Network: Computation in Neural Systems
December 2006; 17(4): 419–445

informa
healthcare

Self-organizing path integration using a linked continuous attractor and competitive network: Path integration of head direction

SIMON M. STRINGER & EDMUND T. ROLLS

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

(Received 25 May 2006; accepted 11 September 2006)

Abstract
A key issue is how networks in the brain learn to perform path integration, that is update a represented position using a velocity signal. Using head direction cells as an example, we show that a competitive network could self-organize to learn to respond to combinations of head direction and angular head rotation velocity. These combination cells can then be used to drive a continuous attractor network to the next head direction based on the incoming rotation signal. An associative synaptic modification rule with a short term memory trace enables preceding combination cell activity during training to be associated with the next position in the continuous attractor network. The network accounts for the presence of neurons found in the brain that respond to combinations of head direction and angular head rotation velocity. Analogous networks in the hippocampal system could self-organize to perform path integration of place and spatial view representations.

Keywords: *Continuous attractor networks, self-organization, trace learning, competitive learning, head direction cells, path integration*

Introduction

The orientation or position of an animal within its environment, and where it is looking, are represented by neurons in different brain areas. Examples of such cells include head direction cells in rats (Ranck 1985; Taube et al. 1990a, 1996; Muller et al. 1996) and primates (Robertson et al. 1999) which respond maximally when the animal's head is facing in a particular preferred direction, and hippocampal place cells in rats (O'Keefe & Dostrovsky 1971; McNaughton et al. 1983; O'Keefe 1984; Muller et al. 1991; Markus et al. 1995) that fire maximally when the animal is in a particular location. In primates, hippocampal spatial view cells have been discovered in the hippocampus that code for the location at which a primate is looking (Rolls et al. 1997, 1998; Robertson et al. 1998; Georges-François et al. 1999; Rolls 1999; Rolls & Xiang 2006). A key problem is how can these systems *learn* to perform path integration, i.e., update their spatial representations using velocity

Correspondence: E. T. Rolls, Oxford University, Centre for Computational Neuroscience, Department of Experimental Psychology, South Parks Road, Oxford OX1 3UD, UK. Tel: +44 1865 271348. FAX: +44 1865 310447. E-mail: Edmund.Rolls@psy.ox.ac.uk

ISSN: 0954-898X print / ISSN 1361-6536 online © 2006 Informa UK Ltd.
DOI: 10.1080/09548980601004032

(e.g., idiothetic) signals in the dark. Many models have hardwired synaptic weights (Skaggs et al. 1995; Redish et al. 1996; Zhang 1996). We have developed a series of self-organising path integration models for head direction cells, place cells and spatial view cells (Stringer et al. 2000a, 2000b, 2005). However, these models utilize Sigma–Pi synaptic connections, which may not be biologically plausible as described below.

In this paper, we introduce a network which does not need Sigma–Pi connections, but which uses a competitive layer of cells that self-organize to represent combinations of velocity and the positional state of the agent. The proposal is backed by experimental evidence for the existence of such combination cells in various spatial processing areas of the brain. For example, in the rat head direction system some cells represent a combination of head direction and rotation in a particular direction (Taube et al. 1990b; Sharp 1996; Bassett & Taube 2005). In the hippocampus, cells are responsive to combinations of place and forward velocity (McNaughton et al. 1983). Similarly, in the entorhinal cortex some grid cells (Hafting et al. 2005) are modulated by running speed (Sargolini et al. 2006). By combining such combination responsive neurons with a coupled attractor network, we show how path integration could be performed and how the whole system could self-organize. The same approach could in principle be applied to other path integration systems, which represent, for example, place in the environment or spatial view. In this paper, we show how this principle can be applied to the problem of head direction cells.

Head direction cells in rats (Ranck 1985; Taube et al. 1990a, 1996; Muller et al. 1996) respond maximally when the animal's head is facing in a particular preferred direction. These cells are able to maintain their firing in the absence of visual input. In particular, an important function of these cells is *orientation path integration.* That is, when the animal moves in darkness, the spatial representation may be updated by idiothetic (self-motion) cues (Taube et al. 1996). Similar cells have been found in primates and also show path integration (Robertson et al. 1999). Some cells have been found in rats which represent a combination of head direction and angular rotation velocity (Taube et al. 1990b; Sharp 1996; Bassett & Taube 2005). Computer modelling studies need to address how such combination cells fit into an overall head direction system in the rat, and the role that combination cells play in path integration.

A number of models of head direction cells rely on 'Continuous Attractor' neural networks (CANN) (Skaggs et al. 1995; Redish et al. 1996; Zhang 1996). In a continuous attractor network there are not only associatively modifiable recurrent connections between the nodes, but also the representation is continuous, in that neighbouring nodes have overlapping (typically Gaussian) receptive fields (Amari 1977; Taylor 1999; Rolls & Deco 2002). The state space is thus inherently continuous. These network models are able to maintain a local packet (or bubble) of neuronal activity representing the current head direction in the absence of visual input. In addition, some models can use idiothetic (self-motion) cues to update the representation of head direction as the agent moves in the dark (Skaggs et al. 1995; Redish et al. 1996; Zhang 1996; Xie et al. 2002; Hahnloser 2003). A major limitation of some of the models studied so far is that the connections are pre-specified by the modeller, with for example two types of connections to deal with clockwise and anticlockwise turns (Skaggs et al. 1995; Redish et al. 1996; Zhang 1996). For example, the model of Skaggs et al. (1995) had a head direction cell layer and a layer with cells responding to combinations of head direction and angular velocity, but was not a self-organising model. A self-organising model of head direction cells, in which the synaptic weights in the network are learned through interaction with the environment, was proposed by Stringer et al. (2002a). The model used a 'trace' learning rule to associate a combination of two presynaptic terms, namely the recent

activity within the layer of head direction cells and an angular velocity of rotation signal, with the current postsynaptic activity of head direction cells. The model thus relied on the use of higher order Sigma–Pi synapses, which suffer from a lack of biological plausibility. Sigma–Pi synapses learn an association between a combination of two or more presynaptic inputs and the post-synaptic state. They are thus second (or higher) order synapses. One anatomical arrangement that might achieve this would be for the rotation cells to have presynaptic terminals on every (recurrent collateral) head direction cell synapse onto a head direction cell neuron, but this is not very biologically plausible as an implementation. Another possibility might be closely paired synapses from rotation and head direction synapses onto a very thin dendrite that would allow only local voltage interactions (Koch 1999), but again this is not known to occur. Because Sigma–Pi synapses are not thought to be a feature of many connections in the brain, an important aim of the work described here was to investigate how a self-organising path integration neuronal system might be set up with simpler first order associative synapses, which are common in many brain areas.

In this paper, we investigate a new 2-layer self-organising model of the rat head direction system, which does not rely on the use of biologically unrealistic Sigma–Pi synapses. The model incorporates a layer of head direction cells and a new additional layer of combination cells, which learn to respond to combinations of head direction and rotation. In this paper, we investigate how the synaptic connections may develop in the network using biologically plausible learning rules. During training, connections from the head direction cells and rotation cells to the combination cells self-organise using unsupervised Hebbian competitive learning (described by Hertz et al. (1991) and Rolls & Deco (2002)) to enable the combination cells to learn to represent combinations of head direction and rotation. At the same time, a trace learning rule associates the recent activity within the layer of combination cells with the current activity of the head direction cells. The additional layer of combination cells thus removes the need in the model of Stringer et al. (2002a) to combine separate inputs from separate populations of head direction cells and rotation cells within higher order Sigma–Pi synapses to update the representation within the layer of head direction cells during path integration. The new model described here also is consistent with and helps to account for the presence of neurons in for example the rat presubiculum and dorsal thalamic nucleus that respond to combinations of head direction and angular rotation (Taube et al. 1990b; Sharp 1996; Bassett & Taube 2005). The new model is also prototypical of path integration systems that may be implemented in the brain, including those used for path integration in rat hippocampal place cells (McNaughton et al. 1983; O'Keefe 1984; Muller et al. 1991; Markus et al. 1994) and primate hippocampal spatial view cells (Robertson et al. 1998). We emphasize that all the synaptic connections in the model are completely self-organizing, and the only inputs needed are head direction and head rotation velocity inputs during training, and only head rotation velocity inputs during the path integration.

The model

The model has the neural architecture shown in Figure 1. There are two connected networks. Firstly, there is a network of head direction cells (with firing rate r_i^{HD} for head direction cell i) which represent the current head direction of the agent. The head direction cell network operates as a rotation path integrator using a continuous attractor neural network (CANN), which is a recurrently connected network that can maintain the firing of its neurons to represent any location along a continuous physical dimension representing the state of the agent (Amari 1977; Taylor 1999). Secondly, there is a network of combination cells (with

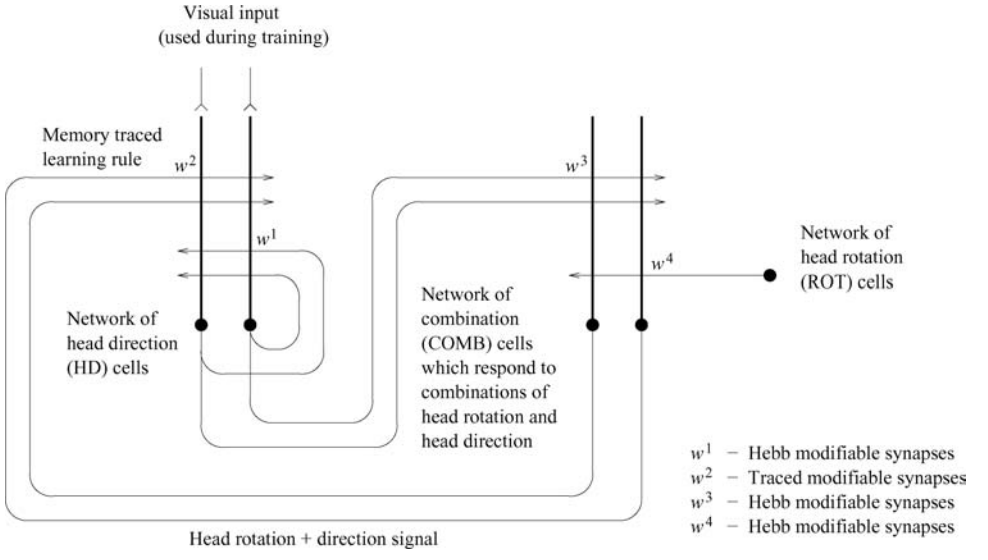

Figure 1. Network architecture for two-layer self-organising neural network model of the head direction system. The network architecture contains a layer of head direction (HD) cells representing the head direction of the agent, a layer of combination (COMB) cells representing a combination of head direction and rotational velocity, and a layer of rotation (ROT) cells which become active when the agent rotates. There are four types of synaptic connection in the network, which operate as follows. The w_{ij}^1 synapses are Hebb-modifiable recurrent connections between head direction cells. These connections help to support stable packets of activity within this continuous attractor layer of head direction cells in the absence of visual input. The combination cells receive inputs from the head direction cells through the Hebb-modifiable w_{ij}^3 synapses, and inputs from the rotation cells through the Hebb-modifiable w_{ij}^4 synapses. These synaptic inputs encourage combination cells by competitive learning to respond to particular combinations of head direction and rotational velocity. In particular, the combination cells only become significantly active when the agent is rotating. The head direction cells receive inputs from the combination cells through the w_{ij}^2 synapses. The w_{ij}^2 synapses are trained using a 'trace' learning rule, which incorporates a temporal trace of recent combination cell activity. This rule introduces an asymmetry into the w_{ij}^2 weights, which plays an important role in shifting the activity packet through the layer of head direction cells during path integration in the dark.

firing rate r_i^{COMB} for combination cell i), which represent combinations of head direction and rotation. The combination cells operate as a competitive network, and develop their firing properties during training. Neurons in the brain that have Gaussian tuning to head direction, and are also tuned for the direction of angular head velocity have been described in for example the dorsal tegmental nucleus (Bassett & Taube 2005).

During the initial learning phase, each head direction cell i receives an external visual input e_i, which carries information about the head direction of the agent. When visual cues are available, these external inputs dominate other excitatory inputs to the head direction cells, and force each head direction cell to respond best to a particular head direction of the agent, with less firing as the agent's head direction moves away from the preferred head direction.

During learning, the recurrent connections w_{ij}^1 from head direction cell j to head direction cell i use associative synaptic modification so that the synaptic strengths between the head direction cells reflect the distance between the directions represented by the head direction cells. The recurrent connectivity implemented by w_{ij}^1 allows the network of head direction cells to operate as a continuous attractor network and support stable patterns of firing in the absence of external visual input, so that the agent can operate in the dark.

The layer of combination cells receives two kinds of synaptic connections: w_{ij}^3 from head direction cell j to combination cell i, and w_{ij}^4 from rotation cell j to combination cell i. During training, the w_{ij}^3 and w_{ij}^4 connections to the layer of combination cells self-organise using a Hebbian competitive learning rule, which enable the layer of combination cells to operate as a competitive network and thus learn to represent combinations of head direction and rotation. That is, after the unsupervised learning, different cells in this network respond to different combinations of head direction and rotation.

The path integration uses the connections w_{ij}^2 from combination cell j to head direction cell i. These connections are learned by an associative learning rule with a short term memory trace in the presynaptic firing, which allows the traced activity within the combination cell network (which represents a combination of the preceding head direction and rotation signal) to be associated with the current head direction. The *principle of operation* is that during training the combination cells that were active in the immediately preceding time become associated with the particular head direction currently represented in the head direction cell continuous attractor network. Then, later on after training, if this particular set of combination cells is active (as a result of activity in the w^3 and w^4 synapses), the head direction cell attractor is moved by the associatively modified w^2 inputs from the combination cells to the next appropriate head direction position. Thus, after training, the firing of a particular cluster of combination cells should stimulate the firing of further head direction cells such that the pattern of activity within the network of head direction cells evolves continuously to faithfully reflect and track the changing head direction of the agent as it rotates.

The activation h_i^{HD} of head direction cell i in the model is governed by

$$\tau \frac{dh_i^{\mathrm{HD}}(t)}{dt} = -h_i^{\mathrm{HD}}(t) + \frac{\phi_1}{C^{\mathrm{HD}\to\mathrm{HD}}} \sum_j \left(w_{ij}^1 - w^{\mathrm{INH}}\right) r_j^{\mathrm{HD}}(t) + e_i + \frac{\phi_2}{C^{\mathrm{COMB}\to\mathrm{HD}}} \sum_j w_{ij}^2 r_j^{\mathrm{COMB}} \tag{1}$$

where the activation h_i^{HD} is driven by the following terms. The term r_j^{HD} is the firing rate of head direction cell j, w_{ij}^1 is the excitatory (positive) synaptic weight from head direction cell j to head direction cell i, and w^{INH} is a global constant describing the effect of inhibitory interneurons within the layer of head direction cells.[1] Further terms in Equation 1 are as follows. The term τ is the time constant of the system. The term e_i represents an external visual input to head direction cell i. When the agent is denied visual input, the term e_i is set to zero. Thus, in the absence of visual input, the key term driving the head direction cell activations in Equation 1 is a sum of inputs from the combination cells $\sum_j w_{ij}^2 r_j^{\mathrm{COMB}}$, where r_j^{COMB} is the firing rate of combination cell j, and w_{ij}^2 is the corresponding strength of connection from this cell.[2] The firing rate r_i^{HD} of head direction cell i is determined from the activation to h_i^{HD} and the sigmoid activation function

$$r_i^{\mathrm{HD}}(t) = \frac{1}{1 + e^{-2\beta\left(h_i^{\mathrm{HD}}(t) - \alpha\right)}}, \tag{2}$$

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

[1] The scaling factor $(\phi_1 / C^{\mathrm{HD}\to\mathrm{HD}})$ controls the overall strength of the recurrent inputs to the layer of head direction cells, where ϕ_1 is a constant and $C^{\mathrm{HD}\to\mathrm{HD}}$ is the number of presynaptic connections received by each head direction cell from other head direction cells.

[2] The scaling factor $\frac{\phi_2}{C^{\mathrm{COMB}\to\mathrm{HD}}}$ controls the overall strength of the combination cell inputs, where ϕ_2 is a constant, and $C^{\mathrm{COMB}\to\mathrm{HD}}$ is the number of connections received by each head direction cell.

The recurrent synapses w_{ij}^{1} in the head direction cell continuous attractor are trained by a local associative (Hebb) rule

$$\delta w_{ij}^{1} = k^{1} r_{i}^{\mathrm{HD}} r_{j}^{\mathrm{HD}}. \tag{3}$$

This rule increases the strength of the synaptic connections between head direction cells that represent nearby head directions of the agent, and which tend to be co-active due to broadly tuned, overlapping receptive fields. To bound the synaptic weights, weight normalization was used. To implement weight normalization, after each timestep of the learning phase, the recurrent synaptic weights between neurons within the continuous attractor network were rescaled to ensure that for each head direction cell i we have

$$\sqrt{\sum_{j} \left(w_{ij}^{1}\right)^{2}} = 1, \tag{4}$$

where the sum is over all head direction cells j. Such a renormalization process may be achieved in biological systems through synaptic weight decay (Oja 1982; Rolls & Treves 1998; Rolls & Deco 2002). The renormalization (Equation 4) helps to ensure that the learning rules are convergent in the sense that the recurrent synaptic weights between neurons within the continuous attractor network settle down over time to steady values.

The learning rule used to update the synapses w_{ij}^{2} can be expressed by

$$\delta w_{ij}^{2} = k^{2} r_{i}^{\mathrm{HD}} \bar{r}_{j}^{\mathrm{COMB}}, \tag{5}$$

where $\bar{r}_{j}^{\mathrm{COMB}}$ refers to a memory trace of the firing r_{j}^{COMB}. The trace value $\bar{r}$ of the firing rate r of a cell is calculated according to

$$\bar{r}(t + \delta t) = (1 - \eta) r(t + \delta t) + \eta \bar{r}(t) \tag{6}$$

where η is a parameter in the interval [0,1] which determines the relative contributions of the current firing and the previous trace. For $\eta = 0$ the trace becomes just the present firing rate, and as η is increased the contribution of preceding firing at times earlier than the current timestep is increased (Földiák 1991; Stringer et al. 2000a, 2000b; Rolls & Deco 2002). The equation simulates exponential decay of the trace value $\bar{r}$.

Possible ways in which such traces of previous neuronal activity could be implemented include short-term memory related firing in networks, and biophysical processes within neurons (see Stringer et al. (2002a)). To bound the synaptic weights, after each timestep of the learning phase, the w_{ij}^{2} synaptic weights were rescaled to ensure that for each head direction cell i we have

$$\sqrt{\sum_{j} \left(w_{ij}^{2}\right)^{2}} = 1, \tag{7}$$

where the sum is over all combination cells j. As noted above, such a renormalization process may be achieved in biological systems through synaptic weight decay (Oja 1982; Rolls & Treves 1998; Rolls & Deco 2002).

The combination cells are driven by synaptic inputs w_{ij}^{3} from the layer of head direction cells and synaptic inputs w_{ij}^{4} from a layer of rotation cells. The rotation cells are external inputs which become active when the agent starts to rotate. There would need to be two separate populations of rotation cells, for clockwise and anticlockwise rotation. Thus, the

activation h_i^{COMB} of combination cell i is governed by

$$\tau \frac{dh_i^{\mathrm{COMB}}(t)}{dt} = -h_i^{\mathrm{COMB}}(t) + \frac{\phi_3}{C^{\mathrm{HD}\rightarrow\mathrm{COMB}}} \sum_j w_{ij}^3 r_j^{\mathrm{HD}} + \frac{\phi_4}{C^{\mathrm{ROT}\rightarrow\mathrm{COMB}}} \sum_j w_{ij}^4 r_j^{\mathrm{ROT}} \quad (8)$$

where the activation h_i^{COMB} is driven by the following terms. The first term driving the activations of the combination cells in Equation 8 is the input from the head direction cells $\sum_j w_{ij}^3 r_j^{\mathrm{HD}}$, where r_j^{HD} is the firing rate of head direction cell j, and w_{ij}^3 is the corresponding strength of the connection from this cell. The second term driving the activations of the combination cells in Equation 8 is the input from the rotation cells $\sum_j w_{ij}^4 r_j^{\mathrm{ROT}}$, where r_j^{ROT} is the firing rate of rotation cell j, and w_{ij}^4 is the corresponding strength of the connection from this cell.[3]

Activity within the network of combination cells is driven by the head direction cells if and only if the rotation cells are also active. If the rotation cells stop firing then the activity in the combination cell network decays to zero according to the time constant τ. The firing rate r_i^{COMB} of combination cell i is determined from the activation h_i^{COMB} and the sigmoid activation function

$$r_i^{\mathrm{COMB}}(t) = \frac{1}{1 + e^{-2\beta\left(h_i^{\mathrm{COMB}}(t) - \alpha\right)}} \quad (9)$$

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

The synaptic weights w_{ij}^3 from the head direction cells to the combination cells are updated during learning according to

$$\delta w_{ij}^3 = k^3 r_i^{\mathrm{COMB}} r_j^{\mathrm{HD}}. \quad (10)$$

To bound the synaptic weights, after each timestep of the learning phase, the w_{ij}^3 synaptic weights were rescaled to ensure that for each combination cell i we have

$$\sqrt{\sum_j \left(w_{ij}^3\right)^2} = 1, \quad (11)$$

where the sum is over all head direction cells j.

The synaptic weights w_{ij}^4 from the rotation cells to the combination cells are updated during learning according to

$$\delta w_{ij}^4 = k^4 r_i^{\mathrm{COMB}} r_j^{\mathrm{ROT}}. \quad (12)$$

To bound the synaptic weights, after each timestep of the learning phase, the w_{ij}^4 synaptic weights were rescaled to ensure that for each combination cell i we have

$$\sqrt{\sum_j \left(w_{ij}^4\right)^2} = 1, \quad (13)$$

where the sum is over all rotation cells j.

During training in the light, the agent rotates on the spot and the synaptic connections are established as follows. Visual input drives the layer of head direction cells according to Gaussian head direction related tuning profiles. (The firing rates of head direction cells in

[3] The scaling factor $\frac{\phi_3}{C^{\mathrm{HD}\rightarrow\mathrm{COMB}}}$ controls the overall strength of the inputs from the head direction cells, where ϕ_3 is a constant, and $C^{\mathrm{HD}\rightarrow\mathrm{COMB}}$ is the number of connections received by each combination cell from head direction cells. The scaling factor $\frac{\phi_4}{C^{\mathrm{ROT}\rightarrow\mathrm{COMB}}}$ for the rotation inputs is defined equivalently.

both rats (Taube et al. 1996; Muller et al. 1996) and macaques (Robertson et al. 1999) are known to be approximately Gaussian.) The layer of combination cells is driven by inputs from the head direction cells and rotation cells. The synaptic weights w_{ij}^1, w_{ij}^2, w_{ij}^3 and w_{ij}^4 are updated according to the simple and local learning rules discussed above. During training, the connections to the combination cells self-organise using competitive learning to enable the second layer of cells to learn to represent combinations of head direction and rotation. That is, the combination cells learn to reflect the current head direction of the agent, but are only active if the agent is rotating. At the same time, the model is able to learn to perform path integration by using a trace learning rule to associate recent activity in the layer of combination cells with the current representation in the layer of head direction cells. After training, the network is able to perform path integration in the absence of visual cues. That is, the network can update the representation in the layer of head direction cells using inputs from the layer of combination cells. A key advantage of this two-layer model over the earlier self-organising model of Stringer et al. (2002) is that the incorporation of the layer of combination cells avoids the need for higher order Sigma–Pi synapses to update the representation in the layer of head direction cells during path integration.

Stabilization of the activity packet within the continuous attractor network of head direction cells when the agent is stationary

As described by Stringer et al. (2002a), the recurrent synaptic weights within the continuous attractor network of head direction cells may be corrupted by a certain amount of noise from the learning regime. This problem is compounded by diluted connectivity. This in turn can lead to drift of the activity packet within the continuous attractor network of head direction cells even when the agent is not moving when there is no external visual input available. Stringer et al. (2002a) proposed that in real nervous systems this problem may be solved 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 voltage dependent ion channels in the brain such as NMDA receptors. In the model presented in this paper, we simulate these effects using an additional non-linearity in the activation functions (2) (such as might be implemented by NMDA receptors, see Wang (1999) and Lisman et al. (1998)) by adjusting the sigmoid threshold α_i for each head direction cell i according to

$$\alpha_i = \begin{cases} \alpha^{\mathrm{HIGH}} & \text{if} \quad r_i < \gamma \\ \alpha^{\mathrm{LOW}} & \text{if} \quad r_i \geq \gamma \end{cases} \tag{14}$$

where γ is a firing rate threshold. This helps to stabilize the current position of the activity packet within the continuous attractor network of head direction cells. The sigmoid slopes are set to a constant value, β, for all cells i. The value of α^{HIGH} was 0 and of α^{LOW} was -40 as shown in Table I, and the difference between these values is less than 2% of the typical activation values of HD neurons in the network.

Training and testing

We next present numerical (simulation) results to prove and illustrate the main properties of the model Figure 1 shows the details of the network architecture used for the simulations. In the simulations there were 1000 head direction cells, 1000 combination cells and 1000 rotation cells. In the model simulations the head direction cells were mapped onto a regular grid

Table I. Simulation parameter values for Experiment 1

Network parameters	
No. HD cells	1000
No. COMB cells	1000
No. ROT cells	1000
No. w^1 connections onto each HD cell	50
No. w^2 connections onto each HD cell	50
No. w^3 connections onto each COMB cell	50
No. w^4 connections onto each COMB cell	50
Firing sparsness in COMB layer	0.05
σ^{HD}	20°
Learning rates k^1, k^2, k^3, k^4	0.001
Trace parameter η	0.9
τ	1.0
ϕ_1	3×10^5
ϕ_2	3×10^7
ϕ_3	5×10^2
ϕ_4	3×10^2
HD cell parameters	
γ	0.5
α^{HIGH}	0.0
α^{LOW}	-40
β	0.1
COMB cell parameters	
α	10.0
β	0.3
Training parameters	
No. training epochs	100

of different head directions, where for each head direction cell i there was a unique preferred head direction x_i of the agent for which the cell was stimulated maximally by the visual cues.

During training, the agent is rotated on the spot. One revolution of the agent through 1000 consecutive head directions from 0–360° constitutes 1 epoch of training. In total, 100 training epochs are performed. As the agent rotates to each head direction, the head direction cells are set to fire according to prescribed Gaussian activity profiles. During the initial learning phase it was assumed that the external input e_i to the head direction cells would dominate all other excitatory inputs. Therefore, in the simulations presented below we employed the following modelling simplification. During the learning phase in which the agent received external visual input, rather than implementing the dynamical Equations 1 and 2, we set the firing rates of the head direction cells according to the following Gaussian response profile

$$r_i^{\mathrm{HD}} = e^{-(s_i^{\mathrm{HD}})^2/2(\sigma^{\mathrm{HD}})^2}, \tag{15}$$

where s_i^{HD} was the difference between the actual head direction x of the agent and the preferred head direction x_i for head direction cell i, and σ^{HD} was the standard deviation. For each head direction cell i, s_i^{HD} was given by

$$s_i^{\mathrm{HD}} = \mathrm{MIN}(|x_i - x|, 360 - |x_i - x|) \tag{16}$$

During training, the activity of the combination cells is driven by the w_{ij}^3 synaptic inputs from the head direction cells and the w_{ij}^4 synaptic inputs from the rotation cells. The aim of the training was to check that standard competitive learning can by self-organization set up the correct COMB cells that respond to particular combinations of the head direction and

rotation velocity inputs, and therefore a standard competitive learning instantiation was used as follows. During learning we set the activations of the combination cells according to

$$h_i^{\mathrm{COMB}} = \frac{\phi_3}{C^{\mathrm{HD}\rightarrow\mathrm{COMB}}} \sum_j w_{ij}^3 r_j^{\mathrm{HD}} + \frac{\phi_4}{C^{\mathrm{ROT}\rightarrow\mathrm{COMB}}} \sum_j w_{ij}^4 r_j^{\mathrm{ROT}}. \tag{17}$$

Next, mutual inhibition between the combination cells implements competition to ensure that there is only a small winning set of combination cells left active. The mutual inhibition (which would be implemented by inhibitory feedback neurons which implement lateral inhibition in the brain) ensures that the sparseness of the activity in the combination cell layer is kept to a fixed value, set in these simulations to 0.05. There are many ways in which competition between cells may be modelled. In the simulations, the competition was achieved by setting to 1 the firing rates of the 5% of the combination cells with the greatest activations h_i^{COMB}, with the firing rates of the remaining 95% of the combination cells set to 0.

The numerical simulations began with the learning phase in which the synaptic weights, w_{ij}^1, w_{ij}^2, w_{ij}^3 and w_{ij}^4, were self-organized. At the start of the learning, the synaptic weights were initialized to random positive values. Then learning proceeded with the agent rotating on the spot. During this, the synaptic weights were updated as described above according to the learning rules (Equations 3, 5, 10, and 12). At the start of training all trace values were initialised to zero.

After the learning phase was completed, the simulations continued with the testing phase, in which the agent had to perform path integration in the absence of visual input. For the testing phase, the full 'leaky integrator' dynamical Equations (1, 2, 8 and 9), were implemented, because in the model the path integration occurs as an evolutionary process that is continuous in time. In the simulations, the differential Equations 1 and 8 were approximated by Forward Euler finite difference schemes with a timestep of 0.2 (which was shown to be sufficiently short to model the dynamics). At the start of the testing phase, all of the firing rates, r_i^{HD}, r_i^{COMB}, and r_i^{ROT}, were set to zero. Then the agent was simulated with visual input available, with the agent in an initial state $x = 72^\circ$, but with no rotation cells active, for 500 timesteps. While the agent maintained this head direction, the visual input term e_i for each head direction cell i was set to a Gaussian response profile identical (except for scaling) to that used for the head direction cell during the learning phase given by Equation 15. Next the visual input was removed by setting all of the e_i terms to zero, and the agent was allowed to rest in the same head direction $x = 72^\circ$ for a further 500 timesteps. This process led to a stable packet of activity within the layer of head direction cells that represented the initial head direction of the agent.

When the agent started to rotate, the firing rates r_i^{ROT} of the rotation cells were set to 1. When the rotation cells fired, the co-firing of the rotation cells and the head direction cells stimulated an activity packet in the layer of combination cells. Then the firing of the combination cells stimulated further head direction cells in the appropriate direction in the head direction cell layer to reflect the altering head direction of the agent given that the agent was rotating. That is, the head direction representation was updated by the signal from the combination cells. As the activity packet within the head direction network moved, the activity packet in the combination network, which was stimulated by the co-firing of the head direction cells and rotation cells, also moved. Thus, the two networks moved in tandem while the rotation cells were active. In this way the network learned to perform path integration. When the rotation cells stopped firing, the driving input to the combination cells disappeared, and the combination cells ceased to fire. Then the activity packet within the

head direction network was no longer updated by the signals from the combination cells, and the head direction representation remained static.

Simulations

Experiment 1: Diluted connectivity for all types of synaptic connection w^1, w^2, w^3 and w^4

In Experiment 1, the 2-layer neural network architecture shown in Figure 1 was simulated with diluted 5% connectivity for all of the four different types of synaptic connection w^l, w^2, w^3 and w^4. Specifically, each head direction (HD) cell received 50 w^1 connections and 50 w^2 connections, while each combination (COMB) cell received 50 w^3 connections and 50 w^4 connections. The diluted connectivity implemented in this simulation (and present in the brain) helps the competitive network to allocate different neurons to different combinations of its inputs (Rolls & Treves 1998; Rolls & Deco 2002). Simulation parameter values for Experiment 1 are given in Table I. Results after training are shown in Figures 2 to 5.

Figure 2 shows the recurrent weights w^1 within the layer of head direction cells after training with the Hebb rule (Equation 3) and weight normalization (Equation 4). Each of the four plots shows the learned recurrent synaptic weights to a different postsynaptic head direction cell from the other 1000 presynaptic head direction cells in the network. In the plots, the 1000 presynaptic head direction cells are arranged in the graphs according to where they fire maximally in the head direction space of the agent when visual cues are available. For each plot, a dashed vertical line is drawn to mark where the postsynaptic head direction cell fires maximally in the light. In each of the four plots, except for the effects of diluted connectivity, the weight profile is approximately symmetric about the head direction at which the postsynaptic head direction cell fires maximally in the light. This will help to support an activity packet within the head direction network during subsequent testing in the absence of visual input.

Figure 3 shows the weights w^3 from the layer of head direction cells to the layer of combination cells after competitive learning with the Hebb rule (Equation 10) and weight normalization (Equation 11). Each of the four plots shows the learned w^3 synaptic weights to a different postsynaptic combination cell from the 1000 presynaptic head direction cells. The 1000 presynaptic head direction cells are arranged in the graphs according to where they fire maximally in the head direction space of the agent when visual cues are available. For each plot, a dashed vertical line is drawn to mark the presynaptic head direction cell from which the postsynaptic combination cell has maximal w^3 weight. Except for the effects of diluted connectivity, each of the w^3 weight profiles is centred on a small cluster of similarly-tuned head direction cells, with a profile that is approximately symmetric about the presynaptic head direction cell from which the postsynaptic combination cell has maximal w^3 weight. Thus, the learned w^3 weights show that the cells in the second layer learn to receive maximal stimulation from particular head directions. Given that the model parameters ϕ_3 and ϕ_4 are tuned to ensure that a strong rotation input through the w^4 synapses is also needed to fire the cells in the second layer, these cells in fact learn to respond to combinations of a particular head direction and rotation velocity.

Figure 4 shows the weights w^2 from the layer of combination cells to the layer of head direction cells after learning with the trace rule (Equation 5) and weight normalization (Equation 7). Each of the four plots shows the learned w^2 synaptic weights from a different presynaptic combination cell to the 1000 postsynaptic head direction cells. The 1000 postsynaptic head direction cells are arranged in the graphs according to where they fire maximally in the head direction space of the agent when visual cues are available. For each plot, a dashed

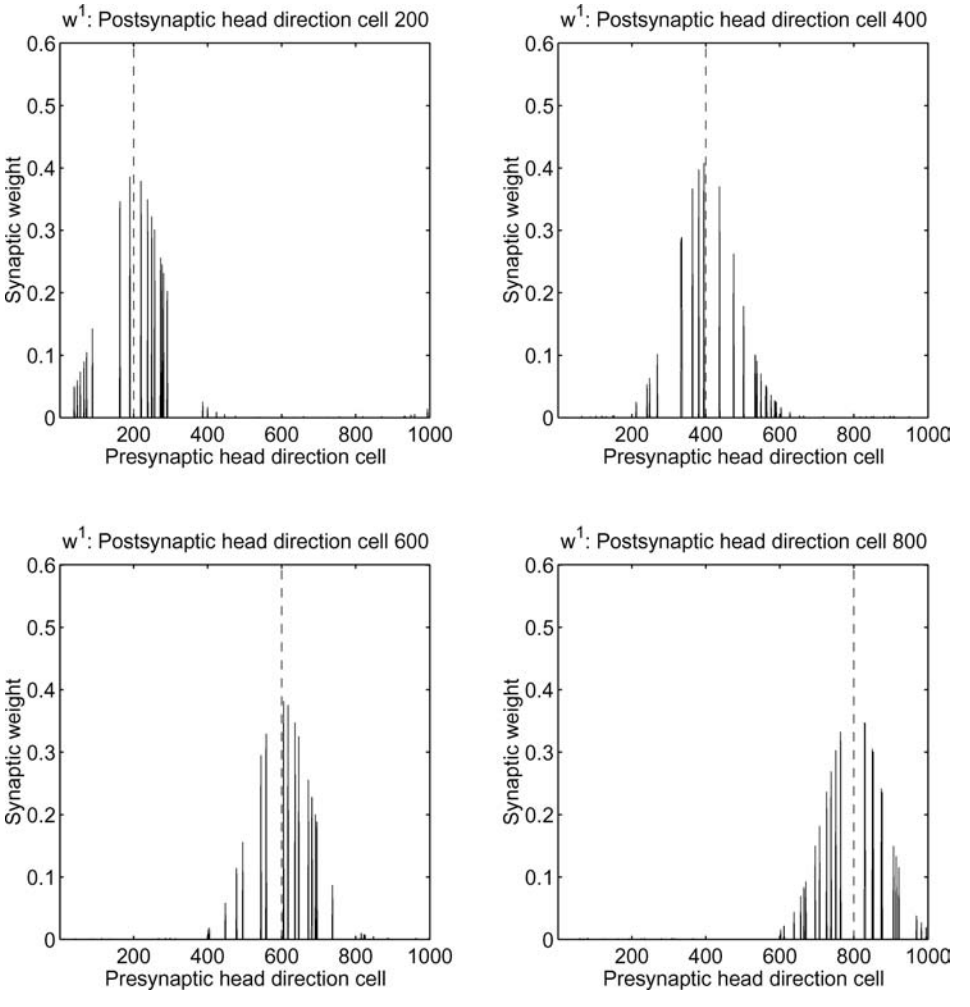

Figure 2. The recurrent weights w^1 within the layer of head direction cells after training with the Hebb rule (Equation 3) and weight normalization (Equation 4). In these simulations, the w^1 weights are trained with diluted connectivity. Each of the four plots shows the learned recurrent synaptic weights to a different postsynaptic head direction cell from the other 1000 presynaptic head direction cells in the network. In the plots, the 1000 presynaptic head direction cells are arranged in the graphs according to where they fire maximally in the head direction space of the agent when visual cues are available. For each plot, a dashed vertical line is drawn to mark where the postsynaptic head direction cell fires maximally in the light. In each of the four plots, except for the effects of diluted connectivity, the weight profile is approximately symmetric about the head direction at which the postsynaptic head direction cell fires maximally in the light. This will help to support an activity packet within the head direction network during subsequent testing in the absence of visual input.

vertical line is drawn to mark the postsynaptic head direction cell from which the presynaptic combination cell has maximal w^3 weight as shown in Figure 3. Thus, the dashed vertical line shows the head direction to which the presynaptic combination cell is tuned. Even with diluted connectivity, it is still evident that the w^2 weight profiles are *asymmetric* about the head direction to which the presynaptic combination cell is tuned. In each plot there is an asymmetric bias to the right, i.e., towards increasing head direction $\theta \in [0, 360]$. This asymmetry is necessary for the network to perform path integration in the absence of visual input.

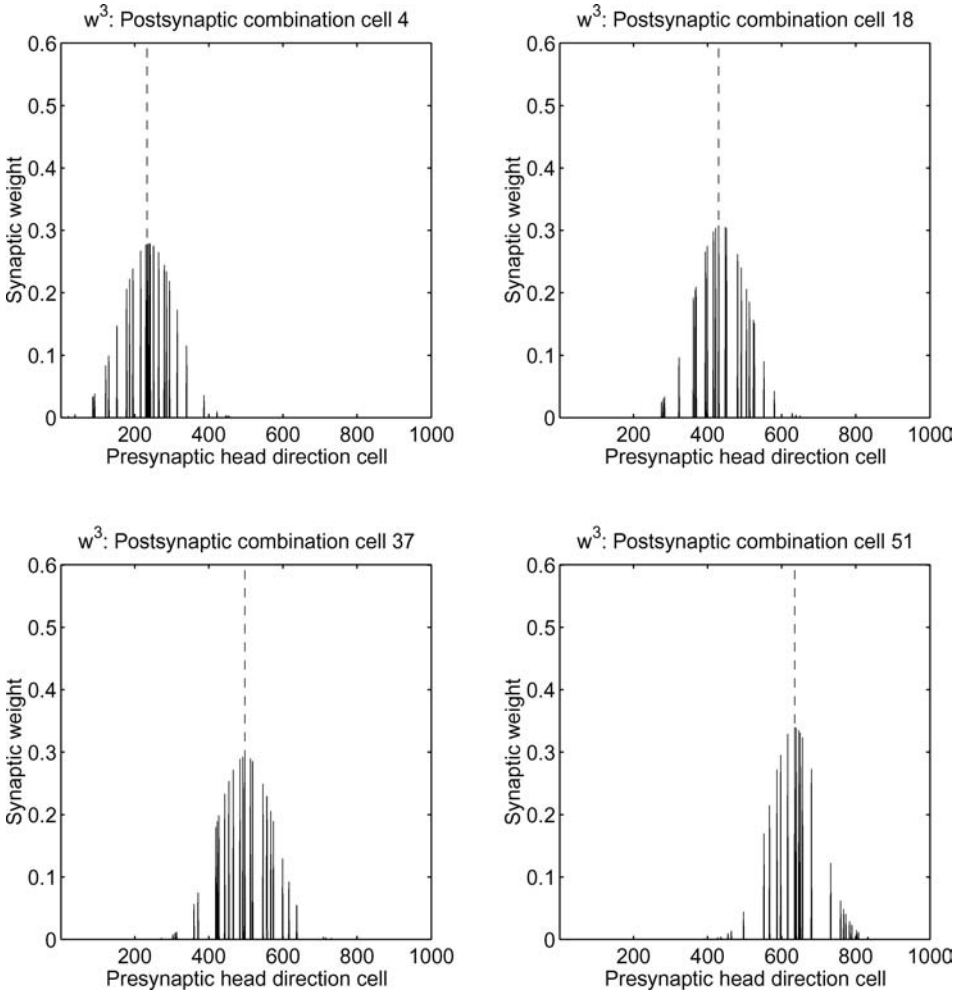

Figure 3. The weights w^3 from the layer of head direction cells to the layer of combination cells after competitive learning with the Hebb rule (Equation 10) and weight normalization (Equation 11). In these simulations, the w^3 weights are trained with diluted connectivity. Each of the four plots shows the learned w^3 synaptic weights to a different postsynaptic combination cell from the 1000 presynaptic head direction cells. The 1000 presynaptic head direction cells are arranged in the graphs according to where they fire maximally in the head direction space of the agent when visual cues are available. For each plot, a dashed vertical line is drawn to mark the presynaptic head direction cell from which the postsynaptic combination cell has maximal w^3 weight. Except for the effects of diluted connectivity, each of the w^3 weight profiles is centred on a small cluster of similarly-tuned head direction cells, with a profile that is approximately symmetric about the presynaptic head direction cell from which the postsynaptic combination cell has maximal w^3 weight. Thus, the learned w^3 weights show that the cells in the second layer learn to receive maximal stimulation from particular head directions. Given that the model parameters ϕ_3 and ϕ_4 are tuned to ensure that a strong rotation input through the w^4 synapses is also needed to fire the cells in the second layer, these cells in fact learn to respond to combinations of a particular head direction and rotation velocity.

When the combination cells are activated by a signal from the rotation cells, the asymmetry in the w^2 weights will drive the activity packet around the continuous attractor network of head direction cells in the direction of the asymmetric weight bias set up through the trace learning rule (Equation 5).

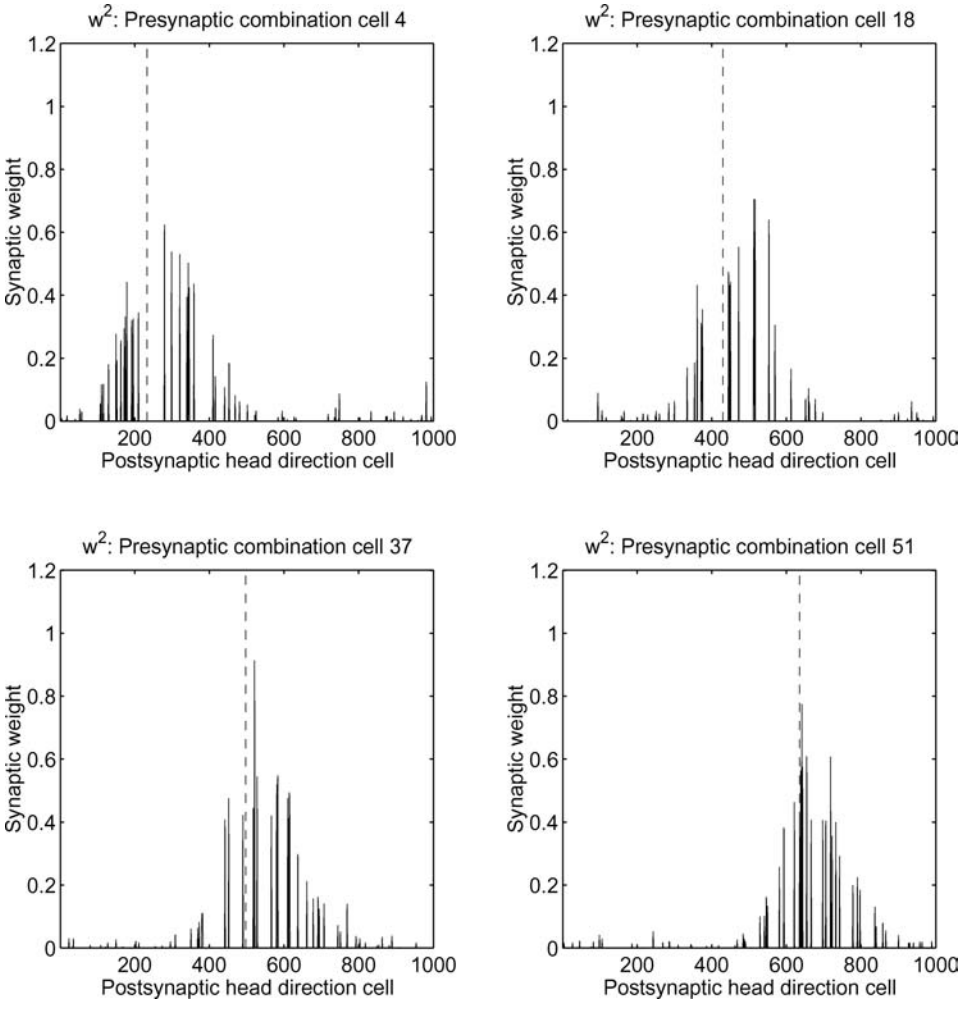

Figure 4. The weights w^2 from the layer of combination cells to the layer of head direction cells after learning with the trace rule (Equation 5) and weight normalization (Equation 7). In these simulations, the w^2 weights are trained with diluted connectivity. Each of the four plots shows the learned w^2 synaptic weights from a different presynaptic combination cell to the 1000 postsynaptic head direction cells. The 1000 postsynaptic head direction cells are arranged in the graphs according to where they fire maximally in the head direction space of the agent when visual cues are available. For each plot, a dashed vertical line is drawn to mark the postsynaptic head direction cell from which the presynaptic combination cell has maximal w^3 weight as shown in Figure 3. Thus, the dashed vertical line shows the head direction to which the presynaptic combination cell is tuned. Even with diluted connectivity, it is still evident that the w^2 weight profiles are *asymmetric* about the head direction to which the presynaptic combination cell is tuned. In each plot there is an asymmetric bias to the right, i.e., towards increasing head direction $\theta \in [0, 360]$. This asymmetry is necessary for the network to perform path integration in the absence of visual input. When the combination cells are activated by a signal from the rotation cells, the asymmetry in the w^2 weights will drive the activity packet around the continuous attractor network of head direction cells in the direction of the asymmetric weight bias set up through the trace learning rule (Equation 5).

Figure 5 shows firing rates during testing in the layer of 1000 rotation cells (top left), the layer of 1000 head direction cells (top right), the layer of 1000 combination cells (bottom left), and the first 20 combination cells (bottom right). Regions of high firing are represented by dark shading. During timesteps 1–200, there was a stable packet of activity in the layer

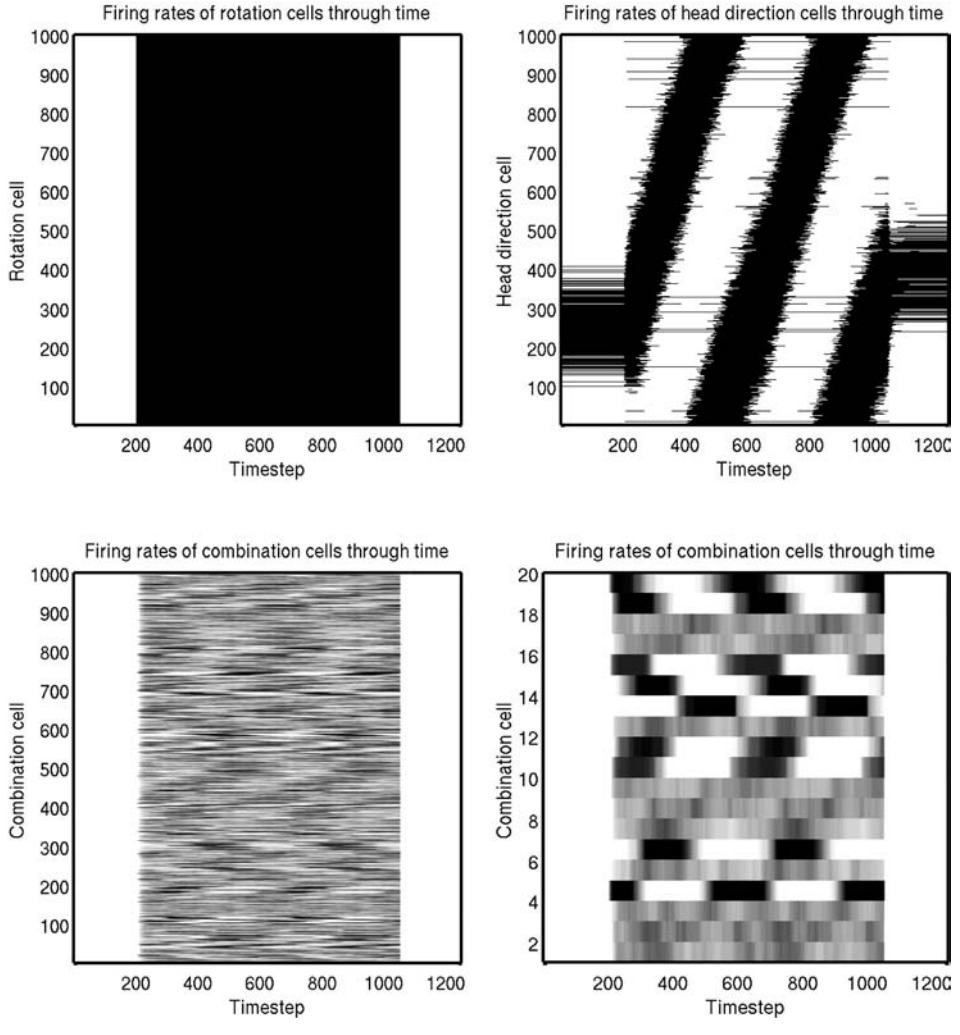

Figure 5. Firing rates during testing in the layer of 1000 rotation cells (top left), the layer of 1000 head direction cells (top right), the layer of 1000 combination cells (bottom left), and the first 20 combination cells (bottom right). Regions of high firing are represented by dark shading. During timesteps 1–200, there was a stable packet of activity in the layer of head direction cells, which was supported by the w^1 recurrent connections. During this period, all of the rotation cells were quiescent and there was no significant activity within the layer of combination cells. During timesteps 201–1050, the rotation cells became active, which in turn drove activity within the layer of combination cells. The combination cells received inputs from the head direction cells through the w^3_{ij} synapses, and inputs from the rotation cells through the w^4_{ij} synapses. The combination cells had learned to respond to particular combinations of head direction and rotational velocity. In particular, the combination cells only became significantly active during the rotation period. Due to the asymmetry in the w^2 connections from the combination cells to the head direction cells, the activity in the layer of combination cells drove the activity packet through the layer of head direction cells. The network thus performed path integration. That is, the representation of head direction within the layer of head direction cells was updated using the rotation signal in the absence of external visual cues. From the bottom right plot, it can be seen that individual combination cells switch on and off according to the current head direction of the agent (as represented by the activity packet within the layer of head direction cells) as the agent rotates through two 360° cycles. From timestep 1051, the rotation cells stopped firing, which caused the combination cells to become quiescent and the activity packet within the layer of head direction cells remained stationary once more.

of head direction cells, which was supported by the w^1 recurrent connections. During this period, all of the rotation cells were quiescent and there was no significant activity within the layer of combination cells. From timestep 201–1050, the rotation cells became active, which in turn drove activity within the layer of combination cells. The combination cells received inputs from the head direction cells through the w^3_{ij} synapses, and inputs from the rotation cells through the w^4_{ij} synapses. The combination cells had learned to respond to particular combinations of head direction and rotational velocity. In particular, the combination cells only became significantly active during the rotation period. Due to the asymmetry in the w^2 connections from the combination cells to the head direction cells, the activity in the layer of combination cells drove the activity packet through the layer of head direction cells. The network thus performed path integration. That is, the representation of head direction within the layer of head direction cells was updated using the rotation signal in the absence of external visual cues. From the bottom right plot, it can be seen that individual combination cells switch on and off according to the current head direction of the agent (as represented by the activity packet within the layer of head direction cells) as the agent rotates through a couple of 360° cycles. From timestep 1051 the rotation cells stopped firing, which caused the combination cells to become quiescent and the activity packet within the layer of head direction cells remained stationary once more.

We now consider how well this path integration system would track continuously varying movements, as might occur in a natural environment. We addressed this by performing simulations in which the head rotation velocity signal varied continuously, to determine whether the path integration would follow this accurately. The left part of Figure 6 shows a linearly increasing rotation speed (shown by the darkness of the rotation cell firing in the top left panel) as a function of time, and that this produces a linear increase in the speed of movement of the activity packet in the head direction network (shown in the lower left panel). The head direction position represented as a result of the path integration thus accurately takes account of the changing velocities. A similar tracking by the path integration process of a linearly decreasing head rotation velocity driving input is shown on the right. In the next paragraph we address quantitatively the relation between the speed of the head direction packet that is the result of the path integration, and the head rotation velocity driving signal. In the Discussion, we consider the effects of the time constants in the system.

An important property of a path integration system would be the ability to compute the correct distance when the speed of the (rotation) inputs is altered. We therefore tested whether after training on even one rotation velocity, the network would generalize correctly to other velocities of the rotation inputs. Figure 7 shows how path integration in the network is able to generalize across different rotation speeds, after training at only one speed. The model assumes that the firing rates of the rotation cells increase approximately linearly with the speed of rotation of the agent. In these simulations, the network was trained with only one speed of rotation, at which the firing rates of the rotation velocity cells (ROT in Figure 1) were set to a value of 1.0. Then, during subsequent testing, the agent was rotated over a range of slower speeds with the firing rates of the rotation velocity cells set within the interval 0.0 to 1.0. The plot shows how the speed of the activity packet within the layer of head direction cells varies with the firing rates of the rotation cells after training. The speed of path integration depends approximately linearly on the firing rates of the rotation cells. If the firing rates of the rotation cells vary linearly with the speed of rotation of the agent, then the speed of path integration in the network increases in correct proportion to the increasing speed of the agent. The path integration network is thus able to generalize successfully across different speeds of rotation.

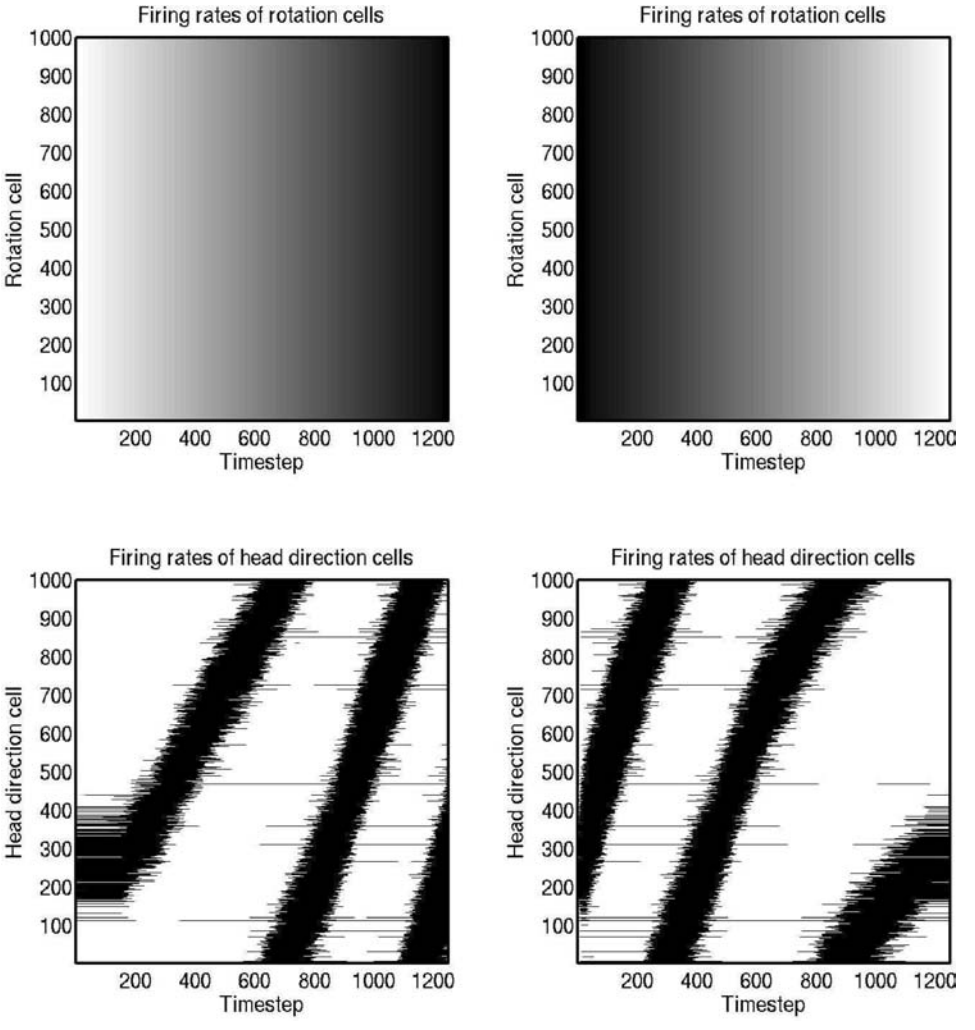

Figure 6. Path integration which tracks the changing speed of the rotation signal. On the left a linearly increasing rotation speed (shown by the darkness of the rotation cell firing in the top left panel) produces a linear increase in the speed of movement of the activity packet in the head direction network (shown in the lower left panel). The lower left panel shows the firing rates of the 1000 head direction cells (by the darkness of the plot) as a function of time. A similar tracking by the path integration process of a linearly decreasing head rotation velocity driving input is shown on the right.

Experiment 2: Full connectivity for connections w^3

In Experiment 2, we show the effects of increasing the level of connectivity for the w^3 connections. In this simulation, the number of w^3 connections received by each combination cell is increased from 50 to 1000. The neural network architecture shown in Figure 1 was simulated with full w^3 synaptic connectivity, while maintaining diluted 5% connectivity for the three other types of connection w^1, w^2 and w^4. Specifically, each head direction (HD) cell received 50 w^1 connections and 50 w^2 connections, while each combination (COMB) cell received 1000 w^3 connections and 50 w^4 connections. The remaining simulation parameter values were as for Experiment 1.

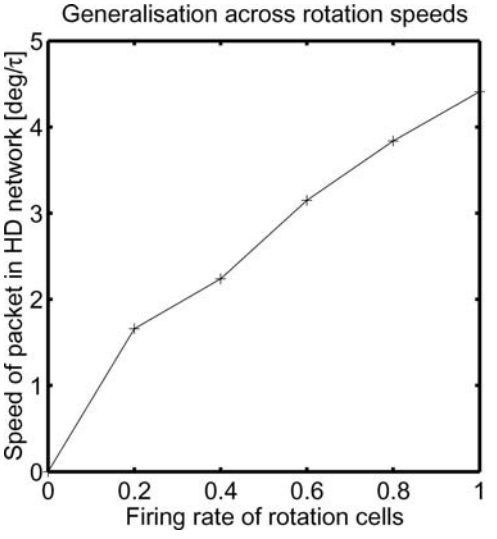

Figure 7. Generalisation across different rotation speeds after training with only a single speed. The plot shows how the speed of the activity packet within the layer of head direction cells varies with the firing rates of the rotation cells during path integration. The speed of path integration depends approximately linearly on the firing rates of the rotation cells. If the firing rates of the rotation cells vary linearly with the speed of rotation of the agent, then the speed of path integration in the network increases in correct proportion to the increasing speed of the agent. The path integration network is thus able to generalise successfully across different speeds of rotation.

Results after training are shown in Figure 8. Each of the four plots shows the learned w^3 synaptic weights to a postsynaptic combination cell from the 1000 head direction cells in the network. Full w^3 connectivity destroys competitive learning in the layer of combination cells, which was responsible for these cells learning to respond to particular combinations of head direction and rotation in the diluted network of Experiment 1. Instead, a quite different style of self-organisation occurs in the layer of combination cells, called *continuous transformation* (CT) learning. Continuous transformation (CT) learning has been proposed as a way in which transform invariant representations may develop in the ventral visual system (Stringer et al. 2006). The learning mechanism is quite general, and may operate in brain areas with the architecture of a competitive network that is trained with continuous patterns.

The continuous transformation learning process operates as follows, and is illustrated in Figure 9 in the context of the HD–COMB network. When the agent is situated in a first head direction, a small group of HD cells is activated. Activity is then propagated through feedforward synaptic connections w^3, from the HD cells to the layer of COMB cells. Next, a small winning set of neurons in the COMB layer will modify (through associative learning) their afferent connections from the HD layer to respond well to the HD cells which are active in that head direction. Because there is spatial continuity as the agent rotates, many of the same HD cells are still active when the agent is in a nearby head direction. These active neurons with their strengthened synapses ensure that the same neurons in the COMB layer will be activated because some of the active afferents are the same as when the agent was in the first head direction. The key point is that if these afferent connections have been strengthened sufficiently while the agent is in the first head direction, then these connections will be able to continue to activate the same neurons in the COMB layer when the agent is situated in overlapping nearby head directions. Then any newly active synapses because of

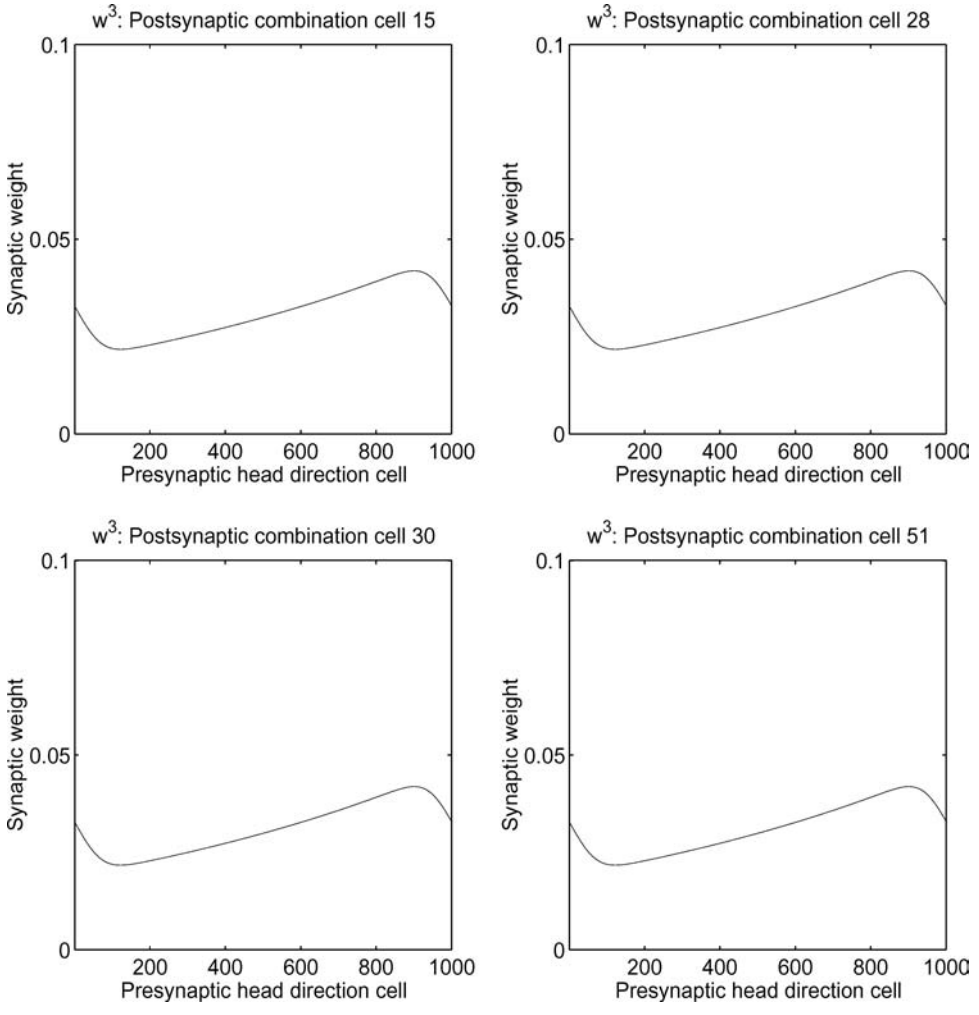

Figure 8. Results of training with full w^3 synaptic connectivity. In this simulation, the number of w^3 connections received by each combination cell is increased from 50 to 1000. That is, each combination cell receives w^3 inputs from all 1000 head direction cells. Conventions as above in Figure 3. Each of the four plots shows the learned w^3 synaptic weights to a postsynaptic combination cell from the 1000 head direction cells in the network. Full w^3 connectivity destroys competitive learning in the layer of combination cells, which was responsible for these cells learning to respond to particular combinations of head direction and rotation in the diluted network. Instead, continuous transformation (CT) learning occurs. This produces cells in the second layer that have strong weights from every head direction, while some other cells remain untrained. The peak in the w^3 weight profiles is merely due to the continuous weight normalisation process, which causes the most recently updated weights to be stronger than weights updated slightly earlier. The lack of head direction specificity in the second layer of cells will cause path integration to fail in the overall network. Thus, diluted connectivity can help to eliminate CT learning, and the elimination of CT learning is a requirement in the w^3 connections for the 2-layer path integration system to self-organise successfully.

the change of head direction of the agent have conjunctive pre- and post-synaptic activity, and these synapses are increased in strength by Hebbian associativity. As can be seen in Figure 9, the process can be continued for subsequent rotations, provided that a sufficient proportion of HD cells stay active between individual changes of head direction. This process

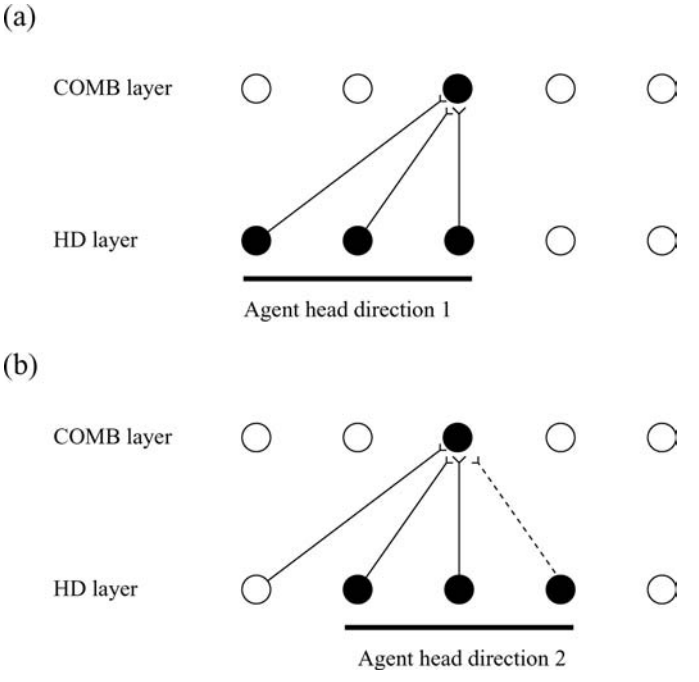

Figure 9. An idealised illustration of how CT learning would function in a network with a single layer of forward synaptic connections w^3 between a layer of HD cells and a layer of COMB neurons. Initially the forward synaptic weights w^3 are set to random values. The top part (a) shows the initial HD pattern with the agent in head direction 1. Activation from the (shaded) active input cells is transmitted through the initially random forward connections to stimulate the cells in the COMB layer. The shaded cell in the COMB layer wins the competition in that layer. The weights from the active HD cells to the active COMB neuron are then strengthened using an associative learning rule. The bottom part (b) shows what happens after the agent has rotated by a small amount to a new partially overlapping head direction 2. As some of the active HD cells are the same as those that were active when the agent was in head direction 1, the same COMB cell is driven by these previously strengthened afferents to win the competition again. The rightmost shaded HD cell activated by the agent being in head direction 2, which was inactive when the agent was in head direction 1, now has its connection to the active COMB cell strengthened (denoted by the dashed line). Thus the same COMB neuron has learned to respond to the two HD patterns that have similar vector elements in common. As can be seen, the process can be continued for subsequent rotations, provided that a sufficient proportion of HD cells stay active between individual rotations.

can lead individual COMB cells to develop response fields over large continuous regions of the head direction space of the agent.

It can be seen in Figure 8 that continuous transformation learning produces cells in the second layer that have strong weights from every head direction, while some other cells remain untrained. Thus, the COMB cells no longer learn to respond to particular head directions, but instead respond significantly to all head directions. The peak which can be seen in each of the w^3 weight profiles shown in Figure 8 is merely due to the continuous weight normalisation process, which causes the most recently updated weights to be a little stronger than weights updated slightly earlier. The lack of head direction specificity in the layer of COMB cells will cause path integration to fail in the overall network. Thus, diluted connectivity can help to eliminate CT learning, and the elimination of CT learning is a requirement in the w^3 connections for the 2-layer path integration system to self-organise successfully.

Experiment 3: Clockwise and anticlockwise rotation

To demonstrate that the network can learn to perform appropriate path integration for different directions of input, in Experiment 3 the 2-layer neural network architecture shown in Figure 1 was simulated with both clockwise and anticlockwise rotation. A crucial requirement tested in this experiment is that different combination cells (COMB) would self-organize so that some encoded a combination of rotation in one direction with a particular head direction, and other cells encoded a combination of rotation in the opposite direction with a particular head direction. Having these different types of combination cells would be crucial for the network to path integrate in different directions given different directions of the velocity input. In this simulation, the model was similar to that used for experiment 1, except that 500 of the rotation cells responded to clockwise rotation, while the remaining 500 rotation cells responded to anticlockwise rotation. During training, the agent was simulated rotating in both clockwise and anticlockwise directions.

Simulation results with clockwise and anticlockwise rotation are shown in Figure 10. The figure shows the firing rates during testing in the layer of 1000 rotation cells (top left), the layer of 1000 head direction cells (top right), the layer of 1000 combination cells (bottom left), and the first 20 combination cells (bottom right). Regions of high firing are represented by dark shading. During timesteps 1–200, there was a stable packet of activity in the layer of head direction cells. During timesteps 201–800, the clockwise rotation cells became active, which in turn drove the activity of those combination cells which had learned to respond to a combination of clockwise rotation and a particular head direction. This activity in the layer of combination cells drove the activity packet through the layer of head direction cells in the clockwise direction. During timesteps 801–1000, no rotation cells were firing, and the packet of activity in the layer of head direction cells was stationary. During timesteps 1001–1600, the anticlockwise rotation cells became active, which in turn drove the activation of those combination cells which had learned to respond to a combination of anticlockwise rotation and a particular head direction. This activity in the layer of combination cells drove the activity packet through the layer of head direction cells in the anticlockwise direction. During timesteps 1601–1800, no rotation cells were firing, and the packet of activity in the layer of head direction cells was again stationary.

From the bottom right plot of Figure 10, it can be seen that individual combination cells switch on and off according to a particular head direction, and whether the agent is rotating in either the clockwise or anticlockwise direction. In particular, combination cells with high firing rate peaks (i.e., dark regions representing high firing rates) either fire during clockwise rotation or fire during anticlockwise rotation, but do not fire during both. These combination cells have thus learned to be selective for the direction of rotation, as well as for particular head directions.

It can be seen from Figure 10 that the performance of the self-organised network model is not as smooth as the simulation results from Experiment 1. With the network having to learn two directions of motion, the weight structure within the network is more noisy and path integration has become less regular. However, these effects may be largely due to the small size of the network used in our simulations, which has only 1000 neurons in each of the layers and 50 connections from any one source. This is much smaller than the number of neurons present in the relevant areas in the brain, and the number of connections per neuron. Scaling up the network to more biologically realistic numbers of neurons should therefore offer smoother performance. In particular, large numbers of combination neurons are needed, so that many different combinations of head direction and rotation speed can be represented. This hypothesis was supported by simulations in which the COMB network

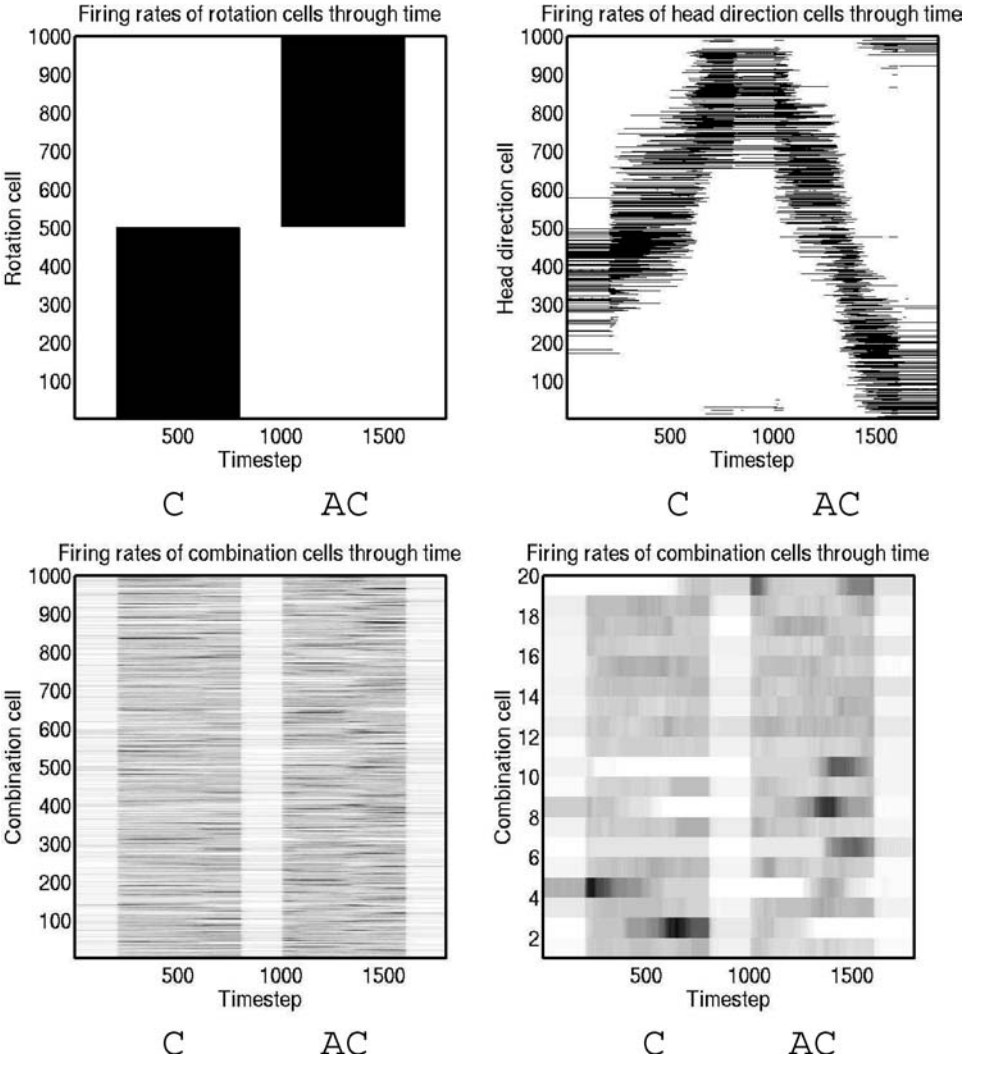

Figure 10. Simulation with clockwise (C) and anticlockwise (AC) rotation. Firing rates during testing in the layer of 1000 rotation cells (top left), the layer of 1000 head direction cells (top right), the layer of 1000 combination cells (bottom left), and the first 20 combination cells (bottom right). Regions of high firing are represented by dark shading. During timesteps 1–200, there was a stable packet of activity in the layer of head direction cells. During timesteps 201–800, the clockwise rotation cells became active, which in turn drove the activation of those combination cells which had learned to respond to a combination of clockwise rotation and a particular head direction. This activity in the layer of combination cells drove the activity packet through the layer of head direction cells in the clockwise direction. During timesteps 801–1000, no rotation cells were firing, and the packet of activity in the layer of head direction cells was stationary. During timesteps 1001–1600, the anticlockwise rotation cells became active, which in turn drove the activation of those combination cells which had learned to respond to a combination of anticlockwise rotation and a particular head direction. This activity in the layer of combination cells drove the activity packet through the layer of head direction cells in the anticlockwise direction. During timesteps 1601–1800, no rotation cells were firing, and the packet of activity in the layer of head direction cells was again stationary. From the bottom right plot, it can be seen that individual combination cells switch on and off according to whether the agent is rotating in either the clockwise or anticlockwise direction, and according to the current head direction of the agent (as represented by the activity packet within the layer of head direction cells).

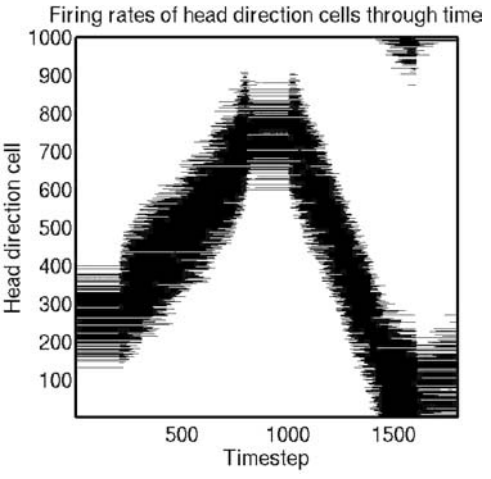

Figure 11. The effect of increasing the size of the combination cell (COMB) network (to 5000) for Experiment 3. The simulation is otherwise for the same conditions as shown in Figure 10. The simulation is with clockwise (C) and anticlockwise (AC) rotation separated by a period of zero velocity input, and the firing rates during testing for the 1000 head direction cells are shown. These results should be compared with the top right part of Figure 10.

size for the simulations illustrated in Figure 10 was increased by a factor of 5, which was found to make the activity packet in the head direction cell network move more smoothly when the path integration was being performed. This is illustrated in Figure 11 which should be compared with figure 10 (top right). For the simulations shown in Figure 11 the part of the network that was increased in size was the number of neurons in the combination cell network (while maintaining the same 5% level of connectivity), indicating that the number of combination cells must be sufficient to cover the head direction x rotation velocity space smoothly.

Discussion

A key advantage of this two-layer model over earlier self-organising models (Stringer et al. 2002a) is that the incorporation of the second layer of cells representing combinations of head direction and rotation avoids the need for multiplicative, higher order Sigma–Pi synapses. This helps to make the new model described here biologically plausible, as the learning uses first order synapses and simple associative local learning rules.

The new model described here also is consistent with and helps to account for the presence of neurons in the rat presubiculum and dorsal tegmental nucleus that respond to combinations of head direction and angular rotation (Taube et al. 1990b; Sharp 1996; Bassett & Taube 2005).

The model described here was intended to illustrate the principle by which competitive learning could be used in one network to build the position and velocity combination neurons that provide the useful signal source to be used as an input to drive the continuous attractor to represent a new position. The model was therefore kept simple. We showed that with different rotation inputs, for example clockwise and counterclockwise, the network can be trained to move in either direction, with this depending on the competitive self-organisation in the combination cell layer of neurons that respond to combinations of current head direction and rotation in either one or the other direction. We also showed that in this type of network,

the position represented in the continuous attractor can be driven at different speeds by different magnitudes of the rotation inputs, even when the network has been trained on only one velocity (Stringer et al. 2002a, 2003). Moreover, the path integration performed by the network linearly reflects the rotation cell input, as shown in Figure 7. Further, the system can arbitrarily track continuous variations in the rotation cell input, as shown in Figure 6. The dynamics of the system do not lead to significant delays or overshoots on the time scale of Figure 5 as shown in that Figure, but it is useful to consider the dynamics further in relation to how they might operate in a biologically implemented network. In the model, there is a single time constant τ, which determines how long it takes a layer of the network to respond to its input. What would be the order of τ in the brain? If we consider the combination cell layer in Figure 1, we see that as a competitive network, the dominant factor in determining how soon the output neurons fire is the time constant of the synapses (Rolls & Deco 2002). Indeed, a single layer network of this type including the feedback competition implemented by the GABA neurons takes in the order of the time constant of the synapses to settle (Panzeri et al. 2001), which will be in the order of 10 ms for AMPA synapses. The situation in the head direction cell network is similar. The combination cell input, working through the asymmetrical synapses illustrated in Figure 4, biases up the firing rates of the continuous attractor on one side of the current activity packet in a direction towards which the attractor should move. This process is dynamically that described for a one-layer network, and the time it takes for the HD cells to respond to the input that causes the packet of activity to move will be of order τ. Thus the delay dynamics of the system shown in Figure 1 will not be a great deal longer than τ, and in terms of brain time, the path integration system might take 20–30 ms to respond. However the overall time over which the integration occurred, and thus the distance covered, would not be affected by the time constants as the delays would be similar for starting and stopping. In practice, the small delays in the time it takes the integration process to start and stop would be small in relation to the time it takes an animal to rotate.

The new model is also prototypical of path integration systems that may be implemented in the brain, including those used for path integration in rat hippocampal place cells (McNaughton et al. 1983; O'Keefe 1984; Muller et al. 1991; Markus et al. 1994) and primate hippocampal spatial view cells (Robertson et al. 1998; Rolls & Xiang 2006; Rolls & Kesner 2006). For these systems to operate in the way we now propose, there would need to be in these other systems neurons that respond to combinations of for example place and running velocity and head direction. It has been shown that hippocampal place cells can have their activity modulated by running speed and direction (McNaughton et al. 1983), and can even be tuned to particular velocities (Wiener 1989). The current theory provides an account of the presence of such cells in the hippocampus. The origin of these signals, and indeed a possible brain region for path integration for place to be implemented, is the dorsocaudal medial entorhinal cortex, which contains spatial grid cells that are modulated by head direction (Hafting et al. 2005; Sargolini et al. 2006).

There are a number of ways in which a temporal trace of preceding neuronal activity in a presynaptic term could be implemented in the brain. One is that the binding period of glutamate in the NMDA channels, which may last for 100 or more ms, may implement a trace rule by producing a narrow time window over which the *average* activity at each presynaptic site affects learning (Rolls 1992; Rhodes 1992; Rolls & Kesner 2006).

The possibility that head direction cells and path integration are performed via the excitatory connections from the lateral mammillary nucleus (LMN) to the dorsal tegmental

nucleus (DTN) which in turn inhibits the lateral mammillary nucleus, and without using recurrent excitatory connections, has been modelled by Song and Wang (2005) and Boucheny et al. (2005), but in contrast to the present, no self-organizing learning process for the synaptic weights needed in the model has been described. The model described here does not need recurrent excitatory connections in the combination cell (COMB) network, and this would fit with architectures such as those of the DTN and LMN. Moreover, if the angular head velocity (ROT or AHV) input is zero, the firing of the combination cells in the network described here does decrease, as is the case of for example DTN neurons (Bassett & Taube 2005). The network described here does need recurrent connections in part of the system, the HD network, to perform the functions of maintaining some head direction cell firing when the animal is not moving, and of performing the path integration, and this part of the system could be in a brain region other than the DTN and LMN but which is connected with these brain areas. However, the network described here has a primary aim of showing how path integration systems could self-organize without using Sigma–Pi synapses, by including a network with neurons that by competitive learning build combination neurons for velocity and position in the state space. The networks we have in mind where such competitive learning could easily be introduced to replace Sigma–Pi connections include networks involved in the path integration of the place where an animal is located (Stringer et al. 2002a), and of spatial view (Stringer et al. 2005).

A new concept introduced in this paper is that a competitive network can be used to help a path integration system, by allowing neurons by self-organising learning to respond to combinations of head direction and head rotation velocity. The use of these combination cells allows the head direction cell path integration network to use only first order synapses, that is synapses with one presynaptic terminal acting at each post-synaptic site, rather than a combination of two presynaptic terminals interacting multiplicatively with a post-synaptic site, which is a second order, Sigma–Pi, synapse. Some of the implications for neuroscience are as follows. First, it is predicted that neurons will be found that respond to combinations of head direction and head rotation velocity. Such neurons are present in the nervous system, for example in the presubiculum and dorsal tegmental nucleus (Taube et al. 1990b; Sharp 1996; Bassett & Taube 2005). A point made in this paper is that these neurons would be useful as part of a self-organizing path integration system. Second, it is predicted that first order synapses may be sufficient in a path integration system, with no need for Sigma–Pi synapses. Third, it is predicted that there will be associatively modifiable connections from the Head Direction to the combination cells. Fourth, there should be competitive interaction between the combination cells, implemented for example through a population of inhibitory neurons. Fifth, there should be diluted (sparse) connectivity from the HD cells to the combination cells, to avoid the Continuous Transformation learning effect shown in Figure 8. Sixth, there should be a short term memory trace in the associative synaptic modification implemented at the combination to HD cell synapses w^2. This could be implemented by a process as simple as the long unbinding time constant of glutamate at NMDA synapses (Spruston et al. 1995; Hestrin et al. 1990), or it could be implemented by a transmission delay in the connections from the combination cells to the w^2 synapses in the order of 10–20 ms.

Finally, we suggest that path integration implemented in the generic way described here could be performed in other brain systems, including the hippocampal spatial view system of neurons which respond when a primate looks at a particular location in space, and which are updated by idiothetic eye movements made in the dark (Robertson et al. 1999; Rolls 1999; Rolls & Xiang 2006).

References

Amari S. 1977. Dynamics of pattern formation in lateral-inhibition type neural fields. Bio. Cybernetics 27:77–87.

Bassett J, Taube JS. 2005. Head direction signal generation: ascending and descending information streams. In Wiener SI. Taube JS. editors. *Head Direction Cells and the Neural Mechanisms of Spatial Orientation,* MIT Press, Cambridge, MA pp. 83–109.

Boucheny C, Brunel N, Arleo, A. 2005. A continuous attractor network model without recurrent excitation: maintenance and integration in the head direction cell system. J. Computational Neurosci 18:205–227.

Foldiak P. 1991. Learning invariance from transformation sequences. Neural Comput 3:194–200.

Georges-Francois P, Rolls ET, Robertson RG. 1999. Spatial view cells in the primate hippocampus: allocentric view not head direction or eye position or place. Cerebral Cortex 9:197–212.

Hafting T, Fyhn M, Molden S, Moser MB, Moser EI. 2005. Microstructure of a spatial map in the entorhinal cortex. Nature 436:801–806.

Hahnloser RH. 2003. Emergence of neural integration in the head-direction system by visual supervision. Neuroscienee 120:877–891.

Hertz J, Krogh A, Palmer RG. 1991. Introduction to the theory of neural computation. Addison Wesley, Wokingham, U.K.

Hestrin S, Sah P, Nicoll R. 1990. Mechanisms generating the time course of dual component excitatory synaptic currents recorded in hippocampal slices. Neuron 5:247–253.

Koch C. 1999. Biophysics of Computation. Oxford University Press, Oxford.

Lisman JE, Fellous JM, Wang XJ. 1998. A role for NMDA-receptor channels in working memory. Nature Neurosci 1:273–275.

Markus EJ, Barnes CA, McNaughton BL, Gladden VL, Skaggs W. 1994. Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. Hippocampus 4:410–421.

Markus EJ, Qin YL, Leonard B, Skaggs W, McNaughton BL, Barnes CA. 1995. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. J. Neurosci 15:7079–7094.

McNaughton BL, Barnes CA, O'Keefe J. 1983. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Exp Brain Res 52:41–49.

Muller RU, Kubie JL, Bostock EM, Taube JS, Quirk GJ. 1991. Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In Paillard J. editor. *Brain and space,* Oxford University Press, Oxford, pp. 296–333.

Muller RU, Ranck JB, Taube JS. 1996. Head direction cells: properties and functional significance. Current Opin Neurobiology 6:196–206.

Oja E. 1982. A simplified neuron model as a principal component analyser. J. Mathematical Biology 15:267–273.

O'Keefe J. 1984. Spatial memory within and without the hippocampal system. In Seifert W. editor. *Neurobiology of the hippocampus,* Academic Press, London, pp. 375–403.

O'Keefe J, Dostrovsky J. 1971. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat. Brain Res 34:171–175.

Panzeri S, Rolls ET, Battaglia F, Lavis R. 2001. Speed of feedforward and recurrent processing in multilayer networks of integrate-and-fire neurons. Network: Comput Neural Systems 12:423–440.

Ranck Jr. JB. 1985. Head direction cells in the deep cell layer of dorsolateral presubiculum in freely moving rats. In Buzsaki G. Vanderwolf CH. editors. *Electrical activity of the archicortex,* Akadémiai Kiadó, Budapest.

Redish AD, Elga AN, Touretzky DS. 1996. A coupled attractor model of the rodent head direction system. Network: Comput Neural Systems 7:671–685.

Rhodes P. 1992. The open time of the NMDA channel facilitates the self-organisation of invariant object responses in cortex. Soc Neurosci Abstracts 18:740.

Robertson RG, Rolls ET, Georges-François P. 1998. Spatial view cells in the primate hippocampus: Effects of removal of view details. J Neurophysiology 79:1145–1156.

Robertson RG, Rolls ET, Georges-François P, Panzeri S. 1999. Head direction cells in the primate pre-subiculum. Hippocampus 9:206–219.

Rolls ET. 1992. Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. Phil Trans Royal Soc 335:11–21.

Rolls ET. 1999. Spatial view cells and the representation of place in the primate hippocampus. Hippocampus 9:467–480.

Rolls ET, Deco G. 2002. Computational neuroscience of vision. Oxford University Press, Oxford.

Rolls ET, Kesner RP. 2006. A theory of hippocampal function, and tests of the theory. Prog in Neurobiology 79:1–48

Rolls ET, Robertson RG, Georges-François P. 1997. Spatial view cells in the primate hippocampus. Eur J. Neurosci 9:1789–1794.

Rolls ET, Treves A. 1998. Neural Networks and Brain Function. Oxford University Press, Oxford.
Rolls ET, Treves A, Robertson RG, Georges-François P, Panzeri S. 1998. Information about spatial view in an ensemble of primate hippocampal cells. J. Neurophysiology 79:1797–1813.
Rolls ET, Xiang J.-Z. 2006. Spatial view cells in the primate hippocampus, and memory recall. Rev Neuroscis 17:175–200.
Sargolini F, Fyhn M, Hafting T, McNaughton BL, Witter MP, Moser MB, Moser EI. 2006. Conjunctive representation of position, direction, and velocity in entorhinal cortex. Science 312:758–762.
Sharp PE. 1996. Multiple spatial-behavioral correlates for cells in the rat postsubiculum: multiple regression analysis and comparison to other hippocampal areas. Cerebral Cortex 6:238–259.
Skaggs WE, Knierim JJ, Kudrimoti HS, McNaughton BL. 1995. A model of the neural basis of the rat's sense of direction. In Tesauro G. Touretzky DS. Leen TK. editors. *Advances in neural information processing systems,* Vol. 7, MIT Press, Cambridge, Massachusetts, pp. 173–180.
Song P, Wang X-J. 2005. Angular path integration by moving "hill of activity": a spiking neuron model without recurrent excitation of the head-direction system. J. Neurosci 25:1002–1014.
Spruston N, Jonas P, Sakmann B. 1995. Dendritic glutamate receptor channel in rat hippocampal CA3 and CA1 pyramidal neurons. J. Physiology 482:325–352.
Stringer SM, Perry G, Rolls ET, Proske JH. 2006. Learning invariant object recognition in the visual system with continuous transformations. Biolo Cybernetics 94:128–142.
Stringer SM, Rolls ET, Trappenberg TP. 2005. Self-organizing continuous attractor network models of hippocampal spatial view cells. Neurobiology of Learning and Memory 83:79–92.
Stringer SM, Rolls ET, Trappenberg TP, De Araujo IET. 2002a. Self-organizing continuous attractor networks and path integration: Two-dimensional models of place cells. Network: Comput Neural Systems 13:429–446.
Stringer SM, Rolls ET, Trappenberg TP, De Araujo IET. 2003. Self-organizing continuous attractor networks and motor function. Neural Networks 16:161–182.
Stringer SM, Trappenberg TP, Rolls ET, De Araujo IET. 2002b. Self-organizing continuous attractor networks and path integration: One-dimensional models of head direction cells. Network: Comput Neural Systems 13:217–242.
Taube JS, Goodridge JP, Golob EG, Dudchenko PA, Stackman RW. 1996. Processing the head direction signal: a review and commentary. Brain Res Bull 40:477–486.
Taube JS, Muller RU, Ranck Jr. JB. 1990a. Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. J. Neurosci 10:420–435.
Taube JS, Muller RU, Ranck Jr. JB. 1990b. Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. J. Neurosci 10:436–447.
Taylor JG. 1999. Neural 'bubble' dynamics in two dimensions: foundations. Biolo Cybernetics 80:393–409.
Wang XJ. 1999. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. J. Neurosci 19:9587–9603.
Wiener SI, Paul CA, Eichenbaum H. 1989. Spatial and behavioural correlates of hippocampal neuronal activity. J. Neurosci 9:2737–2763.
Xie X, Hahnloser RH, Seung HS. 2002. Double-ring network model of the head-direction system. Physical Review E 66:041902.
Zhang K. 1996. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. J. Neurosci 16:2112–2126.