observation vectors” which are lists of the positions of the points with respect to all the LPBs.

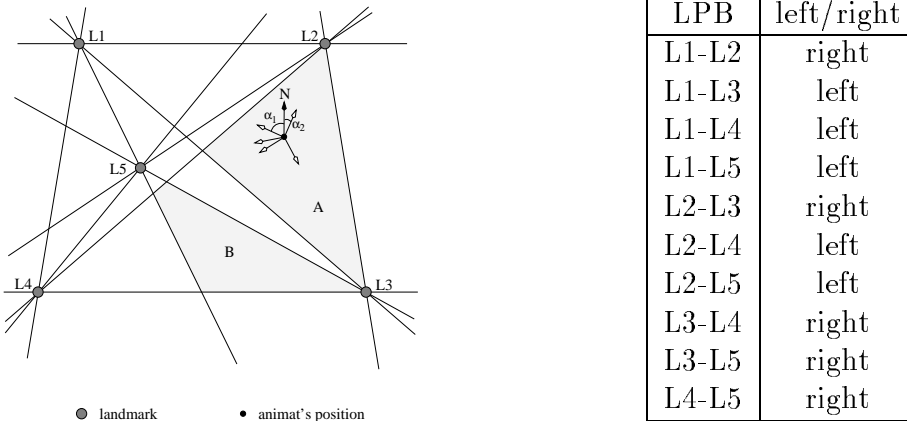


Figure 41: An environment in the model of Penna and Wu. The set of landmarks L1 to L5 are connected comprehensively by lines called Landmark-Pair Boundaries (LPB). (Left) Partitioning an environment into observation regions (OR). Each OR is enclosed within segments of LPBs. There is an absolute reference direction (North). The system registers the orientations of all the landmarks with respect to this reference. (Right) The observation vector for the region the animat is currently situated in (A). The animat is currently to the right of $\overrightarrow{L_1L_2}$ ($\alpha_1 > \alpha_2$), and to the left of $\overrightarrow{L_1L_3}$, etc.

The navigation system receives as inputs the allocentric bearings α_i of the landmarks (angle relative to some fixed and known “North”), so that each observation vector, and thus each OR, can be uniquely described by the relative bearings of each pair of landmarks. The rule for this is:

$$\begin{aligned} \alpha_i - \alpha_j \in [0, \pi] &\rightarrow \text{the current location is right (+1) of the oriented LPB defined by } L_i \text{ and } L_j. \\ \alpha_i - \alpha_j \in [-\pi, 0] &\rightarrow \text{the current location is left (-1) of the oriented LPB defined by } L_i \text{ and } L_j. \end{aligned}$$

The neural network proposed by Penna and Wu is a three-layer feedforward network. This network computes in which OR the animat currently is, from the orientations of the landmarks with respect to North.

The input layer consists of N nodes, where N is the number of landmarks in the environment, and the activity of each node is proportional to the bearing of the corresponding landmark (α_i). The intermediate layer computes the position of the animat with respect to each LPB (the activity of a node is for instance -1 for “left” and +1 for “right”). There are thus $N(N-1)/2$ nodes. The output layer represents the entire set of observation regions. Each node in the output layer computes the dot product between the vector corresponding to the pattern of activity in the intermediate layer and the observation vector corresponding to the OR the node represents. For instance, assume the animat is in region A, as shown in Figure 41. The pattern of activity in the intermediate layer can be represented by $(+1, +1, -1, -1, +1, -1, -1, +1, -1, +1)$ because the animat is to the right of $\overrightarrow{L_1L_2}$

($\alpha_1 > \alpha_2$), and to the left of $\overrightarrow{L_1 L_3}$, and so on. The node representing region A will have maximal activity (10), because its observation vector exactly matches the vector coded by the intermediate layer. In contrast, the node representing region B will have an activity of 8 because its observation vector is (+1,+1,-1,-1,+1,-1,-1,+1,-1,+1), which differs from the observation vector of region A by its second (L1-L3) and 9th (L3-L5) component. Consequently, the activities of the other nodes carry the “topological distance” information: the difference with the maximal activity is proportional to the number of LPBs to cross to move from one region to the current region (2 to move from region B to region A).

The number of output nodes depends not only on N, but also on the spatial layout of the N landmarks, which is not given a priori. Therefore, Penna and Wu propose an algorithm by which the navigation system dynamically builds the output layer of the neural network through exploration, adding a new node and computing the appropriate weights, whenever it crosses a new LPB.

Penna and Wu also propose alternative coding schemes for the input layer. Instead of partitioning the environment into regions defined by LPBs, they partition it into sectors with respect to compass directions and around each landmark. Figure 42 shows an example with four sectors. Each OR can then be uniquely characterized by either of two types of observation vectors. The first type of observation vectors consists of the identities of the sectors that each landmark occupies (e.g. OR 14 is defined by all the locations from where L_1 is seen in sector 1, L_2 in sector 1, L_3 in sector 2, and so on). The other type of observation vectors consists of the number of landmarks seen in each sector (e.g. OR 14 is defined by 2 landmarks in sector 1, 2 in sector 2, 1 in sector 3 and none in sector 4). The second type of coding is attractive because the landmarks don't need to be distinguishable. This is one of the rare models that do not need landmark recognition for efficient navigation (the only other example reviewed in this paper is the snapshot model by Cartwright and Collett, described in Section 3.1.2.1).

A neural network is built, similarly to the previous one, with as many input nodes as there are landmarks (first case) or sectors (second case) and as many output nodes as there are ORs. The input nodes encode the corresponding observation vectors, the number of the sector for each landmark in the first case and the number of landmarks in each sector in the second case.

The navigation system represents the topological layout of the regions through the activities of the nodes corresponding to observation regions. Planning a path requires selection of the sequence of observation regions to traverse and thus, the sequence of LPBs to cross. A guidance navigation strategy can then be used to cross each LPB. For instance, to cross $\overrightarrow{L_1 L_2}$ from left to right, the animat has to move so that α_1 becomes superior to α_2 .

For planning, Penna and Wu present a recursive graph search algorithm. To go from the start point to the goal, an external mechanism “simulates” the next movement and chooses the best. In practice, the search examines all the regions neighboring the origin (which can be determined by the nodes firing at rates just below the maximum), simulates being in each one and then scans the activity in the node representing the goal region. The animat then moves to the neighboring region that led to the maximum activity at the goal (in principle, this process resembles the spreading of activation seen in the model of Schmajuk and Thieme [127] since the choice is made upon the comparison of the activity at the goal node that would be evoked by the several different possible moves).

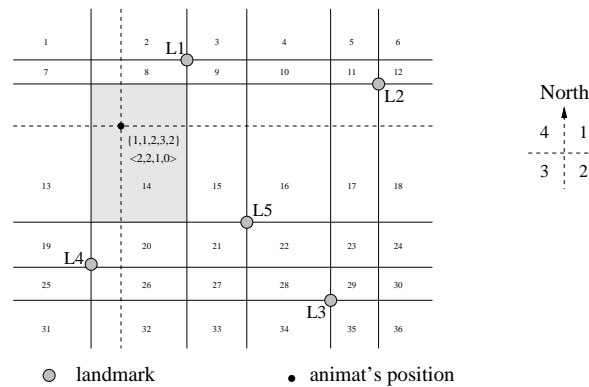


Figure 42: Alternative definitions of observation regions (OR). Penna and Wu here partition the environment at any given observation point along directions relative to a reference North. The labeling of the 4 quadrants is shown to the right. LPBs are replaced by lines generated from each landmark and in each of the specified directions. These lines enclose observation regions (OR). The animat divides the environment around it into the same sectors. Each OR is then uniquely defined by the sector each (recognizable) landmark is in, or by the number of (indistinguishable) landmarks in each sector. For instance, OR n°14 (shaded area) is defined by the fact that landmark 1 is in sector 1, landmark 2 in sector 1, landmark 3 in sector 2, and so on (vector between braces). It can also be defined by the fact that there are 2 landmarks in sector 1, 2 in sector 2 and so on (vector between brackets).

The output nodes of the neural network proposed by Penna and Wu are active not only in the observation regions they represent (where the activity is maximal) but also in neighboring regions. Hippocampal place cells are active in their place fields and silent everywhere else, but cells recorded in the entorhinal cortex or in subiculum can have spatial selectivity with very large place fields [114, 133]. This model thus suggests new analysis for entorhinal and subicular cells' activities. The question is whether their distribution of activities can carry topological information, as in this model.

3.3.3.7 Bachelder and Waxman (1994)

Bachelder and Waxman [4, 3] propose a navigation system which they also implemented on a real mobile robot. The robot successfully wandered in an open-field environment while recognizing “places” on the basis of visual and movement information. Although the current robot merely follows a predetermined trajectory and doesn't actually “navigate”, the system could integrate a topological navigation strategy in a straightforward manner, as described below.

As in Gaussier and Zrehen's model (Section 3.2.2.3), the robot is equipped with an actual 3D object recognition system. The 3D object recognition is based on the Seibert-Waxman neural recognition system [130] that learns to identify objects as seen from several different viewpoints. During the first phase of learning, several objects are presented to the immobile animat and an unsupervised neural classifier (an ART network [22]) creates “categories” corresponding to particular views. The network also learns the transformation of the image from one viewpoint to another. Figure 43 shows how a graph can represent the information stored in this network.

In the second phase of learning, the animat explores an open-field environment where only visually-learned objects are visible. As the animat moves, the object recognition system gives the landmark identities to a spatial representation module. In this module, places are defined by the landmark configuration, that is, the identities of the landmarks and their bearings with respect to an absolute directional reference. These allocentric bearings are computed from the egocentric bearings and the head-direction of the animat, indicated by its internal compass. In practice, only one object is seen




Figure 43: Aspect Graph: representations of the appearance of an object from different viewpoints. Such a graph can capture how the view changes, with the links between characteristic views representing canonical movements or transformations.

and recognized at a time. The robot rotates through 360° at the same place and looks at each object in order to build a “panoramic view” of the environment from its current location. A place is then associated with this view and the starting orientation.

Figure 44 shows the functional organization of the navigation system. There are two modules for navigation: one for place recognition and one for the topological map. As in the model proposed by Gaussier and Zrehen (Section 3.2.2.3), the place recognition module of the navigation system receives three types of information: (i) the head-direction of the animat with respect to an absolute reference direction (compass), (ii) the egocentric bearing of the landmark currently perceived, and (iii) the identity of the landmark currently perceived. As the visual scene is scanned to recognize each of the landmarks, the system stores the conjunction of these three types of information in a short-term memory (STM, bottom block in the place recognition module). The resulting activation pattern in the STM block represents the panoramic view that defines a place — which is categorized by an ART network.

The second module is essentially an associative network that learns what movement leads from one place to another: the transition of activity from one place node to another is mediated through the action of a movement representation node. Here, as in McNaughton’s formal model [69], it is the conjunction of a given place and a specific movement that is associated with the corresponding neighboring place. The STM block in this second module retains information concerning the identity of the current place as the animat moves to the next neighboring place. Thus, there is simultaneous activation of the conjunction of the current place node and the movement node, and the next place node. The association is learned through a Hebbian rule.

This navigation system could generate predictions of action consequences (a certain movement at a given place leads to the next place) and subsequently enable path planning. Indeed, it encodes the same information as the model of Schölkopf and Mallot [129], that is, a topological graph where nodes are places and edges are movements required to go from one place to the next. Consequently, planning for a topological navigation behavior could be done through a graph search or by propagating prediction, as Schmajuk and Thieme [127] have done with their model.

Bachelder and Waxman’s system has been tested in a real physical environment and succeeded, dealing with noise, difficult viewing conditions and self-motion uncertainties. However, it awaits

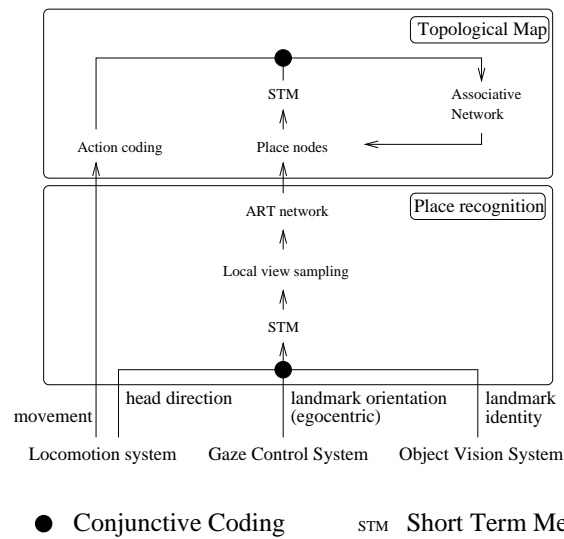


Figure 44: Functional organization of Bachelder and Waxman's animat. The lower-level modules provide the navigation system with four types of information: the identities of the landmarks visible from the current location (all around the animat in a 360° view), the relative orientations of these landmarks, the animat's head-direction, and the last movement made (egocentrically referred). The first three types of information are used by the place recognition module. The last type of information is integrated with the place representation into a topological map that represents how movements lead from one place to another. After (Bachelder and Waxman, 1994)

further implementation of path planning abilities.

3.4 Metric navigation

3.4.1 Behavioral experiments

3.4.1.1 The capability to make detours requires distance measurements

In order to make economical detours, to avoid unforeseen obstacles, and to take shortcuts, the animat must take into account metric information. The first question is how to represent distances, and in particular, distances of viewed objects from the animat. Zipser [168] and Collett et al. [33] utilized the size of the retinal image of objects. This of course requires object recognition and the measured quantity must be scaled by the actual dimensions of the object. Distances can also be measured on the basis of self-movements by using parallax or by counting the number of steps taken between places.

The following experiment by Carr and Watson dramatically demonstrates that, as a rat becomes more familiar with an environment, it relies less on its perceptual input and more on its internal representation. Carr and Watson (as reported by Gallistel [42, p.96]) trained rats in a maze shown schematically in Figure 45. At each trial, the rats were put at the beginning of an arm and required to make 2 turns and retrieve food at the end of the other arm. The length of the long arms could be modified. Surprisingly, the over-trained rats ignored the visual information indicating that the length of the arm had been shortened, and bumped into the “new” wall. In a second experiment, the length of the arm was longer, but the rats ignored the visual information again and entered the cul-de-sac, bumping into the wall at the end of it. However, the actual strategies used are not clear. The rat may have relied on its dead reckoning system and computed the distance traveled before turning from an internal metric representation of the maze. On the other hand, these errors could have been due to a motor response, that is, the animal learned a sequence of movements to reach the goal, without using any representation of the maze configuration (e.g. run 10 steps forward, turn right, then run 3 steps forward, turn left and finally run 10 steps forward).

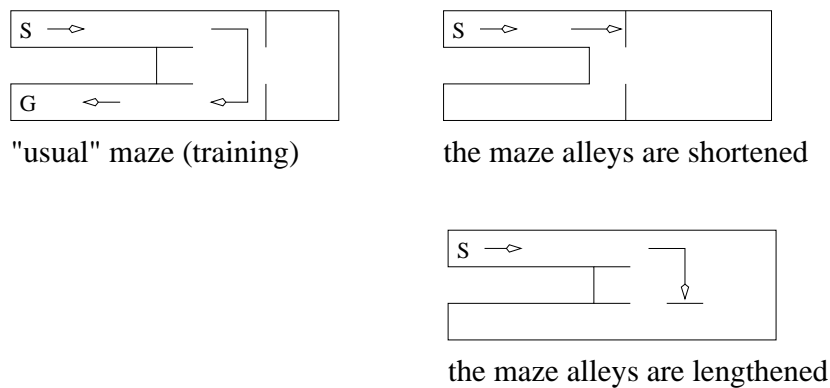


Figure 45: Does the rat learn the length of corridors in a maze? Schematic description of the experiment conducted by Carr and Watson. (Left) The rat was trained to run from a starting position (S) at one end of an alley to a goal position (G) at the end of an other alley, by making two turns. (Right) It was later tested on two other similar mazes where the alleys were either shortened (top) or extended (bottom). In the first case, the rat readily bumped into the new wall. In the second case, the rat turned before the end of the arm, and bumped into the end of the cul-de-sac. In other words, the animal seemed not to notice the change in distances. (Trajectory shown by arrows.)

This ambiguity was addressed in experiments performed by Poucet et al. [111] on cats required to solve a path selection problem. Again, the topological organization of the environment was not

relevant, but the lengths of the paths were.²⁵ As shown in Figure 46, an obstacle was placed so the angular deviation from the direct line to the goal could have one of several relations to path length. The obstacle was either a transparent or an opaque screen, so that the direct line to the goal was not always apparent. When the screen was opaque, cats first tested both paths, and subsequently chose the shortest one. However, when the transparent screen made the goal visible, cats always chose the path requiring the smallest initial turn, regardless of its length. This is another example of perceptual dominance, as described in Section 3.3.1. Thus, the cats chose a non-optimal path when the goal was visible, despite the fact that they were able to compare distance information and had already taken the paths' lengths into account for their choice when the goal was not visible. These experiments demonstrated the fact that cats are able to attend to metrical properties of the environment, independently of specific motor responses, but the choice of this strategy depends upon the available sensory information.




Figure 46: Overhead views of four experimental layouts used by Poucet et al. to test path selection by cats. An obstacle — a transparent or an opaque screen (solid line) — is put between the starting location and the goal location, so that the animal has to choose between two alternative paths (dashed lines with arrows). In experiment 1, the smallest deviation from a straight path led to the shortest trajectory. In experiment 2, both initial angles provide the same path length. In experiment 3, the turning angles are identical and the best choice is based upon distance comparison derived by comparing self-movement information from previous trials. In experiment 4, angular deviation and distance are contradictory. However, cats persist in taking the less divergent path when the goal is visible, that is, the animal's choice is based upon the angular deviation of the path from the direct line to the goal (arrow). After (Poucet et al., 1983)

Note that the choice of the straightest path whenever the goal was visible could be interpreted as a guidance strategy. In this case, without actually experiencing the whole path, the animal follows the smallest deviation from the visible goal location, which determines the choice of the path [111].

Path selection could be based on other non-metric information, such as time to reach the goal since the task consists of comparing different path lengths. However, Blancheteau and Le Lorec [13] reported that rats in their experiments selected paths more on the basis of distance than on the basis of time.

²⁵the same type of experiments were conducted, with similar results, on dogs [25].

The next section will deal with other experiments demonstrating that animals can make shortcuts across unfamiliar terrain to shorten path-lengths, an evidence for representation of metric information.

3.4.1.2 Shortcuts

Chapuis and Varlet [25] tested dogs on their navigational abilities in an outdoor environment. Dogs were shown two feeding sites, A and B, by being led along the paths DA and DB (Figure 47). They were subsequently released at point D to search for food. They ran down the path DA (the shortest one) and then took a shortcut to B, instead of coming back to D and following the known DB path. Two typical shortcuts were observed. The first one is a true directional one, leading directly to B — with minor errors in orientation (Fig. 47b). The second one is an “inside shortcut”, leading to a point on the DB path, between D and B (Fig. 47c). The “outside shortcut” — leading to a point on the line from D to B, but beyond B (Fig. 47c) — was rare. Moreover, dogs almost never followed the training path, that is, returned to D and then ran to B.




Figure 47: Experiments by Chapuis and Varlet with dogs. (a) Dogs were led along the path DA-AD-DB-BD and subsequently released at D. In order to reach food in both places A and B, they would tend to take the DAB path, thus short-cutting from A to B. (b) The first typical test path observed is a true directional shortcut. (c) The second one is an “inside shortcut”. (d) The third one is an “outside shortcut” and was rarely observed. After (Chapuis, 1988)

These results suggest that dogs are able to estimate the direction to a distant and unmarked point on the basis of a metric representation built from the integration of movement and visual information acquired during earlier but incomplete exploration of the environment. Deviations from true directional shortcuts may be due to a wrong estimation of the direction. However, since inside shortcuts were more often observed than outside shortcuts, these deviations could correspond to a “safety” strategy. Indeed, if the animal makes a bad estimation of the direction but aims at a point on the known trajectory (between D and B), it can correct itself as soon as it is in a known area. In contrast, if the animal tries an outside shortcut, it can miss the goal without being able to re-localize itself.

Menzel [75] tested metric spatial representations in chimpanzees. An animal was first carried about an enclosure and permitted to see the experimenter hiding some food in 18 different locations (Figure 48). The animal and the experimenter then left the enclosure and, later on, the animal was allowed to search for the food. It succeeded in retrieving most of the food but the order of retrieval was different from the order in which they were hidden. This is interesting because the animal took a novel path, including shortcuts, to go from one cache to another. It *seemed* to “optimize” the total path length. This suggests that the animal built a representation of the metric layout of the different pieces of food and computed a trajectory that was shorter than the one used to hide the

food. However, there seems to be no simple description yet of a mechanism controlling such a search behavior.




Figure 48: Overhead view of a typical trajectory (line with arrows) taken by a chimpanzee to retrieve pieces of food (circles) that were previously hidden in its presence (the number in the circles indicate the order in which pieces of food were hidden). The animal takes a novel trajectory but it also missed some rewards. After (Menzel, 1987)

3.4.2 Computational models

As stated in Section 2.6, metric navigation can be performed, from a computational point of view, by manipulating vectors that represent locations relative to some origin in a coordinate frame to be specified. The models described below use implicit coordinate frames that can be global and Euclidean [155], Euclidean but fragmented [166], or local and barycentric [112]. The output of the navigation system is a pair of coordinates. The existence of coordinate frames in the brain and the brain's ability to manipulate vectors are still controversial, but the following models suggest how metric spatial representations based on coordinate frames might be built from exploration.

3.4.2.1 Wan, Touretzky and Redish (1994)

Wan, Touretzky and Redish [155] propose a model which uses an underlying vector representation in a Euclidean reference frame. Their model was tested on different simulated behavioral tasks but also was partly implemented on a mobile robot [94]. The model operates an adaptive conjunction of multimodal information, coupling path integration, visual place recognition and updating of the head-direction representation.

A functional block diagram representing the architecture of the model is given in Figure 49. The hippocampal formation is presumed to occupy the “Place Code” module, a site of convergence for visual, directional, and position estimate information.

The visual input to the navigation system is the same as in most models and consists of the distances r_i and the relative bearings θ_i (in an egocentric reference frame) of all the landmarks i . The landmarks are assumed to be known a priori, and are visible from every location in the environment.

The head-direction representation is updated primarily by integrating vestibular input that codes the angular velocity of the head. Drifts resulting from accumulated integration errors can be corrected with visual information (see Section 3.2.1.2). Following McNaughton et al. [73], Wan et al. introduce a “local view” module in their model. This helps transform egocentrically sensed information to allocentric representations. Indeed, at any given place k and for any given landmark i , there is a constant linear relationship represented by Φ_{ki} between the head-direction Φ_h , an external reference

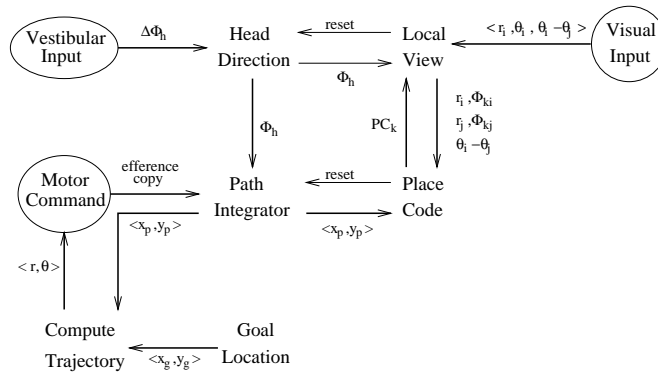


Figure 49: Functional block diagram of Wan et al.’s model. It shows how multimodal information might be combined and how different representations can update one another. Place is represented as the conjunction of multimodal inputs, coupling path integration, visual information, and head-direction. The path integrator outputs the position of the animat in Cartesian coordinates ($\langle x_p, y_p \rangle$) with respect to an a priori global reference frame. Visual information consists of the distances r_i and the relative bearings θ_i of the landmarks i . The relative bearings are also transformed into absolute bearings Φ_{ki} on the basis of information about the current place k . The head-direction (Φ_h) representation is updated by integrating the angular velocity $\Delta\Phi_h$, and is reset by place recognition when there is drift (PC_k). After (Touretzky et al., 1994)

direction and the relative bearings of the landmark i . This relationship is given by: $\Phi_{ki} = \Phi_h + \theta_i$. The “local view” module makes this linear association and stores the allocentric bearing of all the landmarks viewed at each place in its “Allocentric Bearing Memory” (Figure 50). Note that this system supposes that there are as many Allocentric Bearing Memories as there are combinations of place representations k and landmarks i .

When the animat is disoriented or when the drift in head-direction updating is too high, the head-direction representation can be recovered from visual input and place information by using the same linear relationship: $\Phi_h = \Phi_{ki} - \theta_i$.

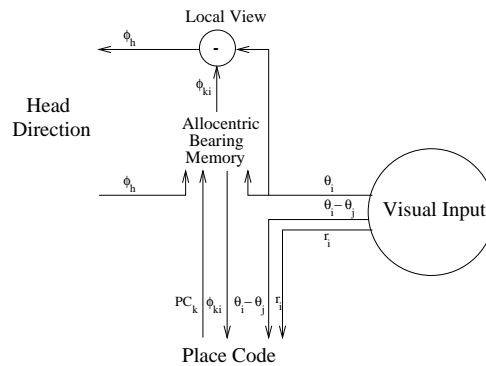


Figure 50: In Wan et al.’s model, the animat learns the relationship between egocentric and allocentric bearings of specific cues as viewed from a set of locations. Information about the current (estimated) head-direction (*input* Φ_h), the current (estimated) position (PC_k) and the egocentric bearings of landmarks (θ_i) converge into the Allocentric Bearing Memory which computes the head-direction independent (allocentric) bearings of these landmarks (Φ_{ki} for landmark i at place k). This module can also correct the head-direction representation by using the memorized allocentric bearings, the representation of the current place as well as visual information. Goal locations are defined in Cartesian coordinates (these coordinates can be provided by the path integrator). When the animat has localized itself and defined its goal, it can compute the direction and distance to the goal, thus generating the appropriate motor command. After (Wan et al., 1994a,b)

The path integrator estimates the Cartesian coordinates of the animat with respect to an arbitrary reference location, on the basis of movement information. This reference location is chosen to be the location of a selected object, a goal location (visible or not), or an entry point in the environment. Movement information consists of the efference copy of the motor commands as well as changes in the head-direction representation.

Wan et al. define a “place” as a region in the environment where all the different sensory measurements match a corresponding set of memorized values. There are 7 such types of quantities: the distances and the allocentric bearings of two selected landmarks, the (egocentric) bearing difference between two other selected landmarks, and the estimated (x,y) position relative to an identified reference point. The landmarks are randomly chosen when a place unit is recruited (see below) and are fixed thereafter. Each place unit “computes” its output activity as the product of the 7 Gaussian functions of these measurements. The mean of each Gaussian function is set to the value of the corresponding measurement at the place, while the standard deviation has fixed values.

One interesting property is that the conjunction of all these quantities is “opportunistic”, or “fuzzy”, in the sense that terms in the product are ignored whenever the corresponding measurement is unavailable, that is, when there is no visual input (in darkness), when there is no head-direction (disorientation) to determine allocentric bearings, or when there is no positional input (the animat arrives in an unfamiliar environment). However, a competition mechanism between place units ensures that, even if the overall activity is lowered due to absent information, one unit will win and be active, and this unit will stably represent the location it is tuned to. Figure 51 shows that the place field doesn’t keep quite the same form nor size when some terms drop from the input conjunction, but that the location of the field’s center remains quite consistent.

Each place unit corresponds to a specific (x,y) position with respect to a chosen reference point. This position is updated by the path integrator and subsequently learned by the place unit. The activity of a place unit (the product of the Gaussians) is maximal when the animat returns to the position learned by the unit and decreases as the animat moves further away. Thus, each place unit gives a position estimate of the animat. Similar to the interpolation between view cell representations in the model of Zipser (Section 3.2.2.1), the navigation system proposed by Wan et al. estimates the animat’s position by computing the weighted average of the positions given by each place unit. In order to deal with incomplete sensory information and ambiguously defined places, Wan et al. extend this population coding mechanism with a “dynamic relaxation process” [150]. First, there is a dynamic thresholding of the activity of place units. This mechanism limits the number of active units but also enables the activation of “enough” units when information is sparse and activities are low. Secondly, there is a recursive mechanism which inhibits some of the active units, which yields a clearer estimate of the animat’s position (each place unit’s corresponding position must be “sufficiently near” the global estimate obtained by the weighted average computation). This process is almost equivalent to a competition among units, such as those already used by Sharp [131] or Burgess et al. [19].

In the simulation, during the exploratory phase when the spatial representation is constructed, all environmental features can be detected. The low activities in the existing place units signify that none of them represents the current location (a fortiori when there is no place unit, at the start of the simulation). New units are then recruited whenever all the existing units have low activities. Each such new unit learns the values of the 7 parameters characterizing the current location [150].

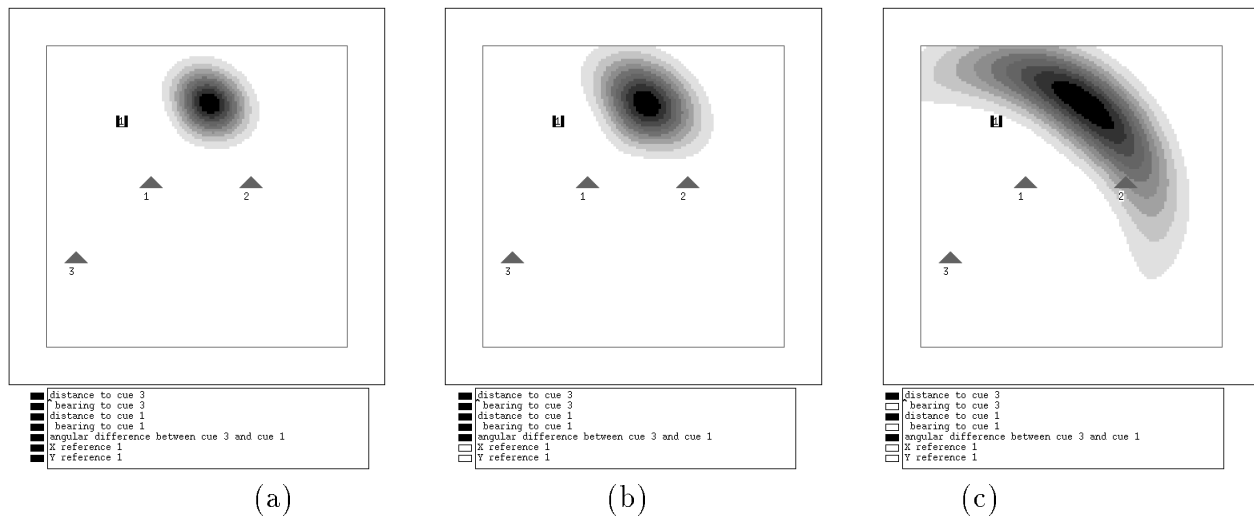


Figure 51: In Wan et al.’s model, the position information coded by the place unit activity smoothly degrades as the unit receives less sensory information. The figures show an overhead view of a square arena with three landmarks (triangles) and an arbitrarily chosen reference point (filled square). The same place unit’s activity map was simulated in different conditions, as described at the bottom of the figures (filled squares are available information). The firing rate (arbitrary units) is shown in gray scale, with black as the highest. (a) When there is complete information, (b) When the path integrator has lost its estimate of the current coordinates with respect to the reference point, (c) When there is no estimate of the current (x,y) position and there is no information from the head-direction module. This place unit was tuned to: the distances and bearings to landmarks 1 and 3, the angular difference between landmark 1 and landmark 3, and the (x,y) position with respect to reference point 1. The place field widens but the center of the field stays stable. Place fields are even more stable if the competition with other place units was taken into account, which is not the case here. (simulations by Trullier)

The goal location is defined by its (x,y) coordinates, relative to the specified reference point. The spatial representation of the environment by the place units is independent of the goal location, but both are in the Cartesian reference frame. Consequently, planning is not at the neural level but at the mathematical level of vector manipulations, as in Cartwright and Collett’s model [24]. This makes the neurobiological validation of the model more difficult.

This model accounts for several electrophysiological data not accounted for by previous models:

- the persistence of place cell activities in darkness [113]. This model assumes that in darkness, the place cell activities are driven solely by the path integrator (using vestibular information), although spatial resolution is diminished.
- the fact that hippocampal place cell activity is also modulated by head direction in rats performing in radial-arm mazes [71]. This model assumes that there is a change in the reference point for the path integrator between an inward and an outward movement of the animal.
- the updating and resetting of head-direction cells by visual information [143].

Redish and Touretzky [115, 116] assume that if some landmarks are indistinguishable, then place units are not tuned to the landmark identities but to their types (e.g. “cylinder”, or “corner of the cue-card”). The model then also accounts for some behavioral data in experiments with ambiguous goal locations (experiments by Cheng [28], by Collett et al. [33], and by Sharp et al. [134]). Instead of comparing the different possible choices — as Collett et al. had suggested it might occur in the gerbil

brain where each possibility was tallied — this navigation system lets the competitive mechanism and the relaxation process cause one solution to emerge.

All of the different modules of this model are described functionally. It would be interesting to elaborate a biologically plausible system that would implement the same computations. In particular, this model currently needs an external mechanism to “read” the estimated current position of the animat. Indeed, the relaxation process involves measuring how far from the estimated average position the position coded by each place unit is.

In Wan et al.’s model, movement and head-direction signals yield an (x,y) position that drives the place units. In other words, there is an intermediate representation of position. Knierim et al. [54] suggest instead that (self-)movement information directly modulates the interactions between head-direction representation and place representation.

3.4.2.2 Worden (1992)

An alternative approach is that the environment can be represented by adjacent sets of triangles, each defined by 3 landmarks. Worden [166] proposes a spatial representation based on “fragments”, or independently stored sets of such triangles.²⁶ Each fragment corresponds to the metric relationships between the three landmarks that define the triangle. Worden suggests that the non-geometric properties (e.g. color, surface texture) of these landmarks are also memorized — so that a fragment can be recognized along with the geometric relations between the landmarks. This enables recall of the triangle. Navigation is consequently based on combining, or “fitting” fragments to form a local geometric map (Figure 52).

There is a vector representation underlying the process. Indeed, forming the local map requires translating and rotating the appropriate fragments in an absolute reference frame in order to fit together the landmarks that are common to different triangles.




Figure 52: Navigating by fragment fitting. The animat, wishing to reach landmark L4 from landmark L2, sees landmarks L1 and L3. It retrieves fragments A and B and reconstructs a local map that relates L1-L2-L3 to L1-L3-L4. The position of L4 with respect to the current position ($\overrightarrow{L_2L_4}$) can be computed through vector manipulations $\overrightarrow{L_2L_1} + \overrightarrow{L_1L_4}$, or $\overrightarrow{L_2L_3} + \overrightarrow{L_3L_4}$ After (Worden, 1992)

According to Worden, the role of the hippocampus is to perform the geometric transformations required for translating, rotating and fitting fragments, and Worden attributes specific functions to each element of the hippocampal circuit.²⁷ Several distinct mechanisms are necessary: (i) encoding of vectors and directions, (ii) encoding of landmark identities (non-spatial attributes), (iii) associative

²⁶Thinus-Blanc [146] also suggests the use of such fragments.

²⁷O’Keefe and Nadel [101] earlier assumed an Euclidean representation in the hippocampus and O’Keefe [97] had proposed operations each element of the hippocampal circuit might perform.

(and parallel) retrieval of fragments, and (iv) parallel multiple matching. If the animat is not at a landmark location, it has to self-localize relative to the fragment containing its current location. Worden's model also uses vector representation to specify the animat's position with respect to this "central" fragment.

In contrast with other similar formal models, this model has been simulated. However, details are not given here because the actual implementation was formal and didn't use biologically plausible mechanisms such as neural networks.

3.4.2.3 Prescott (1994)

Prescott [112] proposes an idea similar to the latter — representing space by fragments — but implemented it with a neural network. Furthermore, the system doesn't need to match fragments because it first considers sets of four landmarks (instead of three), encoding the relative position of one landmark with respect to the fragment defined by the three others. Prescott advocates the notion of multiple partial representations of space and proposes that a global "map" would be reconstructed only when necessary. Furthermore, this doesn't necessarily require an independent neural module that would actually represent the map, since the relationships between fragments (in the sense defined in the previous model) can be implicitly encoded.

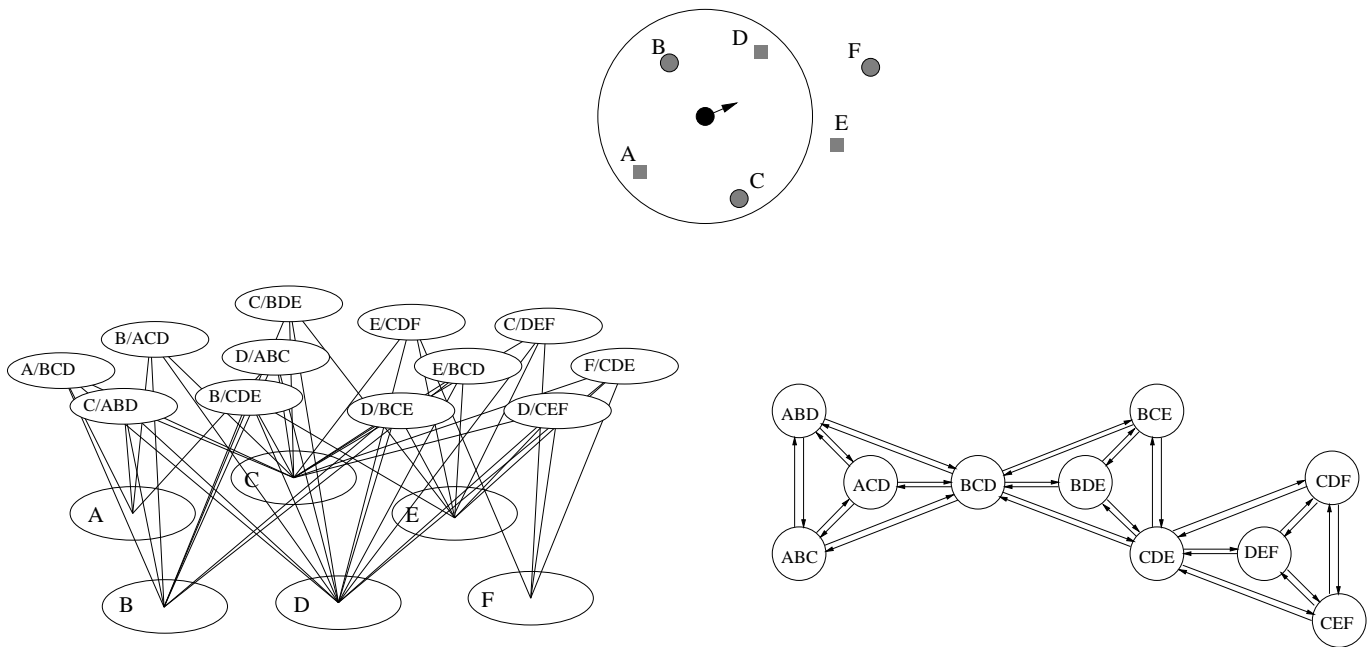


Figure 53: The internal spatial representation of Prescott's animat. (Top) An environment with 6 landmarks (grey objects labeled A to F). The animat is shown as a black circle with the arrow showing its heading direction. The outer circle indicates the limited field of view of the animat. (Bottom left) The animat constructs a relational network by learning the position of each given landmark with respect to a fragment (e.g. landmark A with respect to the triangle BCD). This network thus encodes the coordinate transformations between fragments. (Bottom right) The adjacency graph between fragments is implicitly encoded in this representation. After (Prescott, 1994)

Figure 53 summarizes the principles of the model. As the animat explores the environment, it learns the spatial relationships between the 4 most proximal landmarks in its field of view, that

is, the position of one landmark (e.g. A) with respect to the fragment defined by the three others (e.g. triangle BCD), in barycentric coordinates. The position of any landmark can be estimated with this spatial representation by successive transformations from the fragment currently seen to the relational node containing the landmark. For instance, the position of the currently invisible landmark E can be computed from its relation to fragment BCD, which is currently visible, and the position of the currently invisible landmark F can be further computed from its relation to fragment CDE.

Navigation consists in searching the (implicit) topological graph and planning a path from fragment to fragment. However, this navigation system exploits a metric navigation strategy, and not simply a topological navigation strategy, because the underlying vector representation enables the generation of new trajectories.

The main computational advantage of such a model is that the spatial representation is robust to noisy sensory inputs since the “coordinates” of the landmarks are not fixed but are dynamically estimated as the animat moves and sees different landmarks.

Prescott proposes his model from general biological considerations but he does not try to explicitly model actual neurobiological data, such as place cells. However, there are some formal theories that consider the hippocampus as a module representing “relations” — whether spatial or non-spatial — (e.g. [37]). Since Prescott assumes, as in the other models presented in this section, the existence of an underlying coordinate frame (relational coding through barycentric coordinates instead of vector coding in Cartesian coordinates), it would be interesting to test whether the hippocampus also can code such relations in an implicit underlying frame of reference.

4 Discussion

4.1 Summary

Table 2 is a summary of the main characteristics of the navigation systems described in the previous sections. It will enable comparisons and lead to proposals for future research approaches. The main characteristics can be grouped according to four main criteria:

- **The inputs:** What information is fed into the navigation system.

There are basically two types of inputs: types of information can be given as quantitative values such as distances or angles (**RV**) or can be expressed by the activation of specific “neurons” (**F**) that are selective for features such as ranges of distances. (Note that most models that use visual information assume a 360° view of the environment.) Thus, for models that build up a place representation, distances and egocentric bearings of landmarks are represented either way, often along with the landmarks’ identities. The landmarks’ identities can be directly given as inputs (**G**), or can be the output of a modeled vision system (**OV**). Visual information can also be expressed as specific patterns of a priori input cells (**a priori views**). For models that assume an a priori place representation, each place is represented by a feature-selective “neuron” (**F**) that is assumed to recognize it: the set of these “neurons” function as symbolic code.

There are also other types of information, such as a compass sense, movement information or ultrasonic range sensors (**US**).

- **The internal representation:** What the basic functional element of the navigation system is, what spatial features of the environment are learned and how the goal location is associated with the internal spatial representation.

Some models describe the information processing through explicit algorithms (**AC**), that is, mathematical relations between inputs and outputs. Some other models use neuron-like elements (**NL**), trying to model the neurocomputations with biologically-plausible elements. Furthermore, Burgess et al. [19] try to replicate the hippocampal architecture. Lastly, some other models only describe the navigation systems in general terms, with functional modules (**FM**) or symbolically (**symb**).

There are three popular learning mechanisms: (i) storing specific parameters (**MP**) such as distances, snapshots or barycentric coordinates; (ii) recruiting cells (**RC**) and storing the characteristics of the corresponding place; and (iii) updating the synaptic weights of the artificial neural networks, usually by competitive and Hebbian learning (**CH**) but also by more specific mechanisms (basic weight updating (**WU**), competitive learning (**CL**), ART networks (**ART**)).

The goal representation can be a stored set of parameters (**SP**). In most models, it is central to the spatial representation. A vector (**V**) usually is directly associated with each place, giving the direction to the goal, but it can also be completely implicit in the spatial representation (**imp**). The only model that builds up a goal-independent place representation and exhibits latent learning is the model by Burgess et al. [19], which recruits “goal cells” (**RC**). Models that fall into topological or metric navigation strategies usually build up a representation of spatial relationships from a given representation of individual places. Thus, a goal location is a posteriori easily defined as a given place. In metric representations of the environment, the

goal location can be defined by a set of coordinates (\mathbf{x}, \mathbf{y}) or by one specific landmark position (**LI**).

- **The output:** How the output to the motor system is expressed (**output**). If there is planning and if there are several solutions, how the “optimal” solution is selected (**planning**).

Most models indicate a direction of movement in the form of a vector — computed in different ways — (population coded vector (**PV**), direction to the goal (**V**)), in the form of a direction — given by a compass or defined as one of the n neighboring cells to go to in a grid-like environment (**D**), or in the form of a turning angle (**T**) prior to a forward movement. Models that have an a priori place representation define the animat’s action as reaching a given place (**PI**). Some models’ output is a command to a lower-level navigation strategy, typically a guidance instruction (**GG**).

As for planning, models use two possibilities: a classical graph search (**GS**) or a propagation of activation (**PA**).

- **The testbed:** The models have been tested either by computer simulations (**CS**) or on a mobile robot (**RI**). Their validity is restricted to the environment in which they were tested: a small open-field environment (**OF**) where all the landmarks are visible from every location, a large-scale environment (**LS**) where there are landmarks out of the range of perception — also either without obstacles or with walls (**indoor**), or a maze-like environment (**M**) where the configurations of places as well as possible movements are highly structured.

4.2 Principal suggested directions of future research

4.2.1 The inputs to the navigation system

The hippocampus is the site of convergence for multimodal inputs as evidenced by the anatomy [140] or by the influence of different modalities on the activity of hippocampal cells [163]. It is implicated not only in navigation but also in episodic memory or more generally the association of diverse types of information [119, 120]. For the problem of navigation, sensor fusion, that is, the process of taking into account different information provided by different sensors, has become a popular topic in robotics [65, 35]. This principle of multimodal convergence is re-emphasized in O’Keefe’s [96] definition: “a place cell is a cell which constructs the notion of a place in an environment by connecting together several *multisensory inputs* each of which can be perceived when the animal is in a particular part of an environment.” The sensory inputs to most models described in this paper are based on vision. A few also integrate self-movement information, a step in the direction of greater adaptiveness, such as the robots of Mataric [68], Kuipers and Byun [58], Bachelder and Waxman [4]. These latter still need to be developed to deal with problems like cue-conflict situations.

Wan et al. [155] took this one step further by suggesting how information provided by three different sensors (vision, head direction, and movement information) could be integrated and used to update different spatial representations. However, their model is at the functional level of description and their report [94] about a successful implementation of the model on a mobile robot does not give many details.

Stein and Meredith [138] provide evidence and propose a theory about the convergence of multimodal information in the superior colliculus, which is involved in spatial orienting (including ocular saccades) and attention of movement. In the same spirit, stress should be put in future research on

models	L identities	L distances	L ego. bearings	compass	a priori places	a priori views	mvt information	other sensors	architecture	learning	goal	output	planning	testbed	environment
guidance															
Cartwright and Collett (1983)		RV	RV	C					AC	MP	SP	PV		CS	OF
Benhamou et al. (1994)	G	RV	RV						NL+AC	WU	SP	D		CS	OF
Wilkie and Palfrey (1987)	G	RV							AC	MP	SP	D		CS	OF
place recognition-triggered response															
Zipser (1985-86)	G	RV	F						NL	RC	V	PV		CS	OF
Cartwright and Collett (1987)		RV	RV	C					AC	MP	V	V		CS	LS
Krakauer (1995)	G	RV	RV	C					NL	CH	V	V		CS	LS
Gaussier and Zrehen (1994)	OV		F	RV		F			NL	RC	V	V		CS	LS
Burgess et al. (1994)	G	F	F	F					NL	CH	RC	PV		CS	OF
Blum and Abbott (1994)					F				NL	CH	imp	PI		CS	OF
Brown and Sharp (1995)	G	F	F	F					NL	CH	imp	T		CS	OF
topological navigation															
Muller et al. (1991)					F				NL	CH	PI	PI	GS	CS	OF
Levenick (1991)					F		?		NL	CH	PI	PI	PA	CS	LS
Kuipers and Byun (1991)				RV				US	symp	MP	PI	GG	GS	CS	LS indoor
Mataric (1990)				RV		O		US	NL	CL	PI	D	PA	RI	LS indoor
Schmajuk and Thieme (1992)					F	F			NL	CH	PI	PI	PA	CS	M
Schölkopf and Mallot (1994)						F			NL	CH	PI	PI	GS	CS	M
Penna and Wu (1993)	G		F	C					NL	RC	PI	GG	GS	CS	OF
Bachelder and Waxman (1994)	OV		F	RV		F			NL	ART	not implemented			RI	OF
metric navigation															
Wan et al. (1994)	G	RV	RV	A			RV		FM	MP	(x,y)	V		CS	OF
Worden (1992)	G	RV	RV	RV					symp	MP	LI			CS	LS
Prescott (1994)	G	RV	RV	RV					NL	MP	LI	V		CS	LS

Table 2: Summary of the characteristics of the different navigation systems described in the paper. A: real value of the derivative; AC: analytical computations; ART: ART network; C: constant orientation; CH: competitive and Hebbian learning; CL: creating links; CS: computer simulation; D: direction of movement; ego.: egocentric; G: given; FM: functional modules; G: given; GG: command to the guidance strategy (edge to follow or LPB to cross); GS: AI graph search; imp: implicit; L: Landmark; LI: landmark identity; LS: large-scale environment; M: maze-like environment; MP: memorizing parameters; mvt: movement; NL: neuron-like elements; O: odometry; OF: small open-field environment; OV: output from a vision system; PA: propagation of activity; PI: place identity; PV: population coded vector; RI: robotic implementation; RV: real value; RC: recruiting cells; SP: set of parameters; symp: symbolic (not detailed); T: turning angle; US: ultrasonic range sensors; V: direction to the goal (vector); WU: weight updating; (x,y): coordinates;

what the mechanisms of sensor fusion might be in the hippocampus and further models should try to take this property more into account to provide greater robustness and flexibility.

4.2.2 The internal representation of space

Since animals can generate metric detours and shortcuts, this indicates that they use metric spatial relations in planning their paths. However, this doesn't necessarily mean that these animals have a global Euclidean representation of space, as assumed by formal cognitive map theories (for instance Gallistel [42]). Lieblisch and Arbib [63], for instance, insist on not ascribing Cartesian coordinates to the neural representation²⁸

Neural information processing can be functionally described with vector representations [42, chapter 14]. The activity of a neuron can be characterized by a set of numbers (e.g. firing frequency, phase of firing) and combinations across a population of neurons can yield other numbers. These numbers, ordered in a fixed way, give vectors that can be treated as points in a high-dimensional space — the “neural space”. Physical space and neural space could be related. The activity patterns of neurons might be “physical embodiment of the abstraction that we call vectors” to represent such entities as position of a point within a reference frame, or time-of-occurrence of an event. However, physiological evidence for how such vector representations might be manipulated (e.g. addition, comparisons) is sparse, although there are some computational suggestions (e.g. [151]). Thus, although the neuronal activity can be described with vectors that might be interpreted as representations of physical entities or states, it is not clear what significance those activity patterns have for neighboring areas of the brain.

Menzel also stresses the fact that Euclidean space is a concept which “humans invented to describe the conditions in which they live” [76]. The concept might help describe and explain how animals represent space and navigate but this doesn't imply that the brain implements the concept physiologically and anatomically.

Furthermore, as Wehner [159] demonstrates with a few examples from insect behavior, “our own formulations of the problems discussed [spatial orientation in insects] are not at all representative of how the animals tackle their tasks.” So what does neural space correspond to? We need to radically change the point of view and look for new or forgotten concepts. As Oatley [95] suggests when he describes how Polynesian navigators find their way without “paper maps”, the apparent complexity of the mental processes involved in animal spatial navigation may simply stem from the fact that the problem is tackled from a wrong angle.

4.2.3 Learned features of the internal spatial representation

The computational models can be grouped into systems that deal with small environments where every landmark is visible from everywhere (e.g. Zipser [169]; Burgess et al. [19]; Gaussier and Zrehen [44] and those that deal with large-scale environments (e.g. Mataric [68]; Schmajuk and Thieme [127]; Schölkopf and Mallot [129]). Models of the first group build place representations, whereas models of the second group often assume the existence of place representations and reconstruct the topological

²⁸They propose a formal theory (thus not described in detail here) where space is represented through a graph with nodes corresponding to places (situations) and edges to sensorimotor features that lead from one situation to the next. Such an architecture doesn't easily include metric properties [108].

relationships between these representations. There is thus a need for a model that can both build a place representation and determine the neighborhood relationships between these places.

Considering the kind of representation (place or topological) that the model builds leads to another possible categorization of the models, according to their learning algorithms. There are two types of learning: (i) recruiting place cells (e.g. Zipser [169]; Gaussier and Zrehen [44]; Mataric [68]; Wan et al. [155]) and (ii) modifying synaptic weights in an a priori given neural network (e.g. Burgess et al. [19]; Schmajuk and Thieme [127]; Schölkopf and Mallot [129]; Bacheider and Waxman [4]).

However, there is no evidence that a spatial representation is learned by the hippocampus in the first place. Indeed, Hill [52] reported that most of the recorded hippocampal cells showed specific spatial firing immediately after the animal was introduced in a new environment. This suggests that place fields do not develop during exploration but are directly mapped onto the environment. McNaughton et al. [70] suggest that the connectivity within the hippocampus is such that the place cell ensemble codes for a priori metric maps, where the metrics would correspond to self-movement information. In other words, in the absence of any external information, the proprioceptive information is sufficient to update the place coding by the hippocampal cells. The sensory information is then needed only to calibrate the internally-generated estimate of position, that is, learning would consist in mapping the internal representation to physical features in the environment. For instance, when the rat is introduced into a new environment, a new a priori map, that is, a new subpopulation of hippocampal cells would be selected; as the rat explores this environment, the distribution of activity in the place cell population would change in register with the animal's movements; each place cell would then be associated with the sensory inputs characterizing the corresponding place field; consequently, when the rat is reintroduced into this environment, the sensory features would trigger the appropriate place cell, and thus, the appropriate map.

Assume, then, that the rat is unable to learn the layouts of new environments. Every time the animal is introduced into a given environment, a new a priori map would be selected at random. Place cells would be active for each session in this environment, but the active hippocampal subpopulations would be different. The animal could never recognize the environment, yet it would "map" it each time. This hypothesis could thus be tested by showing that blocking learning in the hippocampus (for instance with NMDA antagonists) permits normal functioning of place cell firing, and by observing the characteristics of the corresponding place fields.

Even if the latter hypothesis turns out to be true, it is not yet clear what kind of synaptic connectivity is needed for such a priori maps. Indeed, one essential requirement is that when the animal returns to a place, the activity distribution within the hippocampal neuronal ensemble should also return to the corresponding place cells, on the basis of movement information only (Figure 54). No theory, to our knowledge, accounts for this. Moreover, the experimental evidence that two different environments can be coded by different yet overlapping populations of place cells with different neighboring properties of the place fields is problematic (cf. Section 3.3.2.1). Indeed, the activation of the current place cell and a signal corresponding to a specific movement are not sufficient to characterize which place cell should be activated next (Figure 55). What is needed is a means to disambiguate between similar activity profiles occurring in different environments.

Moreover, it is not yet clear the mechanisms by which movement information is used to update the place representation or to shift the activity from one place to another, although suggestions are made in some of the models described (Mataric [68]; Schölkopf and Mallot [129]; Bacheider and Waxman [4]; Wan et al. [155]). How can vestibular or kinesthetic information help predict

the future pattern of firing of place cells? There seem to be two main hypotheses (Figure 56). The updating can occur either in the hippocampus, or upstream of the hippocampus. In the first case, modulatory connections signaling movement would let propagate signals between direction-dependent neighboring place cells. This is what is proposed for instance by Schölkopf and Mallot [129]. In the second case, the highly processed information getting into the hippocampus would already be updated by movement information in the absence of extrinsic sensory input. The place cells would then be activated by the updated configuration of sensory input representations.

A similar process occurs in the head-direction representation (cf. Section 3.2.1.2). Models of how movement information could help update this representation have recently been published (e.g. Blair [11]; Zhang [167]). Their principle is similar to the idea put forth by Droulez and Berthoz [34] to suggest how saccadic eye movements can be executed toward memorized targets without the need for a coding of the target location in allocentric coordinates. Droulez and Berthoz proposed a “Dynamic Memory” that consists of a population of topographically connected neurons. The initial activity is generated by visual input. But in the absence of subsequent visual input, each neuron tends to activate the neurons that represent neighboring locations in the direction of the movement. Thus, the ensemble activity shifts in register with eye movements, so that the target location encoded by the activity profile in this neural ensemble can be appropriately updated.

One advantage of studying the head-direction representation is that it is one-dimensional (azimuth). Simulations are less time-consuming. These ideas should in principle apply to the place representation, which is two-dimensional. However, direct implementation of the extended algorithm seems not to be straightforward. Indeed, Zhang [167] presents preliminary results on place representations that are not as powerful as his results on the head-direction representation. One problem yet to be overcome stems from the fact that head-direction is limited to values between 0° and 360° , whereas there are no a priori limits to positions encoded in Cartesian coordinates. In other words, since neural ensembles are limited, the “edges” of the neural representation should correspond to physical edges in the environment.

4.2.4 Computing the motor output

The planning modules in computational models which contain them (e.g. Muller et al. [91]; Mataric [68]; Schmajuk and Thieme [127]; Schölkopf and Mallot [129]) are always more or less a graph search algorithm. For this kind of mechanism to apply, there is usually a need for an external mechanism detecting when the graph search has reached the goal location (decidability of the path planning problem) and then computing and comparing the different path lengths or some other criterion (completeness of the problem). Mataric [68] as well as Schmajuk and Thieme [127] implement the graph search by a propagation of activity. If a solution exists, the animats are able to select the best one but cannot tell the difference between the fact that the goal is too distant or the fact that the goal is unreachable from the current place, in the context of the current topological representation. Moreover, in Mataric’s model, the activity is triggered from the node representing the goal; this would not occur in the hippocampus since place cells are not selective for goals.

This leads to the same question as one posed by the model of Burgess et al. [19]: are there neurons in the animal’s brain (most promisingly downstream of hippocampus) that are selective for the identity of the goal currently aimed at, or neurons that are selective for a spatial component of this goal (e.g. direction or distance to it)?




Figure 54: Movement in the physical environment and movement in the spatial representation. Let us assume that there are a priori maps encoded by hippocampal place cells. In a given environment, a subpopulation of place cells (top, circles labeled from 1 to 7) correspond to a set of place fields (bottom, areas labeled 1 to 7, with dashed arrows indicating correspondence). As the animal moves within the environment, from place field 1 to place field 7 and back to place field 1, the activity in the place cells should shift in register (open arrows), on the basis of movement information only, from place cell 1 to place cell 7 and back to place cell 1.




Figure 55: Two given place cells in two different environments. Circles represent hippocampal neurons (top of both figures). If place cells A and B have overlapping place fields in one environment (left), they do not necessarily have overlapping place fields in a second environment (right). Thus, a signal corresponding to the same movement in the two environments (open arrow at the bottom of both figures) should lead to a different shift in the place cell activity (open arrows between the place cells). In one case, activity in cell A and the given movement lead to activity in cell B; in the other case, they lead to activity in cell C.

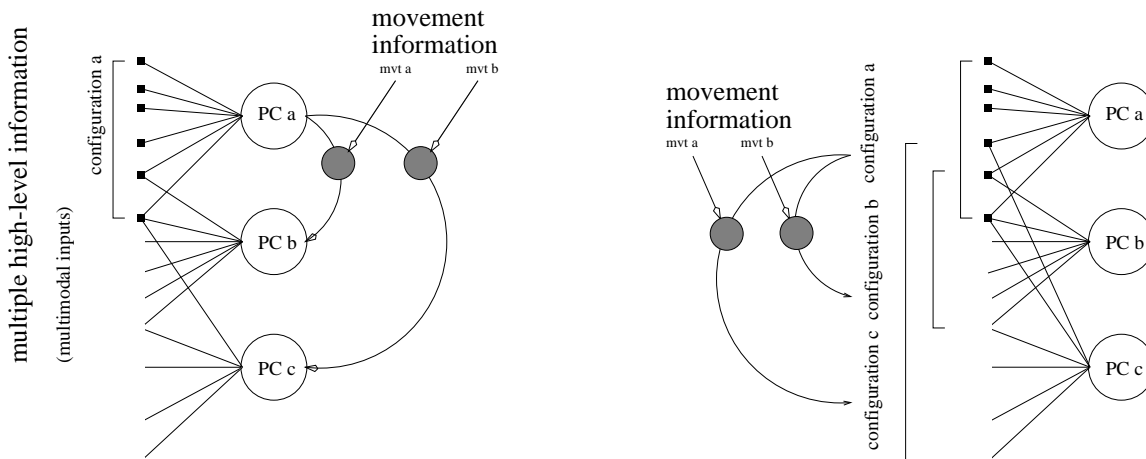


Figure 56: How does self-movement information contribute to change the place cell activity? There are two main hypotheses. (Left) Each place cell (PC) can be activated by a specific configuration of sensory inputs (filled squares are currently active units, open squares are inactive ones). In the absence of the extrinsic sensory input, however, the combination of position information and self-movement information (convergence on shaded circles) shifts the activity from one place cell to another. (Right) Self-movement information is used to update the configuration of extrinsic sensory inputs, which in turn activates the corresponding place cells.

The output of the navigation systems reviewed in this paper are essentially vectors (at least implicitly). This goes along with the assumption that animals perform vector manipulation [160]. However, as discussed above (Section 4.2.2), how vector representations are manipulated by the brain is unclear. Moving in a direction defined by a vector with respect to a reference frame requires comparing the current heading with respect to a reference direction and the desired heading. Most models rely on an internal compass. It would first be interesting to elaborate a system which could use environmental features to determine its heading. Then, a mechanism must be devised, by which the directional output of the navigation system can be transformed into an actual motor output command.

4.2.5 Multiple environments

Since different experimental settings and rooms are represented by distinct subpopulations of place cells (sometimes with overlap), as discussed in Section 3.3.2.1, it seems reasonable to postulate that animals have a particular spatial representation for each environment, instead of using one global cognitive map. However, there is to our knowledge no computational model to date that simultaneously represents multiple environments.

The fact that two-thirds of all hippocampal units isolated by Thompson and Best [147] were behaviorally silent in the three different environments that their rats were exposed to can be interpreted in two ways. On the one hand, these observations mean that there remained many cells that can be “recruited” to represent other environments. On the other hand, since as many as 30% of the isolated hippocampal units are involved in as few as 3 different environments, we could say that each unit participates in the representation of many different environments. Consequently, the simplest model for representing multiple environments, that is, using distinct populations of neurons for different environments, cannot hold. On the contrary, “sparse coding” [67] is what seems to occur, although the results of Thompson and Best [147] indicate that place cells in the hippocampus are not sufficiently sparse.

It is still unknown how a given environment is recognized and the corresponding representation is retrieved. In particular, if building a spatial representation of a given environment involves changing synaptic weights between hippocampal neurons, as it is widely assumed in the computational models, how do synaptic changes involved in learning in one environment affect representations of another environment? Sparse coding partially palliates the problem of interference between different representations [152, 153] but when the animal is introduced into a familiar environment, the correct subpopulation of place cells *along with* the appropriate set of modified connections should be “activated”. (This would require an external gating mechanism.)

In the model proposed by Muller et al. [92] (cf. Section 3.3.3.1), for instance, an external mechanism is required, either to selectively activate the appropriate connections, or to “identify inconsistent synaptic [weights]” by reconstructing a physical map (Figure 34). McNaughton et al. [70] advocate the existence of multiple maps for different environments and different scales (like in an atlas), but without specifying what kind of mechanism might select one map instead of another.

4.2.6 Conclusions

In summary, we have presented four different types of navigation strategies which enable different spatially-oriented behaviors. In particular, we have shown that “cognitive maps”, in the commonly-

used sense of the term, are not necessary for many kinds of navigation tasks. Many of the existing models that deal only with special types of environments could be made more robust by diversification or by adding on modules more suitable for other types of environments. Animats equipped with all four modules of navigation systems would be able to deal with a wide range of environments, with one of the modules exploited preferentially in each particular situation. The opportunistic recruitment of lower level processes in our hierarchy of navigation strategies could increase efficiency.

In our review of the existing biologically-based artificial navigation systems, we showed that many could reproduce some of the properties exhibited by animals and this was verified by successful implementations on mobile robots. There remain, however, several important milestones ahead of us. For example, none of the models could be shown to actually be capable of generating both *metric* detours and *metric* shortcuts. In order to build models that are able to do this, we will continue to look for inspiration to new discoveries about anatomical and physiological mechanisms employed by the brain to manipulate vector information in a specific coordinate frame. In this way we can continue to apply what can be learned about the actual circuit processing in the nervous system of navigating animals to fashion algorithms and even architectures for future robots.

References

- [1] K. P. Able. Mechanisms of orientation, navigation and homing. In S. A. Gauthreaux Jr., editor, *Animal Migration, Orientation and Navigation*, pages 283–373. Academic Press, 1980.
- [2] D. G. Amaral and M. P. Witter. The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, 31(3):571–591, 1989.
- [3] I. A. Bachelder and A. M. Waxman. Mobile robot visual mapping and localization: A view-based neurocomputational architecture that emulates hippocampal place learning. *Neural Networks*, 7(6/7), Special issue on Neurodynamics and Behaviour:1083–1099, 1994.
- [4] I. A. Bachelder and A. M. Waxman. A neural system for qualitative mapping and navigation in visual environments. In Gaussier and Nicoud [43], pages 266–277.
- [5] A. G. Barto, S. J. Bradtke, and S. P. Singh. Learning to act using real-time dynamic programming. *Artificial Intelligence, Special Volume: Computational Research on Interaction and Agency*, 72(1):81–138, 1995.
- [6] A. G. Barto and R. S. Sutton. Landmark learning: An illustration of associative search. *Biological Cybernetics*, 42:1–8, 1981.
- [7] S. Benhamou, P. Bovet, and B. Poucet. A place navigation algorithm based on elementary computing procedures and associative memories. In Cliff et al. [29], pages 206–213.
- [8] S. Benhamou, P. Bovet, and B. Poucet. A model for place navigation in mammals. *Journal of Theoretical Biology*, 173:163–178, 1995.
- [9] A. T. D. Bennett. *The ecology of spatial memory of European jays*. PhD thesis, University of Oxford, 1991. chapter 1.
- [10] R. Biegler and R. G. M. Morris. Landmark stability is a prerequisite for spatial but not discrimination learning. *Nature*, 361:631–633, 1993.
- [11] H. T. Blair. A thalamocortical circuit for computing directional heading in the rat. In Touretzky et al. [149], pages 152–158.
- [12] H. T. Blair and P. E. Sharp. Anticipatory head direction signals in anterior thalamus: Evidence for a thalamocortical circuit that integrates angular head motion to compute head direction. *Journal of Neuroscience*, 15(9):6260–6270, 1995.
- [13] M. Blancheteau and A. Le Lorec. Raccourci et détour chez le rat: Durée, vitesse et longueur des parcours. *L'Année Psychologique*, 72:7–16, 1972.
- [14] K. I. Blum and L. F. Abbott. A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, 8:85–93, 1996.
- [15] C. R. Breese, R. E. Hampson, and S. A. Deadwyler. Hippocampal place cells: Stereotypy and plasticity. *Journal of Neuroscience*, 9(4):1097–1111, 1989.

- [16] M. A. Brown and P. E. Sharp. Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and the Nucleus Accumbens. *Hippocampus*, 5:171–188, 1995.
- [17] T. H. Brown, E. W. Kairiss, and C. L. Keenan. Hebbian synapses: Biophysical mechanisms and algorithms. *Annual Review of Neuroscience*, 13:475–511, 1990.
- [18] N. Burgess, J. O’Keefe, and M. Recce. Using hippocampal ‘place cells’ for navigation, exploiting phase coding. In S. J. Hanson, C. L. Giles, and J. D. Cowan, editors, *Advances in Neural Information Processing Systems 5*. Morgan Kaufmann, 1992.
- [19] N. Burgess, M. Recce, and J. O’Keefe. A model of hippocampal function. *Neural Networks*, 7(6/7):1065–1081, 1994.
- [20] G. Buzsaki, L. S. Chen, and F. H. Gage. Spatial organization of physiological activity in the hippocampal region: Relevance to memory formation. *Progress in Brain Research*, 83:257–268, 1990.
- [21] R. Caminiti, P. B. Johnson, and A. Urbano. Making arm movements within different parts of space: Dynamic aspects in the primate motor cortex. *Journal of Neuroscience*, 10(7):2039–2058, 1990.
- [22] G. A. Carpenter and S. Grossberg, editors. *Pattern Recognition by Self-Organizing Neural Networks*. The MIT Press, 1991. (Part III: Adaptive Resonance Theory).
- [23] B. A. Cartwright and T. S. Collett. Landmark learning in bees. Experiments and models. *Journal of Comparative Physiology A*, 151:521–543, 1983.
- [24] B. A. Cartwright and T. S. Collett. Landmark maps for honeybees. *Biological Cybernetics*, 57:85–93, 1987.
- [25] N. Chapuis. *Les opérations structurantes dans la connaissance de l’espace chez les mammifères: Détour, raccourci et retour*. Thèse de Doctorat d’Etat-ès-Sciences, Université d’Aix-Marseille II, 1988.
- [26] L. L. Chen, L.-H. Lin, C. A. Barnes, and B. L. McNaughton. Head-direction cells in the rat posterior cortex II. Contributions of visual and ideothetic information to the directional firing. *Experimental Brain Research*, 101:24–34, 1994.
- [27] L. L. Chen, L.-H. Lin, E. J. Green, C. A. Barnes, and B. L. McNaughton. Head-direction cells in the rat posterior cortex I. Anatomical distribution and behavioral modulation. *Experimental Brain Research*, 101:8–23, 1994.
- [28] K. Cheng. A purely geometric model in the rat’s spatial representation. *Cognition*, 23:149–178, 1986.
- [29] D. Cliff, P. Husbands, J.-A. Meyer, and S. W. Wilson, editors. *From Animals to Animals 3: Proceedings of the Third International Conference on Adaptive Behavior*. The MIT Press/Bradford Book, 1994.

- [30] T. S. Collett. The use of visual landmarks by gerbils: Reaching a goal when landmarks are displaced. *Journal of Comparative Physiology A*, 160:109–113, 1987.
- [31] T. S. Collett. Landmark learning and guidance in insects. *Philosophical Transactions of the Royal Society of London B*, 337:295–303, 1992.
- [32] T. S. Collett and J. Baron. Biological compasses and the coordinate frame of landmark memories in honeybees. *Nature*, 368:137–140, 1994.
- [33] T. S. Collett, B. A. Cartwright, and B. A. Smith. Landmark learning and visuo-spatial memories in gerbils. *Journal of Comparative Physiology A*, 158:835–851, 1986.
- [34] J. Droulez and A. Berthoz. The concept of dynamic memory in sensorimotor control. In D. R. Humphrey and H. G. Freund, editors, *Motor Control: Concepts and issues*, pages 137–161. John Wiley & Sons, 1991.
- [35] H. F. Durrant-Whyte. *Integration Coordination and Control of Multi-Sensor Robot System*. Kluwer Academic Publishers, 1988.
- [36] F. C. Dyer. Bees acquire route-based memories but not cognitive maps in a familiar landscape. *Animal Behaviour*, 41:239–246, 1991.
- [37] H. B. Eichenbaum. Thinking about cell assemblies. *Science*, 261:993–994, 1993.
- [38] H. B. Eichenbaum, S. I. Wiener, M. L. Shapiro, and N. J. Cohen. The organization of spatial coding in the hippocampus: A study of neural ensemble activity. *Journal of Neuroscience*, 9:2764–2775, 1989.
- [39] A. S. Etienne, S. Joris-Lambert, C. Dahn-Hurni, and B. Reverdin. Optimizing visual landmarks: Two- and three-dimensional minimal landscapes. *Animal Behaviour*, 49:165–179, 1995.
- [40] L. Feng, J. Borenstein, and H. R. Everett. “Where am I?": Sensors and methods for autonomous mobile robot positioning. Technical report, University of Michigan, 1994.
- [41] L. N. Foner and P. Maes. Paying attention to what’s important: Using focus of attention to improve unsupervised learning. In Cliff et al. [29], pages 256–265.
- [42] C. R. Gallistel. *The Organization of Learning*. The MIT Press/Bradford Book, 1990.
- [43] P. Gaussier and J.-D. Nicoud, editors. *Proceedings of the PerAc Conference: from Perception to Action*. IEEE, 1994.
- [44] P. Gaussier and S. Zrehen. Complex neural architectures for emerging cognitive abilities in an autonomous system. In Gaussier and Nicoud [43], pages 278–289.
- [45] V. V. Gavrilov, S. I. Wiener, and A. Berthoz. The properties of hippocampal cells observed during passive displacements of rats on a mobile robot. *European Journal of Neuroscience*, 7:137, 1994.
- [46] V. V. Gavrilov, S. I. Wiener, and A. Berthoz. Enhanced hippocampal theta EEG during whole body rotations in awake restrained rats. *Neuroscience Letters*, 197:239–241, 1995.

- [47] V. V. Gavrilov, S. I. Wiener, and A. Berthoz. Whole body rotations enhance hippocampal theta rhythmic slow activity in awake rats passively transported on a mobile robot. *Annals of the New York Academy of Sciences*, 781:385–398, 1996.
- [48] A. P. Georgopoulos, A. B. Schwartz, and R. E. Kettner. Neuronal population coding of movement direction. *Science*, 233:1416–1419, 1986.
- [49] M. A. Goodale. Neural mechanisms of visual orientation in rodents: Targets versus Places. In A. Hein and M. Jeannerod, editors, *Spatially Oriented Behavior*, chapter 3. Springer Verlag, 1983.
- [50] J. L. Gould. The locale map of honey bees: Do insects have cognitive maps? *Science*, 232:861–863, 1986.
- [51] D. O. Hebb. *The Organization of Behavior*. Wiley, 1949.
- [52] A. J. Hill. First occurrence of hippocampal spatial firing in a new environment. *Experimental Neurology*, 62:282–297, 1978.
- [53] A. J. Hill and P. J. Best. Effects of deafness and blindness on the spatial correlates of hippocampal unit activity in the rat. *Experimental Neurology*, 74:204–217, 1981.
- [54] J. J. Knierim, H. S. Kudrimoti, and B. L. McNaughton. Place cells, head direction cells, and the learning of landmark stability. *Journal of Neuroscience*, 15(3):1648–1659, 1995.
- [55] T. Kohonen. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69, 1982.
- [56] D. C. Krakauer. Simple connectionist models of spatial memory in bees. *Journal of Theoretical Biology*, 172:149–160, 1995.
- [57] B. J. Kuipers. The “map in the head” metaphor. *Environment and Behavior*, 14(2):202–220, 1982.
- [58] B. J. Kuipers and Y. T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems*, 8:47–63, 1991.
- [59] B. J. Kuipers and T. S. Levitt. Navigation and mapping in large-scale space. *AI Magazine*, 9(2):25–43, 1988.
- [60] B. Leonard and B. L. McNaughton. Rat: Conceptual, behavioral, and neurophysiological perspectives. In R. P. Kesner and D. S. Olton, editors, *Neurobiology of comparative cognition*, chapter 13. Lawrence Erlbaum Associates, 1990.
- [61] J. R. Levenick. NAPS: A connectionist implementation of cognitive maps. *Connection Science*, 3(2):107–126, 1991.
- [62] T. S. Levitt and D. T. Lawton. Qualitative navigation for mobile robots. *Artificial Intelligence*, 44:305–360, 1990.
- [63] I. Lieblich and M. A. Arbib. Multiple representations of space underlying behavior. *The Behavioral and Brain Sciences*, 5:627–659, 1982.

- [64] Long-Ji Lin. *Reinforcement Learning for Robots Using Neural Networks*. PhD thesis, Carnegie Mellon University, 1993. Tech report CMU-CS-93-103.
- [65] R. C. Luo and M. G. Kay. Multisensor integration and fusion in intelligent systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(5):901–931, 1989.
- [66] P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson, editors. *From Animals to Animats 4: Proceedings of the Fourth International Conference on Adaptive Behavior*. The MIT Press/Bradford Book, 1996.
- [67] D. Marr. Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London B*, 262(841):23–81, 1971.
- [68] M. J. Mataric. A distributed model for mobile robot environment-learning and navigation. Technical Report TR 1228, MIT AI Lab, 1990.
- [69] B. L. McNaughton. Neuronal mechanisms for spatial computation and information storage. In Nadel et al. [93], chapter 9.
- [70] B. L. McNaughton, C. A. Barnes, J. L. Gerrard, K. Gothard, M. W. Jung, J. J. Knierim, H. Kudrimoti, Y. Quin, W. E. Skaggs, M. Suster, and K. L. Weaver. Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *Journal of Experimental Biology*, 199:173–, 1996.
- [71] B. L. McNaughton, C. A. Barnes, and J. O’Keefe. The contribution of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 52:41–49, 1983.
- [72] B. L. McNaughton, C. A. Barnes, G. Rao, J. Baldwin, and M. Rasmussen. Long-term enhancement of hippocampal synaptic transmission and the acquisition of spatial information. *Journal of Neuroscience*, 6(2):563–571, February 1986.
- [73] B. L. McNaughton, L. L. Chen, and E. J. Markus. ‘Dead reckoning’, landmark learning, and the sense of direction: A neurophysiological and computational hypothesis. *Journal of Cognitive Neuroscience*, 3(2):190–202, 1991.
- [74] B. L. McNaughton, B. Leonard, and L. Chen. Cortical-hippocampal interactions and cognitive mapping: A hypothesis based on reintegration of the parietal and inferotemporal pathways for visual processing. *Psychobiology*, 17(3):236–246, 1989.
- [75] E. W. Menzel. Chimpanzee spatial memory organization. *Science*, 182:943–945, 1973.
- [76] E. W. Menzel. Behavior as a locationist views it. In P. Ellen and C. Thinus-Blanc, editors, *Cognitive Processes and Spatial Orientation in Animal and Man, vol. 1: Experimental Animal Psychology and Ethology*, pages 55–72. Martinus Nijhoff, 1987.
- [77] J.-A. Meyer and A. Guillot. Simulation of adaptive behavior in animats: Review and prospect. In Meyer and Wilson [80].
- [78] J.-A. Meyer and A. Guillot. From SAB90 to SAB94: Four years of animat research. In Cliff et al. [29].

- [79] J.-A. Meyer, H. L. Roitblat, and S. W. Wilson, editors. *From Animals to Animats 2: Proceedings of the Second International Conference on Adaptive Behavior*. The MIT Press/Bradford Book, 1993.
- [80] J.-A. Meyer and S. W. Wilson, editors. *Proceedings of the First International Conference on Adaptive Behavior: from Animals to Animats*. The MIT Press/Bradford Book, 1991.
- [81] R. Miller. *Cortico-Hippocampal interplay and the representation of contexts in the brain*. Springer-Verlag, 1991.
- [82] M. Mishkin, L. G. Ungerleider, and K. A. Macko. Object vision and spatial vision: Two cortical pathways. *Trends in Neuroscience*, 6(10):414–417, 1983.
- [83] H. Mittelstaedt and M.-L. Mittelstaedt. Homing by path integration. In F. Papi and H. G. Wallraff, editors, *Avian Navigation*, pages 290–297. Springer Verlag, 1982.
- [84] M.-L. Mittelstaedt and S. Glasauer. Idiothetic navigation in gerbils and humans. *Zoologische Jahrbucher. Physiologie*, 95:427–435, 1991.
- [85] S. J. Y. Mizumori and J. D. Williams. Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *Journal of Neuroscience*, 13(9):4015–4028, 1993.
- [86] R. G. M. Morris. Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12:239–260, 1981.
- [87] R. G. M. Morris, P. Garrud, J. N. P. Rawlins, and J. O’Keefe. Place navigation impaired in rats with hippocampal lesions. *Nature*, 297:681–683, 1982.
- [88] R. U. Muller, E. M. Bostock, J. S. Taube, and J. L. Kubie. On the directional firing properties of hippocampal place cells. *Journal of Neuroscience*, 14:7235–7251, 1994.
- [89] R. U. Muller and J. L. Kubie. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7:1951–1968, 1987.
- [90] R. U. Muller and J. L. Kubie. The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, 9(12):4101–4110, 1989.
- [91] R. U. Muller, J. L. Kubie, and R. Saypoff. The hippocampus as a cognitive graph (abridged version). *Hippocampus*, 1(3):243–246, July 1991.
- [92] R. U. Muller, M. Stead, and J. Pach. The hippocampus as a cognitive graph. *Journal of General Physiology*, 107:663–694, 1996.
- [93] L. Nadel, L. A. Cooper, P. Culicover, and R. M. Harnish, editors. *Neural Connections and Mental Computations*. The MIT Press/Bradford Book, 1989.
- [94] I. Nourbakhsh, S. Morse, C. Becker, M. Balabanovic, E. Gat, R. Simmons, S. Goodridge, H. Potlapalli, D. Hinkle, K. Jung, and D. Van Vactor. The winning robots from the 1993 robot competition. *AI Magazine*, pages 51–62, Winter 1993.

- [95] K. Oatley. Mental maps for navigation. *New Scientist*, pages 863–866, 19 December 1974.
- [96] J. O’Keefe. A review of the hippocampal place cells. *Progress in Neurobiology*, 13:419–439, 1979.
- [97] J. O’Keefe. Computations the hippocampus might perform. In Nadel et al. [93], chapter 8.
- [98] J. O’Keefe. The hippocampal cognitive map and navigational strategies. In J. Paillard, editor, *Brain and Space*, chapter 17. Oxford University Press, 1991.
- [99] J. O’Keefe and D. H. Conway. Hippocampal place units in the freely moving rat: Why they fire where they fire. *Experimental Brain Research*, 31:573–590, 1978.
- [100] J. O’Keefe and J. Dostrovsky. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 34:171–175, 1971.
- [101] J. O’Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford, 1978.
- [102] J. O’Keefe and M. L. Recce. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3):317–330, 1993.
- [103] J. O’Keefe and A. Speakman. Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, 68:1–27, 1987.
- [104] S. M. O’Mara, E. T. Rolls, A. Berthoz, and R. P. Kesner. Neurons responding to whole-body motion in the primate hippocampus. *Journal of Neuroscience*, 14(11):6511–6523, 1994.
- [105] T. Ono, R. Tamura, and K. Nakamura. The hippocampus and space: Are there “place neurons” in the monkey hippocampus? *Hippocampus*, 1(3):253–257, 1991.
- [106] M. A. Penna and J. Wu. Models for map building and navigation. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(5):1276–1301, 1993.
- [107] D. Pierce and B. Kuipers. Learning hill-climbing functions as a strategy for generating behaviors in a mobile robot. In Meyer and Wilson [80], pages 327–336.
- [108] M. Potegal. The special nature of spatial information. *The Behavioral and Brain Sciences*, 5:647–648, 1982. commentary on Lieblisch and Arbib’s target article.
- [109] B. Poucet. Evaluation of connectedness by cats in path-selection problems. *Perceptual and Motor Skills*, 58:51–54, 1984.
- [110] B. Poucet. Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review*, 100:163–182, 1993.
- [111] B. Poucet, C. Thinus-Blanc, and N. Chapuis. Route-planning in cats, in relation to the visibility of the goal. *Animal Behaviour*, 31:594–599, 1983.
- [112] A. J. Prescott. *Explorations in reinforcement and model-based learning*. PhD thesis, Dpt of Psychology, University of Sheffield, 1994.

- [113] G. J. Quirk, R. U. Muller, and J. L. Kubie. The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience*, 10(6):2008–2017, 1990.
- [114] G. J. Quirk, R. U. Muller, J. L. Kubie, and J. B. Ranck. The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *Journal of Neuroscience*, 12(5):1945–1963, 1992.
- [115] A. D. Redish and D. S. Touretzky. Navigating with landmarks: Computing goal locations from place codes. In K. Ikeuchi and M. Veloso, editors, *Symbolic Visual Learning*. Oxford University Press, 1995. to appear.
- [116] A. D. Redish and D. S. Touretzky. Modeling interactions of the rat's place and head direction systems. In Touretzky et al. [149], pages 61–67.
- [117] F. Restle. Discrimination of cues in mazes: A resolution of the 'place-vs-response' question. *Psychological Review*, 64(4):217–228, 1957.
- [118] R. L. Riolo. Lookahead planning and latent learning in a classifier system. In Meyer and Wilson [80], pages 316–326.
- [119] E. T. Rolls. Functions of the primate hippocampus in spatial processing and memory. In R. P. Kesner and D. S. Olton, editors, *Neurobiology of Comparative Cognition*, chapter 12. Lawrence Erlbaum Associates, 1990.
- [120] E. T. Rolls. Functions of the primate hippocampus in spatial and nonspatial memory. *Hippocampus*, 1(3):258–261, 1991.
- [121] E. T. Rolls. A model of the operation of the hippocampus and entorhinal cortex in memory. *International Journal of Neural Systems*, 6(Supp. 1995):51–70, 1995.
- [122] E. T. Rolls and S. M. O'Mara. View-responsive neurons in the primate hippocampal complex. *Hippocampus*, 5:409–424, 1995.
- [123] E. T. Rolls and M. J. Tovee. Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *Journal of Neurophysiology*, 73(2):713–726, 1995.
- [124] J. W. Rudy, S. Stadler-Morris, and P. Albert. Ontogeny of spatial navigation behaviors in the rat: Dissociation of "proximal"- and "distal"-cue-based behaviors. *Behavioral Neuroscience*, 101(1):62–73, 1987.
- [125] D. E. Rumelhart and J. L. McClelland, editors. *Parallel Distributed Processing, Vol 1*. The MIT Press, 1986.
- [126] D. E. Rumelhart and D. Zipser. Feature discovery by competitive learning. In Rumelhart and McClelland [125], pages 151–193.
- [127] N. A. Schmajuk and A. D. Thieme. Purposive behavior and cognitive mapping: A neural network model. *Biological Cybernetics*, 67:165–174, 1992.
- [128] N. A. Schmajuk, A. D. Thieme, and H. T. Blair. Maps, routes, and the hippocampus: A neural network approach. *Hippocampus*, 3(3):387–400, 1993.

- [129] B. Schölkopf and H. A. Mallot. View-based cognitive mapping and path planning. Technical Report TR-7, Max Planck Institut für biologische Kybernetik, 1994.
- [130] M. Seibert and A. M. Waxman. Adaptive 3D object recognition from multiple views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:107–124, 1992.
- [131] P. E. Sharp. Computer simulation of hippocampal place cells. *Psychobiology*, 19(2):103–115, 1991.
- [132] P. E. Sharp, H. T. Blair, D. Etkin, and D. B. Tzanetos. Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. *Journal of Neuroscience*, 15(1):173–189, 1995.
- [133] P. E. Sharp and C. Green. Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *Journal of Neuroscience*, 14(4):2339–2356, 1994.
- [134] P. E. Sharp, J. L. Kubie, and R. U. Muller. Firing properties of hippocampal neurons in a visually symmetrical environment: Contributions of multiple sensory cues and mnemonic processes. *Journal of Neuroscience*, 10(9):3093–3105, 1990.
- [135] W. E. Skaggs and B. L. McNaughton. Replay of neuronal firing sequences in rat hippocampus during sleep after spatial experience. *Science*, 271:1870–1873, 1996.
- [136] W. E. Skaggs, B. L. McNaughton, M. A. Wilson, and C. A. Barnes. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, (In Press), 1996.
- [137] A. Speakman and J. O’Keefe. Hippocampal complex spike cells do not change their place fields if the goal is moved within a cue controlled environment. *European Journal of Neuroscience*, 2:544–555, 1990.
- [138] B. E. Stein and M. A. Meredith. *The merging of the senses*. The MIT Press, 1993.
- [139] R. S. Sutton. Reinforcement learning architecture for animats. In Meyer and Wilson [80], pages 288–296.
- [140] L. W. Swanson, C. Köhler, and A. Björklund. The limbic region. I: The septohippocampal system. In A. Björklund, T. Hokfelt, and L. W. Swanson, editors, *Handbook of chemical neuroanatomy, Vol 5*, pages 125–277. Elsevier, 1987.
- [141] J. S. Taube. Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, 15(1):70–86, 1995.
- [142] J. S. Taube, R. U. Muller, and J. B. Ranck. Head-direction cells recorded from the postsubiculum in freely moving rats: Description and quantitative analysis. *Journal of Neuroscience*, 10(2):420–435, 1990.
- [143] J. S. Taube, R. U. Muller, and J. B. Ranck. Head-direction cells recorded from the postsubiculum in freely moving rats: Effects of environmental manipulations. *Journal of Neuroscience*, 10(2):436–447, 1990.

- [144] C. Thinus-Blanc. La discrimination entre les deux modalités (ouvert/fermé) d'une propriété topologique chez le hamster doré. Effet des modifications métriques. *L'Année Psychologique*, 78:7–27, 1978.
- [145] C. Thinus-Blanc. Animal spatial cognition. In L. Weiskrantz, editor, *Thought without language*, pages 371–395. Clarendon, 1988.
- [146] C. Thinus-Blanc. Mémoire spatiale distribuée. In *Actes du 5ème Colloque de l'ARC*, 1992.
- [147] L. T. Thompson and P. J. Best. Place cells and silent cells in the hippocampus of freely-behaving rats. *Journal of Neuroscience*, 9(7):2382–2390, 1989.
- [148] E. C. Tolman. Cognitive maps in rats and men. *The Psychological Review*, 55:189–208, 1948.
- [149] D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors. *Advances in Neural Information Processing Systems 8*. MIT Press, 1996.
- [150] D. S. Touretzky and A. D. Redish. Landmark arrays and the hippocampal cognitive map. In L. Niklasson and M. Boden, editors, *Current trends in connectionism - Proceedings of the 1995 Swedish Conference on Connectionism*, pages 1–13. Lawrence Erlbaum, 1995.
- [151] D. S. Touretzky, A. D. Redish, and H. S. Wan. Neural representations of space using sinusoidal arrays. *Neural Computation*, 5:869–884, 1993.
- [152] A. Treves and E. T. Rolls. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, 2:189–199, 1992.
- [153] A. Treves and E. T. Rolls. Computational analysis of the role of the hippocampus in memory. *Hippocampus*, 4(3):374–391, 1994.
- [154] C. H. Vanderwolf. Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalography and Clinical Neurophysiology*, 26:407–418, 1969.
- [155] H. S. Wan, D. S. Touretzky, and A. D. Redish. Towards a computational theory of rat navigation. In M. Mozer, P. Smolensky, D. S. Touretzky, J. L. Elman, and A. Weigend, editors, *Proc. of the 1993 Connectionist Models Summer School*, pages 11–19. Lawrence Erlbaum Associates, 1994.
- [156] T. H. Waterman. *Animal Navigation*. Scientific American Library, 1989.
- [157] C. J. C. H. Watkins and P. Dayan. Technical note: Q-learning. *Machine Learning*, 8:279–292, 1992.
- [158] Christopher Watkins. *Learning from Delayed Rewards*. PhD thesis, King's College, 1989.
- [159] R. Wehner. 'Matched filters' - neural models of the external world. *Journal of Comparative Physiology A*, 161:511–531, 1987.
- [160] R. Wehner and R. Menzel. Do insects have cognitive maps? *Annual Review of Neuroscience*, 13:403–414, 1990.

- [161] S. I. Wiener. Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *Journal of Neuroscience*, 13:3802–3817, 1993.
- [162] S. I. Wiener and A. Berthoz. Forebrain structures mediating the vestibular contribution during navigation. In A. Berthoz, editor, *Multisensory Control of Movement*, chapter 25. Oxford University Press, 1993.
- [163] S. I. Wiener, V. A. Korshunov, R. Garcia, and A. Berthoz. Inertial, substratal and landmark cue control of hippocampal CA1 place cell activity. *European Journal of Neuroscience*, 7:2206–2219, 1995.
- [164] D. M. Wilkie and R. Palfrey. A computer simulation model of rats' place navigation in the Morris water maze. *Behavior Research Methods, Instruments, and Computers*, 19(4):400–403, 1987.
- [165] M. A. Wilson and B. L. McNaughton. Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055–1058, 1993.
- [166] R. Worden. Navigation by fragment fitting: A theory of hippocampal function. *Hippocampus*, 2(2):165–188, 1992.
- [167] K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, 16:2112–2126, 1996.
- [168] D. Zipser. A computational model of hippocampal place fields. *Behavioral Neuroscience*, 99:1006–1018, 1985.
- [169] D. Zipser. Biologically plausible models of place recognition and goal location. In Rumelhart and McClelland [125], chapter 23.