

Exploratory Behaviour Depends on Multisensory Integration during Spatial Learning

Denis Sheynikhovich^{1,*}, Félix Grèzes¹, Jean-Rémi King², and Angelo Arleo¹

¹ Adaptive NeuroComputation Group, Unit of Neurobiology of Adaptive Processes, UMR 7102, CNRS — UPMC Univ P6, Paris, 75005, France

`denis.sheynikhovich@upmc.fr`

² CEA, Division of Life Sciences, Institute of Bioimaging, Neurospin, Gif sur Yvette, 91191 France

Abstract. Active exploration is a necessary component of a putative spatial representation system in the mammalian brain. We address the problem of how spatial exploratory behaviour is generated in rodents by combining an artificial neural network model of place coding with a multiobjective evolutionary algorithm that tunes the model parameters so as to maximise the efficiency of environment exploration. A central property of the spatial representation model is an online calibration between external visual cues and path integration, a widely accepted concept in theoretical accounts of spatial learning in animals. We find that the artificially evolved exploration model leads to recurrent patterns of exploratory behaviour in a way observed in experimental studies of spatial exploration in rodents. Our results provide a link between the functional organisation of the biological spatial learning network and the observed high-level patterns of exploratory behaviour.

Keywords: Computational neuroscience, evolutionary robotics, artificial neural networks, spatial cognition, multisensory integration, exploratory behaviour, hippocampal place cells, path integration.

1 Introduction

Exploratory behaviour is a necessary prerequisite of a putative spatial representation system in the mammalian brain [10]. Behavioural studies have shown that in the course of environment exploration, various species have an important predisposition to systematically come back to a well known location, termed *home base*, resulting in looping, or round-trip, behaviour [2,5]. On the neural level, active exploration has been proposed to subserve the construction of a mental representation of space, thought to be learnt by a network of brain structures mediating spatial cognition [10]. According to the cognitive map theory, pyramidal cells in the hippocampal formation (termed *place cells*) code for spatial locations in an environment [10]. Extensive theoretical and experimental studies of the properties and origins of place cell activity showed that their location

* Corresponding author.

selectivity results from multisensory integration of allothetic (i.e. related to the environmental landmark cues) and idiothetic (i.e. related to self-motion cues, or path integration) information [8,1]. Despite the generally implied link between exploratory behaviour and the construction of mental maps [10], the functional relation between specific exploratory patterns and the spatial representation system has not been proposed so far. In this study we test a novel hypothesis that generation of exploratory behaviour is influenced by the functional organisation of the underlying spatial representation network. In particular, we argue that the necessity of keeping the allothetic and idiothetic components of the spatial map coherent in time, leads to the observed round-trip exploratory patterns.

We use a neural network model of hippocampal place cells, which is a simplified version of our previous model that included realistic visual input, grid-cell and place-cell networks [11]. The central property of the used spatial learning model is that a combination of external and self-motion cues is a necessary condition to build robust place field representations. We adopt multiobjective evolutionary algorithms (EA) [7,9] in order to optimally tune the parameters of the model so as to ensure efficient exploration of a novel environment and study the neural properties of the learnt spatial representation in evolved simulated animals. Our results show that, given our model of spatial representation, the round-trip behaviour is optimal in terms of minimisation of the self-localisation error and maximisation of exploration rate (i.e. explored area per unit time). These results link exploratory patterns observed in animals with cognitive mapping theories of hippocampal neural networks. In addition, they may provide an insight into the biological solution for a well known problem in autonomous robotics, that of simultaneous localisation and mapping (SLAM) [3].

2 Methods

2.1 Spatial Behaviour and Hippocampal Place Coding Model

In our simplified model of spatial behaviour, a simulated mouse explores a square environment 1.6×1.6 m in time steps of $\Delta t = 0.125$ s, with a constant speed $v = 16$ cm/s. At each time step, the motor command for the next movement consists of a fixed displacement $\Delta s = v \cdot \Delta t$ and a rotation $\Delta\phi$ given by the exploration controller (see below).

The spatial representation model constructs a representation of space, given information about external landmarks and noisy estimations of self-motion inputs. While the fixed landmarks represent information about the *true position* of the animal in space, tracked by vector \mathbf{s} , an integration of noisy self-motion inputs gives rise to the *perceived position* at time t , $\mathbf{p}^t = (x^t, y^t)$, where:

$$x^t = x^{t-1} + \Delta\hat{s} \cos(\phi^t) \quad (1)$$

$$y^t = y^{t-1} + \Delta\hat{s} \sin(\phi^t) \quad (2)$$

$$\phi^t = \phi^{t-1} + \Delta\hat{\phi} \quad (3)$$

The noisy estimates of displacement and rotation since the previous time step are $\Delta\hat{s} = \Delta s + \xi$ and $\Delta\hat{\phi} = \Delta\phi + \zeta$, respectively, where ξ and ζ are zero-mean

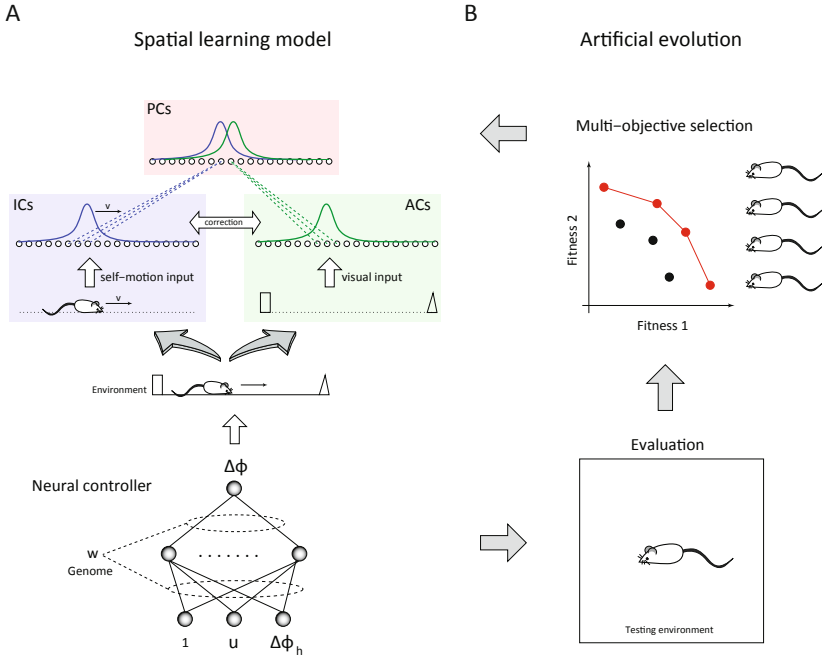


Fig. 1. *The combined modelling and artificial evolution framework.* (A) The spatial learning model (top) consists of idiothetic cells (ICs), allothetic cells (ACs) and place cells (PCs). Connection weights (dotted lines) are learnt during exploration, which is under control of the exploration controller (bottom). (B) The parameters of the exploration controller are optimised in the course of artificial evolution.

Gaussian noise variables with the standard deviation for ξ equal to 50% of the true Δs , and that for ζ equal to 0.1 rad. In the following we omit the explicit time dependence for clarity.

At each time step, the encoding of spatial information occurs in three neural populations, termed *idiothetic cells* (ICs), *allothetic cells* (ACs), and *place cells* (PCs), see Fig. 1A. Activity of an IC is given by:

$$r_i^{IC} = \exp\left(-\frac{\|\mathbf{p} - \mathbf{p}_i^{IC}\|^2}{2\sigma_p}\right) \tag{4}$$

where \mathbf{p} is the perceived position of the animal and \mathbf{p}_i^{IC} is the centre of the cell’s Gaussian receptive field with width $\sigma_p = 10$ cm. The biological counterpart of ICs are subpopulations of entorhinal grid cells [6]. Indeed, a thresholded sum of many grid-cell activities gives rise to an approximately Gaussian profile [8,11].

In contrast to ICs, ACs encode an internal representation of fixed external landmarks, and hence their activity is centred on the true position \mathbf{s} of the animal:

$$r_i^{AC} = \exp\left(-\frac{\|\mathbf{s} - \mathbf{p}_i^{AC}\|^2}{2\sigma_s}\right) \tag{5}$$

where \mathbf{p}_i^{AC} is the centre of the AC's receptive field with $\sigma_s = 10$ cm. A new AC is recruited at each time step during exploration, unless more than 10 ACs are highly active (with threshold 0.9).

Finally, PCs that represent hippocampal place cells in the model, are recruited online similarly to ACs. Once recruited, a PC is connected to all ICs and ACs that are highly active at the moment. Connection weights, w_{ij}^{PC} , from input cell (either IC or AC) j to PC i are set equal to the input activity (i.e. learnt by one-shot Hebbian learning). The activity of the PC is calculated as a vector product of input activities and the weight vector, i.e. $r_i^{\text{PC}} = \sum_j w_{ij}^{\text{PC}} r_j$, where j spans all presynaptic ACs and ICs, and r_j is the activity of the corresponding input cell. Weights are normalised so that maximal activity of a PC is 1.

A central property of the model is the interplay between the idiotactic and allothetic sensory signals. More specifically, the activity of ACs is used to correct the noisy self-motion related information provided by ICs. Whenever the number of highly active ACs exceeds 10, the perceived position \mathbf{p} is corrected by the allothetic input:

$$\begin{aligned}\mathbf{p} &= \mathbf{p} + \alpha(\mathbf{p}^{\text{AC}} - \mathbf{p}) \\ \phi &= \phi + \alpha(\phi^{\text{AC}} - \phi)\end{aligned}\quad (6)$$

where $\mathbf{p}^{\text{AC}} = (x^{\text{AC}}, y^{\text{AC}})$ is the centre of mass of the population of ACs:

$$x^{\text{AC}} = \frac{\sum_i r_i^{\text{AC}} x_i}{\sum_i r_i^{\text{AC}}}; \quad y^{\text{AC}} = \frac{\sum_i r_i^{\text{AC}} y_i}{\sum_i r_i^{\text{AC}}}\quad (7)$$

and (x_i, y_i) are the perceived coordinates stored at the moment of each AC recruitment. ϕ^{AC} is assumed to be equal to the true heading direction of the animal, meaning that the error in the integration of head direction is cumulative, but only until the correction by the allothetic input is made. Note that the true position \mathbf{s} is used only to calculate the AC activity — which roughly corresponds to the activation of the same visual receptors when the same visual cue is observed. In all other formulae it is the *perceived* position that is used, which is subject to the cumulative error.

2.2 Control of Active Exploration by Artificial Evolution

In the model, the direction of the next exploratory movement $\Delta\phi$ (see Section 2.1) is the output of an exploration controller (Fig. 1A) represented by a two-layer neural network with sigmoidal units. Apart from the constant bias equal to 1, the controller's inputs are the uncertainty $u(t)$, given by a threshold-linear function of the time since the last calibration:

$$u(t) = \min\left(\frac{t - t^*}{T}, 1\right)\quad (8)$$

reset to 0 upon each recalibration (i.e. the process described by Eqs. 6,7), and the egocentric homing direction $\Delta\phi_h$:

$$\Delta\phi_h = \phi(t) - \arctan(y/x)\quad (9)$$

where $\phi(t)$ is the current heading direction, and (x, y) is the current perceived position vector. Input values are normalised to be in $[-1, 1]$, while the output value $\Delta\phi$ lies in the range $[-\pi, \pi]$.

The weight vector \mathbf{w} of the exploration controller is evolved according to a multiobjective genetic algorithm [7]. The algorithm starts by generating $N_0 = 1000$ simulated mice with random sets of controller weights (i.e. random genomes). During an evaluation phase, each simulated mouse explores a square environment (Fig. 1B), until 500 PCs are recruited or 1000 time steps passed. A home base is simulated by recruitment of 30 ACs and PCs near the point of entry at the start of the phase. During the following selection phase, the most efficient simulated mice are chosen according to a two-step process. First, two fitness values F_1 and F_2 are calculated. F_1 quantifies the accuracy of the spatial code as the negative mean path integration error:

$$F_1 = -\frac{1}{N_{\text{PC}}} \sum_{i=1}^{N_{\text{PC}}} \|\mathbf{s}_i - \mathbf{p}_i\| \quad (10)$$

where N_{PC} is the number of PCs recruited during the exploration run, \mathbf{s}_i and \mathbf{p}_i are, respectively, the true and perceived positions of the simulated mouse when the place cell i was recruited. F_2 measures the exploration rate as the number of visited quadrants (out of a total 32×32 quadrants, arranged in a grid covering the surface of the arena) per time step. Second, a Pareto-efficiency criterion [7] is used to select the best genomes based on the two fitness values. The selected individuals are used, together with random mutations, to produce the next generation consisting of 100 simulated mice. The evaluation-selection-reproduction cycles are repeated until performance convergence occurs. Artificial evolution simulations were performed using the Sferes2 software platform [9].

3 Results

In the course of artificial evolution, the exploration controller was progressively optimised so as to simultaneously increase the spatial accuracy (F_1) and the exploration rate (F_2), with stabilisation of performance starting after approximately 500 generations (Fig. 2A). For further analysis we distinguish two types of simulated mice: ‘random explorers’, which maximise only exploration rate, and ‘map builders’, which maximise both exploration rate and spatial accuracy (simulated mice that maximise only spatial accuracy always stay in the close vicinity of the home base and thus were not analysed further).

As illustrated by the example traces and the corresponding path integration errors, (Fig. 2B), random explorers quickly lost track of their position, in a sharp contrast to map builders, who tracked their true position with a relatively high accuracy. In order to characterise these differences on the level of a single genome, we compared input-output mappings of the exploration controllers of the random explorer and map builder. In Figure 2C, we show a difference between the next direction of movement and the direction towards the home base as a function of

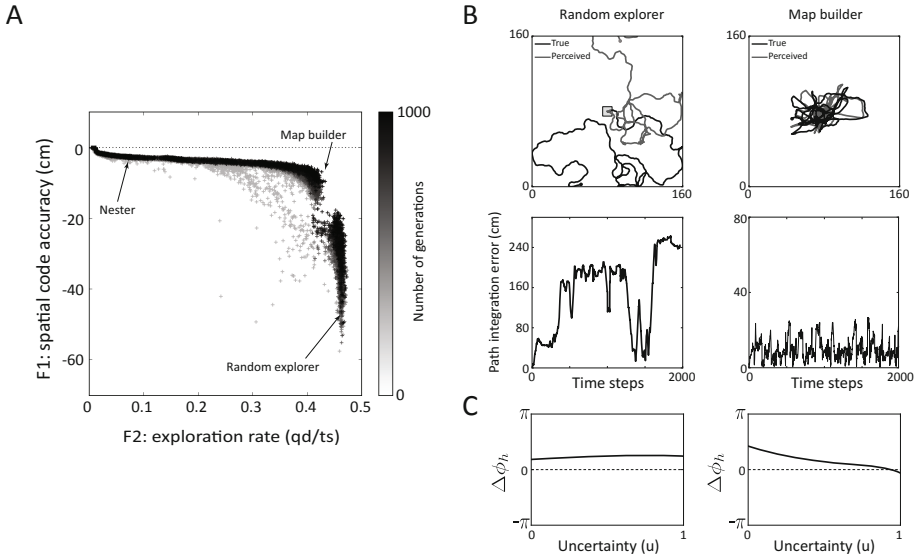


Fig. 2. Evolution of exploratory behaviour. (A) Changes in Pareto front across generations. (B, C) Paths (B, top), path integration errors (B, bottom), and controller input-output mappings (C) for random explorers (left) and map builders (right).

the uncertainty parameter u . Whereas the behaviour of the random explorer is largely independent from the uncertainty, the map builder heads home when the uncertainty is close to its threshold value. These results support our hypothesis that round-trip homing behaviour (as observed in map builders) can result from the need to keep allothetic and idiothetic position estimates coherent in time.

We next checked whether the round-trip behaviour results in a more accurate spatial representation compared to that of random explorers. Indeed, spatial representation on the level of single PCs (Fig. 3A) as well as on the level of the whole PC population (Fig. 3B) are more accurate in map builders than explorers. Multi-peaked receptive fields in random explorers are due to incoherent position estimation by allothetic (ACs) and idiothetic (ICs) cell populations. Finally, we ensured that the differences between random explorers and map builders illustrated in Fig. 2B and Figs. 3A, B are stable across 10 different runs of evolutionary optimisation. We found that (i) map builders have significantly smaller path integration error (Fig. 3C, ANOVA, $F_{1,18} = 72.07$, $P < 0.0001$); (ii) the number of (single-peaked) highly active cells is larger in map builders (Fig. 3D, ANOVA, $F_{1,18} = 190.1$, $P < 0.0001$); (iii) proportion of identified locations, i.e. those with coherent activations of ACs and PCs, is much higher in map builders (Fig. 3E, ANOVA, $F_{1,18} = 152.43$, $P < 0.0001$); and (iv) self-localisation error in these identified locations is smaller in map builders, compared to random explorers (Fig. 3F, ANOVA, $F_{1,18} = 19.46$, $P < 0.001$).

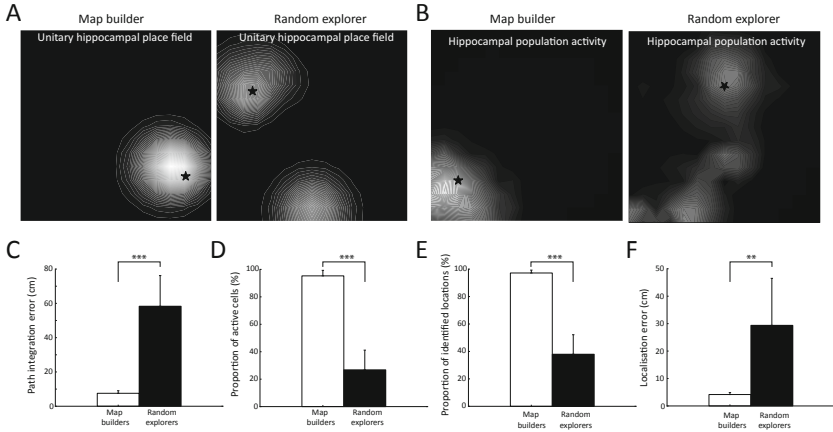


Fig. 3. *Spatial representations in map builders and random explorers.* (A, B) Examples of unitary place fields (A) and of hippocampal population activity profiles (B) for a map builder (left) and a random explorer (right). Colour codes for activity levels from 0 (black) to 1 (white). Black stars mark the true animal positions. (C-F) Bar plots show the mean \pm SD, across representative simulated mice from 10 different evolution runs, of path integration error (C), the proportion of highly active cells (D), the proportion of identified locations (E), and the population-encoded localisation error (F) for map builders (white) and random explorers (black).

4 Discussion

In this study we combined a spatial learning model with an artificial evolution approach to test the hypothesis that exploratory behaviour can be generated based on the functional organisation of the underlying spatial representation network. Our results suggest that round-trip exploration, a basic structural pattern of spatial behaviour observed in animals [2,5], can emerge as a result of an optimisation criterion based on the hypothesis that accurate spatial learning requires a combination of allothetic and idiothetic information [1].

The theoretical concept of combining allothetic and idiothetic inputs to form a stable representation of space is not new and served as a basis for a range of neurocomputational models of rodent behaviour (e.g. [8,11]). However, a fundamental problem with this concept in its current state is that it does not provide a mechanistic explanation of how allothetic and idiothetic cues are combined in a coherent representation during the beginning of exploration in a novel environment. This problem is well studied in the robot navigation field, where it is known as the Simultaneous Localisation and Mapping (SLAM) problem, solution to which is considered to be one of the notable achievements in robotics for the past two decades [3]. However, robotics solutions are generally not biologically plausible and provide poor explanation for the observed animal behaviour. Here, we addressed a similar problem using an artificial neural network model based on widely accepted hypotheses about the functional organisation of the spatial representation system in the mammalian brain.

In this study we used evolutionary algorithms (EAs) to tune the parameters of the exploration controller. The application of EAs in the present case is justified for two reasons. First, it is difficult to use standard gradient descent methods for model training here, since *(i)* the link between the parameters of the controller and the resulting behaviour is not formally defined, and *(ii)* the notion of ‘error’ is hard to define for exploratory behaviour. Second, given that a similar spatial representation system is used by many species (including rats [10] and humans [4]), it might be argued that efficient exploration mechanisms must have appeared at the same time as the spatial representation system itself. On these grounds, an application of EAs is natural for this problem, albeit we recognise that it is rather the computational nature of the problem than its relation to evolution that prompted us to use EAs in the current work.

Acknowledgements. Granted by the French National Research Agency (ANR) Project EvoNeuro, ANR-09-EMER-005-01.

References

1. Arleo, A., Rondi-Reig, L.: Multimodal sensory integration and concurrent navigation strategies for spatial cognition in real and artificial organisms. *J. Integr. Neurosci.* 6(3), 327–366 (2007)
2. Draï, D., Kafkafi, N., Benjamini, Y., Elmer, G., Golani, I.: Rats and mice share common ethologically relevant parameters of exploratory behavior. *Behav. Brain Res.* 125, 133–140 (2001)
3. Durrant-Whyte, H., Bailey, T.: Simultaneous localisation and mapping (SLAM): Part I. the essential algorithms. *IEEE Rob. Autom.* 13, 99–110 (2006)
4. Ekstrom, A.D., Kahana, M.J., Caplan, J.B., Fields, T.A., Isham, E.A., Newman, E.L., Fried, I.: Cellular networks underlying human spatial navigation. *Nature* 425(6954), 184–188 (2003)
5. Fonio, E., Benjamini, Y., Golani, I.: Freedom of movement and the stability of its unfolding in free exploration of mice. *Proc. Natl. Acad. Sci. USA* 106(50), 21335–21340 (2009)
6. Fyhn, M., Molden, S., Witter, M.P., Moser, E.I., Moser, M.B.: Spatial representation in the entorhinal cortex. *Science* 305, 1258–1264 (2004)
7. Horn, J., Nafpliotis, N., Goldberg, D.: A niched pareto genetic algorithm for multi-objective optimization. In: *IEEE WCCI, Congress on Evolutionary Computation*, pp. 82–87 (1994)
8. McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., Moser, M.B.: Path integration and the neural basis of the ‘cognitive map’. *Nat. Rev. Neurosci.* 7(8), 663–678 (2006)
9. Mouret, J.B., Doncieux, S.: Sferes v2: Evolvin’ in the multi-core world. In: *IEEE WCCI, Congress on Evolutionary Computation*, vol. (3), pp. 4079–4086 (2010)
10. O’Keefe, J., Nadel, L.: *The hippocampus as a cognitive map*. Clarendon Press, Oxford (1978)
11. Sheynikhovich, D., Chavarriaga, R., Strösslin, T., Arleo, A., Gerstner, W.: Is there a geometric module for spatial orientation? Insights from a rodent navigation model. *Psychol. Rev.* 116(3), 540–566 (2009)