

Self-organizing continuous attractor network models of hippocampal spatial view cells

S.M. Stringer^a, E.T. Rolls^{a,*}, T.P. Trappenberg^b

^a Department of Experimental Psychology, Centre for Computational Neuroscience, Oxford University, South Parks Road, Oxford OX1 3UD, UK

^b Faculty of Computer Science, Dalhousie University, 6050 University Avenue, Halifax, Nova Scotia, Canada B3H 1W5

Received 24 June 2004; revised 17 August 2004; accepted 20 August 2004

Abstract

Single neuron recording studies have demonstrated the existence of hippocampal spatial view neurons which encode information about the spatial location at which a primate is looking in the environment. These neurons are able to maintain their firing even in the absence of visual input. The standard neuronal network approach to model networks with memory that represent continuous spaces is that of continuous attractor networks. It has recently been shown how idiothetic (self-motion) inputs could update the activity packet of neuronal firing for a one-dimensional case (head direction cells), and for a two-dimensional case (place cells which represent the place where a rat is located). In this paper, we describe three models of primate hippocampal spatial view cells, which not only maintain their spatial firing in the absence of visual input, but can also be updated in the dark by idiothetic input. The three models presented in this paper represent different ways in which a continuous attractor network could integrate a number of different kinds of velocity signal (e.g., head rotation and eye movement) simultaneously. The first two models use velocity information from head angular velocity and from eye velocity cells, and make use of a continuous attractor network to integrate this information. A fundamental feature of the first two models is their use of a ‘memory trace’ learning rule which incorporates a form of temporal average of recent cell activity. Rules of this type are able to build associations between different patterns of neural activities that tend to occur in temporal proximity, and are incorporated in the model to enable the recent change in the continuous attractor to be associated with the contemporaneous idiothetic input. The third model uses positional information from head direction cells and eye position cells to update the representation of where the agent is looking in the dark. In this case the integration of idiothetic velocity signals is performed in the earlier layer of head direction cells.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Spatial view cells; Continuous attractor networks; Self-organization; Trace learning; Idiothetic input; Hippocampus

1. Introduction

In primates, single neuron recording studies have demonstrated the existence of spatial view cells in the primate hippocampus that respond when the monkey is looking towards a particular location in allocentric space (Georges-François, Rolls, & Robertson, 1999; Robertson, Rolls, & Georges-François, 1998; Rolls, 1999; Rolls, Robertson,

& Georges-François, 1997; Rolls, Treves, Robertson, Georges-François, & Panzeri, 1998). Spatial view cells update their spatial representations by idiothetic inputs in the dark, in that if the monkey moves his head and eyes to look at the effective spatial location, then the neurons fire (Robertson et al., 1998). Part of the interest of these primate spatial view neurons is that they represent a place at which a primate is looking, and could therefore be involved in functions such as providing the spatial representation needed for remembering where an object is in space, and in general for representing places at which one is not actually located (Rolls, 1999; Rolls et al., 2002).

* Corresponding author. Fax: +44 1865 310447.

E-mail address: edmund.rolls@psy.ox.ac.uk (E.T. Rolls).

URL: www.cns.ox.ac.uk.

An established approach to modelling neurons which appear to encode spatial information is the continuous attractor neural network (CANN). Continuous attractor networks have previously been used for modelling head direction cells (Hahnloser, 2003; Redish, Elga, & Touretzky, 1996; Sharp, Blair, & Cho, 2001; Skaggs, Knierim, Kudrimoti, & McNaughton, 1995; Zhang, 1996) and place cells (Redish, 1999; Redish & Touretzky, 1998; Samsonovich & McNaughton, 1997; Tsoodyks, 1999). (Head direction cell properties in the rat are described by Taube, Muller, & Ranck (1990), and in macaques by Robertson, Rolls, Georges-François, & Panzeri (1999)). This class of network can maintain the firing of its neurons to represent any location along a continuous physical dimension such as head direction, spatial location, or spatial view. These models use excitatory recurrent collateral connections between the neurons to reflect the distance between the neurons in the state space (e.g., head direction space) of the animal. Global inhibition is used to keep the number of neurons in a bubble of activity relatively constant, and to ensure there is only one activity packet. The properties of these CANNs have been extensively studied, for example by Amari (1977) and Taylor (1999). They can maintain the packet of neural activity (i.e., the set of active neurons that represent a spatial location) constant for long periods. The recurrent weights within a continuous attractor network can be self-organised during learning using a Hebbian associative learning rule, as proposed by Muller, Kubie, and Saypoff (1991). A further property of some continuous attractor models is their ability to perform path integration. This is the ability of the network to update its current spatial representation using idiothetic (self-motion) signals. Recently, we have developed and simulated continuous attractor network models of head direction cells (Stringer, Trappenberg, Rolls, & de Araujo, 2002b) and place cells (Stringer, Rolls, Trappenberg, & de Araujo, 2002a), which are able to self-organise Sigma–Pi idiothetic synaptic connection strengths (responsible for path integration) during training of the network using a *trace* learning rule.

Although, spatial view cells are a conceptually new type of representation now believed to be present in the primate hippocampus, there has so far been no theoretical investigations of the nature of the representations they provide, and of how the representation of places out-there could be updated by self-motion, that is by idiothetic cues. As noted above, spatial view cells do keep firing in the dark, and are updated by eye-position changes made by monkeys (Robertson et al., 1998). Because, as described above, CANNs are the natural model of such spatial representations and have been extended to incorporate hypotheses about how the spatial representations could be updated by idiothetic input, they are the natural choice to model hippocampal spatial view cells. In this paper, we provide the first models

of, and have as one of our important conceptual aims to elucidate, how different types of neural network architecture might aim to account for the data.

The inputs used for the idiothetic update of spatial cells in the primate hippocampus must be very different from those in rats, for in primates the allocentric location is idiothetically updated by eye position as well as by head direction (Robertson et al., 1998), and no eye position update has ever been suggested for the rat. The aim of this paper is to consider how different types of input that reflect eye position and head direction information could be utilized in a path integration CANN. A major challenge is that two idiothetic signals, which separately reflect eye position and head direction, need to be combined in the path integration process, and this has not been treated in previous self-organizing continuous attractor models.

We now provide an overview of the three approaches in this paper. Model 1 utilises recurrent connections between spatial view cells in a continuous attractor network, and updates the spatial view cell firing in the dark using idiothetic velocity signals from body or head rotation cells and eye velocity cells. A key feature of Model 1 is that the synaptic connections to spatial view cells from the different classes of idiothetic cells operate independently. Model 2 is a continuous attractor network model in which the synaptic connections to the spatial view cells from the different classes of idiothetic cells are first combined multiplicatively within higher order Sigma–Pi synapses (Koch, 1999). This allows the network to learn how to shift the activity packet within the layer of spatial view cells in the correct direction for precise combinations of idiothetic signals. Model 3 is a feedforward network model of spatial view cells in which the firing pattern in the spatial view cell layer is updated in the dark by inputs from head direction cells and eye position cells. Here, the key continuous attractor dynamics, where neural activity patterns representing the state of the agent are updated directly from idiothetic velocity signals, occur in an earlier layer of head direction cells which sends inputs to the spatial view cell layer. We note that head direction cells are present in the primate hippocampal system, in the presubiculum (Robertson et al., 1999), as are cells that represent whole body rotation (O'Mara, Rolls, Berthoz, & Kesner, 1994). We note that Model 3 effectively builds neurons that respond to a combination of head direction and eye position, and so reflect gaze angle, and thus also provides a model for cells which reflect gaze angle recorded in the primate parietal cortex (Xing & Andersen, 2000).

The three Models presented in this paper are aimed at showing generically how networks could perform the idiothetic update of a spatial view cell representation. The generic approach leads us to describe three different ways of solving the computational problems involved. The neuron types and computational processes found in

each model are closely related to knowledge of what is represented in different brain regions, as follows. One possible brain region for Models 1 and 2 is the CA3 region of the primate hippocampus, which contains spatial view cells (Georges-Francois et al., 1999; Robertson et al., 1998; Rolls, 1999; Rolls et al., 1997), which show at least some continuing firing in the dark (Robertson et al., 1998). The CA3 cells have a highly developed recurrent collateral system of connections (Ishizuka, Weber, & Amaral, 1990), which show at least associative synaptic modification (Debanne, Gähwiler, & Thompson, 1998). Moreover, whole body rotation cells, which provide the firing represented in Models 1 and 2 by the head rotation cells, are present in the primate hippocampus (O'Mara et al., 1994). Although eye velocity-related neuronal firing has not been reported for the hippocampus, eye velocity is the main form in which signals in the brain related to eye movements are represented (Berthoz, 2000; Berthoz, Grantyn, & Droulez, 1987). Although the CA1 cells of primates do fire better than CA3 cells during update of spatial view cells by eye movements in the dark, this could be because of associative improvement of the representation at the Schaffer collateral (CA3 to CA1) synapses (Robertson et al., 1998). The CA1 region is a less likely region for implementation of Models 1 and 2, in that it has a less well developed recurrent collateral system than does the CA3 system (Rolls & Treves, 1998). It is also possible that Models 1 and 2 are implemented at a stage of neural processing prior to the hippocampus, and that the hippocampal pyramidal cells reflect the inputs received from

this preceding network. It is known that there are spatial view cells in the parahippocampal gyrus (Rolls, 1999; Rolls et al., 1997). With respect to Model 3, if the idiothetic update of the correct spatial view cell firing in the dark is performed in the hippocampus in the way described in that Model, it is relevant that one hippocampal region with feedforward rather than recurrent collateral connectivity is the CA1 cell system. For the cells in CA1 to be updated directly by idiothetic inputs as described for Model 3, the CA1 cells would need to receive head direction cell inputs (which are represented by the firing of neurons in the primate presubiculum which are themselves updated correctly by head rotation in the dark (Robertson et al., 1999)), and eye position inputs.

2. Model 1: The recurrent network model with independent idiothetic inputs

2.1. Architecture and learning in the network

The network architecture is shown in Fig. 1. When visual input is available, each spatial view cell responds with a Gaussian profile to a view of part of the environment, and that part of the environment can be looked at with a given gaze angle which is provided by a combination of head direction and eye position signals. The spatial view cells are connected by recurrent collateral synapses that reflect the distance in the state space (in this case the spatial view) of any two connected cells.

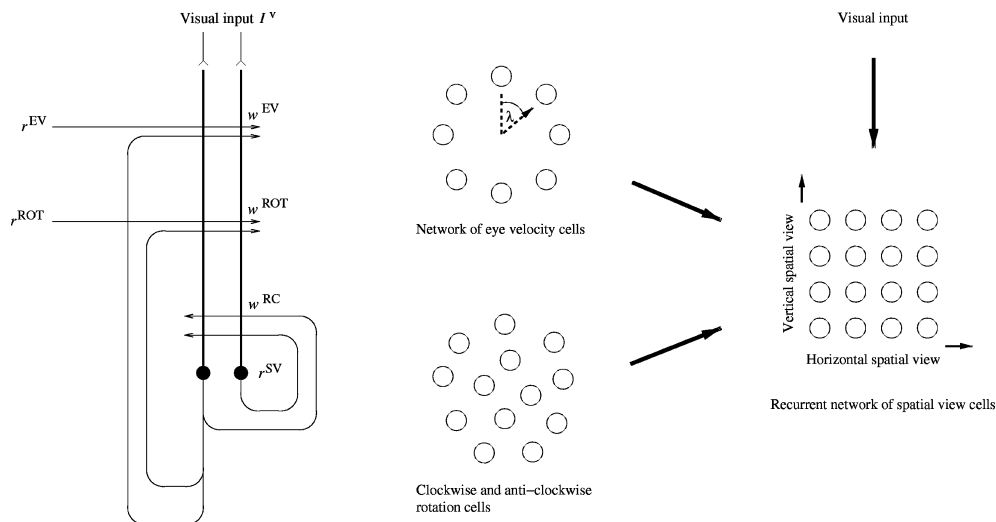


Fig. 1. General network architecture for Model 1, the recurrent continuous attractor network model with independent idiothetic inputs. There is a recurrent layer of spatial view cells in which the cells are mapped onto a regular grid of spatial views. The spatial views are defined by the angle of gaze of the agent. In the light, individual spatial view cells are stimulated maximally when the agent has an angle of gaze corresponding to the position of the cell in the grid. The layer of spatial view cells receives external inputs from three sources: (i) the visual system; (ii) a population of clockwise and anticlockwise rotation cells; and (iii) a population of eye velocity cells. The clockwise and anti-clockwise rotation cells fire according to whether the agent is rotating in the appropriate direction, and with a firing rate that increases monotonically with respect to the speed of rotation. The eye velocity cells fire maximally when an animal moves its eyes in a particular direction, and have firing rates that increase with the speed of movement of the eyes. The velocity of the animal's eyes is denoted by (v, λ) where v is the angular speed of movement of the eyes, and λ is the direction of movement of the eyes. The eye velocity cells are mapped onto a circular grid of directions of eye movement, where each eye velocity cell has a unique direction of eye movement for which the cell fires maximally.

The strengths of these synaptic connections of the continuous attractor network are trained by an associative Hebb rule while the agent explores the environment. The correct synaptic weights are set up because spatial view cells that are close in view space will tend to be coactive (Zhang, 1996, Redish & Touretzky, 1998). That is, during learning the recurrent weights w_{ij}^{RC} from spatial view cell j with firing rate r_j^{SV} to spatial view cell i with firing rate r_i^{SV} are altered according to

$$\delta w_{ij}^{\text{RC}} = k^{\text{RC}} r_i^{\text{SV}} r_j^{\text{SV}}, \quad (1)$$

where $\delta w_{ij}^{\text{RC}}$ is the change of synaptic weight, and k^{RC} is the learning rate constant.

Now we consider the self-organization of the idiothetic synaptic weights w_{ijk}^{ROT} from the head rotation cells. The essence of this learning process is that when the activity packet in the spatial view cell continuous attractor has moved, say, in a clockwise direction, a trace term in a set of synaptic connections between the spatial view cells “remembers” the direction in which the spatial view cells have been activated, and this term is associated by learning with the current head rotation velocity cell input and the firing of the postsynaptic spatial view cell receiving both inputs. Then after learning, if the head rotation velocity signal is present simultaneously with an input from a set of spatial view cells, the asymmetry in the idiothetic rotation synaptic weights in a particular direction with respect to the spatial view cells produces extra activation in one direction in the spatial view cell space, and the packet of neuronal activity moves in the correct direction in the spatial view (state) space. More formally, the synaptic weights w_{ijk}^{ROT} are updated at each timestep during motion through the environment in a self-organizing learning process according to

$$\delta w_{ijk}^{\text{ROT}} = k^{\text{ROT}} r_i^{\text{SV}} \bar{r}_j^{\text{SV}} r_k^{\text{ROT}}, \quad (2)$$

where $\delta w_{ijk}^{\text{ROT}}$ are the changes in the synaptic weights, and where r_i^{SV} is the instantaneous firing rate of the postsynaptic spatial view cell i , \bar{r}_j^{SV} is the trace value of the presynaptic spatial view cell j given by Eq. (3), r_k^{ROT} is the firing rate of rotation cell k , and k^{ROT} is the learning rate associated with this type of synaptic connection. In Eq. (2) \bar{r}^{SV} is a local temporal average or memory trace value of the firing rate of a spatial view cell given by

$$\bar{r}^{\text{SV}}(t + \delta t) = (1 - \eta) r^{\text{SV}}(t + \delta t) + \eta \bar{r}^{\text{SV}}(t), \quad (3)$$

where η is a parameter set in the interval $[0, 1]$ which determines the relative contributions of the current firing and the previous trace.

Next we consider the idiothetic synaptic weights, w_{ijk}^{EV} , to spatial view cells from the eye velocity cells k . The learning phase involves setting up the synaptic weights w_{ijk}^{EV} for all ordered pairs of spatial view cells i and j , and for all eye velocity cells k . As the agent moves its eyes during learning, the synaptic weights w_{ijk}^{EV} are updated at each timestep according to

$$\delta w_{ijk}^{\text{EV}} = k^{\text{EV}} r_i^{\text{SV}} \bar{r}_j^{\text{SV}} r_k^{\text{EV}}, \quad (4)$$

where $\delta w_{ijk}^{\text{EV}}$ is the change of synaptic weight, r_i^{SV} is the instantaneous firing rate of spatial view cell i , \bar{r}_j^{SV} is the trace value of the firing rate of spatial view cell j given by Eq. (3), r_k^{EV} is the firing rate of eye velocity cell k , and k^{EV} is the learning rate associated with this type of connection. Learning rule (4) for the synapses from the eye velocity cells operates in an analogous way to learning rule Eq. (2) for the synapses from the rotation cells.

2.2. The dynamical equations of the spatial view cells

The following equation describes the ‘leaky-integrator’ dynamics of the activation $h_i^{\text{SV}}(t)$ of each spatial view cell i using the above terms:

$$\begin{aligned} \tau \frac{dh_i^{\text{SV}}(t)}{dt} = & -h_i^{\text{SV}}(t) + \frac{\phi_0}{C^{\text{SV}}} \sum_j (w_{ij}^{\text{RC}} - w^{\text{INH}}) r_j^{\text{SV}}(t) \\ & + I_i^{\text{V}} + \frac{\phi_1}{C^{\text{SV} \times \text{ROT}}} \sum_{j,k} w_{ijk}^{\text{ROT}} r_j^{\text{SV}} r_k^{\text{ROT}} \\ & + \frac{\phi_2}{C^{\text{SV} \times \text{EV}}} \sum_{j,k} w_{ijk}^{\text{EV}} r_j^{\text{SV}} r_k^{\text{EV}}. \end{aligned} \quad (5)$$

The term r_j^{SV} is the firing rate of spatial view cell j , w_{ij}^{RC} is the excitatory (positive) synaptic weight from spatial view cell j to spatial view cell i , and w^{INH} is a global constant describing the effect of inhibitory interneurons within the layer of spatial view cells.¹

The further terms in Eq. (5) are as follows. The term I_i^{V} represents a visual input to spatial view cell i , and τ is the time constant of the system. When the agent is in the dark, the term I_i^{V} is set to zero. Thus, in the absence of visual input there are two key input terms driving the cell activations in Eq. (5) as follows. First, there are idiothetic inputs from the rotation cells, $\sum_{j,k} w_{ijk}^{\text{ROT}} r_j^{\text{SV}} r_k^{\text{ROT}}$, where r_k^{ROT} is the firing rate of rotation cell k , and w_{ijk}^{ROT} is the corresponding overall effective strength of connection from this cell.² Second, there is an idiothetic input from the eye velocity cells $\sum_{j,k} w_{ijk}^{\text{EV}} r_j^{\text{SV}} r_k^{\text{EV}}$, where r_k^{EV} is the firing rate of eye velocity cell k and w_{ijk}^{EV} is the corresponding overall effective strength of connection.³

¹ The scaling factor (ϕ_0/C^{SV}) controls the overall strength of the recurrent inputs to the layer of spatial view cells, where ϕ_0 is a constant and C^{SV} is the number of synaptic connections received by each spatial view cell from other spatial view cells.

² For the idiothetic inputs from the rotation cells, the scaling factor $\phi_1/C^{\text{SV} \times \text{ROT}}$ controls the overall strength of the rotation idiothetic inputs, where ϕ_1 is a constant, and the term $C^{\text{SV} \times \text{ROT}}$ is defined as the number of idiothetic connections received by each spatial view cell from couplings of individual spatial view cells and rotation cells.

³ For the idiothetic inputs from the eye velocity cells, the strength of these inputs is controlled by the scaling factor $\phi_2/(C^{\text{SV} \times \text{EV}})$, where ϕ_2 is a constant and $C^{\text{SV} \times \text{EV}}$ is the number of idiothetic connections received by each spatial view cell from combinations of spatial view cells and eye velocity cells.

Once the activations of the spatial view cells have been computed, the firing rates of the cells are then given by the sigmoid transfer function

$$r_i^{\text{SV}}(t) = \frac{1}{1 + e^{-2\beta(h_i^{\text{SV}}(t) - \alpha)}}, \quad (6)$$

where α and β are the sigmoid threshold and slope, respectively.

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

The recurrent synaptic weights within the continuous attractor network may not be formed perfectly if the training is not performed with perfect regularity for every node in the continuous attractor. (In this sense the synaptic weights have static noise.) This in turn can lead to drift of the activity packet within the continuous attractor network of spatial view cells when there are no visual cues available. Although this drift is a property in the dark of rat head direction and place cells (see (Knierim, Kudrimoti, & McNaughton, 1998)) and of spatial view cells in primates (Robertson et al., 1998), the drift needs to be kept small if the network is to be useful. Stringer et al. (2002b) proposed that in real nervous systems the problem of drift may be minimized by enhancing the firing of neurons that are already firing. This might be implemented through mechanisms for short term synaptic enhancement (Koch, 1999), or through the effects of neuronal voltage dependent ion channels such as those influenced by NMDA receptors. The non-linearity of the activation function of neurons with NMDA receptors results in these receptors only contributing to neuronal firing once the neuron is sufficiently depolarized (Wang, 1999). The effect is to enhance the firing of neurons that are already reasonably well activated.

In Model 1 of spatial view cells, we simulate these effects by resetting the sigmoid threshold α_i for each spatial view cell i at each timestep depending on the firing rate of spatial view cell i at the previous timestep. That is, at each timestep $t + \delta t$ we set

$$\alpha_i = \begin{cases} \alpha^{\text{HIGH}} & \text{if } r_i^{\text{SV}}(t) < \gamma, \\ \alpha^{\text{LOW}} & \text{if } r_i^{\text{SV}}(t) \geq \gamma, \end{cases} \quad (7)$$

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

2.4. Simulation results with Model 1

The operation and properties of Model 1 were investigated numerically by simulation. The model param-

Table 1
Parameter values for Model 1

Learning rates k^{RC} , k^{ROT} , and k^{EV}	0.001
Trace parameter η	0.9
τ	1.0
ϕ_0	50,000
ϕ_1	164,500
ϕ_2	1,175,000
w^{INH}	0.06
γ	0.5
α^{HIGH}	0.0
α^{LOW}	-20.0
β	0.1

eters used in the simulations are given in Table 1. During a training phase, the agent either rotated its head, or moved its eyes. For this simulation the agent rotated on one spot, so that the spatial view of the environment corresponded to the gaze angle, which is defined as the algebraic sum of the head direction and eye position values. The training was performed for every spatial view with a fixed eye position during head rotation in both directions with a constant velocity (to train the idiothetic head rotation synaptic weights); and for every spatial view during eye movement in eight principal directions with a constant velocity and a fixed head direction (to train the idiothetic eye rotation synaptic weights). The simulation training procedure was to set $\tau = 1$, and then to set the scaling of the recurrent synaptic weights to produce stable packets of neural activity, and then to set the scaling of the idiothetic weights to move the packet without disrupting it. It was not critical for any of the parameters to be adjusted within a small tolerance. All the Models were robust in this sense.

Fig. 2 shows performance after training. On the left it is shown that if there are no visual and idiothetic inputs, there is a stable firing rate profile within the continuous attractor network of spatial view cells. It is shown in the right plot that the spatial view activity packet moves in the correct direction (in the absence of visual inputs) when the agent: (i) rotates its head clockwise; (ii) moves its eyes upwards; and (iii) simultaneously rotates its head clockwise and moves its eyes upwards. This important result demonstrates that, after the synaptic weights have been set up through learning, the two different types of idiothetic inputs to the layer of spatial view cells, from the rotation cells, and from the eye velocity cells, are able to operate together as the agent rotates and moves its eyes in the dark.

To clarify how the network is able to perform this path integration, we show examples of the recurrent and idiothetic synaptic weights after training in Fig. 3. The recurrent weights between the spatial view cells in the continuous attractor are symmetric with respect to any one neuron. (Neuron 25 was chosen.) This is a condition for stability of the continuous attractor activity

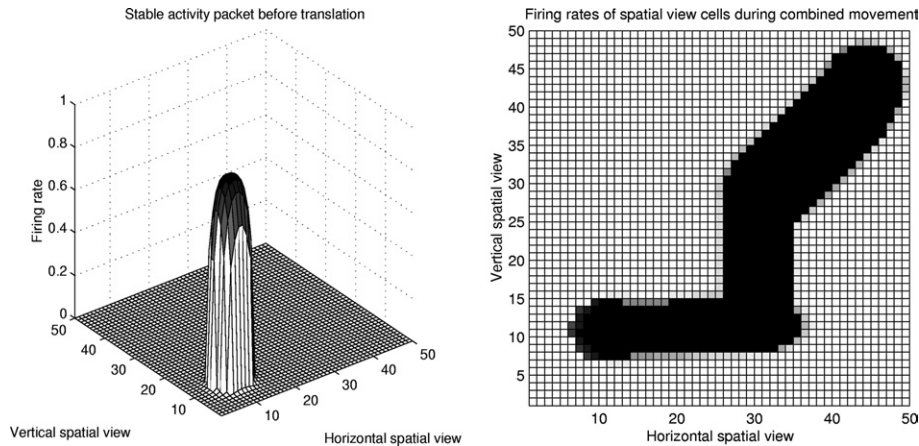


Fig. 2. Firing rate profiles from Model 1 within the continuous attractor network of spatial view cells during the testing phase with the agent in the dark. The grid shows the horizontal and vertical coordinates at which the different spatial view cells have their optimal firing. (Left) Stable firing rate profile within the network of spatial view cells before the agent starts to move. (Right) Maximum firing rates that occurred during movement of the agent in the testing phase. (The maximum firing rate is calculated over all timesteps for each spatial view cell.) First, the agent rotated its head in the clockwise direction. Next the agent moved its eyes vertically upwards. Finally, the agent simultaneously rotated its head in the clockwise direction and moved its eyes vertically upwards.




Fig. 3. Synaptic weight profiles from Model 1 plotted along a horizontal line through the centre of the spatial view area. The plot compares the recurrent and idiothetic synaptic weight profiles as follows. The first graph shows the recurrent weights w_{ij}^{RC} , where the pre-synaptic spatial view cell j is the neuron set to fire maximally when the agent's angle of gaze is at the centre of the spatial view area during the learning phase, and the post-synaptic spatial view cells i are those neurons set to fire maximally at various positions along the horizontal line through the centre of the area. The second graph shows the idiothetic weights w_{ij}^{ROT} from the clockwise rotation cell, where the pre- and post-synaptic spatial view cells j and i are as above. The third graph shows the idiothetic weights w_{ijk}^{EV} from the eye velocity cell which fires maximally when the agent's eyes move horizontally to the right (i.e., $\lambda = 90^\circ$), and where the pre- and post-synaptic spatial view cells j and i are as above.

packet in the absence of any inputs (whether visual or idiothetic). The width of the profile of the synaptic weights from any one neuron reflects the Gaussian tuning to spatial view of the spatial view cells.

The idiothetic weights from the clockwise rotation cell and spatial view neuron 25 are shown, and are

asymmetric. This asymmetry is essential to the process by which for example clockwise rotation inputs drive spatial view cells towards the right of the current activity packet. The idiothetic weights from the rightward eye velocity cell and spatial view neuron 25 are also shown, and are again asymmetric. This asymmetry is essential to the process by which for example rightward eye rotation drives spatial view cells towards the right of the current spatial view cell activity packet. (An account of how this occurs is provided in Section 2.1, and further details of how this type of idiothetic path integration operates in the simpler cases of head direction and place cells are provided by Stringer et al. (2002a) and Stringer et al. (2002b).)

2.5. The problem of training with combined movements

In the above simulations with Model 1, the network was trained with the agent either rotating its head, or moving its eyes, but not performing both kinds of movement together. However, a serious problem occurs when the agent is free to perform more than one movement at a time during learning. For example, consider the case where the agent may rotate its head clockwise and move its eyes vertically upwards at the same time. In this case the network will associate the overall movement of the agent's spatial view, which will be a diagonal path in a direction of 45° as shown in the right plot of Fig. 2, with each of the two separate idiothetic velocity signals, clockwise head rotation, and vertical eye movement. However, the association is made for each kind of idiothetic signal independently. This means that in the above situation, the clockwise rotation cells will develop a tendency to drive the activity packet in a direction of 45° , even when the eye velocity cells are not firing. Sim-

ilarly, the eye velocity cells which are tuned to eye movement vertically upwards will learn to drive the activity packet in a direction of 45°, even when the head rotation cells are not firing. In the next section we present Model 2, which is able to solve the problem of combined movements during learning by associating specific combinations of movement signals with the shift of the activity packet in the continuous attractor network of spatial view cells.

3. Model 2: The recurrent network model with combined idiothetic inputs

3.1. Architecture and learning in the network

The network architecture of Model 2 is shown in Fig. 4. There are two key differences between Model 1 and 2. First, although the layer of spatial view cells receives idiothetic signals from head rotation cells and eye velocity cells, with Model 2 the inputs from the different idiothetic cell types are combined multiplicatively within Sigma–Pi synapses. Second, the idiothetic cells must




Fig. 4. General network architecture for Model 2, the recurrent continuous attractor network model with combined idiothetic inputs. As described for Model 1, there is a recurrent layer of spatial view cells which receives external inputs from three sources: (i) the visual system; (ii) a population of clockwise and anti-clockwise rotation cells; and (iii) a population of eye velocity cells. However, with Model 2 the inputs from the different idiothetic cell types (rotation cells and eye velocity cells) are combined multiplicatively within Sigma–Pi synapses w_{ijkl}^{ROTEV} . Also, for Model 2 the idiothetic cells fire maximally for different speeds of movement, as well as for different kinds of movement. Specifically, the rotation cells fire according to whether the agent is rotating in the appropriate direction, with each rotation cell firing maximally for a specific speed of rotation. Similarly, the eye velocity cells fire maximally when an animal moves its eyes in a particular direction, with each eye velocity cell firing maximally for a particular speed of eye movement.

behave somewhat differently for Model 2. For Model 2, the idiothetic cells must fire maximally for different speeds of movement, as well as for different kinds of movement. Specifically, the head rotation cells fire according to whether the agent is rotating in the appropriate direction, with each rotation cell firing maximally for a specific speed of rotation. Moreover, the model requires the population of rotation cells to cover all possible speeds of rotation, including some cells firing maximally when the agent is not rotating. Similarly, the eye velocity cells fire maximally when an animal moves its eyes in a particular direction, with each eye velocity cell firing maximally for a particular speed of eye movement in a given direction. The model requires the population of eye velocity cells to cover all possible speeds of eye movement, including some cells firing maximally when the agent is not moving its eyes.

During learning the recurrent synaptic weights w_{ij}^{RC} between the spatial view cells are updated according to the same associative (Hebb) rule (1) used in Model 1. However, for Model 2 the different idiothetic signals are combined multiplicatively within higher order Sigma–Pi synapses. Thus, for Model 2 the idiothetic synaptic weights w_{ijkl}^{ROTEV} are updated at each timestep according to

$$\delta w_{ijkl}^{\text{ROTEV}} = k^{\text{ROTEV}} r_i^{\text{SV}} \bar{r}_j^{\text{SV}} r_k^{\text{ROT}} r_l^{\text{EV}}, \quad (8)$$

where $\delta w_{ijkl}^{\text{ROTEV}}$ are the changes in the synaptic weights, and where r_i^{SV} is the instantaneous firing rate of the post-synaptic spatial view cell i , \bar{r}_j^{SV} is the trace value of the presynaptic spatial view cell j given by Eq. (3), r_k^{ROT} is the firing rate of rotation cell k , r_l^{EV} is the firing rate of eye velocity cell l , and k^{ROTEV} is the learning rate associated with this type of synaptic connection. Learning rule (8) is able to associate the co-firing of head rotation cell k and eye velocity cell l with the shift of the activity packet from spatial view cell j to spatial view cell i . This multiplicative combination of the different idiothetic velocity signals within the idiothetic synapses w_{ijkl}^{ROTEV} permits the agent to perform more than one type of movement (e.g., head rotation and eye movement) at a time during learning.

3.2. The dynamical equations of the spatial view cells

We now describe the behaviour of the layer of spatial view cells for Model 2. The activation $h_i^{\text{SV}}(t)$ for each spatial view cell i is updated according to the equation

$$\tau \frac{dh_i^{\text{SV}}(t)}{dt} = -h_i^{\text{SV}}(t) + \frac{\phi_0}{C^{\text{SV}}} \sum_j (w_{ij}^{\text{RC}} - w^{\text{INH}}) r_j^{\text{SV}}(t) + I_i^{\text{V}} + \frac{\phi_1}{C^{\text{SV} \times \text{ROT} \times \text{EV}}} \sum_{j,k,l} w_{ijkl}^{\text{ROTEV}} r_j^{\text{SV}} r_k^{\text{ROT}} r_l^{\text{EV}}. \quad (9)$$

In the absence of visual input, the activity packet within the continuous attractor network of spatial view cells is

driven by the last term of Eq. (9) which represents the idiothetic inputs from particular combinations of head rotation cells and eye velocity cells.⁴ In Eq. (9) the idiothetic signals are combined into a single multiplicative term $r_k^{\text{ROT}} r_l^{\text{EV}}$. This multiplicative form for the idiothetic input to the continuous attractor network of spatial view cells introduces a logical AND function between the different idiothetic signals. Thus, when the network is tested in the dark, the activity packet is shifted according to which exact combination of idiothetic signals are present.

3.3. Simulation results with Model 2

For Model 2 the agent must perform every possible combination of movements (e.g., head rotation in both clockwise and anticlockwise directions, and eye movement in all directions) with its fixation position moving through every visual location in the spatial view area. In particular, each movement combination is defined by not only which kinds of movements (e.g., clockwise head rotation and vertical eye movement) are being made, but also the velocities of the different component movements. However, since the agent may typically only perform combinations of movements from within a relatively small portion of the space of all different combinations of movements, in practice the network would only need to be trained on this small subspace of combinations of movements.

Briefly, the learning phase for Model 2 proceeded in a number of stages, where for each stage of training the agent performed a particular combined movement. The combined movements involved a combination of a head rotation (in either the clockwise or anticlockwise directions), with an eye movement (in one of the eight principle compass directions). However, for each combined movement, the speeds of the separate components (i.e., head rotation or eye movement) may each be either zero or a constant non-zero value. This ensured that the network was trained on different combinations of head rotation and eye movements.

The results of the numerical simulations of Model 2 are shown in Fig. 5, and are similar to those shown earlier for Model 1. Thus the simulations confirmed that Model 2 was able to correctly update the firing in the network of spatial view cells in the absence of visual input, as the agent rotated its head and moved its eyes, even when Model 2 had been trained on concurrent head rotation and eye velocity idiothetic inputs.

⁴ The term $C^{\text{SV} \times \text{ROT} \times \text{EV}}$ is defined as the number of idiothetic connections received by each spatial view cell from couplings of individual spatial view cells, rotation cells, and eye velocity cells.

3.4. The problem of generalisation over different speeds of movement

Although the approach of combining the different kinds of idiothetic signals within higher order Sigma–Pi synapses allows Model 2 to cope with the agent performing combined movements during learning, Model 2 requires significantly more training than Model 1. This is because Model 2 is unable to generalize over different speeds of movement. That is, training Model 2 at one speed of movement, say slow clockwise rotation, will not allow the model to generalize to other speeds of movement, e.g., fast clockwise rotation, during testing. In the next section we present Model 3, which is able to solve the problem of combined movements during learning, without the need for training with all possible combinations of speeds, by incorporating multiple continuous attractor networks, each of which integrates a single kind of idiothetic signal.

4. Model 3: The feedforward network model

4.1. Architecture and learning in the network

Model 3 is a model of spatial view cells in which the firing pattern in the spatial view cell layer is updated in the dark by inputs from head direction cells and eye position cells. In this case, the spatial view cells receive inputs carrying information about the current positional state of the agent (in terms of head direction and eye position) rather than information about the *rate of change* of the state of the agent. The integration of idiothetic velocity signals is performed in an earlier layer as described below, while the spatial view cells are driven by forward inputs carrying information about the positional state of the agent. In particular, spatial view cells are driven by appropriate combinations of head direction and eye position. Model 3, while itself in terms of the architecture illustrated in Fig. 6 is a feedforward model, does rely in the dark on inputs from a layer of head direction cells, where the head direction cells, themselves, operate as a continuous attractor network that integrates idiothetic head rotation velocity signals. The self-organisation and operation of this earlier layer of head direction cells is described in Stringer et al. (2002b).

The general network architecture of Model 3 is as shown in Fig. 6. The layer of spatial view cells receives external inputs from three sources: (i) the visual system; (ii) a population of head direction cells; and (iii) a population of eye position cells.

During learning the recurrent synaptic weights w_{ij}^{RC} between the spatial view cells may be updated according to the same learning rule as described for Model 1. That is, the recurrent synaptic weights w_{ij}^{RC} are updated by the associative (Hebb) rule (1). The synaptic weights w_{ikl}^{D} to

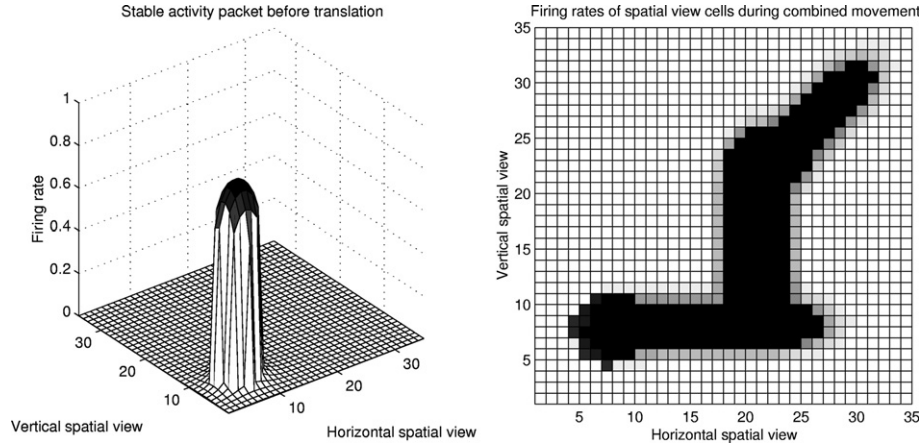


Fig. 5. Firing rate profiles from Model 2 (the recurrent continuous attractor network model with combined idiothetic inputs) from neurons within the continuous attractor network of spatial view cells during the testing phase with the agent in the dark. The grid shows the horizontal and vertical coordinates at which the different spatial view cells have their optimal firing. (Left) Stable firing rate profile within the network of spatial view cells before the agent starts to move. (Right) Maximum firing rates that occurred during movement of the agent in the testing phase. (The maximum firing rate is calculated over all timesteps for each spatial view cell.) First, the agent rotated its head in the clockwise direction. Next the agent moved its eyes vertically upwards. Finally, the agent simultaneously rotated its head in the clockwise direction and moved its eyes vertically upwards.




Fig. 6. General network architecture for Model 3, the feedforward network model for combining head direction and eye position signals to provide spatial view representations. There is a layer of spatial view cells which receives external inputs from three sources: (i) the visual system; (ii) a population of head direction cells; and (iii) a population of eye position cells. Individual head direction cells fire maximally when the agent has a particular head direction θ . The population of head direction cells are mapped onto a circular grid of head directions, where each head direction cell has a unique head direction for which the cell fires maximally. The eye position cells fire maximally for a particular angular position (orientation) of the agent’s eyes, which is denoted by (ω^h, ω^v) , where ω^h and ω^v are the horizontal and vertical angles, respectively. The eye position cells are mapped onto a regular grid of eye positions, where individual eye position cells are stimulated maximally when the eye position of the agent corresponds to the position of the cell in the grid.

spatial view cells from combinations of head direction cells and eye position cells in the earlier layers are updated according to

$$\delta w_{ikl}^{ID} = k^{ID} r_i^{SV} r_k^{HD} r_l^{EP}, \tag{10}$$

where δw_{ikl}^{ID} is the change of synaptic weight, r_i^{SV} is the firing rate of the postsynaptic spatial view cell i , r_k^{HD} is the firing rate of head direction cell k , r_l^{EP} is the firing rate of eye position cell l , and k^{ID} is the learning rate

constant. This rule has the effect of associating a particular combination of head direction and eye position of the agent with a particular spatial view representation.

4.2. The dynamical equations of the spatial view cells

We now describe in more detail the firing behaviours of the spatial view cells for Model 3. In the equations of the model presented below, for the sake of generality we

include recurrent connections within the layer of spatial view cells. However, the recurrent connections are not essential, and the model can operate with or without the recurrent connections. We assume the behaviour of the layer of spatial view cells is governed by the following ‘leaky-integrator’ dynamical equations. The following equation describes the dynamics of the activation $h_i^{SV}(t)$ of each spatial view cell i :

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

The term r_j^{SV} is the firing rate of spatial view cell j , w_{ij}^{RC} is the excitatory (positive) synaptic weight from spatial view cell j to spatial view cell i , and w^{INH} is a global constant describing the effect of inhibitory interneurons within the layer of spatial view cells.⁵ However, a key feature of Model 3 is that the recurrent connections within the layer of spatial view cells are not required, and so ϕ_0 may in fact be set to zero. In this case the model operates in a feedforward manner, with no recurrent connections within the layer of spatial view cells. Further terms in Eq. (11) are as follows. The term I_i^V represents a visual input to spatial view cell i , and τ is the time constant of the system. When the agent is in the dark, the term I_i^V is set to zero. Thus, in the absence of visual input the key input term driving the cell activations in Eq. (11) becomes $\frac{\phi_1}{C^{HD \times EP}} \sum_{k,l} w_{ikl}^{ID} r_k^{HD} r_l^{EP}$, where r_k^{HD} is the firing rate of head direction cell k , r_l^{EP} is the firing rate of eye position cell l , and w_{ikl}^{ID} is the corresponding overall effective strength of the connection from these cells.⁶

Eq. (11) for calculating the cell activation assumes the inputs from the head direction cells and eye position cells operate in a multiplicative manner using Sigma–Pi synaptic connections. Sigma–Pi neurons sum the products of the contributions from two or more sources (see Koch (1999, Section 21.1.1), Rolls & Treves (1998)). The reason for using Sigma–Pi synaptic connections, instead of a more usual additive summation of the inputs, is as follows. Consider, for example, the agent both rotating its head and moving its eyes such that its angle of gaze moves along a horizontal track within the 1×1 unit containment area described above, where 1 unit is equal to, say, 35° . During the learning phase

the agent may fixate at any particular location along the horizontal line either with any head direction or with any horizontal eye position. Therefore, if the model used simple independent Hebbian associative learning between the head direction cells and spatial view cells, and between the eye position cells and spatial view cells, then every head direction cell could become associated with every spatial view cell, and also every eye position cell could become associated with every spatial view cell. In this situation the synapses would be functionally useless, and if the spatial view cells were driven during the subsequent testing phase in the dark simply by the sum of the two separate contributions from the head direction cells and eye position cells, then the resulting spatial view cell firing rates would not be capable of accurately reflecting the true spatial view of the agent. In order for the firing rates of the spatial view cells to reflect the true spatial view of the agent, the spatial view cell activations must respond to the inputs from the head direction cells and eye position cells in a manner somewhat analogous to a logical AND function, and this is achieved by the Sigma–Pi formulation implemented in Eq. (11) above. That is, each spatial view cell will be strongly stimulated only for appropriate combinations of head direction and eye position.

Once the activations of the spatial view cells have been computed, the firing rates of the cells are then given by the sigmoid transfer function (6).

4.3. Simulation results with Model 3

In the case of Model 3, the agent must simply assume every possible combination of head direction and eye position during the learning phase. This is because the learning process for Model 3 simply involves associating every possible combination of head direction and eye position of the agent with the corresponding fixation location. However, it should be noted that the learning described here for Model 3 is only the learning required to self-organise the recurrent synaptic connections w_{ij}^{RC} within the network of spatial view cells, and the idiosyncratic connections w_{ikl}^{ID} to the spatial view cells from the head direction cells and eye position cells. In fact, Model 3 assumes the existence of a separate layer of head direction cells which performs path integration over head rotation inputs, and this could operate as described by Stringer et al. (2002b).

The results of the numerical simulations of Model 3 are shown in Fig. 7, and are similar to those shown earlier for Models 1 and 2. Thus, the simulations confirmed that Model 3 was able to correctly update the firing in the network of spatial view cells in the absence of visual input, as the agent rotated its head and moved its eyes, even when Model 3 had been trained on different combinations of head direction and eye position signals, independently of velocity.

⁵ The scaling factor (ϕ_0/C^{SV}) controls the overall strength of the recurrent inputs to the layer of spatial view cells, where ϕ_0 is a constant and C^{SV} is the number of synaptic connections received by each spatial view cell from other spatial view cells.

⁶ The scaling factor $\phi_1/(C^{HD \times EP})$ controls the overall strength of the inputs from the head direction cells and eye position cells, where ϕ_1 is a constant and $C^{HD \times EP}$ is the number of connections received by each spatial view cell from combinations of head direction cells and eye position cells.

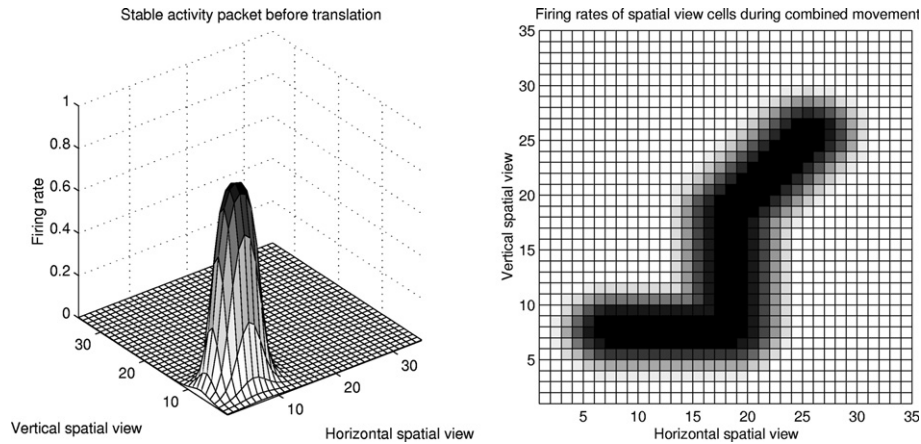


Fig. 7. Firing rate profiles from Model 3 (the feedforward network model for combining head direction and eye position signals to provide spatial view representations) from the spatial view cells during the testing phase with the agent in the dark. The grid shows the horizontal and vertical coordinates at which the different spatial view cells have their optimal firing. (Left) Stable firing rate profile within the network of spatial view cells before the agent starts to move. (Right) Maximum firing rates that occurred during movement of the agent in the testing phase. (The maximum firing rate is calculated over all timesteps for each spatial view cell.) First, the agent rotated its head in the clockwise direction. Next the agent moved its eyes vertically upwards. Finally, the agent simultaneously rotated its head in the clockwise direction and moved its eyes vertically upwards.

5. Discussion

In this paper, we have introduced the first neural network models of spatial view cells in the primate hippocampus. A key difficulty involved in building models of spatial view cells is the general problem of how a continuous attractor network might integrate a number of different kinds of velocity signal (e.g., head rotation and eye movement) simultaneously. This difficulty has not been addressed by earlier self-organizing path integration models of, for example, head direction cells (Redish et al., 1996; Skaggs et al., 1995; Stringer et al., 2002b; Zhang, 1996), and place cells (Redish & Tourtzky, 1998; Samsonovich & McNaughton, 1997; Stringer et al., 2002a). The three spatial view cell models presented in this paper introduce three different ways in which a continuous attractor network could integrate a number of different kinds of velocity signal simultaneously.

In Model 1 (Fig. 1) the continuing firing in the dark, a short term memory operation, of the continuous attractor network is maintained by the recurrent collateral synapses w_{ij}^{RC} . The idiotic inputs from the head rotation cells (with rates r^{ROT}) that co-occur with the changing spatial view cell firing r^{SV} (diagnosed by a trace term in the recurrent collateral synapses of the spatial view cells) in the light become associated together using Sigma-Pi synapses w_{ijk}^{ROT} . In a similar way, the idiotic inputs from the eye velocity cells (with rates r^{EV}) that co-occur with the changing spatial view cell firing r^{SV} (diagnosed by a trace term in the recurrent collateral synapses of the spatial view cells) in the light become associated together using Sigma-Pi synapses w_{ijk}^{EV} . We showed that not only does this idiotic update occur in the dark, but also that the two sets of idiotic

weights could be learned separately, and would later add correctly when both head rotation and eye velocity were present. However, in the simulations with Model 1, the network was trained with the agent either rotating its head, or moving its eyes, but not performing both kinds of movement together. However, a serious problem occurs when the agent is free to perform more than one movement at a time during learning, in that, as described above, combined head and eye movements become associated with each type of idiotic synapse.

Model 2 is a continuous attractor network model which is able to solve the problem of combined movements during learning by associating specific combinations of idiotic movement velocity signals with the shift of the activity packet in the continuous attractor network of spatial view cells. In particular, in Model 2 the synaptic connections to the spatial view cells from the different classes of idiotic cells are combined multiplicatively within higher order Sigma-Pi synapses. This allows the network to learn how to shift the activity packet within the layer of spatial view cells in the correct direction for precise combinations of idiotic signals. However, although the approach of combining the different kinds of idiotic signals within higher order Sigma-Pi synapses allows Model 2 to cope with the agent performing combined movements during learning, Model 2 requires significantly more training than Model 1. This is because Model 2 is unable to generalise over different speeds of movement.

Model 3 is able to solve the problem of combined movements during learning, without the need for training with all possible combinations of speeds, by utilizing earlier networks which perform the path integration on each of the idiotic signals, to produce position signals which can then be by the look-up table (which is essentially

Model 3) look up the correct position on spatial view space (Fig. 6). During training, the Sigma–Pi synapses learn the postsynaptic spatial view firing that occurs for all possible combinations of the eye position and head direction inputs. Model 3 is thus a feedforward model, has no need for recurrent collateral synapses in the spatial view cell network, and does not itself perform path integration. The Sigma–Pi synapses are needed because each particular combination of head and eye position firing must activate the correct spatial view cell, as each spatial view cell is tuned to its own gaze angle (the direction of the spatial view). A disadvantage of Model 3, which highlights an advantage of Models 1 and 2, is that Models 1 and 2 have a continuous attractor spatial view network which can be initially activated by visual cues in a particular environment, so that idiothetic cues can then move the activity packet in the correct direction in that spatial view environment. Model 3 as simulated did not have its recurrent connections in the spatial view network enabled, so there is nothing to keep active in the dark the spatial view environment triggered initially by the visual cues. Thus, in Model 3 particular combinations of eye and head position signals might look up a representation in a different environment, if the spatial view network was trained in several different environments or charts (Battaglia & Treves, 1998; Samsonovich & McNaughton, 1997; Stringer, Rolls, & Trappenberg, 2004; Tsodyks, 1999). Model 3 could be simulated again to utilize the recurrent connections in the spatial view network to hold the environment active in the dark, and then utilize the position signals to look up the next position in that environment without jumping to another chart.

In terms of predictions, Model 1 could be rejected if it can be shown in a particular environment that spatial view path integration learning can take place with simultaneous combinations of eye velocity and head rotation velocity inputs. Model 2 was developed to solve this problem. Model 2 requires different velocities of the eye and head rotation velocity inputs to be represented by different neurons, to operate correctly with different velocities. If the firing rate of a single population of neurons was proportional to the speed of for example head rotation velocity, then Model 2 would not operate correctly for different velocities of input. Neurons tuned to particular velocities are needed. In contrast, Model 1 requires the firing rates of its idiothetic input neurons to be proportional to the velocity.

The actual biophysical mechanisms that are needed to implement the self-organization by learning of the idiothetic connections in all three Models must, necessarily given the computational structure of the problem to be solved, include three or four terms. In Model 1 they are the postsynaptic spatial view cell firing, the presynaptic spatial view cell (traced) firing, and the presynaptic head rotation cell (or eye velocity cell) firing. In

Model 2 they are the postsynaptic spatial view cell firing, the presynaptic spatial view cell (traced) firing, the presynaptic head rotation cell firing, and the presynaptic eye velocity cell firing. In Model 3 they are the postsynaptic spatial view cell firing, the presynaptic head direction cell firing, and the presynaptic eye position cell firing. In all three Models the multiplicative interactions required, namely Sigma–Pi operation or synaptic strength modulation, could be performed by presynaptic contacts. However, multiplicative interactions of the type needed in these Models might be achieved in a number of other biophysically plausible ways described by Goodridge and Touretzky (2000); Jonas and Kaczmarek (1999); Koch (1999, Section 21.1.1); and Rolls and Deco (2002).

We now discuss the learning dynamics of the models. The models presented in this paper use unbounded associative learning rules, where the size of the synaptic weights continues to increase indefinitely during training. In addition, regular training through the spatial view space was provided.⁷ However, the situation becomes more complex when one considers irregular training, in which the gaze of the agent moves erratically and unevenly through the spatial view space. In this case, if the agent's gaze spends more time during training in a particular part of the spatial view space, then (because we specified a non-normalizing associative learning rule) the recurrent weights will be larger in this region, and the activity packet will be attracted there even when the agent's gaze is stationary. We did not address the normalization of synaptic weights to deal with irregular training in this paper because we have addressed it previously in the context of a self-organizing continuous attractor model of head direction cells (Stringer et al., 2002b). In that paper, it was shown that normalizing each synaptic weight vector during irregular training led to a smooth synaptic weight profile within the continuous attractor network which supported good performance, with the small resulting drift of the activity packet being eliminated by implementing the post-synaptic threshold mechanism specified in Eq. (7). Associative long-term heterosynaptic depression is a mechanism that could be involved in such weight normalization (Oja, 1982; Rolls & Treves, 1998).

Hahnloser (2003) has suggested using an error correction process, and a special purpose network architecture (with separate head direction subnetworks required for each direction of idiothetic signal) which can lead to a convergent learning scheme in a one-dimensional head direction cell system. One disadvantage of that ap-

⁷ For Model 1, training once through all possible tracks in spatial view space, first with eye velocity, and second with head rotation, is needed for perfect performance. For Model 2 the training is similar except that all combinations of the two velocity signals are needed during training for perfect performance.

proach is that it is less biologically plausible than the networks described in this paper, because it requires error correction learning (Rolls & Deco, 2002; Rolls & Stringer, 2001). Another disadvantage is that the architecture is dedicated to the particular problem of head direction cells. Moreover, the stability of the activity packets in the two rings of head direction cells, and the ability of the model to perform accurate path integration, are dependent on finely balanced synaptic interactions implemented between the two rings. The two rings are each responsible for driving the activity packets in opposite directions, and so counterbalance each other as they work together to control the movement of the activity packets. The synaptic connections within and between the two rings are self-organised during learning using an error correction learning rule. This error correction rule compares the current firing rate of each head direction cell with the visual signal to that cell (where the visual signal reflects the true head direction of the agent). The required asymmetry in all the synapses arises because a continuous dynamical system involving time delays is used. In contrast Models 1 and 2 described here contain a single CANN, in which each cell receives a number of different types of idiothetic (self-motion) input. This design has enabled the network to be extended easily to more than one type of idiothetic input, as described in this paper, and that would be very difficult for the approach taken by Hahnloser (2003), because every type and even every direction within a type of idiothetic input would appear to need two separate counterbalancing recurrent networks.

The parts of the brain involved in the representation of space may receive a variety of different kinds of idiothetic and motor efference copy signals, which must be integrated together. To date, CANN models of path integration have focussed on head direction cells and place cells, and have not needed to address how multiple self-motion signals could be combined. However, the need for path integration networks which are able to cope with the simultaneous integration of different types of self-motion signal comes to the fore when developing models of spatial view cells found in the primate hippocampus. This paper presents the first CANN models of this kind of cell, and opens up the general question of how CANN path integration networks might learn to integrate multiple simultaneously active self-motion 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

- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27, 77–87.
- Battaglia, F. P., & Treves, A. (1998). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Physical Review E*, 58, 7738–7753.
- Berthoz, A. (2000). *The brain's sense of movement*. Cambridge, MA: Harvard University Press.
- Berthoz, A., Grantyn, A., & Droulez, J. (1987). Some collicular neurons code saccadic eye velocity. *Neuroscience Letters*, 72, 289–294.
- Debanne, D., Gähwiler, B., & Thompson, S. M. (1998). Long term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *Journal of Physiology*, 507, 237–247.
- Georges-François, P., Rolls, E. T., & Robertson, R. G. (1999). Spatial view cells in the primate hippocampus: Allocentric view not head direction or eye position or place. *Cerebral Cortex*, 9, 197–212.
- Goodridge, J. P., & Touretzky, D. S. (2000). Modeling attractor deformation in the rodent head-direction system. *Journal of Neurophysiology*, 83, 3402–3410.
- Hahnloser, R. H. R. (2003). Emergence of neural integration in the head-direction system by visual supervision. *Neuroscience*, 120, 877–891.
- Ishizuka, N., Weber, J., & Amaral, D. G. (1990). Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *Journal of Comparative Neurology*, 295, 580–623.
- Jonas, E. A., & Kaczmarek, L. K. (1999). The inside story: Subcellular mechanisms of neuromodulation. In P. S. Katz (Ed.), *Beyond neurotransmission* (pp. 83–120). New York: Oxford University Press, Chapter 3.
- Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1998). Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells. *Journal of Neurophysiology*, 80, 425–446.
- Koch, C. (1999). *Biophysics of computation*. Oxford: Oxford University Press.
- Muller, R. U., Kubie, J. L., & Saypoloff, R. (1991). The hippocampus as a cognitive graph (abridged version). *Hippocampus*, 1, 243–246.
- Oja, E. (1982). A simplified neuron model as a principal component analyser. *Journal of Mathematical Biology*, 15, 267–273.
- O'Mara, S. M., Rolls, E. T., Berthoz, A., & Kesner, R. P. (1994). Neurons responding to whole-body motion in the primate hippocampus. *Journal of Neuroscience*, 14, 6511–6523.
- Redish, A. D. (1999). *Beyond the cognitive map: From place cells to episodic memory*. Cambridge, Massachusetts: MIT Press.
- Redish, A. D., Elga, A. N., & Touretzky, D. S. (1996). A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems*, 7, 671–685.
- Redish, A. D., & Touretzky, D. S. (1998). The role of the hippocampus in solving the Morris water maze. *Neural Computation*, 10, 73–111.
- Robertson, R. G., Rolls, E. T., Georges-François, P., & Panzeri, S. (1999). Head direction cells in the primate pre-subiculum. *Hippocampus*, 9, 206–219.
- Robertson, R. G., Rolls, E. T., & Georges-François, P. (1998). Spatial view cells in the primate hippocampus: Effects of removal of view details. *Journal of Neurophysiology*, 79, 1145–1156.
- Rolls, E. T. (1999). Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus*, 9, 467–480.
- Rolls, E. T., & Deco, G. (2002). *Computational neuroscience of vision*. Oxford: Oxford University Press.
- Rolls, E. T., Robertson, R. G., & Georges-François, P. (1997). Spatial view cells in the primate hippocampus. *European Journal of Neuroscience*, 9, 1789–1794.

- Rolls, E. T., & Stringer, S. M. (2001). Invariant object recognition in the visual system with error correction and temporal difference learning. *Network: Computation in Neural Systems*, *12*, 111–129.
- Rolls, E. T., Stringer, S. M., & Trappenberg, T. P. (2002). A unified model of spatial and episodic memory. *Proceedings of The Royal Society B*, *269*, 1087–1093.
- Rolls, E. T., & Treves, A. (1998). *Neural networks and brain function*. Oxford: Oxford University Press.
- Rolls, E. T., Treves, A., Robertson, R. G., Georges-François, P., & Panzeri, S. (1998). Information about spatial view in an ensemble of primate hippocampal cells. *Journal of Neurophysiology*, *79*, 1797–1813.
- Samsonovich, A., & McNaughton, B. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, *17*, 5900–5920.
- Sharp, P. E., Blair, H. T., & Cho, J. (2001). The anatomical and computational basis of the rat head-direction cell signal. *Trends in Neurosciences*, *24*, 289–294.
- Skaggs, W. E., Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1995). A model of the neural basis of the rat's sense of direction. In G. Tesauro, D. S. Touretzky, & T. K. Leen (Eds.). *Advances in neural information processing systems* (Vol. 7, pp. 173–180). Cambridge, Massachusetts: MIT Press.
- Stringer, S. M., Rolls, E. T., & Trappenberg, T. P. (2004). Self-organising continuous attractor networks with multiple activity packets, and the representation of space. *Neural Networks*, *17*, 5–27.
- Stringer, S. M., Rolls, E. T., Trappenberg, T. P., & de Araujo, I. E. T. (2002a). Self-organizing continuous attractor networks and path integration: Two-dimensional models of place cells. *Network: Computation in Neural Systems*, *13*, 429–446.
- Stringer, S. M., Trappenberg, T. P., Rolls, E. T., & de Araujo, I. E. T. (2002b). Self-organizing continuous attractor networks and path integration: One-dimensional models of head direction cells. *Network: Computation in Neural Systems*, *13*, 217–242.
- Taube, J. S., Muller, R. U., & Ranck, J. B. Jr., (1990). Head-direction cells recorded from the post-subiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, *10*, 420–435.
- Taylor, J. G. (1999). Neural 'bubble' dynamics in two dimensions: Foundations. *Biological Cybernetics*, *80*, 393–409.
- Tsodyks, M. (1999). Attractor neural network models of spatial maps in hippocampus. *Hippocampus*, *9*, 481–489.
- Wang, X. J. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *Journal of Neuroscience*, *19*, 9587–9603.
- Xing, J., & Andersen, R. A. (2000). Models of the posterior parietal cortex which perform multimodal integration and represent space in several coordinate frames. *Journal of Cognitive Neuroscience*, *12*, 601–614.
- Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, *16*, 2112–2126.