



ANALYZING INTERACTIONS
BETWEEN NAVIGATION
STRATEGIES USING A
COMPUTATIONAL MODEL OF
ACTION SELECTION

Laurent Dollé ^a Mehdi Khamassi ^{a b}
Benoît Girard ^b Agnès Guillot ^a
Ricardo Chavarriaga ^c
IDIAP-RR 08-48

JUNE 2008

REVISED IN MARCH 2008

PUBLISHED IN

Int Conf Spatial Cognition 2008, Freiburg, Germany, September 2008

^a ISIR, FRE2507, Université Pierre et Marie Curie - Paris 6, Paris, F-75016, France
^b LPPA, UMR7152 CNRS, Collège de France, Paris, F-75005, France
^c IDIAP Research Institute, Martigny, Switzerland

ANALYZING INTERACTIONS BETWEEN NAVIGATION STRATEGIES USING A COMPUTATIONAL MODEL OF ACTION SELECTION

Laurent Dollé Mehdi Khamassi Benoît Girard Agnès Guillot
Ricardo Chavarriaga

JUNE 2008

REVISED IN MARCH 2008

PUBLISHED IN

Int Conf Spatial Cognition 2008, Freiburg, Germany, September 2008

Abstract. For animals as well as for humans, the hypothesis of multiple memory systems involved in different navigation strategies is supported by several biological experiments. However, due to technical limitations, it remains difficult for experimentalists to elucidate how these neural systems interact. We present how a computational model of selection between navigation strategies can be used to analyse phenomena that cannot be directly observed in biological experiments. We reproduce an experiment where the rat's behaviour is assumed to be ruled by two different navigation strategies (a cue-guided and a map-based one). Using a modelling approach, we can explain the experimental results in terms of interactions between these systems, either competing or cooperating at specific moments of the experiment. Modelling such systems can help biological investigations to explain and predict the animal behaviour.

Introduction

In natural environments, animals encounter situations where they have to simultaneously learn various means for reaching interesting locations and to select dynamically the best to use. Many neurobiological studies in both rodents and humans have investigated how this selection is performed using experimental paradigms in which several navigation strategies may be learned in parallel. Some strategies may be based on a spatial representation (i.e., inferring a goal-directed action as a function of its location, called *map-based* strategies), whereas other strategies can be based on direct sensory-motor associations without requiring a spatial representation (i.e., *map-free*) [1, 2, 3, 4]. A number of experimental results lead to the hypothesis that these strategies are learned by separate memory systems, with the dorsolateral striatum involved in the acquisition of the map-free strategies and the hippocampus mediating the map-based strategy [5, 6].

However, it is not yet clear whether these learning systems are independent or whether they interact for action control in a competitive or in a cooperative manner. The competition implies that inactivation of one system enhances the learning of the remaining functional system, while the cooperation states that learning in one system would compensate the limitations of the other one [7, 8, 9, 10, 11]. The present work aims at investigating such interactions using a computational model of spatial navigation based on the selection between the map-based and map-free strategies [12]. Besides a qualitative reproduction of the experimental results obtained in animals, the modelling approach allows us to further characterize the competitive or cooperative nature of interactions between the two strategies.

Following our previous modelling efforts [12], we study the interaction between the navigation strategies in the experimental paradigm proposed by Pearce et al. (1998) [13]. In this paradigm, which is a modification of the Morris Hidden Water Maze task [14], two groups of rats (“Control” group of intact animals and “Hippocampal” group of animals with damaged hippocampus) had to reach a hidden platform indicated by a landmark located at a fixed distance and orientation from the platform. After four trials, the platform and its associated landmark were moved to another location and a new session started. The authors observed that both groups of animals were able to learn the location of the hidden platform, but at the start of each new session the hippocampal animals were significantly faster in finding the platform than controls. Moreover, only the control rats were able to decrease their escape latencies within a session. From these results, authors conclude that rats could simultaneously learn two navigation strategies. On the one hand, a map-based strategy encodes a spatial representation of the environment based on visual extra-maze landmarks and self-movement information. On the other hand, a map-free strategy (called by the authors “heading vector strategy”) encodes the goal location based on its proximity and direction with respect to the intra-maze cue [15]. Based on these conclusions, the decrease in the escape latency *within* sessions could be explained by the learning of a spatial representation by intact animals. Furthermore, such learning also suggests that when the platform is displaced at the start of a new session, intact rats would swim to the previous (wrong) location of the platform based on the learned map, whereas hippocampal animals would swim directly to the correct location.

For the modelling purposes, the results of this experiment can be summarized as follows: *(i)* both groups of rats could decrease their escape latencies *across* sessions, but only the control rats improved their performance *within* sessions; *(ii)* the improvement in the performance within each session, observed in the control group, could be attributed to the use of a map-based strategy by these rats; and *(iii)* higher performance of hippocampal rats relative to the controls at the start of each session could be due to the use of the map-free strategy (the only strategy that could be used by the lesioned animals). In other words, the process of choosing the best strategy (i.e. the competition) performed by the control, but not the hippocampal, animals, decreased the performance of controls relative to that of lesioned animals.

We have shown previously that the computational model used in the present study is able to reproduce the behaviour of rats in the experiment of Pearce et al. [12]. In the present paper, we extend these results by performing a further analysis of the interactions between both learning systems

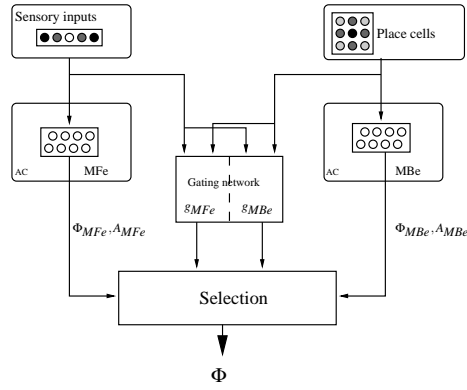


Figure 1: The computational model of strategy selection [12]. The gating network receives the inputs of both experts, and their reward prediction error, in order to compute their reliability according to their performance (i.e., gating values g_k). Gating values are then used with the Action value A_k in order to compute the probability of each expert to be selected. The direction Φ proposed by the winning expert is then performed. See text for further explanations.

at different stages of the experiment, taking into account the three points formulated above. In the following section, we describe the model, the simulated environment and the experimental protocol. Then we present the results and the analyses. Finally, we discuss the results in terms of interactions between systems.

1 Methods and Simulation

1.1 Navigation Model

The neural network computational model is based on the hypothesis of different, parallel learning systems exclusively involved in each strategy, and interacting for behaviour control (Fig. 1). It is composed of two experts, learning separately a map-based strategy and a map-free one (the experts are denoted MBe, and MFe, respectively), both following reinforcement learning rules to acquire their policy, i.e., the way the expert chooses an action given the current state in order to maximize the reward. The model provides a mechanism that selects, at each timestep, which strategy should drive the behaviour of the simulated robot, given its reliability to find the goal. This section briefly describes both navigational experts, their learning process as well as the selection mechanism and the learning mechanism underlying this selection (for a more detailed description see [12]).

The map-free strategy is encoded by the MFe that receives visual signals from *sensory cells* (SI), consisting of a vector of 36 inputs of gray values (one input for every ten degrees), transducing a 360 degrees horizontal 1-D gray-scale image. To simulate the heading vector strategy stated by Pearce et al., the landmark is viewed with an allocentric reference: for example, when the landmark is located to the North with regards to the robot, it will appear in the same area of the camera, whatever might be the orientation of the robot.

The map-based strategy is encoded by the MBe that receives information from a spatial representation encoded in a regular grid of 1600 place cells (PC) with Gaussian receptive fields of width σ_{PC} [16] (values of all model parameters are given in Table 1).

1.1.1 Strategy Learning.

Both experts learn the association between their inputs and the actions leading the robot to the platform, using a direct mapping between inputs (either SI or PC) and directions of movement (i.e.,



Figure 2: (a) A simplified view of ad hoc place cells. Each *circle* represents a place cell and is located at the cell's preferred position (i.e., the place where the cells are most active). Cell activity is color coded from white (inactive cells) to black (highly active cells) (b) The environment used in our simulation (*open circles*: platform locations, *stars*: landmarks).

actions). Movements are encoded by a population of 36 *action cells* (AC). The policy is learned by both experts by means of a neural implementation of Q-learning algorithm [17]. In this algorithm, the value of every state-action pair is learned by updating the synaptic weight w_{ij} linking input cell i to action cell j :

$$\Delta w_{ij} = \eta h_k \delta e_{ij}, \quad (1)$$

where η is the learning rate and δ the reward prediction error. The scaling factor h_k ensures that the learning module is updating its weights according to its reliability (for all the following equations, k is either the MBe or the MFe). Its computation is detailed further below. The eligibility trace e allows the expert to reinforce the state-action couples previously chosen during the trajectory:

$$e_{ij}(t+1) = r_j^{pre} r_i + \lambda e_{ij}(t), \quad (2)$$

where r_j^{pre} is the activity of the pre-synaptic cell j , λ a decay factor, and r_i the activity of the action cell i in the generalization phase. Generalization in the action space is achieved by reinforcing every action weighted by a Gaussian of standard deviation σ_{ac} centered on the chosen action. Each expert suggests a direction of movement Φ_k :

$$\Phi_k = \arctan\left(\frac{\sum_i a_i^k \sin(\phi_i)}{\sum_i a_i^k \cos(\phi_i)}\right), \quad (3)$$

where a_i is the action value of the discrete direction of movement ϕ_i . The corresponding action-value A_k , computed by linear interpolation of the two nearest discrete actions [17].

1.1.2 Action Selection.

In order to select the direction Φ of the next robot movement, the model uses a gating scheme such that the probability of being selected depends not only on the Q-values of the actions (A_k), but also on a gating value g_k . Gating values are updated in order to quantify the expert's reliability according to the current inputs. It takes the shape of a network linking the inputs (place cells and sensory inputs) to the gating values g_k , computed as a weighted sum:

$$g_k = z_k^{PC} r^{PC} + z_k^{SI} r^{SI}, \quad (4)$$

where z_k^{PC} is the synaptic weight linking the PC, with activation r^{PC} to the gate k , idem for z_k^{SI} . Weights are updated in order to approach $h_k = \frac{g_k c_k}{\sum_i (g_i c_i)}$ where $c_k = e^{(-\rho \delta_k^2)}$ ($\rho > 0$), according to the following rule:

$$\Delta z_{kj}^{PC,SI} = \xi (h_k - g_k) r_j^{PC,SI}. \quad (5)$$

Table 1: Parameters of the model

Parameters	Value	Description
N_{PC}	1600	Number of place cells
σ_{PC}	10 cm	Standard deviation of PC’s activity profile
N_{AC}	36	Number of action cells
σ_{AC}	22.5 °	Standard deviation of the enforced activity profile
η	0.015	Learning rate of both experts
λ	0.76	Decay factor of both experts
ξ	0.01	Learning rate of the gating network
ρ	1.0	Decreasing rate in c_k

The next action will be then chosen according to a probability of selection P :

$$P(\Phi = \Phi_k) = \frac{g_k A^k}{\sum_{i \in k} g_i A^i}. \quad (6)$$

If both experts have the same gating value (i.e., *reliability*), then the expert with the highest action value will be chosen. In contrast, if both experts have the same action value, the most reliable expert, i.e., the one with highest gating value, will be chosen.

1.2 Simulated Environment and Protocol

In our simulation, the environment is a square of size equivalent to 200×200 cm, while the simulated robot’s diameter is 15 cm (Fig. 2b). The landmark, represented by a star of diameter 10 cm, is always situated at a constant distance of 30 cm to the North of the platform, whose diameter is 20 cm. These dimensions have been chosen in order to keep similar ratio of distances as in Pearce et al.’s experimental setting (the platform’s size has been scaled up, as the original size (10 cm) was too small and did not allow the experts to learn the task). The number of possible locations of the platform has been reduced from eight to four, in order to compensate the new size of the platform. As in [13], at the beginning of each trial, the simulated robot is placed at a random position at least 120 cm from the platform. The robot moving speed is 10 cm per timestep, meaning that it requires at least 12 timesteps to reach the platform. If it is not able to reach the platform in 150 timesteps, it is automatically guided to it, as were the actual rats. A positive reward ($R = 1$) is provided when the platform is reached.

We performed three sets of 50 experiments. In the first set, both experts (MBe and MFe) are functional (Control group), in the second set only the MFe is activated (Hippocampal group). For the third set of experiments, only the MBe is activated. This “Striatal group” emulates a striatal lesion not included in the original experiment.

1.3 Data Analysis

Performances of different groups were statistically assessed by comparing their mean escape latency (signed-rank Wilcoxon test for matched-paired samples). Moreover, following Pearce’s analysis, we assess learning differences within a session by comparing the performance on the first and fourth trials using the same test as before. Concerning the differences between both groups (i.e., between the first trials of Control and Hippocampal groups, and between their fourth trials), we use a Mann-Whitney test for non matched-paired samples.

To assess strategy changes during the whole experiment, we compare their selection rates at every first and fourth trials of both *early* (first three sessions) and *late* sessions (last three sessions). The selection rate of each expert is recorded on two squares of 0.4 m², centered on the current and on the




Figure 3: Mean escape latencies measured during the first and the fourth trial of each session. (a) Results of the original experiment with rats, reproduced from [13]. (b) Hippocampal group (MFe only) versus Control group (provided with both a MFe and a MBe). (c) Striatal group (MBe only) versus Control group. See text for explanations.

previous platform positions and is computed as the number of times the robot chooses one strategy over the total number of times it goes inside each of these regions.

In order to estimate strategy changes within a trial, the selected strategy at each timestep is recorded. Since trajectories have different lengths, they are first normalized in 10 bins, then we compute the selection rate on each of these bins. The navigational maps of both experts, the preferred orientation at each location of the environment, are also provided in order to illustrate changes in the expert’s learning across trials or sessions.

Finally, we evaluate the influence of the robot’s behaviour when controlled by an expert on the learning of another expert. The averaged heading error across the sessions is computed for the three groups. This heading error corresponds to the difference between the actual direction of movement proposed by the expert and the “ideal” direction, pointing to the current platform location (the heading error will be zero when the robot points towards the platform; an error of one means that the robot moves in the opposite direction). This error is computed in the neighbourhood of the current platform –on a square of 0.4 m^2 – in order to take values from places that are sufficiently explored by the robot. The influence between experts can be assessed by measuring whether the heading error for one of the strategies decreases as a result of the execution of the other strategy.

2 Results

2.1 Learning Across and Within Session

Our model qualitatively reproduces the results obtained in animals (Fig. 3a). As shown in Fig. 3b, both Control and Hippocampal groups are able to learn the task, i.e., their escape latencies decrease with training. Moreover, the performance of the Control group improves within each session, as there is a significant decrement of the escape latency between the first and fourth trials ($p < 0.001$). Finally, as it was the case with rats, escape latencies of the Hippocampal group in the first trial are smaller than the Control group ($p < 0.001$). Concerning the Striatal group, Fig. 3c shows a significant improvement within session for this group, but no learning is achieved across sessions, suggesting a key role of the MFe in the performance improvement across sessions of the Control group.

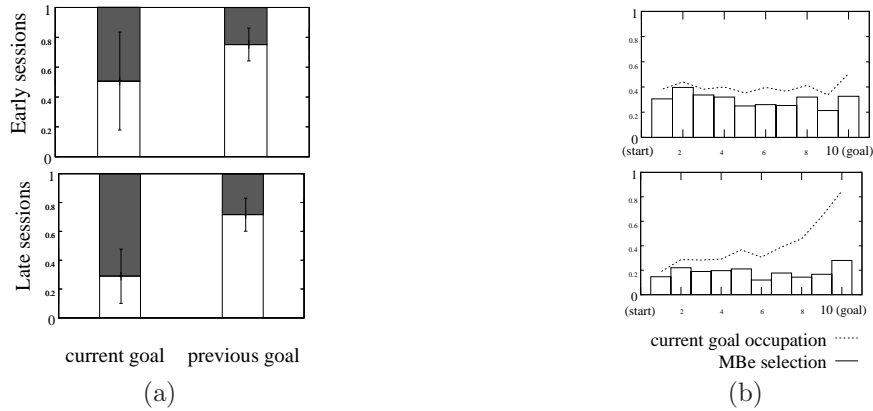


Figure 4: First trials: (a) Selection rates of MBe (*empty boxes*) and MFe (*full boxes*) near the current and the previous platform in early (*top*) and late sessions (*bottom*) (b) Selection rates of MBe and current goal occupation within trial in early (*top*) and late (*bottom*) sessions.

2.2 Role of Interactions Between the MFe and the MBe in the Control Group

First trials: Increase of MFe Selection Across Sessions and Competition Between the MFe and the MBe Within Trials

In the first trial of every session, the platform is relocated, so as the strategy learned by the MBe in the previous session is not relevant anymore. Accordingly, the selection of the MFe expert near the current platform location increases from the early to the late sessions ($p < 0.05$), strongly suggesting a role of the MFe in the latency decrease across sessions that occurs in the Control group (Fig. 4a). Fig. 4a also shows that the MBe is often selected near the previous platform location, suggesting the existence of a competition between both experts. MBe preference does not change within a trajectory and is in average less selected than the MFe (Fig. 4b).

The trajectories (Fig. 5a and 5b) confirm the existence of a competition: the MBe tends to lead the robot to the previous location of the platform – as shown in the navigational maps of this expert (Fig. 5c and 5d) – whereas the MFe has learned to orient the robot towards the appropriate direction, i.e., at the South of the landmark (Fig. 5e and 5f). This result is consistent with the explanation provided by Pearce and colleagues and shows that the competition between the MBe and the MFe is mainly responsible for the poor performances of the Control group in the first trials.

Fourth trials: Cooperation Between the MFe and the MBe Within Trials

At the end of a session, the platform location remained stable during four trials, allowing the MBe to learn its location. According to Pearce’s hypothesis, rats behaviour depends mainly on the *map-based* strategy (involving the hippocampus) that has learned the platform location for this session. However, simulation results show that the Striatal group –controlled by the MBe only– is outperformed by both the Hippocampal and the Control groups, despite a high improvement within sessions (c.f. Fig. 3c). This suggests that the performance of the Control group on the fourth trials cannot be explained exclusively by the MBe expert. Indeed, although this expert leads the agent towards the current goal position, it also leads to the previous goal location as illustrated by its selection rate on both sites (Fig. 6a). In addition, selection rates within a trajectory show a strategy change from the MFe –which is preferred at the beginning of a trial– towards a preference for the MBe at the end of the trajectory (Fig. 6b).

This sequence is visible in typical trajectories (Fig. 7a and 7b). The navigational maps of each expert reveal that the MFe orients the robot towards the South of the landmark (Fig. 7e and 7f), whereas the MBe leads it on the precise location of the platform, only when it is at its vicinity (Fig. 7c and 7d).

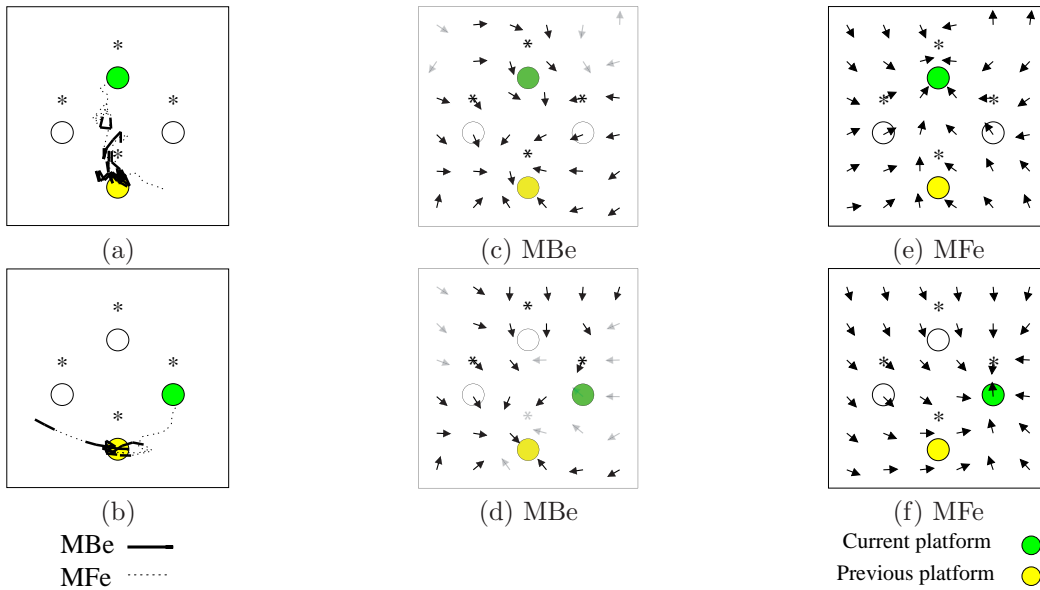


Figure 5: First trials: (a) Trajectory of the robot for the 3rd session (b) Trajectory of the robot for the 9th session. (c) Navigational map of the MBe for the 3rd session (d) Navigational map of the MBe for the 9th session (e) Navigational map of the MFe for the 3rd session (f) Navigational map of the MFe for the 9th session.

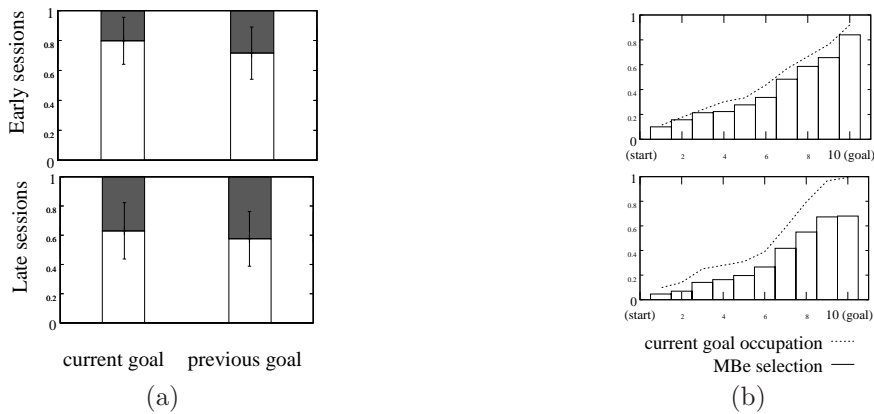


Figure 6: Fourth trials: (a) Selection rates of MBe (*empty boxes*) and MFe (*full boxes*) near the current and the previous platform in early (*top*) and late (*bottom*) sessions. (b) Selection rates of MBe and current goal occupation within trial in early (*top*) and late (*bottom*) sessions.

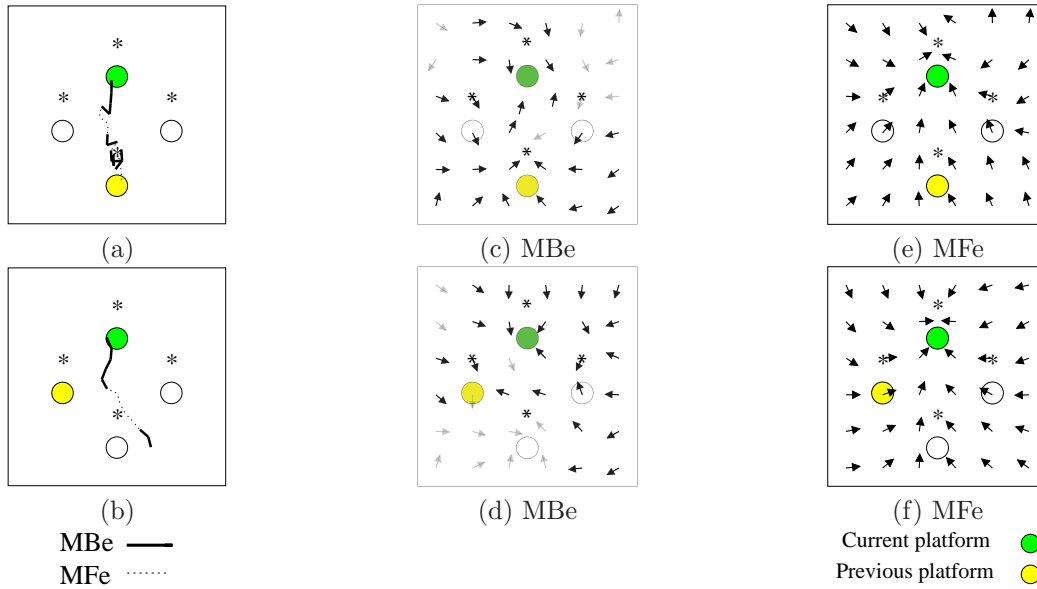


Figure 7: Fourth trials: (a) Trajectory of the robot for the 3rd session (b) Trajectory of the robot for the 11th session. (c) Navigational map of the MBe for the 3rd session (d) Navigational map of the MBe for the 11th session (e) Navigational map of the MFe for the 3rd session (f) Navigational map of the MFe for the 11th session.

This suggests that the experts are cooperating by both adequately participating to the resolution of the task, depending on their reliability at a specific point of the journey. Our findings –pointing out a cooperative interaction at the end of each session– extend Pearce’s hypothesis of MBe dominance in behaviour control.

2.3 Interactions Between MFe and MBe

In simulations of both the Hippocampal and Striatal groups, the inactivation of one expert only prevented it to control the robot’s behaviour, but not to learn. We can thus analyze how the interactions influence the learning of each strategy.

First, looking at the accuracy of both experts in the neighbourhood of the current platform (Fig. 8), we observe that when the robot behavior is driven by the MBe (i.e. Striatal group), the performance of the MFe decreases (Fig. 8c). Second, we observe that MBe performs better in the Control group (Fig. 8a) than in Striatal and Hippocampal groups (Fig. 8b and c), presumably because of the influence of the efficient learning of the MFe (i.e., cooperative interactions).

The navigational maps of MFe are similar –i.e., pointing to the South of the landmark– for the Control, Striatal and Hippocampal groups, despite the difference of performance observed above (Fig. 9c, d and 7f). In contrast, those of the MBe are different: in the Striatal group (Fig. 9a), the MBe is less attracted by the previous platform location than in the Control group (Fig. 7d), whereas it is attracted by the four possible locations in the Hippocampal group (Fig. 9b). The MBe is able to reach the every possible platform location, but only when it is in its vicinity. This suggests that a cooperation between the MFe –leading the robot to the neighbourhood of the current platform– and the MBe –finding the precise location once the robot is there– would perform well and enhance the performance of the robot. Therefore, this particular configuration of the MBe is impaired in the case where the MBe should perform the trajectory alone, but enhanced in the case of a cooperation with the MFe.

We observe that the behavior of the robot when controlled by the MFe, strongly influence the




Figure 8: Average heading error near the current platform for the three groups. Zero means the expert is pointing to the platform, one means a difference of π . (a) Results in the Control group (MBe and MFe activated) (b) Hippocampal group (MFe only) (c) Striatal group (MBe only).

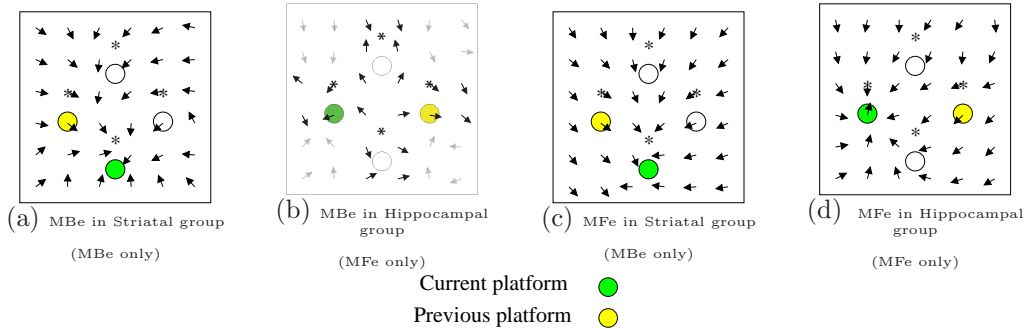


Figure 9: (a) Navigational map of the MBe in the Striatal group in the last session (fourth trial): there are no center of attractions in other platform locations than the current and the previous ones (b) Navigational map of the MBe in the Hippocampal group in the last session (fourth trial): the MBe has learned the ways to go the four possible locations of the platform (c) Navigational map of the MFe in the Striatal group in the last session (fourth trial): it has learned the same kind of map as in the Hippocampal and the Control groups (d) Navigational map of the MFe in the Hippocampal group in the last session (fourth trial): the learned policy is very close to the one in the Striatal group.

MBe. In contrast, the MBe-based behavior has less influence on the improvement of the MFe strategy. Remarkably, activation of both experts (i.e., Control group) do not impair the simultaneous learning of both strategies and allows the MBe to achieve better performance than when this expert is the only one available.

3 Discussion

3.1 Competition and Cooperation

We have been able to reproduce the behaviour of rats in an experiment designed to study interactions between different spatial learning systems. Our simulation results are consistent with the original hypothesis of competitive interaction between map-based (MB) and map-free (MF) strategies at the start of a session when the location of the hidden cue-marked platform suddenly changes [13]. In addition, our model suggests a cooperative interaction during the learning of current location within a session. In these trials, the MF strategy is preferred at the beginning of the journey when the local cue gives information about the general direction to follow, while the robot gets closer to the goal, the MB strategy provides more accurate information about the real platform location and is chosen more often.

Other experimental studies have reported strategy changes during a journey depending on the available information and the animal previous experience [18, 19]. Hamilton et al. [19] reported a change from a *map-based* to a *taxon* strategy when rats were looking for a visible, stable platform. Contrasting to Pearce et al.’s settings, there are no intra-maze cues, and the authors report that rats first used distal landmarks to find a general direction, then approached the platform using a *taxon* strategy. Both, our results and those of Hamilton follow the same rationale, i.e., rats first choose a general direction of movement, and then choose the strategy that allows them to accurately locate the platform. Rat’s head scanning was analyzed in order to estimate the strategy changes. The same approach could be applied to the animal trajectories in Pearce’s paradigm in order to identify whether the strategy change predicted by our model is confirmed by the rats’ behaviour.

3.2 Synergistic Interactions and Dependence of an Expert on Another One

Changes in the heading error, assessed by the evolution of the error in the different experimental groups, suggest synergistic interactions between the two experts. The MFe orients the robot towards the landmark, and the MBe helps the robot to find the platform in the vicinity of the Landmark. If we define an expert as being dependent on another based on its ability to achieve a task alone, we conclude that MBe is dependent on the MFe, as the MBe does not learn the task across sessions. It should be noticed that an opposite relationship –i.e., MFe depending on MBe– has been reported in different experimental conditions (see [11] for a review).

3.3 Further Work

Despite qualitatively reproducing most of the results reported by Pearce et al. [13], our model differs from animal results since a performance improvement was observed within sessions in the Hippocampal group. This difference seems to be mainly due to the learning process of the MFe in cases where, in the previous session, the robot could reach the platform only by following the landmark (for example, if the platform is at the North, as illustrated in Fig. 10). This impairments can also explain the absence of convergence of both groups in the last session.

In contrast to Pearce’s results, no significant difference is found between the fourth trials of the Control and Hippocampal groups. We impute this to the stochastic selection process –i.e., the probabilities associated with a strategy (see section 1.1.2)– which is sometimes sub-optimal. More generally our results might be improved by the use of a dynamically updated hippocampal map, as well as the use of explicit extra-maze cues on which –according to the authors– both strategies were anchored.

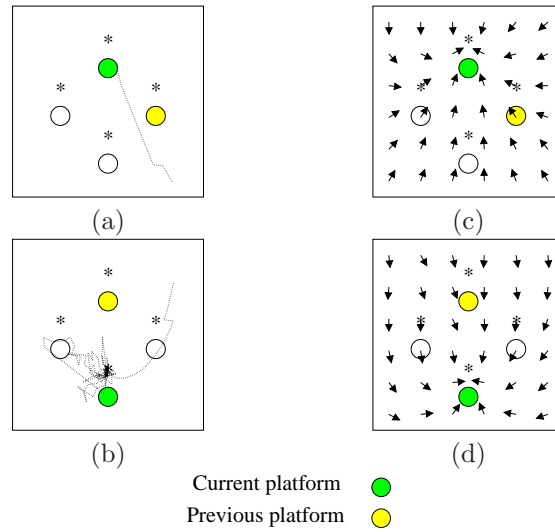


Figure 10: (a) Trajectory at the fourth trial of the 7th session: as the simulated robot mainly went to this platform from the South, direction to the North were reinforced, even at the North of the platform.

(b) Trajectory at the first trial of the 8th session: Starting from North, the robot needs then a longer trial to readjust the direction towards the current platform.

(c) Navigational map of the MFe at the fourth trial of the 7th session : direction to the North were reinforced, even at the North of the platform.

(d) Navigational map of the MFe at the first trial of the 8th session.

In our simulation, these cues were only designed by an absolute reference for the MFe, and an ad hoc cognitive map for the MBe. Finally, models of map-based strategy different than place-response associations, can be taken into account. The place-response strategy currently used in the model associates locations to actions that lead to a single goal location. Therefore, when the platform is relocated, the strategy has to be relearned. An alternative map-based strategy can be proposed such as the relations between different locations are learned irrespectively of the goal location (e.g. a topographical map of the environment). *Planning* strategies can be used to find the new goal location without relearning [3]. The use of computational models of planning (e.g. [20, 21]) as a map-based strategy in our model can yield further insights on the use of spatial information in these types of tasks.

4 Conclusion

What stands out from our results is that our model allowed to analyze the selection changes between both learning systems, while providing information that is not directly accessible in experiments with animals (e.g., strategy selection rate, expert reliability). This information can be used to elaborate predictions, and propose new experiments towards the two-fold goal of further improving our models and expand our knowledge of animal behaviour. It showed also that opposite interactions can happen within a single experiment, and depend mainly on contextual contingencies and practice, as it has been suggested by recent works (e.g., [22, 23]).

Coexistence of several spatial learning systems allows animals to dynamically select which navigation strategy is the most appropriate to achieve their behavioural goals. Furthermore, interaction among these systems may improve the performance, either by speeding learning through collaboration of different strategies, or competitive processes that prevents sub-optimal strategies to be applied. Besides, better understanding of these interactions in animals by use of the modelling approach described in this paper also contributes to the improvement of autonomous robot navigation systems. Indeed, several bioinspired studies began exploring the robotic use of multiple navigation strategies [12, 24, 25, 26], the topic is however far from being fully explored yet.

Acknowledgment

This research was granted by the EC Integrated Project ICEA (Integrating Cognition, Emotion and Autonomy). The authors wish to thank Angelo Arleo, Karim Benchenane, Jean-Arcady Meyer and Denis Sheynikhovich for useful discussions.

References

- [1] Trullier, O., Wiener, S.I., Berthoz, A., Meyer, J.A.: Biologically-based artificial navigation systems : review and prospects. *Progress in Neurobiology* **83**(3) (1997) 271–285
- [2] Filliat, D., Meyer, J.A.: Map-based navigation in mobile robots - i. a review of localisation strategies. *Journal of Cognitive Systems Research* **4**(4) (2003) 243–282
- [3] Meyer, J.A., Filliat, D.: Map-based navigation in mobile robots - ii. a review of map-learning and path-planning strategies. *Journal of Cognitive Systems Research* **4**(4) (2003) 283–317
- [4] Arleo, A., Rondi-Reig, L.: Multimodal sensory integration and concurrent navigation strategies for spatial cognition in real and artificial organisms. *Journal of Integrative Neuroscience* **6** (2007) 327–66

- [5] Packard, M., McGaugh, J.: Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience* **106**(3) (1992) 439–446
- [6] White, N., McDonald, R.: Multiple parallel memory systems in the brain of the rat. *Neurobiology of Learning and Memory* **77** (2002) 125–184
- [7] Kim, J., Baxter, M.: Multiple brain-memory systems: The whole does not equal the sum of its parts. *Trends in Neurosciences* **24**(6) (2001) 324–30
- [8] Poldrack, R., Packard, M.: Competition among multiple memory systems: Converging evidence from animal and human brain studies. *Neuropsychologia* **41**(3) (2003) 245–51
- [9] McIntyre, C., Marriott, L., Gold, P.: Patterns of brain acetylcholine release predict individual differences in preferred learning strategies in rats. *Neurobiology of Learning and Memory* **79**(2) (2003) 177–183
- [10] McDonald, R., Devan, B., Hong, N.: Multiple memory systems: The power of interactions. *Neurobiology of Learning and Memory* **82**(3) (2004) 333–346
- [11] Hartley, T., Burgess, N.: Complementary memory systems: Competition, cooperation and compensation. *Trends in Neurosciences* **28**(4) (2005) 169–170
- [12] Chavarriaga, R., Strosslin, T., Sheynikhovich, D., Gerstner, W.: A computational model of parallel navigation systems in rodents. *Neuroinformatics* **3**(3) (2005) 223–242
- [13] Pearce, J., Roberts, A., Good, M.: Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. *Nature* **396**(6706) (1998) 75–7
- [14] Morris, R.: Spatial localisation does not require the presence of local cues. *Learning and Motivation* **12** (1981) 239–260
- [15] Doeller, C.F., King, J.A., Burgess, N.: Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proceedings of the National Academy of Sciences of the United States of America* **105**(15) (2008) 5915–5920
- [16] Arleo, A., Gerstner, W.: Spatial cognition and neuro-mimetic navigation : A model of hippocampal place cell activity. *Biological Cybernetics* **83**(3) (2000) 287–299
- [17] Strsslin, T., Sheynikhovich, D., Chavarriaga, R., Gerstner, W.: Robust self-localisation and navigation based on hippocampal place cells. *Neural Networks* **18**(9) (2005) 1125–1140
- [18] Devan, B., White, N.: Parallel information processing in the dorsal striatum : Relation to hippocampal function. *Neural Computation* **19**(7) (1999) 2789–2798
- [19] Hamilton, D., Rosenfelt, C., Whishaw, I.: Sequential control of navigation by locale and taxon cues in the morris water task. *Behavioural Brain Research* **154**(2) (2004) 385–397
- [20] Martinet, L.E., Passot, J.B., Fouque, B., Meyer, J.A., Arleo, A.: Map-based spatial navigation: A cortical column model for action planning. In: *Spatial Cognition* (in press). (2008)
- [21] Filliat, D., Meyer, J.: Global localization and topological map-learning for robot navigation. *Proceedings of the seventh international conference on simulation of adaptive behavior on From animals to animats table of contents* (2002) 131–140
- [22] Pych, J., Chang, Q., Colon-Rivera, C., Haag, R., Gold, P.: Acetylcholine release in the hippocampus and striatum during place and response training. *Learning & Memory* **12**(6) (2005) 564–572

- [23] Martel, G., Blanchard, J., Mons, N., Gastambide, F., Micheau, J., Guillou, J.: Dynamic interplays between memory systems depend on practice: The hippocampus is not always the first to provide solution. *Neuroscience* **150**(4) (2007) 743–753
- [24] Meyer, J., Guillot, A., Girard, B., Khamassi, M., Pirim, P., Berthoz, A.: The Psikharpax project: Towards building an artificial rat. *Robotics and Autonomous Systems* **50**(4) (2005) 211–223
- [25] Guazzelli, A., Corbacho, F.J., Bota, M., Arbib, M.A. In: *Affordances, motivations, and the world graph theory*. Volume 6. MIT Press, Cambridge, MA, USA (1998) 435–471
- [26] Girard, B., Filliat, D., Meyer, J.A., Berthoz, A., Guillot, A.: Integration of navigation and action selection in a computational model of cortico-basal ganglia-thalamo-cortical loops. *Adaptive Behavior* **13**(2) (2005) 115–130