

Biomimetic Navigation Models and Strategies in Animats

Olivier Trullier (1,2,3) 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

AICom Journal Vol 10, Number 2 (expected June 1997)

Tel: 01 44 32 36 33 / 01 44 27 16 21

Fax: 01 44 32 39 01

20 pages, 10 figures

Abstract

This paper describes a hierarchy of four navigation strategies — guidance, place recognition-triggered response, topological navigation and metric navigation. Such a hierarchy can be used to categorize models that are inspired by current knowledge about the way animals navigate in their environments. The main mechanisms implemented in each model are described, together with the basic adaptive capacities that the corresponding strategy affords. Because biomimetic models have seldom been implemented in real robots, it is premature to compare their merits with those of traditional engineering solutions to the navigation problem. Nevertheless, the methodological options that such implementations would entail are discussed in the text.

1 Introduction

Animals are living proofs that any system, equipped with proper sensors, proper actuators, and a proper control architecture, can exhibit an adaptive behavior that allows it to survive in environments that can be quite unpredictable and challenging. Such a capacity appears highly desirable in many application areas, notably in robotics, where there is an urgent need for autonomous artifacts that should be able to adapt their behavior to changing circumstances. This need has prompted numerous researchers to devise so-called *animats*, i.e., simulated animals or real robots, whose structure and functionalities are as much inspired from those of animals as possible, and which exhibit some variety of adaptive behavior [25, 24, 7, 15]. Thus, animats can evolve, develop, learn, memorize, plan and communicate [22, 23, 21]. Like animals, instead of being passive reflexive devices, animats often are active data processors that seek useful information in their environment, encode it into internal representations of objects and causal relations, and use these representations for their own benefit in flexible and intelligent ways [20].

The ability to navigate — i.e., to reach as directly as possible any goal location from any starting point, while avoiding passing through specific places - is a basic requisite for an animat's survival. Indeed, without this ability, the animat wouldn't be able to reach energy sources, to avoid encounters with damaging obstacles, or to escape from dangerous hazards. Numerous engineering solutions to the navigation problem have been proposed in the literature, which usually resort to internal geometrical representations of the environment [14, 10]. However, these lead to various implementation difficulties, due to memory or time requirements, as well as sensory and motor errors [29]. Since many animals are proficient at navigating in more efficient and robust ways, it is not surprising that research efforts have recently been aimed at incorporating the relevant neuro-biological data into animat navigation models. This paper provides a summary of such models and categorizes them according to four different navigation strategies — which we call *guidance*, *place recognition-triggered response*, *topological navigation*, and *metric navigation* — that will be successively described and discussed. In this review, we will focus on some of the more successful approaches that are representative of each navigation strat-

egy. A more comprehensive review of these models and strategies, together with a description of the underlying biological data, is to be found in [42].

2 Guidance

The simplest navigation strategy exhibited by animals is guidance, where animals either move toward a goal, or toward a beacon indicating the presence of a goal that can be directly sensed. A variant of this strategy consists in moving along a trajectory that minimizes the discrepancy between the configuration of landmarks sensed in the current place and the (remembered) configuration of the same landmarks sensed at the goal place. Although the guidance strategy eventually entails using some memorized information — i.e., that some sensory information is associated with the goal — it doesn't require that this information explicitly denotes space. In other words, the efficient navigation skills exhibited by some animals, notably insects, do not necessarily require *spatial representations*.

This point is clearly demonstrated by the model of Cartwright and Collett [5], which is inspired from the way bees use nearby landmarks to return to an already known food source. This ability relies upon the use of a remembered *snapshot*, that is, the pattern of activation of their retinas when they were at the food source: bees seem to adjust their flight path by continuously comparing their retinal image of the landmark configuration with their remembered snapshot. This lessens the discrepancy between the two (Figure 1).

PUT FIGURE 1 AROUND HERE

The navigation model of Benhamou et al. [1] calls upon the same principle but, instead of relying on sensory discrepancies at the current location, it involves estimates of the (memorized) discrepancies that would be found at nearby locations. The animat then moves toward the location of least discrepancy.

Likewise, Wilkie and Palfrey [45] propose a “simple perceptual memory-matching model”, according to which an animat moves toward or away from each of a set of landmarks, depending upon the difference between the memorized distance from the goal location to the landmark and the current distance from the animat to the landmark.

3 Place recognition-triggered response

Navigation strategies of animals with sophisticated adaptive skills, notably rodents, entail the use of spatial representations that seem to be encoded in the *hippocampus* and *para-hippocampal formations* in the brain [30] and to involve so-called *place cells* and *head direction cells*. The activity of the former depends on the animal's current location in relation to the spatial layout of surrounding landmarks,

irrespective of the direction faced (for a review, see [27]), whereas the activity of the latter depends upon the absolute orientation of the head, regardless of the animal's location [37, 38].

Several systems have employed elements bearing, to varying degrees, resemblance to place cells and head direction cells. These implement what we call a place recognition-triggered response navigation strategy: according to this strategy, an animat is able to explore its environment and to memorize in various specific places, each characterized by a specific landmark configuration, the direction in which a given goal is situated. Later, the animat can use the information provided by surrounding landmarks to recognize the place it is currently situated in and to move in the memorized direction to the goal.

The model of Burgess, Recce and O'Keefe [4] implements a 5-layer feedforward neural network (Figure 2) that largely reproduces the rat hippocampal architecture. The first layer represents the highly processed information the hippocampus receives as input and consists of sensory neurons that discharge selectively when respective associated landmarks are visible at a certain distance from the animat. The next three layers respectively correspond to *entorhinal cortical cells* (EC), to pyramidal *place cells* (PC) of hippocampal regions CA3 and CA1, and to *subicular cells* (SC). The connectivity of the network is adjusted by competitive learning during a preliminary exploratory phase. At the end of this process, neurons in the three hidden layers act as place cells each characterized by a *place field* — i.e., the region of space in which the animat is situated when the cell is active. Note that EC and SC cells are not considered as place cells in the classical sense. We will thus talk of their *activity fields* instead of their place fields.

The neural network also models experimental data from this same group showing a relation between the timing of CA1 neuronal discharges relative to the phase of the theta rhythm — a sinusoidal electro-encephalographic oscillation, ranging between 4Hz and 10Hz, that occurs essentially when locomotor behaviors are performed — and the position of the rat within the place field. Thus each model place cell tends 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.

After the spatial representation has been built up by the network during a first phase of exploration, it turns out that, when the animat moves in its environment, each place cell is selectively activated, depending upon the position and orientation of the animat relative to the landmarks. The ensemble pattern of activity uniquely characterizes the place the animat is currently situated and has been proposed to constitute a *cognitive map* after Tolman [39]. See Poucet [31] for a discussion of this concept.

During a second phase of exploration, the animat learns the locations of reward sites. This learning involves 8 *goal cells* that are recruited in the last layer of Burgess et al.'s network at each reward site in the environment (one cluster of 8 cells for each goal, as shown in Figure 2). The activity of each goal cell is tuned to one of 8 directions — arbitrarily labeled North, North-East, East, and so on — given by a corresponding head-direction cell. When the animat encounters a reward, it rotates

around slowly. As it looks into each of the 8 directions, the SC neurons that fire at a late phase with respect to the theta rhythm are those with activity fields that lie ahead of the animat in the corresponding direction. For example, when the animat looks North, the activity fields of all the SC neurons that fire during the late phase of the theta cycle lie to the North of the reward site. Each goal cell in the cluster makes associative connections to all the active SC neurons in the corresponding direction. Subsequently, the activity field of each goal cell amounts to the sum of the activity fields of the associated SC neurons, and thus lies ahead of the animat in the associated direction from the goal. 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 3 illustrates the activity fields of two goal cells. Subsequently, when the animat is for example in location A in Figure 3, cells representing “South”, “South-East” and “East” of the goal location are more active 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 4 shows that this model can represent multiple goals. Obstacles can also be taken into account by considering them as negative goals.

PUT FIGURE 2 AROUND HERE

PUT FIGURE 3 AROUND HERE

PUT FIGURE 4 AROUND HERE

The animat of Zipser [47] also calls upon place cells and goal cells. Place cells’ activities are tuned to the distance and bearings of given landmarks seen from the corresponding places. Goal cells are associated with place cells and code the direction of the goal from the corresponding places in the form of vectors. The direction in which the animat should move in order to reach the goal is given by the weighted average of the vectors provided by the goal cells, each vector weighted by the activity level of the associated place cell.

The navigation model of Gaussier and Zrehen [9] doesn’t assume that landmarks are known and recognized a priori. Instead, this model is built upon a collection of neural networks that each perform some active perception-recognition process through feedback loops. During a preliminary learning phase, this process associates the characterization and recognition of a given landmark to the sequence of ocular

saccades that links the most salient points on the object. Likewise, when the animat explores its environment and encounters a reward site, it wanders around this site and characterizes a number of surrounding places from which the goal is visible by the sequence of movements (eye, head, or body rotation) required to orient its perceptual systems toward the landmarks that are observable in each place. Each place thus characterized is associated with the corresponding direction to the goal.

Place cells that are used in the model of Blum and Abbott [2] are connected through modifiable synapses. During exploration, when the animat reaches a goal, a reinforcement learning process is triggered. This enhances the synaptic weights between cells with overlapping place fields that the animat has recently traveled through. At the end of the exploratory phase, the shortest path taken to the goal is associated with the highest synaptic weights between corresponding place cells. Subsequent navigation to the goal entails moving successively through place fields that correspond to strongly connected place cells.

In the model of Brown and Sharp [3], place cells and head direction cells converge onto two clusters of motor output cells whose activity make the animat move a small step forward and then turn right or left, respectively. During exploration, connections between sensory cells and place cells, as well as connections between head direction cells and motor output cells, are submitted to competitive learning processes, according to which only one place cell is active when the animat is in a given place. This information, together with that provided by the currently active head direction cell, permits only one cell to become active in each motor output cluster. After learning, the most active of the two output cells correctly indicates the direction to turn in order to approach the goal.

Finally, although the model of Krakauer [12] doesn't aim at simulating hippocampal place cells, it illustrates how bees might use distances and bearings of given landmarks to navigate according to a place recognition-triggered response strategy. In this hierarchical version of the snapshot model of Cartwright and Collett [5] mentioned above, a first-level neural network receives filtered input from distal landmarks and learns to have only one output channel active at a given time. This channel represents the "catchment area" [6] of one specific snapshot. This output activates a second-level neural network that receives a filtered input from proximal landmarks, and is trained through back-propagation to produce a vector signal directed towards the goal.

Such implementations of place recognition-triggered response strategies endow an animat with minor adaptive capacities. In particular, its ability to self-estimate its position in the environment relies upon information that is often distributed across numerous place cells in its "cognitive map", and the accidental destruction of a few critical place cells would cause a graceful degradation of the animat's navigation capacities. Additionally, implementations like that of Burgess et al. exhibit the highly advantageous capacity of *latent learning*, that is, in the current context, the capacity to learn a spatial representation even in the absence of any reward [39]. However, the usefulness of place recognition-triggered response strategies is severely challenged when, along its path to the goal, an animat encounters an obstacle that was not present at exploration time. In such circumstances, its only chance to avoid

getting lost and to eventually reach its goal consists in wandering randomly around the current place, until it arrives at another already known place, from which it could resume its trajectory toward the goal (Figure 5).

PUT FIGURE 5 AROUND HERE

4 Topological navigation

A major improvement over the previous navigation strategy calls upon learning and exploitation of the topological links that characterize the animat's environment. During exploration, when the animat moves from one place to another, a link is drawn between the place cells that are associated with these places and this link is labeled by the associated move. Thus, later on, because the animat has encoded which moves are likely to lead from one place to another, it can plan a trajectory to the goal, according to a *topological navigation* or *route following* strategy.

Such a strategy has been implemented by Schmajuk and Thieme [35]. Their model of animat navigation relies on a neural network with specific dynamical properties of activity propagation that can represent the topological layout of places and thus enable path planning. It is based on two modules, one that encodes the topological representation and one that selects movements on the basis of predictions generated by the first module (Figure 6).

PUT FIGURE 6 AROUND HERE

As in many models belonging to this category, place cells are given a priori. The navigation system's inputs consist of the identity of the current place and the identities of the adjacent places (called "views") that can be seen from the current place. As shown in Figure 7, 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 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. More precisely, the 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. This weight is decreased if the animat is at P but doesn't perceive view V (that is, the corresponding nodes are not simultaneously activated). When learning is completed,

these synaptic weights represent the transformation from places to views and, since views correspond to adjacent places, synaptic weights are equivalent to topological links 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, if one assumes the animat is in place 1 of Figure 7, nodes p2 and p3 will first be activated since places 2 and 3 are “viewed” from place 1. The activity of p2 will feed back to drive nodes p1, p3, and pG, while p3 drives nodes p1 and p2 (through strongly weighted synapses). Thus, the existence of a path leading to the goal from p1 can be inferred. However, because these recurrent signals shouldn’t create interferences with local signals perceived at the current place — so that a distinction can be kept between multi-step goal finding and single-step prediction of adjacent places — Schmajuk and Thieme make the assumption that there are two rates of change of the transmitted signals, fast and slow¹ (the dynamics of the activities of cells are governed by differential equations). Fast signals are involved in goal predictions, slow signals in prediction of the following step. The recurrent connections involve an intermediate layer of neurons (labeled *d* in Figure 7) 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 cannot contradict the stronger next-place prediction signals. This ensures that no false detection of mismatch between predictions and views will occur. For instance, node pG is activated by the prediction of p2 from place 1, but there is no contradiction between pG then being weakly activated and place G not being seen from place 1. Thus, the network operates as a non-recurrent network for slow-changing signals — thereby preserving the correct internal representation of the local environment — and as a recurrent network for fast-changing signals — thereby allowing the inference of paths to remote places. Moreover, navigational computations proceed at a faster pace than real movement.

PUT FIGURE 7 AROUND HERE

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

¹According to Schmajuk and Thieme, such an assumption is biologically plausible. However, there is presently no evidence for the co-existence of these two types of signals from neurophysiological data.

selection module” of Figure 7. It spreads through the network and the recurrent connections, until it eventually activates the node predicting the goal. The fast-changing signal is also 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 goal-prediction pG node. Thus, the connection weights r correspond to the topological distance between the respective places and the goal. The Action Selection module of Figure 6 compares these weights. The shortest path between the current place and the goal (in the number of intermediate places) starts at the place which generated the strongest signal at 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, it means either that the goal is far away, or that the spatial representation is incomplete, and in this case, the Action Selection module generates a random movement.

The model of Muller et al. [28] also encodes topological links in synapses that connect place cells. Place cells are given a priori and synaptic weights are learned during random exploration. When it moves from one place to another at constant velocity, the corresponding place cells fire with a temporal delay that is proportional to the distance between their place fields. The closer the place fields, the shorter the delay, and the greater the enhancement of the weight of the synapse(s) connecting these cells. According to this mechanism, the synaptic weights eventually become proportional to the inverse of the distance between place fields, a metric quantity. Once learning is completed, routes from the current location to a given goal can be planned by a classical graph search algorithm. The shortest path corresponds to the set of connected place cells for which the sum of the synaptic weights is the greatest.

Schölkopf and Mallot’s [36] navigation system relies upon a neural network with a visual input layer, a movement input layer, and a map layer made of interconnected place cells. These connections code the topological links between places and the information flowing through them is gated by modulatory connections drawn from units in the movement layer, according to *presynaptic facilitation* mechanisms that are known to exist in the brain (e.g. [44]). The “cognitive map” is built according to two learning processes occurring during exploration. The first makes a given place cell fire when a specific pattern is provided as visual input. The second makes the same place cell fire when the animat “expects” this place to be reachable from another one by a certain movement decision. When the map is built, path planning is performed according to a straightforward procedure: (i) the visual input pattern corresponding to the starting place is activated, thus activating the appropriate place cell; (ii) different possible movements out of this place are simulated by activating the corresponding movement nodes, letting the activation spread through facilitated topological links for one time step, and recording where the movement leads to; (iii) the search is restarted from the new place until the goal is reached.

In Kuipers and Byun’s [13] approach, the functional role of place cells and synapses

is played by nodes and edges within a graph. These nodes and edges encode many different types of information, including metric, which are used by the navigation system. In particular, each node is characterized by the local maximum of a so-called *distinctiveness measure*, that is by a function of the different sensor readings of the animat. These correspond to landmarks. For instance, a given node on the animat's cognitive map can refer to the center of a T-junction in a maze and be characterized by a function that evaluates the differences between readings from 3 range sensors spaced at 120° from each other. When the differences are minimized, the 3 readings are equal and the animat is at the junction. Likewise, a given edge can refer to the fact that the animat moved, under the control of a "move-along-the-wall-on-right" strategy, from one particular location to another. An important feature of this system is that it can generate predictions and plans for exploration to disambiguate places with similar sensory attributes. It is however limited by its locomotion strategy, according to which it can only follow walls, follow corridors or move toward a specific remote landmark. The latter functionality entails moving towards the maximum of the corresponding distinctiveness measure — a typical low-level guidance strategy. When the "cognitive map" is built, optimal path planning toward a goal is performed by graph search using the available metric information.

Mataric's navigation system [16] is similar to the preceding in many respects. It allows a robot to explore its environment by following its walls and to use the information provided by its sensors and its effectors to characterize different places it encounters in its environment. This information is used to build a topological representation that records the possible moves from place to place. Thus, while the robot moves in its environment, it records that a given place A is passed through while it moves in an eastward direction and while its sonars detect similar nearby obstacles on its right and on its left — thus suggesting that place A is a corridor facing east. Likewise, it records that a right turn leads from this place to another place B, which is passed through while the robot moves in a south-eastward direction and only while its right sonars detect a nearby obstacle — thus indicating that the corridor is adjacent to a wall to the right which is oriented south-easterly. In other words, representations of places like A and B in the form of nodes in the robot's cognitive map resemble place cells. Because these nodes also provide information about the physical length of each place (i.e. corridor or wall), the shortest path leading to any given goal from the current place can be generated by initiating a spreading activation process throughout the graph of interconnected nodes, in all directions from the goal. Insofar as the speed of this process depends upon the length of the places coded by the nodes through which it travels, the direction from which goal-initiated activation first arrives in the current place node indicates in which direction to move in order to reach the goal.

Besides the previously mentioned possibilities of graceful degradation afforded by eventually decentralized implementations, trajectory planning endows topological navigation with additional highly adaptive capacities. In the same spirit as the *experience-in-the-head* metaphor [8], trajectory planning makes it possible to run internal simulations that, being decoupled from overt behavior, are efficient and avoid the hazards of dangerous encounters. In particular, topological navigation allows

an animat to plan a detour in the presence of a new obstacle (Figure 8). However, it should be noted that such a detour necessarily entails passing through already known places, and there is no *insight* [11, 40] of a trajectory through unexplored areas.

PUT FIGURE 8 AROUND HERE

5 Metric navigation

Finally, a fourth navigation strategy, called metric navigation or survey mapping, seems to be used by animals such as dogs, monkeys and men. This strategy, which calls upon the representation of metric spatial relations between two or more places on a plane seems to imply the reading of a 2-D map, or a “view from above” of the whole environment. Naturally, the actual mechanisms that are involved may have nothing in common with this interpretation. Nonetheless, all the computational models resort to vector manipulations and systems of coordinates.

Metric navigation is implemented in the model of Wan et al. [43] that integrates multi-modal information involving signals brought by efferent copies,² signals brought by the integration of the angular velocity of the head (sensed by vestibular organs), and exteroceptive signals brought by the visual sensory system (Figure 9). In this model, the correspondence between the animat’s spatial representation and its position in the environment is coded by the activity of *place code units*. This activity is determined by the product of seven Gaussian functions, respectively tuned during exploration to the distances and allocentric bearings of two selected landmarks, to the egocentric bearing difference between two other selected landmarks, and to the estimated Cartesian coordinates of the animat’s position relative to an identified reference point. Should one or more of the terms in this calculation be missing, they are simply excluded from the product. Thus, this model accounts for the experimental evidence for persistence of place cell activity in darkness [33], because information brought by the path integrator mechanism is likely to substitute for the missing visual information. The local view module [19] used in this model transforms the (egocentrically sensed) visual information into an allocentric representation that is stored in the place code module, by simply summing the head direction and the egocentric bearing. Conversely, by subtracting the egocentric bearing provided by the visual input from the allocentric bearing registered in the place code module, the local view module makes it possible to update the estimate of the head direction. Thus, the model also accounts for the experimental evidence of re-setting of head direction representations by (new) visual information [38]. Likewise,

²Motor subsystems that drive muscles can map an “efference copy” or “corollary discharge” back onto the perceptual system, which can thus distinguish self-induced movements from world-induced movements [18].

the place code module used in this model makes it possible to correct the Cartesian coordinates of the animat, thus compensating for drift errors that arise from dead-reckoning. Finally, these updated Cartesian coordinates can be used by a vector manipulation module to plan a trajectory to any given goal.

PUT FIGURE 9 AROUND HERE

Navigation tasks performed by Worden's formal model [46] rely on the hypothesis that the hippocampus stores memories of environmental landmark relations as a large number of independent *fragments*. A typical fragment encodes metric relationships between a triplet of prominent landmarks in some region, together with non-geometric properties of these landmarks, such as smells and colors. According to this hypothesis, navigation entails forming a local geometric map by translating and rotating appropriate fragments in an absolute reference frame in order to fit together the landmarks that are common to different fragments. This "jigsaw puzzle" assembly process ends when the map includes both the animat's current place and goal place. Then, the direction to the goal is given by the difference in positions between the current place and the goal within this network of fragments in the absolute reference frame.

In the model of Prescott [32], space is also represented by fragments, which constitute local coordinate frameworks characterized by groups of three salient landmarks. In this model, specific locations are redundantly coded by their relations to a number of fragments, as *object-units* and *beta-units* within a neural network. The activations of the former represent the locations in egocentric references of specific goals and salient landmarks. The latter each receive inputs from three object-units and adapt their connection strengths to match their output to the activation of a fourth object-unit. In other words, beta-units predict the coordinates of a fourth landmark with respect to the fragment defined by three landmarks. Thus, the animat can determine the location of any target landmark when it is within sight of any group of three others. Again, this entails triggering a spreading activation process. However, because landmark positions are encoded redundantly, the system may generate multiple solutions for a given navigation task, in which case a simple heuristic is used to select among alternatives. A computational advantage of this model is that the spatial representation is robust with regard to noisy sensory inputs, because the coordinates of the landmarks are not fixed but are dynamically estimated as the animat moves about and recognizes different landmarks.

Besides the particular adaptive capacities of the run-time error correcting processes that are implemented in the models of Wan et al. and of Prescott, metric navigation affords an animat considerable advantages enhancing the chances for survival, since it allows *metric detours* and *metric shortcuts* through unexplored areas (Figure 10). It also makes possible to compute the shortest path leading to any goal from the current place.

PUT FIGURE 10 AROUND HERE

6 Summary and Conclusion

The very diversity of these biomimetic models raises new questions for further psychobiological research in order to fill the gaps in our understanding of the strategies and mechanisms that animals actually use to navigate. In this perspective, it is worth noting that several models [47, 45, 35, 2], which implement different navigation strategies, all aim to explain the same experimental results obtained with rats in the Morris water maze [26]. It is also clear that, if the physiology of the hippocampus is partially known, its functional role — whether it is specialized in the storage of spatial representations or whether it acts as a general associative memory (e.g. [34]) — is still under debate, without mentioning the uncertainties about what information is carried by signals it receives as input and delivers as output. Likewise, the question of how the same neural circuitry could store and retrieve spatial information concerning multiple environments is troublesome. In addition, the criteria for selection of the landmarks a given animal uses to localize or orient itself in a given environment is seldom perfectly known. Finally, there are uncertainties on the criteria for the landmarks to be recognized independently in succession, combined into a single percept or considered simultaneously as an ordered configuration of individuals.

To cope with these shortcomings, each model incorporates a number of *ad hoc* mechanisms that ensure its overall navigation functionalities. Regardless of the biological plausibility of these mechanisms, it appears that their implementations in the control architecture of a mobile robot have not yet been attempted, with the exception of Mataric’s well documented results [17] and of a preliminary report concerning Wan et al.’s model [41]. Accordingly, it is premature to speculate about the relative merits of biomimetic approaches to the robot navigation problem as compared to traditional engineering solutions.

Be it as it may, it is clear from the above review that the decision to implement a specific model will depend upon various methodological options at each of four levels:

- inputs: what information is fed into the navigation system.

Basically, three types of inputs are used in the above-mentioned models: (1) as the values of some environmental variables (e.g. [5]), (2) as the activation levels of some “neurons” that are responsive to a certain range of values of some environmental feature (e.g. [4]), or (3) as activations of some “neurons” depending on the presence of some environmental feature (e.g. [36]). For models that build up a place representation, distances and egocentric bearings of landmarks are provided either way, often together with landmarks’ identities

(e.g. [47]). For models that rely on an a priori place representation, each place is represented by a “neuron” that is assumed to recognize it (e.g. [35]).

- internal representations: how the control architecture is implemented, what spatial features are learned, and how the goal location is managed.

In numerous models, the control architecture is implemented as a neural network (e.g. [2]). Other models implement mathematical relations (e.g. [45]) or functional modules (e.g. [43]) that relate their inputs and outputs. Still other models call upon symbolic computations (e.g. [46]). Concerning learning, three basic mechanisms are usually used: (i) storing specific parameters such as distances, bearings or snapshots (e.g. [13]); (ii) recruiting place cells and fixing their input-output response (e.g. [9]); (iii) updating the synaptic weights of a neural network, usually through a form of Hebbian learning (e.g. [3]). As for the goal representation, it can be a stored set of values (e.g. [5]), or a “goal cell” (e.g. [4]), or a set of coordinates (e.g. [43]), or a specific landmark representation (e.g. [32]).

- outputs: how planning is performed, what information is sent to the motor system.

When planning is performed, it entails either a classical graph search (e.g. [28]), or a spreading activation process (e.g. [16]). The output of most models is a direction of movement, given as a vector (e.g. [12]), as a compass direction (e.g. [1]), or as a turning angle (e.g. [3]). Models that use an a priori place representation usually define the animat’s action as that of reaching a given place (e.g. [36]). In some cases, the model’s output is a command to a lower-level navigation strategy, typically a guidance instruction (e.g. [13]).

- environment: to which environment are the model’s navigation capacities restricted.

The navigation capacities of some models are restricted to maze-like environments, in which the configurations of places, as well as movements, are highly constrained (e.g. [35]). Likewise, some navigation systems only work in small open-field environments, where all the landmarks are visible from every location (e.g. [4]). Other models make navigation in large-scale environments possible, even when some landmarks are out of the animat’s range of perception (e.g. [32]).

Acknowledgements

The authors would like to thank Alain Berthoz and Sidney Wiener for helpful discussions about this work and Sidney Wiener for reading through the manuscript. This work was supported by the GIS Sciences de la Cognition.

References

- [1] Benhamou, S., Bovet, P., and Poucet, B. (1994): A place navigation algorithm based on elementary computing procedures and associative memories. In Cliff et al. [7] pp. 206–213.
- [2] Blum, K. I. and Abbott, L. F. (1996): A model of spatial map formation in the hippocampus of the rat. *Neural Computation* 8, 85–93.
- [3] Brown, M. A. and Sharp, P. E. (1995): 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.
- [4] Burgess, N., Recce, M., and O’Keefe, J. (1994): A model of hippocampal function. *Neural Networks* 7(6/7), 1065–1081.
- [5] Cartwright, B. A. and Collett, T. S. (1983): Landmark learning in bees. Experiments and models. *Journal of Comparative Physiology A* 151, 521–543.
- [6] Cartwright, B. A. and Collett, T. S. (1987): Landmark maps for honeybees. *Biological Cybernetics* 57, 85–93.
- [7] Cliff, D., Husbands, P., Meyer, J.-A., and Wilson, S. W., Eds. (1994): *From Animals to Animats 3: Proceedings of the Third International Conference on Adaptive Behavior*. The MIT Press/Bradford Book.
- [8] Craik, J. (1943): *The Nature of Explanation*. Cambridge University Press.
- [9] Gaussier, P. and Zrehen, S. (1994): Complex neural architectures for emerging cognitive abilities in an autonomous system. In Gaussier, P. and Nicoud, J.-D., Eds., *Proceedings of the PerAc Conference: from Perception to Action*. pp. 278–289. IEEE.
- [10] Hwang, Y. K. and Ahuja, N. (1992): Gross motion planning — a survey. *ACM Computing Surveys* 24(3), 219–291.
- [11] Köhler, W. (1927): *The mentality of apes*. Routledge and Kegan Paul. Reprinted 1976, Liveright.
- [12] Krakauer, D. C. (1995): Simple connectionist models of spatial memory in bees. *Journal of Theoretical Biology* 172, 149–160.
- [13] Kuipers, B. J. and Byun, Y. T. (1991): A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems* 8, 47–63.
- [14] Latombe, J.-C. (1991): *Robot Motion Planning*. Kluwer Academic.
- [15] Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J., and Wilson, S. W., Eds. (1996): *From Animals to Animats 4: Proceedings of the Fourth International Conference on Adaptive Behavior*. The MIT Press/Bradford Book.

- [16] Mataric, M. J. (1990): A distributed model for mobile robot environment-learning and navigation. Technical Report TR 1228 MIT AI Lab.
- [17] Mataric, M. J. (1992): Integration of representation into goal-driven behavior-based robots. *IEEE Transactions on Robotics and Automation* 8(3), 304–312.
- [18] McCloskey, D. I. (1981): Corollary discharges: Motor commands and perception. In Brooks, V. B., Ed., *Handbook of Physiology - The Nervous System. II. Motor control*. American Physiological Society.
- [19] McNaughton, B. L. (1989): Neuronal mechanisms for spatial computation and information storage. In Nadel, L., Cooper, L. A., Culicover, P., and Harnish, R. M., Eds., *Neural Connections and Mental Computations*. chapter 9. The MIT Press/Bradford Book.
- [20] Meyer, J.-A. (1995): The animat approach to cognitive science. In Roitblat, H. and Meyer, J.-A., Eds., *Comparative Approaches to Cognitive Science*. The MIT Press.
- [21] Meyer, J.-A. (1996): Artificial life and the animat approach to artificial intelligence. In Boden, M., Ed., *Computational Psychology and Artificial Intelligence*. Academic Press.
- [22] Meyer, J.-A. and Guillot, A. (1991): Simulation of adaptive behavior in animats: Review and prospect. In Meyer and Wilson [25].
- [23] Meyer, J.-A. and Guillot, A. (1994): From SAB90 to SAB94: Four years of animat research. In Cliff et al. [7].
- [24] Meyer, J.-A., Roitblat, H. L., and Wilson, S. W., Eds. (1993): *From Animals to Animats 2: Proceedings of the Second International Conference on Adaptive Behavior*. The MIT Press/Bradford Book.
- [25] Meyer, J.-A. and Wilson, S. W., Eds. (1991): *Proceedings of the First International Conference on Adaptive Behavior: from Animals to Animats*. The MIT Press/Bradford Book.
- [26] Morris, R. G. M. (1981): Spatial localization does not require the presence of local cues. *Learning and Motivation* 12, 239–260.
- [27] Muller, R. U., Kubie, J. L., Bostock, E. M., Taube, J. S., and Quirk, G. J. (1991): Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In Paillard, J., Ed., *Brain and Space*. chapter 17. Oxford University Press.
- [28] Muller, R. U., Kubie, J. L., and Saypoff, R. (1991): The hippocampus as a cognitive graph (abridged version). *Hippocampus* 1(3), 243–246.
- [29] Nehmzow, U. (1995): Animal and robot navigation. *Robotics and Autonomous Systems* 15, 71–81.

- [30] O'Keefe, J. and Nadel, L. (1978): *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford.
- [31] Poucet, B. (1993): Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review* 100, 163–182.
- [32] Prescott, A. J. (1995): Spatial representations for navigation in animats. *Adaptive Behavior* 4(2), 85–123.
- [33] Quirk, G. J., Muller, R. U., Kubie, J. L., and Ranck, J. B. (1992): The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *Journal of Neuroscience* 12(5), 1945–1963.
- [34] Rolls, E. T. (1991): Functions of the primate hippocampus in spatial and non-spatial memory. *Hippocampus* 1(3), 258–261.
- [35] Schmajuk, N. A. and Thieme, A. D. (1992): Purposive behavior and cognitive mapping: A neural network model. *Biological Cybernetics* 67, 165–174.
- [36] Schölkopf, B. and Mallot, H. A. (1995): View-based cognitive mapping and path planning. *Adaptive Behavior* 3(3), 311–348.
- [37] Taube, J. S., Muller, R. U., and Ranck, J. B. (1990): Head-direction cells recorded from the postsubiculum in freely moving rats: Description and quantitative analysis. *Journal of Neuroscience* 10(2), 420–435.
- [38] Taube, J. S., Muller, R. U., and Ranck, J. B. (1990): Head-direction cells recorded from the postsubiculum in freely moving rats: Effects of environmental manipulations. *Journal of Neuroscience* 10(2), 436–447.
- [39] Tolman, E. C. (1948): Cognitive maps in rats and men. *The Psychological Review* 55, 189–208.
- [40] Tolman, E. C. and Honzik, C. H. (1930): “Insight” in rats. *University of California Publications in Psychology* 4, 215–232.
- [41] Touretzky, D. S., Wan, H. S., and Redish, A. D. (1994): Neural representation of space in rats and robots. In Zurada, J. M., II, R. J. M., and Robinson, C. J., Eds., *Computational Intelligence. Imitating Life*. IEEE.
- [42] Trullier, O., Wiener, S. I., Berthoz, A., and Meyer, J.-A. (1997): Biologically-based artificial navigation systems: Review and prospects. *Progress in Neurobiology*. Sous Presse.
- [43] Wan, H. S., Touretzky, D. S., and Redish, A. D. (1994): Towards a computational theory of rat navigation. In Mozer, M., Smolensky, P., Touretzky, D. S., Elman, J. L., and Weigend, A., Eds., *Proc. of the 1993 Connectionist Models Summer School*. pp. 11–19. Lawrence Erlbaum Associates.
- [44] White, E. L. (1989): *Cortical Circuits*. Birkhäuser.

-
- [45] Wilkie, D. M. and Palfrey, R. (1987): A computer simulation model of rats' place navigation in the Morris water maze. *Behavior Research Methods, Instruments, and Computers* 19(4), 400–403.
 - [46] Worden, R. (1992): Navigation by fragment fitting: A theory of hippocampal function. *Hippocampus* 2(2), 165–188.
 - [47] Zipser, D. (1986): Biologically plausible models of place recognition and goal location. In Rumelhart, D. E. and McClelland, J. L., Eds., *Parallel Distributed Processing, Vol 1*. chapter 23. The MIT Press.

Figure captions

Figure 1

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. (c) The resulting movement, which globally reduces the discrepancy between the snapshots, is computed from the summed vector (filled arrow). After (Cartwright and Collett, 1983)

Figure 2

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)

Figure 3

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.

Figure 4

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)

Figure 5

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.

Figure 6

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)

Figure 7

(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 6. The synaptic weights (filled and empty 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)

Figure 8




(a) With the place recognition-triggered response strategy there can be an ensemble of intersecting routes. The animat is able to go only 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. 5). (b) In contrast, if the animat merges its representations of routes into a topological representation, the animat can then 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.


Figure 9

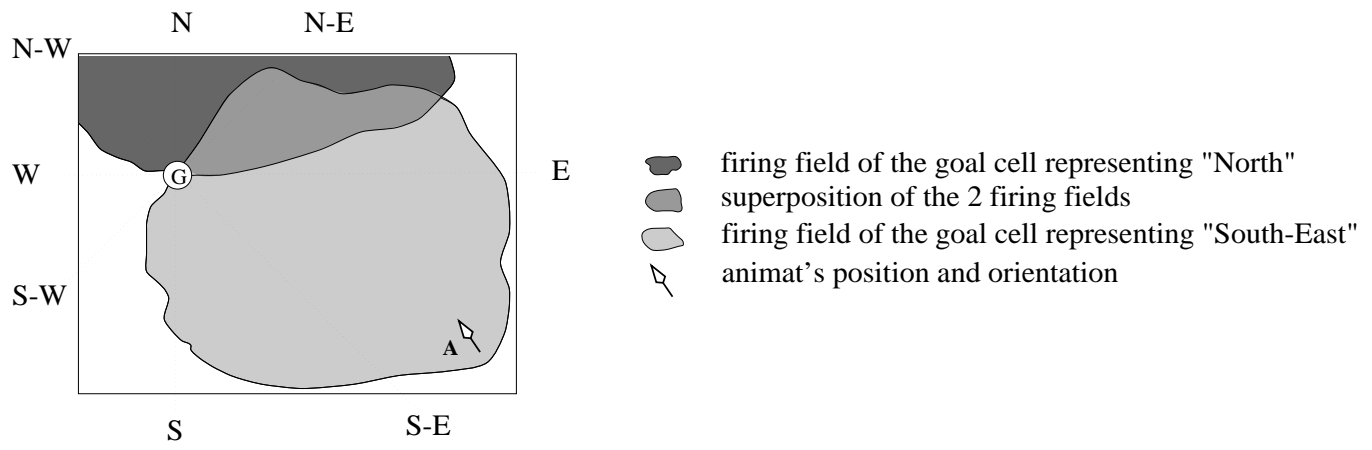
Functional block diagram of Wan et al.'s model. It shows how multimodal information might be combined and how different representations can update one another. Place is represented as the conjunction of multimodal inputs, coupling path integration, visual information, and head-direction. The path integrator 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)


Figure 10


(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.

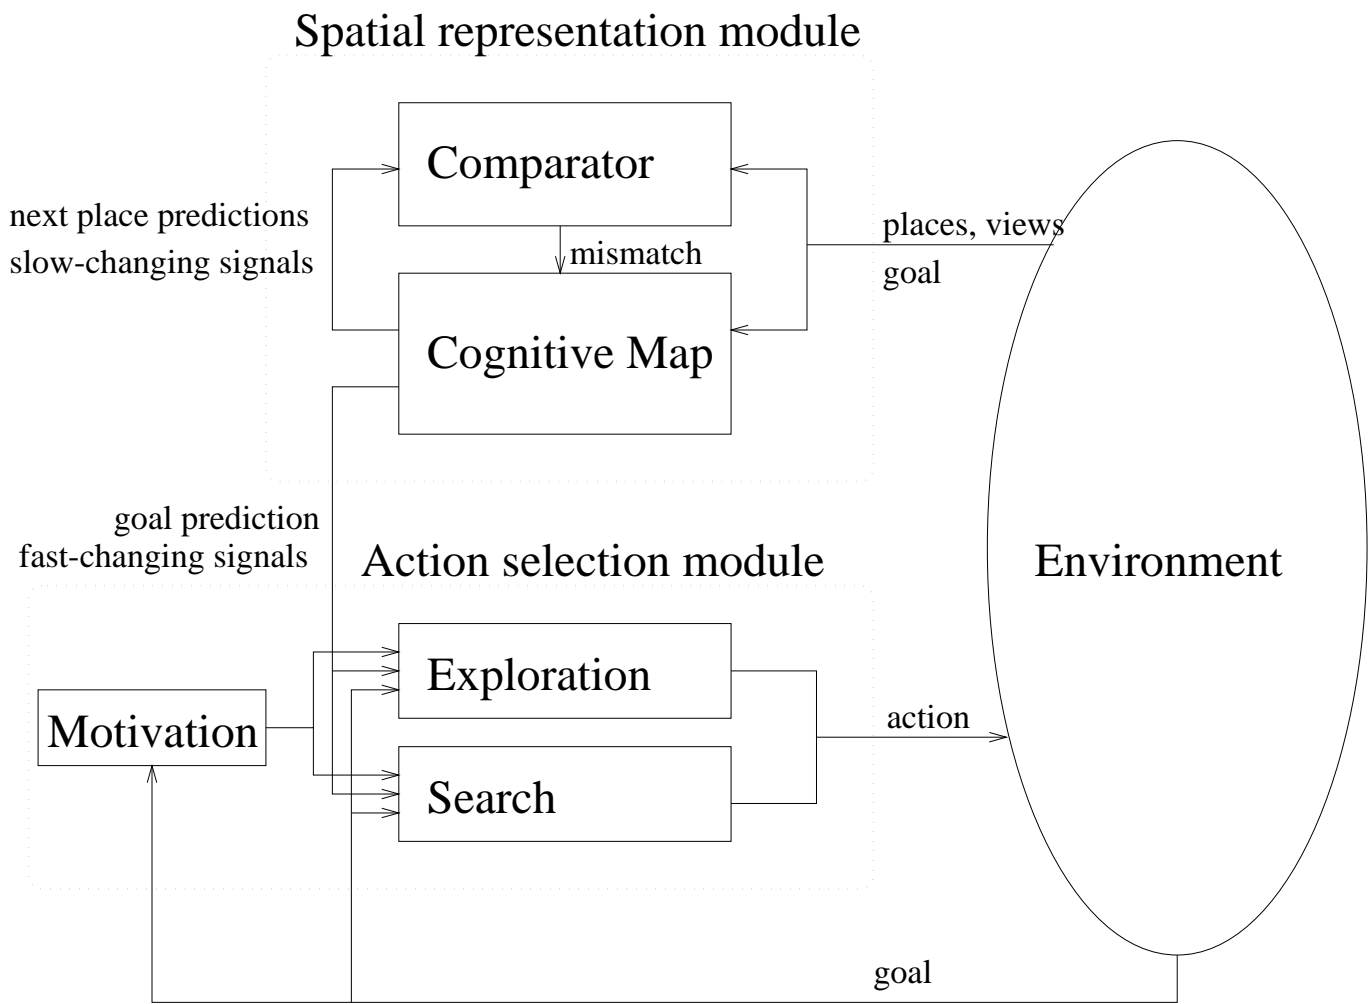


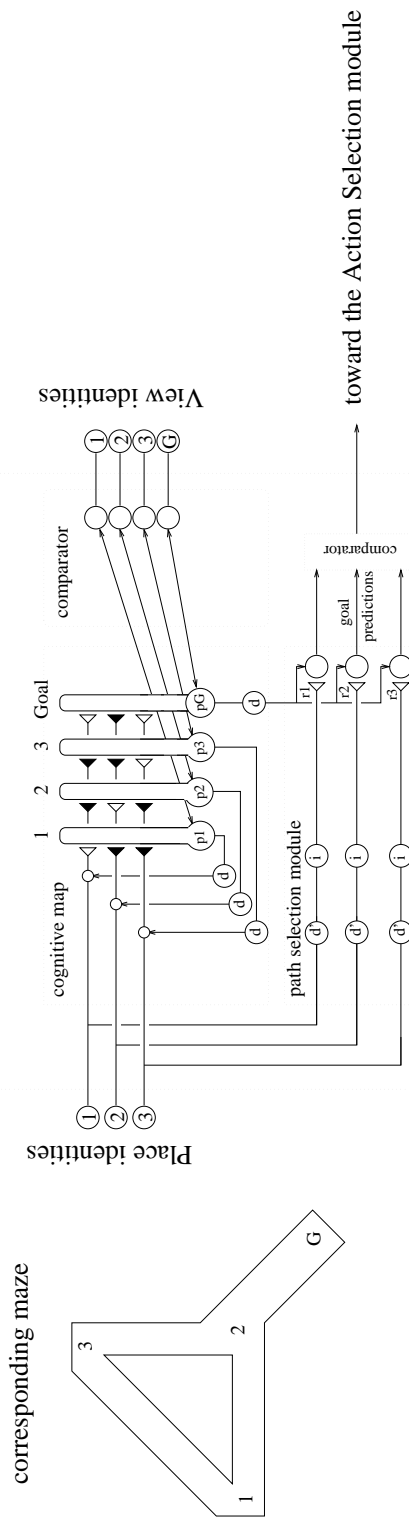






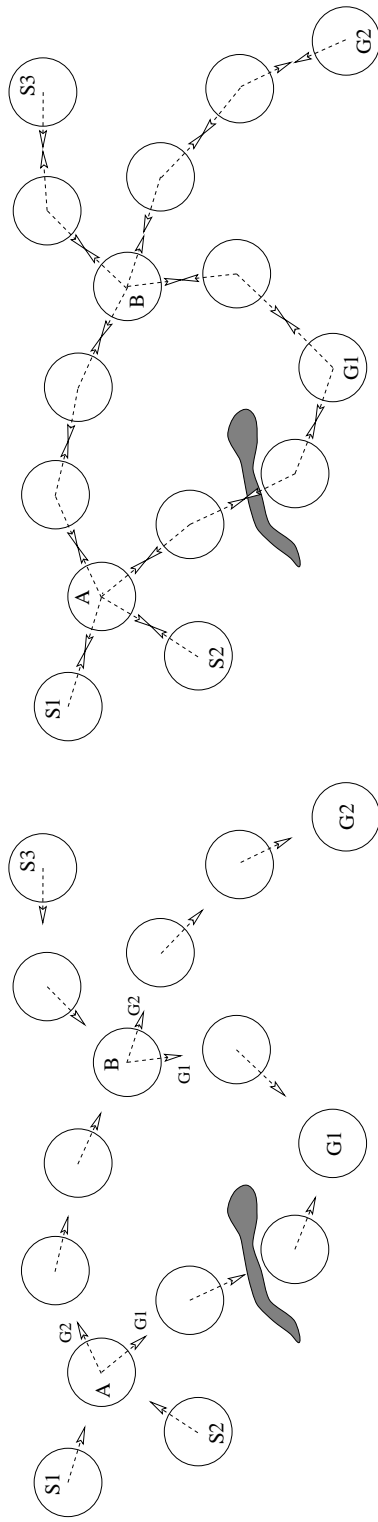






Spatial representation module

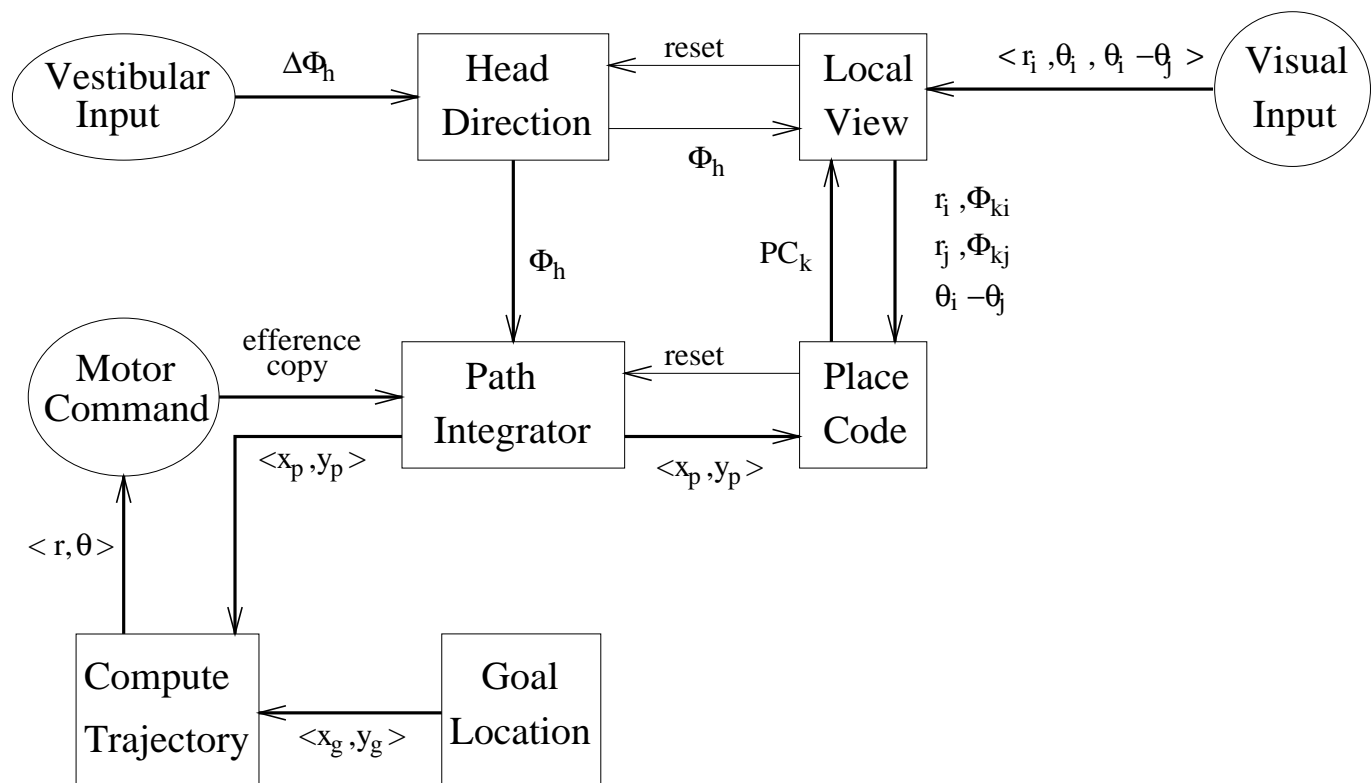
- ◀ high synaptic weight coding for the adjacency relation from one place to another.
(◀ modifiable synaptic weight)
- Ⓟ unit outputting the derivative of its input
- Ⓢ leaky integrator
- Ⓟ place-prediction nodes

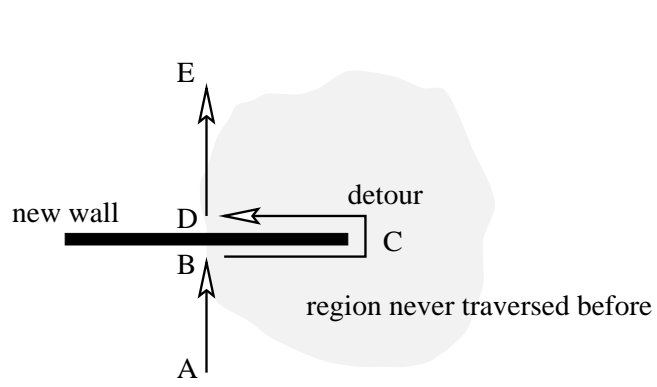


(a) collection of routes to goals G1 and G2

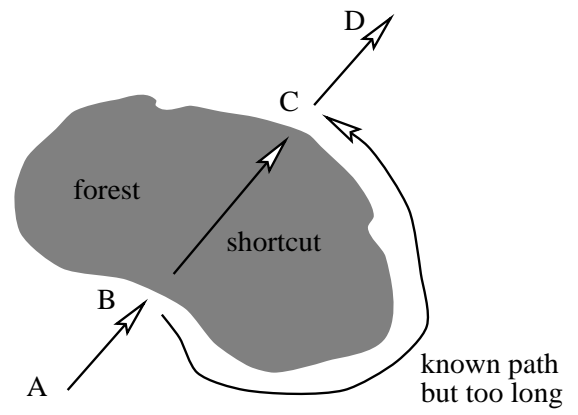
(b) topological representation derived from the routes

- place
- direction of movement to reach the corresponding goal
- direction of travel between two "adjacent" places
- new obstacle





(a)



(b)