Online Correction of Orientation Estimates Using Spatial Memory in a Neuromorphic Head Direction System

Tarek M. Massoud
Electrical and Computer Engineering Dept.
Institute for Systems Research
University of Maryland
College Park, MD 20742, U.S.A.
tmassoud@umd.edu

Timothy K. Horiuchi
Electrical and Computer Engineering Dept.
Institute for Systems Research
Neuroscience and Cognitive Science Program
University of Maryland
College Park, MD 20742, U.S.A.
timmer@isr.umd.edu

***Abstract*— Many animals are known to maintain an internal estimate of their orientation in the environment. In the absence of external sensory cues, this estimate inevitably exhibits drift. When sensory information is available, associations between sensory landmarks and the internal estimate can be used to correct for drift. In this paper we present a neuromorphic system to model such associations between sensory landmarks in the environment (as provided by sonar) and the activity of a hardware-based head direction cell system (HDS) that continuously integrates angular velocity signals to maintain an estimate of the orientation. These associations are shown to correct for drift errors that are encountered in the HDS.**

I. INTRODUCTION

Spatial navigation is an important skill upon which depends the survival of many animals. In most animals, the neural mechanisms used to perform navigation are not well understood but many discoveries in mammals suggest the use of a combination of internal estimates of their position in space (e.g., odometry using place cells, head direction cells, and grid cells) and available sensory information in the environment (e.g. vision, audition, olfaction, etc…)[1, 2]. In the absence of external sensory cues, animals can navigate successfully, but errors accumulate over time. When sensory cues *are* present in the environment, such drifts in navigation have not been observed [3-5]. These findings suggest that, when animals do not have sensory signals for navigation, they rely on internal estimates of their position in the environment to navigate and that this estimate is noisy. In contrast, the presence of navigational sensory cues is used to continuously correct for drift.

This problem represents one of the challenges facing biologically-inspired robotics and there has been an increasing effort to solve it [6]. In this paper we present a mixed hardware and software system that offers a biologically-plausible model of how the brain could integrate spatial memory (e.g., echolocation sensing of the environment) with an internal estimate of its position in space (e.g., head orientation as estimated by our neuromorphic head direction (HD) cell system; described in full detail in [7]) to keep the noisy estimate of the orientation aligned with the environment.

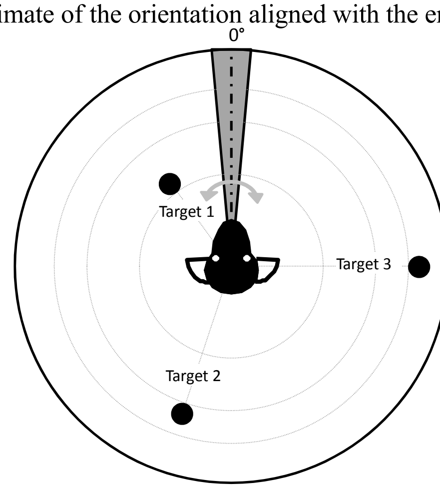

Fig 1. A simple sonar transducer is mounted on a rotating platform from which the rotation velocity can be measured. The grey cone represents the effective field of view of the sonar. For simplicity in this demonstration, the targets are classified based on their radial distance from the head.

II. SYSTEM MODEL

Our demonstration system is a mixed software-hardware system that solves a one-dimensional version of the problem described earlier; maintaining an internal estimate of head orientation following rotations. Fig 1 shows a schematic of the setup we are using, we have a sonar head mounted on a rotating platform from which we can measure the rotation velocity. The rotation velocity signal is integrated by the HD system to continuously update the estimate of the head

This work was supported by a grant from the Air Force Office of Scientific Research (FA95500710446) and the National Science Foundation (CCF0347573)

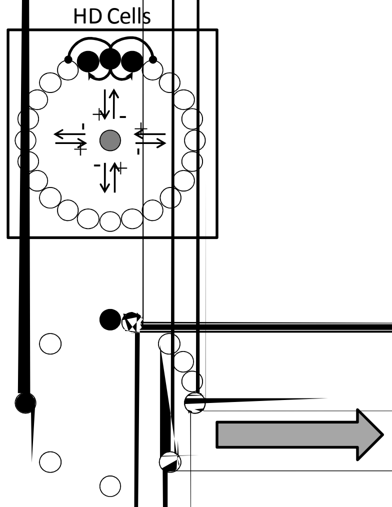
HD Cells

When the sonar does not detect any objects, the "no object" cell in the object cell group will be active and the activity in the conjunctive neurons directly reflects the HD neurons. The learning rule weakens the connections from inactive object cells to the conjunctive cell layer (OC synapses). In parallel with this learning, the weak one-to-one excitatory connection from the object cell layer to the expectation cell layer excites the expectation layer's "no object" cell. Once this occurs, Hebbian learning strengthens the connection between the active conjunctive layer cells and the active expectation layer "no object" cells, thus learning that no object is at this orientation.

When the sonar system detects the presence of a known object (recognized but not seen in the environment before), the object cell corresponding to the sensed object will be active. Hebbian learning begins to strengthen the OC synapse from the active object cell to the active conjunctive cell *and* drives any conjunctive cells that were previously associated with this object cell (if any). If the OC projection from the active object cell projects to a different set of conjunctive neurons than those that are active, more than one set of conjunctive neurons can become active, leading to a reset condition. (We discuss this in more detail below.) The projections from the conjunctive cells to the expectation cells (CE synapses) are designed to be much stronger than the ones from the object cells to the expectation cells. Additionally, the winner-take-all function within the expectation cell group will ensure that the expectation cells will normally be driven by the conjunctive cells.

As the sonar platform system is rotated, errors in the HD estimates begin to accumulate. During normal operation we can identify three distinct cases:

1. The same cells are active in both the expectation and object cell groups, meaning that the HD estimate is aligned with its memory of the sensory experience.
2. An object cell is active but a different expectation cell is active in the expectation cell group (typically the "no object" cell), meaning that the sensory system has found an object while the spatial memory did not expect that object based on its current estimate of orientation. If the object already has an association in memory, we allow this memory to override the current HD estimate of orientation.
3. An expectation cell is active, but the "no object" cell is active in the object cell group, meaning that we know that the HD estimate of orientation has drifted but there isn't enough information yet to fix the estimate. In this case, we simply suppress Hebbian learning so that previously learned associations for the currently estimated orientation are not lost.

III. Results

We have tested the system using different combinations of object number and orientations. For all of these cases, the system was able to successfully learn the environment and correct for accumulating errors in the HD estimate of orientation. In this section we show the activity in the different groups of cells for the three cases discussed in section II. The results presented here are from an experiment with 4 objects in the environment placed at (0˚, 337.5˚, 315˚, and 292.5˚).

A. Properly Aligned.

In this case, the expectation cell and object cell groups are aligned in their activity, implying that the HD estimate of orientation is either aligned with the actual orientation or there is insufficient data to suspect otherwise. For these two cases Fig 3 shows an example of the activity in the system.

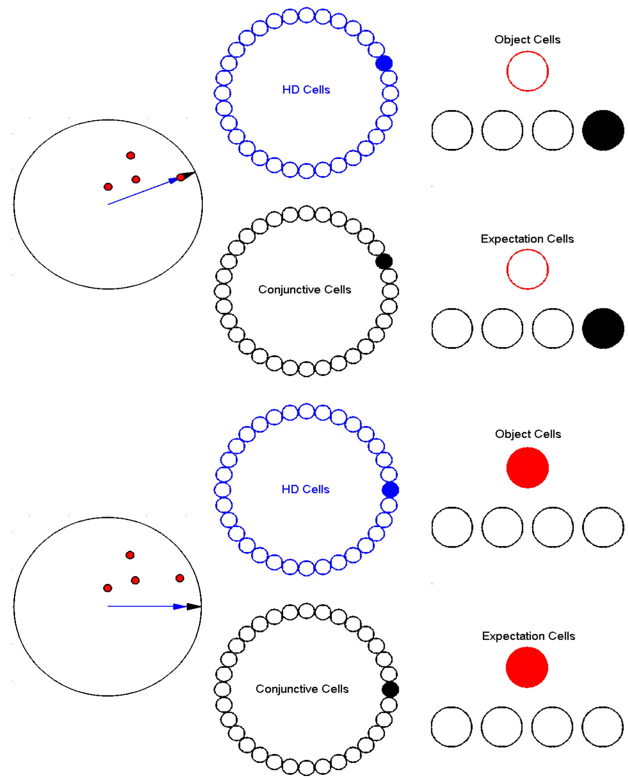

Fig 3. This figure shows the activity in the system when the HD estimate is aligned with the actual orientation in space. The left panel shows a schematic for the arena with the targets as red circles. The black arrow represents the actual head position and the blue arrow is the position as estimated by the HDS. The center panel shows the HD and Conjunctive cells. Although in practice the HD system activates a contiguous group of four neurons when indicating a location, for simplicity we only show the activity of one cell active for each position on both networks. The right panel shows the Object and Expectation cells, the top cell (in red) is the "no object" cell and each of the bottom cells represents one of the four targets. (a) shows the case when the head is pointing towards target #1 and (b) shows the case when the head is not pointing towards any target.

B. Disoriented.

In this case, it is known that the HD estimate of orientation is not accurate, however, there is not enough information to correct the error. The system suppresses the learning process to avoid losing previously acquired memories. Fig. 4 shows the activity of the cells in this case.

C. Reset Condition.

Here, the system knows that the HD estimate of orientation is not accurate *and* can recall the previously learned orientation for that object. This information is stored in the OC synapses which are used to abruptly move (or "reset") the activity in the HD system to a remembered location. Fig. 5

shows an example for this case where the HD system was estimating the head to be pointing at 270˚ where no object is placed, however the sensory data show that the head is pointing towards target 4 which is placed at 292.5˚ so the HD is reset to the remembered orientation.

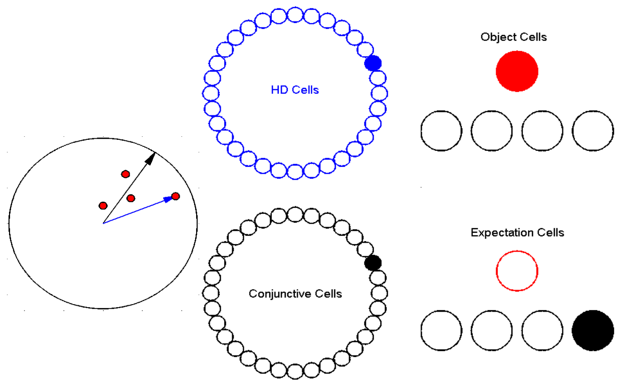

Fig. 4. The system is disoriented. Based on the current HD estimate, the system was expecting to see object #4, however the live sensory data shows no target in sight.

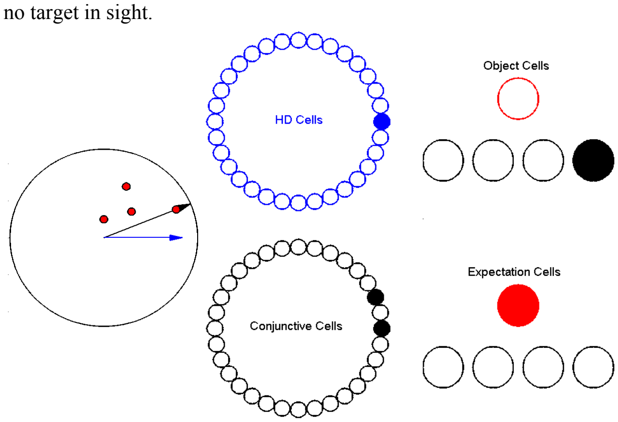

Fig. 5. The system will reset. Based on the HD estimate of orientation, the system was not expecting to see a target, however, the sensory data show the presence of object #4. The system will reset the HD system to point toward the position of object #4. Note that the activity in the conjunctive cells reflect both the HDS' estimate and the stored orientation.

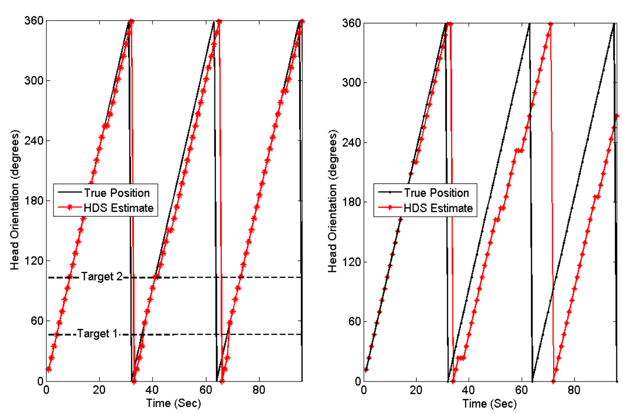

Fig 6. Left panel shows the results from an experiment with 2 targets present at (45˚ and 100˚), as the HDS drifts, spatial memories of the targets locations are used to reset it to the accurate head position. Right panel shows the case with no correction, the HDS' estimate accumulates error with time with no means to correct it.

D. System Performance.

In this section we show the performance of the system in the presence of learning and compare it to the case where no learning was applied. Fig 6 shows the results of an experiment where we had 2 targets in the space at (45˚ and 100˚). In both cases, the HD estimate of the position is aligned with the actual position at the beginning of the experiment. Due to the noisy integration in the HD system, the estimate accumulates some integration errors and begins drifting away from ground truth. In the case without learning, the errors accumulate and the orientation error grows larger over time. In contrast, the learning case shows that the first encounter was sufficient to reset the HD system to the same orientation in future encounters.

IV. Conclusion

In this paper we have proposed an extension to an existing neuromorphic head-orientation odometry system that uses spatial memory to improve and maintain the accuracy of its estimate. Sensory cues in the environment are associated with different orientations and are capable of correcting errors when these cues are encountered again. Beyond the specific corrections provided by the memory, these memories could also provide longer time-constant corrections in integration gain. This system is built as part of a biologically inspired navigation system that can be mounted on a mobile robot and it represents one important mechanism (of many) that could explain how biological neurons with all of their variability and noise can be used reliably for navigation in everyday environments.

Acknowledgment

We would like to thank MOSIS for the chip fabrication.

References

[1] E. I. Moser, *et al.*, "Place Cells, Grid Cells, And The Brain's Spatial Representation System," *Annual Review Of Neuroscience,* vol. 31, pp. 69-89, 2008.

[2] B. L. McNaughton, *et al.*, "Path Integration And The Neural Basis Of The "Cognitive Map"," *Nature Reviews Neuroscience,* vol. 7, pp. 663-678, 2006.

[3] J. S. Taube and J. P. Bassett, "Persistent Neural Activity In Head Direction Cells," *Cerebral Cortex,* vol. 13, pp. 1162-1172, 2003.

[4] D. P. Goodridge JP, Worboys KA, Golob EJ, Taube JS., "Cue control and head direction cells.," *Behav Neurosci,* vol. 112, pp. 741-69, 1998.

[5] R. H. Hahnloser, "Emergence Of Neural Integration In The Head-Direction System By Visual Supervision," *Neuroscience,* vol. 120, pp. 877-891, 2003.

[6] P. Stratton, *et al.*, "Automatic Calibration of a Spiking Head-Direction Network for Representing Robot Orientation," presented at the Australian Conference on Robotics and Automation, Sydney, Australia, 2009.

[7] T. M. Massoud and T. K. Horiuchi, "A Neuromorphic VLSI Head Direction Cell System," *Circuits and Systems I: Regular Papers, IEEE Transactions on,* vol. 99, 2010.

[8] K. A. S. Hertz J.A., Palmer R.G., *Introduction To The Theory Of Neural Computation, Volume I*: Westview Press, 1991.

[9] F. Sargolini, *et al.*, "Conjunctive Representation of Position, Direction, and Velocity in Entorhinal Cortex," *Science,* vol. 312, pp. 758-762, May 5, 2006 2006.

[10] W. E. Skaggs, *et al.*, "A Model Of The Neural Basis Of The Rat's Sense Of Direction," *Advances In Neural Information Processing Systems,* vol. 7, pp. 173-180, 1995.