

Self-organizing continuous attractor networks and path integration: two-dimensional models of place cells

S M Stringer, E T Rolls¹, T P Trappenberg and I E T de Araujo

Oxford University, Department of Experimental Psychology, South Parks Road,
Oxford OX1 3UD, UK

E-mail: Edmund.Rolls@psy.ox.ac.uk

Received 5 November 2000, in final form 25 June 2002

Published 27 August 2002

Online at stacks.iop.org/Network/13/429

Abstract

Single-neuron recording studies have demonstrated the existence of neurons in the hippocampus which appear to encode information about the place where a rat is located, and about the place at which a macaque is looking. We describe ‘continuous attractor’ neural network models of place cells with Gaussian spatial fields in which the recurrent collateral synaptic connections between the neurons reflect the distance between two places. The networks maintain a localized packet of neuronal activity that represents the place where the animal is located. We show for two related models how the representation of the two-dimensional space in the continuous attractor network of place cells could self-organize by modifying the synaptic connections between the neurons, and also how the place being represented can be updated by idiothetic (self-motion) signals in a neural implementation of path integration.

1. Introduction

Place cells which respond when the animal is in a particular location are found in the rat hippocampus (O’Keefe and Dostrovsky 1971, McNaughton *et al* 1983, O’Keefe 1984, Muller *et al* 1991, Markus *et al* 1995) and spatial view cells that respond when the monkey is looking towards a particular location in space are found in the macaque hippocampus (Rolls *et al* 1997, Georges-François *et al* 1999, Robertson *et al* 1998). As the rat moves without visual input in the dark, the place cells that are firing change based on idiothetic (self-motion) cues to represent the new place. Similarly, spatial view cells still respond when the monkey moves his eyes to look towards the same location in the dark (Robertson *et al* 1998). The spatial view cells are tuned to the two spatial dimensions of the horizontal and vertical dimensions of the

¹ Author to whom any correspondence should be addressed.

location in space of the spatial view, and can be updated by the horizontal and vertical positions of the eyes in the dark (Robertson *et al* 1998). In this paper we describe networks that could represent the current place where a rat is located, and how this network could be updated by self-motion cues. The theory described shows how the necessary synaptic connections could be set up by self-organizing learning. The theory also leads to models for how path integration is performed by the brain.

Continuous attractor neural networks (CANNs) provide a strong candidate for implementing the type of memory required. Two-dimensional (2D) continuous attractor networks can maintain the firing of their neurons at any location in a continuous physical 2D space such as the place where an animal is located (Samsonovich and McNaughton 1997, Taylor 1999) by using excitatory recurrent collateral connections between the neurons which reflect the distance between the neurons in the state space of the agent. What is not clear is how the necessary synaptic connections could be set up. Samsonovich and McNaughton (1997) have described a model which uses essentially a look-up table to map from one place to another place using head direction and whole body motion velocity signals, with no suggestion about how the look-up table would be set up. In this paper we develop 2D models of how place cells could be updated using idiothetic inputs, developing the 2D model from some of the ideas used to develop 1D models of how a head direction cell network could self-organize (Stringer *et al* 2002). Key issues we model are how the continuous attractor network that represents places could learn to represent the distances between different places; how the continuous attractor network could be updated in the dark by idiothetic head direction and self-motion signals; and how drift can be minimized in the 2D CANN. The models provide a foundation also for understanding how the spatial view cell system could self-organize in primates.

2. Models of 2D continuous attractors

2.1. Continuous attractor models of place cells: the neural representation of the location of a stationary agent

In this section we present a continuous attractor model of place cells that is able to support a stable activity packet representing the static location of a stationary agent. The model shows how the network could self-organize so that the synaptic connections between the neurons in the continuous attractor network can be set to strengths that reflect the distance in the state space between the places represented by the firing of the neurons. The model is composed of a recurrent network of place cells, which receives inputs from the visual system. In particular, there is a single type of modifiable synaptic connection: the recurrent connections within the network of place cells. These connections are established during an initial learning phase during which the agent moves throughout the containment area with visual input available. During the learning phase we assume that the visual inputs dominate all other excitatory inputs to the place cells. The ways in which visual cues might stimulate the place cells have been demonstrated in neurophysiological studies (McNaughton *et al* 1983, O'Keefe 1984, Muller *et al* 1991, Markus *et al* 1995), and modelled in theoretical investigations (O'Keefe and Burgess 1996, de Araujo *et al* 2001). The initial learning phase with visual inputs available is able to set up the recurrent synaptic connections between the neurons in the continuous attractor network such that, during the subsequent testing phase without visual cues, the network is able to stably represent the static location of a stationary agent.

The behaviour of the continuous attractor network of place cells is governed during the testing by the following 'leaky-integrator' dynamical equations. The following equation

describes the dynamics of the activation h_i^P of each place cell i :

$$\tau \frac{dh_i^P(t)}{dt} = -h_i^P(t) + \frac{\phi_0}{C^P} \sum_j (w_{ij}^{RC} - w^{INH}) r_j^P(t) + I_i^V, \quad (1)$$

where r_j^P is the firing rate of place cell j , w_{ij}^{RC} is the excitatory (positive) synaptic weight from place cell j to cell i , w^{INH} and ϕ_0 are constants, C^P is the number of synaptic connections received by each place cell from other place cells, I_i^V represents a visual input to place cell i , and τ is the time constant of the system. When the agent is in the dark, then the term I_i^V is set to zero. The firing rate r_i^P of cell i is determined from the activation h_i^P and the sigmoid function

$$r_i^P(t) = \frac{1}{1 + \exp[-2\beta(h_i^P(t) - \alpha)]}, \quad (2)$$

where α and β are the sigmoid threshold and slope, respectively. The equations (1) and (2) governing the internal dynamics of the 2D continuous attractor network of place cells are of an identical form to the corresponding equations (1) and (2) in Stringer *et al* (2002), which govern the behaviour of the 1D continuous attractor network of head direction cells. Thus, it is only the external inputs to these networks that determine the dimension of the space represented by the cell response properties. For the network of place cells to behave as a continuous attractor network, the recurrent synaptic weight profile must be self-organized during the learning phase in a similar manner to that described by Stringer *et al* (2002) for the head direction cell models. The details of how this is done for the place cell models are given later.

The dynamical equations (1) and (2) are used to model the behaviour of the place cells during testing. However, we assume that when visual cues are available, the visual inputs I_i^V stimulate the place cells to fire when the agent is at specific locations within the environment. Therefore, in the simulations presented below, rather than implementing the dynamical equations (1) and (2) during the learning phase, we set the firing rates of the place cells according to typical Gaussian response profiles as observed in neurophysiological studies (McNaughton *et al* 1983, O'Keefe 1984, Muller *et al* 1991, Markus *et al* 1995).

2.2. Self-organization of the recurrent synaptic connectivity in the continuous attractor network of place cells to represent the topology of the environment

To model a biologically plausible way of setting up the synaptic weights between the neurons in the continuous attractor network of place cells, we used an associative (Hebb-like) synaptic modification rule. The rationale is that neurons close together in the state space (the space being represented) would tend to be co-active during learning due to the large width of the firing fields, so that after the associative synaptic modification the synaptic strength between any two neurons represents the distance between the places represented in the state space of the agent (cf Redish and Touretzky 1998 and Stringer *et al* 2002). During learning the spatial fields are forced onto each neuron by for example visual inputs.

In the models proposed here the agent visits all parts of the environment. During the learning phase the visual input drives the place cells such that they fire maximally at particular locations. Hence, each place cell i is assigned a unique location (x_i, y_i) in the environment at which the cell is stimulated maximally by the visual cues. Then the firing rate r_i^P of each place cell i is set according to the following Gaussian response profile

$$r_i^P = \exp[-(s_i^P)^2 / 2(\sigma^P)^2], \quad (3)$$

where s_i^P is the distance between the current location of the agent (x, y) and the location at which cell i fires maximally (x_i, y_i) , and σ^P is the standard deviation. s_i^P is given by

$$s_i^P = \sqrt{(x_i - x)^2 + (y_i - y)^2}. \quad (4)$$

An associative learning rule that may be used for updating the weights w_{ij}^{RC} from place cell j with firing rate r_j^P to place cell i with firing rate r_i^P is the Hebb rule (Zhang 1996, Redish and Touretzky 1998)

$$\delta w_{ij}^{RC} = k r_i^P r_j^P \quad (5)$$

where δw_{ij}^{RC} is the change of synaptic weight and k is the learning rate constant. This rule operates by associating together place cells that tend to be co-active, and this leads to cells which respond to nearby locations developing stronger synaptic connections. A second rule that may be used to update the recurrent weights w_{ij}^{RC} is the trace rule

$$\delta w_{ij}^{RC} = k \bar{r}_i^P \bar{r}_j^P \quad (6)$$

where \bar{r}^P is the trace value of the firing rate of a place cell given by

$$\bar{r}^P(t + \delta t) = (1 - \eta)r^P(t + \delta t) + \eta\bar{r}^P(t) \quad (7)$$

where η is a parameter set in the interval $[0, 1]$ which determines the contribution of the current firing and the previous trace. See Stringer *et al* (2002) for more details about the significance of η .

2.3. Stabilization of the activity packet within the continuous attractor network when the agent is stationary

As described for the head direction cell models (Stringer *et al* 2002), the recurrent synaptic weights within the continuous attractor network will be corrupted by a certain amount of noise from the learning regime because of the irregularity it introduces, because for example cells in the middle of the containment area receive more updates than those towards the edges of the area. This in turn can lead to drift of the activity packet within the continuous attractor network of place cells when there are no visual cues available even when the agent is not moving. We propose that in real nervous systems this problem may be solved by enhancing the firing of neurons that are already firing, as suggested by Stringer *et al* (2002). This might be implemented through mechanisms for short-term synaptic enhancement (Koch 1999), or through the effects of voltage-dependent ion channels in the brain such as NMDA receptors (Lisman *et al* 1998). We simulate these effects by resetting the sigmoid threshold α_i at each timestep depending on the firing rate of place cell i at the previous timestep. That is, at each timestep $t + \delta t$ we set

$$\alpha_i = \begin{cases} \alpha^{\text{HIGH}} & \text{if } r_i^P(t) < \gamma \\ \alpha^{\text{LOW}} & \text{if } r_i^P(t) \geq \gamma \end{cases} \quad (8)$$

where γ is a firing rate threshold. This helps to reinforce the current position of the activity packet within the continuous attractor network of place cells. The sigmoid slopes are set to a constant value, β , for all cells i .

2.4. Continuous attractor models of place cells with idiothetic inputs: the neural representation of the time-varying location of a moving agent

In the basic continuous attractor model of place cells presented above, we did not address the issue of path integration, that is, the ability of the network to track and represent the time-varying location of a moving agent in the absence of visual input. A possible solution to the

problem of how the representation of place cells might be updated in the dark is provided by two possible inputs:

- (i) head direction cells whose activity in the dark may be updated by rotation cells, and
- (ii) idiothetic cues carrying information about the forward velocity of the agent.

Such information could in principle be used to update the activity within a network of place cells in the absence of visual input (Samsonovich and McNaughton 1997). In this section we present a model of place cells, Model 2A, that is able to solve the problem of path integration in the absence of visual cues through the incorporation of idiothetic inputs from head direction cells and forward velocity cells, and which develops its synaptic connections through self-organization. When an agent is moving in the dark, the idiothetic inputs are able to shift the activity packet within the network to track the state of the agent. A closely related model, Model 2B, is presented later in section 4. Models 2A and 2B employ mechanisms similar to those employed in the 1D head direction cell Models 1A and 1B of Stringer *et al* (2002), respectively.

The general neural network architecture is shown in figures 1 and 2. There is a recurrent continuous attractor network of place cells which receives three kinds of input:

- (i) visual inputs from the visual system used to force place cell firing during learning,
- (ii) idiothetic inputs from a network of head direction cells, and
- (iii) idiothetic inputs from a population of forward velocity cells.

For the place cell models presented below, we assume that the synaptic connectivity of the network of head direction cells has already become self-organized, as described by Stringer *et al* (2002), such that the head direction cells are already able to accurately represent the head direction of the agent during the learning and testing phases. This is achieved in the simulations presented below by fixing the firing rates of the head direction cells to reflect the true head direction of the agent, i.e. according to the Gaussian response profile

$$r_i^{\text{HD}} = \exp[-(s_i^{\text{HD}})^2 / 2(\sigma^{\text{HD}})^2], \quad (9)$$

where s_i^{HD} is the difference between the actual head direction x (in degrees) of the agent and the optimal head direction x_i for head direction cell i , and σ^{HD} is the standard deviation. s_i^{HD} is given by

$$s_i^{\text{HD}} = \text{MIN}(|x_i - x|, 360 - |x_i - x|). \quad (10)$$

This is the same Gaussian response profile that was applied to the head direction cells during learning for the 1D continuous attractor models of head direction cells presented in Stringer *et al* (2002). The firing rates of head direction cells in both rats (Taube *et al* 1996, Muller *et al* 1996) and macaques (Robertson *et al* 1999) are known to be approximately Gaussian. However, an additional collection of cells required for the place cell models, that was not needed for the head direction cell models, is the population of forward velocity cells. These cells fire as the agent moves forward, with a firing rate that increases monotonically with the forward velocity of the agent. Whole body motion cells have been described in primates (O'Mara *et al* 1994). The basic continuous attractor model of place cells described by equations (1) and (2) does not include idiothetic inputs. However, when the agent is moving and the location is changing with time, and there is no visual input available, then idiothetic inputs (from head direction and forward velocity cells) are used to shift the activity packet within the network to track the state of the agent. We now show how equation (1) may be extended to include idiothetic inputs from the head direction and forward velocity cells.

The first place cell model, Model 2A, is similar to Model 1A of Stringer *et al* (2002) in that it utilizes Sigma-Pi neurons. In this case, the dynamical equation (1) governing the

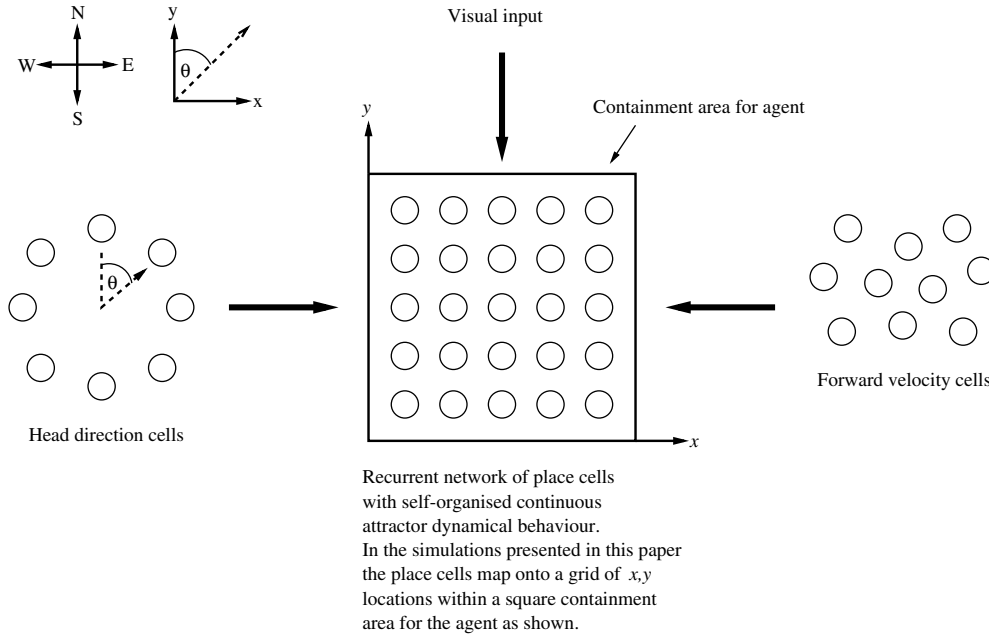


Figure 1. General network architecture for 2D continuous attractor models of place cells. There is a recurrent network of place cells which receives external inputs from three sources: (i) the visual system, (ii) a population of head direction cells, and (iii) a population of forward velocity cells. The place cells are distributed throughout a square containment area for the agent, where each place cell fires maximally when the agent is at a particular location in the area denoted by the Cartesian position coordinates x, y .

activations of the place cells is now extended to include inputs from the head direction and forward velocity cells in the following way. For Model 2A, the activation h_i^P of a place cell i is governed by the equation

$$\tau \frac{dh_i^P(t)}{dt} = -h_i^P(t) + \frac{\phi_0}{C^P} \sum_j (w_{ij}^{RC} - w^{INH}) r_j^P(t) + I_i^V + \frac{\phi_1}{C^{P \times HD \times FV}} \sum_{j,k,l} w_{ijkl}^{FV} r_j^P r_k^{HD} r_l^{FV}, \quad (11)$$

where r_j^P is the firing rate of place cell j , r_k^{HD} is the firing rate of head direction cell k , r_l^{FV} is the firing rate of forward velocity cell l and w_{ijkl}^{FV} is the corresponding overall effective connection strength, ϕ_0 and ϕ_1 are constants, and $C^{P \times HD \times FV}$ is the number of idiothetic connections received by each place cell from combinations of place cells, head direction cells and forward velocity cells. The first term on the right of equation (11) is a decay term, the second describes the effects of the recurrent connections in the continuous attractor, the third is the visual input (if present) and the fourth represents the effects of the idiothetic connections implemented by Sigma-Pi synapses. Thus, there are two types of synaptic connection to place cells:

- (i) recurrent connections from place cells to other place cells within the continuous attractor network, whose effective strength is governed by the terms w_{ij}^{RC} , and
- (ii) idiothetic connections dependent upon the interaction between an input from another place cell, a head direction cell, and a forward velocity cell, whose effective strength is governed by the terms w_{ijkl}^{FV} .




Figure 2. Neural network architecture for 2D continuous attractor models of place cells. There is a recurrent network of place cells with firing rates r^P , which receives external inputs from three sources: (i) the visual system, (ii) a population of head direction cells with firing rates r^{HD} , and (iii) a population of forward velocity cells with firing rates r^{FV} . The recurrent weights between the place cells are denoted by w^{RC} , and the idiotropic weights to the place cells from the forward velocity cells and head direction cells are denoted by w^{FV} .

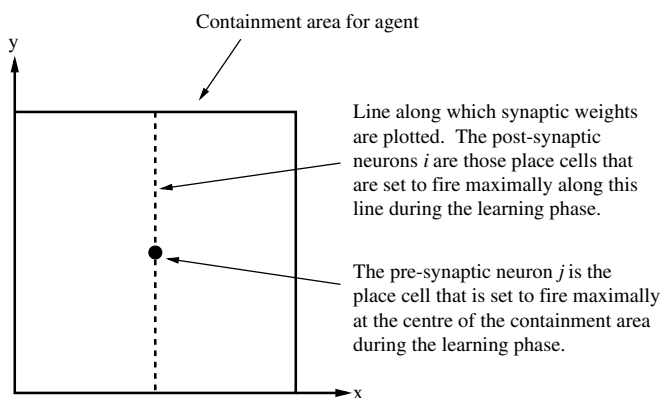


Figure 3. The line through containment area along which synaptic weights are plotted in figure 4. (The containment area is as shown in figure 1, where x and y denote the Cartesian position coordinates.) The weights plotted in figure 4 are for synapses with the following pre-synaptic and post-synaptic neurons. All of the synapses have the same pre-synaptic neuron j , which is the place cell set to fire maximally at the centre of the containment area during the learning phase. The post-synaptic neurons i are those place cells set to fire maximally at various positions along the dashed line through the centre of the containment area.

At each timestep, once the place cell activations h_i^P have been updated, the place cell firing rates r_i^P are calculated according to the sigmoid transfer function (2). Therefore, the initial learning phase involves the setting up of the synaptic weights w_{ij}^{RC} and w_{ijkl}^{FV} .

2.5. Self-organization of synaptic connectivity from idiothetic inputs to the continuous attractor network of place cells

In this section we describe how the idiothetic synaptic connections to the continuous attractor network of place cells self-organize during the learning phase such that, when the visual cues are removed, the idiothetic inputs are able to shift the activity packet within the network of place cells such that the firing of the place cells is able to continue to represent the location of the agent. A qualitative description occurs first, and then a formal quantitative specification of the model. The proposal is that during learning of the synaptic weights, the place where the animal is located is represented by the post-synaptic activation of a place cell, and this is associated with inputs to that place cell from recently active place cells (detected by a temporal trace in the presynaptic term), and with inputs from the currently firing idiothetic cells. In this network, the idiothetic inputs come from head direction and forward velocity cells. The result of this learning is that if idiothetic signals occur in the dark, the activity packet is moved from the currently active place cells towards place cells that during learning had subsequently become active in the presence of those particular idiothetic inputs.

During the learning phase the response properties of the place cells, head direction cells and forward velocity cells are as follows. As the agent moves through the environment, the visual cues drive individual place cells to fire maximally for particular locations, with the firing rates varying according to Gaussian response profiles of the form (3). Similarly, as discussed for the head direction cell models, we assume the visual cues also drive the head direction cells to fire maximally for particular head directions, with Gaussian response profiles of the form (9). Lastly, the forward velocity cells fire if the agent is moving forward, and with a firing rate that increases monotonically with the forward velocity of the agent. In the simulations performed later, the firing rates of the forward velocity cells during the learning and testing phases are set to 1 to denote a constant forward velocity.

For Model 2A, the learning phase involves setting up the synaptic weights w_{ijkl}^{FV} for all ordered pairs of place cells i and j , and for all head direction cells k and forward velocity cells l . At the start of the learning phase the synaptic weights w_{ijkl}^{FV} may be set to zero or random positive values. Then the learning phase continues with the agent moving through the environment with the place cells, head direction cells and forward velocity cells firing according to the response properties described above. During this, the synaptic weights w_{ijkl}^{FV} are updated at each timestep according to

$$\delta w_{ijkl}^{\text{FV}} = \tilde{k} r_i^{\text{P}} \bar{r}_j^{\text{P}} r_k^{\text{HD}} r_l^{\text{FV}} \quad (12)$$

where $\delta w_{ijkl}^{\text{FV}}$ is the change of synaptic weight, r_i^{P} is the instantaneous firing rate of place cell i , \bar{r}_j^{P} is the trace value of the firing rate of place cell j given by equation (7), r_k^{HD} is the firing rate of head direction cell k , r_l^{FV} is the firing rate of forward velocity cell l , and \tilde{k} is the learning rate. If we consider two place cells i and j that are stimulated by the available visual cues to fire maximally in nearby locations, then during the learning phase cell i may often fire a short time after cell j depending on the direction of movement of the agent as it passes through the location associated with place cell j . In this situation the effect of the above learning rule would be to ensure that the size of the weights w_{ijkl}^{FV} would be largest for head direction cells k which fire maximally for head directions most closely aligned with the direction of the location represented by cell i from the location represented by cell j . The effect of the above learning rules for the synaptic weights w_{ijkl}^{FV} should be to generate a synaptic connectivity such that the co-firing of (i) a place cell j , (ii) a head direction cell k , and (iii) a forward velocity cell l , should stimulate place cell i where place cell i represents a location that is a small translation in the appropriate direction from the location represented by place cell j . Thus, the co-firing

Table 1. Parameter values for Model 2A.

| | |
|---------------------------|------------|
| σ^{HD} | 20° |
| σ^{P} | 0.05 |
| Learning rate k | 0.001 |
| Learning rate \tilde{k} | 0.001 |
| Trace parameter η | 0.9 |
| τ | 1.0 |
| ϕ_0 | 50 000 |
| ϕ_1 | 1000 000 |
| w^{INH} | 0.05 |
| γ | 0.5 |
| α^{HIGH} | 0.0 |
| α^{LOW} | -20 |
| β | 0.1 |

of a set of place cells representing a particular location, a particular cluster of head direction cells, and the forward velocity cells, should stimulate the firing of further place cells such that the pattern of activity within the place cell network evolves continuously to faithfully reflect and track the changing location of the agent.

3. Simulation results

The experimental set-up is as illustrated in figures 1 and 2. The neural network architecture of the simulated agent consists of a continuous attractor network of place cells, which receives inputs from the visual system (during the learning phase), a population of head direction cells, and a population of forward velocity cells (in fact only a single forward velocity cell is simulated). During the simulation, the agent moves around in a square 1 unit \times 1 unit containment area. This area is covered by a square 50 \times 50 grid of nodes. At each node i there is a single place cell that is set to fire maximally at that location (x_i, y_i) during the learning phase with a Gaussian response profile given by equation (3). The standard deviation σ^{P} used for the place cell Gaussian response profiles (3) is 0.05 units. This gives a total of 2500 place cells within the place cell continuous attractor network of the agent. In addition, due to the large computational cost of these simulations, we include only eight head direction cells. The head direction cells $k = 1, \dots, 8$ are set to fire maximally for the eight principal compass directions in a clockwise order as follows: $k = 1$ fires maximally for head direction 0° (North); $k = 2$ fires maximally for head direction 45° (North-East); and so on up to $k = 8$ which fires maximally for head direction 315° (North-West). The head direction cells are set to fire maximally for these directions according to the Gaussian response profile (9). The standard deviation σ^{HD} used for the head direction cell Gaussian response profiles (9) is 20° . These head direction cell response profiles are implemented for both the initial learning phase in the light and the subsequent testing phase in the dark. Finally, we include only 1 forward velocity cell, which is the minimal number required for the model to work. For the simulations discussed below we present results only for the Sigma-Pi Model 2A; further simulations (not shown here) have confirmed that Model 2B gives very similar results. The model parameters used in the simulations of Model 2A are given in table 1.

The numerical simulation begins with the initial learning phase in which the recurrent and idiothetic synaptic weights, which are initialized to zero at the beginning of the simulation, are self-organized. During this phase, visual cues are available to the agent to help guide the self-organization of the weights. However, as described earlier, we do not model the

visual cues explicitly and the simulations are simplified in the following manner. During the learning phase, rather than implementing the dynamical equations (2) and (11) for Model 2A explicitly, we set the firing rates of the place cells to typical Gaussian response profiles in accordance with the observed behaviours of such cells in physiological studies, as described above. The learning phase then proceeds as follows. Firstly, the agent is simulated moving along approximately 50 equi-distant parallel paths in the northwards direction. These paths fully cover the containment area. That is, each of the parallel paths begins at the boundary $y = 0$ and terminates at the boundary $y = 1$. In addition, the separate paths are spread evenly through the containment area, with the first path aligned with the boundary $x = 0$ and the last path aligned with the boundary $x = 1$. The agent moves along each path with a constant velocity for which the firing rate r^{FV} of the forward velocity cell is set to 1. (Although regular training was used, this is not necessary for such networks to learn usefully, as shown by results with 1D continuous attractor networks, and in that the trace rule can help when the training conditions are irregular (Stringer *et al* 2002).) Each path is discretized into approximately 50 steps (locations) at which the synaptic weights are updated. At each step the following three calculations are performed:

- (i) the current position of the agent is calculated from its location, head direction and speed at the previous timestep;
- (ii) the firing rates of the place cells, head directions cells and the forward velocity cell are calculated as described above; and
- (iii) the recurrent and idiothetic synaptic weights are updated according to equations (6) and (12).

Further model parameters are as follows. The learning rates, k and \tilde{k} , for both the recurrent and idiothetic weights are set to 0.001. Also, the trace learning parameter η used in the trace update equation (7) is set to 0.9. This value of η is used in both learning rules (6) and (12) for both the recurrent and idiothetic synaptic weights. A similar procedure is then repeated for each of the seven other principal directions in a clockwise order. This completes the learning phase.

Examples of the final recurrent and idiothetic synaptic weights are shown in figure 4. The plot on the left compares the recurrent and idiothetic synaptic weight profiles as follows. The first graph shows the recurrent weights w_{ij}^{RC} where the pre-synaptic neuron j is the place cell set to fire maximally at the centre of the containment area during the learning phase, and the post-synaptic neurons i are those place cells set to fire maximally at various positions along the dashed line through the centre of the containment area shown in figure 3. It can be seen that the learning has resulted in nearby cells in location space, which need not be at all close to each other in the brain, developing stronger recurrent synaptic connections than cells that are more distant in location space. Furthermore, it can be seen that the graph of the recurrent weights is symmetric about the central node and is approximately a Gaussian function of the distance between the cells in location space. This is important for stably supporting the activity packet at a particular location when the agent is stationary. The second graph shows the idiothetic weights w_{ijk}^{FV} where the pre- and post-synaptic neurons j and i are as above, and the head direction cell k is 1 (North). Here it can be seen that the idiothetic weight profile for $k = 1$ (North) is highly asymmetric about the central node, and that the peak lies to the north of the central node. This is essential to enable the idiothetic inputs to shift the activity packet within the network of place cells in the correct direction when the agent is moving northwards. The plot on the right compares three different idiothetic synaptic weight profiles as follows. The first, second and third graphs show the idiothetic weights w_{ijk}^{FV} where the pre- and post-synaptic neurons j and i are as above, and the head direction cells k are respectively $k = 1$ (North),




Figure 4. Synaptic weight profiles plotted along the dashed curve through the centre of the containment area shown in figure 3. Left: the plot on the left compares the recurrent and idiothetic synaptic weight profiles as follows. The first graph shows the recurrent weights w_{ij}^{RC} where the pre-synaptic neuron j is the place cell set to fire maximally at the centre of the containment area during the learning phase, and the post-synaptic neurons i are those place cells set to fire maximally at various positions along the dashed curve through the centre of the containment area shown in figure 3. The second graph shows the idiothetic weights w_{ijk}^{FV} where the pre- and post-synaptic place cells j and i are as above, and the head direction cell k is 1 (North). Right: the plot on the right compares three different idiothetic synaptic weight profiles as follows. The first, second and third graphs show the idiothetic weights w_{ijk}^{FV} where the pre- and post-synaptic place cells j and i are as above, and the head direction cells k are respectively $k = 1$ (North), $k = 3$ (East) and $k = 5$ (South).

$k = 3$ (East) and $k = 5$ (South). Here it can be seen that the idiothetic weight profiles for $k = 1$ (North) and $k = 5$ (South) are both asymmetric about the central node, but that while for $k = 1$ the peak lies to the north of the central node, for $k = 5$ the peak lies to the south of the central node. In addition, for $k = 3$ (East), the idiothetic weight profile is in fact symmetric about the central node, with a lower peak value than for either $k = 1$ or 5. Hence, the effect of the learning rule for the idiothetic weights has been to ensure that sizes of the weights w_{ijk}^{FV} are largest for head direction cells k which fire maximally for head directions most closely aligned with the direction of the location represented by cell i from the location represented by cell j .

After the learning phase is completed, the simulation continues with the testing phase in which visual cues are no longer available, and the continuous attractor network of place cells must track the location of the agent solely through the idiothetic inputs. For the testing phase, the full dynamical equations (2) and (11) for Model 2A are implemented. At each timestep the following four calculations are performed:

- (i) the activations h_i^P of the place cells are updated;
- (ii) the firing rates r_i^P of the place cells are updated according to the sigmoid function (2);
- (iii) the firing rates r_k^{HD} of the head direction cells are set according to their Gaussian response profiles (9), as described above;
- (iv) the firing rate r_l^{FV} of the forward velocity cell is set to 1 if the agent is moving, and set to 0 if the agent is stationary.

The testing phase then proceeds as follows. Firstly, the activations and firing rates of the place cells are initialized to zero. Then the agent is placed in the containment area for 500 timesteps

at the location $x = 0.2, y = 0.2$, with visual input available. While the agent rests at this position the visual input terms I_i^V for each place cell i in equations (11) is set to a Gaussian response profile identical (except for a constant scaling) to that used for place cells during the learning phase given by equation (3). Next, the visual input is removed by setting all of the terms I_i^V to zero, and then the agent is allowed to rest in the same location for another 500 timesteps. This process leads to a stable packet of activity within the continuous attractor network of place cells. The agent was then moved along a three stage multi-directional track through the containment area as follows:

- (i) during the first stage the agent moves eastwards for 150 timesteps;
- (ii) during the second stage the agent moves northwards for 150 timesteps; and
- (iii) during the third stage the agent moves North–Eastwards for 150 timesteps.

Between each successive pair of stages, the agent was kept stationary for 100 timesteps to demonstrate that the place cell network representation can be stably maintained for different static locations of the agent. During these periods of rest and movement, the firing rates of the place cells, head direction cells and the forward velocity cell were calculated as described above. Further parameter values are as follows. The parameter τ is set to 1.0, ϕ_0 is set to 50 000, ϕ_1 is set to 1000 000, and w^{INH} is set to approximately 0.05. The sigmoid transfer function parameters are as follows: α^{HIGH} is 0.0, α^{LOW} is -20.0 , β is 0.1, and γ is 0.5. Finally, the timestep was approximately 0.2.

Figure 5 shows the firing rate profiles within the continuous attractor network of place cells that occurred during the testing phase. In the top left is shown a stable firing rate profile within the continuous attractor network of place cells just before the agent starts to move on the first stage of its track. That is, the plot shows the network activity after the initial visual input has stimulated activity within the network for 500 timesteps and the network has then been allowed to settle for a further 500 timesteps without visual input available. The activity packet is stable, allowing the place cell network to reliably maintain a representation of the static location of the stationary agent. However, the stability of the activity packet relies on the reinforcement of the firing of those neurons that are already highly active by setting α^{LOW} to be, for example, -20.0 . If the sigmoid parameter α^{LOW} is set to 0 (so that there is no reinforcement of the firing of those neurons that are already highly active), then the activity packet within the continuous attractor network is not stably maintained. In this case, the activity packet drifts quickly towards the centre of the containment area in location space. This is due to a particular kind of systematic inhomogeneity that accumulates in the recurrent weights of the continuous attractor network during the learning phase because, with the current learning regime, the agent will spend a greater amount of time near to the central nodes of the containment area than nodes at the boundaries of the containment area. This means that place cells that fire maximally when the agent is near the centre of the containment area will receive more recurrent synaptic updates during learning than neurons that fire maximally near to boundaries. Hence the neurons that represent the centre of the containment area will tend to develop stronger recurrent weights, and will therefore tend to attract the activity packet away from the boundaries of the containment area. However, this drift of the activity packet can be prevented by enhancing the firing of those neurons that already have a relatively high firing rate by setting α^{LOW} to be, for example, -20.0 .

In the top right of figure 5 the maximum firing rates that occurred during the testing phase as the agent moved along a fixed track within the containment area are shown. The maximum firing rate is calculated over all timesteps for each x, y location. This plot clearly shows the path taken by the agent. It can be seen that the activity packet can accurately track the path of the agent in the containment area, even as the agent changes direction. In the simulation




Figure 5. Firing rate profiles within the continuous attractor network of place cells. Top left: stable firing rate profile within the network of place cells before the agent starts to move. Top right: maximum firing rates that occurred during the testing phase as the agent moved along a fixed track within the containment area. The maximum firing rate is calculated over all timesteps for each x , y location. This plot clearly shows the path taken by the agent. Bottom left: maximum firing rates that occurred during the testing phase as the agent moved along a fixed track, but with the maximum firing rate now calculated over all y nodes for each x location and for each timestep. Bottom right: maximum firing rates that occurred during the testing phase as the agent moved along a fixed track, but with the maximum firing rate now calculated over all x nodes for each y location and for each timestep.

presented here, the agent was kept stationary for 100 timesteps between successive stages of the track. However, in further simulations (not shown here) it has been found that this resting period for the agent between different directions is not necessary for the place cell network to accurately track the location of the agent. The bottom left of figure 5 shows the maximum firing rates that occurred during the testing phase as the agent moved along the same fixed

track, but with the maximum firing rate now calculated over all y nodes for each x location and for each timestep. In the bottom right of figure 5 the maximum firing rates are shown that occurred during the testing phase as the agent moved along the fixed track, but with the maximum firing rate now calculated over all x nodes for each y location and for each timestep. The bottom two plots of figure 5 reveal more detail about the movement of the activity packet within the continuous attractor network through time. Firstly, it can be seen that for each interval of 100 timesteps between successive stages of movement of the agent along its track, the activity packet is able to stably represent the static location of the stationary agent. This may be seen, for example, from timesteps 150 to 250 where the activity packet remains static in both of the bottom two plots of figure 5. Another detail revealed by the bottom two plots of figure 5 is the elongation of the activity packet in the direction of motion of the agent when the agent is moving. For example, for timesteps 1–150 when the agent is moving eastwards in the x direction, the projection of the activity packet onto the x -axis at any particular timestep (bottom left) is broader than the projection of the activity packet onto the y -axis (bottom right). A further interesting observation is that during the subsequent resting stage (timesteps 150–250), the resting activity packet appears to keep the shape from the preceding movement stage. That is, for timesteps 150–250 when the agent is stationary, the projection of the activity packet onto the x -axis at any particular timestep (bottom left) is still broader than the projection of the activity packet onto the y -axis (bottom right). This is due to the short term memory effect implemented by the nonlinearity in the activation function of the neurons that could reflect the operation of NMDA receptors effected by equation (8). This phenomenon can be seen again for timesteps 400–500 where the agent is again stationary, but the projection of the activity packet onto the x -axis at any particular timestep (bottom left) is now less broad than the projection of the activity packet onto the y -axis (bottom right). This reflects the fact that in the preceding timesteps 250–400 the agent was moving northwards in the y -direction.

4. Model 2B

The second place cell model, Model 2B, is similar to Model 1B of Stringer *et al* (2002) in that it incorporates synapse modulation effects into the calculation of the neuronal activations in the recurrent network. In this case, the dynamical equation (1) governing the activations of the place cells is now extended to include inputs from the head direction and forward velocity cells in the following way. For Model 2B, the activation h_i^P of a place cell i is governed by the equation

$$\tau \frac{dh_i^P(t)}{dt} = -h_i^P(t) + \frac{\phi_0}{C^P} \sum_j (\tilde{w}_{ij}^{\text{RC}} - w^{\text{INH}}) r_j^P(t) + I_i^V, \quad (13)$$

where r_j^P is the firing rate of place cell j , and where $\tilde{w}_{ij}^{\text{RC}}$ is the modulated strength of the synapse from place cell j to place cell i . The modulated synaptic weight $\tilde{w}_{ij}^{\text{RC}}$ is given by

$$\tilde{w}_{ij}^{\text{RC}} = w_{ij}^{\text{RC}} \left(1 + \frac{\phi_2}{C^{\text{HD} \times \text{FV}}} \sum_{kl} \lambda_{ijkl}^{\text{FV}} r_k^{\text{HD}} r_l^{\text{FV}} \right) \quad (14)$$

where w_{ij}^{RC} is the unmodulated synaptic strength set up during the learning phase, r_k^{HD} is the firing rate of head direction cell k , r_l^{FV} is the firing rate of forward velocity cell l , and $\lambda_{ijkl}^{\text{FV}}$ is the corresponding modulation factor. Thus, there are two types of synaptic connection between place cells:

- (i) recurrent connections from place cells to other place cells within the continuous attractor network, whose strength is governed by the terms $\tilde{w}_{ij}^{\text{RC}}$, and

- (ii) idiothetic connections from the head direction and forward velocity cells to the place cell network, which now have a modulating effect on the synapses between the place cells, and whose strength is governed by the modulation factors $\lambda_{ijkl}^{\text{FV}}$.

As for Model 2A, once the place cell activations h_i^{P} have been updated at the current timestep, the place cell firing rates r_i^{P} are calculated according to the sigmoid transfer function (2).

The initial learning phase involves the setting up of the synaptic weights w_{ij}^{RC} and the modulation factors $\lambda_{ijkl}^{\text{FV}}$. The synaptic weights w_{ij}^{RC} and the modulation factors $\lambda_{ijkl}^{\text{FV}}$ are set up during an initial learning phase similar to that described for Model 2A above, where the recurrent weights are updated according to equation (5), and the modulation factors $\lambda_{ijkl}^{\text{FV}}$ are updated at each timestep according to

$$\delta\lambda_{ijkl}^{\text{FV}} = \tilde{k} r_i^{\text{P}} \bar{r}_j^{\text{P}} r_k^{\text{HD}} r_l^{\text{FV}} \quad (15)$$

where r_i^{P} is the instantaneous firing rate of place cell i , \bar{r}_j^{P} is the trace value of the firing rate of place cell j given by equation (7), r_k^{HD} is the firing rate of head direction cell k , r_l^{FV} is the firing rate of forward velocity cell l , and \tilde{k} is the learning rate.

5. Discussion

In this paper we have developed 2D continuous attractor models of place cells that can learn the topology of a space. We have also shown how they are able to self-organize during an initial learning phase with visual cues available such that, when the agent is subsequently placed in complete darkness, the continuous attractor network is able to continue to track and faithfully represent the state of the agent using only idiothetic cues. The network thus performs path integration. The motivation for developing such self-organizing continuous attractor models stems from the problems of biological implausibility associated with current models which tend to rely on pre-set or hard-wired synaptic connectivities for the idiothetic inputs, which are needed to shift the activity packet in the absence of visual input. In this paper we have presented models that operate with biologically plausible learning rules and an appropriate architecture to solve the following two problems.

The first problem is how the recurrent synaptic weights within the continuous attractor should be self-organized in order to represent the topology of the 2D environment (or more generally the 2D state space of the agent). In the models presented here, a continuous attractor network is used for the representation, and it is trained with associative synaptic modification rules such as those shown in equations (5) and (6) to learn the distances between the places represented by the firing of neurons based on the coactivity of neurons with tapering (e.g. Gaussian) place fields produced by the visual input.

The self-organizing continuous attractor networks described in this paper are in some senses an interesting converse of Kohonen self-organizing maps (Kohonen 1989, 1995, Rolls and Treves 1998). In Kohonen maps, a map-like topology is specified by the connections being strong between nearby neurons, and weaker between distant neurons. This is frequently conceptualized as being implemented by short-range fixed excitatory connections between nearby neurons, and longer range inhibition (implemented for example by inhibitory interneurons). The Kohonen map operates in effect as a competitive network but with local interactions of the type just described (Rolls and Treves 1998). The result is that neurons which respond to similar vectors or features in the input space become close together in the map-like topology. In contrast, in the self-organizing continuous attractor networks described here, there is no map-like topology between the spatial relations of the neurons in the continuous attractor and the space being represented. Instead, the representations in the continuous attractor of different

locations of the state space are randomly mapped, with the closeness in the state space being represented by the strength of the synapses between the different neurons in the continuous attractor, and not by the closeness of the neurons in the network. A major difference between the two networks is that a Kohonen map has a dimensionality that is preset effectively by whether the neighbour relations are described between neurons in a 1D array (a line), a 2D array (a 2D space), etc. Whatever the dimensionality of the inputs that are to be learned, they will be represented in such (typically low order) spaces. In contrast, in the self-organizing continuous attractor networks described here, the continuous attractor takes on the dimensionality of the inputs. In this way, there is no presetting of the dimensionality of the solutions that can be represented, and the system will self-organize to represent as high a dimensionality as is present in the input space (subject to the number of connections onto each neuron in the continuous attractor from other neurons in the continuous attractor). Self-organizing continuous attractor networks thus provide a much more flexible system for representing state spaces, including for example capturing the geometry and the complex spatial arrangements of an irregular environment. Indeed, the recurrent connections in the continuous attractor network described here could learn the topological relations even in highly irregular and warped spaces, such as spaces partially intersected by barriers. Self-organizing continuous attractor networks thus provide a flexible system for representing state spaces, including for example in the case of place cells capturing the geometry and the complex spatial arrangements of a cluttered natural environment. We also note that, as in the continuous attractor networks described here, nearby cells in the hippocampus do not represent nearby places, and indeed there is no clear topology in the way that place cells are arranged in the rat hippocampus (O'Keefe and Conway 1978, Muller and Kubie 1987, Markus *et al* 1994). Indeed, with respect to place cells in the rat hippocampus, it is not feasible that the rat could be born with an innate synaptic encoding of the different environments it will encounter during its lifetime. Thus, a recurrent synaptic connectivity that reflects the relationship between neuronal responses and the state of an agent may need to arise naturally through learning and self-organization, for example by modifying the strengths of connections based on the similarity in the responses of the neurons.

The second problem is how idiothetic synaptic connections could be self-organized in such a way that self-motion of the agent moves the activity packet in the continuous attractor to represent the correct location in state space. Without a suggestion for how this is achieved, a hard-wired model of this process must effectively rely on a form of 'look-up' table to be able to move the activity packet in the correct way (Samsonovich and McNaughton 1997). The apparent absence of any spatial regularity in the cell response properties of the continuous attractor networks makes such innate hard-wiring unlikely, as discussed by Stringer *et al* (2002). In this paper we present two models that can self-organize to solve the problem of the idiothetic update of the representation of the current position of the agent in a 2D state space.

The neural architecture implied by Model 2A utilizes Sigma-Pi synapses with four firing rate terms, as seen in equations (11) and (12). If this seems complicated, one can note that if there was a population of cells that represented combinations of linear motion and head direction (e.g. moving North fast), then two of the terms, r_l^{FV} and r_k^{HD} would collapse together, leaving only three firing rate terms for modified versions of equations (11) and (12). An equivalent reduction can be made for Model 2B. Biophysical mechanisms that might be able to implement such three-term (or even possibly four-term) synapses have been discussed by Stringer *et al* (2002), and include, for three-term synapses, presynaptic contacts.

We note that it is a property of the models described in this paper as well as in the companion paper on 1D models of head direction cells (Stringer *et al* 2002) that they move the current representation at velocities that depend on the magnitude of the driving idiothetic input (which reflects the linear velocity of the agent for the 2D models, and the angular velocity of the agent

for head direction cells). This occurs even when the network is tested with magnitudes of the idiothetic inputs with which it has not been trained. For the models of head direction cells presented, it was found that the relation between the idiothetic driving input and the velocity of the head direction representation in the continuous attractor network is approximately linear if NMDA receptor non-linearity is not used to stabilize the network (see figure 7 of Stringer *et al* (2002)), and shows a threshold non-linearity if NMDA receptor like non-linearity is included in the neuronal activation function (see figure 10 of Stringer *et al* (2002)).

The models described here have been applied to place cells in rats. However, the models are generic, and can be applied to other problems. For example, the spatial view cells of primates are tuned to the two spatial dimensions of the horizontal and vertical dimensions of the location in space of the spatial view, and can be updated by the horizontal and vertical positions of the eyes in the dark (Robertson *et al* 1998). The 2D models described here provide a foundation for understanding how this spatial view system could self-organize to allow idiothetic update by eye position. We note that very few idiothetic Sigma-Pi synapses would suffice to implement the mechanism for path integration described in this paper. The reason for this is that the introduction of any asymmetry into the continuous attractor functional connectivity will suffice to move the activity packet. The prediction is thus made that the connectivity of the idiothetic inputs could be quite sparse in brain systems that perform path integration.

In this paper we show how path integration could be achieved in a system that self-organizes by associative learning. The path integration is performed in the sense that the representation in a continuous attractor network of the current location of the agent in a 2D environment can be continuously updated based on idiothetic (self-motion) cues, in the absence of visual inputs. The idiothetic cues used to update the place representation are from head direction cell firing and from linear whole body velocity cues. We note that whole body motion cells are present in the primate hippocampus (O'Mara *et al* 1994) and that head direction cells are present in the primate presubiculum (Robertson *et al* 1999). In the companion paper (Stringer *et al* 2002) we showed how a continuous attractor network representing head direction could self-organize to allow idiothetic update by head rotation signals. Together these two proposals provide much of what is needed for an agent to perform path integration in a 2D environment. In particular, an agent that implemented the proposals described in these two papers could continuously update its head direction and position in a 2D space in the dark using head rotation and linear velocity signals.

Acknowledgments

This research was supported by the Medical Research Council, grant PG9826105, by the Human Frontier Science Program, and by the MRC Interdisciplinary Research Centre for Cognitive Neuroscience.

References

- de Araujo I E T, Rolls E T and Stringer S M 2001 A view model which accounts for the spatial fields of hippocampal primate spatial view cells and rat place cells *Hippocampus* **11** 699–706
- Georges-François P, Rolls E T and Robertson R G 1999 Spatial view cells in the primate hippocampus: allocentric view not head direction or eye position or place *Cereb. Cortex* **9** 197–212
- Koch C 1999 *Biophysics of Computation* (Oxford: Oxford University Press)
- Kohonen T 1989 *Self-Organization and Associative Memory* 3rd edn (Berlin: Springer)
- Kohonen T 1995 *Self-Organizing Maps* (Berlin: Springer)

- Lisman J E, Fellous J M and Wang X J 1998 A role for NMDA-receptor channels in working memory *Nature Neurosci.* **1** 273–5
- Markus E J, Barnes C A, McNaughton B L, Gladden V L and Skaggs W 1994 Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input *Hippocampus* **4** 410–21
- Markus E J, Qin Y L, Leonard B, Skaggs W, McNaughton B L and Barnes C A 1995 Interactions between location and task affect the spatial and directional firing of hippocampal neurons *J. Neurosci.* **15** 7079–94
- McNaughton B L, Barnes C A and O'Keefe J 1983 The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats *Exp. Brain Res.* **52** 41–9
- Muller R U and Kubie J L 1987 The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells *J. Neurosci.* **7** 1951–68
- Muller R U, Kubie J L, Bostock E M, Taube J S and Quirk G J 1991 Spatial firing correlates of neurons in the hippocampal formation of freely moving rats *Brain and Space* ed J Paillard (Oxford: Oxford University Press) pp 296–333
- Muller R U, Ranck J B and Taube J S 1996 Head direction cells: properties and functional significance *Curr. Opin. Neurobiol.* **6** 196–206
- O'Keefe J 1984 Spatial memory within and without the hippocampal system *Neurobiology of the Hippocampus* ed W Seifert (London: Academic) pp 375–403
- O'Keefe J and Burgess N 1996 Geometric determinants of the place fields of hippocampal neurons *Nature* **381** 425–8
- O'Keefe J and Conway D H 1978 Hippocampal place units in the freely moving rat: why they fire where they fire *Exp. Brain Res.* **31** 573–90
- O'Keefe J and Dostrovsky J 1971 The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat *Brain Res.* **34** 171–5
- O'Mara S M, Rolls E T, Berthoz A and Kesner R P 1994 Neurons responding to whole-body motion in the primate hippocampus *J. Neurosci.* **14** 6511–23
- Redish A D and Touretzky D S 1998 The role of the hippocampus in solving the Morris water maze *Neural Comput.* **10** 73–111
- Robertson R G, Rolls E T and Georges-François P 1998 Spatial view cells in the primate hippocampus: effects of removal of view details *J. Neurophysiol.* **79** 1145–56
- Robertson R G, Rolls E T, Georges-François P and Panzeri S 1999 Head direction cells in the primate presubiculum *Hippocampus* **9** 206–19
- Rolls E T, Robertson R G and Georges-François P 1997 Spatial view cells in the primate hippocampus *Eur. J. Neurosci.* **9** 1789–94
- Rolls E T and Treves A 1998 *Neural Networks and Brain Function* (Oxford: Oxford University Press)
- Samsonovich A and McNaughton B 1997 Path integration and cognitive mapping in a continuous attractor neural network model *J. Neurosci.* **17** 5900–20
- Stringer S M, Trappenberg T P, Rolls E T and de Araujo I E T 2002 Self-organizing continuous attractor networks and path integration: one-dimensional models of head direction cells *Network: Comput. Neural Syst.* **13** 217–42
- Taube J S, Goodridge J P, Golob E G, Dudchenko P A and Stackman R W 1996 Processing the head direction signal: a review and commentary *Brain Res. Bull.* **40** 477–86
- Taylor J G 1999 Neural 'bubble' dynamics in two dimensions: foundations *Biol. Cybern.* **80** 393–409
- Zhang K 1996 Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory *J. Neurosci.* **16** 2112–26