

2005 Special issue

Spatial view cells in the hippocampus, and their idiothetic update based on place and head direction

Edmund T. Rolls *, Simon M. Stringer

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

Abstract

Single neuron recording studies have demonstrated the existence of 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 neural networks. Stringer, Rolls and Trappenberg (2005) have recently shown how idiothetic (self-motion) inputs could update the activity packet of neuronal firing within a two-dimensional continuous attractor neural network of spatial view cells. However, this earlier study examined only the simplified situation in which the agent could rotate on the spot or move its eyes. In this paper we show how spatial view cells could be driven by head direction and place cells, themselves idiothetically updated. The head direction and place neurons are remapped by a competitive network with expansion recoding which self-organises so that different neurons represent different combinations of head direction and the place where the agent is located. The combination cells are then mapped by pattern association involving long-term synaptic potentiation but also long-term homosynaptic depression to spatial view cells, which during training are driven by the spatial view. After training, the spatial view cells are updated in the dark by the idiothetically driven head direction and place cells.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Spatial view cells; Place cells; Head direction cells; Continuous attractor neural networks; Self-organization; Competitive networks

1. Introduction

In this paper we first review the properties and models of primate hippocampal spatial view cells and of presubicular head direction cells. Then we consider how the spatial representations in these models could be updated by idiothetic (self-motion) signals. Then we present a new model of how spatial view cells could be updated by self-motion even when the animal is located in different parts of an environment, and how the connections in the model could self-organize as a result of learning. Part of the interest of the new model presented here is that the idiothetic update can be performed without multiplicative Sigma-Pi synapses, which were used in earlier models (Stringer, Trappenberg, Rolls, & De Araujo, 2002; Stringer, Rolls, Trappenberg, & De Araujo, 2002; Stringer et al., 2005; Stringer, Rolls, Trappenberg, & De Araujo, 2003; Stringer, Rolls, & Trappenberg, 2004), but instead with Hebbian synapses. The new model thus also shows

how these earlier models could be extended to utilize just Hebbian synapses.

1.1. Spatial view cells and head direction cells in primates: neurophysiology and models

1.1.1. Head direction and place cells

Single neuron recording studies have demonstrated the existence of neurons which appear to encode information about the position or orientation of an animal with respect to its environment, and which are able to maintain their response properties even in the absence of visual input. Examples of such cells include head direction cells in rats (Muller, Ranck, & Taube, 1996; Ranck, 1985; Taube, Muller, & Ranck, 1990; Taube, Goodridge, Golob, Dudchenko, & Stackman, 1996) and primates (Robertson, Rolls, Georges-François, & Panzeri, 1999) which respond maximally when the animal's head is facing in a particular preferred direction, and hippocampal place cells in rats (Markus, Qin, Leonard, Skaggs, McNaughton and Barnes, 1995; Muller, Kubie, Bostock, Taube & Quirk, 1991; McNaughton, Barnes, & O'Keefe, 1983; O'Keefe & Dostrovsky, 1971; O'Keefe, 1984) that fire maximally when the animal is in a particular location. The ways in which visual cues might stimulate the rodent place cells have been

* Corresponding author. Tel.: +44 1865 271348; fax: +44 1865 310447.

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

URL: <http://www.cns.ox.ac.uk>.

demonstrated in neurophysiological studies (Markus et al., 1995; McNaughton et al., 1983; Muller et al., 1991; O’Keefe, 1984), and modelled in theoretical investigations (De Araujo, Rolls, & Stringer, 2001; O’Keefe, & Burgess, 1996).

1.1.2. Spatial view cells

In primates, hippocampal spatial view cells have been discovered that respond when the monkey is looking towards a particular location in space (Rolls, Robertson, & Georges-François, 1997; Rolls, Treves, Robertson, Georges-François, & Panzeri, 1998; Robertson, Rolls & Georges-François, 1998; Rolls, 1999). Hippocampal spatial view cells code for the location at which a primate is looking, and hence code for particular locations in allocentric (world-based rather than egocentric) space (Georges-François, Rolls, & Robertson, 1999; Rolls et al., 1998). An important property of primate spatial view cells is their ability to maintain their spatial properties for periods of up to several minutes in the dark. For example, some spatial view cells respond to remembered spatial views in that they respond when the view details are obscured, and use idiothetic cues including eye position and head direction to trigger this memory recall operation (Robertson et al., 1998). Another idiothetic input that drives some primate hippocampal neurons is linear and axial whole body motion (O’Mara, Rolls, Berthoz, & Kesner, 1994), and in addition, the primate presubiculum has been shown to contain head direction cells (Robertson et al., 1999).

Part of the interest of spatial view cells is that they could provide the spatial representation required to enable primates to perform object-place memory, for example remembering where they saw a person or object, which is an example of an episodic memory, and indeed similar neurons in the hippocampus respond in object-place memory tasks (Rolls, Miyashita, Cahusac, Kesner, Niki and Feigenbaum, 1989; Rolls, Xiang, & Franco, 2005). Associating together such a spatial representation with a representation of a person or object could be implemented by an autoassociation network implemented by the recurrent collateral connections of the CA3 hippocampal pyramidal cells (Rolls, 1989; Rolls, 1996; Rolls & Treves, 1998; Treves & Rolls, 1992; Treves & Rolls, 1994). Some other primate hippocampal neurons respond in the object-place memory task to a combination of spatial information and information about the object seen (Rolls et al., 1989). Further evidence for this convergence of spatial and object information in the hippocampus is that in another memory task for which the hippocampus is needed, learning where to make spatial responses conditional on which picture is shown, some primate hippocampal neurons respond to a combination of which picture is shown, and where the response must be made (Cahusac, Rolls, Miyashita, & Niki, 1993; Miyashita, Rolls, Cahusac, Niki, & Feigenbaum, 1989), and this task and an arbitrary stimulus-motor response association task are impaired by damage to the hippocampal system (Brasted, Bussey, Murray, & Wise, 2003; Wise & Murray, 1999). A further interesting type of convergence found in the primate hippocampus is of spatial view and reward value, for spatial view-reward neurons have now been found that allow

the particular rewards available at different locations in a scene to be recalled (Rolls, 2005; Rolls & Xiang, (2005a)).

1.1.3. A model of the storage and retrieval of memories by the hippocampal system

These neurophysiological analyses are complemented by neuronal network models of how the hippocampus could operate to store and retrieve large numbers of memories (Rolls, 1987; Rolls, 1989; Rolls, 1996; Treves & Rolls, 1992; Treves & Rolls, 1994; Rolls & Treves, 1998). One key hypothesis (adopted also by McClelland, McNaughton and O’Reilly (1995)) is that the hippocampal CA3 recurrent collateral connections which spread throughout the CA3 region provide a *single autoassociation network* that enables the firing of *any* set of CA3 neurons representing one part of a memory to be associated together with the firing of any other set of CA3 neurons representing another part of the same memory (cf. Marr (1971)). The number of patterns p each representing a different memory that could be stored in the CA3 system operating as an autoassociation network would be as shown in Eq. (1) (see Rolls and Treves (1998); Rolls and Deco (2002))

$$p \approx \frac{C^{\text{RC}}}{a \ln\left(\frac{1}{a}\right)} k \quad (1)$$

where C^{RC} is the number of synapses on the dendrites of each neuron devoted to the recurrent collaterals from other CA3 neurons in the network, a is the sparseness of the representation, and k is a factor that depends weakly on the detailed structure of the rate distribution, on the connectivity pattern, etc., but is roughly in the order of 0.2–0.3. Given that C^{RC} is approximately 12,000 in the rat, the resulting storage capacity would be greater than 12,000 memories, and perhaps up to 36,000 memories if the sparseness a of the representation was as low as 0.02 (Treves & Rolls, 1992, 1994). Another part of the hypothesis is that the very sparse (see Fig. 1) but powerful connectivity of the mossy fibre inputs to the CA3 cells from the dentate granule cells is important during learning (but not recall) to force a new, arbitrary, set of firing onto the CA3 cells which dominates the activity of the recurrent collaterals, so enabling a new memory represented by the firing of the CA3 cells to be stored (Rolls, 1989; Rolls, 1987; Treves & Rolls, 1992). The perforant path input to the CA3 cells, which is numerically much larger but at the apical end of the dendrites, would be used to initiate recall from an incomplete pattern (Rolls & Treves, 1998; Treves & Rolls, 1992). The prediction of the theory about the necessity of the mossy fibre inputs to the CA3 cells during learning but not recall has now been confirmed (Lassalle, Bataille, & Halley, 2000). A way to enhance the efficacy of the mossy fibre system relative to the CA3 recurrent collateral connections during learning may be to increase the level of acetyl choline by increasing the firing of the septal cholinergic cells (Hasselmo, Schnell, & Barkai, 1995).

Another key part of the quantitative theory is that not only can retrieval of a memory to an incomplete cue be performed

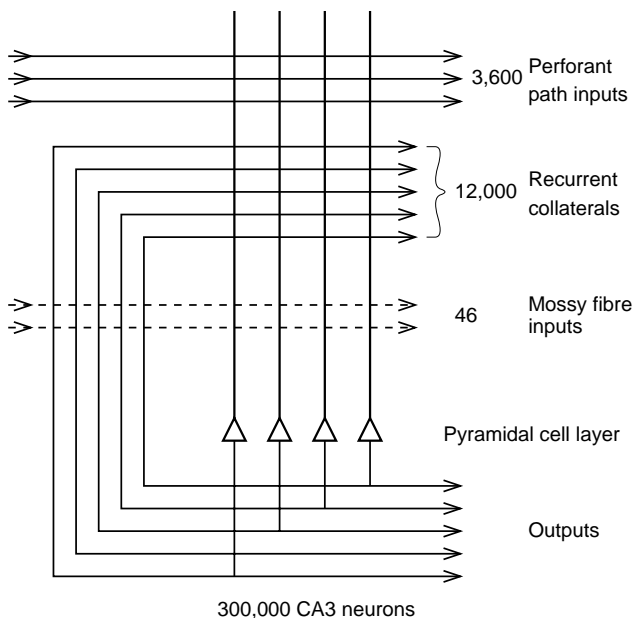


Fig. 1. The numbers of connections onto each CA3 cell from three different sources in the rat. (After Treves & Rolls, 1992; Rolls & Treves, 1998.)

by the operation of the associatively modified CA3 recurrent collateral connections, but also that recall of that information to the neocortex can be performed via CA1 and the hippocampocortical and cortico-cortical backprojections (Treves & Rolls, 1994; Rolls, 1996; Rolls, 2000; Rolls & Treves, 1998) shown in Fig. 2. In this case, the number of memory patterns p^{BP} that can be retrieved by the backprojection system is

$$p^{\text{BP}} \approx \frac{C^{\text{BP}}}{a^{\text{BP}} \ln\left(\frac{1}{a^{\text{BP}}}\right)} k^{\text{BP}} \quad (2)$$

where C^{BP} is the number of synapses on the dendrites of each neuron devoted to backprojections from the preceding stage (dashed lines in Fig. 2), a^{BP} is the sparseness of the representation in the backprojection pathways, and k^{BP} is a factor that depends weakly on the detailed structure of the rate distribution, on the connectivity pattern, etc., but is roughly in the order of 0.2–0.3. The insight into this quantitative analysis came from treating each layer of the backprojection hierarchy as being quantitatively equivalent to another iteration in a single recurrent attractor network (Treves & Rolls, 1994; Treves & Rolls, 1991). The need for this number of connections to implement recall, and more generally constraint satisfaction in connected networks, provides a fundamental and quantitative reason for why there are approximately as many backprojections as forward connections between the adjacent connected cortical areas in a cortical hierarchy.

The number of memories that could be recalled by the backprojections would be maximal if the backprojection representations were sparse, as shown above in Eq. (2). In a recent investigation of the sparseness of the representation provided by populations of neurons in the inferior temporal visual cortex, it has been found that the representations are not

very sparse, with $a=0.76$ (Franco, Rolls, Aggelopoulos, & Jerez, 2005). While less is known of the sparseness in the primate subiculum, entorhinal cortex, and perirhinal cortex (Hölscher, Rolls, & Xiang, 2003; Rolls, Franco, & Stringer, 2005), in rodents the representation does not appear to be very sparse in the subiculum (Sharp & Green, 1994), but can be as sparse in the entorhinal cortex as in the hippocampus (Fyhn, Molden, Witter, Moser, & Moser, 2004). Even if the representations are not very sparse, the number of memories that can be retrieved is still reasonable, as shown by the following. The calculation performed by Treves and Rolls (1991) and illustrated in their Fig. 5a (reproduced in Rolls and Treves (1998) Fig. A4.2) deals with a number of cases of dilution of the connectivity and the probability distribution of the firing rates in each pattern. Eq. (2) only summarizes this calculation rather approximately, and reference to the graphs is needed for the details, as the capacity is somewhat different for the fully connected and diluted connectivity cases. (In the brain, the connectivity will be diluted.) For the highly diluted limit the storage capacity does not decrease much below the value it has for $a=0.2$ as a increases. For realistic connectivity with $C/N \approx 0.1$ (where C is the number of connections per neuron, and N is the number of neurons), the number of patterns p^{BP} continues to decrease with increasing a , but remains reasonable.

Another aspect of the theory is that the operation of the CA3 system to implement recall, and of the backprojections to retrieve the information, would be sufficiently fast, given the fast recall in associative networks built of neurons with continuous dynamics (see Rolls and Deco (2002); Rolls and Treves (1998)).

1.1.4. Continuous attractor models of spatial representations

A key challenge is to understand how classes of cells such as head direction cells, place cells and spatial view cells, can maintain their response properties when the animal is in darkness, with no visual input available to guide and update the firing of the cells. A class of network that can maintain the firing of its neurons to represent any location along a continuous physical dimension such as head direction is a ‘Continuous Attractor’ neural network (CANN). A number of researchers have modelled head direction and place cells with ‘continuous attractor’ neural networks (Battaglia & Treves, 1998; Redish, Elga, & Touretzky, 1996; Redish & Touretzky, 1998; Samsonovich & McNaughton, 1997; Skaggs, Knierim, Kudrimoti, & McNaughton, 1995; Stringer, Trappenberg, Rolls, & De Araujo, 2002; Stringer, Rolls, Trappenberg, & De Araujo, 2002; Stringer et al., 2005; Stringer et al., 2004; Zhang, 1996) which are able to maintain a localised packet of neuronal activity representing the current state of the animal. In these networks the excitatory recurrent collateral connections between the neurons reflect the distance between the neurons in the state space (e.g. head direction space) of the animal. Then, global inhibition is used to keep the number of neurons in a bubble of activity relatively constant, and to help to ensure that there is only one activity packet. The properties of these

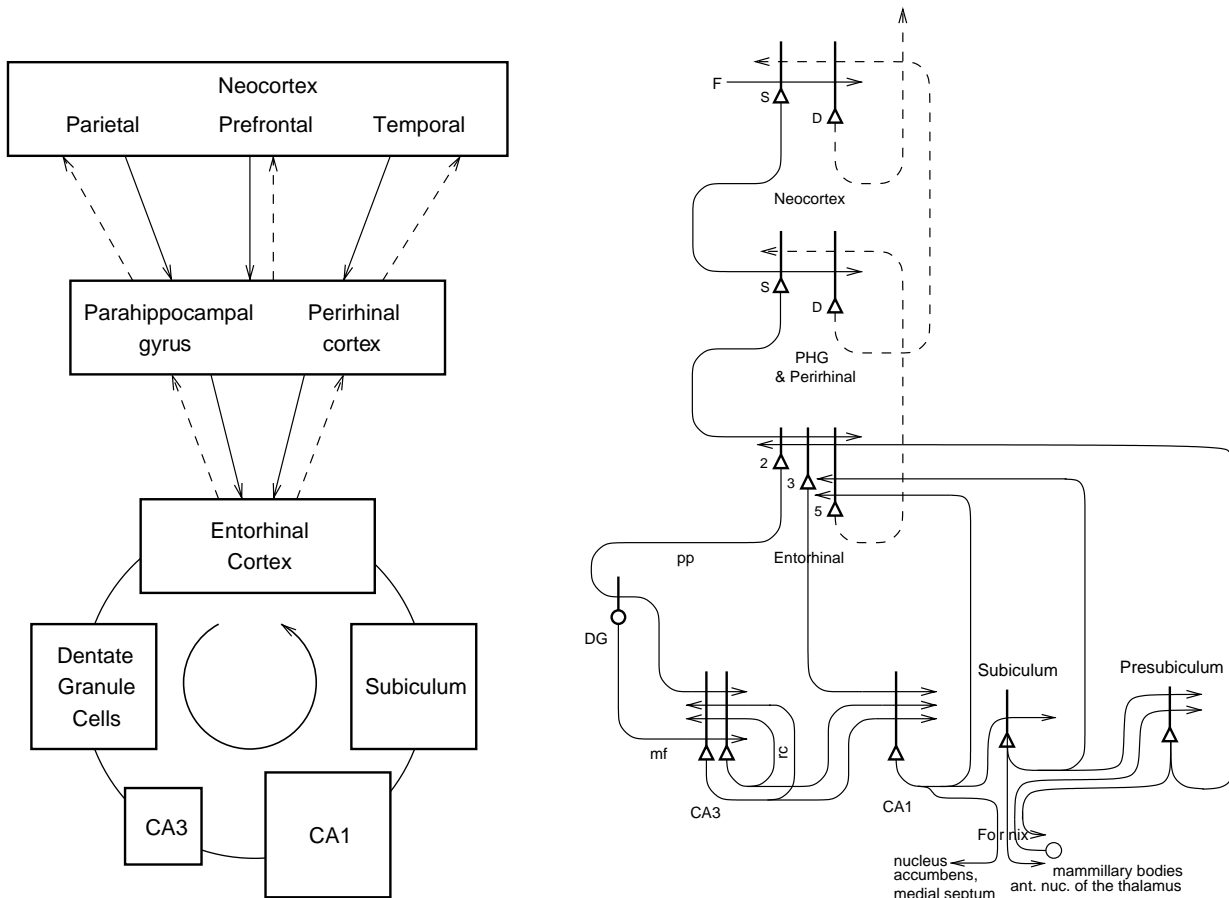


Fig. 2. Forward connections (solid lines) from areas of cerebral association neocortex via the parahippocampal gyrus and perirhinal cortex, and entorhinal cortex, to the hippocampus; and backprojections (dashed lines) via the hippocampal CA1 pyramidal cells, subiculum, and parahippocampal gyrus to the neocortex. There is great convergence in the forward connections down to the single network implemented in the CA3 pyramidal cells; and great divergence again in the backprojections. Left: block diagram. Right: more detailed representation of some of the principal excitatory neurons in the pathways. Abbreviations: D, Deep pyramidal cells; DG, dentate granule cells; F, forward inputs to areas of the association cortex from preceding cortical areas in the hierarchy. mf, mossy fibres; PHG, parahippocampal gyrus and perirhinal cortex; pp, perforant path; rc, recurrent collaterals of the CA3 hippocampal pyramidal cells; S, superficial pyramidal cells; 2, pyramidal cells in layer 2 of the entorhinal cortex; 3, pyramidal cells in layer 3 of the entorhinal cortex; 5, pyramidal cells in the deep layers of the entorhinal cortex. The thick lines above the cell bodies represent the dendrites.

continuous attractor networks have been extensively studied, for example by Amari (1977); Taylor (1999).

1.1.5. Networks that combine spatial and discrete information

It has now been shown that attractor networks can store both continuous patterns and discrete patterns, and can thus be used to store for example the location in (continuous, physical) space where an object (a discrete item) is present (Rolls, Stringer, & Trappenberg, 2002). In this network, when events are stored that have both discrete (object) and continuous (spatial) aspects, then the whole place can be retrieved later by the object, and the object can be retrieved by using the place as a retrieval cue. Such networks are likely to be present in parts of the brain such as the hippocampus which receive and combine inputs both from systems that contain representations of continuous (physical) space, and from brain systems that contain representations of discrete objects, such as the inferior temporal visual cortex. The combined continuous and discrete attractor network described by Rolls et al. (2002) shows that in brain regions where the spatial and object processing streams

are brought together, then a single network can represent and learn associations between both types of input. Indeed, in brain regions such as the hippocampal system, it is essential that the spatial and object processing streams are brought together in a single network, for it is only when both types of information are in the same network that spatial information can be retrieved from object information, and vice versa, which is a fundamental property of episodic memory.

1.1.6. Self-organizing models of the idiothetic update of head direction and place cells

In most continuous attractor models of head direction and place cells the synaptic connections have been fixed or hard-wired; that is, the synaptic weights have been set by formula to an appropriate value. This leads to a lack of biological plausibility, and to a lack of flexibility in terms of being able to build representations of complex or warped state spaces of the agent. These issues were addressed in Stringer, Trappenberg, Rolls, and De Araujo (2002); Stringer, Rolls, Trappenberg, and De Araujo (2002), where *self-organizing* continuous attractor

network models were developed for head direction and place cells respectively. In these papers it was shown how the synaptic connectivity of the continuous attractor networks could self-organize during an initial learning phase in the light, such that, during the subsequent testing phase in the dark, the networks were able to support a stable packet of neuronal activity representing the current state (e.g. head direction) of the agent which could be updated by idiothetic inputs. A fundamental feature of these self-organizing continuous attractor network models was their use of ‘trace’ learning rules which incorporate a form of temporal average of recent cell activity. Such rules are able to build associations between different patterns of neural activities that tend to occur in temporal proximity (Földiák, 1991; Rolls & Stringer, 2001; Wallis and Rolls, 1997).

1.1.7. Previous models of primate hippocampal spatial view cells

The ideas presented in Stringer, Trappenberg, Rolls and De Araujo (2002); Stringer, Rolls, Trappenberg and De Araujo (2002) were further developed to seek to explain the firing properties of spatial view cells in Stringer et al. (2005). In Stringer et al. (2005) two models of spatial view cells were proposed that were based on self-organizing continuous attractor networks. Both of the spatial view models presented by Stringer et al. (2005) self-organized through modification of their synaptic connections such that the spatial view being represented by the continuous attractor networks could be updated by idiothetic (self-motion) signals. In the first of these models, the recurrent network model, the firing pattern in the spatial view cell layer is updated in the dark by idiothetic velocity signals from rotation cells and eye velocity cells. In the second of these models, the feedforward model, the firing pattern in the spatial view cell layer is updated in the dark by positional information from head direction cells and eye position cells. However, the two models presented in Stringer et al. (2005) examined only the simplified situation in which the agent could rotate on the spot or move its eyes.

2. A new model of the idiothetic update of spatial view cells

In this paper we extend the work presented in Stringer et al. (2005) to address the more general case in which the agent is able to move through the environment visiting different locations. In this paper we present a model of spatial view cells that enables the spatial view cells to respond correctly to the spatial view even when the agent is located in different places in its environment which result in different head directions being required to see the spatial view. A feature of this new model is that the idiothetic update can be performed without multiplicative Sigma-Pi synapses, which were used in our earlier models (Stringer, Trappenberg, Rolls, & De Araujo, 2002; Stringer, Rolls, Trappenberg, & De Araujo, 2002; Stringer et al., 2005; Stringer et al., 2003; Stringer et al., 2004), but instead with Hebbian synapses. The new model thus also shows how these earlier models could be extended to utilize just Hebbian synapses.

2.1. The architecture of the model

The neural network architecture of the new model is shown in Fig. 3. There is a spatial view cell, which is driven by visual input from a particular spatial view in the environment during training in the light. (The model is described for simplicity for a single spatial view cell, though in practice there would be a whole population, each tuned to a different spatial view, as in the primate hippocampus (Rolls, 1999).) The spatial view cell also receives inputs from a layer of combination cells, which represent particular combinations of head direction and place. The combination cells receive inputs from a layer of head direction cells and a layer of place cells. We note that although we ourselves have not found place cells in the primate hippocampus (Rolls & Xiang, (2005b)) we acknowledge that they could be present (Ekstrom, Kahana, Caplan, Fields, Isham and Newman, 2003; Fried, MacDonald, & Wilson, 1997; Ono, Nakamura, Nishijo, & Eifuku, 1993), though very careful testing is required to prove that a cell responds to the place where the primate (including human) is located, and not to what is being looked at (Rolls & Xiang, (2005b)). Thus the inclusion of place cells in the current model is not unreasonable. In addition, we believe that the current model might apply to rodents, in the sense that rat place cells have very directional properties when the rat is running down an alley with a small view of the world straight ahead. Under these conditions, the neurons may respond when the rat is running in one direction, but not in the other (Gothard, Skaggs, & McNaughton, 1996), and in this sense might be thought of as spatial view cells. To test this suggestion, it would be interesting to measure whether rat place cells respond when the rat is running in an alley pointing to the same spatial view, but situated in different places in the environment. In addition, we note that one of the types of neuron that forms a key component of the present model, neurons that respond to a combination of place and head direction, have been found in the rat (Cho & Sharp, 2001; Sharp & Green, 1994; Sharp, 1996). Independently of these points, an additional part of the value of this model is that it shows how the Hebbian synapses for the idiothetic update of position in a state space can self-organize by learning, and thus has application to models that use other variables in the mapping system.

In the model, the combination cells operate as a self-organizing competitive network in which different neurons come as a result of the learning to respond to particular combinations of head direction and place as the agent moves to different places in the environment and adopts different head directions. Competitive networks are described by Hertz, Krogh, and Palmer (1991); Rolls and Treves (1998), and Rolls and Deco (2002), and involve competition between the neurons by for example mutual inhibition, and for neurons with relatively high firing after the competition, an associative learning rule to make the weight vectors move towards the pattern that is activating a neuron. It was found that to encourage a wide variety of different head direction/place combination cells to be set up, diluted connectivity from the place and head direction cells to the combination cells was

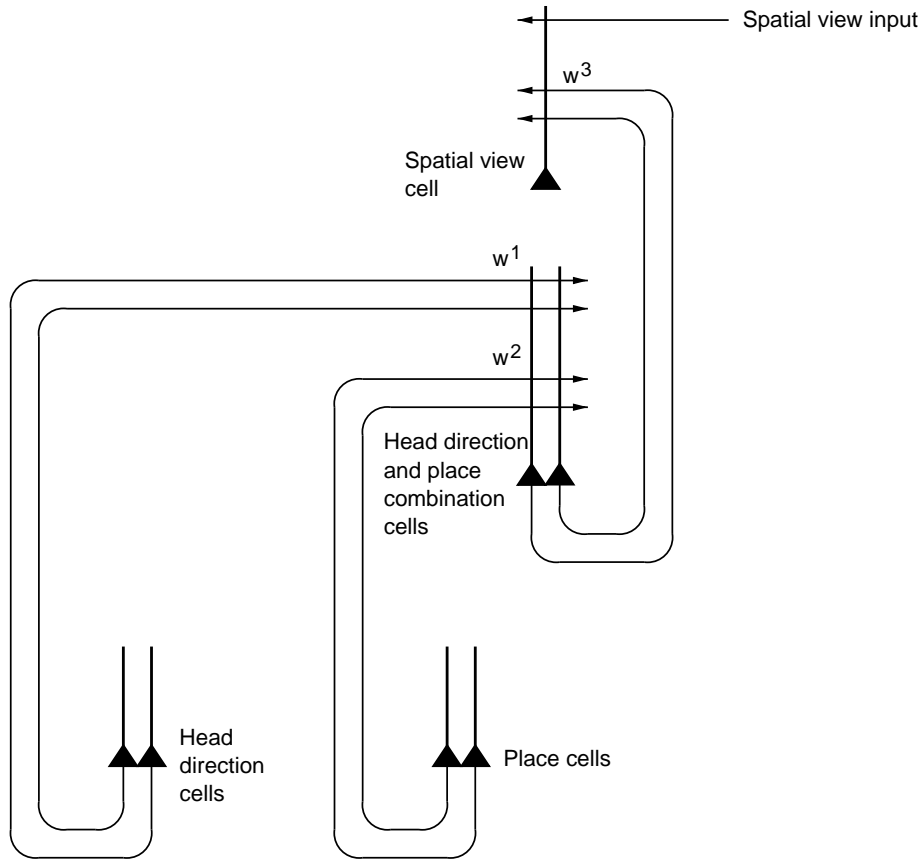


Fig. 3. Neural network architecture of the model. There is a spatial view cell, which is driven by visual input from a particular spatial view in the environment during training in the light. The spatial view cell also receives inputs from a layer of combination cells, which represent particular combinations of head direction and place. The combination cells receive inputs from a layer of head direction cells and a layer of place cells. There are three sets of modifiable synaptic weights, which self-organise during training: w^1 , w^2 and w^3 .

helpful. Once the combination cells had self-organised by competitive learning, the spatial view cells are forced into activity by a particular spatial view in the light, and by pattern association learn which input patterns from the combination cells are associated with the current spatial view. It was found that for this pattern association learning system to operate well, not only were increases in synaptic weights needed, but so were decreases, in particular corresponding to homosynaptic long-term depression (LTD). Homosynaptic LTD corresponds to decreasing the synaptic weight if there is no post-synaptic activity (because the neuron is not being driven by the current spatial view) but there is presynaptic activity. The homosynaptic LTD allows the spatial view cell(s) to learn not to respond to the activity of combination cells that are not associated with the current spatial view. Homosynaptic LTD has been found in the hippocampus (Stanton & Sejnowski, 1989).

The details of the model are as follows.

2.2. Cell types and properties

The head direction cells have firing which reflect the current head direction of the agent. In the simulations, each head direction cell i has a unique head direction θ_i for which the cell fires maximally. The firing rate r_i^{HD} of each head direction cell i

may be set to be the following Gaussian function of the displacement of the head from the optimal firing direction of the cell (Muller et al., 1996; Taube et al., 1996; Robertson et al., 1999)

$$r_i^{\text{HD}} = \exp\left(-\frac{(s_i^{\text{HD}})^2}{2(\sigma^{\text{HD}})^2}\right), \quad (3)$$

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

$$s^{\text{HD}} = \text{MIN}(|\theta_i - \theta|, 360 - |\theta_i - \theta|). \quad (4)$$

It is known that the firing rates of head direction cells in both rats (Muller et al., 1996; Taube et al., 1996) and macaques (Robertson et al., 1999) are approximately Gaussian.

The place cells have firing which reflect the current location of the agent within the square environment. Place cells 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

$$r_i^{\text{P}} = \exp\left(-\frac{(s_i^{\text{P}})^2}{2(\sigma^{\text{P}})^2}\right), \quad (5)$$

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

$$s_i^P = \sqrt{(x_i - x)^2 + (y_i - y)^2}. \quad (6)$$

The combination cells receive inputs from the head direction cells and place cells. The activation of each combination (COMB) cell i is given by

$$h_i^{\text{COMB}} = \sum_j w_{ij}^1 r_j^{\text{HD}} + \sum_j w_{ij}^2 r_j^{\text{P}} \quad (7)$$

where w_{ij}^1 are the weights from head direction cell j to combination cell i , and w_{ij}^2 are the weights from place cell j to combination cell i . The firing rates of the combination cells with the highest activations are set to 1, while the firing rates of the remaining combination cells are set to 0. In the simulations, the top 1 per cent of cells were kept active, and the sparseness a was therefore 0.01. In the simulations, diluted connectivity of these connections was used, to facilitate the setting up of neurons that responded to different combinations of head direction and place. The numbers of cells, and the numbers of synaptic connections between them, are given in Table 1.

During training in the light, the spatial view cell is driven by visual input from a particular spatial view in the environment. In the simulations, the spatial view cell is driven by visual input from a location on the North wall $(x=0.5, y=1.0)$. This gives rise to a spatial view field on the North wall about this position. The firing rate r^{SV} of the spatial view cell is set to be a Gaussian function of the displacement of the current head direction (i.e. gaze angle) of the agent from the head direction at which the agent would be looking towards the centre of the spatial view field. Of course, the head direction at which the agent would be looking towards the centre of the spatial view field depends on the current location of the agent. For example, it can be seen that when the agent is in location 1 $(x=0.25, y=0.75)$, the spatial view cell fires maximally when the agent has a head

direction of 45 degrees. However, when the agent is in location 2 $(x=0.75, y=0.75)$, the spatial view cell fires maximally when the agent has a head direction of 315°. Therefore, during training in the light, the firing rate of the spatial view cell is given by

$$r^{\text{SV}} = \exp\left(-\left(s^{\text{SV}}\right)^2 / 2\left(\sigma^{\text{SV}}\right)^2\right), \quad (8)$$

where s^{SV} is the absolute difference between the actual head direction θ (in degrees) of the agent and the head direction at which the agent would be looking towards the centre of the spatial view field, and σ^{SV} is the standard deviation.

The spatial view cell also receives inputs from the layer of combination cells. During testing in the absence of visual input, the activation of the spatial view cell is driven by the inputs from the combination cells, and is given by

$$h^{\text{SV}} = \sum_j w_j^3 r_j^{\text{COMB}} \quad (9)$$

where w_j^3 are the weights from combination cell j to the spatial view cell. Once the activation of the spatial view cell has been computed according to Eq. (9), the firing rate of the cell is then given by the sigmoid transfer function

$$r^{\text{SV}} = \frac{1}{1 + e^{-2\beta(h^{\text{SV}} - \alpha)}}, \quad (10)$$

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

2.3. Training procedure and learning rules

An underlying assumption of the model is that during learning, when visual cues are available to the agent, the visual inputs dominate other excitatory inputs to the spatial view cell. In particular, in the light, the spatial view cell is stimulated maximally by a particular view. The initial learning phase with visual cues available is able to set up the synaptic weights such that, during subsequent testing in the dark, the spatial view cell maintains the same firing properties that it had in the dark. In particular, the spatial view cell fires maximally when the agent is looking towards the centre of the spatial view field.

The training and testing environment is a unit square containment area as shown in Fig. 4. During training and testing, the agent is rotated clockwise through 360° on the spot in two different locations within the square containment area. The spatial view cell is driven by visual input from a location on the North wall $(x=0.5, y=1.0)$. This gives rise to a spatial view field on the North wall about this position. When the agent is in location 1 $(x=0.25, y=0.75)$, the agent will be looking towards the centre of the spatial view field and the spatial view cell will fire maximally when the agent has a head direction of 45°. However, when the agent is in location 2 $(x=0.75, y=0.75)$, the agent will be looking towards the centre of the spatial view field and the spatial view cell will fire maximally when the agent has a head direction of 315°.

Table 1

Parameter values for model

| | |
|--|-----------------------|
| Number of head direction cells | 2500 |
| Number of place cells | $50 \times 50 = 2500$ |
| Number of combination cells | 2500 |
| Number of training epochs | 50 |
| Learning rate k^1 | 0.001 |
| Learning rate k^2 | 0.001 |
| Learning rate k^3 | 0.001 |
| Number of w^1 connections received by each combination cell | 50 |
| Number of w^2 connections received by each combination cell | 50 |
| Number of w^3 connections received by each spatial view cell | 2500 |
| Sparseness of firing rates in layer of combination cells | 0.01 |
| Standard deviation σ^{HD} | 10° |
| Standard deviation σ^{P} | 0.1 |
| Standard deviation σ^{SV} | 10° |
| Sigmoid threshold α | 0.14 |
| Sigmoid slope β | 20.0 |

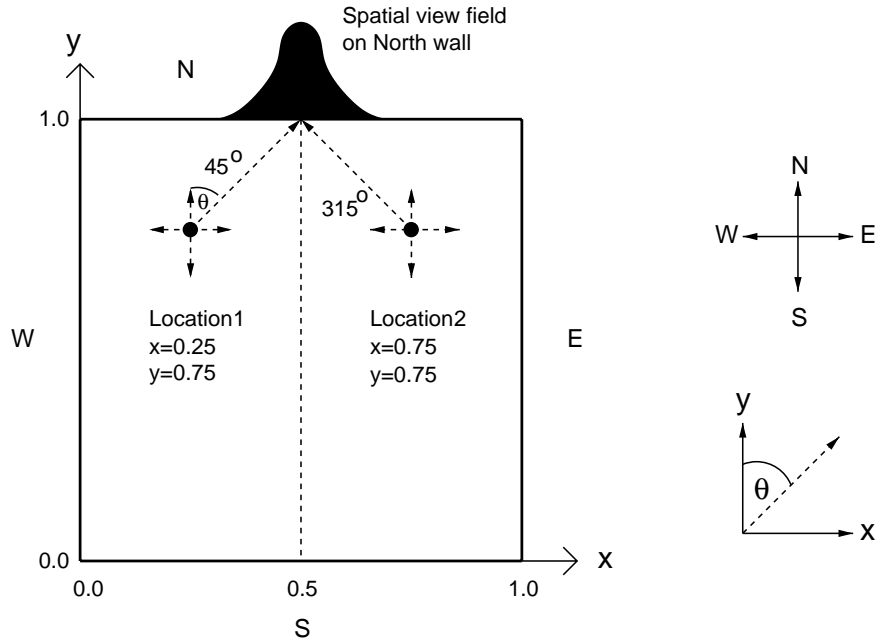


Fig. 4. The training and testing environment is a unit square containment area. During training and testing, the agent is rotated clockwise through 360° on the spot in two different locations within the square containment area. The spatial view cell is driven by visual input from a location on the North wall ($x=0.5, y=1.0$). This gives rise to a spatial view field on the North wall. When the agent is in location 1 ($x=0.25, y=0.75$), the agent will be looking towards the centre of the spatial view field and the spatial view cell will fire maximally when the agent has a head direction of 45° . However, when the agent is in location 2 ($x=0.75, y=0.75$), the agent will be looking towards the centre of the spatial view field and the spatial view cell will fire maximally when the agent has a head direction of 315° degrees.

One epoch of training constitutes one revolution in each of the two training locations. In the simulations, the agent performed 50 training epochs.

During training, the weights w_{ij}^1 from head direction cell j to combination cell i were updated according to an associative Hebb learning rule

$$\delta w_{ij}^1 = k^1 r_i^{\text{COMB}} r_j^{\text{HD}} \quad (11)$$

where r_i^{COMB} is the firing rate of combination cell i , r_j^{HD} is the firing rate of head direction cell j , δw_{ij}^1 is the change of synaptic weight, and k^1 is the learning rate constant.

Similarly, the weights w_{ij}^2 from place cell j to combination cell i were updated according to an associative Hebb learning rule

$$\delta w_{ij}^2 = k^2 r_i^{\text{COMB}} r_j^{\text{P}} \quad (12)$$

where r_i^{COMB} is the firing rate of combination cell i , r_j^{P} is the firing rate of place cell j , δw_{ij}^2 is the change of synaptic weight, and k^2 is the learning rate constant.

After each timestep of the learning phase, the w_{ij}^1 and w_{ij}^2 weights are rescaled according to

$$\sqrt{\sum_j (w_{ij}^1)^2 + \sum_j (w_{ij}^2)^2} = 1, \quad (13)$$

where the sums are over all head direction cells and place cells j . This ensures that the weight vectors of all combination cells are normalised to length 1. This is important for the layer of combination cells to operate effectively as a competitive

network (Hertz et al., 1991; Rolls and Deco, 2002; Rolls and Treves, 1998).

The weights w_j^3 from combination cell j to the spatial view cell are updated according to a learning rule which combines associative LTP with homosynaptic LTD

$$\delta w_j^3 = k^3 (r^{\text{SV}} - \langle r^{\text{SV}} \rangle) r_j^{\text{COMB}} \quad (14)$$

where r^{SV} is the firing rate of the spatial view cell, $\langle r^{\text{SV}} \rangle$ is the average value of r^{SV} which was set to 0.5 in the simulations, r_j^{COMB} is the firing rate of combination cell j , δw_j^3 is the change of synaptic weight, and k^3 is the learning rate constant.

3. Simulation results

Values for the model parameters used in the simulations are given in Table 1.

Fig. 5 shows the firing rate of the spatial view cell during training in the light. Results for location 1 are presented in the left column, while results for location 2 are presented in the right column. Top row: firing rate of spatial view cell as a function of head direction θ during training in the light. It can be seen that when the agent is in location 1 ($x=0.25, y=0.75$), the spatial view cell fires maximally when the agent has a head direction of 45° . However, when the agent is in location 2 ($x=0.75, y=0.75$), the spatial view cell fires maximally when the agent has a head direction of 315° . Bottom row: firing rate of spatial view cell as a function of position on the wall perimeter, towards which the agent is looking during training in the light. It can be seen that when the agent is in either location 1 or 2, the spatial view cell fires maximally when the agent is looking

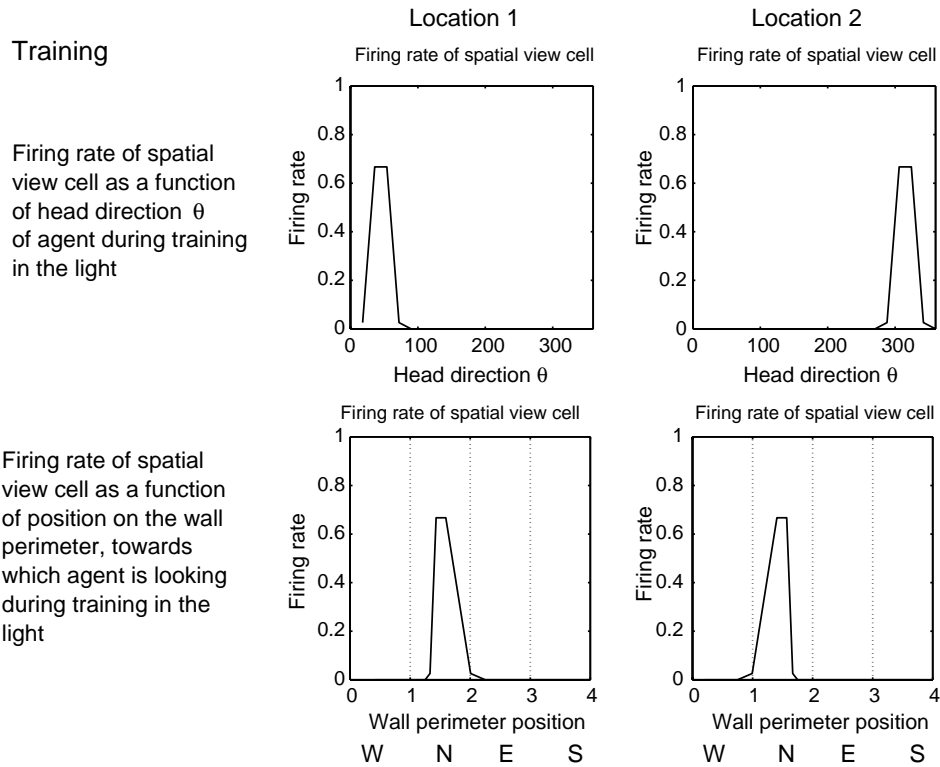


Fig. 5. Firing rate of a spatial view cell during training in the light. Results for location 1 are presented in the left column, while results for location 2 are presented in the right column. Top row: firing rate of the spatial view cell as a function of head direction θ during training in the light. It can be seen that when the agent is in location 1 ($x=0.25, y=0.75$), the spatial view cell fires maximally when the agent has a head direction of 45° . However, when the agent is in location 2 ($x=0.75, y=0.75$), the spatial view cell fires maximally when the agent has a head direction of 315° . Bottom row: firing rate of the spatial view cell as a function of location on the wall, towards which the agent is looking during training in the light. It can be seen that when the agent is in either location 1 or 2, the spatial view cell fires maximally when the agent is looking towards the centre of the spatial view field on the North wall ($x=0.5, y=1.0$).

towards the centre of the spatial view field on the North wall ($x=0.5, y=1.0$).

Fig. 6 shows the firing rate of spatial view cell during testing in the dark. Results for location 1 are presented in the left column, while results for location 2 are presented in the right column. Top row: firing rate of spatial view cell as a function of head direction θ during testing in the dark. Bottom row: firing rate of spatial view cell as a function of position on the wall perimeter, towards which the agent is looking during testing in the dark. By comparing with Fig. 5, it is evident that during testing in the dark, the firing rates of the spatial view cell are similar to what they were during training in the light. In particular, the spatial view cell fires maximally when the agent is looking towards the centre of the spatial view field on the North wall ($x=0.5, y=1.0$). The plots in the left and right columns are with the agent in two different locations, and these activate different sets of head direction and place combination cells. The plots are not symmetrical just due to small imperfections in the way in which the head direction and place combination cells self-organize.

4. Discussion

After the model has been trained, the spatial view cells can be updated by whatever combination of place cell and head direction cells is currently active. The head direction and place

cells can themselves be idiotically updated by movements made in the dark, in for example the ways described by Stringer, Rolls, Trappenberg, and De Araujo (2002); Stringer, Trappenberg, Rolls, and De Araujo (2002). Thus the model shows how spatial view cells could learn the correct combinations of head direction and place cell firing to account for the spatial view cells being updated by idiotic movements in the dark (Robertson et al., 1998).

It was found that to enable the combination neurons to self-organise so that different neurons responded to different combinations of place cell and head direction cell firing a number of factors were useful. First, if each neuron in the combination layer received approximately equal numbers of inputs from the head direction and place cell populations, then this increased the chance that combination cells were formed. If by chance connectivity a neuron in the combination layer received inputs mainly from head direction cells, then it was likely to self-organize to respond to head direction and not to a combination of head direction and place. Having neurons genetically specified to receive approximately a given number of input connections from a given source may be one solution to this (Rolls & Stringer, 2000). Second, having diluted random connectivity from the place and head direction cells to the combination cells helped each combination cell to form representations that were distinct from each other. As noted by Rolls and Deco (2002); Rolls and Treves (1998), the diluted

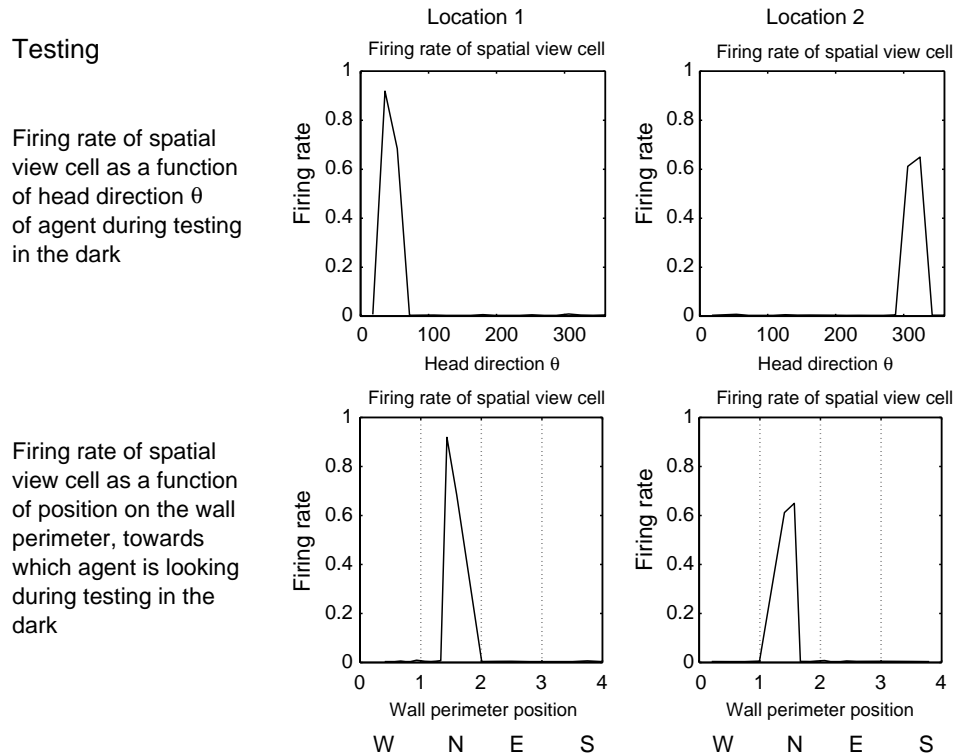


Fig. 6. Firing rate of spatial view cell during testing in the dark. Results for location 1 are presented in the left column, while results for location 2 are presented in the right column. Top row: firing rate of the spatial view cell as a function of the head direction θ during testing in the dark. Bottom row: firing rate of the spatial view cell as a function of the location on the wall towards which the agent is looking during testing in the dark. By comparing with Fig. 5, it is evident that during testing in the dark, the firing rates of the spatial view cell are similar to those during training in the light. In particular, the spatial view cell fires maximally when the agent is looking towards the centre of the spatial view field on the North wall ($x=0.5$, $y=1.0$), independently of the place where the agent is located.

connectivity typical of forward connections in the brain, where the connectivity may rarely be over 5%, may be advantageous for this type of computation. Third, with approximately equal numbers of inputs from the place cell and head direction cell sources to a given combination cell, it was further important that the sparseness of each of these inputs was approximately the same. Very unequal sparsenesses again made the self-organization unbalanced, in that the input with the less sparse representation tended to dominate the types of cell that formed in the combination layer. The magnitude of the inputs to the combination cells have to be approximately comparable, and this could be achieved in a number of ways involving the number of connections from each source, their mean strength, and their sparseness. Fourth, it was useful to keep the sparseness of the representation in the combination cell layer fairly low, so that many different types of combination neuron formed. Fifth, it was important, as in all competitive networks (Hertz et al., 1991; Rolls & Deco, 2002; Rolls & Treves, 1998), to make sure that no cells dominated the learning, by using some form of weight normalization on each neuron in the competitive network. Vector length normalization was used in the simulations described here, but more biologically plausible methods utilizing heterosynaptic long-term depression are available (Rolls & Deco, 2002; Rolls & Treves, 1998). Sixth, it is a key concept of the architecture shown in Fig. 3 that combination cells, which act as expansion recoding devices (Rolls & Deco, 2002; Rolls & Treves, 1998), are necessary to

solve the mapping from place cell and head direction cell representations to spatial view cell representations, because otherwise the inputs are too compressed. The expansion recoding provided by the combination cells enables spatial view cells to learn using simple associative learning rules to respond correctly to the different combinations of place cell and head direction cell firing that together define the conditions when a particular spatial view is currently visible.

Once the combination cells had self-organised by competitive learning, the spatial view cells are forced into activity by a particular spatial view in the light, and by pattern association learn which input patterns from the combination cells are associated with the current view. It was found that for this pattern association learning system to operate well, not only were increases in synaptic weights needed, but so were decreases, in particular corresponding to homosynaptic long-term depression (LTD). Homosynaptic LTD corresponds to decreasing the synaptic weight if there is no post-synaptic activity (because the neuron is not being driven by the current spatial view) but there is presynaptic activity. The homosynaptic LTD allows the spatial view cell(s) to learn not to respond to the activity of combination cells that are not associated with the current spatial view. This is important, as some cells in the combination layer self-organised to respond just to place, or just to head direction, and this type of representation is not useful for the mapping to spatial view, for which combinations of head direction and place are important.

Inherent in the use of homosynaptic LTD is the concept that if a visual input for a particular spatial view is not being received by the spatial view neuron, then it must be forced into low activity. This could be arranged during learning by having strong inputs from spatial views to the spatial view neurons, and strong mutual inhibition between the spatial view neurons, so that if they are not being activated by a visual input then they would be inhibited by other neurons into having low activity. Training the combination cell to spatial view cell connections without LTD, or with heterosynaptic LTD, did not produce such efficient learning of combination cell to spatial view cell connectivity as training with homosynaptic LTD.

We note that there are reports of some neurons responding to combinations of head direction and place in rodents, found for example when rodents traverse a linear runway. In this situation, a place cell may fire when the rat is running in one direction, and can see one view of the environment, but may not fire when the rat is running in the opposite direction and has a different view of the environment (McNaughton et al., 1983; Muller, Bostock, Taube, & Kubie, 2004). In a sense, these neurons described in rodents almost code for spatial view, except that they are not place independent, as are primate spatial view cells. By adding the final pattern association stage to the combination cells, the network shown in Fig. 3 associates together different place-specific head direction/place combination cells that fire when the animal has a particular spatial view from different places in the environment. In addition, hippocampal neurons that respond to the place where a macaque is located have been described (Hori, Tabuchi, Matsumura, Tamura, Eifuku and Nishijo, 2003; Ludvig, Tang, Gohil, & Botero, 2004), and such place neurons, if their responses are independent of spatial view, could provide the representation of place that is part of the model described here.

The relation between rodent place cells in directional environments such as a runway, in which the place cells have (head) directional properties, and primate spatial view cells, deserves further comment. A primate spatial view cell responds when the primate is looking at a given location in the environment, and does this when the monkey is in different places in the environment, provided that the primate is looking at the spatial view location (Robertson et al., 1998; Rolls, 1999; Rolls et al., 1997; Rolls et al., 1998). It would be very interesting to try an experiment with rats in which the rat was running in a runway towards a particular wall of the room, and a place field was located at the halfway point of the runway arm. The runway arm would be displaced laterally to a different place in the room, but reoriented so that it still faced towards the same wall location, which would be mainly what would be visible to the rat when running up the runway. It would then be interesting to know whether the rat place cell still responded, when the rat was approximately halfway along the runway, but in a new place in the room. If so, this would draw out the similarity between spatial view cells in primates, and place cells when dominated by extramaze cues in rodents. Further, such a result would be in line with the hypothesis and model of De Araujo et al. (2001), which shows that the difference in the field of view of rodents (which may be 270°) and primates

(which with the fovea has high precision for a small viewing angle), may help to account for apparent differences between rodent place cells and primate spatial view cells. In particular, if the rodent place cell encoded the location of several cues separated by a large angle in the environment, this would effectively define a place. If a primate spatial view cell also encoded the location of several cues in the environment but now separated by a small angle, this would effectively define a spatial view cell. By performing the experiment just described with a rodent running in a narrow runway, effectively the visual environment would be dominated by the extramaze cues on for example the wall at the end of the runway, and under these conditions, a rat place cell might be seen to be in some respects quite similar to a primate spatial view cell.

Finally, we note that spatial view cells may have a number of functions. One could be in navigation, in which forming associations between a particular spatial view and whole body movements (which are known to be represented in the primate hippocampus (O'Mara et al., 1994), and an example of which might be turning to the right), would be helpful in route finding. This would utilize the autoassociative memory properties of the hippocampus described in the first part of this paper. A second could be in the memory for where objects have been seen in space, a prototypical exemplar of event or episodic memory. Evidence that supports this is that during the performance of an episodic memory task, some macaque hippocampal neurons code for the spatial location, some for the object being shown, and some, combination neurons, responding to particular combinations of an object and the place where it was shown 'out there' in space, in the room (Rolls, Xiang, & Franco, 2005). A third appears to be in the memory of the particular rewards associated with different locations in a scene, in that some hippocampal neurons respond to the spatial locations in a number of scenes at which one type of reward is located, and other neurons respond to the locations in a number of scenes where a different type of reward is located (Rolls & Xiang, 2004). Moreover, the majority of these neurons reverse the location in the scene to which the neurons respond when the rewards available at the different locations reverse. Moreover, these neurons are especially involved in spatial view/scene location associations with rewards, and not with object-reward associations, in that these neurons do not respond differentially to a rewarded and punished stimulus in an object-reward visual discrimination task. Further, scene memory of this type is impaired by excitotoxic lesions of the primate hippocampus (Murray, Baxter, & Gaffan, 1998).

Thus there is increasing evidence on how spatial view neurons in the primate hippocampus may be used in a number of different types of memory task including object and reward associations to places 'out there', and potentially in spatial view to response associations (cf. Miyashita et al. (1989); Cahusac et al. (1993)). The new model presented here shows how primate spatial view cells could be updated in the dark by idiothetic signals derived from idiothetically updated head direction and place cells.

Acknowledgements

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.
- Brasted, P. J., Bussey, T. J., Murray, E. A., & Wise, S. P. (2003). Role of the hippocampal system in associative learning beyond the spatial domain. *Brain*, *126*, 1202–1223.
- Cahusac, P. M. B., Rolls, E. T., Miyashita, Y., & Niki, H. (1993). Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task. *Hippocampus*, *3*, 29–42.
- Cho, J., & Sharp, P. E. (2001). Head direction, place and movement correlates for single cells in the rat retrosplenial cortex. *Behavioural Neuroscience*, *115*, 3–25.
- De Araujo, I. E. T., Rolls, E. T., & Stringer, S. M. (2001). A view model which accounts for the response properties of hippocampal primate spatial view cells and rat place cells. *Hippocampus*, *11*, 699–706.
- Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., et al. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*, 184–188.
- Földiák, P. (1991). Learning invariance from transformation sequences. *Neural Computation*, *3*, 193–199.
- Franco, L., Rolls, E. T., Aggelopoulos, N. C., & Jerez, J. M. (2005). *Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex*.
- Fried, I., MacDonald, K. A., & Wilson, C. (1997). Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. *Neuron*, *18*, 753–765.
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., & Moser, M.-B. (2004). Spatial representation in the entorhinal cortex. *Science*, *204*, 1258–1264.
- 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.
- Gothard, K. M., Skaggs, W. E., & McNaughton, B. L. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. *Journal of Neuroscience*, *16*, 8027–8040.
- Hasselmo, M. E., Schnell, E., & Barkai, E. (1995). Learning and recall at excitatory recurrent synapses and cholinergic modulation in hippocampal region CA3. *Journal of Neuroscience*, *15*, 5249–5262.
- Hertz, J. A., Krogh, A., & Palmer, R. G. (1991). *Introduction to the theory of neural computation*. Wokingham, UK: Addison-Wesley.
- Hölscher, C., Rolls, E. T., & Xiang, J. Z. (2003). Perirhinal cortex neuronal activity related to long term familiarity memory in the macaque. *European Journal of Neuroscience*, *18*, 2037–2046.
- Hori, E. E., Tabuchi, E., Matsumura, N., Tamura, R., Eifuku, S., Nishijo, H., et al. (2003). Representation of place by monkey hippocampal neurons in real and virtual translocation. *Hippocampus*, *13*, 190–196.
- Lassalle, J. M., Bataille, T., & Halley, H. (2000). Reversible inactivation of the hippocampal mossy fiber synapses in mice impairs spatial learning, but neither consolidation nor memory retrieval, in the Morris navigation task. *Neurobiology of Learning and Memory*, *73*, 243–257.
- Ludvig, N., Tang, H. M., Gohil, B. C., & Botero, J. M. (2004). Detecting location-specific neuronal firing rate increases in the hippocampus of freely moving monkeys. *Brain Research*, *1014*, 97–109.
- Markus, E. J., Qin, Y. L., Leonard, B., Skaggs, W., McNaughton, B. L., & Barnes, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience*, *15*, 7079–7094.
- Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of The Royal Society of London, Series B*, *262*, 23–81.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457.
- McNaughton, B. L., Barnes, C. A., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, *52*, 41–49.
- Miyashita, Y., Rolls, E. T., Cahusac, P. M. B., Niki, H., & Feigenbaum, J. D. (1989). Activity of hippocampal neurons in the monkey related to a conditional spatial response task. *Journal of Neurophysiology*, *61*, 669–678.
- Muller, R. U., Bostock, E., Taube, J. S., & Kubie, J. L. (2004). On the directional firing properties of hippocampal place cells. *Journal of Neuroscience*, *14*, 7235–7251.
- Muller, R. U., Kubie, J. L., Bostock, E. M., Taube, J. S., & Quirk, G. J. (1991). Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In J. Paillard (Ed.), *Brain and Space* (pp. 296–333). Oxford: Oxford University Press.
- Muller, R. U., Ranck, J. B., & Taube, J. S. (1996). Head direction cells: properties and functional significance. *Current Opinion in Neurobiology*, *6*, 196–206.
- Murray, E. A., Baxter, M. G., & Gaffan, D. (1998). Monkeys with rhinal cortex damage or neurotoxic hippocampal lesions are impaired on spatial scene learning and object reversals. *Behavioral Neuroscience*, *112*, 1291–1303.
- O'Keefe, J. (1984). Spatial memory within and without the hippocampal system. In W. Seifert (Ed.), *Neurobiology of the hippocampus* (pp. 375–403). London: Academic Press.
- O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, *381*, 425–428.
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, *34*, 171–175.
- 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.
- Ono, T., Nakamura, K., Nishijo, H., & Eifuku, S. (1993). Monkey hippocampal neurons related to spatial and nonspatial functions. *Journal of Neurophysiology*, *70*, 1516–1529.
- Ranck, J. B. (1985). Head direction cells in the deep cell layer of dorsolateral pre-subiculum in freely moving rats. In G. Buzsáki, & C. H. Vanderwolf (Eds.), *Electrical Activity of the Archicortex*. Budapest: Akadémiai Kiadó.
- 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. (1998). Spatial view cells in the primate hippocampus: Effects of removal of view details. *Journal of Neurophysiology*, *79*, 1145–1156.
- Robertson, R. G., Rolls, E. T., Georges-François, P., & Panzeri, S. (1999). Head direction cells in the primate pre-subiculum. *Hippocampus*, *9*, 206–219.
- Rolls, E. T. (1987). Information representation, processing and storage in the brain: analysis at the single neuron level. In J.-P. Changeux, & M. Konishi (Eds.), *The Neural and Molecular Bases of Learning* (pp. 503–540). Chichester: Wiley.
- Rolls, E. T. (1989). Functions of neuronal networks in the hippocampus and neocortex in memory. In J. Byrne, & W. Berry (Eds.), *Neural models of plasticity: experimental and theoretical approaches* (pp. 240–265). San Diego: Academic Press (chapter 13).
- Rolls, E. T. (1996). A theory of hippocampal function in memory. *Hippocampus*, *6*, 601–620.

- Rolls, E. T. (1999). Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus*, 9, 467–480.
- Rolls, E. T. (2000). Hippocampo-cortical and cortico-cortical backprojections. *Hippocampus*, 10, 380–388.
- Rolls, E. T. (2005). *Emotion explained*. Oxford: Oxford University Press.
- Rolls, E. T., & Deco, G. (2002). *Computational neuroscience of vision*. Oxford: Oxford University Press.
- Rolls, E.T., Franco, L., Stringer, S.M. (2005). The perirhinal cortex and long-term familiarity memory. *Quarterly Journal of Experimental Psychology, B Comparative and Physiological Psychology*, 58B, 234–245.
- Rolls, E. T., Miyashita, Y., Cahusac, P. M. B., Kesner, R. P., Niki, H., Feigenbaum, J., et al. (1989). Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *Journal of Neuroscience*, 9, 1835–1845.
- 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. (2000). On the design of neural networks in the brain by genetic evolution. *Progress in Neurobiology*, 61, 557–579.
- 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.
- Rolls, E. T., & Xiang, J.-Z. (2004). Reward-place neurons in the primate anterior hippocampus. *Society for Neuroscience Abstracts*.
- Rolls, E.T., Xiang, J.-Z. (2005a). Reward-spatial view representations and learning in the primate hippocampus, *Journal of Neuroscience*, 25, 6167–6174.
- Rolls, E.T., Xiang, J.-Z. (2005b). Spatial view cells in the primate hippocampus, and memory recall, *Reviews in the Neurosciences*.
- Rolls, E. T., Xiang, J.-Z., & Franco, L. (2005). Object, space and object-space representations in the primate hippocampus. *Journal of Neurophysiology*, 94, 833–844.
- 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. (1996). Multiple spatial/behavioral correlates for cells in the rat postsubiculum: Multiple regression analysis and comparison to other hippocampal areas. *Cerebral Cortex*, 238–259.
- Sharp, P. E., & Green, C. (1994). Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *Journal of Neuroscience*, 14, 2339–2356.
- Skaggs, W. E., Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1995). In G. Tesauro, D. S. Touretzky, & T. K. Leen, *A model of the neural basis of the rat's sense of direction. Advances in neural information processing systems* (Vol. 7) (pp. 173–180). Cambridge, Massachusetts: MIT Press.
- Stanton, P. K., & Sejnowski, T. J. (1989). Associative long-term depression in the hippocampus: Induction of synaptic plasticity by hebbian covariance. *Nature*, 339, 215–218.
- 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. (2005). Self-organizing continuous attractor network models of hippocampal spatial view cells. *Neurobiology of Learning and Memory*, 83, 79–92.
- Stringer, S. M., Rolls, E. T., Trappenberg, T. P., & De Araujo, I. E. T. (2002). 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., Rolls, E. T., Trappenberg, T. P., & De Araujo, I. E. T. (2003). Self-organizing continuous attractor networks and motor function. *Neural Networks*, 16, 161–182.
- Stringer, S. M., Trappenberg, T. P., Rolls, E. T., & De Araujo, I. E. T. (2002). 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., Goodridge, J. P., Golob, E. G., Dudchenko, P. A., & Stackman, R. W. (1996). Processing the head direction signal: A review and commentary. *Brain Research Bulletin*, 40, 477–486.
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum 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.
- Treves, A., & Rolls, E. T. (1991). What determines the capacity of autoassociative memories in the brain? *Network*, 2, 371–397.
- Treves, A., & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, 2, 189–199.
- Treves, A., & Rolls, E. T. (1994). A computational analysis of the role of the hippocampus in memory. *Hippocampus*, 4, 374–391.
- Wallis, G., & Rolls, E. T. (1997). Invariant face and object recognition in the visual system. *Progress in Neurobiology*, 51, 167–194.
- Wise, S. P., & Murray, E. A. (1999). Role of the hippocampal system in conditional motor learning: Mapping antecedents to action. *Hippocampus*, 9, 101–117.
- 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.