An implementation of the path integrator mechanism of head direction cells for bio-mimetic navigation.

Ankur Sinha and Jack Wang

Abstract—Head direction cells are thought to be an integral part of the neural navigation system. These cells track the agent's current head direction irrespective of the host's location. In doing so, they process a combination of inputs: angular velocity and visual inputs are major effectors; to correctly encode the agent's current heading. There are close to fifteen models of head direction cell systems found in literature today. Very few of these models have been implemented for biomimetic navigation in robots. In this paper, we describe an implementation of the head direction cell system on the robot operating system (ROS) robotic platform as a first step towards a bio-mimetic navigation system for Willow Garage's personal robot 2 (PR2) robot.

I. INTRODUCTION

AVIGATION is a capability animals must necessarily possess to survive. It is well known that even smaller animals, such as rats, have sufficiently well developed navigation systems that enable them to perform tasks such as foraging and finding shelter, which are a necessity. The biological navigation systems possessed by these animals are quite different from the navigation systems that are employed in classical robotics. Classical robotics usually employs mathematical methods that are based on nautical navigation. These methods are generally robust, complicated and maladaptive in that they attempt to track the agent's exact spatial parameters at all times. On the other hand, reviews of biological navigation techniques conducted by Trullier and colleagues [1] and Franz & Mallot [2] document them to be highly adaptive, layered and sufficient. As an example, at simpler levels of biological navigation, an animal does not need to know its exact location and can successfully navigate to a goal location with reference to a prominent landmark in the immediate environment.

In spite of the advantages that biological navigation brings, bio-mimetic systems are seldom used in robotics. Even when they are, they are bio-inspired: they pick certain properties from biological systems, but they do not completely model them. In the past, a major reason behind this was the lack of information about the underlying neural processes that make biological navigation possible. In the past decades, however, the discovery of neurons that are related to navigation: place cells [3], head direction cells [4], [5], grid cells [6] and others have increased our understanding of the subject to such an extent that a cognitive map theory of biological navigation has been proposed [7], [8], [9], [10], [11], [12]. The theory states that via the associations between neurons with specific behaviours, animals maintain a "cognitive map" of their environment.

In this paper, we detail an implementation of a model of the path integrator dynamics of head direction cells proposed by Stringer et al. [13] on the ROS robotic platform [14] as a starting point towards a head direction, place and grid cell based bio-mimetic path integrator navigation system. We also document minor extensions to the original model that improve the navigation system.

In section I-A, we provide an introduction to head direction cells and briefly discuss their computational modelling in section I-B. We then detail our model in section II. In section III, we present our results and briefly discuss challenges and our future work plans in section IV. Finally, we summarize and conclude in section V.

A. Head direction cells

Since their discovery by Ranck [5], head direction cells have been subjected to a great amount of research. Head direction cells are neurons that maintain a firing response that encodes the agent's current head direction. Each individual neuron in a head direction cell system fires maximally when the agent faces a particular direction. This is referred to as the *preferred direction* of the particular neuron. A set of such neurons, the preferred directions of which together encompass all 360 degrees, therefore, acts like an internal compass, tracking the head direction of the animal as it moves around. For detailed information on the properties of head direction cells, we refer the reader to Taube's comprehensive reviews [15], [16], [17].

Head direction cells possess two navigational capabilities. First, they exhibit path integrator properties similar to an inertial navigation system (INS) [18]: head direction cells integrate angular velocity information to calculate the agent's current heading [19], [20]. Second, head direction cells have been found to anchor to salient landmark cues [21], [22], [23], [24] which makes it possible for them to contribute to a landmark based navigation system.

Path integrator systems inherently suffer from drift. Since the value at each iteration is calculated on the basis of the result at the previous iteration, the error component in their outputs is also integrated at each iteration. This implies that the error in their calculations continuously increases over time. In order to limit the error such that the computed value is usable, inertial navigation systems are generally coupled with other absolute navigation systems: landmark systems [25] or global positioning systems (GPS) [26], for example. The landmark or GPS systems serve to reset the

Ankur Sinha and Jack Wang are with the Faculty of Engineering and Information Technology, The University of Technology, Sydney, Australia (email: ankur.sinha@student.uts.edu.au, jianguo.wang@uts.edu.au).

We're most grateful to Dr. Xun Wang at the "Magic" Lab, The University of Technology, for his inputs and assistance with the PR2 robot.

system to an accurate value at regular intervals, ensuring that the drift is kept in check. Head direction cells, similarly, suffer from drift in the absence of visual cues and are reset when a familiar, previously observed landmark is sighted again by the agent [17].

B. Computational modelling of head direction cells

Various computational models of the head direction cell system have been proposed in literature. These models are identical in that they all attempt to replicate the biological data that is known about head direction cells: firing rate profiles, neural processing, underlying architecture.

With the discovery of more and more information on head direction cells, computational models have also become more detailed. Early models only concentrated on how the path integrator mechanism might process angular velocity information to calculate the current heading. McNaughton et al. proposed a look up table based method as the underlying mechanism [27]. Skaggs et al. replaced the look up table with a ring attractor where correctly set up synaptic connections between neuron sets would permit a single packet of activity to encode head direction [28]. The discovery of predictive head direction cells in the anterior thalamic nucleus first led Blair and Sharp to model these neurons on the NEU-RON [29] simulator [20]. In their model, they demonstrated how the predictive head direction neurons lead current head direction neurons as had been observed. Similarly, Redish et al. proposed a coupled attractor model that also featured both predictive and current head direction neurons [30]. Their model employed more than one set of synaptic connections between neurons sets. Zhang also put forward an analytical model that used multiple synaptic connections between neurons sets [31]. The implication, that quick synaptic changes are required for the functioning of the neuron sets deems both these models less biologically plausible.

All of the above mentioned models used synaptic weights that were predefined. Stringer et al. proposed a system based on Hebbian learning to explain how the synapses may be set up by self organization [13]. They only applied the method to model a path integrator system. Kyriacou extended the model to include kinesthetic inputs [32]: he implemented the model on a robot and used an omnidirectional camera to simulate visual inputs in the system. In another work, Kyriacou documents how an evolutionary algorithm [33] may be used to ascertain the parameters that controlled the dynamics of such a network [34].

Other models include but are not limited to: Arleo & Gerstner's implementation on the mobile Khepera robot [35]; Song & Wang's implementation that used a lower level spiking neuron architecture [36]; Degris and colleagues' implementation of Song & Wang's work on to the Pekee robot as part of the Psikharpax project [37]; Goodridge & Touretzky's model of the anticipatory characteristics of head direction cells that did not use an attractor network following observations that GABA containing interneurons that are necessary for attractor formation were absent in the anterior dorsal nucleus [38]; and Zeidman & Bullinaria's extension



Fig. 1: The head direction cell set schematic.

to Goodridge & Touretzky's model to include optic flow information [39].

II. METHODS: THE MODEL

A. Structure

Our model is based on the self organizing model proposed by Stringer et al. [13]. Their original model presented a biologically plausible method of setting up synaptic weights in the head direction cell system, via Hebbian learning [40]. It only applied the method to the path integrator system of head direction cells and did not discuss visual inputs and landmark navigation. We've incorporated vision cells into their model in an attempt to extend it to project both vestibular and visual inputs on the head direction cell set as has been found.

The head direction cell system here makes use of three sets of neurons as shown in the Figure 1. The head direction cell set is a fully connected recurrent network, i.e., each neuron HD_i is connected to every other neuron HD_j via synapses, the weights of which are denoted as w_{ij}^{HD} . Here, HD_i and HD_j are the *pre-synaptic* and *post-synaptic* neurons with firing rates r_i^{HD} and r_j^{HD} respectively. The head direction cells, therefore, form a one dimensional circular array. While cells with adjacent preferred directions appear next to each other in our implementation (and most models), this is not how head direction cells are found in the brain. In fact, as long as the synaptic connections are set up appropriately, the location of neurons is irrelevant to the functioning of the system.

Rotation cells ROT_k , with firing rates r_k^{ROT} , are also connected to each head direction cell via synapses denoted $w_{ijk}^{HD_ROT}$. Here, these synapses are not simply synapses between a singular head direction cell and a rotation cell. Rather, they're *effective synapses* between the pre-synaptic neuron HD_i , the post-synaptic neuron HD_j and the rotation cell ROT_k . The two rotation cells in Figure 1 fire depending upon clockwise and anti-clockwise rotation inputs respectively as shown in Figure 6.

The visual cell set represents an abstraction of the visual processing system and each visual cell VIS_l , with firing rate r_l^{VIS} , is connected to every head direction cell HD_i via synapses $w_{il}^{HD_vVIS}$.

B. Dynamics

Our model is a firing rate based model, which provides a level of detail that is sufficient for our system. The activation, h_i^{HD} , of each head direction cell HD_i is given at time t by:

$$\tau \frac{\mathrm{d}h_{i}^{HD}(t)}{\mathrm{d}t} = -h_{i}^{HD}(t) + \frac{\phi_{0}}{C^{HD}} \sum_{j} ((w_{ij}^{HD} - w^{INH})r_{j}^{HD}(t)) + \frac{\phi_{1}}{C^{HD\times ROT}} \sum_{jk} (w_{ijk}^{HD-ROT}r_{j}^{HD}(t)r_{k}^{ROT}(t)) + \frac{\phi_{2}}{C^{HD\times VIS}} \sum_{jl} (w_{jl}^{HD-VIS}r_{l}^{VIS}(t))$$
(1)

Here, τ is the time constant while ϕ_0 , ϕ_1 , ϕ_2 , C^{HD} , $C^{HD\times ROT}$, $C^{HD\times VIS}$ and w^{INH} are tunable parameters. These parameters control the effect the respective inputs have on the head direction cell attractor. w^{INH} represents global inhibition that the GABAergic interneurons exert on the system. The combination of local excitation of head direction cells and the global inhibition gives the system continuous attractor characteristics [41]. We use a Euler stepper method to integrate equation (1).

The firing rate of each head direction neuron is a sigmoid function of its activation:

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

where α and β are constants. Figure 5 shows firing rate profiles exhibited by the head direction cell set during a test run. Due to the regular learning employed in this implementation, the firing rates of all head direction cells are similar. This isn't the case in biology, where the firing rates of head direction cells vary from one to another.

The synapses between all neuron sets are set up using Hebbian learning:

$$\Delta w = k \times (r^{post} \times r^{pre}) \tag{3}$$

Here, Δw is the change in synaptic weight. k is the *learning* rate of the synapse. r^{pre} and r^{post} are the firing rates of the pre-synaptic and post-synaptic neurons respectively. This learning rule does not, however, include synaptic depression, or bounding of synaptic weights. We use a competition based normalization rule to bound our synapses:

$$\hat{W} = \frac{W}{|W|} \tag{4}$$

Here, W is a matrix containing the individual synaptic weights between neurons of the form:

$$W = \begin{pmatrix} w_{0,0} & w_{0}, 1 & \cdots & w_{0,n} \\ w_{1,0} & w_{1,1} & \cdots & w_{1,n} \\ \vdots & \vdots & \ddots & \vdots \\ w_{n,0} & w_{n,1} & \cdots & w_{n,n} \end{pmatrix}$$

|W| is the norm of the W matrix and \hat{W} is the normalized synaptic weight. It is worth noting that the above



Fig. 2: w_{ii}^{HD} values after training.

normalization departs from the Hebbian learning requirement of *locality* [42] - when a normalization method such as (4) is used, the final values of the individual synaptic connections between the neuron pairs become dependent on all the synaptic strengths in the complete neuron set. However, the classical Hebbian rule states that the synaptic weights between two neurons should only be affected by the activity of the two concerned neurons. We briefly discuss other formulations of Hebbian learning in section IV.

III. EXPERIMENTAL PROCEDURE AND RESULTS

We implemented the model based on the ROS [14] platform which provides support for a number of robots, including the PR2. ROS provides underlying tools that enable us to develop offline simulations and then move the code asis on to the robot. For development and testing, we collected data bags from the inertial measurement unit (IMU) sensors of the PR2 robot to run our simulations. We used a hundred head direction cells to cover the 360° direction space. We used two rotation cells, one each for clockwise and anti clockwise rotation, and a single visual cell (Figure 1). The values of constants used in our implementation are given in table I.

The system runs in three phases:

A. Setting up of synaptic weights to appropriate values

During this first phase, we set up the synaptic weights in the network to their appropriate values. The network is initialized with all synaptic weights as zero, implying that no learning or association has taken place between the sets of neurons. In order to set up both the internal head direction cell synapses w_{ij}^{HD} and the effective rotation synapses w_{ijk}^{HD-ROT} , we simulate rotation in the system in both clockwise and anti-clockwise directions. Each head direction cell is assigned a preferred direction such that the set encompasses the complete 360° range: 0°, 3.6°, ..., 356.4°. As done previously in literature, we model the head direction cell firing profile as a Gaussian:

$$r_i^{HD} = exp(-\frac{1+\Delta S^2}{2\sigma^{HD^2}}) \tag{5}$$

 ΔS is the angular distance between the current head direction and the head direction cell's preferred direction:

$$\Delta S = \min(|x|, |360 - x|) \tag{6}$$

where, for each neuron HD_i with preferred direction $x_i^{preferred}$, for a head direction θ

$$x = \theta - x_i^{preferred} \tag{7}$$

 σ^{HD} is a constant that controls the width of the Gaussian profile, and consequently, controls the angular width that a head direction cell is active in. Note that equation (5) is only used to calibrate the synapses to their correct values. Equation (2) is used thereafter.

In order to calibrate the network's synapses, one of the rotation cells is activated, simulating either clockwise or anticlockwise rotation. Simultaneously, the firing rate profile, as obtained by equation (5), is simulated such that each head direction cell fires maximally, i.e., the system is simulated to face the preferred direction of each head direction neuron in the set, one neuron after the other. The firing rate profile of the head direction neuron set shifts according to (5). The simultaneous firing of the neuron sets modifies the synapses between them according to the Hebbian learning rule discussed in section II.

Since we permit each neuron to fire maximally only once during both the clockwise and anti-clockwise iterations, our training method is uniform. The uniformity can be seen in Figure 2, which shows the recurrent synaptic weights between head direction cells, w_{ij}^{HD} , before normalization. Observe that normalization will not modify the nature of the graph. Also note that as a result of Hebbian learning from the simulated Gaussian firing rate profiles, cells with preferred directions near each other have stronger synaptic connections than cells that are far apart, as is expected.

B. Initializing the network to an initial direction

Once the synaptic weights are set up appropriately, a packet of activity must be stimulated in the attractor. This packet of activity is the initial or reference heading of the agent. An initial packet of activity is forced on the system by projecting the required profile on to the head direction cell attractor from the visual input for a short period of time. The firing of the single vision cell can be simulated by simply setting its firing rate to the maximum value, 1, in equation (1):

$$r_1^{VIS} = 1 \tag{8}$$

Further, to set the initial direction to the preferred direction of head direction cell *i*, we simply set the synapses between the vision cell and head direction cells, $w_{1,j}^{HD_-VIS}$, to the synaptic weight between head direction cell *i* and every other head direction cell *j*, w_{ij}^{HD} . This is equivalent to setting up the association between the visual feature cell and the head direction cell set using equation (3).



Fig. 3: Stabilization of initial activity packet



Fig. 4: Sub-optimal parameters do not result in a stable activity packet

$$w_{1,j}^{HD-VIS} = w_{ij}^{HD} \tag{9}$$

Since w_{ij}^{HD} has already been set up in a way that will cause the set of head direction neurons to form a packet of activity peaking at head direction neuron *i* during the training step, the assignment plainly ensures that projections from the visual cell will form a similar activity packet that also peaks at the chosen head direction neuron. The particular neuron is *associated* to the visual feature cell in this way. Note that this method cannot be extended to a set of multiple visual feature cells that will project a firing rate profile instead of a single projection. In such a case, an association will have to be made by simulating the set of visual feature cells and head direction cells similar to the method described in the previous section.

During our experiments, we observed that forcing an initial packet of activity did not guarantee a functional system. The attractor should maintain the packet of activity in the absence of external inputs. The parameters ϕ_0 and w^{INH} that effect the recurrent behaviour of the network must be fine tuned to ensure that the activity packet stabilizes as shown in Figure 3, the other outcome being the activity packet *flattening out* as shown in Figure 4.¹

¹We are most grateful to Dr. Simon Stringer for his input on this subject.



Fig. 5: Firing rates, r_i^{HD} , obtained from equations 1 and 2 during test run.

The continuous attractor, if set up properly, permits the packet of activity to lie in a state of neutral equilibrium, like a ball lying on a perfectly horizontal table surface.

C. Running the system with angular velocity data

Figure 6 shows the behaviour of the network when tested with angular velocity data. The firing rates of the two rotation cells are a sigmoid function of the angular velocity inputs similar to equation (2). The figure shows that the head direction indeed responds to angular velocity inputs. Of special interest is the graph's behaviour at time=1300 which shows circular attractor nature of the network - since synaptic weights between head direction neurons with preferred directions 356.4° and 0° have been trained such that these two neurons are adjacent, the firing profile translates seamlessly off the "end" of our one dimensional attractor to its "beginning". This is a consequence of equation (6) which ensures that the head direction neuron with preferred direction 356.4° has the same synaptic connection with its two adjacent cells that have 352.8° and 0° as their preferred directions respectively. This confirms the continuous attractor nature of the network. Figure 4 also illustrates the ring nature of the one dimensional attractor network - it shows an initial packet of activity where the firing rate peaks at head direction neuron 0 and symmetrically reduces in both directions. We tested the system with several bags of data and observed encouraging results.

The rate at which the head direction firing profile moves depends on the strength of the rotation neurons' projections on it. In the current configuration, the system has not been optimised to correctly map the rotations of the agent, i.e., the movement of the head direction activity profile does not reflect the true rotation of the agent in the world frame. This isn't because it cannot be done: the accuracy can be improved fine tuning the value of ϕ_1 which controls the strength of the rotation neurons' projections on to the head direction cells. However, as we briefly discuss in the next sub section, our current work focusses on associating head direction cells to salient features in the environment for landmark navigation, and since this will function as a drift correction mechanism in itself, we've permitted the system to drift and have put off

TABLE I: Constants used in the implementation

Number of head direction cells	100
Number of rotation cells	2
Number of vision cells	1
α	1.5
β	3
ϕ_0	1000
ϕ_1	2000
ϕ_2	1000
C^{HD}	100
$C^{HD \times ROT}$	200
$C^{HD \times VIS}$	100
k	1
w^{INH}	0.02
σ^{HD}	10

optimization of these constants.

D. Correcting drift using salient visual cues: preliminary tests

Figure 7 shows the results of a simulation where the head direction cell set is assumed to incur drift and, a projection from the visual feature cell, which would be caused by the agent observing a familiar feature, is used to correct this drift. The mechanism behind this is similar to the mechanism used to initialize the network as described in section III-B. In order for the visual inputs to override the activity packet maintained by path integration, the projections must be strong enough. We discuss our ideas on using multiple visual features in the next section.

IV. DISCUSSION

Researchers have attempted to implement bio-inspired navigation systems in the past, although the level of inspiration and the extent of implementation has varied. The Psikharpax project [43], [44], for example, attempted to create an artificial rat and implemented head direction and place cells as its basis for navigation. RatSLAM [45], on the other hand, uses similar "pose cells" to propose a solution to the simultaneous localization and mapping (SLAM) [46] problem. (For more examples, see Franz & Mallot's review [2].)

Bio-mimetic navigation provides two closely related research areas. The first is to improve our understanding of biological navigation. This is done at different levels, for example, via behavioural studies or neuron recordings. The second, computational modelling, serves as a tool to verify collected information and proposed theories, while providing alternative navigation systems that can, in the future, be used in robotics. Even though bio-mimetic systems are not yet considered mature enough for use in task oriented robotics ahead of classical navigation techniques, it is accepted that even smaller animals such as ants and rats possess navigational capabilities that are superior to classical robotic navigation techniques. To be able to reliably mimic these biological methods would be a great stride in the field of autonomous navigation.



Fig. 6: Test runs with angular velocity data from the PR2 robot

In spite of the complexity of the underlying neural systems that makes bio-mimetic systems difficult to implement, systems that achieve a high level of similarity can and have be designed. The model of the head direction system implemented here, for example, deviates from known information on head direction cells in certain aspects but does still sufficiently carry out its intended navigational function.

Challenges and future work

The implementation of the path integrator half of the head direction cell system is the starting point in our attempt to develop a bio-mimetic navigation system for the PR2 robot using the ROS platform. The information this provides is not yet sufficient to carry out actual navigation. In the absence of a set of salient visual cues, the system will continue to drift as it runs, like any other INS [25]. An INS is generally coupled with other input sources, such as GPS [26], that reset the accumulated drift at regular intervals. The head direction system can similarly correct drift by detecting familiar visual features in the environment as was briefly demonstrated in section III-D. Our next goal, is therefore, to associate the head direction cell system to environmental cues that will reset the head direction system to its associated direction whenever they are observed. Since the system will associate with a set of cues, it will be capable of approaching these cues. The agent will be able to carry out the local navigation strategies: search, direction-following, aiming and guidance, as enumerated by Franz & Mallot [2].

A further goal of ours is to include other known neuron sets that are related to navigation: grid cells [6], place cells [3] and boundary cells [47]. Grid cells form a regular triangular grid which acts as a spatial map to encode the agents location. The neuron set uses path integrator mechanisms similar to head direction cells². Place cells too encode the agent's location and have been found to associate with visual features similar to head direction cells. However, they form place fields instead of a regular grid like map that is exhibited by grid cells. The last, boundary cells, respond to environmental boundaries. Implementing a system that includes all these neuron sets will provide obstacle, location and heading information to the agent along with some information on visual features that they associate with. This information is enough for a simple landmark based system.

An important part of the model is the Hebbian learning rule mentioned in equation (3). The formula that we've used currently is the simplest mathematical formulation of a Hebbian synaptic modification rule. As is visible, this rule does not provide for synaptic saturation. As long as the presynaptic and post-synaptic neurons fire simultaneously, the synapse between them will continue to strengthen. While this formulation covers strengthening of synapses by long term potentiation (LTP) [48], [49], it does not implement the flip side: long term depression (LTD) [50]. Just as LTP causes strengthening of synapses when presynaptic and postsynaptic neurons fire nearly simultaneously, LTD causes weakening of synapses if such simultaneous firing does not occur. The inability of the learning rule to provide for synaptic saturation and weakening makes it less biologically plausible. It also makes the implementation more difficult: if the synapses do not saturate at a known value, it is difficult to ascertain constant values for parameters that control the projection of inputs on to head direction cells: ϕ_0 , ϕ_1 , ϕ_2 . During

²Our paper detailing the addition of grid cells to this head direction cell model has been accepted at ICANN 2014 to be presented at the conference in Hamburg, Germany from September 15 - 19, 2014.



Fig. 7: Correcting drift using a single visual cue. (Note that the graphs for t=40 and t=50 overlap)

our simulations, we discovered that unbounded synapses constitute one of the cases where the projections on the head direction cells increase to such an extent that all neurons begin to fire maximally. Since the maximum firing rate in our model is the same for all head direction cells, this also causes the firing rate profile to flatten out and the system ceases to provide a peak that denotes the current head direction.

A number of formulations of the Hebbian rule have been proposed in literature [51], [52]. Such rules incorporate modifications to provide for LTD and synaptic saturation. One such modification is to normalize the synaptic weights, as shown in equation (4). Normalization is a competition based method: if synaptic efficacy increases, it must be at the expense of other synapses [42]. Other formulations include gating of the synaptic changes by either presynaptic or postsynaptic activity. While we did attempt to use gated rules with saturation in our model, we were unable to find a set of parameters that provided the required dynamics. Since the recurrent synapses between the head direction cells and the synapses between rotation cells and head direction cells remain largely unchanged during the running of the system, we decided to use the normalization rule and fix the synaptic weights after initial calibration of the system. For association with visual features, however, the formulation that is used to modify synapses between the head direction and visual cells will need to incorporate weakening of synapses via LTD so that the system can disassociate with (forget) features that are no longer present in the environment over a period of time.

The inclusion of multiple visual features presents a challenge also. In order for the system to associate with more than one feature, our implementation of visual cells should be able to differentiate between features so that they can each be encoded uniquely by our population of vision cells. We are yet to decide on a method to obtain this since modelling the biological visual system may be too complex for our purposes.

V. SUMMARY AND CONCLUSION

An implementation of a model of the head direction cell system on the ROS platform was detailed in this paper.

The implementation integrates angular velocity inputs using neural mechanisms to track the orientation of the robot. Certain extensions to the original model required to enable integration of visual inputs were also discussed. Visual inputs provide sensory information that is used to correct drift in the path integrator mechanism. The anchoring of head direction cells to visual landmarks makes it possible for the system to also provide a feature based navigation method. While a landmark navigation system will be limited to navigating between locations that are visible to the agent, they can be coupled with more complex map based methods to provide a layered navigation system. Various challenges and their tentative solutions were discussed as work to be done in the future to follow up the path integrator system.

Biological navigation follows a layered approach. The navigation technique employed depends on the task to be completed by the organism. For example, it is inefficient to use metric mapped navigation to approach a goal state that is visible to an agent. In such cases, simple navigation techniques such as piloting or landmark navigation would be quicker and more efficient. With this implementation of the path integrator dynamics of a head direction cell ensemble, we move closer to developing a landmark based navigation system as the simplest component of a layered bio-mimetic system for the PR2 robot.

REFERENCES

- O. Trullier, S. I. Wiener, A. Berthoz, and J.-A. Meyer, "Biologically based artificial navigation systems: Review and prospects," *Progress* in neurobiology, vol. 51, no. 5, pp. 483–544, 1997.
- [2] M. O. Franz and H. A. Mallot, "Biomimetic robot navigation," *Robotics and autonomous systems*, vol. 30, no. 1, pp. 133–153, 2000.
- [3] J. O'Keefe, "Place units in the hippocampus of the freely moving rat," *Experimental neurology*, vol. 51, no. 1, pp. 78–109, 1976.
- [4] R. U. Muller, J. B. Ranck, and J. S. Taube, "Head direction cells: properties and functional significance," *Current opinion in neurobiol*ogy, vol. 6, no. 2, pp. 196–206, 1996.
- [5] J. Ranck Jr, "Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats," in *Society for Neuroscience Abstracts*, vol. 10, 1984.
- [6] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser, "Microstructure of a spatial map in the entorhinal cortex," *Nature*, vol. 436, no. 7052, pp. 801–806, 2005.
- [7] R. U. Muller, M. Stead, and J. Pach, "The hippocampus as a cognitive graph." *The Journal of general physiology*, vol. 107, no. 6, pp. 663– 694, 1996.
- [8] B. L. McNaughton, F. P. Battaglia, O. Jensen, E. I. Moser, and M.-B. Moser, "Path integration and the neural basis of the cognitive map", *Nature Reviews Neuroscience*, vol. 7, no. 8, pp. 663–678, 2006.
- [9] O. Trullier and J.-A. Meyer, "Animat navigation using a cognitive graph," *Biological Cybernetics*, vol. 83, no. 3, pp. 271–285, 2000.
- [10] J. O'Keefe, "The hippocampal cognitive map and navigational strategies." 1991.
- [11] L. Nadel and L. MacDonald, "Hippocampus: Cognitive map or working memory?" *Behavioral and neural biology*, vol. 29, no. 3, pp. 405– 409, 1980.
- [12] J. O'Keefe and L. Nadel, *The hippocampus as a cognitive map*. Clarendon Press Oxford, 1978, vol. 3.
- [13] S. Stringer, T. Trappenberg, E. Rolls, and I. E. T. Araujo, "Selforganizing continuous attractor networks and path integration: onedimensional models of head direction cells," *Network: Computation in Neural Systems*, vol. 13, no. 2, pp. 217–242, 2002.
- [14] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, 2009.

- [15] J. S. Taube, "Head direction cells and the neurophysiological basis for a sense of direction," *Progress in neurobiology*, vol. 55, no. 3, pp. 225–256, 1998.
- [16] S. I. Wiener and J. S. Taube, *Head Direction Cells and the Neural Mechanisms of Spatial Orientation (Bradford Books)*. The MIT press, 2005.
- [17] J. S. Taube, "The head direction signal: origins and sensory-motor integration," Annu. Rev. Neurosci., vol. 30, pp. 181–207, 2007.
- [18] B. Barshan and H. F. Durrant-Whyte, "Inertial navigation systems for mobile robots," *Robotics and Automation, IEEE Transactions on*, vol. 11, no. 3, pp. 328–342, 1995.
- [19] H. T. Blair, B. W. Lipscomb, and P. E. Sharp, "Anticipatory time intervals of head-direction cells in the anterior thalamus of the rat: implications for path integration in the head-direction circuit," *Journal* of neurophysiology, vol. 78, no. 1, pp. 145–159, 1997.
- [20] H. T. Blair and P. E. Sharp, "Anticipatory head direction signals in anterior thalamus: evidence for a thalamocortical circuit that integrates angular head motion to compute head direction," *The journal of neuroscience*, vol. 15, no. 9, pp. 6260–6270, 1995.
- [21] J. P. Goodridge and J. S. Taube, "Preferential use of the landmark navigational system by head direction cells in rats." *Behavioral Neuroscience; Behavioral Neuroscience*, vol. 109, no. 1, p. 49, 1995.
- [22] J. J. Knierim, H. S. Kudrimoti, and B. L. McNaughton, "Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells," *Journal of neurophysiology*, vol. 80, no. 1, pp. 425–446, 1998.
- [23] —, "Place cells, head direction cells, and the learning of landmark stability," *The journal of neuroscience*, vol. 15, no. 3, pp. 1648–1659, 1995.
- [24] M. B. Zugaro, A. Arleo, C. Djean, E. Burguiere, M. Khamassi, and S. I. Wiener, "Rat anterodorsal thalamic head direction neurons depend upon dynamic visual signals to select anchoring landmark cues," *European Journal of Neuroscience*, vol. 20, no. 2, pp. 530–536, 2004.
- [25] S. Sukkarieh, "Low cost, high integrity, aided inertial navigation systems for autonomous land vehicles," Ph.D. dissertation, The University of Sydney, 2000.
- [26] S. Panzieri, F. Pascucci, and G. Ulivi, "An outdoor navigation system using gps and inertial platform," *Mechatronics, IEEE/ASME Transactions on*, vol. 7, no. 2, pp. 134–142, 2002.
- [27] B. L. McNaughton, L. Chen, and E. Markus, ""dead reckoning", landmark learning, and the sense of direction: A neurophysiological and computational hypothesis," *Journal of Cognitive Neuroscience*, vol. 3, no. 2, pp. 190–202, 1991.
- [28] W. E. Skaggs, J. J. Knierim, H. S. Kudrimoti, and B. L. McNaughton, "A model of the neural basis of the rat's sense of direction," in Advances in Neural Information Processing Systems 7: Proceedings of the 1994 Conference, vol. 7. MIT Press, 1995, p. 173.
- [29] M. L. Hines and N. T. Carnevale, "The neuron simulation environment," *Neural Computation*, vol. 9, no. 6, pp. 1179–1209, 1997.
- [30] D. A. Redish, A. N. Elga, and D. S. Touretzky, "A coupled attractor model of the rodent head direction system," *Network: Computation in Neural Systems*, vol. 7, no. 4, pp. 671–685, 1996. [Online]. Available: http://informahealthcare.com/doi/abs/10.1088/0954-898X_7_4_004
- [31] K. Zhang, "Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory," *The journal* of neuroscience, vol. 16, no. 6, pp. 2112–2126, 1996.
- [32] T. Kyriacou, "An implementation of a biologically inspired model of head direction cells on a robot," in *Towards Autonomous Robotic Systems*, ser. Lecture Notes in Computer Science, R. Gro, L. Alboul, C. Melhuish, M. Witkowski, T. Prescott, and J. Penders, Eds. Springer Berlin Heidelberg, 2011, vol. 6856, p. 66. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-23232-9_7
- [33] H.-G. Beyer and H.-P. Schwefel, "Evolution strategies-a comprehensive introduction," *Natural computing*, vol. 1, no. 1, pp. 3–52, 2002.
- [34] T. Kyriacou, "Using an evolutionary algorithm to determine the parameters of a biologically inspired model of head direction cells," *Journal* of computational neuroscience, vol. 32, no. 2, pp. 281–295, 2012.
- [35] A. Arleo and W. Gerstner, "Spatial cognition and neuro-mimetic navigation: A model of hippocampal place cell activity," *Biological Cybernetics*, vol. 83, no. 3, pp. 287–299, 2000.
- [36] P. Song and X.-J. Wang, "Angular path integration by moving hill of activity: a spiking neuron model without recurrent excitation of the head-direction system," *The journal of neuroscience*, vol. 25, no. 4, pp. 1002–1014, 2005.

- [37] T. Degris, L. Lacheze, C. Boucheny, and A. Arleo, "A spiking neuron model of head-direction cells for robot orientation," in *Proceedings* of the Eighth Int. Conf. on the Simulation of Adaptive Behavior, from Animals to Animats, 2004, pp. 255–263.
- [38] J. P. Goodridge and D. S. Touretzky, "Modeling attractor deformation in the rodent head-direction system," *Journal of neurophysiology*, vol. 83, no. 6, pp. 3402–3410, 2000.
- [39] P. Zeidman and J. A. Bullinaria, "Neural models of head-direction cells," *From associations to rules: Connectionist models of behavior* and cognition, pp. 165–177, 2008.
- [40] D. Hebb, "The organization of behavior: A neuropsychological theory," 1949.
- [41] T. Trappenberg, "Continuous attractor neural networks," *Recent developments in biologically inspired computing*, pp. 398–425, 2003.
- [42] W. Gerstner and W. M. Kistler, Spiking neuron models: Single neurons, populations, plasticity. Cambridge university press, 2002.
- [43] J.-A. Meyer, A. Guillot, B. Girard, M. Khamassi, P. Pirim, and A. Berthoz, "The psikharpax project: Towards building an artificial rat," *Robotics and autonomous systems*, vol. 50, no. 4, pp. 211–223, 2005.
- [44] J.-A. Meyer, A. Guillot, P. Pirim, A. Berthoz, and F. BEV, "Psikharpax: An autonomous and adaptive artificial rat," in *Proceedings of ISR2004*, the 35th International Symposium on Robotics, Paris, 2004.
- [45] M. J. Milford, G. F. Wyeth, and D. Prasser, "Ratslam: a hippocampal model for simultaneous localization and mapping," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 1. Ieee, 2004, pp. 403–408.
- [46] S. Thrun, "Simultaneous localization and mapping," *Robotics and cognitive approaches to spatial mapping*, pp. 13–41, 2008.
- [47] C. Lever, S. Burton, A. Jeewajee, J. O'Keefe, and N. Burgess, "Boundary vector cells in the subiculum of the hippocampal formation," *The journal of neuroscience*, vol. 29, no. 31, pp. 9771–9777, 2009.
- [48] T. V. Bliss and T. Lmo, "Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path," *The Journal of physiology*, vol. 232, no. 2, pp. 331–356, 1973.
- [49] T. Teyler and P. DiScenna, "Long-term potentiation," Annual review of neuroscience, vol. 10, no. 1, pp. 131–161, 1987.
- [50] M. Ito, "Long-term depression," Annual review of neuroscience, vol. 12, no. 1, pp. 85–102, 1989.
- [51] W. Gerstner and W. M. Kistler, "Mathematical formulations of hebbian learning," *Biological Cybernetics*, vol. 87, no. 5-6, pp. 404–415, 2002.
- [52] K. D. Miller and D. J. MacKay, "The role of constraints in hebbian learning," *Neural Computation*, vol. 6, no. 1, pp. 100–126, 1994.