An Implementation of a Biologically Inspired Model of Head Direction Cells on a Robot

Theocharis Kyriacou*

Research Institute for the Environment, Physical Sciences and Applied Mathematics (EPSAM), Keele University, Keele, UK
t.kyriacou@cs.keele.ac.uk

Abstract. A biologically inspired model of head direction cells is presented and tested on a small mobile robot. Head direction cells (discovered in the brain of rats in 1984) encode the head orientation of their host irrespective of the host's location in the environment. The head direction system thus acts as a biological compass (though not a magnetic one) for its host. Head direction cells are influenced in different ways by idiothetic (host-centred) and allothetic (not host-centred) cues. The model presented here uses the visual, vestibular and kinesthetic inputs that are simulated by robot sensors. Three test cases are presented that cover different state combinations of the inputs. The test results are compared with biological observations in previous literature.

Keywords: biologically inspired robot navigation, head direction cells.

1 Introduction: Biologicaly Inspired Robot Navigation

Biologically inspired navigation methods can perhaps be put in two broad categories. In the first, methods draw only from observations of animal behaviour without considering the underlying cognitive mechanisms that play part in navigation. Tolman in 1948 (see [19]) was among the first who conducted experiments with rats that allowed such observations but more recent work using rodents and insects is presented for example in [2] and [22].

In contrast, "bottom-up" approaches to bio-inspired models of navigation make use of knowledge obtained by observing the brain activity of animals while they perform navigational tasks (see for example [11]). During such experiments, the activity of a few brain cells can be recorded by means of microelectrodes. This gives some clues as to how a navigational mechanism is implemented in the brain. Bio-inspired models in this category make quite a lot of extrapolations that try to fill in the gaps.

Three types of brain cells called *place cells* (see [13]), *head direction cells* (see [18]) and *grid cells* (see [9]) have been discovered (mostly from experiments conducted on rats) and are thought to play a significant role in animal navigation.

* The author would like to thank his colleagues Charles Day and John Butcher for the useful discussions he had with them during the work presented in this paper.

R. Groß et al. (Eds.): TAROS 2011, LNAI 6856, pp. 66–77, 2011.

The work in this paper concentrates on head direction cells (or HD cells) and presents a model of these cells that is inspired by biological observations. A more detailed description of HD cells and an overview of previous work in modelling them is presented below.

1.1 Head Direction Cells

The most recent comprehensive review of neurophysiological observations related to HD cells is found in [23]. Here below, the main characteristics of the HD system are outlined.

Head direction cells were first discovered in 1984 by Ranck Jr. and more detailed findings on them were published in 1990 by Taube and colleagues (see [18]). An HD cell fires maximally when the animal's head points in a particular direction. This is called the *preferred head direction* of the particular cell. The HD system includes a population of HD cells with preferred head directions distributed through 360°. The activity of an HD cell does not depend on the location of the animal in the environment. The head direction cell system can thus be thought of as being a biological head compass (though not a magnetic one) that is influenced by several senses. When an animal is placed in a new environment the preferred head direction of each cell in the HD system quickly settles to an arbitrary value. The system maintains this alignment for the specific environment even if the animal is removed and re-introduced back to the same environment. This alignment will only be reset (thus the environment will be treated as a new, previously unseen one) if several weeks have passed before the animal is re-introduced in the environment ([18]). The visual sense is the major input that helps to align the HD system when the animal is introduced in a previously visited environment. Strong visual cues (for example large and prominent landmarks) can influence the preferred head direction of HD cells (see experiments and observations in [12] and [18]). Another major input to the HD system is the vestibular sense. This allows the animal to maintain a correct head direction for some time after visual input is removed (by switching off the lights for example). Apart from the most important two inputs to the HD system mentioned above, other cues also play a part in influencing the preferred head direction of HD cells. These include olfactory cues (see [8]) and cues that are involved in self-locomotion (motor, kinesthetic and proprioceptive) (see [17] for a review). Cues to the HD system are classified in two categories: they can be *allothetic*, i.e. not self-centred (for example visual and olfactory) or *idiothetic*, i.e. self-centred (for example vestibular and kinesthetic). When two or more inputs to the HD system are in conflict the preferred head direction of the HD cells depends mainly on two factors: (a) the relative influence to the HD system of each of the inputs in conflict and (b) the extent of the conflict (see example in [6]). Conflicts are extensively discussed in [23]. See also [17] for a more concise review.

1.2 Models of Head Direction Cells

Several attempts have been made to model the HD system. Redish and colleagues in [14] and Goodridge and Touretzky in [7] present anatomically faithful models but these only use the vestibular sense. In both cases the neural network weights are prescribed (i.e. not obtained by training). The models are however tested using real data obtained from experiments with rats. Models using both visual and vestibular inputs are presented in [21], [16] and [24]. In all three cases the model weights are also prescribed and tests are carried out in simulated environments. Also, only in [24] were input conflict situations simulated and compared with biological observations. Only a few examples exist in the literature of HD system models applied to real robotic agents. Of these, the most notable are presented in [1] and [4]. Both models incorporate visual and vestibular inputs but again, in both cases the models weights are prescribed.

In the work presented here, the head direction cell model by Stinger and colleagues in [16] is extended by incorporating in it the kinesthetic input (in addition to the visual and vestibular input). No previous work has been found that incorporates all three inputs to an HD cell model. Furthermore, the proposed model is trained and tested using real robot data. Conflict situations between the three inputs to the model are also presented and compared with biological observations.

Details of the model presented in this paper are given in section 2. The experimental procedure, testing and results are presented in section 3 and discussed in section 4. A summary and concluding remarks are presented in section 5.

2 Methods: The Model

The HD system model presented here is using a continuous attractor neural network (see introduction by Trappenberg in [3]). A continuous attractor neural network (CANN) is a network of interconnected nodes. In a fully connected network each node is connected via weighted connections to every other node including itself. A CANN is thus a form of recurrent network. The operation of the network is such that nodes in close association excite each other (via excitatory connections). The amount of excitation being proportional to the degree of association between the nodes. On the other hand, nodes that are less associated with each other are connected with inhibitory connections. The connection weights from one node to all other nodes are often prescribed by a Gaussian function with its tails below zero (the negative weight values for the inhibitory connections). A CANN gives rise to a self-sustained "hill" of excitation (the attractor) in the network. If the network is perfectly symmetrical about each node (both in connectivity and weight values to other nodes) the attractor will be stationary when the network has no external influences. External input stimuli that temporarily distort the network's symmetry (by biassing the activation of nodes) can cause the attractor to move.

In the case of the HD system model presented here the CANN network is comprised of 360 nodes (the HD cells) and it is fully connected. Each node is

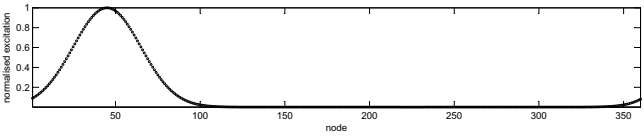

Fig. 1. The state of the HD system when it points at 45°

associated with a preferred head direction and it is therefore most active when the subject (the host of the HD system) is facing in that direction. The nodes can be thought to be arranged in a circle in order of preferred head direction[1]. Figure 1 shows the state of the HD system when it points to 45° (from an arbitrary reference direction).

The state (i.e. the excitation of each node) of the network implementation presented here is a function of the previous state of the system and three inputs: visual, vestibular and the kinesthetic. Figure 2(left) is a partial diagram that shows how these inputs are connected to the HD cells. The reader is referred to the caption of the figure for a detailed explanation.

The activation h_i^{HD} at time t of a head direction cell i in the model presented here can be determined using the following differential equation:

$$\tau \frac{\mathrm{d}h_i^{\mathrm{HD}}(t)}{\mathrm{d}t} = -h_i^{\mathrm{HD}}(t) + \theta^{\mathrm{RC}} \sum_j (w_{ij}^{\mathrm{RC}} - w^{\mathrm{INH}}) r_j^{\mathrm{HD}}(t) + \theta^{\mathrm{VIS}} I_i^{\mathrm{VIS}}(t)$$
$$+\theta^{\mathrm{VES}} \sum_{jk} w_{ijk}^{\mathrm{VES}} r_j^{\mathrm{HD}}(t) r_k^{\mathrm{VES}}(t) + \theta^{\mathrm{KIN}} \sum_{jk} w_{ijk}^{\mathrm{KIN}} r_j^{\mathrm{HD}}(t) r_k^{\mathrm{KIN}}(t) \quad (1)$$

where $r_i^{\mathrm{HD}}(t)$ is the firing rate (excitation) of HD cell i given by the sigmoid function:

$$r_i^{\mathrm{HD}}(t) = \frac{1}{1 + e^{-2\beta h_i^{\mathrm{HD}}(t)}} \quad (2)$$

τ is the time constant of the system, w_{ij}^{RC} is the recurrent connection weight from HD cell j to HD cell i, I_i^{VIS} is the visual input to HD cell i, r_k^{VES} is the firing rate of vestibular sensor cell k, w_{ijk}^{VES} is the weight value of the connection from HD cell j to VCL cell i that is associated with the vestibular sensor cell k. Similarly, r_k^{KIN} is the firing rate of kinesthetic sensor cell k and w_{ijk}^{KIN} is the weight value of the connection from HD cell j to KCL cell i that is associated with the kinesthetic sensor cell k. The factors θ^{RC}, θ^{VIS}, θ^{VES} and θ^{KIN} control the influence of the recurrent, visual, vestibular and kinesthetic inputs respectively to the HD cells. Finally, w^{INH} is a constant negative offset to the recurrent connection weights that serves as a quick way to make the connection weights

[1] This arrangement however is not necessary. In fact, in the brain, HD cells have not been found to be arranged in any particular order that relates to their preferred head direction.

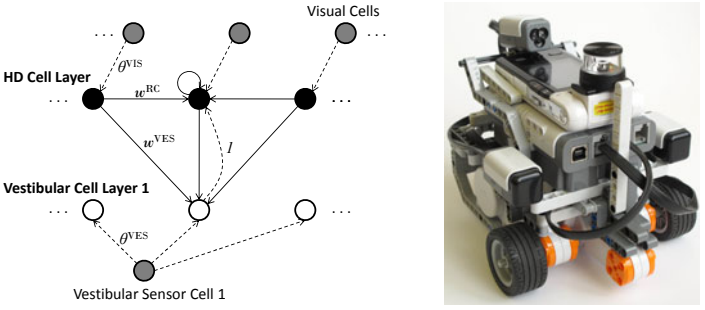

Fig. 2. Left: A partial diagram of the network of the HD model. The HD cell layer contains the CANN network nodes. These are the HD cells and they are fully interconnected. The visual input is a 360-element vector that is directly (one-to-one) connected to the HD layer. The excitation of the HD layer nodes is the output of the system and it indicates the direction in which the HD system is pointing at a given time. A layer labelled Vestibular Cell Layer 1 (VCL1) is also a 360-node layer that comprises of sigma-pi neurons. These are neurons that compute the sum of products of their inputs (see [10]). Each node in the VCL1 layer accepts input from all HD cells. In addition, each node in the VCL1 layer accepts input from a single cell (called the Vestibular Sensor Cell 1) that simulates the clockwise vestibular sense. Similarly, another layer (not shown for clarity) called Vestibular Cell Layer 2 also contains 360 sigma-pi neurons and it is connected to the HD layer in the same way as VCL1. This layer is associated with a cell that simulates the counter-clockwise vestibular sense (Vestibular Sensor Cell 2). In the same fashion, two more cell layers (KCL1 and KCL2 - also not shown) exist for the kinesthetic sense. In addition to the input from the HD cells each of these two layers accepts input from two cells correspondingly that simulate clockwise and counter-clockwise kinesthetic senses. Even though three HD cells are shown in the diagram above, only the connections to the middle one are explicitly drawn for clarity. Connections shown with solid lines are trained whereas those with dashed lines are not. Some of the weight connections are labelled but without their subscripts for clarity. **Right:** The LEGO® MINDSTORMS® NXT robot used for the experiments presented here. The robot is equipped with an on-board omnidirectional video camera (above the NXT brick) a gyroscopic sensor and an acceleration sensor (pictured to the left and right above the wheels). For locomotion, the robot uses two active wheels (seen at the front) and a dummy castor (not visible in the picture).

between distant cells negative (inhibitory). The weights w^{RC}, w^{VES} and w^{KIN} can be prescribed using Gaussian-like functions like for example in [21], [16] and [24] but this method is not biologically plausible and conveniently ignores the effects of noise to the weight values. Here, in order to train the network weights, data is collected from the robot as explained below.

3 Experimental Procedure and Results

The model described above was trained and tested using a LEGO® robot with an on-board omnidirectional video camera (see figure 2(right)). The robot is further

Fig. 3. Snapshot from the omni directional video camera on-board the robot

equipped with a gyroscopic sensor and an acceleration sensor. For locomotion, the robot uses two active wheels and a dummy castor. For the purposes of the work presented here, the visual input to the HD system was provided by processing the image from the omnidirectional video camera on the robot (see figure 3 for an example of a video image). The vestibular input was provided directly by the gyro sensor. One vestibular sensor cell (see figure 2(left)) was driven by the raw gyro signal and the other by the inverted version of the gyro signal. The two kinesthetic inputs were provided by differentiating the signals from the odometric sensors in the motors driving each wheel of the robot. The robot was controlled by a PC via USB connection in order to achieve maximum possible data transfer rates when reading the robot's sensors. The video during each recording session was independently recorded on the video camera and during post-processing it was time-corresponded with the robot's data. This setup allowed a 10Hz sampling rate in all data sources (gyro, motor position, video). Two sets of data were collected using the above setup. The first (training set) was used to train the network weights of the HD model and the second (the test set) was used to test the model.

3.1 Training the Model

The duration of the training set was 873 seconds (14.55 minutes). During this time the robot was programmed to continuously rotate in a random direction with a constant rotational speed of approximately 35 degrees/second. In order to obtain the expected output of the system at time t_n (i.e. training pattern p_n), initially, the direction of the robot (based on a world reference frame) was extracted from the video data by finding the maximum correlation (along the abscissa of the video image) between the video frame taken at t_n and the first captured video frame taken at t_0. As the robot was only rotating on the spot during the experiments described here, this provided a convenient way of establishing the world-based orientation of the robot. After that, p_n was created by translating a Gaussian function (of standard deviation σ) so that it would be centred at the the robot's orientation at t_n.

Training of the recurrent weights w^{RC} was achieved using the Hebbian learning rule:

$$\delta w_{ij}^{\mathrm{RC}} = k^{\mathrm{RC}} r_i^{\mathrm{HD}} r_j^{\mathrm{HD}} \tag{3}$$

where k^{RC} is the learning rate. Training of the vestibular weights was achieved using the following rule:

$$\delta w_{ijk}^{\mathrm{VES}} = k^{\mathrm{VES}} r_i^{\mathrm{HD}} \bar{r}_j^{\mathrm{HD}} r_k^{\mathrm{VES}} \tag{4}$$

where $\bar{r}^{\mathrm{HD}}$ is a historical trace value of r^{HD} and is given by:

$$\bar{r}_j^{\mathrm{HD}}(t + \delta t) = (1 - \eta) r^{\mathrm{HD}}(t + \delta t) + \eta \bar{r}_j^{\mathrm{HD}}(t) \tag{5}$$

A similar rule to the one given by equation 4 was also used in order to train the kinesthetic weights w_{ijk}^{KIN}.

Note that equation 4 (through the use of equation 5) considers together current and past values of the state of the HD system. It is this feature that gives rise to idiothetic weight profiles that bias the CANN attractor to move with the right speed and in the right direction according to the idiothetic inputs. In equation 5, η is a parameter that dictates the influence of the current and previous state of the network in the trace rule (equation 4) and δt is the time delay between the two states considered.

It can be seen from the equations defining the HD model presented here that there are many parameters (τ, θ^{RC}, w^{INH}, δt, η, β, etc.) that define the behaviour of the system. Importantly, these parameters are not always orthogonal to each other. As in previous work, these parameters (listed in table 1) have been obtained here by trial-and-error. It should be stressed however that once determined, the same parameter values were used in all the test cases presented below.

3.2 Testing the Model

Test data was collected for 102.9 seconds (1.715 minutes). During this session the robot was programmed to continuously rotate on the spot and in the same direction with a constant rotational speed of approximately 35 degrees/second. The same test data set was used for the three cases presented below.

Test Case 1. For this test case the visual input was turned off for $30s < t < 60s$ and $t > 90s$. In addition, the vestibular and kinesthetic inputs were also turned off for $t < 10s$ and $t > 80s$. These input suppressions were done manually by overwriting the data during the "off" windows by zeroes. Figure 4 shows a set of plots with common time scale that display the inputs (plots 3-5) and output (plot 1) to and from the HD model during this test case. Plot 2 compares the actual head direction (obtained from video as described above) and the model-predicted head direction.

Table 1. The parameter values used in the HD system model presented

τ	0.5s	δt	0.1s	β	0.2
k^{RC}	0.01	θ^{RC}	3	η	0.9
k^{VES}	0.01	θ^{VES}	0.0195	σ	20°
k^{KIN}	0.01	θ^{KIN}	0.0025	w^{INH}	0.4
		θ^{VIS}	30		

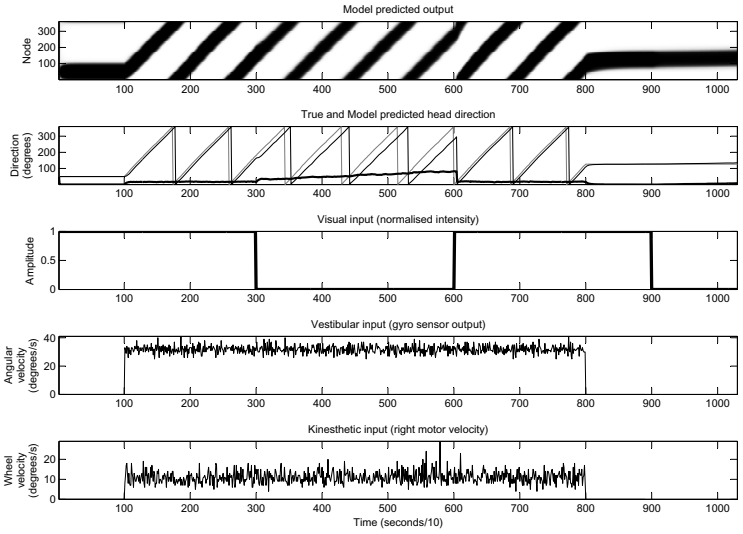

Fig. 4. Data describing test case 1. Plot 1: the model output. This is a surface plot viewed from the top. The black colour value is proportional to the excitation of each node. Plot 2: The robot's true head direction (grey line) and the model-predicted head direction (black line) with the absolute difference between them (wide black line). Note that the model-predicted head direction is the crest of the surface in plot 1. Plot 3: The normalised intensity of the visual input. When this is 0 the robot is in the dark (no visual input). Plot 4: The vestibular input (gyro sensor signal). Plot 5: The kinesthetic input (right motor velocity). Note that only the right motor velocity is plotted for clarity. The left motor velocity is very similar to this but in the opposite direction.

For $t < 10s$ the robot is stationary and only receives visual input ("lights on"). The attractor of the network is centred around 50°. This is where the robot "thinks" it is facing and this is in agreement with its true orientation. For $10s < t < 30s$ the robot moves under its own volition while the lights are still on. The expected orientation follows closely the actual orientation of the robot over a bit more than two full revolutions. For $30s < t < 60s$ the robot keeps rotating but in the dark (no visual input). The output still follows the true orientation of the robot but there is a small drift (with the model output lagging). At $t = 60s$ visual input is re-introduced (lights turn on). The small discrepancy between the actual and model-predicted orientation is quickly rectified as the model uses the visual cue to reset its state. For $60s < t < 80s$ the model continues to follow the true orientation of the robot. At $t = 80s$ the robot stops moving (no idiothetic inputs) and 10 seconds later the lights are also turned off. For $t > 80s$ the model-predicted and actual orientations of the robot are stationary and in agreement.

Test Case 2. During this test case the visual and vestibular inputs are the same as in test case 1 but the kinesthetic input is set to zero throughout. Figure 5 only

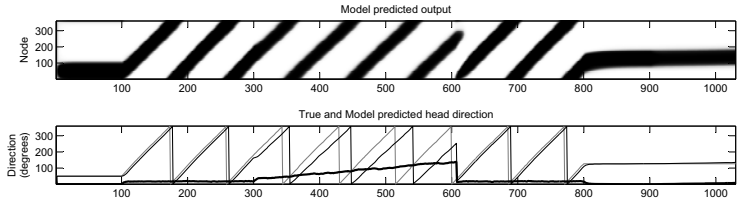

Fig. 5. Data describing test case 2. Plot 1: the model output. Plot 2: the true (grey line) and the model-predicted head direction (black line) with the absolute error (wide black line). Visual and vestibular inputs are the same as in plots 3 and 4 respectively in figure 4. The kinesthetic input is zero at all times and therefore not shown here.

shows the model-predicted output plot and the plot comparing the true and model-predicted head directions.

Since kinesthetic input is always zero here, a non-zero vestibular input implies that an external force is moving the robot. Initially the lights are on and the robot is first stationary ($t < 10s$) and then it is being moved ($10s < t < 30s$). The model follows the true orientation of the robot just like in test case 1 despite the conflict between the kinesthetic input and the other two senses (visual and vestibular)[2]. The same happens for $t > 60s$. The notable difference between this test and the previous one is visible in the interval $30s < t < 60s$. Here, it is as if the robot is being rotated in the dark. The drift between the model-predicted outpout and true orientation is greater than that in test case 1. This leads to a greater error at $t = 60s$ (when visual input is switched back on) but again this is quickly rectified by the system.

Test Case 3. In order to observe the effects of a persistent conflict between the visual and vestibular senses, during this test case the visual input was kept the same as before (see plot 3 in figure 4) but the vestibular input was (manually) inverted (see figure 6). The kinesthetic input was again kept at zero throughout this test case.

The most interesting observation here is during $10s < t < 30s$ and $60s < t < 80s$ when the visual and vestibular inputs are in persistent conflict. Even though the model mostly follows the true orientation of the robot, plot 1 in figure 6 is more informative of the state of the HD system during this test. The plot shows that there is a periodic effect during which the output is influenced more by the vestibular input than at other times. What happens when visual input is removed during $30s < t < 60s$ should not be surprising as the vestibular input is the only input to the system and therefore it solely drives the model-predicted orientation. Note that when the lights come on at $t = 60s$, it is only by coincidence that the expected head direction almost coincides with the actual head direction.

[2] Note that the word *conflict* in this text is used to imply any form disagreement between two inputs and not just one where two inputs oppose each other.

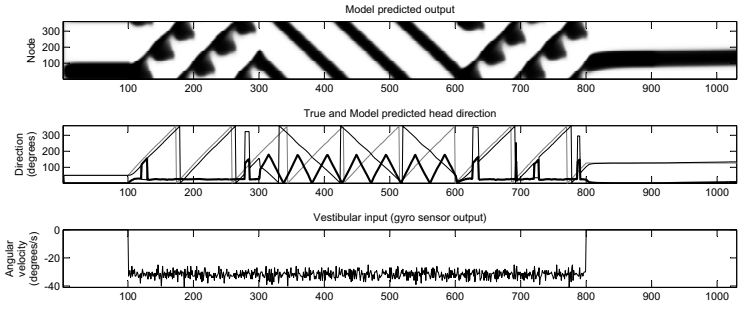

Fig. 6. Data describing test case 3. Plot 1: the model's output. Plot 2: true (grey line) and model-predicted head direction (black line) with the absolute error (wide black line). Plot 3: The vestibular input which is the inverted version of that shown in figure 4. Visual input is the same as in plots 3 of figure 4. The kinesthetic input is zero at all times and therefore not shown here.

4 Discussion

The qualitative observations in the test cases above are in agreement with biological observations made in experiments with animals (mainly rodents). More specifically it is shown that the visual input, when available, helps to set or reset the HD system to the true orientation ([17]). The visual input helps to maintain an absolute point of reference of the true orientation. When all inputs are in agreement the HD system follows the true orientation of the robot but when the visual input is turned off a slight drift is observed (see biological observations in [15]). The idiothetic inputs thus help to maintain a relative point of reference of the true orientation. In the absence of allothetic cues (i.e. visual input) a greater drift is observed when only vestibular input is used (i.e. the robot is being rotated by an external force) than when both idiothetic inputs are present (i.e. when the robot is self-rotating). Biological observations of this are presented in [5] who performed behavioural experiments with hamsters. When no input is present the HD system maintains a stable orientation. A very small drift of the output can be noticed in figures 4, 5 and 6 when all inputs are switched off ($t > 90s$). Observations of this drift and its time scales in rats are presented in [15]. It is speculated here that, in the case of a biological HD system, the synaptic weights cannot be perfectly symmetrical about each HD cell (noise, uneven number of connections etc). Similarly here, due to the training of the model using real (noisy) data from the robot the same effect is observed. Finally, when a persistent conflict is introduced between the visual and vestibular inputs, the output appears to follow the visual input but shows periodic tendency to follow the vestibular input. This effect is a consequence of the chosen values of the model parameters. The parameter values were not chosen to deliberately construct it but also, no attempts have been made to change or remove it by varying the model parameters. Conflicts between the visual and vestibular inputs to the HD system have been presented (after biological observations) in

quite a few places (summarised in [17]). To the author's best knowledge however, the temporal effects of conflicts such as the one presented in test case 3 have not been documented before. It would be interesting therefore to recreate such conflict conditions with real animals in order to determine the validity of the model presented here.

5 Summary and Conclusion

A biologically inspired model of the head direction system was presented and implemented on a small mobile robot. The model takes three inputs (visual, vestibular and kinesthetic) that are among those which provide the most influencing cues to the biological HD system [17]. The model was trained and tested using real data obtained from the robot. Three test cases were conducted and the results obtained were compared with biological observations made in previous documented work.

According to Trullier and colleagues ([20]) biological navigation methods may not always produce the best, most mathematically optimal solution to a navigation problem but they are fast, flexible and adaptive. What is best and most mathematically optimal however depends on the employer of a particular navigational skill. We know little about even the simplest of organisms in nature and it could therefore really be that for a particular organism, their navigation strategy is the best in all respects. Modelling biological mechanisms of navigation helps us understand better the remarkably complex systems in nature. Besides the information value however, the understanding of how these mechanisms evolved, rather than just what they do and how they do it, may lead us to more generalised principles of designing artificial navigation systems that would be the best and most optimal for their intended application.

References

1. Arleo, A., Gerstner, W.: Spatial orientation in navigating agents: Modeling head-direction cells. Neurocomputing 38-40(1-4), 1059–1065 (2001)
2. Biegler, R., Morris, R.G.M.: Landmark stability is a prerequisite for spatial but not discrimination learning. Nature 361, 631–633 (1993)
3. de Castro, L.N., Von Zuben, F.J. (eds.): Recent Developments in Biologically Inspired Computing. Idea Group Publishing, USA (2005)
4. Degris, T., Lachèze, L., Boucheny, C., Arleo, A.: A spiking neuron model of head-direction cells for robot orientation. In: Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior, from Animals to Animats, pp. 255–263. MIT Press, Cambridge (2004)
5. Etienne, A.S., Maurer, R., Saucy, F.: Limitations in the assessment of path dependent information. Behaviour 106, 81–111 (1988)
6. Etienne, A.S., Maurer, R., Séguinot, V.: Path integration in mammals and its interaction with visual landmarks. Journal of Experimental Biology 199, 201–209 (1996)

7. Goodridge, J.P., David, Touretzky, D.S., Jeremy, P.: Modeling attractor deformation in the rodent head-direction system. Journal of Neurophysiology 83, 3402–3410 (2000)
8. Goodridge, J.P., Dudchenko, P.A., Worboys, K.A., Golob, E.J., Taube, J.S.: Cue control and head direction cells. Behavioral Neuroscience 112(4), 749–761 (1998)
9. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B.B., Moser, E.I.: Microstructure of a spatial map in the entorhinal cortex. Nature 436(7052), 801–806 (2005)
10. Mel, B.W., Koch, C.: Sigma-pi learning: on radial basis functions and cortical associative learning. In: Touretzky, D.S. (ed.) Advances in Neural Information Processing Systems, vol. 2, pp. 474–481. Morgan Kaufmann Publishers Inc., San Francisco (1990)
11. Morris, R.G., Garrud, P., Rawlins, J.N., O'Keefe, J.: Place navigation impaired in rats with hippocampal lesions. Nature 297(5868), 681–683 (1982)
12. Muller, R.U., Kubie, J.L., Ranck, J.B.: Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. Neuroscience 7(7), 1935–1950 (1987)
13. O'Keefe, J., Dostrovsky, J.: The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. Brain Research 34(1), 171–175 (1971)
14. Redish, A.D., Elga, A.N., Touretzky, D.S.: A coupled attractor model of the rodent head direction system. Network: Computation in Neural Systems 7(4), 671–685 (1996)
15. Mizumori, S.J., Williams, J.D.: Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. Neuroscience 13(9), 4015–4028 (1993)
16. Stringer, S., Trappenberg, T., Rolls, E., de Araujo, I.: Self-organizing continuous attractor networks and path integration: one-dimensional models of head direction cells. Network: Computation in Neural Systems 13(2), 217–242 (2002)
17. Taube, J.S.: Head direction cells and the neurophysiological basis for a sense of direction. Progress Neurobiololy 55(3), 225–256 (1998)
18. Taube, J., Muller, R., Ranck Jr., J.: Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantitative analysis. Neuroscience 10(2), 420–435 (1990)
19. Tolman, E.C., Ritchie, B.F., Kalish, D.: Studies in spatial learning. i. orientation and the short-cut. Journal of Experimental Psychology 36, 13–24 (1946)
20. Trullier, O., Wiener, S., Berthoz, A., Meyer, J.: Biologically-based artificial navigation systems: Review and prospects. Progress in Neurobiology 51, 483–544 (1997)
21. Skaggs, W.E., Knierim, J.J., Kudrimoti, H.S., McNaughton, B.L.: A model of the neural basis of the rat's sense of direction. Advances in Neural Information Processing Systems 7, 173–180 (1995)
22. Wehner, R., Menzel, R.: Do insects have cognitive maps? Annual Review of Neuroscience 13, 403–414 (1990)
23. Wiener, S.I., Taube, J.S. (eds.): Head direction cells and the neural mechanisms of spatial orientation. MIT Press, Cambridge (2005)
24. Zeidman, P., Bullinaria, J.A.: Neural models of head-direction cells. In: French, R.M., Thomas, E. (eds.) From Associations to Rules: Connectionist Models of Behavior and Cognition, pp. 165–177 (2008)