A MODEL OF PREFRONTAL COLUMNAR ORGANIZATION FOR MULTISCALE SPATIAL PLANNING

Louis-Emmanuel Martinet[1,2] and Angelo Arleo[1]
[1]CNRS - UPMC Univ Paris 6, UMR 7102, F-75005, Paris, France
[2] CNRS - UPMC Univ Paris 6, UMR 7222, F-75005, Paris, France
louis-emmanuel.martinet@upmc.fr

ABSTRACT
An important issue in spatial memory is the learning of abstract representations suitable for navigation planning. To address this problem, we have already developed a planning system inspired by the columnar organization of the mammalian cortex [1]. This model provides a neuromimetic architecture capable of learning topological spatial representations and planning goal-directed actions. The work presented here deals with the ability to encode multiscale representations of the environment, in order to solve large maze tasks. This is shown by validating the model on a multiscale version of the Tolman & Honzik's detour task [2]. Simulation results demonstrate that the performances of the planning system are invariant with respect to the scale of the maze. A series of statistical analyses is provided to characterise the neural activities subserving spatial planning. It is shown that the structural properties of the environment are encoded by the discharges of the location-selective neurones of the model. Complementing this purely spatial coding, the activity of another class of neurones in the model integrates both spatial and reward-dependent information suitable for navigation planning.

KEY WORDS
hippocampus, neuromimetic modelling, prefrontal cortex, spatial navigation planning.

1. Introduction

According to experimental evidence, spatial navigation planning is likely to rely on a distributed neural network spanning limbic and cortical brain structures. This network includes the hippocampus, which mediates spatial representations, and neocortical structures, such as the prefrontal cortex, which participate to the elaboration of abstract contextual descriptions (e.g., accounting for motivation-dependent memories and action cost/risk constraints). We have built a columnar cortical model [1] to provide a neuromimetic architecture suitable for spatial navigation planning, and based on the interaction between the hippocampus and the prefrontal cortex. The planning process is based on an activation-diffusion mechanism, propagating reward-related information from the goal position through the entire topological network [1]. This propagation enables the system to plan action sequences (i.e., trajectories) from the current position towards the goal. The activation-diffusion mechanism produces an exponential decrease of the intensity of the goal signal that propagates along the topological graph [1]. To prevent the system from planning failures in the presence of large scale environments (where locations exist in which the propagating signal is likely to reach the noise level and decision making becomes random) the current model also learns topological representations whose resolution is adapted to the complexity of the environment (to account for structural regularities as corridors). A review of theoretical discussions on hierarchical cognitive maps can be found in [3]. McNamara et al. [4] have suggested that human can solve difficult spatial problems by building a hierarchical cognitive map including multiple representations of the same environment at different spatial scales. Moreover, animals may be able to chunk available information and build hierarchical representations to facilitate learning [5-9]. Recently, multiscale spatial representations have been identified at the neural level. For example in the entorhinal cortex, Hafting et al. [10] have shown that grid cells have spatial fields forming a grid of variable resolution. Kjelstrup et al. [11] have provided neural recordings of place cell activities in a large maze, supporting the same multiscale coding property in the hippocampus. In our model, we suggest that this kind of multiscale representations should also be found in the neocortical areas such as the prefrontal cortex, commonly associated with high-level cognitive processes.

2. Methods

2.1 Topological Map Learning and Action Planning with a Column Model

Existing cortical column models (from earlier to most recent ones, e.g. [12-14]) focus on either the cytoarchitecture of the column or the functional aspect of columnar computation. Our model lies between these two extremes, i.e. it attempts to relate the columnar organization to the behavioral responses based on a bioinspired (highly simplified) neural network model. The basic components of our column model and its learning principles have been previously presented [1]. To summarise, an unsupervised learning scheme is employed to make each column encode a specific spatial location $s \in S$. Within a column, a set of minicolumns are selective for all the state-action pairs

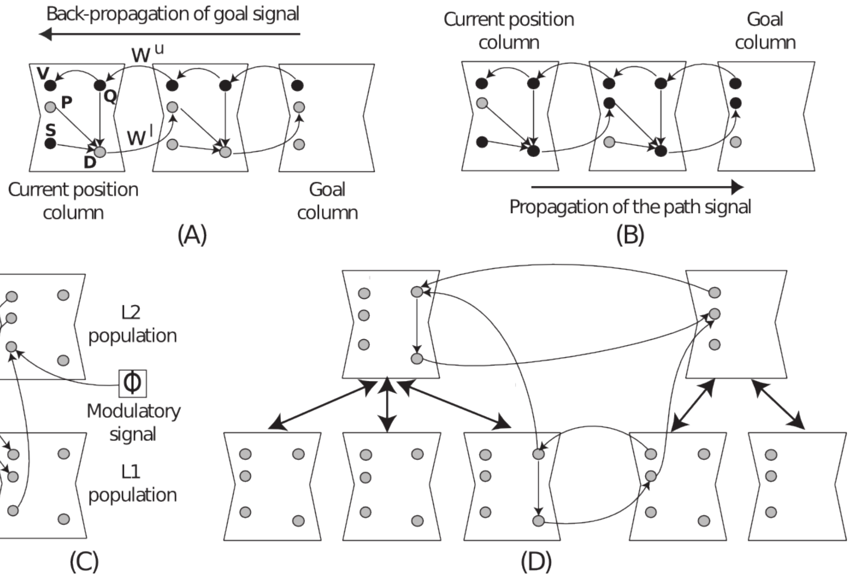

Figure 1: **(A-B)** The cortical model and the implementation of the activation-diffusion process. Columns consist of three supragranular layer units (S, P, V) and a set of minicolumns containing a supragranular (Q) and an infragranular (D) layer unit. Black neurones are firing whereas gray ones are silent. **(A)** back-propagation of the motivational signal through the network of Q and V neurones. **(B)** forward-propagation of the goal-directed action signal through the P and D neurones. **(C)** Top-down and bottom-up connections between a L_1 column (bottom) and a L_2 column (top). Φ is a modulatory signal indicating variation in the high level context. **(D)** Topological connections are also learnt in the L_2 network (on this picture, connections detailed in (C) are summarized by a bidirectional arrow).

$(s, a_{1\ldots N}) \in S \times A$ experienced by the animat at location s. During navigation planning, all the minicolumns of a column compete with each other to locally infer the most appropriate goal-directed action. Compared to our previous model [1], the columnar structure has been refined in order to provide a better understanding of the dynamics of the planning system and to improve its biological plausibility. In the model presented here (Fig. 1A), a column consists of three computational units S, P and V and a set of minicolumns, each of which consists of two units Q and D. S neurones are meant to encode a compact state-space representation from the location-selective activities of hippocampal place cells [15]. The simulated place cells provide the system with a continuous distributed and redundant allocentric state-space representation S [16-18]. Q and V neurones are responsible for encoding respectively the quality (i.e. the efficiency) of an action given a state and the value of a state regarding its distance to the goal. D neurones integrate spatial and reward-related information to code for the best local decision in their discharges. P neurones are used to propagate the path signal encoding the plan from a given position to the goal. The discharge of these units simulates the mean firing activity of a population of cortical neurones either in supragranular layers II-III (for S, P, V and Q neurones), or in infragranular layers V-VI (for D neurones).

The planning process mediated by the columnar network (see example in Fig. 1A-B) is inspired by Burnod's activation-diffusion mechanism [19]. During trajectory planning, the unit V of the column corresponding to the goal location is activated via a motivational signal. Then, this reward-related activity is back-propagated through the network via the V and Q units (Fig. 1A). Q neurones convey this goal-related information to D units, where it is integrated with the spatial information coming from S and P units. When the back-propagated goal signal reaches the column selective for the current position s, the D unit becomes active and triggers the forward propagation of a goal-directed signal through projections $\mathbf{w}^l$ (Fig. 1B). Notice that each $\mathbf{w}^u$ synapse attenuates the back-propagated goal signal. Thus, the smaller the number of synaptic relays, the stronger the goal signal received by the Q neurones of the column corresponding to the current location s. Since the receptive fields of the model columns are distributed uniformly over the environment by the unsupervised learning scheme [1], the intensity of the goal signal at a location s is roughly proportional to the distance of the target. Thus, goal-related metrical information is encoded implicitly by the network, which is is fundamental in order to select the shortest pathway to the target.

2.2 Encoding Multiscale Spatial Representations

Let us denote population L_1 the previous cortical column population receiving spatial inputs from the hippocampus. A second population L_2 of columns is learnt by the current model to encode a large scale map adapted to the size of the environment. The learning algorithm is based on a measure that can define the boundaries between the high scale states. Here, we use a very simple mechanism suited for mazes with corridors, but the overall principle remains the same. A signal Φ is introduced to encode a change in the egocentric locomotion: Φ = 1 when the animat is going straight and Φ = 0 when it turns. This signal conveys relevant information to extract subpart of corridors in a

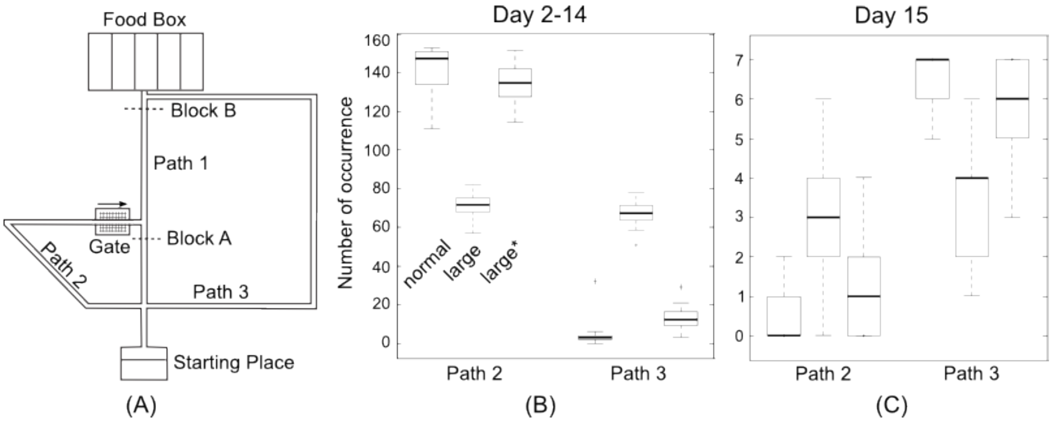

Figure 2: **(A)** Tolman & Honzik's maze (adapted from [2]). The gate near the second intersection prevented rats from going from right to left. **(B-C)** Behavioral results for the Tolman & Honzik's maze represented as the mean number (averaged over 40 animats) of transits through P2 and P3 during Day 2-14 (B) or Day 15 (C). Several sizes of the maze are used: normal and large (four time bigger). In the normal and large conditions, no multiscale learning was used, unlike in the large* condition.

maze. L_2 columns and minicolumns are the same generic computational units as in L_1 network but they are receiving afferents from L_1 columns modulated by the gating signal Φ (Fig. 1C). This ``boundary'' signal introduces a locomotion-dependent bias in the spatial selectivity of S neurones, such that the morphological properties of the environment (e.g., alleys in a maze) can be encoded by the L_2 topological map explicitly. An unsupervised growing network scheme is being employed to recruit L_2 columns similarly to the L_1 population. Additional top-down connections are created from L_2 to L_1 so that the former population can exert a top-down modulation on the P and V neurones of the L_1 population (Fig. 1C), enabling the planning process at the level of L_1 to cope with the decreasing back-propagating signal. This is achieved simply by enhancing the transfer function of P and V units in L_1 with a positive factor. Because the size of high scale states will not be homogeneous as opposed to the state representation in L_1, a more flexible topological learning must be employed to account for the distance between any state and the goal. To solve this issue, two sets of bottom-up weights are used to convey the goal-distance information estimated at the level of the L_1 network by the activity of its Q and P units (Fig. 1D). This input is used to learn the lateral weights $\mathbf{w}^l$ and $\mathbf{w}^u$ in the population L_2, so that the activity of a V unit in L_2 is correctly correlated with the distance of the high scale state to the goal thanks to the information encoded in L_1. In other words, planning computations propagated at the level of the L_1 network are available in the L_2 network which uses them to estimate correct goal-distance information. Thus there is a bi-directional (bottom-up and top-down) flow of information between the two populations of columns of the model, making it possible to encode the environment at multiple scales and to solve large maze planning tasks.

3. Results

3.1 Spatial Behaviour in a Detour Task

In order to validate our multiscale navigation planning system, we chose the classical experimental task proposed by Tolman & Honzik [2], as in our previous work [1]. The main objective of this behavioural protocol was to demonstrate that rodents undergoing a navigation test were able to show some ``insights'', e.g. to predict the outcome of alternative trajectories leading to a goal location in the presence of blocked pathways. The original Tolman & Honzik's maze and protocol are shown in Fig. 2A. Here we extended its principle by using multiple size of the same maze to test the ability of the model to produce multi-scale topological maps and to solve detour tasks in increasingly larger mazes. Two versions of the Tolman & Honzik's maze were thus used: the classical one and a large one which was four times bigger than the original. For their experiments, Tolman & Honzik used 10 rats with no previous training. In our simulations, we examined a set of 40 simulated animats for each experimental condition. In the *classical* and the *large* conditions, the top-down influence of the L_2 population was discarded to show how the size of the maze progressively impaired the performance of animats in the absence of a compensatory neural adaptation. We also ran a set of 40 experiments in the large maze allowing the top-down influence of the high-scale cognitive map over the planning process (*large* condition*). Here we focus on the multiscale aspect of the task, because we have already shown in [1] that the cortical column could reproduce the original results in Tolman & Honzik's normal maze. We assessed the statistical significance of the results by means of an ANOVA analysis (the significant threshold was set at 10^{-2}, i.e. $p < 0.01$ was considered significant).

Day 1. During the first 12 training trials, the animats learnt the topology of the maze and planned their navigation trajectory in the absence of both block A and B. Similar to Tolman & Honzik's findings, our results in all conditions (normal, large and large*) show that the model learnt to select the shortest pathway P1 significantly more frequently than the alternative trajectories P2, P3 (ANOVA, $p < 0.01$ for all mazes).

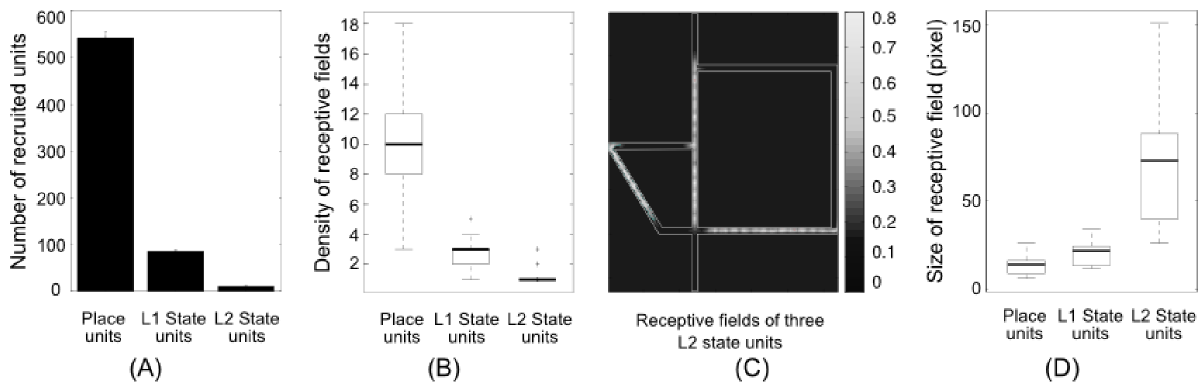

Figure 3: **(A-B)** Spatial properties for HP population and S population of L_1 and L_2 networks: number of active units during the task (A) and spatial density of the receptive fields (B). **(C)** Samples of receptive fields of three units from L_2 population. **(D)** Size of receptive fields of single neurone responses for HP cells, L_1 cortical S units and L_2 cortical S units.

However, for the large condition (but not for large*), the size of this maze began to induce few mistakes, as indicated by a lower median value of Path 1 selection.

Days 2-14. During this training phase (consisting of 156 trials), a block was introduced at location A, which forced the animats to update their topological maps dynamically, and to plan a *detour* to the goal. P1 was ignored in this analysis (similarly to Tolman & Honzik's analysis) because blocked. The results reported by Tolman & Honzik provided strong evidence for a preference for the shortest *detour* path P2. Consistently, in our simulations (Fig. 2B) we observed a significantly larger number of transits through P2 compared to P3 for normal and large* cases (ANOVA, $p < 0.01$), but this was hardly significant for the large condition with a mean number of selected P3 very closed to P2 (ANOVA, $p < 0.01$). This low performance was very closed to the behavior of an animat turning randomly toward Path 2 or Path 3.

Day 15. Seven probe trials were performed during the 15th day of the simulated protocol, by removing the block A and adding a new block at location B. This manipulation aimed at testing the ``insight'' working hypothesis: after a first run through the shortest path P1 and after having encountered the unexpected block B, will animats try P2 (wrong behaviour) or will they go directly through P3 (correct behaviour)? According to Tolman & Honzik's results, rats behaved as predicted by the insight hypothesis, i.e. they tended to select the longer but effective P3. Our probe test simulation results are shown in Fig. 2C. Similar to rats, the animats exhibited a significant preference for P3 compared to P2 (ANOVA, $p < 0.01$) for normal and large* cases. However this probe test was a failure for the large condition, where the number of P3 choices was not significantly different from the number of P2 choices (ANOVA, $p < 0.68$). Taken together, these results clearly show an impaired performance proportional to the size of the maze, which can be overcome thanks to an adaptive multiscale representation fitting the structure of the maze and providing a top-down modulation of the activation-diffusion mechanism.

3.2 Analyses of Neural Activities

One aim of the model is to provide a functional framework for some of the neuronal responses observed in the prefrontal cortex during performance of spatial memory tasks in rats. We demonstrate how the activity of simulated cortical units can be interpreted as elements of a functional circuit which guide the actions of an agent on the basis of delayed reward. A series of analyses, partially based on the same theoretical tools as in [1], was done to characterise the neural activities subserving the behavioural responses of the system. The set of stimuli S consisted of the places visited by the animat. For the analyses, the continuous two-dimensional input space was discretized, with each location $s \in S$ defined as a 5 x 5 cm square region of the environment. Neural activities were recorded from three location-selective populations during the large Tolman & Honzik's task: HP cells and S units from L_1 and L_2 population. In our previous work, we have shown that the cortical column model (i.e., the L_1 population) was able to build a more compact spatial representation storing the main part of the spatial information [1]. Here we focus on the spatial properties of the L_2 population compared to L_1 and HP neurones: *(i)* fewer units of L_2 are necessary to represent the same environment (Fig. 3A, ANOVA, $p < 0.01$), *(ii)* their receptive fields are less redundant (Fig. 3B, ANOVA, $p < 0.01$) [1,20], i.e. fewer neurones of population L_2 were, on average, responding to a given stimulus s simultaneously. These results suggest that the L_2 cortical column network was able to provide an even sparser state-space population coding than L_1 population. In a second series of analyses, we focused on the activity of single cells, and we recorded the receptive fields of the three types of units. What is mostly remarkable is the firing properties of L_2 state neurones (Fig. 3C): after learning, the activity of these units capture some structural properties of the environment (i.e., corridors organization). A quantitative analysis was performed: the mean size of place fields (computed as the number of contiguous pixels with the firing rate above the grand mean rate plus the standard deviation [21]) was indeed significantly bigger than for L_2 units (Fig. 3D, ANOVA, $p < 0.01$). These results receive support from

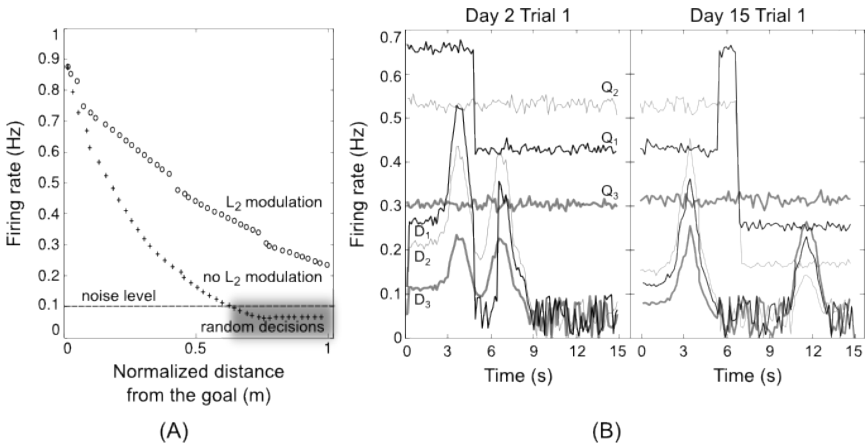

Figure 4: **(A)** Effect of the top-down modulation exert by population L_2 V units over the propagating activity at the level of L_1 V unit population. **(B)** Activities of 3 pairs of (Q,D) units belonging to the same minicolumn for two phase of the protocol where a block is introduced in the maze.

experimental data by [21]: these authors measured the field size of place cells and prefrontal neurones in rats solving a navigational task, showing that the former were significantly smaller than the latter. We also used an information theoretic analysis [22]: the *mutual information* MI(S;R) between neural responses R and spatial locations S allowed us to quantify the spatial information content of a neural code, i.e. how much could be learnt about the animat's position s by observing the neural responses r. It was evaluated for single units as well as for a whole population of neurones (in that case, r was a vector of firing rates), and the ratio between these two values was used to assess the level of sparseness of spatial information. The results of our information theoretic analysis are consistent with the properties described in the previous paragraph). Indeed, L_2 state units responding to a broader range of spatial stimuli, their single neurone mutual information is much higher (ANOVA, $p < 0.01$). The spatial mutual information computed for the whole population of place cells, L_1 and L_2 state units demonstrates a larger information content for the HP population (ANOVA, $p < 0.01$). This indicates that, for the binning procedure applied in this analysis, the place cell population is far more precise to encode a position (because of its high redundancy). In comparison, state neurones in L_2 population encode a very coarse spatial information. This is coherent with our initial goal of building a more compact representation accounting only for the main properties of the environment (here the corridors). Finally, when computing the information sparseness (i.e. the ratio between population information and the sum of single cell information), it appears that the information content was more redundant for place cells (ANOVA, $p < 0.01$), meaning that many of them encoded the same information. Although loosing a part of the population spatial information, the cortical population achieved a better coding scheme, maximizing the coding role of each units, particularly for the L_2 population.

Our second objective here is to show how information relevant for planning are encoded in the neural network. It is first necessary to demonstrate that the V population of the cortical model encodes a measure of distance to the goal. This property is shown on the Fig. 4A, with or without the effect of the high-scale cortical population. This study observed the effect of the top-down modulation exert by population L_2 V units over the propagating activity at the level of L_1. Indeed, we remind that one motivation for this extension of the model was the possibility to deal with large environments. We have shown behaviourally that the model was able to adapt to them. Fig. 4A is a direct evidence of the neural effect of this top-down modulation. Without any modulation, the strength of V units discharge fall exponentially with the distance of the column from the goal position. At a given point, this fast decreasing propagating activity will reach the neural noise level. From that point, only random decisions will be made because there will not be any correlation left between the firing activity and the real distance to the goal (e.g., the low performance on Day 15 for the *large* condition). When a top-down modulation is present, the decreasing effect becomes piecewise linear, each subpart corresponding to a high scale zone encoded by a L_2 column. Reward-related V units and location-selective S units convey their information into the D neurones which integrate them into activities reflecting the selection of action (Fig. 4B). We remind that each minicolumn of the model is supposed to encode a specific state-action pair *(s,a)*. As such, Q units encode the distance to the goal if a is selected at s, and D units integrate spatial information indicating the current position with this reward information. It can be seen on Fig. 4B at $t = 6s$ of the Day 2 Trial 1 that the animat has updated its connectivity in the cortical network to represent the presence of the block A. Thus, the previous best choice Path 1, represented by the best pair (Q_1, D_1) at $t \sim 4s$ is not correct anymore at $t \sim 7s$: Path 2 is now the best alternative as shown by the best pair (Q_2, D_2). The same mechanism occurs on Day 15 Trial 1, with Path 3 represented by (Q_3, D_3) becoming the best choice. Taken together, these analyses demonstrate how the patterns of activity formed by V, Q, D & P neurones can encode relevant task information such as distance-to-goal and best action, leading to an efficient decision-making. We suggest that these patterns could be related to Miller and Cohen's proposal [23]: cognitive control stems from the active maintenance of patterns of activity in the prefrontal cortex that represent goals and the means to achieve them.

4. Conclusion

The model presented here addresses how spatial planning behavior can be mediated by populations of neurons organized in columnar structures. It enables the encoding of cognitive maps whose resolution fits the structure of the environment (e.g., corridors). As a consequence, the model is provided with a better adaptability in large mazes (e.g., in the presence of a maze four times larger than the original Tolman & Honzik's one), thanks to a top-down modulation regulating the activation-diffusion process. It should be noted that encoding multiscale maps is not the only solution to solve behavioral tasks in large mazes. However this approach is useful to address the issue of learning multiscale spatial representations, as found in the brain (e.g., [10,11]). Moreover, the model unravels the possible links between the single unit level and the behavioural level relevant to the learning of the task (e.g., to the selection of the shortest path to the reward). Our neural response analysis suggests how the interplay between the simulated hippocampus and prefrontal cortex can yield to the encoding of manifold information pertinent to the spatial planning function, including for example distance-to-goal correlates. The model is currently being validated by comparing simulated neural response patterns against electrophysiological recordings from the prefrontal cortex of freely moving rats [24]. This comparative study aims at providing new insights on the interaction between the hippocampus and the prefrontal cortex. In addition, an ongoing work in coordination with experimentalists [24] attempts to study the learning processes related to spatial memory, such as declarative memory consolidation occurring during sleep. This will possibly lead to testable predictions about the formation of memory traces relevant to spatial behaviour.

References

[1] L.E. Martinet, J.B. Passot, B. Fouque, J.A. Meyer & A. Arleo, Map-based spatial navigation: A cortical column model for action planning, *LNAI, 5248,* 2008, 39-55.
[2] E.C. Tolman & C.H. Honzik, "Insight" in rats, *Univ. Calif. Publ. Psychol., 4*(14), 1930, 215-232.
[3] N. Schmajuk & H. Voicu, Exploration and navigation using hierarchical cognitive maps. In: *Animal Spatial Cognition: Comparative, Neural, and Computational Approaches*. (2006)
[4] T.P. McNamara, J.K. Hardy & S.C. Hirtle, Subjective hierarchies in spatial memory, *J. Exp. Psychol. Learn. Mem. Cogn., 15*(2), 1989, 211-227.
[5] N.L. Dallal & W.H. Meck, Hierarchical structures: chunking by food type facilitates spatial memory, *J. Exp. Psychol. Learn. Mem. Cogn., 16*(1), 1990, 69-84.
[6] S.B. Fountain & J.D. Rowan, Coding of hierarchical versus linear pattern structure in rats and humans, *J. Exp Psychol. Learn. Mem. Cogn., 21*(3), 1995, 187-202.
[7] T. Macuda & W.A. Roberts, Further evidence for hierarchical chunking in rat spatial memory, *J. Exp. Psychol. Learn. Mem. Cogn., 21*(1), 1995, 20-32.
[8] W.H. Meck & C.L. Williams, Perinatal choline supplementation increases the threshold for chunking in spatial memory, *Neuroreport, 8*(14), 1997, 3053-3059.
[9] W. Roberts, Spatial memory in the rat on a hierarchical maze, *Learn. Motiv., 10,* 1979, 117-140.
[10] T. Hafting, M. Fyhn, S. Molden, M.B. Moser & E.I. Moser, Microstructure of a spatial map in the entorhinal cortex, *Nature, 436*(7052), 2005, 801-806.
[11] K.B. Kjelstrup, T. Solstad, V.H. Brun, T. Hafting, S. Leutgeb, M.P. Witter, E.I. Moser & M.B. Moser, Finite scale of spatial representation in the hippocampus, *Science, 321,* 2008, 140-143.
[12] J. Szentágothai, The 'module-concept' in cerebral cortex architecture, *Brain Res., 95,* 1975, 475-496.
[13] J.C. Eccles, The modular operation of the cerebral neocortex considered as the material basis of mental events, *Neurosci., 6,* 1981, 1839-1856.
[14] H. Markram, The blue brain project, *Nat. Rev. Neurosci., 7,* 2006, 153-160.
[15] J. O'Keefe & L. Nadel, The Hippocampus as a Cognitive Map, *Oxford Univ. Press.,* 1978.
[16] A. Arleo & W. Gerstner, Spatial orientation in navigating agents: modeling head-direction cells, *Neurocomput., 38*(40), 2001, 1059-1065.
[17] A. Arleo, F. Smeraldi & W. Gerstner, Cognitive navigation based on nonuniform gabor space sampling, unsupervised growing networks, and reinforcement learning, *IEEE Trans Neural Netw, 15*(3), 2004, 639-51
[18] D. Sheynikhovich, R. Chavarriaga, T.S. Strosslin, A. Arleo & W. Gerstner, Is there a geometric module for spatial orientation? insights from a rodent navigation model, *Psychol. Rev., 116,* 2009, 540-566.
[19] Y. Burnod, Organizational levels of the cerebral cortex: an integrated model, *Acta Biotheoretica, 39,* 1991, 351-361.
[20] B. Willmore & D.J. Tolhurst, Characterizing the sparseness of neural codes, *Netw.: Comput. Neural Syst., 12,* 2001, 255-270.
[21] V. Hok, E. Save, P.P. Lenck-Santini & B. Poucet, Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex, *Proc. Natl. Acad. Sci. USA., 102*(12), 2005, 4602-4607.
[22] W. Bialek, F. Rieke, R. de Ruyter van Steveninck & D. Warland, Reading a neural code, *Science, 252*(5014), 1991, 1854-1857.
[23] E.K. Miller & J.D. Cohen, An integrative theory of prefrontal function, *Annu. Rev. Neurosci., 24*(1), 2001, 167-202.
[24] A. Peyrache, M. Khamassi, K. Benchenane, S.I. Wiener, F.P. Battaglia, Replay of rule learning related neural patterns in the prefrontal cortex during sleep, *Nat. Neurosci., 12,* 2009, 919-926.