Biol. Cybern. 86, 15–28 (2002)

Biological
Cybernetics

From view cells and place cells to cognitive map learning: processing stages of the hippocampal system

P. Gaussier[1], **A. Revel**[1], **J. P. Banquet**[2], **V. Babeau**[1]

[1] Neuro-cybernetic team, Image and Signal Processing Laboratory (ETIS), Cergy Pontoise University and ENSEA, 6 av du Ponceau, 95014 Cergy Pontoise, France
[2] Neurosciences and Modelisation Institute, CREARE, Jussieu, Paris, France

Received: 28 June 1999 / Accepted in revised form: 26 April 2001

Abstract. The goal of this paper is to propose a model of the hippocampal system that reconciles the presence of neurons that look like "place cells" with the implication of the hippocampus (Hs) in other cognitive tasks (e.g., complex conditioning acquisition and memory tasks). In the proposed model, "place cells" or "view cells" are learned in the perirhinal and entorhinal cortex. The role of the Hs is not fundamentally dedicated to navigation or map building, the Hs is used to learn, store, and predict transitions between multimodal states. This transition prediction mechanism could be important for novelty detection but, above all, it is crucial to merge planning and sensory–motor functions in a single and coherent system. A neural architecture embedding this model has been successfully tested on an autonomous robot, during navigation and planning in an open environment.

1 Introduction

Since the initial observations of severe anterograde amnesia following medial temporal lobe resection (Scoville and Milner 1957), the hippocampus (Hs) has been known as a very important structure for human and primate memory. The main structures within the medial temporal lobe are the hippocampal region (the hippocampal field, the dentate gyrus, and the subiculum) and the adjacent entorhinal, perirhinal, and parahippocampal cortices (Zola-Morgan et al. 1989; Suzuki et al. 1993). Electrical recording of neurons activity in rodent Hs have also shown the existence of "place cells" that fire for specific places in the environment (O'Keefe and Nadel 1978; O'Keefe 1991). O'Keefe and Nadel proposed that the Hs works like a Cartesian map of the environment, with each place cell coding for a very limited and well defined area (O'Keefe and Nadel 1978). Many experiments over the following two decades have confirmed the importance of the rat Hs in navigation tasks, but the more precise experiments have also shown a wide variety of sometimes contradictory results. First, less selective and more noisy place cells have been found in the entorhinal cortex (EC) just before the Hs. Second, some spatial tasks are not impaired by hippocampal lesions (Whishaw et al. 1995; Murray and Mishkin 1998). Third, the cell place fields seem to be less precise during passive navigation, depending upon the task context. All these points suggest that the idea of the Hs as a Cartesian map must be revised, and that the Hs contributes to diverse tasks.

Moreover, other studies concerning the understanding of conditioning mechanisms have shown that the Hs was also activated during the acquisition of classical conditioning (the rabbit's nictitating membrane conditioning, for instance), and that lesions of the Hs disrupt acquisition of long-latency conditioned responses (Berger and Thompson 1978; Solomon et al. 1986; Thompson 1990). In rats, the Hs has also been shown to be implied in encoding some learning aspects of instrumental conditioning (Han et al. 1995; Baxter et al. 1997). Lesions to the dorsal Hs selectively impair the ability of rats to represent the causal relationship between an action and its consequence (Corbit and Balleine 2000). More puzzling is the fact that hippocampectomized rats are modestly but significantly impaired only for long delays in nonmatching to sample tasks (Dudchenko et al. 2000). More generally, Hs lesions seem to impair nonspatial relational tasks (Bunsey and Eichenbaum 1996) as well as nonspatial conditioning tasks involving long temporal delays. But the Hs is only one of the many circuits for the elaboration of conditioned responses (Berger and Thompson 1978; Thompson 1990). The Hs may be critical in situations in which the relevant stimuli do not occur contiguously in time (Thompson 1990). These works on the role of Hs in memory processes have also enlighted the importance of the memory capacities of the perirhinal cortex (Pr) and parahippocampus (Ph) (Aggleton and Brown 1999), which are directly connected to the EC. Therefore, a strict opposition between rodent Hs solely devoted to a

---

Correspondence to: P. Gaussier

spatial representation of the environment (O'Keefe and Nadel 1978; Burgess et al. 1994; Wan et al. 1994; Arleo and Gerstner 2000) and primate Hs mostly involved in declarative memory (Milner et al. 1968; Reiss and Taylor 1991; Grossberg 1996) is not tenable anymore. There is clearly a need to overcome the limitations of these models and to deal with the complexity of the whole problem. Indeed, if place-cell-like activities can be explained as a side effect of a memory system, it is also possible that nonspatial task impairments are a consequence of an underestimated spatial component of the behavior.

In Sect. 2, we propose a model in which the Ph is used to build and store a robust and invariant representation of the visual environment. Prehippocampal "visual place cells" or "directional place cells" appear at the level of the EC as the result of the recognition of particular Pr–Ph configurations merging, for example "what" and "where" information: the Ph is used to store experienced configurations while the place recognition is performed in the EC and the dentate gyrus (DG). The robotic experiments in a simple open area will show that visual information are sufficient to build neurons with a very large place field that is useful to return to the vicinity of their associated location without the need to learn a map of the environment.

In Sect. 3, the outputs of the prehippocampal place cells modeled in Sect. 2 are used for planning according to a "cognitive map". Considering robotics and biological constraints in complex action-selection problems, we will show that it is impossible to associate a node representing a place with the action to be performed in the context of a multigoal problem using a single map. We will propose that the Hs learns the transitions between states defined by the merging of multimodal information.

In Sect. 4, we detail a model of hippocampal prediction cells in which the Hs participates in the selection of a transition state allowing to choose the correct action during planning. More precisely, we suppose a sharp place recognition with strong competition and a short-term memory storage can be performed by the DG from the large and/or directional place cells of the EC (see Banquet et al. (1997) and Banquet (1998) for more details). Next, CA3 pyramidal cells can learn to predict the next places that can be reached according to the recognition of the current location (DG activity). Hence, CA3 pyramidal cells express transitions between places and CA1 pyramidal cells try to select the most appropriate transition according to a particular motivation (information coming from the prefrontal cortex through EC medial connections, for instance). The selected transition can easily be associated with an action by means of the nucleus accumbens or by a back projection to the posterior parietal cortex (PPC) via the EC or the subiculum and the EC. Moreover, the selected transition can also be used in the subiculum to update the memory of the previous head direction. The presence – in the Hs – of neurons that react when the animal is at a particular place does not contradict this model. Place-cell activity results of a snapshot of a system dealing with transitions between multimodal states, related to navigation since the animal is performing navigation task. The plausibility of the model will be discussed in Sect. 5. The two parts of the model have been simulated on a real mobile robot using a two-dimensional CCD camera [see Gaussier et al. (1999a) for a complete description of the visual place recognition system].

2 Prehippocampal place cells

The nature of the visual information used by animals for place recognition or homing behavior is at the center of an old debate. Two main theories are proposed. In the first one, the panorama is supposed to be recognized as a whole and matched directly to some memorized panorama (model of insect navigation, Wehner et al. (1996)). This theory is opposed to a second one in which panoramas should be decomposed into objects and relations between objects matched to a learned model (model of human vision, for example Oram and Perret (1994)). These two points of view can be reconciled if we consider that several simple low-level features can be extracted in parallel, and that a sequential scene analysis is also important for high-level discrimination. For instance, it appears that the whole field luminance is not used much by normal rats (nor primates) (Goetsch and Issac 1983; Kolb et al. 1983). Rats must be able to use some kind of landmark identification but "it is not clear, however, how these distant cues are processed. They might be fixated sequentially or processed in peripheral vision..." (Kolb and Tees 1990, p 293). Indeed there are also some suggestions that the rat temporal cortex may be involved in peripheral and not central vision (Kolb and Tees 1990). But there exists in the rodent lateral posterior cortex an area called Te2 (Krieg's area 36, Zilles's area Te2) that has some resemblance to the inferotemporal cortex of the monkey. It projects to the perirhinal region (Deacon et al. 1983) and receives projections from the EC (Kosel et al. 1982). In the monkey, the perirhinal area receives a large input from higher visual areas, while the Ph receives larger input from the parietal and cingulate cortices (Suzuki et al. 1993). A similar topography of inputs distinguishes the perirhinal and the postrhinal areas of the rat (Deacon et al. 1983) (equivalent to primate Ph; in rats, some of the memory-processing properties of cells in Te1 (anterior inferotemporal cortex) area are very similar to the features of cells in the parahippocampal region, suggesting that at least part of this area may be more reasonably considered as an adjunct to the parahippocampal region, with its role limited to visual memory processing (Eichenbaum et al. 1994)). A major similarity between perirhinal and parahippocampal (postrhinal) areas is that both project heavily to parts of the EC. These structures are a major intermediary in communication between the hippocampal formation and the neocortex.

Lesions to the PPC involve a deficit in perceiving the spatial relations between objects, such that the animals do not orient themselves accurately with respect to spatial information. This deficit is reminiscent of similar deficits in humans and nonhuman primates (Kolb and Tees 1990). The PPC deficit is likely to be dissociable from the frontal or hippocampal deficits in that there is no additional memory deficit after PPC lesions. From the anatomical data, it is not possible to say that complex object recognition cannot be used for navigation, but the ablation of the Te2 region does not perturbate navigation tasks like those in the Morris swimming pool. On the contrary, it suppresses the capability of associating the recognition of a complex object with an action, like jumping in the direction of the object or avoiding it. It is thus possible to conclude that there are in rodents – like in monkeys – two distinct visual processing routes: one for spatial guidance and one for the kind of visual analysis required for object recognition (Mishkin et al. 1983; Kolb and Tees 1990). Obviously, all these features can be much more complex in the primate or human brain, with the possibility of more cognitive function in the PPC (e.g., spatial rotation and map visualization). In conclusion, animals such as rodents and primates can use – at the same time – more or less complex (or integrated) visual information. We will consider that visual place recognition is the result of a more or less complex merging of visual information including "object" recognition ("what" pathway in mammals) and "object" location ("where" pathway). Obviously the notion of "object" may be supposed to be more complex for primates than for rodents, but our model suggests that it may not be fundamentally different (Lashley (1938) for neurobiological and behavioral justifications). In our model, low-level visual treatments are simply "object/feature" detectors using features such as gradient information to detect a transition between two areas. The useful information is the azimuth of the "object/feature" according to a given absolute direction. The proposed model considers the parietal and temporal cortical areas as inputs to the Pr and Ph. In fact, the Pr or the Ph could be the place where local configurations of different kinds of information (like "what" and "where" information; Mishkin et al. (1983)) are merged and memorized for the time of the visual scene exploration. The merging of the "what" and "where" information requires some kind of units calculating the product between the recognition level of a given landmark and information about its angular displacement from the learned configuration (sigma – pi units, for example Rumelhart and Zipser (1985)). The same holds if we suppose that other contextual information can be used – they will have a modulation effect on more spatial information.

Next, Pr–Ph areas are connected to and from the EC, which is itself the input of the Hs proper. The EC is also connected to other cortical structures such as the prefrontal cortex, which is known as a very important component of planning capabilities (Goldman-Rakic et al. 1984). Hence, the EC seems to have an important role in the filtering and integration processing of the hippocampal inputs and outputs (Jones 1993; Rolls and Treves 1998), and could correspond to a first level of place cells.

In the present model, the activity of the "place cells" is computed by the neural network proposed in Fig. 1. The activity of a Pr–Ph neuron is the product of the recognition level of a given feature (distant landmark, for instance) and its angular distance to the learned azimuth. The activity of a place cell is a normalized sum of Pr–Ph cells which were activated during place learning [see Gaussier et al. (1999a) for more details]. Our merging and place recognition models have been implemented on a mobile robot in order to test their efficiency. This system allows to build a panoramic image up to 300° wide (1200 × 288 pixels).

The interest of these experiments for our understanding of animal visual navigation lies in the opportunity to record the response of each neuron that codes for a particular place (we may consider them as artificial "place cells" – see below) in a completely controlled situation. Figure 3 shows the activity of 25 neurons associated with 25 different places according to the recognition level of the landmarks and the information about their azimuths. Learning was supervised with each neuron associated with a particular location in a 5 × 5 grid of the room used for the experiments (see plan, Fig. 2). The different level lines show that a neuron may be activated even if the current view is far from the learned one, and that there are no local minima. Figure 4 shows what happens after a strict competition. For a given location, only the winner neuron remains activated (it keeps its previous value while the others are reset to 0). This activity looks very much like the shape of place-cell activity recorded in the rat hippocampus (O'Keefe and Nadel 1978) (we will come back to this kind of comparison in Sect. 5). The results show the model allows segregation of the place locations in a real environment from solely visual information.

The most interesting result is that the gradient of the place recognition level can be used for homing behavior very far away from the learned location – more than 2–3 m in our case – without local minima problems. In previous works, we presented a simple conditioning network which allowed our robot to return to a given location without the need to learn a map of the envi-

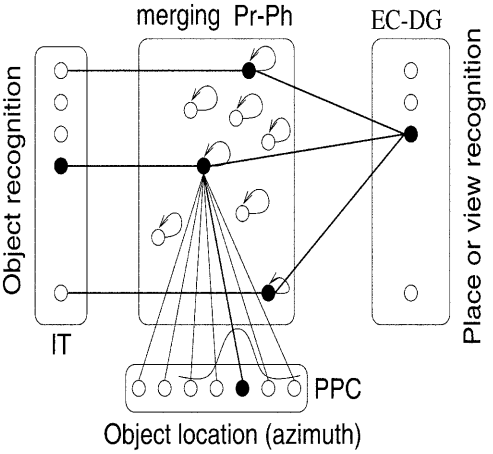

Fig. 1. Merging of "what" and "where" information for place recognition in the case of an high-level visual system. The lateral diffusion allows measured of the difference between the learned azimuth and the current azimuth

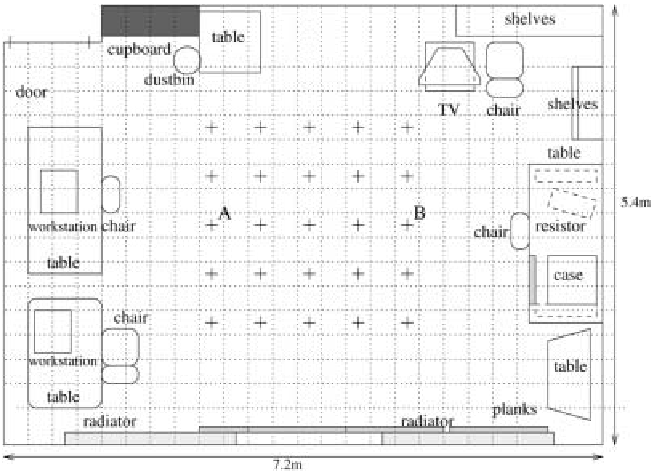

Fig. 2. The room in which the experiments are performed. The *crosses* represent the places where the robot has learned (see Figs. 3 and 4)

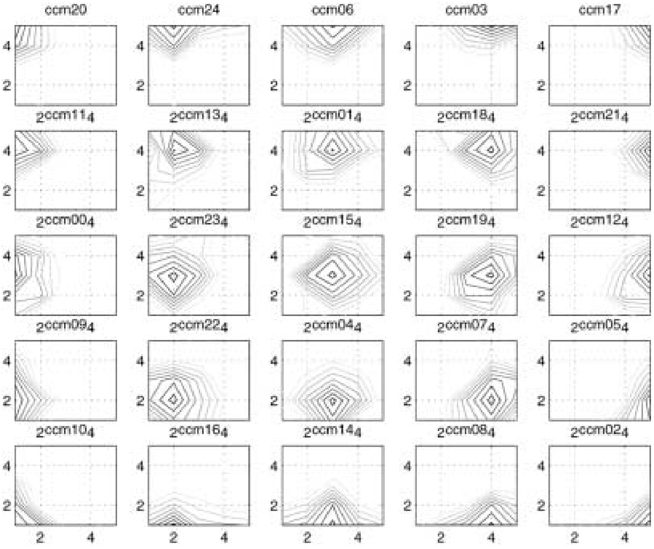

Fig. 3. Twenty five neurons, 25 places, 5 measures averaged per place, and a diffusion of 35°. Neurons are considered to be isolated (activation only comes from the direct input). Learning was supervised: each neuron is associated with a particular location in a 5 × 5 paving of the room in which the experiments have been carried out. Each *rectangle* is a map of the room. The *curves* show the activity of the neurons corresponding to the crosses of Fig. 2

ronment in the case when the starting position belongs to the same open visual environment as the goal. These results confirms the experimental results of Whishaw (Whishaw 1991; Whishaw et al. 1995; Whishaw and Jarrard 1996), showing that rats do not need a cognitive map to solve the Morris swimming pool.

A last robotics experiment can be useful for understanding some differences between rodent and monkey perception of their visual environment. In this new experiment, the visual field of the system is limited to 180°. Then, the activity of the neurons associated with the recognition of the "what and where" representation are no longer place cells (see Fig. 5), but have very similar properties to the "view cells" recorded by Rolls (Rolls and O'Mara 1995; Rolls and Treves 1998) in the monkey Hs when the monkey perceives a particular view. The cell responses are relatively invariant of the distance and location. This experiment shows something that was already true in the previous experiment but was more difficult to see: our place cells are directional and their directionality depends on the field of view. For large fields of view (about 300°) and for a given location, the activation of a place cell decreases a little when the robot orientation is at 180° to the orientation used for learning (if the landmarks are uniformly distributed). For smaller visual fields, it happens more easily that landmarks disappear from the field of view which reduces the cell activity more.

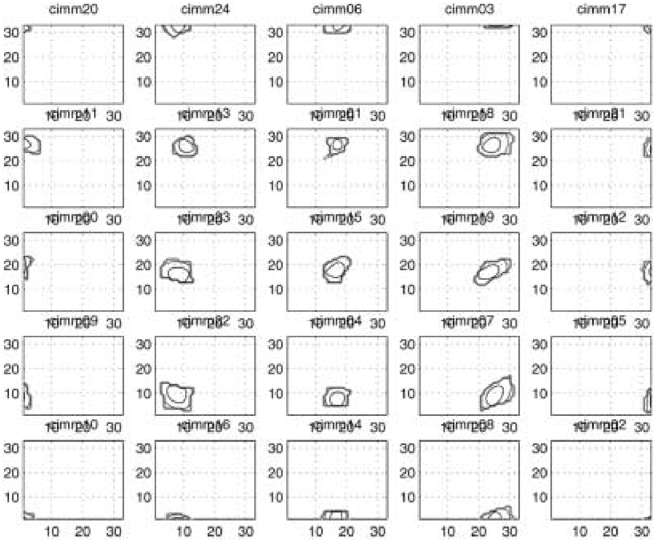

Fig. 4. Same setup as in Fig. 3. A competition between neurons is introduced. Only the winner neuron remains activated while the others are set to 0 (no activity enhancement used)

But it remains striking that the place cells found in the rat Hs are restricted to an area of about 10–20 cm, while we insist on the possible availability of a continuous place recognition information that can be used over several meters (at least 2 m – 10 times the diameter of the observed place fields). In a simple open environment, the sole use of the discrete place recognition performed in the Hs means that the animal is only able to recognize its actual location and cannot have any idea about its distance to a particular location like a goal, or the platform in the case of the Morris swimming pool. As a result, models of rat navigation based on hippocampal place cell always involve internal maps (Burgess et al. 1994; Trullier et al. 1997; Arleo and Gerstner 2000). Unfortunately to come back to a goal location, an animat using these models will have to already know a path linking the current location with the goal location. A priori generalization is then not possible, and if mammals were all using only this kind of strategy they would have less navigation insight than insects (no latent learning for instance). On the contrary, using a perception–action architecture is sufficient to learn a minimum of three place-cell/action associations in the immediate neigh-

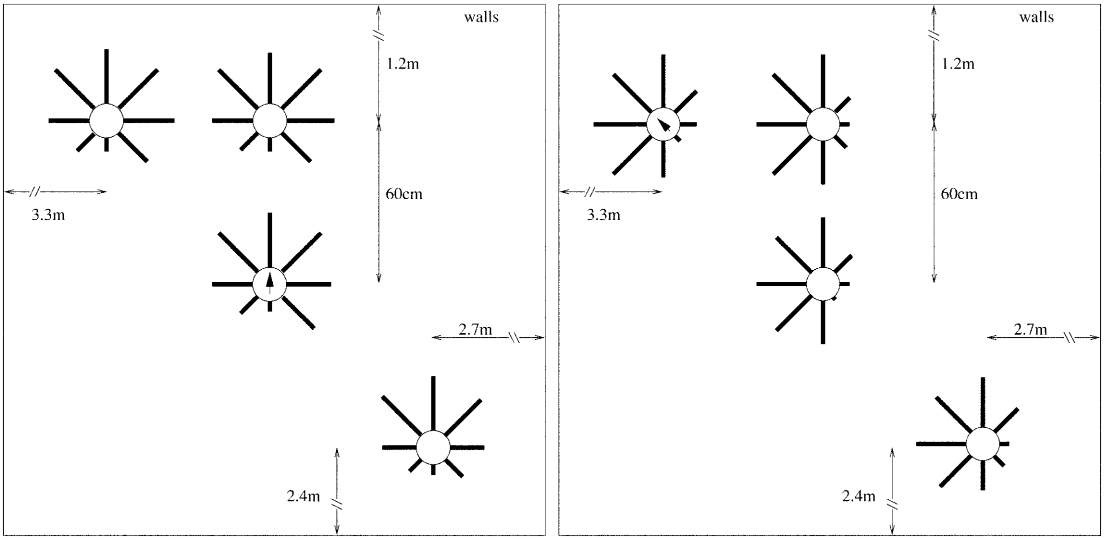

Fig. 5. Record of two "view cells" according to the robot location and orientation in our experiment room (*arrow* represents the position and direction of the learned view). The *black bars* represent the neuron or "view cell" activation in the direction of view. The robot uses the same algorithm as for place recognition. The robot field of view is simply reduced from 300 to 180°

borhood of a goal, allowing the robot to come back to the goal from any location in the open space around the goal (even if the robot never experienced this location before – see Gaussier and Zrehen (1995)).

3 Using cognitive maps

For more complex tasks, each set of place/action association can be linked with a basin of attraction, and is thus potentially reachable (provided it is recognized, and that there is an open space between the place where the animat actually is and the one it wants to reach – see Gaussier et al. (1999a). Therefore, if some of those places are associated with specific motivations ("hunger" or "thirst" for instance – see Fig. 6), the modulation of the recognition by the motivation level allows the animat to reach the appropriate place. This simple action-selection mechanism still works even if there are obstacles on the trajectory, provided a reflex obstacle avoidance system competing with the goal attraction mechanism (Gaussier et al. 1999a).

The main difficulties arise when the animat cannot perceive directly where the goal is, for example when there is a wall or a huge object. As the animat is not in the goal visual environment, it cannot generalize the "good" movement to perform from what it has learned next to the goal. An intermediate step consisting of learning how to reach places from where it is possible to see the goal is necessary. Behaviorists have proposed that it could be performed by learning what they call "conditioning chains": they suggest that from place to place, the movements leading to the goal could be reinforced (it is interesting to notice that this mechanism is similar to Q-learning techniques and has been applied in hippocampal model (Arleo and Gerstner 2000)). This explanation can still stand even if, at given places, there are several possible routes (but a single motivation). For instance, if in Fig. 6 the animat finds a new goal G2 which is closer than G1, the route can be reinforced more efficiently since it is closer (shorter delay between reinforcement signals). Yet, in the case where G1 and G2 are activated simultaneously, the "conditioning chain" explanation does not work anymore. For instance, if G2 is associated with "thirst", links BC and BD will be reinforced (maybe link BD will be reinforced more since G2 is closer) and there would be no way to distinguish the action leading to G1 from the action leading to G2. One can think of using a modulation mechanism, but in this case there should be as many place cells or action coding as motivations.

In fact, the ultimate problem of the above model is that it does not explain latent learning. Indeed, in the

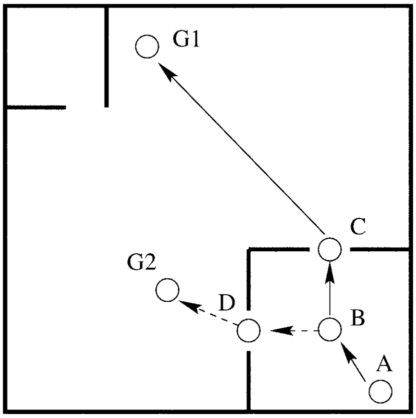

Fig. 6. Example of a complex environment showing the need for a planning mechanism. *Circles* (and the associated *letter*) represent learned places and *arrows* show the movement leading from one place to another. Two goals *G1* and *G2* associated with two different motivations are also present. Conditioning alone cannot explain the capacity of the animat to choose between two possible movements in *B* (the *dotted arrows* represent the alternative movements): a planning or anticipation mechanism must be added

1930s, Tolman performed a maze experiment with rats, which consisted of comparing the results of rats that were always rewarded with another group of rats that was only rewarded after 11 days (Tolman 1948). According to the behaviorism paradigm, rats of the first group should have had much better results than rats of the second group. Yet, after the 11th day, the rats of the second group had results equivalent to those of the first group. Thus, it seems that although they were not rewarded before the 11th day, rats have learned (“latent learning”) an internal representation of the maze (called a “cognitive map”). The definition of the “cognitive map” is controversial. In particular, some consider that this kind of map must be in Cartesian coordinates (Gallistel 1993). Schmajuck proposed a definition which is, in our opinion, more appropriate (Schmajuk and Thieme 1992):“… The cognitive map … is a topological map, i.e., it represents only the adjacency, but not distances or directions, between places… The cognitive map allows the combination of information about spatially adjacent places, and thereby the inference of the connections to remote places.”

3.1 Building a cognitive map

Let us come back to our “place cell” model: how is it possible to *learn* a “cognitive map” starting from our implementation of the recognition of places? In fact, this is rather easy if we consider that neurons of the place recognition map are fully interconnected. Then, when a place is recognized, it is coded on a neuron i. When a new place appears it is coded on j, and a simple Hebbian rule allows the learning of the time relationship (and thus the topological relationship) between these two situations. If x_i is the activity of neuron i and $\bar{x}_i$ its corresponding short-term memorization (time constant τ), $\bar{x}_i$ can be computed as follow:

$$\tau \cdot \frac{\mathrm{d}\bar{x}_i(t)}{\mathrm{d}t} = -\bar{x}_i(t) + x_i(t) \tag{1}$$

The Hebbian learning rule is given by:

$$\frac{\mathrm{d}W_{ij}}{\mathrm{d}t} = -\lambda W_{ij} + \left(C + \frac{\mathrm{d}R}{\mathrm{d}t}\right) \cdot (1 - W_{ij}) \cdot \bar{x}_i \cdot x_j \tag{2}$$

where C is a positive constant (speed of the associative learning) and λ is the decay factor. The weight modification also depends on the variation of the reinforcement signal R. If the reinforcement decreases while the robot is moving from one place to the other, then the associated link will be less reinforced than in the normal case. $(1 - W_{ij})$ is a saturation term to ensure $W_{ij} \leq 1$. By generalizing the mechanism, a graph of the spatial relationships between places is constructed.

Planning also requires learning of the link between the recognition of a given place and the satisfaction of a drive (that will be used as a goal). For this purpose, it must be considered that the recognition of this place is associated with the activation of a given “motivational” neuron. The links between the motivational neuron and the recognition of the situation can be reinforced using a Hebbian rule. Then, a solution for the planning process can be the back-propagation of the motivation activity within the map (as for resistive grids – Bugmann et al. (1995)). For a proper functioning of the planning process, the activity of a neuron in the cognitive map (x_i) must be a function of its topological distance (number of intermediate places, for example) to the goal. We propose:

$$x_i = \max_j \left(W_{ij} \cdot x_j\right) \tag{3}$$

We assume that neuron activity is bounded by 1 and that the values of the weights are bounded by a value $W_{\max}$ which is lower than 1 (see above). Before stabilization of the algorithm, there must be several iterations (if spiking neurons are used, the max operator can be replaced by a simple turn-off mechanism such as proposed by (Thorpes et al. 1996)). The two phases of planning algorithm are:

1. Initialization: i_0 is activated by a drive, $x_{i_0}(0) \leftarrow 1$ and $x_i(0) \leftarrow 0, \forall\, i \neq i_0$.
2. Back-propagation of the drives: Do $\forall\, i$, $x_i(t+1) \leftarrow \max\left(W_{ij} \cdot x_j(t)\right)$ while the net activity is not stable.

The use of the cognitive map could be the following: after back-propagating the motivations from the goals to all the subgoals or known places, the system tries to recognize its current location (the most activated “place cell” is the nearest learned place from the current robot location). Then, it selects as a subgoal the most activated node on the “cognitive map” directly linked with its current location. To reach that place, a gradient-following technique can be used (moving in the direction that maximizes the activity of the place-cell coding for the subgoal that the system wants to reach). When the system enters in the vicinity of the next subgoal, the process is repeated (selection of the next subgoal) until it reaches the nearest goal relative to its starting location. This neural algorithm is formally equivalent to Bellman’s shortest-graph-distance algorithm (Bellman 1958), and it has been shown to be very efficient in simulation to solve complex action-selection problems including opportunistic choices and contradictory motivations (Gaussier et al. 1999b; Quoy et al. 1999).

If the place-recognition neurons were directly connected to each other with recurrent connections so they can be also used as a cognitive map, the meaning of the neuron activity would become ambiguous: the neuron activity would represent a mixing between the distance to the associated learned places and the distance to the goal (the feedforward and backward computations would be mixed together).

A solution consists of separating the process into two functional levels so as to separate the two information flows: one corresponding to the “goal” level, another corresponding to the recognition of the current place. During learning, the information must flow “bottom-up” from the recognition level ($P_k(x, y)$ activity) to the “goal” level, so as to allow cognitive map learning (one-to-one connection between the neurons of the two levels). Conversely, during planning, the information

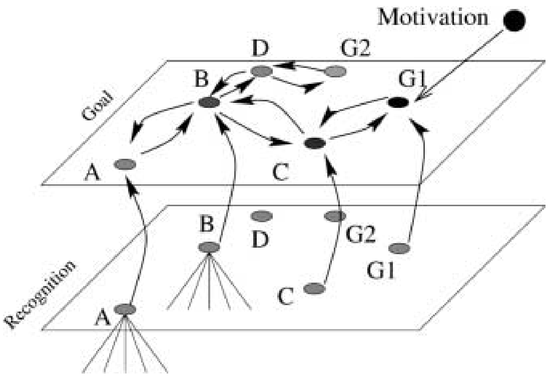

Fig. 7. Building of a "cognitive map". During learning, the information goes from recognition to "goal" level to allow the learning of the map (topology learning between *A*, *B*, *C*, *G1*, etc). During planning, the activation of a motivation is back-propagated on the map

coming from the "goal" level must go "top-down" in order to select the action to perform (see Fig. 7). Interestingly, this functional organization corresponds to the neurobiological model of cortical columns proposed by Burnod (1989).

3.2 Transition cells and action selection

We have not yet considered action learning and selection aspects. Animal studies have shown that rats are able – at "strategic locations" (T-junctions of a maze for instance) – to try for a short distance the different possible ways, and come back to the junction, and finally choose the more appropriate action to reach the goal (Tolman 1932). This mechanism, known as "vicarious trial and error", can be understood as a way of estimating the local gradient on the cognitive map (Schmajuk and Thieme 1992; Scholkopf and Mallot 1995). But these systems have to rely on an external mechanism to analyze the cognitive map activity and to decide the correct movement. In fact, in these systems the action only results from a gradient descent and cannot be controlled willingly. To overcome this problem and the limitations of gradient-descend techniques (need for vicarious exploration) it is important the system can learn to select a particular movement in a given situation. Our problem is then about the "embodiment" of the cognitive map in the sensory–motor system; i.e., the problem of connecting it correctly to the "place cells" and to the movement-selection mechanism.

A first idea could be to decide that each place-recognition cell has to be linked with the movement to the next place. Yet, during planning, the system must use the information coming from the goal level to perform the action sequence allowing reaching this goal. If one place is linked to several actions, the correct action cannot be selected. For instance, if the animat starts from A on Fig. 6, when it arrives in B there is no way to choose which action must be performed (B is linked with two movements – see Fig. 8).

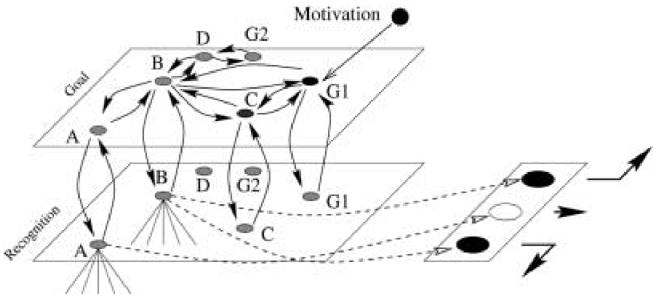

Fig. 8. Planning is impossible by using only steady states recognition. Indeed, a situation can be linked with two different movements, and it is thus impossible to decide which action must be selected

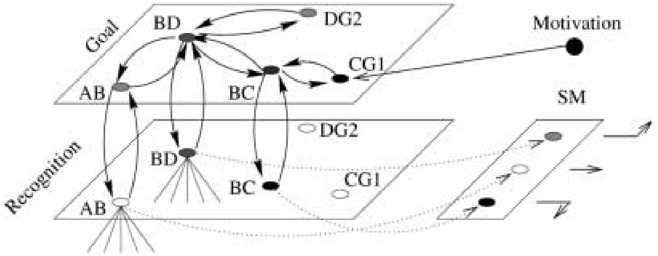

Fig. 9. Use of the transitions for actions planning. At a given place, there are several possible transitions which are likely to be recognized by the recognition system. Thanks to the back-propagation of the motivation activity, the recognition can be biased in order to propose the appropriate movement to reach the goal (impossible with steady states; see Fig. 8)

As there are action-selection problems due to the association of two movements with a single place recognition, the solution is to build a representation which can disambiguate movement selection. Besides, this representation must take into account both the start and end points. A solution consists of building a representation of the *transition between two situations* (see Fig. 9). Let AB be the internal representation of the transition between A and B. The associated action is the movement allowing to go from A to B (learned using a conditioning rule). In situation B, for instance, the animat may turn left (arrives in C) or turn right (arrives in D). Transitions BC and BD will be created and respectively associated with "turn left" and "turn right" movements.

During exploration, the "recognition" level creates an internal representation for each transition between places and associates the movement from one to the other. At the "goal" level, connections between representations are learned so as to create a graph of the topological relationships between the transitions. When the animat reaches the goal, it learns the association between the recognition of the last transition and the motivation satisfaction. During planning, the motivation back-propagation towards the current state allows the activation of the graph nodes, thereby indicating the movement that is necessary to perform in order to reach the goal. The planning mode requires (at variance with the exploration mode) deciding what movement must be performed to reach the goal. Therefore, the system must allow a prediction of the place(s) which can be reached

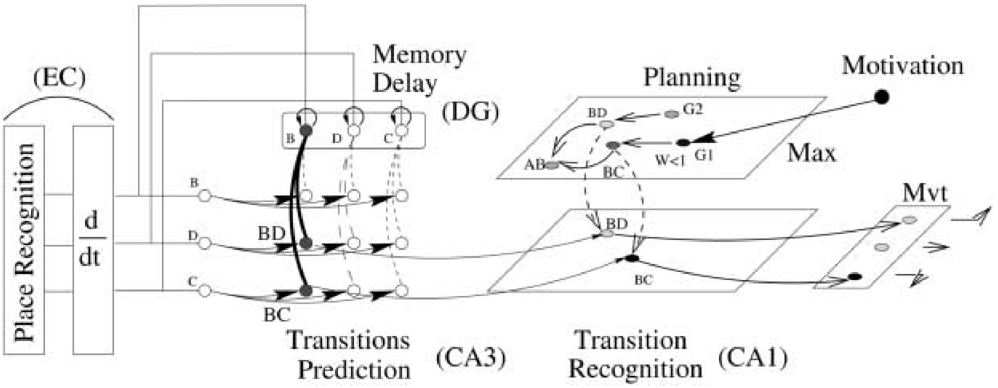

Fig. 10. Overview of the planning architecture using transition detection: the recognized place and the past recognized place trigger the learning of one neuron in the transition prediction group (*CA3*). When only the memory of the previous place (dentate gyrus *DG*) is available, the same group (*CA3*) predicts all the possible transitions. Finally, these transitions are proposed to the transition recognition group (*CA1*) and the planning system can bias the recognition so as to perform the relevant movement

from the current location. The idea of this representation has been inspired by a neurobiological model of timing and temporal sequences learning in the Hs (Banquet et al. 1997; Banquet 1998; Gaussier et al. 1998b). As the internal representation is elaborated from transition recognition and not only from scene recognition, an internal representation is necessarily linked with only a single movement. The neural implementation (see Fig. 10) consists in a group of neurons which are, on the one hand, linked with the derivative of the current place recognition and, on the other hand, with the memory of the previous input. The activity of transition–prediction neurons is simply obtained by adding the activity of each input and then thresholding the result. The initial weights and the threshold must be chosen so that the activity of a single input cannot activate the neuron, whereas both can. The prediction of transitions can easily be achieved assuming that the link between the delayed input and the map can be reinforced enough (by Hebbian learning, for instance), in order to be able to activate the neuron when only the delayed input is presented (simulation results can be found in Revel et al. (1998)). This model has a lot of points in common with the model developed by Grossberg and Merril to explain precise timing capability of the Hs in classical conditioning experiments (Schmajuk 1991; Grossberg 1996), and extends it to sequence learning and navigation tasks.

4 Hippocampal prediction cells

In the Hs, the information coming from the superficial layers of the EC projects to the granular cells of the DG as well as to the pyramidal cells of the CA3 region. The mossy fibers connecting DG granular cells to CA3 pyramidal cells can force the activation of the pyramidal cells and are known as a fundamental element of the Hs learning and memory capabilities (long-term potentiation of the mossy fibers has been widely studied (Derrick and Martinez 1996)). CA3 recurrent connections are supposed to be very important for pattern completion (Marr 1971), but since this property is not crucial for our problem, it will not be simulated here (Mizumori et al. 1989). CA3 pyramidal cells project to CA1 pyramidal cells which also receive connections from the median layer of the EC. In our model, we suppose the direct EC-to-CA1 connections provide activations that can be used to control CA1 cell learning and bias CA1 competition in order to select one transition according to the drive back-propagation on the cognitive map (from the parietal and prefrontal cortex).

CA1 pyramidal cells are themselves connected to the nucleus accumbens (belonging to the basal ganglia and presumably involved in action modulation) and to the subiculum in which head-direction cells have been found. The subiculum also receives direct information from the median layers of the EC (that could be also a relay from the prefrontal information used to select the most interesting transition according to a given motivation) which back-projects information to the deep layers of the EC (for more details of Hs anatomy, see Amaral and Witter (1989)).

It has been shown that rats with medial prefrontal lesions have impaired performances for a variety of delay-type tests including a delayed response (Kolb and Tees 1990). They have also difficulties performing spatial maze tasks (Becker et al. 1980; Kolb et al. 1983; Herrmann et al. 1985) Thus, it seems that the prefrontal cortex plays a major role in planning abilities. Our proposal is that the cognitive map could be elaborated in the prefrontal cortex but that the internal representations of transitions could be coded in CA3/CA1. Indeed there is neurobiological evidence that CA1 has extensive direct connections to the prefrontal cortex and indirect back-connections via the EC (which is consistent with our model (Ferino et al. 1987)). However, it has also be shown that rats with medial perfrontal lesions succeed in solving Morris swimming pool experiments. It does not contradict our model, since navigation in the Morris pool can be explained simply by a direct association between "prehippocampal place cells" and actions (Gaussier and Zrehen 1995; Gaussier et al. 1997, 1998a, 1999a). After hippocampal damage, monkeys and rats are also unim-

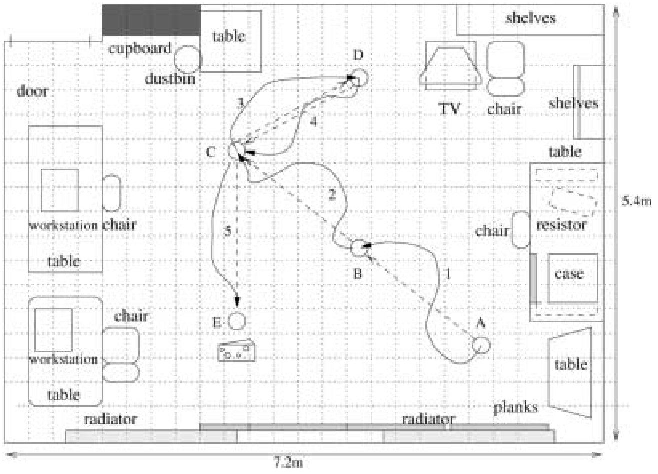

Fig. 11. The robot path during the exploration and the latent learning of the cognitive map. *Curved lines* indicate the real robot trajectory and *straight lines* represent learned transitions linked to the path integration between the two successively recognized or learned places

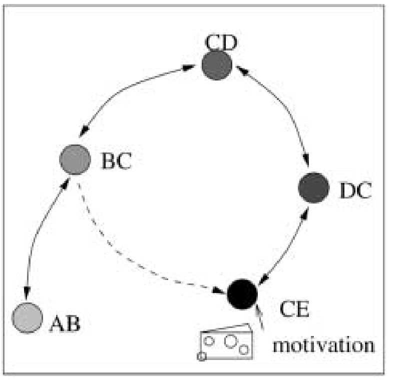

Fig. 12. Cognitive map built after the exploration of the environment shown Fig. 11. Note that there is no direct connection between the neuron coding the transition BC and the neuron coding the transition CE (because it was not experienced by the robot during the exploration). With the transition prediction system, the robot can nevertheless choose the transition CE after the transition BC because of its opportunistic behavior (recognition of C allows it to predict CD and CE)

paired when required to remember two locations over a 20 s delay (Mumby 1992; Murray and Mishkin 1998). Moreover, rats with fimbria–fornix lesions display a place response in a swimming pool with a dissociation between "*getting there*" and "*knowing where*" (Whishaw and Jarrard 1996). These results are confirmed by the experiments of Dudchenko et al. (2000) suggesting that hippocampectomized rats could recognize – but could not inhibit – their approach to previously visited locations. The capability of reaching the platform even after DG granular cell degeneration (even if the rats performances are impaired) could be explained by the existence of some kind of place cells in the EC (Armstrong et al. 1993) (DG granular cells degeneration induces an impairment in intramaze cue rotation). The impairment of the performances could result from the incapacity either to merge odometric information and/or to plan a short trajectory, because of an inability to predict the possible transitions from the current state. We have tested this architecture in the open environment detailed in Fig. 11.

In a first phase, the robot explores the environment at random and, at the same time, learns the places A, B, C, D, and E when the activation of the already learned places becomes smaller than a given threshold (vigilance level). The cognitive map is elaborated over time, creating a new node for each new transition and linking nodes as they are consecutively encountered. In the present experiment, the robot learns and links together the transitions AB, BC, CD, DC, and CE (see Fig. 12). The movement associated with a given transition is obtained after path integration of the robot movements from the previous activated place to the next one. The path integration is reset when the robot discovers a new place or arrives at a known place.

It can be noticed that a single path has been learned in this very simple experiment. Nevertheless, it appears to be a very interesting case to show the capability of opportunistic behavior of the algorithm. Indeed, according to the graph, there is no reason for the robot to go directly from C to E – it should go first to D.

After this short exploration of the room, the robot is brought back to the starting point and the motivational neuron is activated. Because of the learning, the motivation is back-propagated on the cognitive map. At the same time, the robot recognizes from the sole visual panorama its current location (the most activated place cell). This recognition is used to "predict" all the possible transitions from the current location. When several transitions are possible the goal back-propagation on the cognitive map allows biasing the competition between the predicted transitions (through the top-down connections). The best transition according to the planning system is then selected. When the robot starts from A, only the transition AB is possible so the robot moves in the direction AB until it recognizes B. Then, the transition BC is proposed and the robot moves in that direction. Later, when the robot arrives in C, it predicts from C the transition CD but also the transition CE. Because the planning node CE is activated more than the node CD (since it is nearer to the goal), the recognition of the transition CE wins and the robot triggers the movement in the direction CE (top-down bias of the planning system on the transition recognition system). Hence, it performs a "shortcut" in the graph. This new path (the link between BC and CE) is learned on the cognitive map after the robot arrives in E.

The current robot experiments deal with navigation between several rooms. In simple cases, like moving from room A to the corridor and selecting room B or C according to a particular motivation, the system works quite correctly. But a lot of work remains to be done: first, to select correctly the places to be learned (for instance to learn to pass correctly through a door) and next, to deal with visually ambiguous locations that appear in long corridors or in dark conditions (the place recognition has to take into account odometric infor-

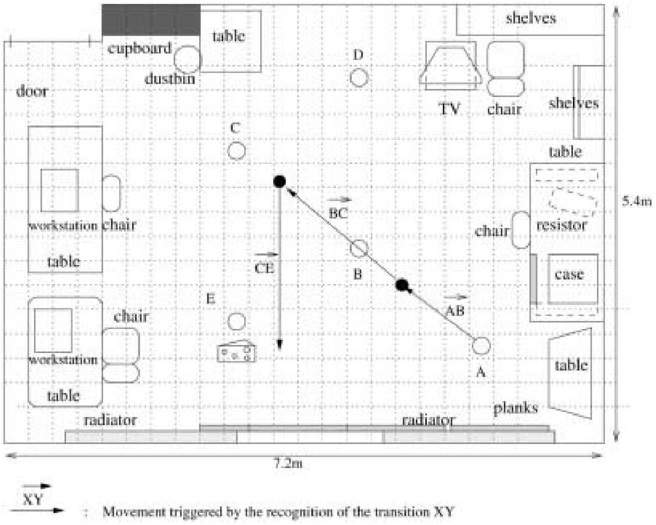

Fig. 13. Sequence of actions when the robot starts again from *A* and uses the diffusion of the motivation on the cognitive map to plan its actions

mation). Finding a low-cost way to introduce odometric information in the planning system will also be very important for enabling real short cuts such as going directly from A to E in the experiment presented Fig. 13.

5 Discussion

In this paper, we have proposed a model of the processing stages of the hippocampal system with a functional distinction between prehippocampal regions and the Hs proper.

First, we have shown the idea of "prehippocampal place cells" is supported by some theoretical considerations but also by electrophysiological recordings and a model based on anatomical data. The possibility that animals directly use the perceived panorama as a whole cannot be discarded, but the limitations of such a global comparison between stored and perceived panoramas prohibits its use in decision making when the distance to the learned location is large (due to a very low tolerance to object displacements, and variations in lighting and shadows and object scale). According to our model, "place cells" might react over very long distances and would not appear at first sight as place-specific because of their wide range of responsiveness. This is confirmed by experimental results showing that receptive fields of neurons in the EC are spatially selective but that these neurons are more active and noisy than CA neurons (Quirk et al. 1990) where the classical "place cells" are usually recorded. Our model is also consistent with the idea that rats with hippocampal damage are able to recognize previously learned locations but are unable to avoid their approach to them (Dudchenko et al. 2000). Moreover, our model also accounts for "view cells" found in the monkey Hs if the system field of view is reduced from 320 to 180°. Hence many data support the way that information is merged in our model to account for place or view recognition. Other information such as apparent size or landmark distance (Burgess et al. 1994) could also be used to improve the system performances. But, as pointed out by Zipser (1985) the sole distance information cannot be used to distinguish between mirror situations. Information about the relative locations of the objects from each other is very important. Moreover, distance information obtained from vision is less precise than azimuth information. One of the main arguments of the Burgess, Recce, and O'Keefe model in favor of the use of distance information for place recognition was accounting for the results showing that an expansion in a single direction of the experimental setup induces a splitting in two of the place fields (Brugess et al. 1994). Their explanation is that if a neuron takes distance information from two different walls and if the distance between the two walls increases, then the neuron will have two maxima of activation from the two locations associated with the learned distances. If place-cell activation is computed according to our model, it is not possible to account for the place-field splitting. If the distant landmarks are the corners of the box then one place field remains robustly associated with the same location even after an anisotropic dilatation of the box. The robustness is due to the way angles ("where" information) are measured. (Note that the place field also remains robust if a given landmark is taken as the reference to measure the angles. This kind of robustness of place fields in our model is interesting for robotics applications but disagrees with the biological data). If the model is modified so that the angles are measured from one landmark to the other then the robustness disappears and a splitting in 2 of the neuron place field can be observed (see Fig. 14).

Next, in the second part of the paper we have proposed that the Hs is mainly devoted to the learning and the prediction of transitions between steady states which are crucial when planning abilities are required (navigation in a maze-like environment, for instance). The Hs must not be thought of as devoted only to navigation. It rather seems to be a structure involved in a more generic function of integration of spatiotemporal information and transition detection. This view is consistent with models of the Hs as a novelty detector (Kohonen 1984; Knight 1996) – an error in the prediction can be detected and the input considered as a novelty. In the same way as in the model of Hasselmo et al. (1995), if an unpredicted event appears when CA3–CA1 neurons try to predict the possible futures, there is a mismatch on CA3 and CA1 that could be detected by septal neurons and thus modifies the cholinergic (ACh) modulation and express novelty detection (it also provides learning stabilization). These results are also consistent with the fact that lesions to the medial septum destroy both the theta rhythm and the ability of rats to solve a spatial task (Wilson 1978).

Figure 15 summarizes the second part of our model (the Hs as a prediction system). DG granular cells integrate information linked with the directional place cells of the EC to build place cells with no preferen-

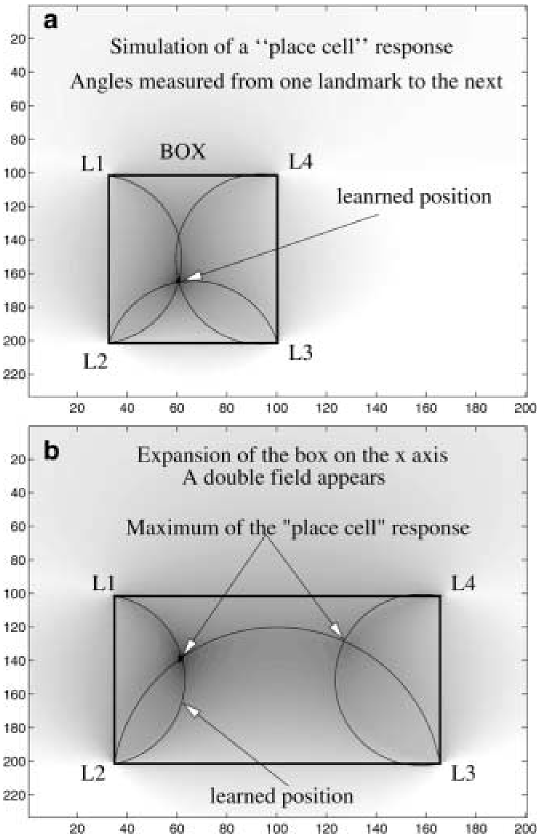

Fig. 14. **a** One place field before box dilatation when angles are measures from one landmark to the next. **b** Splitting in two of the place field after dilatation. Areas associated with higher activities are represented by *darker areas*. The *circle arcs* represent the location where there is no error on a given angle or neuron input

tial orientation (classical place cells). CA3 pyramidal cells build transition–prediction cells. The model prediction is that place fields associated with neurons in the DG are smaller than these associated with CA3 neurons. Indeed, in our model CA3 neurons react not only when the animat is at a particular location, say B, but also when it moves from B to C or from B to D (see Fig. 15). If transition selection really happens in CA1, it should also be possible to notice smaller place fields in CA1 than in CA3 (because of transition selection according to the goal). These properties seems to be confirmed by experimental results. The place/direction specificity of complex spike cells is significantly higher in CA1 than in CA3 [twice as selective on average – McNaughton et al. (1983)] and DG cells have a very limited place field (Jung and McNaughton 1993).

What is the role of the recurrent CA3 connections? Are they used as an auto-associative memory (for the completion of distributed patterns)? Or are they used as an hetero-associative memory that could predict a sequence of possible transitions? In the present model, both solutions are possible but it would be simpler to imagine the CA3 recurrent connections as an auto-associative memory. Otherwise, the problem would be to control the diffusion of activation, first during the learning phase and then during the exploitation phase. This is complex to manage but not impossible to imagine. Models such as those of Jensen, Sharp, or Trullier are very interesting in the sense that they try to use the theta rhythm as a form of short-term memory (Sharp 1991; Jensen and Lisman 1996; Trullier and Meyer 1998). But do their results justify the complexity of their system (eg., the difficulty of getting back the stored information or of adding new information on-line)? The fact that recurrent CA3 connections code for the cognitive map may not be incompatible with our model since ACh modulation can allow different kinds of functioning modes on the same physical map (Hasselmo et al. 1995). In any case, an important unsolved question in that case would be understanding how the prefrontal cortex and the Hs really interact, and what are their respective roles in planning and action-selection process.

In contrast to the previous models, our model can explain Whishaw's statement (Whishaw et al. 1995) that: "the hippocampus is necessary for navigation between spatial location ("getting there") but not for

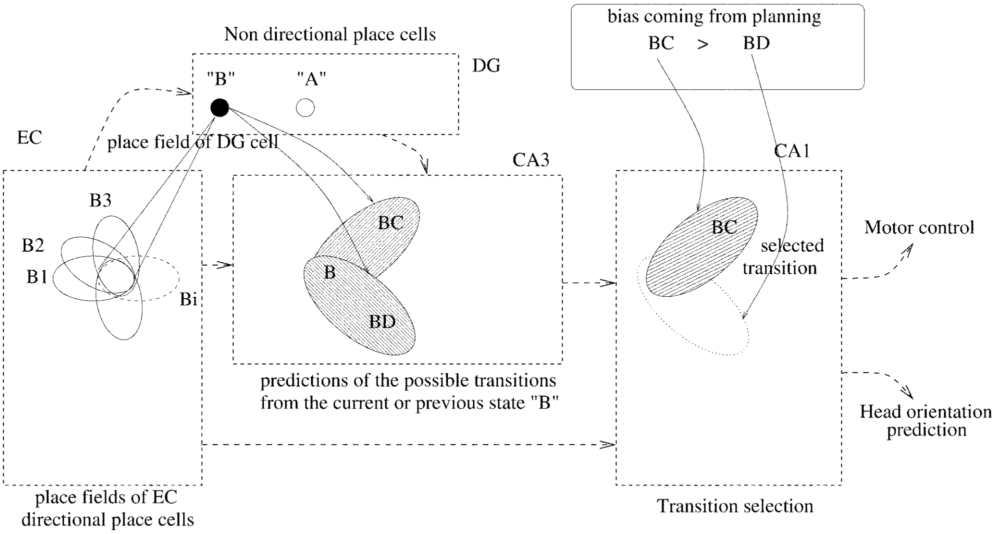

Fig. 15. Relative size of the different place fields in the rodent hippocampus: according to neurobiological evidence, there could be directional place cells in the entorhinal cortex (*EC*). The dentate gyrus (*DG*), due to its time integration capability, could integrate this information to build place cells without orientation preference (classical place cells). CA3 could build transition cells as described in our model. According to the different neurons' functionality, the place fields associated with DG granular cells should be smaller than the place fields of the CA3 pyramidal cells, which are larger than the place fields of CA1 pyramidal cells

the remembering the spatial location of reinforcement based on ambient cues, or for the recognition such as spatial location ("knowing where")." Our model is also confirmed by recent experimental studies showing that the Hs has the ability to represent the causal relationship between an action and its consequence (Corbit and Balleine 2000). An important aspect concerns the recording in rodent Hs of a regular theta rhythm during voluntary movements which is very low or non-existent if no movements (such as grooming activities) are present. The phase precession of neuron spikes during the theta rhythm suggests that hippocampal unit activity predicts the animal's future location on a short (about 120-ms) time scale (Muller and Kubie 1989; Armstrong et al. 1993) and this is also coherent with our model. More detailed simulations should be performed to test all the consequences of phase precession in our model.

Another interest of the present model is that it accounts correctly for the implication of the Hs in conditioning, as reported by Thompson (1990): "Neuronal unit activity in the hippocampus increases markedly within trials early in training. These increases in unit activity in region CA1 and CA3 form a predictive 'model' of the amplitude time course of the learned behavioral response (CR) and happen before the CR in trials, but only under conditions where behavioral learning occurs".

The meaning of the robotics results presented here opens many questions if we compare them to animal visual navigation. First of all, our results – like others obtained in neurobiology these last years – confirm that visual information is very important for place recognition. In contradiction to McNaughton and others (McNaughton 1989; Wan et al. 1994), it appears to us simpler to imagine a model of place cells based mainly on visual information but where the idiothetic information would be used to improve the performance, or to maintain a consistent activity when the light is off or when the environment is visually nondiscriminant (corridor situations or very poor visual environments like those used for most of the neuropsychological studies about navigation). Another argument in favor of the pre-eminence of visual information is the fact that rodents reset their odometric information according to visual information if there is a strong contradiction between them (Etienne et al. 1998). To go further, a more detailed study is required into how integrating conditioning and planning systems is achieved. We will face the problem of understanding the respective roles of the prefrontal cortex and the basal ganglia for the control of action selection and motivated planning. Robotics arguments will be very important in deciding between the different possibilities and testing the coherence of these complex cognitive models.

Acknowledgements. We would like to thank O. Trullier, M. Franz, and N. Burgess for the interesting exchanges we had in Tübingen about hippocampal modeling and visual navigation. Many thanks also to G. Bugmann and I. Fijalkow for their constructive comments on the first drafts.

References

Aggleton JP, Brown MW (1999) Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. Behav Brain Sci 22: 425–489

Amaral DG, Witter MP (1989) The three-dimensional organization of the hippocampal formation: a review of anatomical data. Neuroscience 31: 571–591

Arleo A, Gerstner W (2000) Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity. Biol Cybern 83: 287–299

Armstrong JN, McIntyre DC, Neubort S, Sloviter RS (1993) Learning and memory after adrenalectomy-induced hippocampal dentate granule cell degeneration in the rat. Hippocampus 3: 359–371

Banquet JP, Gaussier P, Dreher JC, Joulain C, Revel A, Günther W (1997) Space-time, order, and hierarchy in fronto-hippocampal system: a neural basis of personality. In: Matthews G (ed) Cognitive science perspectives on personality and emotion. North Holland, Amsterdam, pp 123–189

Banquet JP, Gaussier P, Contreras-Vidal JL, Burnod Y (1998) A neural network model of memory, amnesia and cortico hippocampal interactions. In: Park R, Levin D (eds) Fundamentals of neural network modelling for neuropsychologists. MIT Press, Cambridge, Mass., pp 77–121

Baxter MG, Hollan PC, Gallager M (1997) Disruption of decements in conditioned stimulus processing by selective removal of hippocampal cholinergic input. J Neurosci 17: 5230–5236

Becker JT, Walter JA, Olton DS (1980) Object discrimination by rats: the role of the frontal and hippocampal systems in retention and reversal. Physiol Behav 24: 33–38

Bellman RE (1958) On a routing problem. Q Appl Math 16: 87–90

Berger TW, Thompson RF (1978) Neuronal plasticity in the limbic system during classical conditioning of the rabbit nictitating membrane response. I. The hippocampus. Brain Res 145: 323–346

Bugmann G, Taylor JG, Denham MJ (1995) Route finding by neural nets. In: Taylor JG (ed) Neural Networks. Waller, Henley-on-Thames, pp 217–230

Bunsey M, Eichenbaum H (1996) Conservation of hippocampal memory function in rats and humans. Nature 379: 255–257

Burgess N, Recce M, O'Keefe J (1994) A model of hippocampal function. Neural Netw 7: 1065–1081

Burnod Y (1989) An adaptive neural network: the cerebral cortex. Masson, Paris

Corbit LH, Balleine BW (2000) The role of the hippocampus in instrumental conditioning. J Neurosci 20: 4233–4239

Deacon TW, Einchenbaum H, Rosenberg P, Eckmann KW (1983) Afferent connections of the perirhinal cortex in the rat. J Comp Neurol 220: 168–190

Derrick BE, Martinez JL (1996) Associative, bilateral modifications at the hippocampal mossy fibre-CA3 synapse. Nature 381: 429–434

Dudchenko PA, Wood ER, Eichenbaum H (2000) Neurotoxic hippocampal lesions have no effect on odor span and little effect on odor recognition but produce significant impairements on spatial span, recognition, and alternation. J Neurosci 20: 2964–2977

Eichenbaum H, Otto T, Cohen NJ (1994) Two functional components of the hippocampal memory system. Behav Brain Sci 17: 449–517

Etienne A, Maurer R, Berlie J, Reverdin B, Rowe T, Georgakopoulos J, Séguinot V (1998) Navigation through vector addition. Nature 396: 161–164

Ferino F, Thierry AM, Glowinski J (1987) Anatomical and electrophysiological evidence for a direct projection from ammon's horn to the medial prefrontal cortex in the rat. Exp Brain Res 65: 421–426

Gallistel CR (1993) The organization of learning. MIT Press, Cambridge, Mass

Gaussier P, Zrehen S (1995) Perac: a neutral architecture to control artificial animals. Robot Auton Syst 16: 291–320

Guassier P, Joulain C, Zrehen S, Banquet JP, Revel A (1997) Visual navigation in an open environment without map. In: Proceedings of the International Conference on Intelligent Robots and Systems, Grenoble, France, 7–11 September, pp 545–550

Gaussier P, Joulain C, Banquet JP (1998a) Motivated animal navigation: a visually guided approach. In: Pfeiferr R, Blumberg B, Meyer JA, Wilson SW (eds) Proceedings of Fifth International Conference on the Simulation of Adaptive Behavior. MIT Press, Cambridge, Mass., pp 234–239

Gaussier P, Moga S, Quoy M, Banquet JP (1998b) From perception-action loops to imitation processesses: a bottom-up approach of learning by imitation. Appl Artif Intell 12: 701–727

Gaussier P, Joulain C, Banquet JP, Leprêtre S, Revel A (1999a) The visual homing problem: an example of robotics/biology cross fertilization. Robot Auton Syst 30: 155–180

Gaussier P, Leprêtre S, Quoy M, Revel A, Joulain C, Banquet JP (1999b) Experiments and models about cognitive map learning for motivated navigation. In: Demiris J, Birk A (eds) Interdisciplinary approaches to robot learning. (Robotics and intelligent systems series, vol 24) World Scientific, Singapore, pp 53–94

Goetsch VL, Isaac W (1983) The effect of occipital ablation on visual sensitivity in young and old rats. Physiol Psychol 11: 173–177

Goldman-Rakic PS, Selemon LD, Schwartz ML (1984) Dual pathways connecting the dorsolateral prefrontal cortex with the hippocampus formation and the parahippocampal cortex in the rhesus monkey. Neuroscience 12: 719–746

Grossberg S, Merrill JWL (1996) The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. J Cogn Neurosci 8: 257–277

Han JS, Gallagher M, Holland P (1995) Hippocampal lesions disrupt decrements but not increments in conditioned stimulus processing. J Neurosci 15: 7323–7329

Hasselmo ME, Schnell E, Berke J, Barkai E (1995) A model of the hippocampus combining self-organization and associative memory function. In: Tesauro G, Touretzky D, Leen T (eds) Advances in Neural Information Processing Systems. MIT Press, Cambridge, Mass., pp 77–84

Herrmann T, Poucet B, Ellen P (1985) Spatial problem solving in the rat following medial frontal lesion. Physiol Psychol 13: 21–25

Jensen O, Lisman JE (1996) Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. Learn Mem 3: 243–256

Jones RSG (1993) Entorhinal-hippocampal connections: a speculative view of their function. Trends Neurosci 16: 58–64

Jung MW, McNaughton BL (1993) Spatial selectivity of unit activity in the hippocampal granular layer. Hippocampus 3: 165–182

Knight RT (1996) Contribution of human hippocampal region to novelty detection. Nature 383: 256–259

Kohonen T (1984) Self-organization and associative memory. Springer, New York

Kolb B, Tees RC (1990) The cerebral cortex of the rat. MIT Press, Cambridge, Mass

Kolb B, Sutherland RJ, Wishaw IQ (1983) A comparison of the contributions of the frontal and parietal cortex to spatial localization in rats. Behav Neurosci 87: 13–27

Kosel KC, Hoesen GW, Rosene DL (1982) Non-hippocampal cortical projections from the entorhinal cortex in the rat and rhesus monkey. Brain Res 244: 201–213

Lashley KS (1938) The mechanism of vision. XV. Preliminary studies of the rat's capacity for detailed vision. J Gen Psychol 18: 123–193

Marr D (1971) Simple memory: a theory for archicortex. Phil Trans R Soc Lond B 262: 23–81

McNaughton BL (1989) Neuronal mechanisms for spatial computation and information storage. In: Nadel L, Cooper L, Culicover P, Harnish R (eds) Neural connections and mental computations. MIT Press, Cambridge, Mass, pp 285–349

McNaughton BL, Barnes CA, O'Keefe J (1983) The contribution of position, direction and velocity to single unit activity in the hippocampus of freely-moving rats. Exp Brain Res 52: 41–49

Milner B, Corkin S, Teuber HL (1968) Further analysis of the hippocampal amnesia syndrome: 14-year follow-up study of H.M. Neuropsychologia 6: 215–234

Mishkin M, Ungerleider LG, Macko KA (1983) Object vision and spatial: two cortical pathways. Trends Neurosci 6: 414–417

Mizumori SJY, McNaughton BL, Barnes CA, Fox KB Preserved spatial coding in hippocampal ca1 pyramidal cells during reversible suppression of CA3c output: evidence for pattern competition in hippocampus. J Neurosci 9: 3915–3928

Muller R, Kubie JL (1989) The firing of hippocampal place cells predicts the future position of freely moving rats. J Neurosci 9: 4101–4110

Mumby DG, Wood ER, Pinel JPJ (1992) Object-recognition memory is only mildly impaired in rats with lesions of the hippocampus and amygdala. Psychobiology 23: 26–36

Murray EA, Mishkin M (1998) Object recognition and location memory in monkeys with excitotoxic lesions of the amygdala and hippocampus. J Neurosci 18: 6568–6582

O'Keefe J (1991) The hippocampal cognitive map and navigational strategies. In: Paillard J (ed) Brain and space. Oxford University Press, Oxford, pp 273–295

O'Keefe J, Nadel N (1978) The hippocampus as a cognitive map. Clarendon, Oxford

Perret DI, Oram MW (1994) Modelling visual recognition from neurobiological constraints. Neural Netw 7: 945–972

Quirk GJ, Muller RU, Kubie JL (1990) The firing of hippocampal place cells in the dark depends on the rat's recent experience. J Neurosci 10: 2008–2017

Quoy M, Gaussier P, Leprêtre S, Revel A, Joulain C, Banquet JP (1999) A neural model for the visual navigation and planning of a mobile robot. In: Proceedings of the Fifth European Conference on Advances in Artificial Life, Lausanne, Switzerland, 13–17 September, pp 319–323

Reiss M, Taylor JG (1991) Storing temporal sequences. Neural Netw 4: 773–787

Revel A, Gaussier P, Leprêtre S, Banquet JP (1998) Planification versus sensory-motor conditioning: what are the issues? In: Pfeifer R, Blumberg B, Meyer JA, Wilson SW (eds) Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior. MIT Press, Cambridge, Mass., pp 129–138

Rolls ET, O'Mara SM (1995) View-responsive neurons in the primate hippocampal complex. Hippocampus 5: 409–424

Rolls ET, Treves A (1998) Neural networks and brain function. Oxford University Press, New York

Rumelhart DE, Zipser D (1985) Feature discovery by competitive learning. Cogn Sci 9: 75–112

Schmajuk NA (1991) A neural network approach to hippocampal function in classical conditioning. Behav Neurosci 105: 82–110

Schmajuk NA, Thieme ND (1992) Purposive behavior and cognitive mapping: a neural network model. Biol Cybern 67: 165–174

Schölkopf B, Mallot HA (1995) View-based cognitive mapping and path planning. Adapt Behav 3: 311–348

Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. J Neurol Neurosurg Psychiatry 20: 11–21

Sharp PE (1991) Computer simulation of hippocampal place cells. Psychobiology 19: 103–115

Zola-Morgan SM, Squire LR, Amaral DG, Suzuki WA (1989) Lesions of perirhinal and parahippocampal cortex that spare the amygdala and hippocampal formation produce severe memory impairment. J Neurosci 9: 4355–4370

Solomon PR, Vander Schaaf ER, Thompson RF, Weisz DJ (1986) Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. Behav Neurosci 5: 729–744

Suzuki WA, Zola-Morgan S, Squire LR, Amaral DG (1993) Lesions of the perirhinal and parahippocampal cortices in the monkey produce long lasting memory impairment in the visual and tactile modalities. J Neurosci 13: 2430–2451

Thompson RF (1990) Neural mechanism of classical conditioning in mammals. Phil Trans R Soc Lond B 329: 161–170

Thorpe S, Fize D, Marlot C (1996) Speed of processing in the human visual system. Nature 381: 520–522

Tolman EC (1932) Purposive behavior of animals and men. Irvington, New York

Tolman EC (1948) Cognitive maps in rats and men. Psychol Rev 55: 189–208

Trullier O, Meyer JA (1998) Animat navigation using a cognitive graph. In: Pfeifer R, Blumberg B, Meyer JA, Wilson SW (eds) Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior. MIT Press, Cambridge, Mass., pp 213–222

Trullier O, Wiener SI, Bethoz A, Meyer JA (1997) Biologically based artificial navigation systems: review and prospects. Prog Neurobiol 51: 483–544

Wan HS, Touretzky DS, Redish AD (1994) Towards a computational theory of rat navigation. In: Mozer M, Smolensky P, Touretzky DS, Elman JL, Weigend A (eds) Proceedings of the 1993 Connectionist Models Summer School. Erlbaum, Hillsdale, N.J., pp 11–19

Wehner RH, Michel B, Antonsen P (1996) Visual navigation in insects: coupling of egocentric and geocentric information. J Exp Biol 199: 129–140

Whishaw IQ (1991) Latent learning in a swimming pool place task by rats: evidence for the use of associative and not cognitive mapping processes. Q J Exp Psychol B 43: 83–103

Whishaw IQ, Jarrard LE (1996) Evidence for extrahippocampal involvement in place learning and hippocampal involvement in path integration. Hippocampus 6: 513–524

Whishaw IQ, Cassel JC, Jarrard LE (1995) Rats with fimbria-fornix lesions display a place response in a swimming pool: a dissociation between getting there and knowing where. J Neurosci 15: 5779–5788

Wilson J (1978) Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. Science 201: 160–163

Zipser D (1985) A computational model of hippocampal place fields. Behav Neurosci 99: 1006–1018