

Calibration of the head direction network: a role for symmetric angular head velocity cells

Peter Stratton · Gordon Wyeth · Janet Wiles

Received: 13 May 2009 / Revised: 18 February 2010 / Accepted: 17 March 2010 / Published online: 31 March 2010
© Springer Science+Business Media, LLC 2010

Abstract Continuous attractor networks require calibration. Computational models of the head direction (HD) system of the rat usually assume that the connections that maintain HD neuron activity are pre-wired and static. Ongoing activity in these models relies on precise continuous attractor dynamics. It is currently unknown how such connections could be so precisely wired, and how accurate calibration is maintained in the face of ongoing noise and perturbation. Our adaptive attractor model of the HD system that uses symmetric angular head velocity (AHV) cells as a training signal shows that the HD system can learn to support stable firing patterns from poorly-performing, unstable starting conditions. The proposed calibration mechanism suggests a requirement for symmetric AHV cells, the existence of which has previously been unexplained, and predicts that symmetric and asymmetric AHV cells should be distinctly different (in morphology, synaptic targets and/or methods of action on postsynaptic HD cells) due to their distinctly different functions.

Keywords Head direction · Attractor · Calibration · Angular head velocity · Symmetric cells

Action Editor: Barry Richmond

P. Stratton · G. Wyeth · J. Wiles
School of Information Technology and Electrical Engineering,
The University of Queensland,
Queensland, Brisbane, Australia

P. Stratton (✉) · J. Wiles
Queensland Brain Institute, The University of Queensland,
Queensland, Brisbane, Australia
e-mail: stratton@itee.uq.edu.au

1 Introduction

1.1 Overview

Systems that interact with an environment must be calibrated if robust, stable interaction is required. The calibration requirement holds equally for man-made and natural systems, and includes, for example, the neural systems used for reaching and grasping, locomotion, and navigation through the world. The neural substrates of mammalian navigation ability are gradually being revealed, starting several decades ago with the discovery of place cells in rats (O'Keefe and Dostrovsky 1971) which fire when an animal is in particular well-defined positions in an environment, and later with the discovery of grid cells (Hafting et al. 2005), which fire in a repeating triangular tessellation as the rat moves. The head direction (HD) system (Taube et al. 1990; Taube 2007) is postulated to be a functionally vital input to the place and grid cell regions of the brain, providing the directional information required to update the animal's location when it moves. The directional information provided by the HD system must accurately reflect the animal's motion through the world; that is, the system must be calibrated. The following sections describe the anatomy of the mammalian HD system (section 1.2), reviews current models of system functionality, in particular showing why accurate calibration of the HD system is critical for its correct operation (section 1.3), and then presents our new model of HD calibration (section 1.4).

1.2 Anatomy of the head direction system

Head direction neurons are so-called because they fire preferentially when an animal is facing in specific directions relative to invariant cues in the animal's environment

(Taube 2007). To date, such neurons have been discovered in rodents and monkeys, in several brain regions including the lateral mammillary nuclei (LMN), the antero-dorsal thalamic nuclei (ADN), and the postsubiculum (PoS), with smaller numbers being found in many other regions (Taube and Bassett 2003). The HD system is comprised of a number of small, tightly-connected subcortical structures through which information flows into the cortex, with reciprocal connections between some regions. The dorsal tegmental nuclei (DTN) receive vestibular and motor efference input. The characteristic neuron in the DTN is the angular head velocity (AHV) cell. There are two categories of AHV neuron—*symmetric* AHV neurons which fire to indicate head turning speed in both turning directions equally and *asymmetric* AHV neurons which indicate turning speed in one direction preferentially. More than half the AHV cells are symmetric and the computational function of these cells is unknown (Taube and Bassett 2003). The DTN sends connections to the lateral mammillary nucleus (LMN); in models of the HD system these connections are presumed to update the current head direction representation in the LMN, based on how fast and in which direction the head is turning as represented by the *asymmetric* AHV cells.

Each HD neuron typically has one preferred direction at which it will fire maximally, with firing tapering off as the head moves away from this preferred direction. All directions are represented approximately equally across a population of HD neurons. The combined effect of the firing of all the HD neurons is to give a unique activity pattern called the *bump* or *hill of activity* across the HD neuron population for any given direction the animal faces (Skaggs et al. 1995; Redish et al. 1996; Goodridge and Touretzky 2000; Xie et al. 2002; Song and Wang 2005). The peak of the bump represents the animal's current direction. As the animal moves, the bump translates in a systematic way through the HD neuron population such that the peak continues to represent the current head direction. When the animal stops moving, the bump is maintained in place; those neurons which are firing to represent the current direction will continue to fire at about the same rate as long as the animal's direction remains the same (Taube and Bassett 2003).

1.3 Continuous attractor models of the head direction system

In computational modelling of the head direction system, the head direction network is typically constructed as a *continuous attractor* (Sharp et al. 2001). An attractor is a dynamical system that tends towards one or more equilibrium states over time; a *continuous* attractor is stable in a continuum of states instead of a set of discrete states. For

HD networks, the attractor states are stable firing patterns able to represent any possible head direction across the population of HD neurons (hence the need for it to be a *continuous* attractor). These states are created by the interaction between two sets of weights, one excitatory and one inhibitory.

- 1) Excitatory connections strongly connect neurons with similar preferred directions, so that groups of neurons with similar preferred directions are coopted into firing together. The group of firing neurons is the *activity bump*.
- 2) Inhibitory connections strongly connect neurons with dissimilar preferred directions, so that groups of neurons with dissimilar preferred directions cannot fire together.

Despite the required excitatory connections not yet being found anatomically, the HD system exhibits attractor-like dynamics (Sharp et al. 2001). Any attractor that can represent a continuum of states, such as head direction, is prone to *drift*; that is, if the interactions between, or input to, head direction neurons are subject to any noise or inhomogeneity, the activity bump will not be stable for all directions and will instead shift through the neuron population. A slight imbalance between excitation and inhibition, for example, will likely induce movement of the attractor representation. Given the stochastic nature of neural migration and axonal growth using concentration gradients of diffusing chemicals during brain development (Scott and Luo 2001; Dickson 2002), it is most unlikely that connections can be perfectly homogenous and unbiased from birth across the entire HD neuron population, and arguably even less likely that it could function adequately without calibration throughout the entire lifetime of an animal.

The major models of HD networks assume that the network connections are predetermined and static, already perfectly balanced and calibrated (Skaggs et al. 1995; Redish et al. 1996; Goodridge and Touretzky 2000; Xie et al. 2002; Song and Wang 2005); we call this class of models the *non-adaptive attractor models*. Models in this class assume that the synaptic efficacies in the HD system are set perfectly from the outset of system operation and never need fine tuning or calibration. However, these models require exquisitely precise synaptic connectivity to maintain stable HD activity in all situations. Since it is unlikely that the HD system could be so precisely-wired from birth, or that static HD connectivity could suffice indefinitely, a calibration mechanism is likely required for both initial wiring and tuning of the system, and ongoing re-tuning in the face of injury and age-related degeneration of system components.

A small number of studies have considered the calibration issue; we call the class of models presented in these studies the

adaptive attractor models. Zhang (Zhang 1996) showed that there is a calculable solution to the weights that provides stable HD activity under stationary and translation conditions, and showed how a form of delta rule could learn these weights on the condition that an appropriate training signal is supplied, most likely a visual signal. Stringer et al. (Stringer et al. 2002) and Hahnloser (Hahnloser 2003) independently proposed mechanisms whereby the HD system could be trained by visual input, but in both cases the training algorithms rely on learning to mimic HD neuron responses to highly localised visual input that is already clearly and uniquely associated with specific head directions. Success of the HD training in these proposed algorithms critically depends on these associations; if the associations are incorrect or imperfect, then the imperfections will be carried through to the final trained HD network weights. It is unclear how such well-defined directional associations between visual landmarks and head directions could be established before the HD system is fully functional (a chicken-and-egg dilemma). At best, this assumption would require that the HD system already be ‘partially functional’ as is assumed in the current study. However, the assumed HD associations with visual landmarks would in that case likely be imperfect, and there is no method in these previous models of removing these imperfections from the final HD representations. We propose that a calibration mechanism that is independent of visual input must therefore exist.

1.4 Current work—calibrating the head direction attractor weights

The specific aim of this study was to investigate the role that symmetric AHV neurons may play in the head direction calibration mechanism, and the synaptic learning rules that are hypothesised to be involved in this process. In this study, the HD connections were assumed to be 1) initially approximately correct, but prone to both 2) random error and 3) systematic bias that must be corrected by the connection weight calibration. We will now consider each of these assumptions in turn.

1) Connections are initially approximately correct. Nominal connection strength is a function of the difference in preferred directions of the HD cells. These connections can be approximately set up as the brain is developing—using concentration gradients of diffusing chemicals (Scott and Luo 2001; Dickson 2002)—or through a random wiring process where HD cells are more likely to connect to proximal neighbours and less likely to connect to those that are more distal. These initial connections are approximate, and the resulting networks would require further tuning by a calibration mechanism such as the one presented in this study.

2) Connections have random error. A multitude of complementary and competitive forces drive synapse creation, growth and strength (Cohen-Cory 2002), and the balance between these forces changes continuously based on both intrinsic and extrinsic factors. Stochastic variation in synaptic efficacies is a result of the interactions in such a complex system.

3) Connections have systematic bias. Initial wiring of the HD network is not likely to be perfectly balanced, as discussed above. Any such bias will result in drift in the HD bump over time and the inability of the network to equally track turns in opposing directions.

In order for the HD activity to accurately and consistently represent head direction through time, three stability requirements must be fulfilled.

1) The activity bump must remain still when vestibular input indicates there is no head movement.

2) The bump must move at the same speed in both directions given turns of the same speed.

3) A continuous turn through 360° must return the bump to its starting point.

The learning mechanism hypothesised here uses vestibular and motor efference input, represented by the firing of the symmetric AHV neurons, as an informative signal that indicates whether the head is turning and therefore whether the HD activity bump should be moving. One learning mechanism is able to fulfil, with no dependence on visual input, the first two of the stability requirements. We conjecture that for the fulfilment of the third requirement, external visual input (or possibly directional auditory, tactile or olfactory input) is a necessity. However, unlike earlier models of HD attractor calibration which require many landmarks at defined positions, our algorithm would require only one landmark and its position does not need to be predefined. Effectively, the calibration algorithm for stability requirements 1 and 2 turns the multi-parameter tuning problem of training all HD weights into a single-parameter problem¹, which is to tune the gain on the symmetric AHV signal in order to control the absolute speed of the bump in the network (i.e. requirement 3). As this study concerns a potential non-visual mechanism of HD calibration, the visually-modulated mechanism needed for requirement 3 is not considered further here.

2 Methods

2.1 Network model architecture

The adaptive attractor network representing the HD system was configured similarly to non-adaptive attractor models

¹ The authors wish to thank an anonymous reviewer for suggesting this interpretation of the calibration algorithm.

(Skaggs et al. 1995; Redish et al. 1996; Goodridge and Touretzky 2000; Xie et al. 2002; Song and Wang 2005). It contained populations of four distinct types of neurons; head direction (HD) cells, left turn angular head velocity (AHV) cells, right turn AHV cells, and symmetric AHV cells (see Fig. 1). The symmetric AHV cells were an added component not found in the above non-adaptive attractor models. Each set of cells in the other populations was connected in a ring so that head direction activity “wrapped around” and restarted after a full rotation.

The relevant subcortical brain regions for the adaptation mechanisms hypothesised by this model are DTN, containing the symmetric and asymmetric AHV cells, and LMN, containing the HD cells to which the AHV cells directly project. The thalamic and cortical regions which are also known to contain HD cells are downstream of DTN and LMN and contain a much lower percentage of AHV cells (Taube and Bassett 2003), and hence were not included in the model.

The untrained HD connection weights were initialised with random noise and a systematic bias. The bias was implemented as a shift in the postsynaptic target cells of the recurrent excitatory HD connections. A uniform clockwise shift of just five cells, out of a total of 100 HD cells in the network, resulted in a significant continuous drift in the HD activity bump (see Fig. 2). In order to satisfy the stability requirements for the attractor, this drift had to be corrected by the adaptation mechanism.

The symmetric AHV cells in the model were represented by a single neuron that connected to every HD cell (see Fig. 1). The simplification of having a single symmetric AHV cell was feasible since all symmetric AHV cells are

functionally equivalent; having many of these cells connected to the HD neurons with weaker synapses is functionally identical to having one such neuron with strong synapses.

2.2 Neuron and synapse models

Neurons were modelled as leaky integrate-and-fire (LIF) cells (Stein 1967) (for parameters see Table 1).

Synaptic currents were modelled as fast rise, slow decay currents (for parameters see Table 2). For simplicity and speed of simulation, this rise/decay profile was an amalgam of AMPA/GABA_A (fast rise, fast decay) and NMDA/GABA_B (slow rise, slow decay) currents, both of which are required for recurrently-sustained neural activity (Wang 1999).

The rate of change of neuron j 's membrane potential, v_j , is given by

$$\dot{v}_j = -C_j(l_j + s_j) \quad (1)$$

where C_j is the capacitance of the cell (C_{excit} for excitatory cells and C_{inhib} for inhibitory cells—see Table 1), l_j is the leak current given by

$$l_j = G_{\text{leak}}(v_j - V_{\text{rest}}) \quad (2)$$

where G_{leak} and V_{rest} are constants (see Table 1) and s_j is the synaptic input current given by

$$s_j = \sum_i W_{j,i} p_{j,i} (v_j - V_i) \quad (3)$$

for neurons $i=1$ to 100 synapsing on (connected to) neuron j . V_i is the synaptic current reversal potential (V_{excit} for




Fig. 1 Head direction network. HD cells excite their close neighbours strongly and more distant neighbours less strongly, and this self-excitation creates the HD activity bump. The HD cells also excite the asymmetric AHV cells using a similar neighbourhood relation. The left turn AHV cells then project back to the HD cells with an offset in one direction (“leftwards”) and the right turn AHV cells project back to the HD cells with offset in the other direction (“rightwards”). These

AHV projections back to the HD cells are inhibitory, and the combined effect of these two populations is to constrain growth of the bump to be within the bounds of the offset distances in both directions. Input from the symmetric AHV cells is used to calibrate the HD connections. DTN: Dorsal tegmental nucleus, LMN: Lateral mammillary nucleus

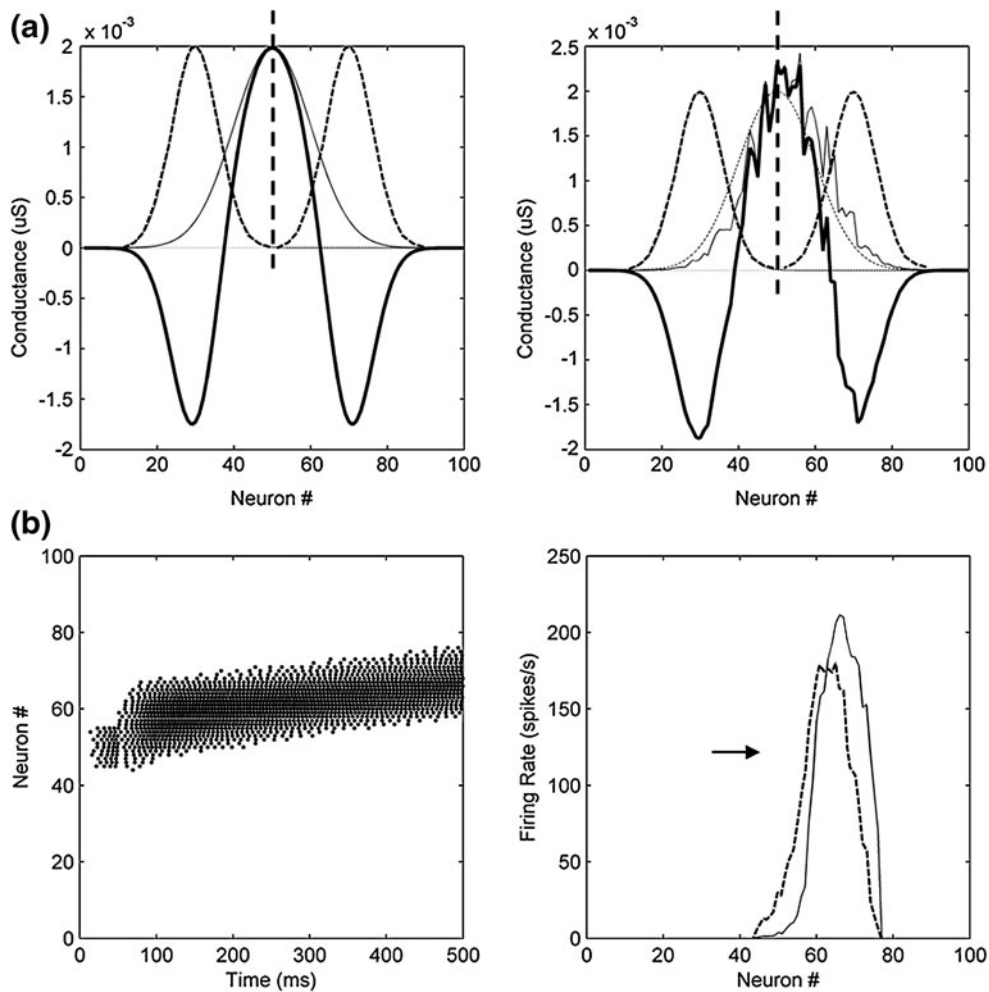


Fig. 2 Initial network state. (a) Connection weights for ideal (left) and biased, noisy (right) networks. Shown are the functions for all connections from neuron #50 to all the neurons in the network. Thin solid line—HD weights; thin dashed lines—left and right turn weights (inhibitory); thick solid line—summed total synaptic weights. LEFT: Ideal Gaussian weights with no systematic HD connection shift and no weight noise. RIGHT: Systematic HD shift of 5 units to the right and noise added to the weights. Notice that the total synaptic weight (thick black line) is off-centre and is less inhibitory on the right side due to the systematic shift. For reference, the ideal Gaussian HD weights are also shown as a dot-dash line, with the centre marked by the thick vertical dashed line. For all simulations $n=100$, $G_{excit}=G_{inhib}=$

$0.002 \mu\text{S}$, $r_{HH}=13$, $r_{LH}=r_{RH}=7$, and $o_{inhib}=26$. (b) Activity bump created by noisy, shifted weights as shown in (a). LEFT: An activity bump is initiated at cell 50 (out of $n=100$ cells in the HD network). It is then left to run for 500 ms of simulated time, during which it drifts through approximately 10 cells or 36° . RIGHT: The current bump firing rate (solid black line) is trailed slightly by the current short-term exponential moving average firing rate (dashed line). The direction of bump drift movement is indicated by the arrow. The drift must be corrected by the calibration mechanism in order to address part 1 of the stability requirements (i.e. keeping the attractor still when vestibular input indicates there is no head movement)

excitatory presynaptic cells and V_{inhib} for inhibitory—see Table 2), $W_{j,i}$ is the synaptic weight from neuron i to neuron j (in the range $[0..G_{excit}]$ or $[0..G_{inhib}]$) and $p_{j,i}$ is the proportion of synaptic channels currently open, given by

$$p_{j,i} = \delta(t - t_i)(1 - p_{j,i})P_{syn} - \frac{p_{j,i}}{\tau_{syn}} \quad (4)$$

where t_i is the time of the last spike from presynaptic neuron i and δ is the Dirac delta function.

Simulations were conducted using forward Euler integration with a 1 ms timestep.

Table 1 Leaky integrate-and-fire cell parameters

V_{rest}	Resting potential	-70 mV
V_{thresh}	Spike threshold	-52 mV
V_{reset}	Spike reset potential	-59 mV
G_{leak}	Leak conductance	0.02 μS
C_{excit}	Capacitance of excitatory (HD) cells	0.5 nF
C_{inhib}	Capacitance of inhibitory (AHV) cells	0.25 nF

Table 2 Synapse parameters

G_{excit}	Max. excitatory conductance	0.002 μ S
G_{inhib}	Max. inhibitory conductance	0.002 μ S
V_{excit}	Excitatory current reversal potential	0 mV
V_{inhib}	Inhibitory current reversal potential	-90 mV
P_{syn}	Synaptic channel opening probability	0.2
τ_{syn}	Synaptic current time constant	100 ms

2.3 Network connections

The equations governing connection strengths between cells are as follows: There are $n=100$ HD cells in the network. These cells are connected by recurrent excitatory weights, given by a Gaussian function of the distance between the cells (see Fig. 2). The strength of the connection W^{HH} from HD cell i to HD cell j is given by

$$W_{j,i}^{HH} = G_{excit} e^{-d_{j,i}^2/2r_{HH}^2} f_{noise}(\cdot) \quad (5)$$

where r_{HH} is the distance, in number of cells, of one standard deviation of the Gaussian (set to $n/8$ for all trials) and d is the circular distance (*i.e.* with wrapping at the ends) from cell j to cell i , given by

$$d_{j,i} = \min[\text{mod}(j - (i - o_i), n), \text{mod}((i - o_i) - j, n)] \quad (6)$$

where o_i is a systematic shift which can be applied to the direction of the Gaussian weights (set to $0.05n$). It is this shift which, when non-zero, can cause a continuous drift in the attractor bump and which can be corrected for by the adaptation algorithms presented in this study. $f_{noise}(\cdot)$ is a noise function defined as

$$f_{noise}(\cdot) = 1 + \lambda \cdot \text{gauss}(\cdot) \quad (7)$$

where λ is a noise weighting (between 0 and 0.2, varied between trials) and $\text{gauss}(\cdot)$ is a Gaussian random process with zero mean and unit standard deviation.

The HD cells are also directly connected to the left and right turn asymmetric AHV cells with excitatory connections. There are n left turn and n right turn AHV cells. These connections create a corresponding bump of activity in the AHV cells, and the strength of these connections is a Gaussian function of the distance between the HD and AHV cells, as follows: the connection strength W^{LH} from HD cell i to left-turn (LT) cell j is given by

$$W_{j,i}^{LH} = G_{excit} e^{-d_{j,i}^2/2r_{LH}^2} \quad (8)$$

where r_{LH} is the distance, in number of cells, of one standard deviation of the Gaussian (set to $r_{HH}/1.8$ for all trials), and d is the circular distance (*i.e.* with wrapping at the ends) from cell i to cell j , given by

$$d_{j,i} = \min[\text{mod}(j - i, n), \text{mod}(i - j, n)] \quad (9)$$

Similarly, the strength of the connection W^{RH} from HD cell i to right-turn (RT) cell j is given by

$$W_{j,i}^{RH} = G_{excit} e^{-d_{j,i}^2/2r_{RH}^2} \quad (10)$$

where d is as above and $r_{RH} = r_{LH}$.

The AHV cells connect back to the HD cells with offset inhibitory connections (see Figs. 1 and 2); the offset is to the “left” for all left-turn AHV cells and to the “right” for all right-turn AHV cells. These are the connections that directly stop the bump activity from spreading throughout the HD network, and are also the connections that allow controlled translation of the bump activity through selective inhibition of one or the other of the asymmetric AHV cell populations, which weakens inhibition to one side of the HD activity bump, and therefore causes it to move in that direction. The connection strength W^{HL} from left-turn cell i to HD cell j is given by

$$W_{j,i}^{HL} = G_{inhib} e^{-d_{j,i}^2/2r_{HL}^2} \quad (11)$$

where d is the circular distance (*i.e.* with wrapping at the ends) from cell j to cell i , given by

$$d_{j,i} = \min[\text{mod}(j - (i - o_{inhib}), n), \text{mod}((i - o_{inhib}) - j, n)] \quad (12)$$

where o_{inhib} is the turn connection offset from the centre of the HD activity bump (for all trials set to $2r_{HH}$). The turn connection offset is how far the inhibitory connections from the left and right AHV cell populations are offset where they project back into the HD cells.

Similarly, the connection strength W^{HR} from right-turn cell i to HD cell j is given by

$$W_{j,i}^{HR} = G_{inhib} e^{-d_{j,i}^2/2r_{RH}^2} \quad (13)$$

where d is the circular distance (*i.e.* with wrapping at the ends) from cell i to cell j , given by

$$d_{j,i} = \min[\text{mod}(j - (i + o_{inhib}), n), \text{mod}((i + o_{inhib}) - j, n)] \quad (14)$$

Notice the change in sign for the offset term (o_{inhib}) compared to Eq. (12), which offsets the inhibition from the right-turn AHV cell population in the opposite direction to the left-turn cells.

3 Connection weight adaptation

The connection weight adaptation addresses the first and second calibration requirements, namely:

- 1) Keeping the attractor still when there is no vestibular input; this is *no-turn calibration*.
- 2) Moving the attractor at the same speed in both directions given turns of the same speed; this is *equal-turn calibration*.

HD cells in the model maintain a representation of the instantaneous change in their firing rates, which corresponds to how fast the activity bump is moving through the HD network. The HD cells also receive input from the symmetric AHV cells which convey how fast the head is currently physically turning. If the movement of the HD bump does not correspond to the head movement indicated by the symmetric AHV cells, the recurrent HD connection strengths are adjusted so that the bump more closely follows the indicated movement. Note that the connections from the AHV cells are not modified; these connections only carry the physical head movement information. Rather it is the connections that maintain the bump dynamics—the recurrent excitatory HD connections—that are adjusted.

3.1 HD connection weight updates

The change in the strength of the connection W from presynaptic HD cell i to postsynaptic HD cell j is given by:

$$\Delta W_{j,i} = \alpha \Delta r_i (|\Delta r_j| - AHV_{sym}) \tag{15}$$

where α is the learning rate, r_i is the current firing rate of cell i , $|\dots|$ denotes the absolute value of a term and AHV_{sym} is the current input from the symmetric AHV cells. The instantaneous change in HD cell i 's firing rate is calculated as the difference between its current instantaneous firing rate and its short term average firing rate:

$$\Delta r_i = r_i - m_s(r_i) \tag{16}$$

where $m_s(\cdot)$ is the short-term average, calculated as an exponential moving average with a time constant of 10 ms. The current instantaneous firing rate of cell i , r_i , is given by:

$$r_i = \begin{cases} \frac{1}{t_s - t'_s} & \text{when } \delta(t - t_s) \neq 0 \\ \frac{1}{t_s - t'_s} e^{-(t-t_s)/\tau} & \text{when } \delta(t - t_s) = 0 \end{cases} \tag{17}$$

where t_s is the time of the last spike, t'_s is the time of the second-last spike (hence $t_s - t'_s$ is the most recent interspike interval (ISI)), τ is the firing rate time constant (100 ms) and δ is the Dirac delta function.

Intuitively, the weight adaptation works as follows. If the bump is moving correctly at the speed indicated by the AHV input, then AHV_{sym} and the absolute change in neuron j 's firing rate r_j are equal, so the term $|\Delta r_j| - AHV_{sym}$ is zero and no weight change takes place.

If the bump is moving too quickly, then AHV_{sym} is less than the absolute change in neuron j 's firing rate, so the

term $|\Delta r_j| - AHV_{sym}$ is positive. In this case, connections are weakened from those presynaptic HD cells whose firing rates are decreasing (Δr_i is negative), and strengthened from those presynaptic HD cells whose firing rates are increasing. HD cells with decreasing firing rates are always on the trailing edge of the bump, while cells with increasing firing rates are always on the leading edge of the bump (see Fig. 3). Weakening the ‘forward’ weights from cells on the trailing edge (where the bump just came from) retards the movement of the bump because it reduces the impetus of the bump to move forward, while strengthening the ‘backward’ weights from cells on the leading edge (where the bump is moving to), has the same effect because it is trying to stabilise the bump in its current position. Thus when the bump is moving too quickly, the net effect of the weight updates is to slow it down. The converse reasoning is applied to show that when the bump is moving too slowly, the net effect of the weight updates is to speed it up.

For the special case of the head being stationary (*i.e.* $AHV_{sym}=0$) then any movement of the bump (*i.e.* $|\Delta r_j|$ positive) is analogous to the bump moving too quickly. Weights from cells on the trailing edge of the bump are weakened, and weights from cells on the leading edge are strengthened, retarding further movement of the bump. Once the bump stops moving, $|\Delta r_j|=0$ and no further weight changes occur.

The learning rate α was set to 10^{-10} in the case of symmetric AHV input indicating that the head was not moving ($AHV_{sym}=0$), and was set to 10^{-9} during head turns. The slower learning rate during no-turn calibration ensured that the HD attractor connections did not preferentially learn to hold the bump still in all situations. Connection weight normalisation was applied to each postsynaptic neuron every 1,000 ms of simulated time. The normalisation ensured that weights remained bounded. Evidence suggests that normalisation does occur in nervous systems for both total synaptic weight (Royer and Pare 2003) and mean firing rate (Turrigiano and Nelson 2004).

3.2 Training regimen

Given the network and connections as described above, an activity bump was initiated by stimulating the central HD neuron and the neurons within a given distance of the centre for 100 ms. Such stimulation caused enough neurons to begin spiking for the activity bump to be self-sustaining around the central neuron. The network was then subjected to random vestibular input comprised of interspersed periods of turns and no turns. Each period lasted a random duration from zero to 3 s of simulated time, after which a new vestibular input was randomly chosen for the next period. Consecutive periods with the same vestibular input (e.g. two turns in the same direction) were permissible.




Fig. 3 Synaptic updates by the adaptation rules. **(a)** Learning to slow the activity bump down when it is moving too fast. The graph shows the current bump across the HD cell population (*solid line*) and the short term exponential moving average (*dashed line*) which lags slightly behind. The bump is moving too quickly to the left (indicated by the *straight arrow*). Because the bump is moving too quickly, the change in postsynaptic neuron j 's firing rate is greater than the symmetric AHV input, and the term $|\Delta r_j| - AHV_{sym}$ is positive. Therefore the connection weights from presynaptic neurons i whose firing rates are increasing, on the leading edge of the bump, are increased (marked with the *plus symbols* and *thick arrow*), and from those whose firing rates are decreasing, on the trailing edge of the

bump, are decreased (*minus symbols* and *thin arrow*). These weight updates mean that the bump propagates forward more slowly next time it moves through this region of the HD network. **(b)** The converse case, learning to speed the bump up when it is moving too slowly. The bump is moving slowly to the right (*straight arrow*). If the firing rate of a neuron is changing too slowly (*i.e.* the bump is moving too slowly), the connection strengths from other HD cells whose firing rate is reducing need to be increased (*thick arrow*), while connections from those cells whose firing rate is increasing need to be reduced (*thin arrow*). These connection strength updates cause the bump to propagate ahead faster next time it reaches the cells in the 'plus' region, therefore increasing the speed of the bump at that point

Networks were trained for 1,200 s (20 min) of simulated time.

4 Results

The attractor weights were initialised with noisy weights and a uniform bias as shown in Fig. 2(a) (right), inducing a continuous clockwise drift of approximately $60^\circ/\text{s}$. The network's angular representation was tested before (see Fig. 4(a)) and after training. Also, to compare turning rates in both directions, the *turn rate error* was calculated before and after training as follows: The network was subjected to vestibular input indicating a turn in one direction, and the network turn angle noted, followed by a turn of identical speed and duration in the opposite direction. The baseline turning rate was calculated as the average of the two turn angles as represented in the HD network. The absolute difference between either angle and the average was then expressed as a percentage, and deemed the turn rate error:

$$\text{turnrate}_{err} = \text{abs}\left(\frac{100(\theta_1 - (\theta_1 + \theta_2)/2)}{(\theta_1 + \theta_2)/2}\right) \quad (18)$$

where θ_1 and θ_2 are the measured turn angles. Therefore a turn rate error of 0% means that turns were of equal speed in both directions, while an error of 33% means that a turn in one direction proceeded twice as fast as in the other.

Prior to training, the uniform bias in the HD weights induced a continuous drift in the bump position and caused turns in the clockwise direction to be much faster than anticlockwise, giving a turn rate error of 70% (see Fig. 4 (a)). After training, the drift was removed and the turn rate error was reduced to 0.5% (see Fig. 4(b)). The adaptive attractor network successfully removed the bias from, and corrected for the noise in, the HD weights, resulting in an HD bump that was stable in all positions and translated through the network uniformly in both directions.

5 Discussion

The results have shown how head direction network connections need not be perfectly pre-wired with static weights. Instead an adaptive mechanism can adjust the weights to compensate for both systematic shift and noise. This mechanism is consistent with current knowledge about the HD system and is compatible with existing (*i.e.* attractor-based) computational models (Skaggs et al. 1995; Redish et al. 1996; Goodridge and Touretzky 2000; Xie et al. 2002; Song and Wang 2005).

Whilst asymmetric AHV cells are presumed by all models of the HD system (including the current model) to be used to update the position of the HD bump during head turns, the symmetric cells cannot be used for this purpose since, by definition, they do not differentiate between turns

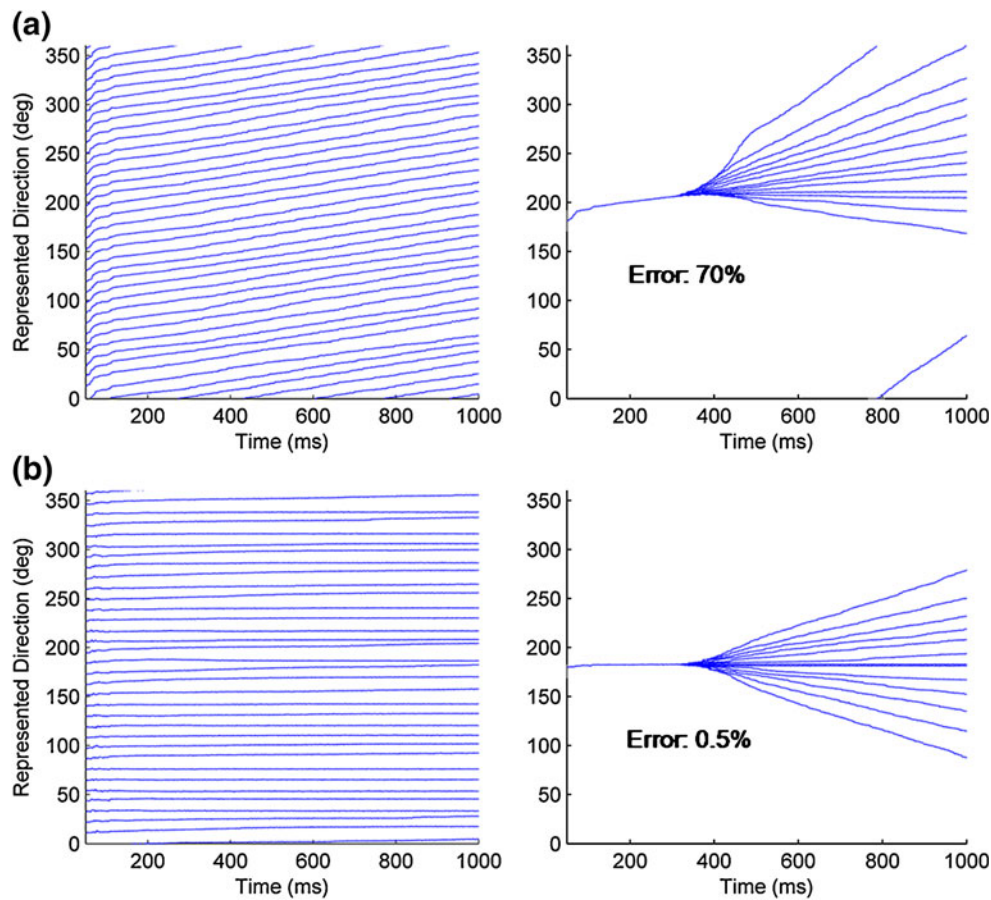


Fig. 4 Drift and turn rate over time. **(a)** Initial conditions prior to training. *LEFT*: The initial weights were noisy and shifted to the right, and the shift induced a continuous drift in the bump position. To create this graph, an attractor bump was initiated with its centre at each cell position in turn, and was then left to run for 1,000 ms of simulated time, during which it drifted continuously to the right. At each simulation time-step, the position of the bump was calculated using population vector decoding (Georgopoulos et al. 1986). With ideal weights (noiseless, no shift), cell n represented a head direction of 360° and cell 1, $360/n^\circ$; however due to the shift, the bump

of either direction. Until now it has been difficult to explain why more than half of the AHV cells should be of the symmetric type (Taube and Bassett 2003). In our adaptive attractor model, because connection updates are calculated and applied on an HD cell-by-cell basis from information provided by the symmetric AHV cells, every HD cell must have at least one connection from a symmetric AHV cell. Any HD cell that does not have such a connection does not receive the information required to update its excitatory connections as per Eq. (15). There must therefore be sufficient symmetric AHV cells to ensure that every HD cell is very likely to receive symmetric AHV input. Asymmetric AHV cells, on the other hand, need not strictly connect to every HD cell. Because HD cells receive all their excitation from other HD neighbours, and a single HD cell (or even a small group of HD cells) cannot sustain its own

translated about 60° in 1,000 ms, and if left long enough would complete a full revolution of the HD network in just over 6 s. *RIGHT*: To test the integration of head turns over time, a turn of a given speed was initiated at 300 ms, first in one direction and then in the other. Multiple turning speeds were tested and the bump position recorded in each case. The asymmetry in turning velocity is clear. Turn rate error (see text) is marked on the graph. **(b)** Results after training. *LEFT*: The continuous drift has been eliminated; the bump stayed at rest at each position it was initiated. *RIGHT*: The turn rate error was reduced to 0.5%

activity, then provided the bump as a whole is receiving sufficient movement impetus, the entire bump will be coerced into moving, including any HD cells not receiving direct asymmetric AHV input. The adaptive attractor theory predicts that there should be more connections from symmetric than asymmetric AHV cells to the HD network. Such connectivity could be accomplished with a greater number of connections *per se*, or else with a greater number of symmetric cells, as is known to be the case in the DTN of the rat (Taube and Bassett 2003).

Note that if the head direction bump is drifting in one direction and the head is turning in the other at a rate *slower* than the drift (so that the resultant bump movement is in the direction counter to the current actual head turn) then initially the calibration will erroneously reinforce the turn in the wrong direction. This effect is temporary though, as it

is counteracted by calibration that occurs when the head is not moving, as well as at those times when the head is turning faster than the drift. Even for fast initial drift velocities, the calibration at those times when the head is not moving eventually slows down the bump drift sufficiently that calibration during head turns is able to correctly calibrate the HD network.

The calibration information provided to the network by the symmetric AHV cells cannot be gained by simply summing the inputs from asymmetric AHV cells with opposite turning preferences. Asymmetric AHV cells have a broad range of gains, meaning that different cells fire at different rates for same-speed head turns. Two asymmetric cells that appear to be firing at the same rate for turns in opposite directions could be doing so because the turns are at different speeds. To use asymmetric AHV cells in place of symmetric cells would require each HD cell to connect to two (or more) asymmetric cells with exactly matching gains but representing turns in opposite directions. However for HD cells to find matching asymmetric cells they would need access to information about absolute turn speed, which is the information provided by symmetric AHV cells. The symmetric AHV cells are therefore vital for the functioning of this calibration mechanism.

Attractor networks are difficult to calibrate since even a slight connectivity imbalance can cause drift that accumulates catastrophically over time. The presence of symmetric AHV cells, which provide an accurate, reliable absolute head rotational velocity signal exactly where it would need to be to enable calibration of the head direction network as hypothesised in this study, is intriguing. The question that logically follows—how are the symmetric AHV cells created or calibrated?—is a separate issue which nevertheless bears comment. If these cells indeed require calibration, then it could be achieved via a combination of vestibular input, motor efference from head rotation movements, and proprioception in the neck (spine, muscles and skin). However, it is also possible that calibration is not necessary, since if a symmetric AHV cell samples a sufficient number of outputs from randomly selected asymmetric cells, then its own activity profile would be symmetric. In this case, random sampling would be sufficient despite minor spatial biases since, unlike for HD cells, there is no requirement for spatial wiring of AHV cells.

There are three major requirements for the adaptation rule as per Eq. (15) to function.

- 1) HD neurons must have a representation of their current firing rate.
- 2) HD neurons must have a representation of the *change* in their firing rate (represented in Eq. (16) by the difference between the current firing rate and the short term average firing rate).
- 3) Weight change in one type of connection must be modulated by presynaptic activity impinging on the postsynaptic cell from a different synapse to the one where the weight change is effected. Specifically, the HD weight updates must take into account the vestibular input from the symmetric AHV cells.

We will now consider each of these requirements from the perspective of evidence which may support their implementation in nervous systems.

- 1) It is known that some neurons try to maintain a fixed long term average firing rate (Baddeley 1997; Turrigiano and Nelson 2004), and it is therefore possible they also have a representation of their current firing rate. A candidate mechanism for this representation is intracellular calcium concentration (Feng and Li 2003).
- 2) A recent discovery showed that even very simple animals can have senses which are sensitive to the change in concentration of a chemical, rather than the absolute concentration (Suzuki et al. 2008). Therefore if current firing rate is represented as a chemical concentration, then the direction and magnitude of a change in this concentration could presumably also be detected.
- 3) We conjecture that, since the synapses involved in this adaptation all have a common postsynaptic cell, they are able to influence each other through either their spatial proximity or through signalling mechanisms inside the cell. Synapses which are close together are known to be able to influence each other. For instance, potentiation in one synapse can lead to potentiation of adjacent synapses (Engert and Bonhoeffer 1997). Also, the locations of synapses between different brain regions are known in some cases to be very specific (for example the high laminar specificity of most hippocampal connections). Therefore modulation of synaptic connection strength by signals received on another, perhaps spatially proximal, synapse cannot be ruled out. Further, there is evidence that well-established phenomena such as spike-timing-dependent plasticity (STDP) can be strongly modulated by factors such as proximity of the synapse from the soma, to the extent that the direction of synaptic efficacy update can be reversed in distal synapses (Caporale and Dan 2008). In this case the distance of the synapse from the postsynaptic soma has a significant modulating influence on the synaptic efficacy changes. It therefore seems reasonable to conjecture that synaptic input from a second synapse could modulate another synapse's efficacy in the manner implemented by Eq. (15).

A number of experiments could be conducted to test the theory of calibration of head direction by symmetric AHV cells.

- 1) To demonstrate that the head direction system is in fact adaptive, a rat could be placed in an environment where all visible cues rotate slowly and continuously in one direction. The theory predicts that the combination of slow continuous rotation, and AHV input indicating that no rotation is taking place, would result in an adaptive update of HD connections to slow down or stop rotation of the HD bump². When the physical rotation is later stopped, the adapted connections should then cause a temporary continuous counter-rotation of the bump, before re-adaptation restores stability.
- 2) Asymmetric and symmetric AHV cells have very different functions in our adaptive attractor model. The former act as standard inhibitory cells, whereas the latter have a modulatory effect only. We would therefore expect that these cells will have different morphologies, different synaptic target areas, use different neurotransmitters and/or have different means of action on the postsynaptic cell.
- 3) If there is an identifiable neurochemical difference between symmetric and asymmetric AHV cells, it may be possible to interrupt the operation of just the symmetric cells via administration of a specific antagonist. The adaptive attractor theory predicts that the HD system would continue operating normally due to the normal functioning of the asymmetric cells but that calibration would not be possible, so over time the alignment of head direction with true heading will deteriorate, and may ultimately result in a continuous drift in the HD bump position.

6 Conclusion

Classic models of the head direction (HD) system use asymmetric angular head velocity (AHV) cells to move the activity bump in the direction of head movement. We propose a novel calibration mechanism, which we call the adaptive attractor model, in which the weights between HD neurons in an attractor network are calibrated by symmetric AHV neurons. The calibration corrects both drift and turn

rate, and we demonstrate that it is able to precisely calibrate the HD network.

We suggest that this calibration function is a plausible reason for the existence of the symmetric AHV cells, since every HD cell needs at least one connection from symmetric cells in order to receive the head turning velocity information required for it to adapt its connections with neighboring HD cells. If this conjectured function for the symmetric AHV cells is correct, we predict that symmetric and asymmetric AHV cells should be distinctly different (in morphology, synaptic targets and/or methods of action on postsynaptic HD cells) due to their distinctly different functions.

Acknowledgements The authors wish to thank two anonymous reviewers for suggestions which greatly improved the quality of this manuscript.

References

- Baddeley, R. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings: Biological Sciences*, 264(1389), 1775–1783.
- Caporale, N., & Dan, Y. (2008). Spike timing-dependent plasticity: A hebbian learning rule. *Annual Review of Neuroscience*, 31, 25–46.
- Cohen-Cory, S. (2002). The developing synapse: Construction and modulation of synaptic structures and circuits. *Science*, 298(5594), 770.
- Dickson, B. J. (2002). Molecular mechanisms of axon guidance. *Science*, 298(5600), 1959–1964.
- Engert, F., & Bonhoeffer, T. (1997). Synapse specificity of long-term potentiation breaks down at short distances. *Nature*, 388(6639), 279–284.
- Feng, J., & Li, G. (2003). The relationship between neuronal calcium concentration and firing rate during stochastic synaptic inputs. *Journal of Theoretical Biology*, 223(3), 367–375.
- Georgopoulos, A. P., Schwartz, A. B., et al. (1986). Neuronal population coding of movement direction. *Science*, 233(4771), 1416–1419.
- Goodridge, J. P., & Touretzky, D. S. (2000). Modeling attractor deformation in the rodent head-direction system. *Journal of Neurophysiology*, 83(6), 3402–3410.
- Hafting, T., Fyhn, M., et al. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436, 801–806.
- Hahnloser, R. H. R. (2003). Emergence of neural integration in the head-direction system by visual supervision. *Neuroscience*, 120(3), 877–891.
- 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(1), 171.
- Redish, A. D., Elga, A. N., et al. (1996). A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems*, 7(4), 671–685.
- Royer, S., & Pare, D. (2003). Conservation of total synaptic weight through balanced synaptic depression and potentiation. *Nature*, 422(6931), 518–522.
- Scott, E. K., & Luo, L. (2001). How do dendrites take their shape? *Nature Neuroscience*, 4, 359–365.
- Sharp, P. E., Blair, H. T., et al. (2001). The anatomical and computational basis of the rat head-direction cell signal. *Trends in Neurosciences*, 24(5), 289–294.

² This experiment assumes that slow rotation of the visual environment induces a slow rotation of head direction representation. The alternative outcome, that the animal simply becomes disoriented, could possibly be avoided by using prominent visual cues and an extended training period, prior to testing, during which there is no cue rotation.

- Skaggs, W. E., Knierim, J. J., et al. (1995). A model of the neural basis of the rat's sense of direction. *Advances in Neural Information Processing Systems*, 7, 51–58.
- 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. *Journal of Neuroscience*, 25(4), 1002–1014.
- Stein, R. B. (1967). The frequency of nerve action potentials generated by applied currents. *Proceedings of the Royal Society of London - Series B, Biological Sciences (1934–1990)*, 167 (1006), 64–86.
- Stringer, S. M., Trappenberg, T. P., et al. (2002). Self-organizing continuous attractor networks and path integration: One-dimensional models of head direction cells. *Network: Computation in Neural Systems*, 13(2), 217–242.
- Suzuki, H., Thiele, T. R., et al. (2008). Functional asymmetry in *Caenorhabditis elegans* taste neurons and its computational role in chemotaxis. *Nature*, 454(7200), 114.
- Taube, J. S. (2007). The head direction signal: Origins and sensory-motor integration. *Annual Review of Neuroscience*, 30, 181–207.
- Taube, J. S., & Bassett, J. P. (2003). Persistent neural activity in head direction cells. *Cerebral Cortex*, 13(11), 1162–1172.
- Taube, J. S., Muller, R. U., et al. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, 10(2), 420–435.
- Turrigiano, G. G., & Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5(2), 97–107.
- Wang, X. J. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *Journal of Neuroscience*, 19(21), 9587–9603.
- Xie, X., Hahnloser, R. H. R., et al. (2002). Double-ring network model of the head-direction system. *Physical Review E*, 66(4), 41902(1–9).
- Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, 16(6), 2112–2126.