

Learning and motivational couplings promote smarter behaviors of an animat in an unknown world

Philippe Laroque^a, Mathias Quoy^a, Philippe Gaussier^a

^aUniversité de Cergy-Pontoise - ETIS, 6, av. du ponceau, 95014 Cergy-Pontoise Cedex, France
quoy@ensea.fr

<http://www-etis.ensea.fr/>

This short article reports on statistical results of an animat behavior in an unknown environment using a cognitive map. We study the coupling effect between the motivation for eating and the one for drinking. Some smart behaviors are not caused by a sophisticated “intelligent” algorithm, but only through coupling of the motivations. Adding a learning rule on the links of the cognitive map allows to reinforce particular paths, and to forget others. This also leads to reinforce the same “smart” behaviors.

1. Introduction

The cognitive map approach to planning is not new [39,2,28,37,5,38,11,40,19]. However, adding a learning procedure on the links of the cognitive map enables to reinforce particular paths, and to forget others, allowing for an adaptation to a changing perception of the environment, as it is gradually discovered by the animat.

As far as planning is concerned, most models use graphs giving how to go from one place to the other [2,15,28,30,37,5,11,38,40,19,20,24]. The differences between them lay in the construction of this graph, the way it is used in order to find the shortest path, how it works when the environment changes or when there are different simultaneous goals to be achieved. Other works use rule based algorithms (hierarchical classification [17]) or “potential fields” [3], or even classical A^* algorithms. The classical functional approach relies on the *a priori* definition of desired behaviors [13] or on a complete world modeling [41]. These classical programming techniques may allow to produce navigation models exhibiting behaviors similar to ours. To sum up, the navigation problem coupled with a planning mechanism may be solved by dedicated efficient techniques. We want to stress here that the navigation/planning part of our system is only a piece of a more complex system which should be able to recognize objects [26], to imitate [29,1], and finally to communicate.

So, from the beginning, we have to integrate the fact that our final goal is more than “just” navigating and planning. In particular, we stress that it is important to take into account the dynamical aspects of internal and external couplings.

These dynamical aspects are explored here as the coupling between the animat’s motivations (eat, drink, rest ...), and their interaction with the environment. So, we will study how this coupling affects the animats behavior. In the following, we first present the experiment, then our model and the results achieved. Finally, we conclude.

2. Experiment

In our experiments, the animat is “living” in an unknown environment with obstacles, “food” and “water” sources and a “nest” to rest. The animat is switching between two behaviors: exploring the environment and planning. The animat has three contradictory motivations: eating, drinking and resting. Each of these motivations is linked with a satisfaction level decreasing over time, and increasing when on the proper supply source. When the level is below some threshold, the animat triggers the planning behavior for reaching the needed source.

Most experiments have been carried out in a maze environment. This enables us to control the kind of paths used by the animat by varying the size of the arms of the maze. A corridor is 7 animat’s

width wide. The animat needs approximately 50,000 steps for exploring the whole maze.

The animat only gets two informations. A visual input gives the existence of a particular landmark. A compass gives the azimuth of these landmarks relatively to the magnetic north (top of the environment in the computer simulation). Hence a particular location is defined by the set of (landmark, azimuth). There is no Cartesian map of the environment, nor use of a square paving as in some other planning algorithms (Q-learning for instance).

In the right arm of the T-maze, we place a food source. In the left one, there are both food and water. In the bottom of the vertical arm, there is a nest, where the animat rests. It is the initial starting location. The need for going to rest increases twice as fast as the ones for drinking or eating.

3. Model

The cognitive map is inspired by works on the possible use of “Place Cells” of the rat [31]. Place Cells are particular neurons of the CA3 region of the rat’s hippocampus. These neurons are typical for a particular location in the environment (they only fire when the rat comes to the corresponding location). Hence the various places of a room for instance are coded by Place Cells becoming active when the rat comes into its response field. Some recent works of spatial navigation include the simulation of some biologically plausible neuronal structures, and in particular the use of “place cells” as they may exist in the hippocampus [12,36,6,16,7,40,35,22].

In the beginning, the animat has no knowledge about the visual forms it can encounter, except that there are some landmarks. The set of landmarks (together with their azimuth) the animat can see at a given moment determines a “known” location. This location is coded in a particular neuron on a probabilistic topological map (PTM). We call this neuron a Place Cell. Initially, the weights of the connections in this map are randomly generated, so any neuron can win. When a new configuration is coded on a given Place Cell, the weights of its neighbors are updated with a

probability depending on the Euclidean distance between the weight vector of the winner and that of each neighbor. That allows for a neuron “near” the winner to be able to react to a situation similar to the newly learned situation. A configuration is considered as “new” when the activation of the winner comes under a given threshold. The new set of (landmark, azimuth) is then learned on a new Place Cell. In that way, the exploration of the environment leads to the creation and connection of a various number of Place Cells.

$$Act_{P_i} = 1 - \frac{\sum_{k=1}^{N_i} V_{i,k} \cdot f(|\Theta_{i,k} - \theta_k|, v_k)}{\pi N_i} \quad (1)$$

In that equation, N_i is the number of visible landmarks when the robot is at the learned place i (or cell P_i). $\Theta_{i,k}$ represents the learned value of the azimuth of landmark k from the learned place i . θ_k is the value of the same landmark azimuth for the current robot location. All angles are expressed in radians and measured from an absolute direction (the north for instance, given by a compass). $|\Theta_{i,k} - \theta_k|$ is computed modulo π , $V_{i,k}$ and v_k are set to 1 when the landmark k is seen from the learned location i and from the current robot location respectively (0 otherwise). When the system exhibits landmark recognition problems, we can have $V_{i,k} = 1$ and $v_k = 0$ (landmark learned but not recognized). f is a non linear function solving these landmark visibility problems:

$$f(\theta, v_k) = \begin{cases} \theta & \text{if } v_k = 1 \\ \pi & \text{if } v_k = 0 \end{cases}$$

The error associated with one landmark azimuth is maximum when the landmark cannot be found ($f(\theta, 0) = \pi$). Before learning, all neurons synaptic weights are set to a random value, so the neurons response is very low. When learning is triggered (for instance when the goal is discovered), the most activated neuron is selected to learn that place ($\Theta_{i,k} \leftarrow \theta_k$). Next, according to eq. 1 the activation of that neuron will grow and tend to be maximum ($Act = 1$) when the robot moves in the direction of the learned location (the azimuths θ_k associated with the current location

are close to the stored $\Theta_{i,k}$). By using a simple gradient technique, it is theoretically possible to reach the goal wherever the starting point of the robot is [34].

So, the exploration of the environment leads to the building of a various number of Place Cells. At this time, there is no cognitive map, but only a collection of Place Cells. We have shown that it is already sufficient for reaching a particular goal when each Place Cell is associated with a movement to perform [23]. When thinking of a real robot, the first problem is to reduce the information flow coming from the CCD camera. It may be done by a parallelizing visual information processing [21,32]. But the main reduction is due to the fact that a particular behavior will not be associated with a whole scene, but only with the combination of local snapshots taken around particular points (landmarks) [25,14]. Thus a view will be reduced to a set of snapshots. These snapshots are taken around points of high contrast change. We could also take other techniques such as comparing panoramic views by computing their displacement field [18].

For more complex tasks, when several goals have to be achieved or when the system has to learn to navigate in a complex visual environment (the different goals do not belong to the same visual environment), the choice of the route to follow cannot only rely on simple sensory-motor associations. We tackle this problem by using a cognitive map. The map is built by linking together two Place Cells reached successively. This builds a graph which roughly represents the topology of the environment in figure 1.

This map alone is not enough for planning. It has to be associated to a motivational system so that the most interesting places (food and water sources) can be linked with the appropriate motivation. The activity of a neuron is now modulated by the motivation activity, so that we have:

$$Pot_i = Act_i * Motiv_i \quad (2)$$

Where $Motiv_i$ is the the motivation arriving on neuron i .

Now, when a particular need has to be fulfilled, the associated motivation triggers the activation




Figure 1. The animat in a T-Maze whose left arm is 1.6 longer than the right one.

of the appropriate neuron in the cognitive map. The diffusion of this activity along the graph until the actual location of the animat defines the route to follow for reaching the goal [15,11,10]. The neural diffusion we use is equivalent to the Bellman-Ford algorithm [8,34]. The activity of each neuron of the graph is the maximum of its neighbors times the connection weight (see 3). So the activity of a neuron at a “distance” of N links from the activate goal is W^N , if all weights have the same value W . Taking the maximum enables to select one of the goals and suppresses local minima problems [24]

The kind of map built (figure 2) this way is fixed in the sense that the connection between the neurons may not change. Hence, for instance, when the environment changes (a door opens, somebody’s walking ...) the map may not be used anymore. So it is necessary to be able to modify the links and/or to create/delete neurons when appropriate. For the time being, we have only investigated the first possibility. So, we have defined a hebbian learning rule reinforcing the links between neurons often activated successively and decreasing them if the link is not used. Note that this rule is not a “one shot learning”. The links

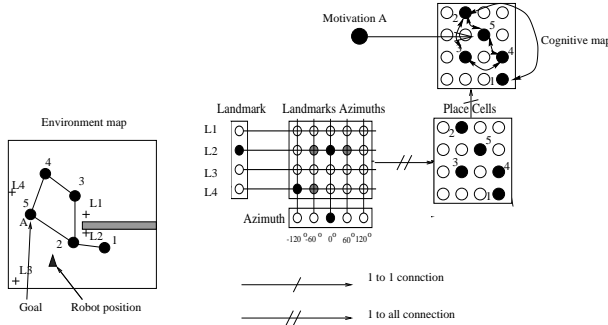


Figure 2. Global sketch of the architecture. An animat in a simulated environment (on the left) and the corresponding landmark/azimuth fusion, place cell map and cognitive map. The cognitive map is projected on the environment in order to see where the place cells have been learned, and that are the links between them.

have to be used several times before significantly increasing or decreasing. So when a change occurs in the environment, it may take some time before the new map is updated.

Let W_{ij} be the weight associated with the fact that from place P_i it is possible to reach directly place P_j , its learning rule is the following:

$$\frac{dW_{i,j}}{dt} = -\lambda W_{i,j} + (1 + \frac{dR}{dt}) \cdot (1 - W_{i,j}) \cdot \overline{Pot_{P_i}} \cdot Pot_{P_j} \quad (3)$$

Pot_{P_j} is the activity of place cell P_j . It is 0 if the animat is not at location P_j and non zero and positive otherwise (the exact value depends on the recognition of landmarks and may be also modulated by some motivation). Increase of weights $W_{i,j}$ occurs when Pot_{P_i} and Pot_{P_j} are non zero. Since the animat can not be at both locations simultaneously, we have to maintain the activity of place cell P_i until P_j is reached. This is performed by a time integration of the Pot_{P_i} values represented in the equation by $\overline{Pot_{P_i}}$. $\overline{Pot_{P_i}}$ decreases with time and can be used as a raw measure of the distance between i and j . λ is a very low positive value. It allows to forget unused links ($Pot_{P_i} = 0$ and $Pot_{P_j} = 0$). The term

$\frac{dR}{dt}$ corresponds to the variation of an external reinforcement signal (negative or positive) that appears when the animat enters or leaves a “difficult” or “dangerous” area.

4. Results

We are studying the way the animat behaves as the size of the left arm increases. In the beginning of each experiment, we choose the ratio “left arm/right arm”. The animat starts from its nest, which is located at the bottom of the vertical bar of the T-maze (fig.1). There are two phases in the experiment:

- the “discovery” phase, in which the animats has to discover the three sources. It then can choose either the left or the right food source when it is hungry.
- after its paths are stabilized (after 50 “meals”), we record the animat choices for the next 100 “meals”.

Each experiment is run during 50,000 steps, and has been tested 30 times. We give below the average choice rate for the left food source as the left arm length increases.

Whatever the size of the left arm, the animat has to go in it because it contains the only water source. However, it may eat either in the left or in the right arm. The basic expected behavior is to go to eat in the nearest food source when the eating motivation is triggered. As the animat is more often in its nest, it will go to the food source which is nearer to the nest. So as the left arm size increases, the right food source will be preferred. A “smarter” behavior would be to go to eat near the place where the water source is. Indeed, this could save energy because if the animat gets thirsty while traveling to the food source, it would be also on the way to the water source. Note that there is no rule in our system able to foresee such a behavior.

On a symmetrical maze, we observed that, after some time, the animat almost always chooses to eat in the left arm. This is due to the *coupling* after some time of the eating and drinking motivations. This coupling occurs as soon as the animat eats just after (or before) drinking, so that

the two levels are approximately the same. The effect of this coupling, is reported in a first experiment. We compare the animat behavior when (des-)activating the coupling between food and water needs. This is achieved by “forcing” it to drink each time it eats, so that it is never thirsty when it is hungry again. We make the left arm grow from 1.0 to 1.8 times the initial size. The results are graphically represented in fig.3.

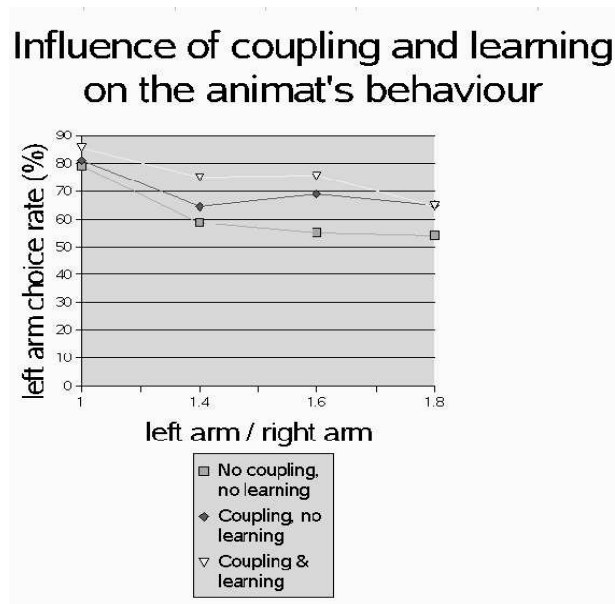


Figure 3. Influence of coupling and learning on the left arm selection rate. The top curve (triangles) is with coupling and learning. The intermediate curve (diamonds) is with coupling and without learning, and the bottom curve (squares) is without coupling and learning. Enabling coupling leads to choose more often the left arm. Allowing learning still increases the choice of the left arm.

As can be seen on the results, coupling leads to an increased preference for the left food source. However, we chose to only reduce, and not fully suppress, this coupling since the need for water

still happens from time to time. If the need for water was completely suppressed, the left arm selection rate would decrease from 50% to 0 as the size of the left arm increases. So this kind of smart behavior may be achieved by the sole coupling between internal needs and the environment.

Now, we activate the learning rule on the links of the cognitive map. This rule increases the weight between two Place Cells activated successively. So, the more used the path, the higher the link. Conversely, a path less often used will see its link value decrease. As a consequence, when a motivation is activated, the source, which is nearer in terms of number of links from the actual position, is not necessarily the one providing the highest activity. So the animat may choose to go to a place which is farther away in terms of Euclidean distance, but closer in terms of cognitive map representation. So there is a remapping of the perception of the environment depending on the experience of the animat. In particular, paths often used will be reinforced, so that they will be perceived as “shorter”.

So, as shown in the following experiments (fig. 3), the consequence of the learning rule is to keep choosing the left arm for eating, thus still following the “smart” behavior.

5. Conclusion and perspectives

The planning mechanism alone is not enough to account for smart behaviors. Their emergence, in our experiment, is first due to the coupling between the different motivational variables. This emphasizes the need for taking into account the internal dynamics of the animat and its coupling with the environment. Adding a learning rule on the planning map contributes to reinforce this behavior.

The learning rule we have described may be used in the more general context of a changing environment. Indeed, when the environment changes, new place cells are learned corresponding to the new landmark configuration. Their attraction basin extends over the previously created ones. In the extreme case, a place cell learned in the old environment may not fire at all in the new one. Since the place cell is not activated, it

is as if it does not exist. So it may neither be linked to other place cells anymore, nor may the links be reinforced. On the contrary, the learning rule enables to forget them progressively. However, if the old environment is presented again, the old place cells fire again, so that their links are again valid. So several different “layers” of cognitive maps can appear in the same physical N.N. structure [33]. They may be linked together through some place cells, and may be activated successively when the environment they are coding is presented again. However, recent neurobiological findings rather support the evidence of maps stored for a short time at the hippocampal level which are transferred in for long time storage [27,9]. When the animal is brought back in the original environment, the same cognitive map appears, as if it were “uploaded” from the cortical areas.

The animat has to keep trying for some time before building a new efficient planning map. So, even if it may be a good solution to use a very low passive decay (λ parameter) to store several different maps (memory effect), it also slows the process of finding new pathways, when one may not be used anymore. Indeed, the hebbian learning rule we have chosen needs some time before significantly changing the weights. Hence, in order to react faster to a change in the environment, it would be necessary to introduce an active decay mechanism decreasing unused links.

This navigation model does not take into account a path integration mechanism. This mechanism allows an animal to find the direct path to a particular location even if this movement has not been experienced before (short-cut). Our mechanism only allows to reinforce known paths and is not able to try a way direct to the goal if it is not an existing path.

It is found that a deadlock situation occurs when a place field of a neuron farther from the goal overlaps the trajectory to the goal (fig. 4). Thus when the animat goes into this place field, it recognizes a location farther away from the goal (though topologically nearer on the planning map) and goes back from where it came. When it leaves the place field, it goes back towards the goal again. It follows the same trajectory as pre-




Figure 4. Dead-lock behavior. On the left, the animat is at position 1 heading for position 2 which corresponds to the most active neighbor. Once at 2, the animat goes to 3. However, while traveling to 3, the animat goes through the area where place cell 1 fires (figure on the left). Hence the animat goes back to position 1 and begins to circle in a deadlock situation.

viously, and so eventually ends in the same overlapping place field. We are now back to the previous situation where the animat recognizes a location which is (for him) farther away from the goal. So we have now a deadlock situation. Creating a new place cell in this situation does not resolve the problem since it would only be linked with place cell 1 (see fig. 4). Thus, for the moment when a dead-lock situation is detected, the animat switches to a random walk for some time before resuming the planning behavior. It should be noted that this kind of deadlock behavior mainly occurs because some links are missing: ie. the cognitive map is not explored completely. For instance, a link between neuron 1 and 3 in figure 4 would lead to a correct planning behavior. Path integration could also be used for going directly from neuron 1 to neuron 3.

Finally, we have only considered in this work Place Cells in the classical approach: coding for a particular location in the environment. In the robotics experiments we have also carried out, we use these neurons for coding *transitions* between locations. This enables to associated such a neuron with the movement performed and given by a path integration mechanism [4]. So path integration and transition coding are the two next improvements to the present simulation platform.

REFERENCES

1. P. Andry, P. Gaussier, S. Moga, A. Revel, and J. Nadel. Imitation: Learning and communication. In *The Sixth International Conference on the Simulation of Adaptive Behaviour, SAB 2000*, Paris, september, 2000.
2. M.A. Arbib and I. Lieblch. Motivational learning of spatial behavior. In J. Metzler, editor, *Systems Neuroscience*, pages 221–239. Academic Press, 1977.
3. R.C. Arkin. Motor schema-based mobile robot navigation. *International Journal of Robotics Research*, pages 92–112, 1987.
4. V. Babeau, P. Gaussier, C. Joulain, A. Revel, and J.P. Banquet. Merging visual place recognition and path integration for "cognitive" map learning. In *The Sixth International Conference on the Simulation of Adaptive Behaviour, SAB 2000*, Paris, september, 2000.
5. I.A. Bachelder and A.M. Waxman. Mobile robot visual mapping and localization: A view-based neurocomputational architecture that emulates hippocampal place learning. *Neural Networks*, 7(6/7):1083–1099, 1994.
6. I.A. Bachelder and A.M. Waxman. A view-based neurocomputational system for relational map-making and navigation in visual environments. *Robotics and Autonomous System*, 16(2-4), December 1995.
7. J.P. Banquet, P. Gaussier, J.C. Dreher, C. Joulain, and A. Revel. *Cognitive Science Perspectives on Personality and Emotion*, chapter Space-Time, Order and Hierarchy in Fronto-Hippocampal System: A Neural Basis of Personality. Elsevier Science BV Amsterdam, 1997.
8. R.E. Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16:87–90, 1958.
9. Bruno Bontempi, Catherine Laurent-Demir, Claude Destrade, and Robert Jaffard. Time-dependent reorganization of brain circuitry underlying long-term memory storage. *Nature*, 400:671–675, 1999.
10. G. Bugmann. *Basic Concepts in Neural Networks: A survey*, chapter Chap 5: A Connectionist Approach to Spatial Memory and Planning. Perspectives in Neural Networks. Springer, 1997.
11. G. Bugmann, J.G. Taylor, and M.J. Denham. Route finding by neural nets. In J.G. Taylor, editor, *Neural Networks*, pages 217–230, Henley-on-Thames, 1995. Alfred Waller Ltd.
12. N. Burgess, M. Recce, and J. O'Keefe. A model of hippocampal function. *Neural Networks*, 7(6/7):1065–1081, 1994.
13. R. Chatila. Deliberation and reactivity in autonomous mobile robots. *Robotics and Autonomous System*, 16(2-4):197–211, December 1995.
14. T.S. Collett and J. Zeil. The selection and use of landmarks by insects. *Orientation and Communication in Arthropods*, pages 41–65, 1997.
15. C.I. Connolly, J.B. Burns, and R. Weiss. Path planning using laplace's equation. In *International Conference on Robotics and Automation*, pages 2102–2106, 1990.
16. M.J. Denham and J. Boitano. A model of the interaction between prefrontal cortex, septum and the hippocampal system in the learning and recal of goal-directed sensory-motor behaviours. Technical Report NRG-96-01, University of Plymouth - School of computing, 1996.
17. J.Y. Donnart and J.A. Meyer. Learning reactive and planning rules in a motivationnally autonomous animat. *IEEE Transactions on Systems, Man and Cybernetics-Part B*, 26(3):381–395, 1996.
18. M.O. Franz, B. Schölkopf, and H.H. Bülthoff. Homing by parameterized scene matching. In *Proc. 4th European Conf. on Artificial Life, MIT Press, Cambridge*, pages 236–245, 1997.
19. M.O. Franz, B. Schölkopf, H.A. Mallot, and H.H. Bülthoff. Learning view graphs for robot navigation. *Autonomous Robots*, 5:111–125, 1998.
20. M.O. Franz, B. Schölkopf, H.A. Mallot, and H.H. Bülthoff. Where did i take that snapshot? scene-based homing by image matching. *Biol. Cybern.*, 79:191–202, 1998.
21. O. Gallet, P. Gaussier, and J.P. Cocquerez. A parallel singularity amplifier model for focusing visual attention. In B. Dresp and C. Bonnet, editors, *European Conference on Visual Perception, ECVP'96*, page 74, Strasbourg, September 1996. Congrès Pasteur.
22. P. Gaussier, C. Joulain, J.P. Banquet, S. Leprêtre, and A. Revel. The visual homing problem: an example of robotics/biology cross fertilization. *Robotics and Autonomous Systems*, 30:155–180, 2000.
23. P. Gaussier, C. Joulain, S. Zrehen, J.P. Banquet, and A. Revel. Visual navigation in an open environment without map. In *International Conference on Intelligent Robots and Systems - IROS'97*, pages 545–550, Grenoble, France, September 1997. IEEE/RSJ.
24. P. Gaussier, S. Leprêtre, M. Quoy, A. Revel, C. Joulain, and J.P. Banquet. *Interdisciplinary approaches to robot learning*, volume 24, chapter Experiments and models about cognitive map learning for motivated navigation, pages 53–94. Robotics and Intelligent Systems Series, World Scientific, ISBN 981-02-4320-0, 2000.
25. P. Gaussier and S. Zrehen. Navigating with an animal brain : a neural network for landmark identification and navigation. In *Proceedings of Intelligent Vehicles*, pages 399–404, Paris, 1994.
26. S. Leprêtre, P. Gaussier, and J.P. Cocquerez. From navigation to active object recognition. In *SAB 2000*, Paris, France, september 2000.
27. N. Ludvig. Place cells can flexibly terminate and develop their spatial firing. a new theory for their function. *Physiology and Behavior*, 67(1):57–67, 1999.
28. M.J. Matarić. A distributed model for mobile robot environment learning and navigation. In *Technical report TR 1228 MIT IA Lab*, 1990.
29. S. Moga and P. Gaussier. A neuronal structure for learning by imitation. In F. Mondada D. Floreano, J.D. Nicoud, editor, *European Conference on Artificial life, ECAL99*, volume 1674, pages 314–318, Lausanne, september 1999.

30. U. Nehmzow and T. Smithers. Mapbuilding using self-organising networks. In J.A. Meyer and S. Wilson, editors, *From Animals to Animats: SAB'91*, Cambridge, MA, 1991. MIT Press.
31. J. O'Keefe and N. Nadel. *The hippocampus as a cognitive map*. Clarendon Press, Oxford, 1978.
32. M. Quoy, S. Moga, P. Gaussier, and A. Revel. Parallelization of neural networks using pvm. In J. Dongarra, P. Kacsuk, and N. Podhorszki, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pages 289–296, Berlin, 2000. Lecture Notes in Computer Science, Springer, no 1908.
33. A. Redish and D. Touretzky. Separating hippocampal maps. In K. Jeffery, N. Burgess, and J. O'Keefe, editors, *The hippocampus and parietal foundations of spatial cognition*, Monte Verita, Ticino, Switzerland, 1999. Oxford University Press.
34. A. Revel, P. Gaussier, and J.P. Banquet. Taking inspiration from the hippocampus can help solving robotics problems. In *European Symposium on Artificial Neural Networks*, Bruges, Belgium, Avril 1999. IEEE.
35. Alexei Samsonovich and Bruce McNaughton. Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, 17:5900–5920, 1997.
36. N.A. Schmajuk and H.T. Blair. Place learning and the dynamics of spatial navigation: a neural network approach. *Adaptive Behavior*, 1:353–385, 1992.
37. N.A. Schmajuk and A.D. Thieme. Purposive behavior and cognitive mapping: a neural network model. *Biological Cybernetics*, 67:165–174, 1992.
38. B. Schölkopf and H.A. Mallot. View-based cognitive mapping and path-finding. *Adaptive Behavior*, 3:311–348, 1995.
39. E.C. Tolman. Cognitive maps in rats and men. *The Psychological Review*, 55(4), 1948.
40. O. Trullier, S.I. Wiener, A. Berthoz, and J.A. Meyer. Biologically based artificial navigation systems: review and prospects. *Progress in Neurobiology*, 51:483–544, 1997.
41. Uwe R. Zimmer. Robust world-modelling and navigation in a real world. *Neurocomputing*, 13:247–260, 1996.