

ABSTRACT

Title of Document: MODELING THE BAT SPATIAL
NAVIGATION SYSTEM: A
NEUROMORPHIC VLSI APPROACH

Tarek M Massoud, Doctor of Philosophy, 2012

Directed By: Professor Timothy K. Horiuchi
Department of Electrical and Computer
Engineering

Autonomously navigating robots have long been a tough challenge facing engineers. The recent push to develop *micro-aerial* vehicles for practical military, civilian, and industrial use has added a significant power and time constraint to the challenge. In contrast, animals, from insects to humans, have been navigating successfully for millennia using a wide range of variants of the ultra-low-power computational system known as the brain. For this reason, we look to biological systems to inspire a solution suitable for autonomously navigating micro-aerial vehicles. In this dissertation, the focus is on studying the neurobiological structures involved in mammalian spatial navigation. The mammalian brain areas widely believed to contribute directly to navigation tasks are the Head Direction Cells, Grid Cells and Place Cells found in the post-subiculum, the medial entorhinal cortex, and the hippocampus, respectively. In addition to studying the neurobiological structures

involved in navigation, we investigate various neural models that seek to explain the operation of these structures and adapt them to neuromorphic VLSI circuits and systems. We choose the neuromorphic approach for our systems because we are interested in understanding the interaction between the real-time, physical implementation of the algorithms and the real-world problem (robot and environment). By utilizing both analog and asynchronous digital circuits to mimic similar computations in neural systems, we envision very low power VLSI implementations suitable for providing practical solutions for spatial navigation in micro-aerial vehicles.

MODELING THE BAT SPATIAL NAVIGATION SYSTEM: A NEUROMORPHIC
VLSI APPROACH

By

Tarek Mohamed Anwar Massoud

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2012

Advisory Committee:
Professor Timothy K. Horiuchi, Chair / Advisor
Professor Pamela Abshire
Professor Perinkulam Krishnaprasad
Professor Jonathan Simon
Professor Cynthia Moss

© Copyright by
Tarek Mohamed Anwar Massoud
2012

Dedication

To my parents, my wife (Hoda), my little angel (Lina), and my late grand-mother.

Acknowledgements

I wish to express my deep gratitude to my advisor Dr. Timothy Horiuchi for his invaluable guidance and support during the course of this work. I really enjoyed working under his supervision and am thankful to him for being always there when I needed help or support on the professional level and on the personal level. I learned a lot from his passion about science and sincerity about academic excellence.

I would like to thank Dr. Pamela Abshire, Dr. Perinkulam Krishnaprasad, Dr. Jonathan Simon, and Dr. Cynthia Moss for serving on my dissertation committee. My special thanks go to Dr. Shihab Shamma, who served on my dissertation committee and provided advice and encouragement until the very end but was not able to attend my defense.

I would like to thank all the wonderful people I met at the Telluride Neuromorphic Engineering Workshop over the years for very informative discussions and teaching me the spirit of team work. I would like to extend special thanks to Dr. Giacomo Indiveri, Dr. Elisabetta Chicca, Dr. Tobi Delbruck, and Dr. Ralph Etienne-Cummings for their help and beneficial discussions.

I would like to thank the wonderful labmates I had over the years, Matt Cheely, Rock Shi, Hisham Abdallah, Chetan Bansal, and Matt Runchey.

I would like to thank Peter Mariani for his help preparing for the thesis defense.

I would like to thank the great group of friends who supported me over the years, the great times we had here helped me overcome home sickness, loneliness and retain my sanity all the way.

Last but not least, I thank my parents who encouraged me to learn and who have always and continue to support me, and I would like to acknowledge the understanding, patience, encouragement, and love of my wife. Without the love and support of my family I could not have made it.

Table of Contents

Dedication	ii
Acknowledgements	iii
Table of Contents	v
List of Figures	ix
Chapter 1 : Introduction.....	1
1.1 Introduction.....	1
1.2 Problem Overview	5
1.3 Thesis Overview	10
Chapter 2 : Biological Background	11
2.1 Head Direction Cell System.....	11
2.1.1 Neuroanatomy of Head Direction Cells.....	12
2.1.2 Neurophysiological Properties of Head Directional Cells.....	14
2.1.3 Head Direction Cell Determinants:.....	15
2.1.4 From Head Direction Cells to Navigation:	16
2.2 Grid Cells	17
2.3 Place Cells.....	19
2.3.1 Neuroanatomy.....	19
2.3.2 Neurophysiological Properties of Place Cells.....	24
2.3.3 Hippocampus' Place Cells (DG, CA1, CA3).....	26
2.3.4 Subicular Place Cells	28
2.3.5 Place Field Determinants	29
2.3.6 The Theta Rhythm and Hippocampal Activity:.....	30
2.3.7 Hippocampal Synaptic Plasticity	31
2.3.8 Memory Space or Space Memory?.....	32
2.3.9 Navigation:.....	33
Chapter 3 : Modeling the Navigation System (State of the Art)	34
3.1 Modeling Head Direction Cells	34

3.1.1	(McNaughton et al., 1991)[43]:	34
3.1.2	(Skaggs et al. 1995)[67]:.....	35
3.1.3	(Redish et al. 1996) [68]:	37
3.1.4	(Goodridge and Touretzky, 2000) [70]:.....	40
3.1.5	(Sharp et al., 2001)[37]:.....	43
3.2	Modeling the Grid Formation:	46
3.2.1	Continuous Attractor Models:.....	46
3.2.2	(Fuhs and Touretzky, 2006)[71]:	47
3.2.3	(McNaughton et al., 2006)[72]:	48
3.2.4	Interference Model, [73, 74]:.....	49
3.3	Modeling the Place Cells	51
3.3.1	“Pre-Grid” Place Cell Models.....	52
3.3.2	Post-Grid Cells Models.....	57
3.4	From Models to Navigation Systems.....	60
3.4.1	(Arleo, 2000; Arleo and Rondi-Reig, 2007) [4, 76]:	60
3.4.2	(Milford, 2008) [10]:.....	61
3.4.3	(Erdem and Hasselmo 2012)[77]:.....	63
Chapter 4 :	Head Direction Cell System.....	65
4.1	Introduction.....	65
4.2	HD System Model.....	67
4.2.1	Stable Activity in the Network	68
4.2.2	Moving the “Bump”	71
4.3	Circuits	75
4.3.1	Synapse Circuit	76
4.3.2	Neuron Circuit	78
4.4	System Implementation	79
4.4.1	Bump Formation	80
4.4.2	Angular Velocity Integration	81
4.5	Testing Results.....	83
4.5.1	Creating a Stable Bump of Activity.....	83
4.5.2	Mean Rate Mode of Operation	84

4.5.3	Synchronized (Bursting) Mode of Operation	90
4.5.4	Controlling the Bump Width	94
4.5.5	System Data	97
4.5.6	Moving the Bump	98
4.5.7	Drift in position estimation	99
4.5.8	Resetting the bump	101
4.5.9	Characterizing the HD System.....	102
4.6	Conclusions.....	106
Chapter 5 :	Online Error Correction in HD System	108
5.1	Introduction.....	108
5.2	System Model	109
5.2.1	Sonar Object Detection and Object Cells	111
5.2.2	Head Direction Cell System.....	111
5.2.3	Conjunctive Cells.....	111
5.2.4	Expectation Cells	112
5.2.5	System Operation.....	113
5.3	Results.....	116
5.3.1	Properly Aligned.....	116
5.3.2	Disoriented.....	116
5.3.3	Reset Condition.....	118
5.3.4	System Performance.	118
5.4	Conclusion	120
Chapter 6 :	Grid Cell System.....	122
6.1	Introduction.....	122
6.2	Modeling the Grid Cells.....	124
6.3	Simulating Grid Formation	124
6.4	The Grid Circuit.....	126
6.4.1	Chip Organization.....	126
6.4.2	The Pixel	128
6.4.3	Projection Network	128
6.4.4	Inhibitory Input	130

6.4.5	Excitatory Input	130
6.4.6	Readout	130
6.4.7	Request Generation.....	131
6.4.8	Reset Generation.....	131
6.5	Chip Results	133
6.6	Modeling the Chip Connectivity.....	136
6.6.1	1-Dimensional Analysis.....	136
6.6.2	2-Dimensional Analysis.....	137
6.7	Moving the Grid Pattern	140
6.8	Conclusion	141
Chapter 7 :	Place Cell System	142
7.1	Introduction.....	142
7.2	Head Direction Cells to Place Cells.....	143
7.3	Grid Cells to Place Cells.....	143
7.3.1	Grid Cells with Similar Velocity Drive and Different Connectivity	144
7.3.2	Grid Cells with Similar Connectivity and Different Velocity Drive	149
7.4	Conclusions.....	152
Chapter 8 :	Conclusions and Future Directions.....	153
Bibliography	156
List of Publications	165

List of Figures

Fig 2.1. HD system Connectivity. The top panel shows the location of different structures exhibiting HD responses in the rat's brain (blue boxes) [14]. The bottom panel shows the functional interconnection between regions with HD cell and neighboring structures, shaded boxes denote regions that are known to contain HD cells. Abbreviations: AD, anterodorsal thalamus; DTN, dorsal tegmental nucleus; HIP, hippocampus; LMN, lateral mammillary nucleus; PoSC, postsubiculum; RsCX, retrosplenial cortex; CX, cortex, modified from [40].	13
Fig 2.2. The receptive fields of 2 HD neurons from the rat postsubiculum, adapted from [14].	14
Fig 2.3: Grid Cells, Trajectory maps (left), rate maps (middle) and spatial autocorrelograms (right) for three cells recorded from the MEC. The peak firing rate for each cell is indicated from [19].	18
Fig 2.4. Anatomical interconnections between the different structures in the hippocampal formation. EC: entorhinal cortex, DG: dentate gyrus, SC: subiculum, paSC: parasubiculum, prSC: presubiculum, the hippocampus proper consists of the DG and CA3-CA1 areas (shaded area). The hippocampal intrinsic loop is shown in bold: EC → DG → CA3 → CA1 → SC → EC adapted from [4].	20
Fig 2.5. Place Cells in the hippocampus: a) Spike locations (<i>red</i>) are superimposed on the animal's trajectory in the recording enclosure (<i>black</i>), adapted from [47]. b) Firing rate maps of place fields from 4 CA1 cells, the peak firing rates vary between 4.8, 22.8 Hz.; adapted from [55].	25
Fig 2.6 Subicular Place Cells recorded in two arenas, it is to note that the cells exhibit topologically related firing in both arenas, adopted from [59].	29
Fig 3.1. Head Direction model from Touretzkey and Skaggs [67].	36
Fig 3.2. Two ring attractor network from [68].	39
Fig 3.3. Schematic of HD Neural System [68].	39
Fig 3.4. Neuronal Model from [70].	40
Fig 3.5. a) Stable bump, b) Bump rotation from [37].	45
Fig 3.6. Grid pattern in the MEC layer forms instantaneously and makes a displacement to the right to follow the motion. In the origin of the white axes, a neuron is not firing at time $t = 0$, fires in the maximum of a grid node at $t = 140$, and is at rest again at $t = 290$ from [71].	47

Fig 3.7. A topographically arranged network serves as a tutor to train an MEC module with no topographical arrangement (*left*). Connectivity in a layer of MEC if the connectivity is rearranged topographically (*center*), the effective geometry is that of a toroidal surface (*right*) from [71]..... 48

Fig 3.8. Sum of three or more linear interference maps; while the animal is running on a linear track, the sum of a somatic (s) and a dendritic (d) oscillation with slightly different frequencies results in an interference pattern exhibiting phase precession and slow periodic spatial modulation (*left*). Combining three or more linear interference maps, responding to different projections of the velocity, results in a grid map (*center*). Simulated grid map after 10 min of a rat’s actual trajectory (*right*) from [73]. 50

Fig 3.9: Sharp model [14]. 55

Fig 3.10. Burgess, Recce, and O’Keefe model, adopted from [55]. 56

Fig 3.11. Multiple grid fields with different scales contribute to form a single peaked place field adopted from [72]. 58

Fig 3.12. Simulated place cell formation in the hippocampal DG cells for the model proposed by Rolls et al., adopted from [75]. 59

Fig 3.13. The architecture of the RatSLAM algorithm, adopted from [10]. 62

Fig 3.14. A schematic for the network used in the (Erdem and Hasselmo) model for goal directed navigation; adapted from [77]. 64

Fig 4.1. The synaptic projection pattern as a function of the relative interneuron distance. The actual strength is determined by a product of this kernel and a global current setting. A positive value indicates an excitatory connection, a negative value indicates an inhibitory connection and “0” indicates no connection. 68

Fig 4.2. A ring of neurons with the interconnections shown for only one neuron for clarity. The neuron in the center of the ring is the global inhibitory neuron. The connectivity allows one group of neurons (“bump”) to be active at a time (shaded in black). The global inhibitory neuron (shaded gray) normalizes the total network activity through inhibitory feedback. Additional neurons and connectivity for moving the bump activity are not shown. 69

Fig 4.3. Connectivity diagram of the bump movement neurons. The arrow head indicates an excitatory connection and the circle-shaped head indicates an inhibitory connection. The network connections that maintain bump activity (see Fig. 4.1) are not repeated here for clarity. 71

Fig 4.4. Schematic diagram showing the sequence of events for moving the bump of activity based on a left velocity signal (active connections are shown in black and the inactive ones in light gray). The top panel shows the stationary case with no velocity

signal, the middle panel shows the activity in the network as a left rotation signal is applied and the bottom panel shows the activity in the network after it moves one neuron to the left in response to the velocity signal..... 72

Fig 4.5. Micrograph of the Neuron Chip. The chip has 32 I&F neurons, each with 14 Pulse Extender (PE) synapses. The activity is read from the chip using an AER interface..... 76

Fig 4.6. Synapse circuit: VbiasN1, VbiasN2 and VbiasN3 control the charging and discharging of C1 and C2 and hence control the width of the pulse extender pulse. Vw is the weight of the synapse and controls the amplitude of the output current; for an excitatory synapse the output current is from the drain of M9, whereas for an inhibitory synapse, the output current is at the drain of M11. 77

Fig 4.7. Neuron circuit slightly modified from the one presented by Indiveri et al [88]. 79

Fig 4.8. Block Diagram of the HD system. The dsPICs are the boxes labeled PIC and the other boxes refer to multiple copies of our neuron chip. 80

Fig 4.9. Testing setup showing the Bump chip, the routing microcontroller, and the serial connection to the PC. 81

Fig 4.10. HD System Connectivity: The neuron chips are labeled in the figure and the other boards are the dsPIC microcontroller implementing the AER routing in the system. 83

Fig 4.11. Left panel, neural response spike rasters of the ring neurons to a constant input current with the global inhibitory neuron suppressed. Multiple bumps of activity are present. Right panel; response of the ring neurons to a constant input current with the global inhibitory neuron providing feedback. The bump is initiated at an unstable location but eventually drifts to a more stable location. With the global inhibition activated, the activity bump occurs with lower firing rates than the case shown in the left panel..... 85

Fig 4.12. Left panel, rasters showing response of bump neurons all of them receiving the same tonic input. Right panel, deviation from mean firing freq by individual neurons..... 87

Fig 4.13. Each trace represents the centroid of the activity of a bump started at each of the 32 possible locations around the ring. The plot shows the presence of some global attractors; for example, the bumps with initial center of mass near neurons 6, 7, 8, 9, and 10 eventually migrate to a centroid near neuron 9. The plot shows the two time intervals where the drift in the bump occurred; the first is the first 0.5 sec of operation; then another window after 2.5 seconds until 5 seconds, after 5 seconds we did not see any other drift in the bump location. A neuron # wrap-around was used for better visualization of the attractor basins. 88

Fig 4.14. Left panel, neural response spike rasters of the ring neurons to a constant input current with the global inhibitory neuron suppressed. Multiple bumps of activity are present. Right panel, response of the ring neurons to a constant input current with the global inhibitory neuron providing feedback. A single, synchronously firing group of neurons emerges. 91

Fig 4.15. Each trace represents the centroid of the activity of a bump started at each of the 32 possible locations around the ring in the synchronized mode of operation. The plot shows that the bumps initiated at all locations are stable (i.e., no drift was detected during the course of operation). The same neuron # wrap-around as in Fig. 13 was used. 92

Fig 4.16. Activity bumps with variable widths of three, four, and five neurons wide. For the three neuron-wide bump, each neuron was connected to 10 others (four excitatory and six inhibitory). For the four neuron-wide bump each neuron was connected to 12 others (six excitatory and six inhibitory), and for the five neuron-wide bump each neuron was connected to 14 others (eight excitatory and six inhibitory). 95

Fig 4.17. System data, for rotating the bump to the left with variable speed. 96

Fig 4.18. Moving the bump around the ring in both directions with variable rotation speed. The bump starts at rest. At 6.2 sec, the system is fed with an increasing leftward rotation velocity until it reaches a maximum of 153.2 deg / sec at 10.4 sec. Starting from 14.8 sec, it is slowly brought back to a full stop at 16.6 sec. In the second part starting from 21.3 sec, an increasing right-ward rotation velocity is fed to the system until it reaches a maximum speed of 123.3 deg / sec at 22.3 sec. At 29.3 sec, the rightward velocity signal is decreased slowly until the system reaches a rest at 33.3 sec. 98

Fig 4.19. The black, blue and green lines show the evolution head position as represented by the centroid of the bump of the activity as the system is driven to perform a 360° rotation around the ring using three different rotation speeds and the dotted gray lines show the expected head position if the same velocities were presented to a perfect velocity integrator system. 100

Fig 4.20. The error in estimating the head position in degrees as coded for using our system vs. using a perfect integrator for the same three examples shown in Fig. 19. 100

Fig 4.21. Resetting the Bump Location: Externally-provided stimulation input is shown in big black dots, spikes from the ring neurons are in gray dots and the inhibitory neuron spikes are in red along the top of the graph. 102

Fig 4.22. Average spiking rates on the rotation chips and the corresponding speed signal Vang. 104

Fig 4.23. The average bump velocity measured in (Neurons per Sec.) as a function of average spiking rates on the rotation chips. The vertical bars represented one standard deviation from the average.....	106
Fig 5.1. 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.....	109
Fig 5.2. System block diagram. The blocks enclosed in the dashed-line box are implemented in software and the other blocks are in hardware. The black arrows indicate predefined non-plastic synaptic connections, the arrows in grey show the plastic synaptic connections, and the white arrows indicate non -plastic connections used as teacher signals to guide the learning process.	110
Fig 5.3. 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.....	117
Fig 5.4. 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.	119
Fig 5.5. 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.....	119
Fig 5.6. 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.	120
Fig 6.1. Response of 4 grid cells from the entorhinal cortex of a rat, modified from [19]......	123
Fig 6.2. Two examples of interconnection weights for neuron at location (0, 0) on a 2D sheet of neurons used to simulate grid cells activity.....	125

Fig 6.3. Firing rates (in Hz) of one simulated grid neuron over a virtual arena using the connection weights schemes shown in Fig 6.2, showing that the spatial frequency of the response grid changes by changing the width of the connection matrix for the neurons on the sheet. The neuron in the left panel is connected to its neighbors using the weight pattern shown in Fig 6.2 (left) and the one in the right panel uses the weights from Fig 6.2 (right panel). 125

Fig 6.4. Micrograph of the Grid Chip. The chip has an array of 16 x 16 analog pixels. Scanner circuits are used to select pixels, allowing the activity of each pixel to be read out in the form of AER pulses 127

Fig 6.5. Diagram showing the connectivity of one pixel. The 16 x 16 array is formed by tiling this pixel. Each pixel receives row and column select lines, is connected to each of its neighbors by six local signal lines two of which go through pFET diffusers (small shaded box), and shares a global output voltage line (Vmem). 127

Fig 6.6. Pixel Circuit. The pixel circuit consists of the central node (between M13 and M14) and four circuit blocks that implement the inhibitory and excitatory projection kernels and a readout circuit. 129

Fig 6.7. Request generation. M25 provides the bias for the source-follower transistor in every pixel (M24 in Fig. 1). The voltage output from the selected pixel is buffered and compared to Vref_diff, to generate a Req_bar signal when the pixel voltage is higher than the spiking threshold. 132

Fig 6.8. Reset generation. The Ack_bar signal coming from the AER system pulls the Reset signal low. The Reset signal can only be raised high by the rising edge of a new clk cycle as the scanner selects the next pixel. 132

Fig 6.9. Raster plot of 0.2 sec of activity recorded from the chip, this raster translates to the patterns shown in Fig 6.10. The rasters show pixels firing at different frequencies which create analog bumps of activity on the chip. 134

Fig 6.10. Reconstructed grid cell circuit activity pattern on the chip at three different spatial frequencies. The brightness of the pixel represents the rate of spikes generated by the pixel (spikes/sec). 135

Fig 6.11. Matlab generated response of a 16 x 1 array of pixels using the model of the interconnection on the chip for an impulse input at pixel 8. ($a = 0.17$, $c = 0.25$) 139

Fig 6.12. Matlab generated response of a 16 x 16 array of pixels using the model of the interconnection on the chip showing the emergence of grid-like patterns on the pixels ($a = 0.17$, $b = 0.25$) 139

Fig 6.13. Condition Number results for the system of Equations described in (4) for the entire range of parameters of a and c 140

Fig 7.1. The response of the three grid cells used to create the place cell. We have three different spatial frequencies for the grid cells' responses..... 146

Fig 7.2. Two examples for place fields that system was able to learn, the place field can be as small as one pixel or as big as 10 x 10 pixels. 147

Fig 7.3. Weight matrices for the two place cells shown in Fig 7.2..... 148

Fig 7.4. The grid cell response used to create the place cells in this experiment. 149

Fig 7.5. Two examples for place fields that system was able to learn, the place field can be as small as one pixel or as big as 10 x 10 pixels. 150

Fig 7.6. Weight matrices for the two place cells shown in Fig 7.5..... 151

Chapter 1 : Introduction

1.1 Introduction

If an animal (e.g., a rat) is placed in a novel environment and is taught to search for food, we see the animal moving around and exploring the environment to achieve this task. During this exploration, it maintains a sense of its starting point (home location) and returns back home successfully after finishing its task. What capabilities are needed in an autonomous robot to perform the same task?

In cognitive neuroscience, navigation is defined as the capability of planning and traveling along a path from the current position to a desired goal [1]. To accomplish this in a novel environment requires the following elements:

- Sensing the external world via multiple sensory modalities (e.g., vision, audition, olfaction, touch) and sensing its internal state through other specialized internal sensory modalities (e.g., vestibular signals, proprioception, motor commands).
- Using the collected information to adapt the behavior to achieve the goal.

Reaching an interesting location (e.g., a food source), returning home, finding shortcuts, and adopting efficient exploration strategies are all tasks that require spatial behavior and interaction with the environment.

Approaching a visible target's location is not difficult if it is in clear view and is uniquely identifiable (e.g. large red star). In this case, a simple target-guidance behavior can be used to do the task (*taxon navigation*) [2, 3]. Taxon navigation can be

understood as simple stimulus-response type of behavior. The agent associates a single motor response to a single stimulus (e.g. orients itself towards the stimulus and moves forward). Another type of navigation based on stimulus-response behavior is *praxic navigation* [3]. In this case, an agent moves towards a goal by executing a specific sequence of motor actions learned through prior training. This strategy is appropriate when the trajectory to the target is identified by a sequence of specific cues. Instead of single orienting responses as in taxon navigation, the agent must learn sequences of stimulus-action associations. In more difficult (and more realistic) situations, target locations are either not directly identified by a specific cue (or sequences of cues) or are simply hidden with respect to the agent's sensory capabilities.

A third form of navigation is goal-oriented navigation or (*locale navigation*) [3-5], which requires more complex spatial learning; the agent builds a representation of the environment, uses available sensory information to locate itself within this representation, and plans its next step to reach the goal. This type of representation is known as the *cognitive map* (Tolman, 1948) and is formed by linking available information about different locations in the environment (from sensors or proprioception) to build a unified representation for navigation. An agent behaving in this way is said to be performing a SLAM (Simultaneous Localization And Mapping) [6] behavior. This agent is able to explore new locations and create maps along the way. Moreover, it maintains an estimate of its position within this dynamically forming map.

All living creatures (e.g., desert ants, bees, rodents, monkeys, and humans) must navigate within their particular ecological niche and with their own unique senses and computational hardware. Is there a common theme in their navigation strategies?

When looking for food, desert ants forage and walk along winding, indirect paths, however, to return home, they take a straight path to the home nest. Experiments with the ant species (*Cataglyphis fortis*) suggest that they perform a form of path integration by maintaining an estimate of the average of all directions in which they have moved, weighted by the distance moved in each direction and thus always have a homing vector pointing towards home [7]. Path integration, however, is a noisy computation and experimental results show that the further the ants forage, the larger the error they suffer in the homing vector. Upon arriving to the expected nest site, if it is not found, they perform a systematic search to find their nest.

The honeybee is another example of an insect capable of foraging for food over very long distances. Typical distances of 2 to 3 Km have been observed with extreme trips up to 13.5 Km [8, 9]. There is considerable debate and modeling [10] on how they are able to navigate over such long distances. At one extreme, some models of honeybee navigation are based on the local route concept; a combination of path integration and route-specific landmark memories. Landmark memories are not stored within a general map, but in temporal sequences, that mirror the order in which landmarks were encountered. On the other extreme, it has been hypothesized that honeybees can learn their environment and build a cognitive map that they can use to

navigate. Experimental results found thus far can be interpreted in light of both theories.

The most extensively studied animals for spatial navigation are rats. The dominant experimental paradigm is to place the rat in one of three environments: an open arena where it walks and searches for food pellets, a “Morris water maze“ in which it swims to find (or remember the location of) a hidden platform, or runs along piecewise-linear T- or star-shaped mazes, while recording from a variety of neurons found in the hippocampal formation [15-18]. These experiments have demonstrated the presence of various types of cells in the hippocampal formation that respond to different aspects of navigational behavior: head direction cells that code for the animal’s head orientation in space, grid cells that code (in a periodic way) for the animal’s 2-D displacement in the arena, and place cells that code for the animal’s position in space [11, 14, 19].

The objective of this work is to study some of the neurophysiological mechanisms underlying spatial navigation behavior in mammals, focusing on the echolocating bat as our model organism. While there is a wealth of qualitative data about the spatial behavior of bats, however, very little data is available on the neurophysiology of the hippocampal formation in bats, however it has been shown in recent experiments that there are cells in the bat’s entorhinal cortex and hippocampus that behave in a very similar way to the grid cells and place cells of the rats [20-24]. We rely on the fact that bats are mammals and have brains very similar to the brains of rodents and therefore we adopt models based on the rodent literature in our work.

As our primary interest in spatial navigation stems from an interest in bat behavior, we adopt sonar as our main external sensory modality.

1.2 Problem Overview

The task of developing autonomously navigating systems (particularly those small and light enough to be mounted on micro-aerial vehicles) is still an open and very challenging problem. Successful solutions will likely be simple, power-efficient and able to adapt its behavior to unexpected changes in the environment, its own body and systems, and its goals. Many different technologies have been proposed for self-localization in small autonomous vehicles, ranging from very advanced and complex (large infrastructure) technologies like GPS to pure Artificial Intelligence (AI) algorithms (e.g., dead reckoning) to biomimetic algorithms and hybrid techniques comprising the use of one or more of the previous methods.

Although solutions using global-positioning satellites might seem straightforward, GPS receivers are rarely a complete solution for flying robots due to their power requirements and their reliance on satellite reception which limit their indoor operation. In any case, GPS does not help with collision avoidance without current, detailed, and accurate maps that include *everything* a flying vehicle might encounter, static or dynamic.

Traditional AI techniques for navigation are based on pre-specifying accurate internal models of the world which allow the agent to solve many tasks in a symbolic or logical fashion. Simple and accurate internal models of the world are, however, extremely hard to develop and most AI solutions are highly sensitive to errors in the

model [10]. In response to these drawbacks, a non-symbolic approach called behavior-based (or reflexive) robotics attempts to operate without much memory or prior assumption about the environment [25]. In this approach, the robot reacts to the complexity of the environment and builds its own minimal, functional representation of the world by means of its own experience and behavior using learning techniques. In most of these examples, robots extract just the information that is needed to solve the various tasks, producing fast, adaptive (but simple) behavior. The principles for designing behavior-based robots are often inspired from basic behaviors observed in biological systems and from neurophysiologically-inspired learning mechanisms. Several behavior-based learning frameworks, such as reinforcement learning and evolutionary techniques, have successfully addressed the problem of designing adaptive systems to autonomously navigate in unpredictable real environments [26]. Most of these systems, however, are based on reactive behavior; agents learn to map incoming stimuli to actions to accomplish their tasks without building any internal spatial model of the environment. The behavior-based paradigm can thus be employed to capture the functions of a simple landmark-guidance system in biological agents, but do not scale up well to more complex navigation tasks.

The issue of building an internal model of the world for autonomous navigation in mobile robotics has produced two principal approaches: the metric paradigm and the topological paradigm [4, 10].

The metric paradigm aims to model the geometrical features of the world accurately [27]. One very popular and intuitive method is the occupancy-grid method. It employs a multidimensional (typically 2D or 3D) tessellation of space into cells,

where each cell stores a probabilistic estimate of its state. The state of each cell represents the probability of occupancy of the corresponding area of the world [28]. The agent must rely on the information gathered by its sensors to build its knowledge of the world and hence the occupancy grid. The task of estimating the world from sensory information depends heavily on the quality of the data provided by the robot's sensors and is accordingly contaminated by noise in those sensors. The probability of occupancy of each grid cell is computed and updated by the information provided from all sensors that have information about the corresponding part of space. As the occupancy grids reproduce the geometrical structure of the environment explicitly, they are easy to conceptualize and manipulate. On the other hand, this approach is limited by its spatial discretization and complexity of algorithms required to manage any non-uniform tiling scheme. Constructing occupancy grids can therefore be very expensive in terms of both memory and time [4]. To adequately model a complex environment, the resolution of the occupancy grid must be matched to the spatial resolution of relevant parts of the environment, which can require tremendous memory capacity and heavy computational loads [4]. An important consideration for the implementation of this approach in fixed computational hardware is the *a priori* uncertainty about the size and resolution of memory needed for general navigation.

Topological maps are more compact representations in which spatial relationships between relevant locations in the environment are modeled by means of a graph structure [29]. Accordingly, a map is a graph of nodes representing landmarks in the environment and links that represent topological adjacency. In most cases,

precise metrical information about the links is not preserved and locations that are not visited are not represented. For example, vision-based topological approaches learn the visual scenes associated with visited places in the environment. The perceptual decision that certain features in the space constitute a landmark is critical and dramatically affects the performance of the system. In some cases, landmarks are based on a sensory-based decision (e.g. visual scenes of critical junctions in a maze environment), but in other algorithms nodes are simply locations that are spaced equally from each other [30]. Topological maps are chosen to be qualitative representations of the world, so that they are neither affected much by quantitative errors in metric information nor by “unexpected” complexities in the environment (i.e., changes in dimensionality, such as going up/down stairs in a building to another level). Additionally, the complexity of the learned graph reflects the complexity of the observed world, leading to a more optimal use of time and memory resources. One limitation for this approach is its reliance on the perceptual systems of the robot (i.e., a place may not be recognized simply due to a change in orientation of the robot compared to the map-associated view. [4]. Compared to the occupancy grid approach, which can be very inefficient (since much of the grid will represent locations of indeterminate occupancy), the topological approach allocates its memory only to state-action trajectories that it has experienced.

Because these two paradigms exhibit complementary strengths and weaknesses, several models have been put forth to integrate both representations into hybrid systems [31, 32]. So far, however, most of these systems are not as robust, flexible, or adaptable as the biological examples we strive to emulate [32].

To provide a better solution for the problem, researchers have begun adopting biologically-inspired models which directly emulate mammalian navigation abilities [4, 10]. To date, all of the modeling work that falls into this category have adopted biologically-inspired algorithms and have implemented them on general purpose PCs or special purpose digital hardware like FPGAs or DSPs. Although such solutions have benefited from biologically-inspired algorithms, the choice of digital hardware as the platform for implementation leads to systems that suffer from the following drawbacks:

- While digital technology provides a computational environment that is very robust against noise and provides perfect matching between all computational units. These features allow noise-sensitive computational algorithms to operate successfully on digital circuitry. In neurobiology it is generally agreed that the computational elements (neurons) operate somewhat stochastically and suffer from noise both due to inherent mismatch and cross-talk from neighboring neurons. Despite this noisy environment, the neural circuitry in the brain functions properly, implying that neurobiological algorithms are very robust against noise. Similar to biological circuitry, analog circuits suffer from unpredictable mismatch between its elements; this feature forces the designer to confront the mismatch issue and develop solutions that will be biologically-plausible.
- Power consumption is an important consideration if the objective is to build micro- aerial vehicles; digital circuits are known to use more power to obtain a large noise margin. , whereas neuromorphic analog hardware

examples commonly consume power around three to four orders of magnitude less than traditional approaches with digital circuits.

Some neurophysiological data revealing the presence of structures in the brain that contribute to the SLAM behavior in mammals became available after the completion of this work and were not taken into account when designing these systems.

In the work presented here, we target a biologically-realistic solution for the SLAM problem by adopting biologically-inspired computational algorithms and use neuromorphic analog VLSI circuits to construct the system.

1.3 Thesis Overview

The thesis is organized as follows: in Chapter 2 we review the neurophysiological data from brain structures involved in spatial navigation. We then discuss various proposed models for these neural structures in Chapter 3. Chapter 4 addresses our implementation of a neuromorphic head direction cell system. In Chapter 5 we discuss how sensory information can be linked with head direction information to correct for drift errors of the head direction cell system. In Chapter 6 we discuss our implementation of a neuromorphic grid cell system. In Chapter 7 we discuss how grid cells can be used to create neuromorphic place cells. Finally, Chapter 8 concludes this work with a summary of the findings and proposed future directions for this work.

Chapter 2 : Biological Background

Space and how it is represented in the brain has been the focus of much research for many years. The discovery of place cells by John O'Keefe and colleagues in 1971 [11] opened the door to understanding how the brain might provide this capability. Recording from the hippocampi of freely moving rats exploring an environment containing food, liquid and other objects, they discovered neurons in which activity was correlated with the location of the rat in the environment. These "Place Cells" [11] fire only when the animal is in a specific region of the environment, the *place field* of the cell. With neurons exhibiting spatially correlated activity, the hippocampus seemed to be a good structure for storing a *map* of space, a discovery that supported the theory proposed by Tolman in 1948 which suggested that rodents navigate using cognitive maps [33]. Since their discovery, the hippocampus and neighboring brain structures have been the subject of a huge body of research, leading to the discovery of other neurons which also show spatially-correlated activity, head direction cells and grid cells [12, 19]. In this chapter, we will review the neuroanatomy and neurophysiology of the main structures that have shown spatially-correlated activity in the brain of the rat (i.e., Head Direction Cells, Grid Cells, and Place Cells).

2.1 Head Direction Cell System

Head Direction (HD) cells were first recorded by Ranck (1984) from the rat postsubiculum (poSC). They are a population of neurons that individually respond to different orientations of the animal's head with respect to its environment. Subsequent studies have shown the presence of HD neurons in other areas of the

brain of the rat as well [14]. Furthermore, HD cells have been found in the brains of other mammals such as monkeys [34].

The latest neurophysiological data show the presence of HD cells in the following brain regions [14]:

- The deep layers of the postsubiculum (poSC) [12].
- The laterodorsal nucleus (LDN) of the thalamus [35].
- The dorsal striatum (caudate nucleus) [36].
- The retrosplenial cortex (RsCX) [37].
- The anterodorsal nucleus (ADN) of the thalamus [38].
- The lateral mammillary nucleus (LMN) [39].

The dorsal tegmental nucleus (DTN) [14].

2.1.1 Neuroanatomy of Head Direction Cells

While the functional relationship between the different brain areas containing HD cells is not fully known, the connectivity between some of the structures that are thought to participate in computing the HD signals are shown in Fig 2.1 [14]. Current theories suggest that cortical structures containing HD cells such as poSC, and RsCX mainly support allothetic-based orientation (information about the environment obtained through the location of the spatial cues; ‘Landmark Orientation’), whereas subcortical structures mostly support idiothetic-based orientation (based on self-motion signals e.g. angular head velocity integration [40]).

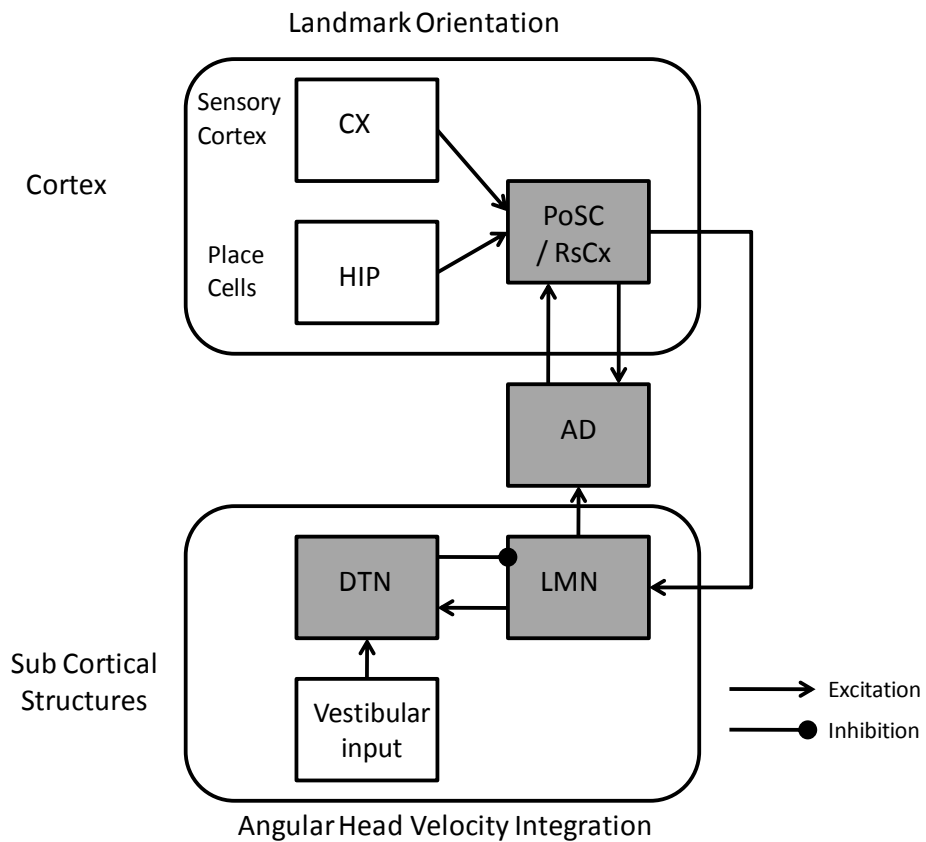



Fig 2.1. HD system Connectivity. The top panel shows the location of different structures exhibiting HD responses in the rat's brain (blue boxes) [14]. The bottom panel shows the functional interconnection between regions with HD cell and neighboring structures, shaded boxes denote regions that are known to contain HD cells. Abbreviations: AD, anterodorsal thalamus; DTN, dorsal tegmental nucleus; HIP, hippocampus; LMN, lateral mammillary nucleus; PoSC, postsubiculum; RsCX, retrosplenial cortex; CX, cortex, modified from [40].

2.1.2 Neurophysiological Properties of Head Directional Cells

The firing activity of a typical head direction cell is found to be maximal if the head of the rat is facing a certain orientation in the environment and drops as the head angle changes away from the preferred orientation. The firing rates of these cells vary as a function of the angle between the midline of the animal's head and its preferred direction. The distribution of activity typically has a Gaussian shape with a peak at the preferred orientation [14, 41]. Fig 2.2 shows a response field for a HD cell. The width of the response field varies between the cells but it is usually on the order of $60^\circ - 90^\circ$ [42].

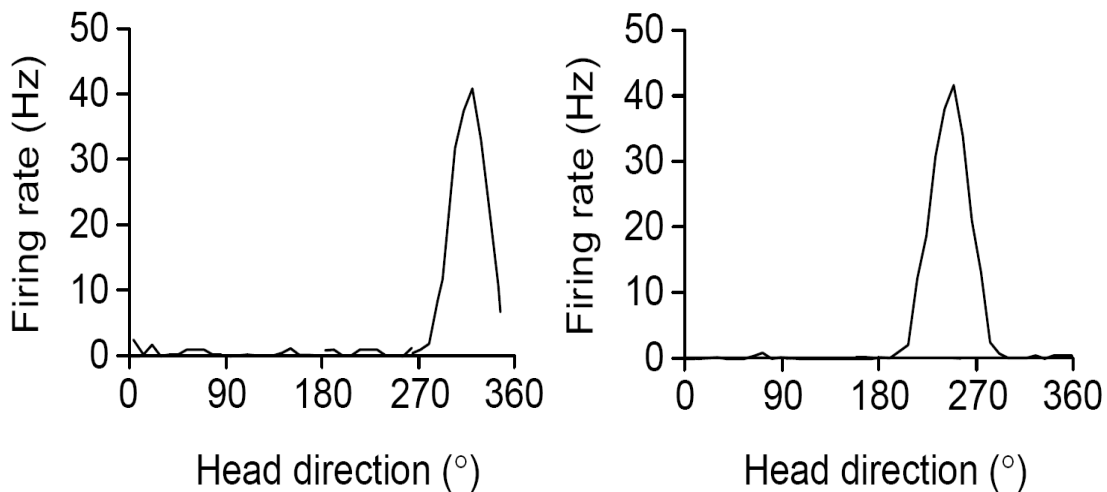


Fig 2.2. The receptive fields of 2 HD neurons from the rat postsubiculum, adapted from [14].

Topological Features: Anatomical data show that head-direction cells are not topographically organized. That is, two distinct head-direction cells with neighboring preferred directions are not necessarily neighboring units within the network [12].

Stability of Representation: One very important feature of HD cells is the stability and consistency of the representation; the difference between preferred directions of different HD cells remains constant, despite cue manipulation, disorientation of the animal, or exposure to different environments [43].

2.1.3 Head Direction Cell Determinants:

Allothetic vs. idiothetic cues: While compass-like, HD cells do not appear to be based on the Earth's magnetic field; neurophysiological data show that HD cells can be controlled by salient cues in the environment [43]. The recorded neurophysiological data show that the activity of HD neurons is sensitive to allothetic cues; the rotation of cue cards on the walls of the environment results in a corresponding rotation in all of the HD cells' preferred directions [14, 43]. The HD cells are also sensitive to idiothetic cues; if an animal enters an environment in complete darkness, the preferred-direction of the head-direction cells are carried forward from the previous environment, and they continue to show normal directional firing [3].

Angular head velocity: While head-direction cells are primarily correlated with the animal's spatial orientation, the head angular velocity plays an important role in modulating the firing activity. Experiments show that most of the cells in the lateral mammillary nuclei (LMN) are strongly correlated with both head direction and angular velocity [39]. In contrast, the anterodorsal nucleus (AND) head-direction cells fire proportionally to the magnitude of the angular velocity and can shift their preferred direction as a function of the head velocity to anticipate the head direction [38, 44, 45].

Direction of Head Rotation: In the postsubiculum (poSC), there are cells whose activity depends on the sign of the angular head velocity as well as the current head orientation. For example, a cell tends to fire more if the head is facing a certain direction and rotating clock wise whereas anti-clockwise rotation result in little or no activity in the same cell [12, 41].

2.1.4 From Head Direction Cells to Navigation:

HD cells are thought to participate in the computation required for spatial navigation because they maintain an ongoing estimate of the orientation of the animal in the environment. They use both allothetic and idiothetic information to maintain the estimate. In total darkness, the HD cells' preferred orientation suffers from drift [46]. This is hypothesized to occur because the data available in total darkness is only the angular head velocity and integrating this information to estimate the orientation may suffer from a variety of systematic errors and biases. If the lights are turned ON, however, the HD cells recover from any drifts and "reset" to their original preferred direction. In addition to the HD cells described above, cells that anticipate the animal's "future" direction have been discovered in the anterodorsal nucleus (AD). During head turning, AD cells shift their preferred direction as if to temporally "anticipate" head directions by a time delay which has been found to be between 0ms and 100ms [14].

Having these characteristics, the HD cells are important components for solving the navigation task since they maintain the current head orientation estimate by both integrating head velocity and by using sensory input to reset the preferred direction in case of drifts or conflicts.

2.2 Grid Cells

In 2005 Hafting and his colleagues reported the presence of cells in the medial entorhinal cortex (MEC) of the rat that fire repeatedly as the animal explores the environment. These cells are organized in a way that each cell has a firing field organized as a hexagonal grid over the environment. These grids are independent of sensory cues in the environment and the firing is solely dependent on the rat being in one of the response fields of the cell. They named these cells *grid cells* [19]. This discovery was striking because these so-called “grid cells” provide a good candidate for path integration circuitry that feeds the hippocampal formation. The regular nature of the activity distribution was verified by spatial autocorrelation analyses, which for all cells showed a tessellating pattern similar to that of the original rate maps. Fig 2.3 shows the trajectory maps, rate maps, and spatial autocorrelograms of cells recorded from the MEC [19].

Grid Structure: The geometrical structure of the grid firing field is a regular hexagon; the central peak of the autocorrelogram of the firing of a single cell surrounded by six equidistant peaks forming the vertices of a regular hexagon.

- Within each firing grid, the distance from the central peak of the autocorrelogram to the nearest six peaks is nearly constant.
- Hexagons of equidistant firing peaks were formed at multiples of the distance to the nearest hexagon, implying that the pattern was regular across the entire field.
- The angular separation of the vertices of the inner hexagon was in multiples of 60 degrees

- The locations of the grid vertices are stable across multiple recording trials

If the environment gets expanded, the number of activity nodes increases, but the density remains constant [19, 47].

Grid Topographical Organization: The grid fields of neighboring neurons share a number of metric properties (e.g. spacing, orientation, and field size), however, some of these features varied for cells at different layers in the MEC; the spatial frequency of the grid increases going from the ventral to the dorsal axis of the MEC [47].

Grid Development in Novel Environments: In novel environments, the grid patterns form and stabilize rapidly, as first as the first passage through the field. This suggests that the grid patterns are based on pre-existing network mechanisms [47].




Fig 2.3: Grid Cells, Trajectory maps (left), rate maps (middle) and spatial autocorrelograms (right) for three cells recorded from the MEC. The peak firing rate for each cell is indicated from [19].

2.3 Place Cells

In the following sections we will review the main features of place cells, showing their anatomical and physiological characteristics and describing the main theories for their functionality. We will then discuss the grid cell system and how previous models have been adapted to explain their relationship to the cognitive map theory.

2.3.1 Neuroanatomy

Location-sensitive neurons have been discovered in the following brain areas of freely-moving rats:

- The hippocampus (DG and CA1 – CA3) [11, 48].
- The entorhinal cortex (EC) [49].
- The subicular structure (Subiculum, Parasubiculum, Presubiculum) [50].

Extensive study of these areas led to the discovery of place cells and grid cells. To characterize their basic functionalities, it is important to study their anatomical interconnections. Fig 2.4 is a simplified representation of the mutual projections between these areas.

The hippocampal formation contains the hippocampus, the entorhinal cortex (EC), the subiculum (SC), the parasubiculum (paSC), and the presubiculum (prSC), whose dorsal part forms the postsubiculum (poSC). The hippocampus includes the dentate gyrus (DG) and the hippocampus proper (or *cornu ammonis* CA) made of 4 subregions CA1 - CA4, but CA1 and CA3 are the most distinguishable areas [51].

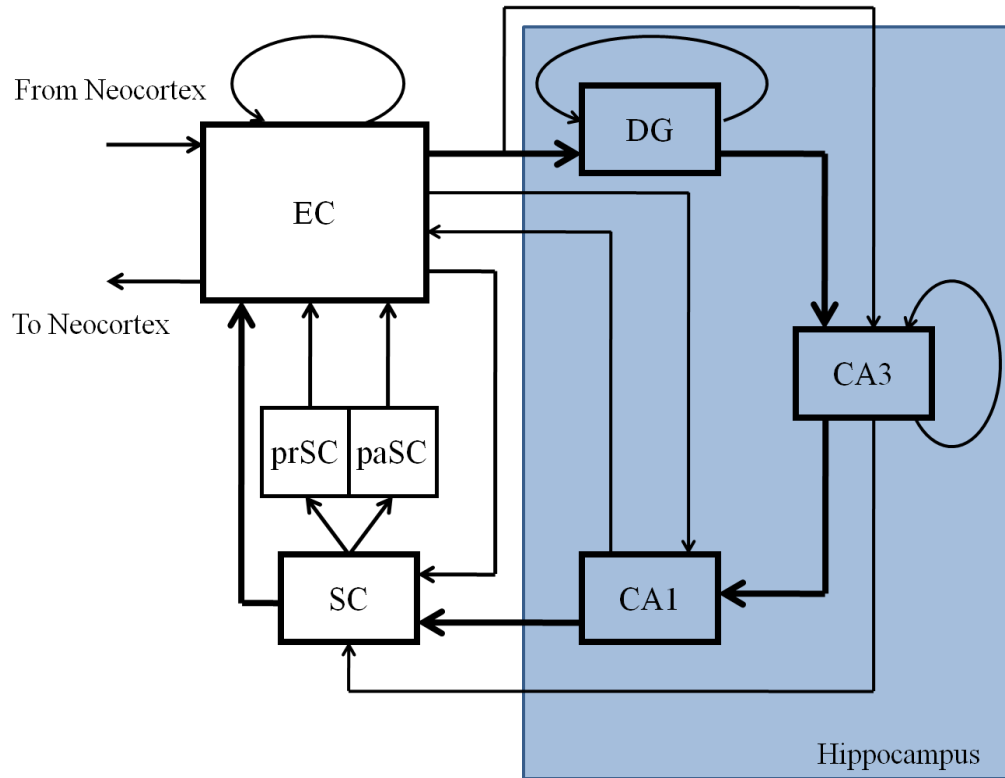


Fig 2.4. Anatomical interconnections between the different structures in the hippocampal formation. EC: entorhinal cortex, DG: dentate gyrus, SC: subiculum, paSC: parasubiculum, prSC: presubiculum, the hippocampus proper consists of the DG and CA3-CA1 areas (shaded area). The hippocampal intrinsic loop is shown in bold: **EC → DG → CA3 → CA1 → SC → EC** adapted from [4].

2.3.1.1 The Hippocampal Formation

Hippocampal afferents: Two major inputs enter the structure [51]:

- The perforant pathway carries highly processed multimodal sensory information coming from neocortical areas projecting through the EC.
- The fornix fiber bundle carries signals from subcortical areas (the thalamus, the hypothalamus, the brainstem, and the amygdala) to the hippocampus. It is thought that this input carries arousal and emotional information which modulate the ensemble hippocampal activity and seem to be responsible for generating the hippocampal theta rhythm.

Intrinsic hippocampal circuit [51, 52]: The propagation of neocortical inputs through the hippocampal formation is believed to be the main processing pathway in this structure:

- The highly processed information from neocortical areas reaches the entorhinal cortex which project to DG granule cells, CA3 and CA1, and the subiculum. Furthermore, EC exhibits intrinsic connections.
- The DG sends efferents to CA3 via the mossy fibers, these projections are very selective: each granule cell projects approximately onto about 14 CA3 cells only. The DG has also intrinsic projections.
- CA3 cells form a dense recurrent network through the Shaffer collaterals. Moreover, they send projections onto CA1 and the subiculum.
- CA1 neurons send their output to entorhinal as well as subicular cells via the angular bundle.
- SC projects onto the entorhinal cortex.

The hippocampal circuit has been approximated by a simple feed-forward loop known as the trisynaptic loop [51]; information enters the loop via EC, proceeds towards the DG, then to CA3 and CA1, and finally arrives at SC which closes the loop by projecting back to the EC.

Hippocampal efferents [51]: The subiculum provides the main output of the hippocampal formation by projecting to the deep layers of the entorhinal cortex (dEC). From dEC, information is sent to a variety of cortical areas. It is important to note that the CA3 and CA1 regions also send an output directly to subcortical areas.

2.3.1.2 The Entorhinal Cortex

The entorhinal cortex EC is considered to be the main “cortical gate” of the hippocampal formation; it receives sensory signals from the neocortex and conveys this information to the hippocampus via the perforant path. It is a six-layered structure divided into superficial (sEC), layers I, II, and III, and deep (dEC), layers IV, V, and VI. In rats, there is also a distinction between the lateral (LEC) and the medial (MEC) areas of the entorhinal cortex [51].

EC afferents [49, 51]:

- Afferents to sEC are thought to form the main source for the information that enters the hippocampus. An important cortical input to sEC is via the perirhinal (peRH) and parahippocampal (paHI) cortices, carrying information from most of the associative areas (visual, auditory, and somatosensory), as well as from the parietal, temporal, frontal, and retrosplenial lobes. Another major input to the rodent sEC comes from the olfactory system.

- The dEC receives cortical afferents from the limbic system, the retrosplenial cortex (which projects almost exclusively onto MEC), and from the frontal cortex. The role of this input on the hippocampal formation activity is not yet clear.

EC efferents [49]: The entorhinal cortex projects mainly to perirhinal (peRH), infralimbic, prelimbic, orbitofrontal, and olfactory cortices. Secondary EC projections reach the temporal, frontal, retrosplenial, occipital, and parietal regions.

EC hippocampal afferents [51]: The EC receives input from CA1 as well as SC, especially layers I – III of the MEC. Another important input to sEC comes from presubicular and parasubicular cortices.

EC hippocampal efferents [51]: The EC projects to the hippocampus via the perforant path that arises from the superficial layers. Layer II mainly synapse onto the dentate gyrus and CA3, but also sends projections to the subiculum SC. Layer III primarily projects to CA1 and SC. EC sends also a (weak) output to the presubicular and parasubicular cortices.

EC intrinsic projections [49]: Internal links connect the deep layers of the entorhinal cortex (dEC) to the superficial ones (sEC). In addition, the lateral entorhinal cortex (LEC) strongly projects to the medial entorhinal area (MEC).

2.3.1.3 The Subicular Complex

It consists of three principal subregions; the subiculum (SC), the presubiculum (prSC), and the parasubiculum (paSC). It is considered to provide the main output of the hippocampal formation back to the EC. Although it exhibits significant

connectivity with other areas, the processing role of this structure in the hippocampal circuit is still not clear [51].

Subicular afferents: SC receives input projections from CA1 and sEC. The presubiculum (prSC) receives afferents from the subiculum, the posterior parietal cortex, the temporal lobes, the retrosplenial cortex, the anterodorsal thalamic nucleus (ADN), the laterodorsal thalamic nucleus (LDN), and the lateral mammillary nuclei (LMN). The parasubiculum (paSC) is reached by projections from the subiculum, the presubiculum, and the retrosplenial cortex.

Subicular efferents: prSC and paSC project to layers III and II of EC, respectively. SC projects to the LEC, MEC, prSC, paSC, the medial prefrontal cortex, the retrosplenial cortex, the septal complex, the nucleus accumbens (NA), the mammillary nuclei, the amygdala (AM), the hypothalamus, and the thalamic nuclei.

2.3.2 Neurophysiological Properties of Place Cells

The spatial plot of mean spiking rate of a place cell forms a two-dimensional place field in space with a peak firing frequency at some preferred location within the field, (usually in the center) and with smoothly falling off edges in all directions [5, 11, 12, 16, 21, 47, 53, 54]. Fig 2.5 (left panel) shows the firing field of a place cell in addition to the trajectory of the rat in the arena. The right panel shows firing rate maps for the place fields of 4 CA1 cells recorded simultaneously with a tetrode in a freely moving rat exploring a square arena [47, 55]. The darkest regions indicate the areas in which the cells respond maximally. When the animal is visiting the white marked areas, the cells remained silent. Place cells tend to cover the environment densely and uniformly with highly overlapping place fields. As a consequence, space

coding may be achieved by taking into account the ensemble firing activity, rather than single cell activity.




Fig 2.5. Place Cells in the hippocampus: a) Spike locations (*red*) are superimposed on the animal's trajectory in the recording enclosure (*black*), adapted from [47]. b) Firing rate maps of place fields from 4 CA1 cells, the peak firing rates vary between 4.8, 22.8 Hz.; adapted from [55].

Many behavioral, neurophysiological, and neuroanatomical studies have been done since the discovery of place cells in 1971 to understand the space coding properties of hippocampal place cells, to study the effect of changes in the environment on the hippocampal representation, and to identify the specific functionality of different anatomical areas. The common experimental setup used is to record from neurons of the hippocampus of animals engaged in a navigation task (e.g. randomly exploring new environments or searching for food). Various environmental setups have been used in such experiments (e.g. water maze, cylindrical and rectangular arenas, T, Y, radial arm mazes) [15, 16, 54]. These experiments have produced a large amount of data that serves to show some insights about spatial

capabilities of rodents [3, 54]. There is general agreement that the hippocampus is involved in spatial behavior, however, there is no agreement on the specific role of the hippocampal formation. Is it used only to represent and learn space [5] or it is involved in higher order episodic memory functions (many of which have space as an inherent constituent [3]). A review of some of the most relevant properties of the hippocampal place cells is presented below.

2.3.3 Hippocampus' Place Cells (DG, CA1, CA3)

Place fields tuning shapes: A typical place field can be roughly approximated by a two-dimensional Gaussian with the peak representing the preferred location for the cell [55]. Cells coding for peripheral locations show crescent-shaped fields hugging the arena walls

Place field formation: Establishing a place field representation in a novel environment takes a relatively short time; experiments show that it takes ~ 10 – 30 minutes of exploration to generate stable place fields, some cells develop their firing fields quicker (a few minutes) than others (30 minutes). It has also been shown in some experiments that cells can exhibit place coding as robust on the first visit to their field as on subsequent visits, i.e. no experience was required to tune up their firing [56].

Place field experience-dependent reshaping: As the animal experiences a route several times, CA1 cells tend to asymmetrically expand their (initially symmetric) field and to shift their field center backwards with respect to the rat's direction of motion. This place field expansion is related only to the experienced environment

and does not transfer to the other place fields formed by the same cells in different environments [57].

Place field directionality: Experimental data shows that place cells have directionally-independent place fields; the firing activity does not depend on head direction when the animal randomly moves over two-dimensional, open environments. On the other hand, place cells have been shown to have directional place fields in experiments where the rat moves along fixed trajectories, such as on linear track mazes [54].

Place field distribution: CA3-CA1 firing fields tend to cover the whole environment uniformly, without differentiating areas with respect to their potential relevance (e.g., home or feeder location). A given place cell can have place fields in several but not all environments. A very large number of cells participate in the representation of an environment and thus a dense population of highly overlapping place fields [5, 16] results.

Place cells topological features:

- CA3-CA1 place cells are not topographically organized, there is no relationship between the physical place field topology and the anatomical place cell arrangement; two cells coding for neighboring locations in space are not necessarily anatomically adjacent, in fact, they are most likely not adjacent [16].
- Experiments with recordings in different environments show that the spatial relationships between place cells and their place fields are not preserved across environments; if two place cells have neighboring place fields in one

environment, they may not have neighboring place fields in a different environment [54].

Place code replaying during sleep [58]: Recordings from the hippocampus of animals that were performing spatial tasks while they are sleeping show that:

- Cells that were activated during the last session are more active during subsequent sleep episodes than others.
- Cells with temporally correlated activity during the most recent sessions, exhibit correlated reactivation during sleep.

2.3.4 Subicular Place Cells

SC neurons, similar to the hippocampal ones, show spatially-dependent firing, however, their firing properties are different from hippocampal neurons; they have the property of maintaining a similar place field topology across distinct environments Fig 2.6 shows examples of 4 subicular place cells recorded in two different arenas [59]:

- When recorded in two geometrically different arenas (e.g., a square box and a cylinder) with diverse visual cues (e.g., different wall colors), a typical cell in SC exhibits similar firing patterns (field location, firing rate, and field size) in the two recording chambers. For instance, if a cell codes for a location near the east wall of the cylinder, it will also respond maximally when the animal is at a location near to the east wall of the square box.
- When recorded in two square environments of different size, a cell in the SC tends to show similar overall firing patterns such that its place field tends to expand or shrink to fit the size of the current environment. SC cells have

broader receptive fields than place cells in the hippocampus proper. Moreover, SC cells show directional tuning even in situations in which CA3-CA1 do not (e.g., open-field spatial tasks). Also, subicular cell activity seems to be strongly modulated by self-motion information.




Fig 2.6 Subicular Place Cells recorded in two arenas, it is to note that the cells exhibit topologically related firing in both arenas, adopted from [59].

2.3.5 Place Field Determinants

Several studies have tried to identify what information controls place cell activity. Most experiments have been performed using visual stimuli, however, there is a hypothesis that the hippocampus does not rely on any one sensory modality and this is supported by experiments done on visually impaired animals or in complete darkness. The hypothesis is that the hippocampus tries to build an allocentric representation of the environment based on complimentary information from multiple sensory modalities [54, 60].

Allothetic determinants: Neurophysiological data show that hippocampal place fields strongly depend on external sensory cues (e.g. visual landmarks in the arena). Rotating cues during the experiment often result in place field rotations.

Idiothetic determinants: Despite their dependence on external sensory signals, place cells exhibit location selectivity in the absence of these cues. This suggests that hippocampal cells are also driven by internal movement-related signals, since these are the only available cues in the absence of sensory cues (i.e., environmental landmarks). A strong support for the idiothetic basis for place fields comes from experiments where:

- Place cells have unchanged place fields after the lighting is switched off.
- When the animal is introduced in the arena in complete darkness, place cells are established and subsequently persist after lighting is provided.

2.3.6 The Theta Rhythm and Hippocampal Activity:

The hippocampal EEG from rats exhibits a regular sinusoidal signal of 7 to 12Hz called the theta rhythm during locomotion (e.g., walking, running, swimming, jumping), as well during passive locomotion of the animal, otherwise, the hippocampus shows a non-rhythmic pattern [17]. The theta rhythm is observed also during sensory scanning as well as in REM sleep. There exists a phase correlation between the theta rhythm and hippocampal place cell firing. As the animal goes through the place field of a cell on a linear path, the theta phase at which the cell discharges shifts systematically; every time the rat enters the field, the cell starts firing at the same phase late in the theta period, then as the animal proceeds through the field, the firing tends to occur earlier and earlier in the cycle. This phase shift

phenomenon is termed *phase precession* and is believed to provide an estimate of the animal's position inside the place field of the cell. DG, CA3 and CA1 cells all exhibit this phase precession [17]. Neuroanatomical results suggest that the medial septum might be involved in modulating temporal processing in the hippocampus. In particular, the cholinergic and GABA-ergic septal efferents seem to be responsible for driving the theta rhythm [61].

2.3.7 Hippocampal Synaptic Plasticity

To accomplish its proposed role, the hippocampus has to provide rapid on-line learning and linking of spatio-temporal patterns extracted from its inputs. Activity-dependent synaptic plasticity in the hippocampal formation offers a suitable neurochemical substrate for implementing associative learning [62].

Long-term potentiation (LTP): is the general term used to describe the synaptic modification that is the basis of hippocampal learning. LTP is defined as a persistent potentiation of the synaptic efficacy that will persist for at least an hour, but can last for hours or days under certain conditions [62]. The most studied mechanism underlying synaptic, long-term potentiation is a NMDA-mediated LTP [63, 64]. It depends on the causal relationship of the presynaptic input spikes and the postsynaptic output spikes. If a spike from the presynaptic (input) neuron precedes a spike by the postsynaptic (output) neuron by a small time interval (milliseconds), the synaptic efficacy between these two neurons is increased. On the other hand, if the postsynaptic neuron precedes the presynaptic neuron, either the synapse is unchanged (no LTP) or long-term depression (LTD) occurs, i.e. the synapse efficacy is reduced.

Pharmacological studies support the idea that the hippocampal NMDA-mediated LTP is relevant for spatial learning; blocking the NMDA receptors results in impaired spatial learning capabilities. Mice that have been genetically engineered to have deficient LTP exhibit unstable CA1 place fields between recording sessions and are severely impaired in learning spatial tasks like the water maze. Recently it has been proposed that NMDA-dependent hippocampal synaptic plasticity plays a more general role in episodic memory rather than merely providing a basis for spatial learning [62].

2.3.8 Memory Space or Space Memory?

Although most of the literature on hippocampal activity comes from spatially-dependent behavior, experiments in other animals have shown that it may be involved in more complex functions [3, 65].

Non-spatial place cell correlates: hippocampal cells have been studied in non-spatial tasks such as in odor and auditory discrimination experiments. Results show that the hippocampus is also involved in this kind of non-spatial process. Experimental data has shown the influence of the *context* on cell activity: a cell responding to a stimulus in a specific task, might not encode the same stimulus in a different context. The fact that place cells may be task-sensitive has been explicitly demonstrated by a series of experiments in which rats were trained to solve different tasks within the same environment[66]. In this experiment, the place field varied with the task (i.e., context). These findings indicate that the hippocampus might play a more general role than just being responsible for tagging a specific location in the environment; hippocampal cells might be involved in a more general class of memory (e.g.,

episodic memory) related to perceptual, behavioral, and reward-related variables. Nevertheless, space might be the contextual framework necessary to encode and link such memories.

2.3.9 Navigation:

In this Chapter we reviewed the biological features of the brain structures known to be involved in spatial navigation, the head direction cells, the grid cells, and the hippocampal place cells. Our focus is to understand the neurobiology underlying spatial navigation to build a neuromorphic system to mimic that behavior and implement it on a behaving robot. Many hypotheses have been postulated to interpret the available biological data and arrange it into a unified theory for how the brain handles the task of spatial navigation.

The interpretation with the strongest evidence from neurobiological data is that the rodents are relying on a cognitive map-like navigation system. Their spatial navigation system relies on both idiothetic and allothetic information to behave. The information about orientation (from head direction cells), displacement (from the grid cells) and the available cues or landmarks in the environment (from the different sensory modalities) are used to create a unified sense of location in space represented by the place cells. Navigation can be performed by relying only on the allothetic or the idiothetic data, however, it can be inferred that the brain uses both sources of information to perform spatial navigation.

Chapter 3 : Modeling the Navigation System (State of the Art)

Since the discovery of place cells in 1971 [11], much research has gone into creating a theory explaining their formation and use. The specific role of place cells in navigation was and still is under debate. Subsequent discoveries such as head direction cells [12] and grid cells [19] shed some light on the neural substrate of spatial behavior, however, there still remains many unanswered questions. In this chapter we will review models proposed for the formation and use of head direction cells, grid cells and place cells with an emphasis on models that are closest to biological realism and are suitable for implementation using neuromorphic VLSI circuits.

3.1 Modeling Head Direction Cells

Several models have been proposed for the creation of the head direction response characteristic. The primary concept governing most of these models is the integration of idiothetic signals (e.g., vestibular and proprioceptive information) to steer the activity of the network between different attractor states. In addition to being a sensory integrator (mathematically integrating head angular velocity to orientation), the head direction system acts as a memory for orientation in the case of deprivation of sensory cues (e.g. navigation in darkness).

3.1.1 (McNaughton et al., 1991)[43]:

Bruce McNaughton and his colleagues proposed one of the earliest models for the head direction cell system in 1991; the model employed a mechanism to integrate vestibular signals to update the directional representation using a linear associative

neural network to functionally model the mathematical integration of angular velocity. The inputs to the associative network are the current orientation and the current angular velocity and the network output provides the next orientation estimate. The model has a population of cells (H) encoding the current orientation, a population of cells (H') encoding the angular velocity, and an intermediate group of neurons in which all possible joint values HH' are represented by a set of linearly independent vectors to ensure linear independence of the inputs. Each HH' cell provides the next angular orientation for a specific heading and a specific angular velocity i.e. the HH' population represents a neural associative look-up table of angular velocity integrals that determines state transition and projects onto H cells properly. To code for the calibration of directional cells by means of external cues, they proposed a population of local-view cells projecting to the H population that rely on Hebbian learning to correlate local-view cell activity to H cell activity.

The hypothesis postulated by McNaughton *et al.* was the first plausible theory explaining both the update mechanism underlying head-direction cells and the influence of extrinsic signals. They addressed the problem on an abstract level, however, without accounting for neural connections and dynamics. Moreover, they do not report any implementation of the model that was tested in a closed-loop system.

3.1.2 (Skaggs et al. 1995)[67]:

Skaggs and colleagues proposed a one-dimensional attractor network scheme to model the head direction cells. In this model, head-direction cells are coupled by intrinsic connections, such that nearby cells are linked by strong excitatory synapses,

and distant cells are connected by strong inhibitory projections. The activity in such a network is represented as an activity bump that is moving around the network in an angular velocity dependent fashion. Fig 3.1 shows a schematic of the proposed network.

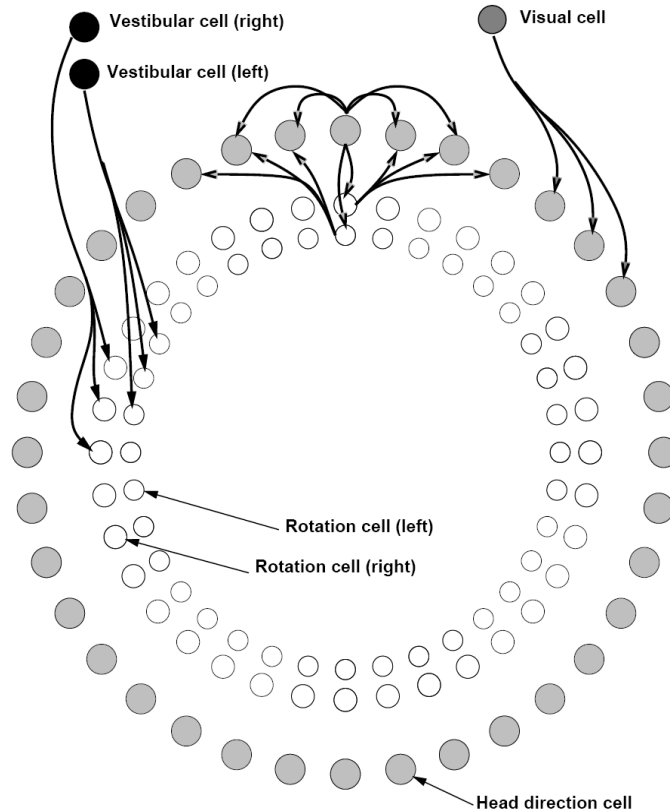


Fig 3.1. Head Direction model from Touretzkey and Skaggs [67].

Due to the attractor model, at each time step the system has a stable representation of orientation consisting of a single localized cluster of active cells, or “bump of activity”. If an excitatory external signal is applied to cells on one side of the peak, the activity will gradually shift towards the side at which the input has been applied. The model has two populations of “rotation” cells; one group is responsible for

clockwise turns and the other for counterclockwise turns. Each rotation cell is active only if the current heading is equal to its preferred direction and the head is turning according to its preferred angular velocity sign. Clockwise rotation cells project to head-direction cells neighboring them on the right. Counterclockwise rotation cells project to head-direction cells neighboring them on the left. During clockwise turns, clockwise rotation cells will excite cells to the right of the current peak, and the bump of activity will shift rightward. The model has also a set of visual feature detectors, each of which responds maximally to a specific visual cue located at a specific angle. These neurons project to head direction cells and Hebbian learning is used to modify synaptic connections. The activity of visual feature detectors is used to correct head-direction activity. This model was the first in adopting the attractor concept and triggered a long line of subsequent models. One notable shortcoming of the original paper was the lack of simulations or implementations to validate its functionality.

3.1.3 (Redish et al. 1996) [68]:

This model uses a coupled attractor network representation for the interaction between the postsubiculum (poSC) and the anterodorsal thalamic nucleus (ATN) which is known to be reciprocally connected in the rat. The poSC and the ATN are both represented by a separate attractor network each consisting of two groups of neurons: One pool of excitatory units E , and one pool of inhibitory units I . Neurons in both pools have evenly distributed preferred directions. Each excitatory cell $e \in E$ with preferred direction θ_e projects strong exciting synapses to neurons in both E and I having preferred directions close to θ_e . Each inhibitory neuron $i \in I$ weakly inhibits all cells in both E and I pools, whereas cells in E and I having preferred directions

close to θ_i are inhibited slightly more. The dynamics of the attractor are such that both excitatory and inhibitory pools are characterized by a single stable state at each time.

Fig 3.2 shows a schematic for the connectivity between the pools of neurons.

The poSC attractor and ATN attractor networks are connected through their excitatory pools by a set of synapses called “matching connections” which connect between neurons having equivalent preferred directions. Moreover, a set of projections, namely left and right-offset connections, is responsible for updating the head-direction representation. Each excitatory unit i in the poSC with preferred direction θ_i has a left-offset connection to excitatory unit j in the ATN such that $\theta_j = \theta_i - 10^\circ$ and has a right-offset connection to excitatory unit k in the ATN such that $\theta_k = \theta_i + 10^\circ$; the offset has been arbitrarily chosen to be 10° in the model and is the same for all units. The weights of the offset synapses are modulated by the head angular velocity. During rightward head turns, the right-offset connections have a strength proportional to the magnitude of angular velocity, whereas left-offset connections have strength zero and the opposite situation occurs during leftward turns. The model has an inhibitory input to ATN coming from the mammillary bodies (MB). This input is assumed to be proportional to the magnitude of the angular velocity. MB cells work as a gain control mechanism to compensate for modulated offset connections. This allows the system to maintain the shape of the hill of AT activity nearly unchanged during rotations (otherwise, the combined input from offset and matching connections would distort it). A schematic of the entire system is shown in Fig 3.3




Fig 3.2. Two ring attractor network from [68].

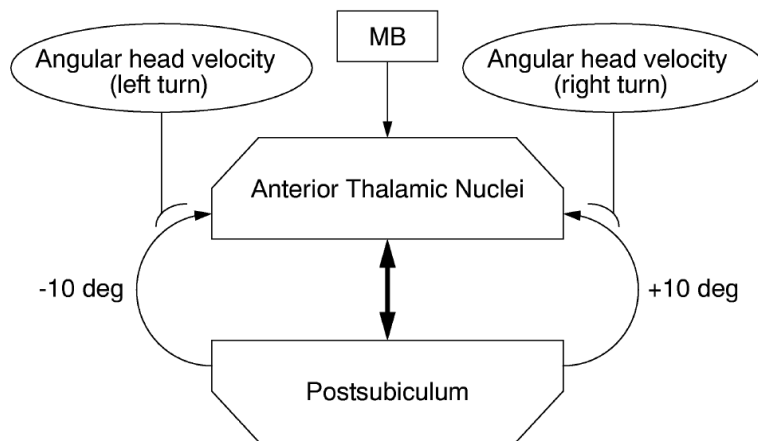


Fig 3.3. Schematic of HD Neural System [68].

This model showed biologically plausible directional tuning curves and was effective in tracking the orientation of the head. ATN activity anticipates poSC activity by a constant lead time of approximately 10ms (capturing the anticipatory property of real anterior thalamic neurons). On the other hand, a major problem for this theory is its failure to explain results from anatomical lesion experiments. After lesions to the postsubiculum (poSC), anterior thalamic cells (ATN) are still directional selective [69]. By contrast, the model predicts that lesions to poSC cells would disrupt the entire system. The model requires the existence of poSC → ATN synapses that can be dynamically modulated by head angular velocity, which is not known to be biologically-plausible.

3.1.4 (Goodridge and Touretzky, 2000) [70]:

This model takes into account the interactions of HD cells in three brain areas: poSC, ATN, and LMN. Fig 3.4 shows the proposed neuronal model.

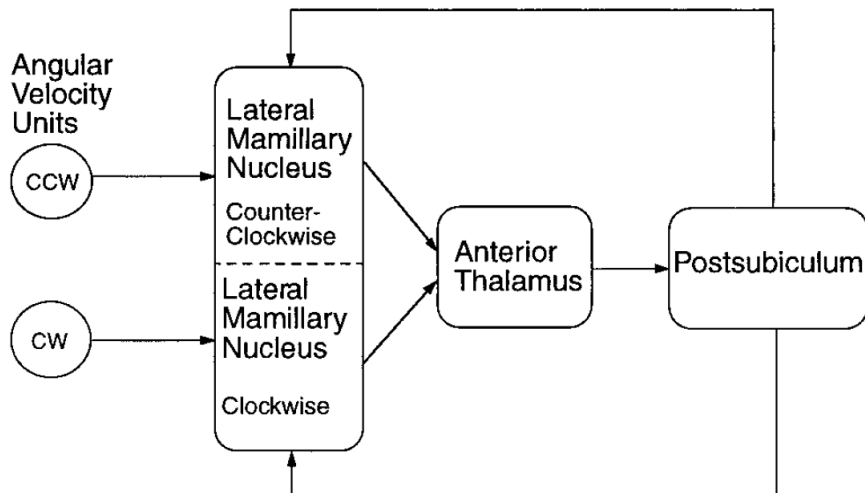


Fig 3.4. Neuronal Model from [70].

The postsubiculum poSC: It is modeled as a ring of neurons that generates attractor dynamics as bumps of activity amongst its neurons. Each element along the ring has a Gaussian tuning curve with its maximum activity at a preferred direction. The neurons in the ring are interconnected with excitatory connections that have weights with a Gaussian distribution in the difference between the preferred directions of the units. The model has a global inhibitory neuron that receives excitation from all the elements in the ring and project back inhibition to all of them. The poSC receives its input from the ATN.

The Lateral Mamillary Nucleus LMN: The LMN is modeled as two independent attractor networks: one for clockwise rotation and the other for counter-clockwise rotation. Each attractor network has lateral excitatory recurrent connections and a global inhibitory unit. The LMN networks receive excitatory input from the angular velocity units which are active during either clockwise head rotation or counter-clockwise head rotation. The LMN units send excitatory projections to the ATN unit.

The Anterior Thalamic Nucleus ATN: the model for ATN in this model is a ring of neurons with no recurrent connections. The inputs to the ATN are excitatory offset projections from the LMN units. A unit i in the ATN with preferred direction of θ° receives input from the unit $i+1$ in the clockwise LMN with preferred direction of $\theta^\circ + \Delta\theta$ and from the unit $i-1$ in the counter-clockwise LMN with preferred direction of $\theta^\circ - \Delta\theta$. The ATN units send their output excitatory projections to units in the poSC attractor network which have the same preferred direction.

To summarize, the poSC units' exhibit attractor dynamics in the form of a bump of activation, the current location of the bump will be on the units which have

the same preferred direction as the current head orientation. If the head is not turning, units in both LMN attractor units will have a stable bump of activity along the neurons with the same head orientation as the current head direction, since there is no angular velocity, the firing rates in the two LMN attractor bumps will be the same. The ATN receives offset excitatory input from LMN units, which will cause a bimodal response on the neurons of the ATN, one coding for the proposed new head orientation had the head been turning clockwise and the other for counter-clockwise. The projection from the ATN units to the poSC will be symmetric and hence the bump on the poSC will not move.

If the head is turning clockwise, the clockwise LMN will be more active than the counter-clockwise LMN and the response in the ATN will be bimodal but not symmetric. In this case, the clockwise ATN will be stronger than the counter-clockwise one, and hence the poSC will receive a stronger input driving clockwise motion of the bump and the bump will turn clockwise.

This model captures a lot of the details of the biology like the interconnection between the poSC, LMN, and ATN. The LMN neurons have been shown to have direction dependent velocity modulation as modeled; the model also captures the poSC feedback projection to the LMN in the rats. The model has some drawbacks, however, the angular velocity units are not biologically realistic, because no such structure has been reported and no bimodal responses in the ATN have been reported so far. Finally, the model does not take into account any head-direction calibration based on external signals, although this could be added.

3.1.5 (Sharp et al., 2001)[37]:

In their model (shown in Fig 3.5), the network consists of two populations of HD cells one excitatory and one inhibitory. As in other models, neurons that represent adjacent directions are located next to one another.

3.1.5.1 Stable and Sustained Activity

The excitatory HD cells provide excitatory input to nearby HD cells, but provide inhibitory input to distant HD cells, via the inhibitory cell layer. In the absence of external influences, the population activity pattern will settle into an attractor state in which one HD cell will fire maximally, and neighboring HD cells will show progressively lower activity levels, as a function of distance from this maximal peak.

3.1.5.2 Angular Velocity Integration

To accomplish path integration, the system must first receive input about the animal's head angular velocity. This is provided by the angular velocity (AV) cells shown in Fig 3.5. These cells are tonically active when the head is still, but increase their firing rate during a preferred turn direction (clockwise or counter-clockwise) and decrease their firing rate during a turn in the opposite direction. This AV input is connected to the attractor network by a set of inhibitory cells that receive input from both excitatory HD cells and from the AV cells as shown in Fig 3.5. Two types of such AV-by-HD cells are shown. One type receives input from excitatory HD cells to the right and excitatory input from AV cells that fire at high rates during clockwise head movement; these cells, in turn, project onto excitatory HD cells that are located to the left. The other type receives input from excitatory HD cells to the left, and from

AV cells that prefer counterclockwise motion. These cells project onto excitatory HD cells located to the right. When the rat does not move, the two types of AV-by-HD cell are equally active, and so the activity packet in the HD cell layer remains stable. When the rat begins to turn clockwise, inhibitory AV-by- HD cells to the left of the activity packet increase their firing rate, while inhibitory cells on the right side decrease their firing rate. This causes the activity packet to shift to the right, estimating the new directional heading (i.e. that which resulted from the clockwise turn). During counter-clockwise turns the opposite is true.

3.1.5.3 HD cell system can be calibrated by environmental landmarks

In a familiar environment, the HD cell firing direction can be set by the position of environmental landmarks. HD cells also receive sensory inputs that can be strengthened in a Hebbian manner, as a result of experience. These sensed landmarks operate to calibrate drift in the integration of head velocity by associating the detection of the landmark with the position of the activity bump, so that in a familiar environment (where many associations have been made) each HD cell will have the same preferred direction each time the rat visits (i.e., landmarks will reset the HD network to the same state each time).

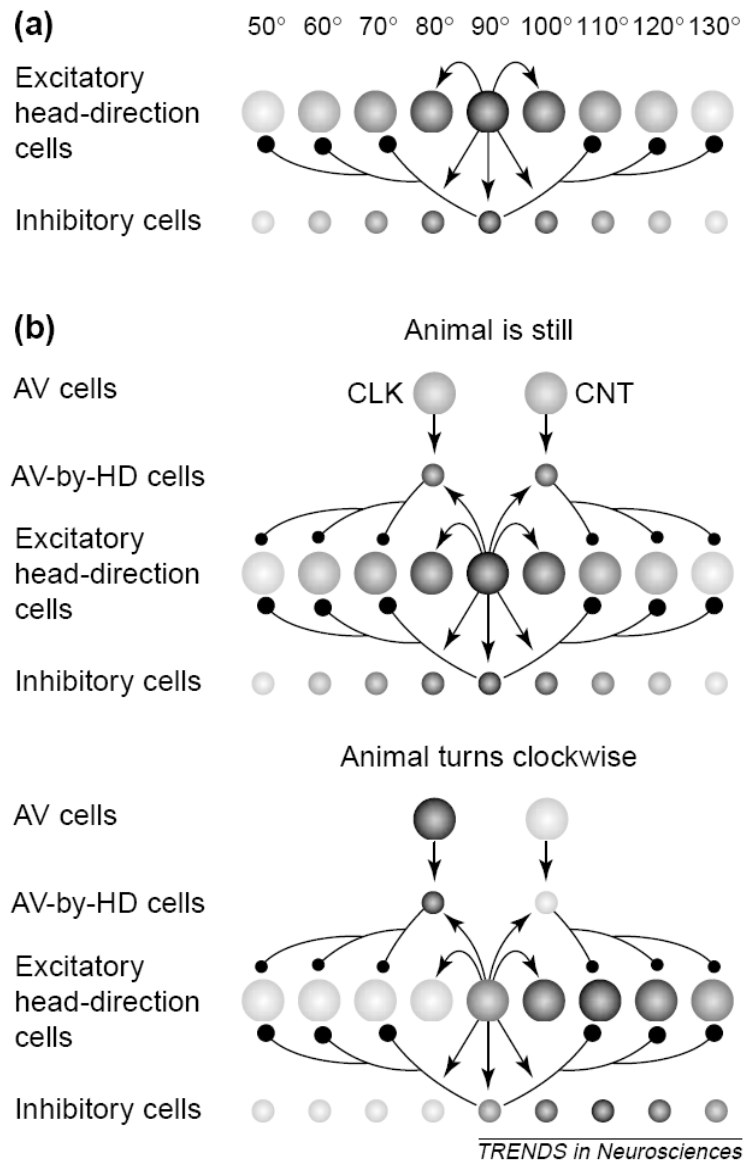


Fig 3.5. a) Stable bump, b) Bump rotation from [37].

3.2 Modeling the Grid Formation:

Following the discovery of grid cells in 2005 [19], many efforts were made to model the neural computation underlying the formation and activity of these grids in the medial entorhinal cortex (MEC). Current models for the grid field formation propose that MEC neurons path-integrate speed and direction provided by self-motion signals, whereas sensory information related to the environment is used for setting the initial parameters of the grid (2-dimensional phase) or calibrating it to correct for the cumulative error associated with the integration of velocity. There are two classes of models for the grid formation; one class suggests that grid formation is based on attractor dynamics in the MEC formation and the other suggests that grid formation is a result of constructive interference of subthreshold inputs to the network's neurons [47].

3.2.1 Continuous Attractor Models:

In this class, a single position in space is represented by an attractor state in the system. The network stores several attractor states associated with different locations and recalls any of them in response to sensory cues or path integration. When a large number of very close positions are represented, a continuous attractor (i.e., a hyperplane) emerges, which then allows a smooth transition between states according to the animal's trajectory. There are two important models in this class.

3.2.2 (Fuhs and Touretzky, 2006)[71]:




Fig 3.6. Grid pattern in the MEC layer forms instantaneously and makes a displacement to the right to follow the motion. In the origin of the white axes, a neuron is not firing at time $t = 0$, fires in the maximum of a grid node at $t = 140$, and is at rest again at $t = 290$ from [71].

In this model, the authors assumed that the MEC is topographically organized (i.e. neighboring cells become active together). The representation of a single location in space forms a stable grid pattern on the neuron layer. This type of activity pattern can be generated if each neuron has mutual excitatory connection with its nearest neighbors, mutual inhibitory connections with neurons at intermediate distance, and no direct connections with those far away. Using different sizes for the excitatory and inhibitory projections, grids of activity with different spatial frequencies appear in the MEC layer. At the border of the layer, the lack of balanced connections at the input of these neurons leads to over excitation, hence the grid pattern is not uniform and additional attenuation is required to overcome that.

To transform the grid pattern of activity in the MEC layer into a allocentric grid firing field for each neuron, the pattern of activity must be moved around on the

network following the movements of the animal. This is done by applying an input to the neuron proportional to the running speed of the animal when the rat runs in a preferred direction, which is different for each neuron. Fig 3.6 shows simulation results for this model, a stable grid pattern of activity forms spontaneously as a rightward motion is introduced the activity in the grid is displaced in the corresponding direction. Increased speed produces increased excitation and a faster displacement of the grid pattern across the neural layer.

This model's main challenges are the apparent lack of topography in grid phases of neighboring MEC neurons and the distortion of activity at the peripheral that was not reported in the neurophysiological literature.

3.2.3 (McNaughton et al., 2006)[72]:




Fig 3.7. A topographically arranged network serves as a tutor to train an MEC module with no topographical arrangement (*left*). Connectivity in a layer of MEC if the connectivity is rearranged topographically (*center*), the effective geometry is that of a toroidal surface (*right*) from [71].

McNaughton and his colleagues proposed in (2006) an alternative model that relies on the assumption that a topographically arranged network is present in the cortex during early development and serves as a tutor to train MEC cells. The tutor

network has scrambled connections with the neurons in the MEC. The synaptic connections between neurons in MEC are modified by Hebbian learning to form connectivity pattern that exhibits attractor dynamics which is not necessarily topographically organized. If the neurons in the MEC were to be rearranged according to their connectivity, a topographically connected pattern similar to the one proposed by Fuhs and Touretzky could be seen Fig 3.7. Because the tutor has the periodicity of a grid, the rearranged MEC network has no borders and resembles the surface of a torus Fig 3.7.

To displace representations along the abstract space of the continuous attractor, the model introduces an additional layer of cells whose firing is modulated by place cells, head direction cells, and neurons that code for speed;, neurons with such properties have been shown to exist. Neurons in this hidden layer receive input from currently firing grid cells and project back selectively to grid cells that fire next along the trajectory. The activation of target cells depends on the current head direction and velocity of the animal. They model the variability of grid spatial frequencies seen along the MEC dorsoventral axis to variation in the projection strength of the speed signal into the layer, the stronger the projection, the higher the spatial frequency of the grid.

3.2.4 Interference Model, [73, 74]:

The principles governing the second class of models are that path integration occurs at the single cell level and is related to the theta phase precession. In 1993, O'Keefe and Reece modeled phase precession as the sum of two oscillatory signals with frequencies around the theta rhythm that are slightly different by an amount

proportional to the running speed. The resulting interference pattern can be decomposed into an oscillation at the mean of the two frequencies, which advances with respect to the theta rhythm and a slow periodical modulation with a phase that integrates the rat's speed and thus reflects its position along the track.

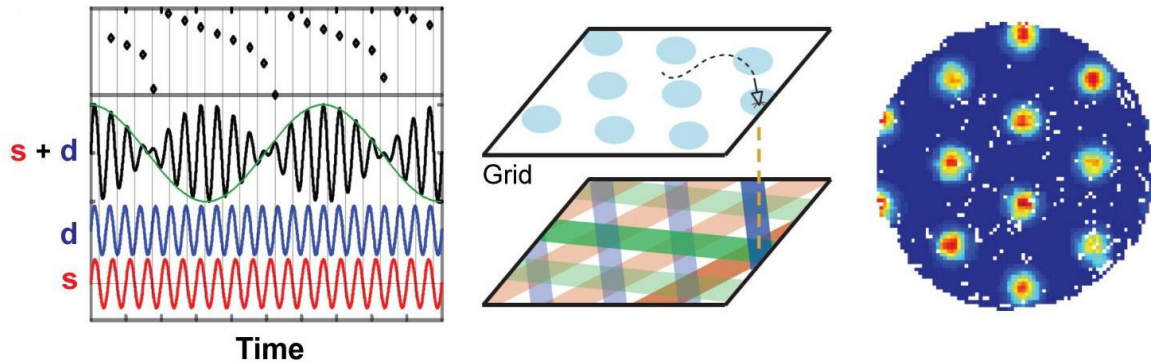


Fig 3.8. Sum of three or more linear interference maps; while the animal is running on a linear track, the sum of a somatic (s) and a dendritic (d) oscillation with slightly different frequencies results in an interference pattern exhibiting phase precession and slow periodic spatial modulation (*left*). Combining three or more linear interference maps, responding to different projections of the velocity, results in a grid map (*center*). Simulated grid map after 10 min of a rat's actual trajectory (*right*) from [73].

Burgess and his colleagues extended the interference model to be two dimensional. At each cell, there is an interaction of a somatic intrinsic oscillator of frequency w_s (\sim theta rhythm) with several dendritic oscillators, each with a frequency equal to w_s plus a term proportional to the projection of the rat velocity in some preferred direction as shown in Fig 3.8. The interference of the somatic signal with each of these dendritic oscillators has a slow modulation that integrates the preferred component of the velocity into a linear spatial interference pattern. As several of these linear patterns combines, a triangular grid map is obtained, provided that their directions differ in multiples of 60° and the phases are set in such a way that all

maxima coincide as shown in Fig 3.8. This choice of parameters could result from a self-organization process maximizing the neuron's overall activity. In this model, the modulation in spacing of grids along the MEC is associated with a gradient in the frequency of subthreshold membrane potential oscillations along this axis.

The main challenge of this model is that it depends on many parameters being chosen accurately, a process that the authors claim can emerge in biology from self-organization. Another challenge for this model is its operation relies entirely on having subthreshold oscillations at each grid cells with precise and stable phases and frequencies since any drift in them will result in the grid cells not operating properly.

3.3 Modeling the Place Cells

The effort to postulate a model for place cell formation and functionality began in 1971, and many models have been proposed trying to explain and replicate known neurophysiological data. Recent discoveries have altered the way the models are proposed, whereas the fact that most place cells show no direction selectivity allowed the models to be independent of head direction cell activity, it is not feasible to ignore the grid cells due the direct projections from the entorhinal cortex to the hippocampus. So the models for place cell formation can be divided into models proposed before the discovery of grid cells (a.k.a., “pre-grid”) and models proposed after the discovery of grid cells (a.k.a., “post-grid”).

3.3.1 “Pre-Grid” Place Cell Models

3.3.1.1 (Sharp, 1991)[14]:

In this model place fields are generated based on local-view pattern classification by competitive learning. The system consists of a three-layer neural network in which all units of one layer project to all units of the next layer through Hebb-like synapses as shown in Fig 3.9.

Cells in the first layer act as metric sensory cells responding to specific stimuli (e.g., distance to a cue). The input layer involves two types of units: Type 1 cells that fire as a function of the distance to specific external cues, and Type 2 cells that encode the distance of a cue as well as its angle relative to the agent’s heading. The activity of type 2 cells therefore depends on the agent’s current location and orientation. Both types of sensory units respond to a specific visual cue by a stochastic assignment done the first time the agent enters an environment.

Each cell in the middle layer receives afferent projections from all units in the sensory cell layer. Connection weights are initialized randomly such that they are all positive and their sum is normalized. There are 60 cells in the middle layer, divided into three winner-take-all clusters. Only one cell per cluster, the one receiving the largest input, can fire at any time. Synaptic weights to each winner cell are strengthened by Hebbian learning. The middle layer in this model corresponds to the entorhinal cortex.

The pattern of activity of cells in the middle layer (i.e., 3 active cells at any time t) is propagated to the output level of the model: the hippocampus. In this level, there is only one cluster of 20 cells whose firing activity is determined according to

the same winner-take-all scheme defined above. Also, synaptic enhancement occurs as before.

Sharp (1991) simulated a circular environment with 8 evenly distributed landmarks on the edge of the arena. At each step a simulated rat computes distances and angles of all 8 cues relative to its current position and heading, respectively. This information is used to drive the cells in the first layer of the model, and is propagated through the network to generate hippocampal place cell firing. Reported receptive fields are similar to real hippocampal place fields. A limitation of the model is that it does not capture data from experiments in the absence of visual landmarks (e.g., when light is extinguished) suggesting that rats are able to maintain place fields even without visual information.

3.3.1.2 (Burgess et al., 1994) [55]:

In this model, the allothetic sensory information activates a neural layer of entorhinal cells (EC), and then propagates through the network to form place fields in CA1-CA3 and in the subiculum (SC). Fig 3.10 shows a diagram of the proposed connectivity in this model.

At the sensory level, there is a population of cells, each cell responds maximally to a sensory cue at a different distance. A discrete set of cues is arranged around the edge of a square arena.

An intermediate layer of entorhinal cells (EC) receive hardwired input from the sensory cells to produce a large firing field, as shown in Fig 3.10. The EC cells project to the place cells (CA1-CA3) layer in the model through binary connections that are formed through Hebbian learning. Cells in the CA1-CA3 layer are arranged

in 5 clusters of 50 units each, and competitive learning is applied within each cluster: At each time step, only the four cells with the largest input fire a number of spikes proportional to their input, the others remain silent. As a consequence of competition, CA1-CA3 cells have smaller place fields than EC cells.

CA1-CA3 units project to SC cells through connections that are formed according to the same Hebbian learning scheme as before. Competitive learning also occurs within the SC layer. However, SC cells are arranged in 10 groups of 25 each, that is, each SC cell has to compete with fewer cells than each CA1-CA3 cell. This results in larger SC place fields. A simple set of goal cells was used to associate directions to goals with the current rat's location as given by the subicular cells.

A limitation of the model is that external landmarks are assumed to be perfectly distinguishable. Also, the approach does not take into account idiothetic information (i.e., path integration) in order to enable the simulated agent to exhibit stable place fields in the absence of external cues.

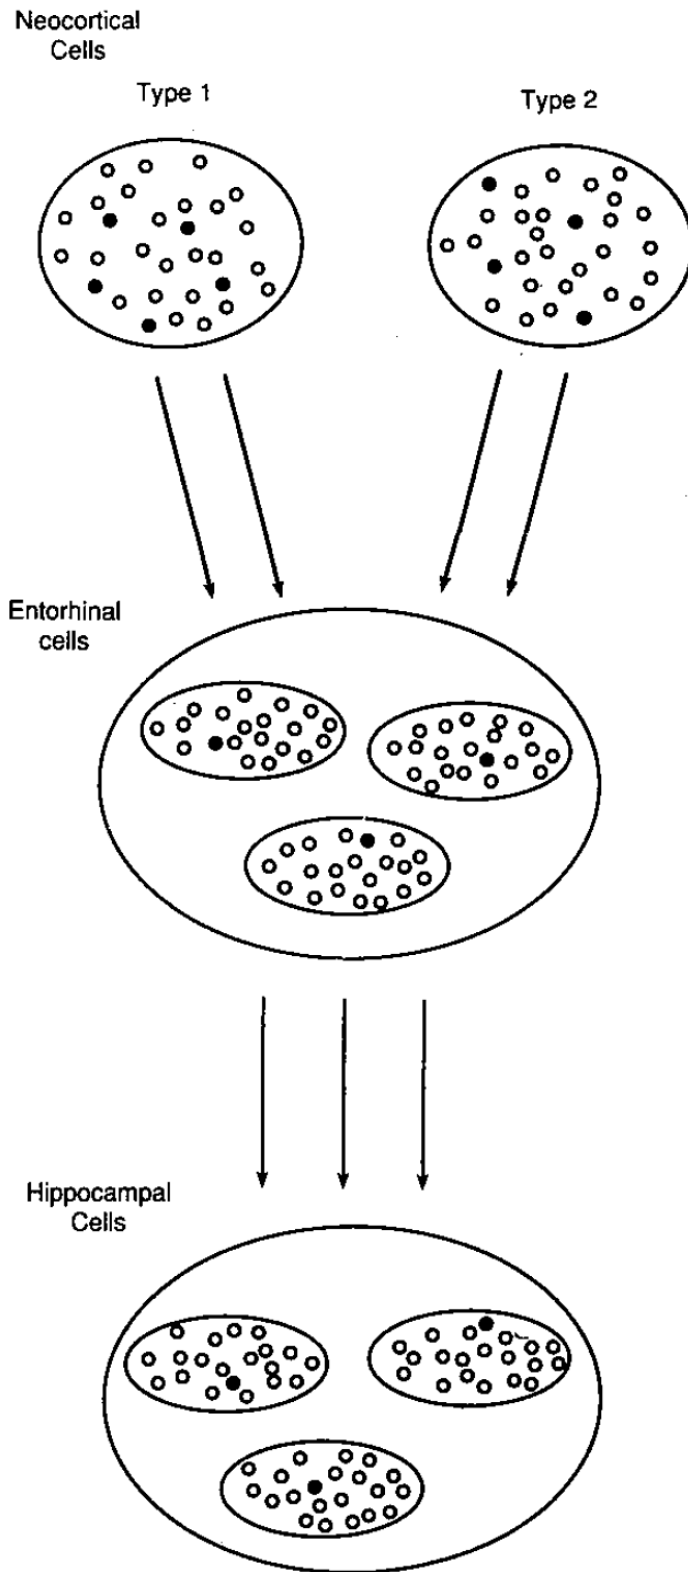


Fig 3.9: Sharp model [14].




Fig 3.10. Burgess, Recce, and O'Keefe model, adopted from [55].

3.3.2 Post-Grid Cells Models

3.3.2.1 (McNaughton et al., 2006) [72]:

By 2006, following the discovery of grid cells, McNaughton and his colleagues adopted a Fourier like analysis to explain the formation of non-periodic place fields in the hippocampus using periodic grids from the entorhinal cortex as shown in Fig 3.11. They assume that the hippocampal activity reflects the summation of the outputs of many grids with different spatial frequencies which leads to a periodic representation in the hippocampus but the cycle for repetition would be very larger than the scale of the largest grid. This assumption enables each position to be expressed by a unique pattern to collective activity. The activation of the hippocampal place field is computed by applying a simple thresholding operation to the summed grid fields as shown in Fig 3.11.

In this model, the authors assumed that the grid cell showed equal response at all active locations in the grid; this assumption is not supported by neurobiological data which shows variability in the response of the a given unit at different grid points.

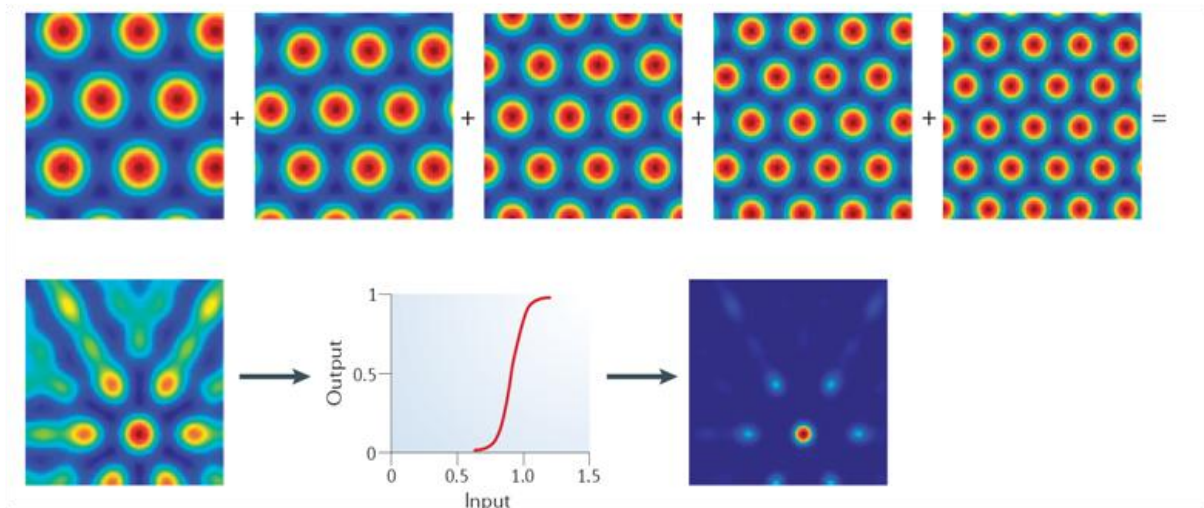


Fig 3.11. Multiple grid fields with different scales contribute to form a single peaked place field adopted from [72].

3.3.2.2 (Rolls et al., 2006) [75]:

In this model, the authors argue that the connections from the EC to the hippocampus need not be hardwired and that starting from random initial connectivity between the two layers and by applying simple learning techniques, place field like responses are achievable in the hippocampal neurons. This model has a group of EC (125) cells exhibiting grid like activities with multiple spatial frequencies. In this model they took into account the fact that each EC unit does not exhibit the same activation in all the vertices of the grid. They present two different algorithms for the learning, both apply competitive hebbian learning, however, in one algorithm the change in weight is a function of the current activity levels of pre and post synaptic neurons, and in the other algorithm, the change in weights is a function of the current and previous activity levels which allowed the formation of wider place fields in the hippocampal layer. Fig 3.12 shows the activity of two EC cells and the activity in a

hippocampal dentate gyrus (DG) neuron before and after training, showing that the hippocampal cell had single peaked place field in the arena after training.

The main limitations of this model are the lack of significant neurobiological evidences that competitive learning is used to establish place field formation in hippocampal neurons.

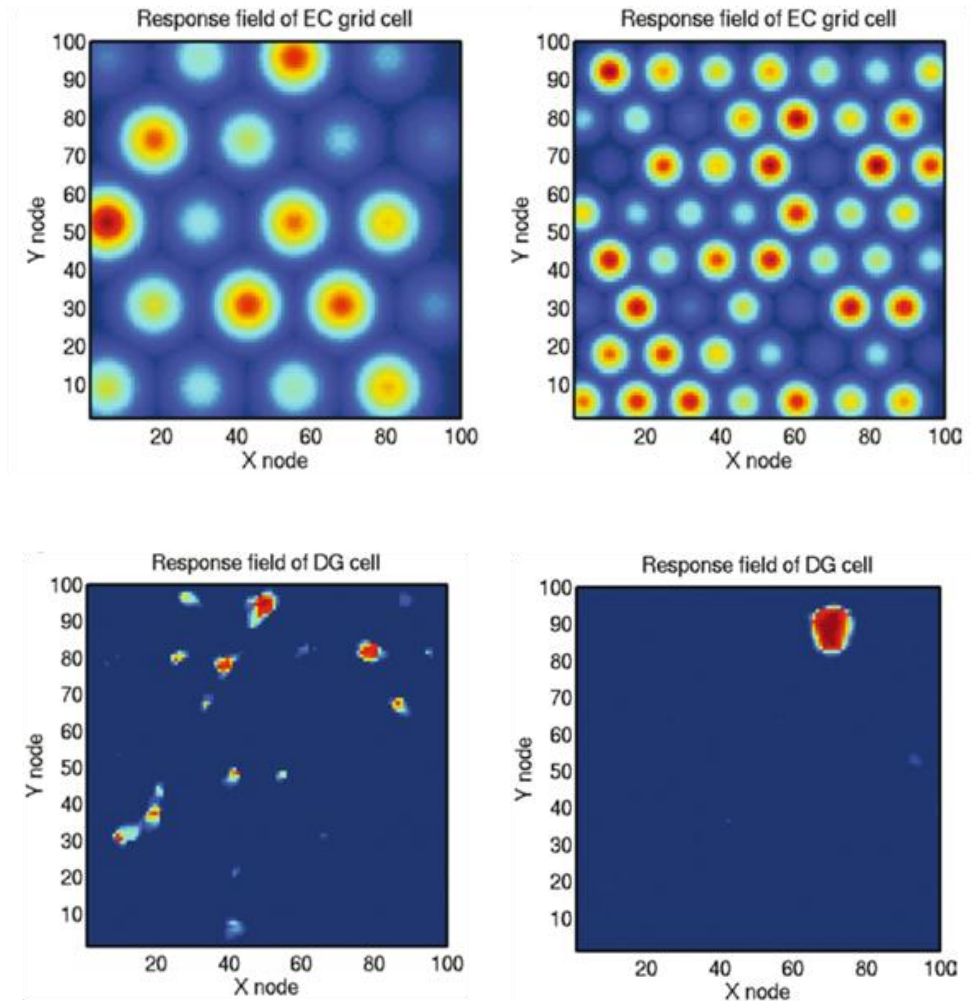


Fig 3.12. Simulated place cell formation in the hippocampal DG cells for the model proposed by Rolls et al., adopted from [75].

3.4 From Models to Navigation Systems

In the previous section, various proposed models for different biological structures known to be involved in spatial navigation were reviewed. In this section, we will review relevant previous efforts to implement biologically-inspired navigation systems on mobile robot systems, as our main objective is to build a neuromorphic spatial navigation system inspired from the bat and develop autonomous navigation capabilities for mobile robotic platforms.

3.4.1 (Arleo, 2000; Arleo and Rondi-Reig, 2007) [4, 76]:

Angelo Arleo implemented a system with place code cells, head direction cells and angular velocity cells. Activity in the head direction cells was shifted by idiothetic input (path integration) and recalibrated by visual input (allothetic calibration). Likewise, activity in the place code cells was shifted around by idiothetic input modulated by the activity in the head direction cells and corrected by visual input. The experimental setup consisted of a Kheperra robot equipped with a camera and wheel encoders in a small 800 by 800 mm arena with bar-coded walls. The room contained one external light source that was effectively used as a North Pole reference (salient cue). The robot starts exploring the environment and forms place cells activity relying only on idiothetic signals, Hebbian learning is used to link the created place cells to the allothetic cues encountered in the space (different views). As the robot spends more time exploring, it develops a growing sense of uncertainty due to drift in the self-motion integration. As the uncertainty passes a certain threshold, a homing vector is created and the robot heads back to the starting point while switching off the learning. As the robot heads back home it uses the learned spatial

memories to calibrate for the integration errors. The computation and processing were performed off-board as the robot is moving in the arena.

The robot was able to stay localized under these conditions although no results were presented for global kidnapping tests or for dealing with highly ambiguous visual input. There was no evaluation of the system's performance without the North Pole reference.

3.4.2 (Milford, 2008) [10]:

Milford worked on building a system relying on path integration and visual signals that he called RatSLAM. Rather than building a system with head direction and place cells he built a unified representation for Cartesian and angular position (x , y , θ) called a pose cell system. Fig 3.13 shows the RatSLAM model, the agent's pose is represented by the activity in a competitive attractor network called the pose cells. Wheel encoder information is used to perform path integration by injecting activity into the pose cells to shift the activity packets. Vision information is converted into a Local View (LV) representation which if familiar, injects activity into the particular pose cells that are associated with that specific local view.

The pose cells are implemented as a competitive attractor network, the network is organized such that it allows only one group of units to be active at a given time coding for the current estimate for (x , y , θ). The path integration process provides pose updates in the absence of visual input using information from the wheel encoders of the robot. The local view module is a collection of neural units that represent what the robot sensors 'see' to the rest of the RatSLAM system. Hebbian

learning is used to associate visual data to pose cells. Data were collected by the robot but the processing was performed using an off-board computer.

This system was able to operate properly in the absence of salient cues in the environments and it was successfully tested for exploring novel and large environments and for generating goal-directed behavior.

Although this system was able to operate successfully it is not suitable for our intended application because of the off-board processing and lack of neurophysiological detail. Moreover, using on-board dedicated hardware to implement this algorithm is not feasible since the pose cell system is an attractor network with infinite number of attractors corresponding to the potential poses for any given environment. This system solved for the ambiguity of position by recruiting a new pose for each new location and then using higher order processing to link different pose cells coding for the same physical location.




Fig 3.13. The architecture of the RatSLAM algorithm, adopted from [10].

3.4.3 (Erdem and Hasselmo 2012)[77]:

Erdem and Hasselmo presented a model for goal-directed navigation using forward linear look-ahead probes of potential trajectories through the environment using a circuit of head direction cells, grid cells, place cells and prefrontal cortex (PFC) cells. Fig 3.14 shows a schematic of the network used in this model to create place cells.

The circuit creates a place cell map using Hebbian modification of connections between prefrontal cortex (PFC) cells. The agent creates a map composed of place cells and PFC cells by random exploration. After exploration, the rat recalls the goal location, picks its next movement direction by forward linear look-ahead probe of trajectories in several directions while stationary to find the one activating PFC cells with the highest reward signal. Each probe direction activates of a static pattern of head direction cells to drive the grid cells to update their phases in a specific direction which in turn drive place cells along the probed look-ahead trajectory. This probing is repeated until the look-ahead trajectory activates the reward signal and the corresponding direction is used to guide goal-finding behavior. This model represents a plausible method for goal directed navigation using the network of head direction cells, grid cells and place cells that can shed some light on how the brain actually uses these circuits for navigation, however, the network is computationally expensive since at every position the agent has to evaluate all the possible trajectories to maximize the reward signal. This model addresses the problem of place cell creation by recruiting new cells when the agent is in a place that is not

represented, from the neurobiology we don't know whether this is a plausible assumption or not.

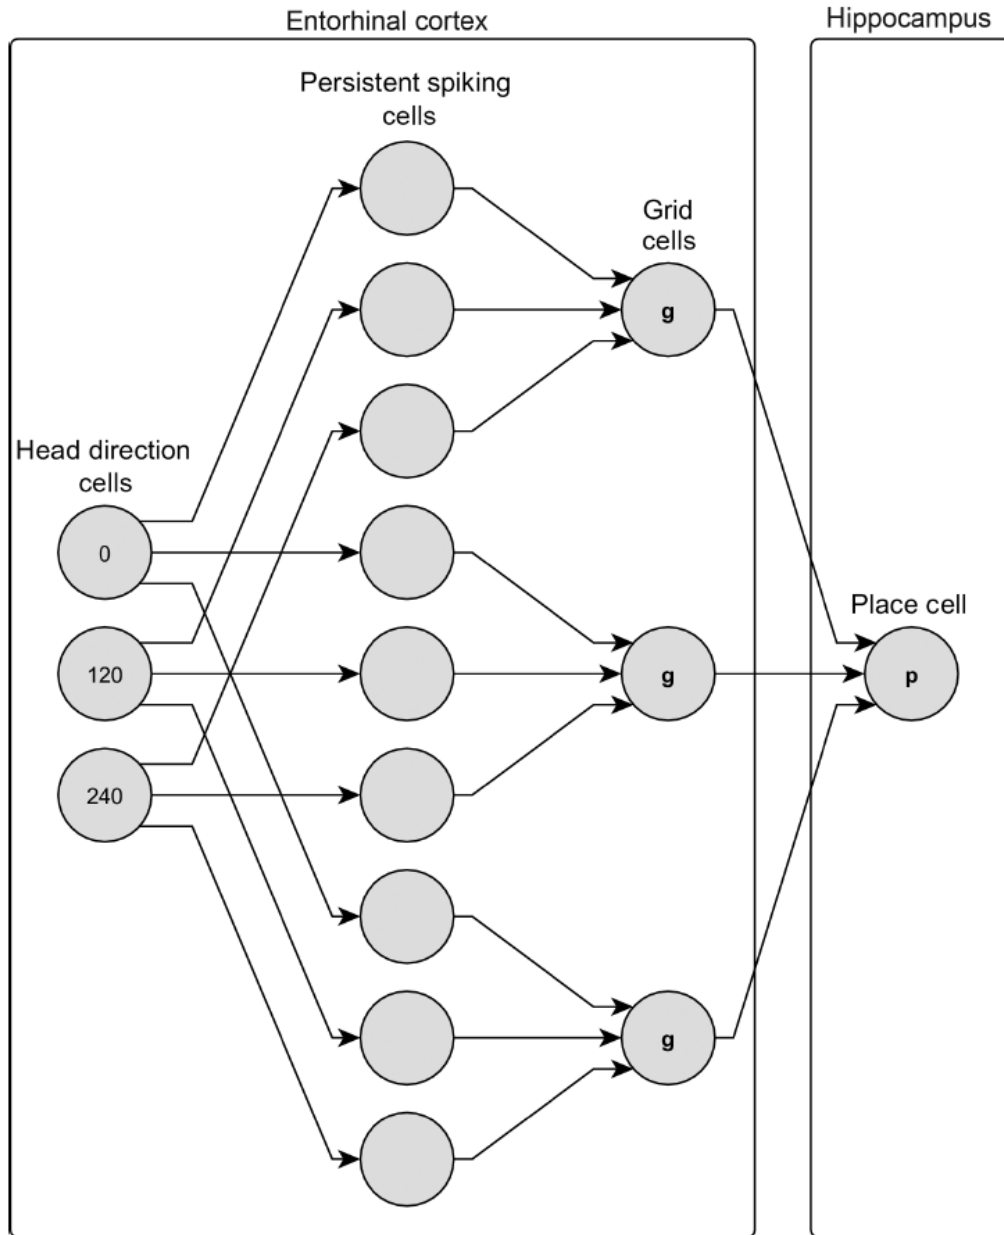


Fig 3.14. A schematic for the network used in the (Erdem and Hasselmo) model for goal directed navigation; adapted from [77].

Chapter 4 : Head Direction Cell System

4.1 Introduction

As we discussed in chapter 2, the hippocampus, the entorhinal cortex, and the postsubiculum are areas of the mammalian brain that have been shown to be heavily involved in spatial navigation related tasks [11, 12, 19, 47]. In rats, where spatial navigation and memory have been studied using two-dimensional environments, neural recordings from the postsubiculum have shown the presence head direction (HD) cells, that are tuned to the orientation of the animal's head in space; a given neuron fires whenever the animal's head is facing its preferred orientation and is silent otherwise [12, 14]. It is widely believed that the activity of these cells is updated using self-motion signals (e.g. head rotation velocity as detected by vestibular sensors), spatial memory (e.g. associating landmarks with particular directions), or both [14, 41, 44, 78]. Beyond simply providing a readout of the current orientation, these neurons collectively maintain an estimate of the head orientation by sustaining its activity and by shifting this activity from one part of the neural population to another to represent changes in orientation. Accurate integration of velocity signals to compute position or orientation, however, requires carefully matched neural and synaptic properties. Although biological neurons and their synaptic interconnections are assumed to be adaptive within any given system to compensate for mismatched characteristics, accurate calibration across all positions and rotation velocities is likely to be an elusive goal. Due to the nonlinearities inherent in all components of the system, without calibration, angular velocity

integration is likely to be nonlinear, location-dependent, and prone to drifting towards attractor locations. On the other hand, experiments with rats show that when stable spatial landmarks are present, the response of the HD neurons does not drift, which led to the conclusion that simple associative learning is used to correct for drift in the activity of the network [13, 41].

Since its discovery in 1984, many theories have been postulated to explain how the HD system works and to provide a biologically-plausible neural circuit for its operation. Nearly all of the proposed models have suggested that the network of HD cells can be described to be a recurrent neural network which exhibits a continuum of stable attractor states. Each attractor is represented by the persistent activity of a subset of the neurons (i.e., a bump of activity) acting as a memory for a certain head orientation. In the absence of any head movement, the network remains in its latest state, maintaining the current estimate of head orientation. When the head rotates, angular velocity information is used by the network to move the bump of activity through the network to represent the new head orientation [43, 67, 68, 79, 80].

Whereas most models for HD cells capture the general behavior of the system by adopting mean rate representations for the activity of neuron populations, some models use spiking neurons to model the HD system [81-83]. The operation of these models, however, relies on large neural populations with matched parameters to implement the HD system which makes them more suitable for software implementation rather than for analog hardware.

In this chapter, we present how a recurrent network of spiking neurons can successfully achieve the same functional behavior of HD cells without using large

neural populations. Our model incorporates modifications to conventional HD models by using a disinhibition-based gating mechanism to allow precise control of the movement of the neural activity and thus angular velocity integration.

4.2 HD System Model

In this work, we are interested in developing a biologically-plausible model for the rat HD system that can be directly mapped onto neuromorphic analog hardware. To model biological HD cells using an artificial system, the system must exhibit some key features of the biological system: 1) the system should be capable of maintaining stable activity in the network (i.e., act as a memory) in the absence of external stimuli, and 2) head motion-related signals should be used to move the activity in the network smoothly from the current location to a new location representing the new orientation of the head in space [68]. The model we present here for maintaining stable bumps of activity is inspired from previously presented HD models in the literature (section I) and this model can be directly implemented in hardware with minimal difficulty. To move the activity bump across the field of HD neurons in response to the head's angular velocity, however, we propose a novel biologically-plausible model.




Fig 4.1. The synaptic projection pattern as a function of the relative interneuron distance. The actual strength is determined by a product of this kernel and a global current setting. A positive value indicates an excitatory connection, a negative value indicates an inhibitory connection and “0” indicates no connection.

4.2.1 Stable Activity in the Network

To create stable, sustained spiking activity in the absence of external stimulation, the HD system is modeled as a ring of recurrently interconnected neurons. Each neuron in the ring has a neighborhood of neurons with which it has fixed mutual connections whereas no direct connections exist with neurons outside of this neighborhood. The weights of the connections between the neuron and its neighborhood has a “Mexican-hat” shape, with two groups of projections: excitatory projections to nearest neighbors with decaying weights as the separation between the

neurons increases, and a group of inhibitory projections that are locally strong then decay with increasing interneuron separation. Fig 4.1 shows the projection pattern used in our experiments. A ring of neurons with such connectivity and a weak, constant input current driving them all to fire will exhibit activity in the shape of one or more groups of active neurons forming bumps of activity [84-86].

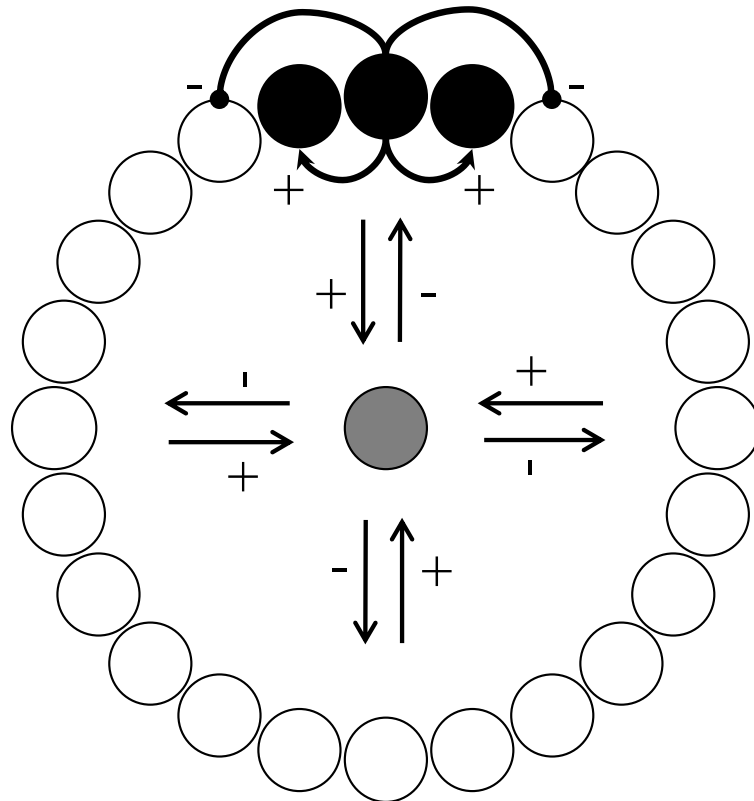


Fig 4.2. A ring of neurons with the interconnections shown for only one neuron for clarity. The neuron in the center of the ring is the global inhibitory neuron. The connectivity allows one group of neurons (“bump”) to be active at a time (shaded in black). The global inhibitory neuron (shaded gray) normalizes the total network activity through inhibitory feedback. Additional neurons and connectivity for moving the bump activity are not shown.

For the HD system operation, a unique representation is needed for each direction in space and thus a single bump of active neurons is the desired pattern to avoid ambiguity. To achieve this requirement, a soft winner-take-all (WTA) function is added to the network, in the form of a global inhibitory neuron which receives excitation from each neuron in the ring and projects inhibition back to all neurons in

the ring. In Fig 4.2 we show an example set of connections for one neuron in the ring with the minimal neighborhood connection width (for clarity in the illustration) of only one excitatory and one inhibitory neuron on each side. This connection pattern is repeated for each neuron. The excitatory and inhibitory neighborhoods for each neuron are programmable (via external spike routing processors) and extend to more than one connection in each pool (e.g., six excitatory and eight inhibitory in Fig 4.1). All neurons are biased to have an excitatory DC current which weakly drives them to become active and produce spikes in the absence of other inputs. The recurrent lateral excitation creates a neighborhood of support (of a size defined by the excitatory connectivity) for bumps of spiking activity that can be sustained without external input. As neurons in the ring produce spiking activity, the global inhibitory neuron is excited and projects inhibition back to each neuron in the ring. This global inhibition limits the activity in the ring to those neurons with the highest spiking rates (induced by recurrent input connections or external inputs). This soft “winner-take-all” function generally allows only one bump of activity to exist on the ring at a given time. If the lateral connections are balanced, the bump of activity will remain at its current location on the ring.

4.2.2 Moving the “Bump”




Fig 4.3. Connectivity diagram of the bump movement neurons. The arrow head indicates an excitatory connection and the circle-shaped head indicates an inhibitory connection. The network connections that maintain bump activity (see Fig. 4.1) are not repeated here for clarity.

In our network, each group of active neurons (attractor state) in the ring is associated with a certain head direction in the environment. As the head is turned (due to head movements or combined head-body movements), the system will receive input signals coding for the angular velocity of the head. This head velocity information is provided to the system as left and right speed signals which drive the bump to move around the ring to represent the new head direction. This corresponds functionally to a mathematical integration of the head’s velocity information to compute the head angle. The neurons and connectivity used to perform this function are shown in Fig 4.3.

The “bump neurons” in Fig 4.3 are the same neurons shown in Fig 4.2 with the bump-layer connections hidden. At each location in the ring a “left-rotation” and a “right-rotation” neuron receive global head rotation signals. For leftward (rightward) rotation, these signals produce spiking activity in just the left-rotation neuron (right rotation neuron) proportional to leftward (rightward) speed.




Fig 4.4. Schematic diagram showing the sequence of events for moving the bump of activity based on a left velocity signal (active connections are shown in black and the inactive ones in light gray). The top panel shows the stationary case with no velocity signal, the middle panel shows the activity in the network as a left rotation signal is applied and the bottom panel shows the activity in the network after it moves one neuron to the left in response to the velocity signal.

To move the bump at a particular speed, spiking activity with a rate proportional to the head turning rate is projected onto the bump neurons at the current location of the bump with a leftward or rightward shift of one neuron. The spiking rate of this projection controls the *rate* at which the bump location will shift by one neuron. To achieve this specific local pattern of projection using only global rotation signals, a disinhibition-based gating mechanism inspired by the operation of the superior colliculus in the saccadic eye movement is used [87]. In this circuit, *active* bump neurons disinhibit the local left-rotation and right-rotation neurons which then respond to the globally-supplied rotation speed inputs. All other rotation neurons in the ring are actively suppressed. As a result, the activity of the disinhibition neurons appears as the negative of the bump activity. Fig 4.4 shows a cartoon example of a movement illustrating the sequence of events for shifting the bump location one position to the left in response to a left-rotation signal. In this figure, we show the active connections in black and the inactive connections in light gray. Initially, (in the top panel) there is a stable bump of activity centered on a group of neurons in the ring of bump circuits. All of the neurons in the disinhibition circuit are active except the ones corresponding to the bump location. With no global velocity signal, however, the rotation neurons are not active. In the center panel, when the head is turning left, all left-rotation neurons receive input, but only those locations corresponding to the current bump location will respond. These active neurons project excitation to bump neurons shifted one position to the left; this projection causes the activation of the neuron to the left of the current bump which was not firing previously. This activation disrupts the balance between the projection kernel

and the global inhibitory projection pattern; this balance allows only a certain number of neurons (3 as shown in Fig 4.4) to be active at a time to form the bump. The global inhibitor neuron is now driven with higher input (4 active bump neurons) thus projects back more inhibition to the bump neurons trying to shut down one of its active neurons. Moreover, the newly activated neuron (now at the left most edge of the active array) interacts with the right most neuron in the current bump location as they mutually inhibit each other with direct synaptic projection. Thus the level of inhibition on the right most neuron and left most neuron is higher than the other neurons in the array. The system would go back to a balanced state only when one of the active bump neurons shuts down. The newly activated left most neuron receives elevated inhibition levels but also receives excitatory input from the left rotation chip whereas the right most neuron receives only an elevated level of inhibition which eventually will shut it down and hence moving the bump one location to the left.

Our model for moving the activity in the HD system is different than previously proposed models [43, 67, 68, 79-81] in that our synaptic projections are fixed and do not need to rapidly change to provide an asymmetrical projection and shift the bump of activity. Furthermore, many previous models move the bump by suppressing its “backside” but in our model, we move the bump by expanding the excitation in the direction of motion.

4.3 Circuits

To build this system, we designed and fabricated a neuron chip consisting of an array of 32 integrate-and-fire (I&F) neurons using a neuron circuit slightly modified from Indiveri et al. [88]. The output activity of these neurons is read using the address-event-representation protocol AER [89]. For its input, each neuron has 14 pulse-extender (PE) synapses (8 excitatory and 6 inhibitory) that use a principle of operation similar to that described by Arthur and Boahen [90]. Each one of these synapses has a unique digital address and an input AER system is used to activate them. We use a resistive ladder to generate the required bias voltages to control the weights of these synapses. Using three externally-supplied DC voltages (the top node, the bottom node, and an intermediate node), seven different DC biases are generated to control the weight of two matching-weight synapses on each neuron. To retain flexibility to investigate different lateral connection patterns between the neurons, no hardwired connections were implemented between the neurons on the chip; a digital processor was used to implement the connections by providing a look-up-table of interconnections and routing spikes back to the chip accordingly. This configuration allows a neuron to be connected with up to 14 neighbors, seven on each side with symmetrical connection weights.

A global inhibitory neuron, that receives excitation from all 32 neurons using equally-weighted synapses, projects inhibition back to all neurons with equal weights. The global inhibitory neuron connections are hardwired on the chip with separate controls for the excitatory weights stimulating the neuron and the inhibitory weights back to the bump neurons. The global inhibitory neuron can also be activated or

suppressed using a separate external control. The membrane voltage of the global inhibitory neuron was externally monitored.

Fig 4.5 shows a micrograph of the neuron chip; the chip is 1.5 mm x 1.5 mm and was fabricated using a commercially-available 0.5 μm CMOS (3-metal, 2-poly) process by the MOSIS chip fabrication service.




Fig 4.5. Micrograph of the Neuron Chip. The chip has 32 I&F neurons, each with 14 Pulse Extender (PE) synapses. The activity is read from the chip using an AER interface.

4.3.1 Synapse Circuit

Fig 4.6 shows the synapse circuit we used in our system. The circuit operates as follows; as an AER Ack pulse (on the order of tens of nanoseconds wide) is received, the capacitor voltage V_{C1} (initially at 0V) charges up to V_{dd} . After the pulse is removed, V_{C1} discharges back to ground through M2 and M3 at a rate controlled by the voltage V_{biasN1} . As V_{C1} rises, M6 turns ON, M5 turns OFF, and V_{C2} discharges to ground through M6 and M7 at a rate set by the voltage V_{biasN2} . When V_{C1} is

discharged close enough to ground, M6 turns OFF, M5 turns ON, and VC2 charges to V_{dd} through M4 and M5 at a rate controlled by V_{biasP2} . During the time that C2 is discharged to ground and then charged back up to V_{dd} , M9 conducts the saturation drain current of M8 (defined by V_w) at the output of the synapse. In an inhibitory synapse, the output current in M9 is mirrored to M11 and is drawn from the membrane capacitance of the postsynaptic neuron which results in the discharging of the membrane capacitance down towards ground. In an excitatory synapse, the current from M9 is pushed directly into the membrane of the neuron charging it up towards V_{dd} . If another AER pulse stimulates the synapse before VC2 fully charges back to V_{dd} , (i.e. while the synapse's output current is still ON) the duration of the output current pulse is extended.




Fig 4.6. Synapse circuit: V_{biasN1} , V_{biasN2} and V_{biasN3} control the charging and discharging of C1 and C2 and hence control the width of the pulse extender pulse. V_w is the weight of the synapse and controls the amplitude of the output current; for an excitatory synapse the output current is from the drain of M9, whereas for an inhibitory synapse, the output current is at the drain of M11.

With these controls, the synapse circuit is configured to produce an output current pulse on the order of a few milliseconds for each received AER spike. The duration of the output current pulse is controlled by the three biases (V_{biasN1} , V_{biasN2} , and V_{biasP2}).

4.3.2 Neuron Circuit

Fig 4.7 shows the neuron circuit, based closely on the neuron by Indiveri et al. [88]. In this schematic, C_{mem} is the membrane capacitance of the neuron, I_{Exc} and I_{Inh} represent the excitatory and inhibitory synaptic currents. The neuron's DC bias current is provided by transistor M1 whose gate is controlled by the voltage V_{inj} . V_{inj} is a global parameter for all the neurons in the chip except for the global inhibitory neuron. Transistor M2 provides the leak current that continuously discharges the membrane capacitance, C_{mem} . This leak current is controlled by V_{leak} , also a global parameter for all neurons on the chip. M6 and M7 are a source follower that raises and adjusts the neuron's threshold voltage for firing a spike. The circuit operates as follows; C_{mem} integrates the input currents of the neuron, resulting in a voltage V_{mem} . When V_{mem} rises above the threshold voltage of the neuron controlled by V_{sf} and the inverter (M10 and M11), V_2 is pulled towards ground which activates the positive feedback current path through M3, M4, and M9 and accelerates the rise of V_{mem} and hence the spiking of the neuron. As V_2 drops to ground, the inverter (defined by M17 and M18) switches state to output high which switches M20 (ON) to pull down on the AER request line ($Req/\$). When the AER system acknowledges the request, it sends back an active-low signal ($Ack/\$) to activate the inverter (defined by M14 and M15) through M12 and M13. Pulling V_3 to a high state resets V_{mem} to

ground using M5. As V_{mem} drops to zero, V_1 , V_2 , and V_3 switch states and the neuron is in the refractory period controlled by the discharging of C_{refr} through M16 with a rate controlled by V_{rfr} .

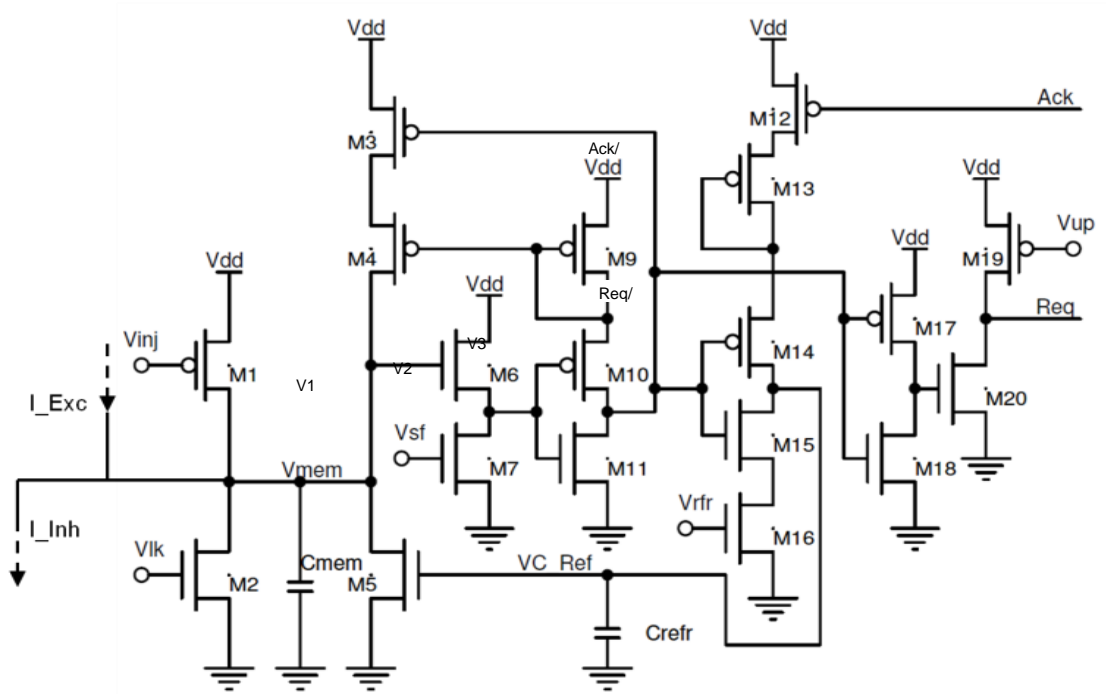


Fig 4.7. Neuron circuit slightly modified from the one presented by Indiveri et al [88].

4.4 System Implementation

To construct the HD system, four neuron chips were used to perform the spike-based computations and a set of dsPIC® microcontrollers (Microchip Inc.) were used to implement the required synaptic connectivity between neurons on the same chip and neurons on different chips. The dsPIC microcontroller was also used to record the activity in the system by sending spike information to a PC using a standard serial

connection. Fig 4.8 shows the block diagram of the neuromorphic HD system we implemented.




Fig 4.8. Block Diagram of the HD system. The dsPICs are the boxes labeled PIC and the other boxes refer to multiple copies of our neuron chip.

4.4.1 Bump Formation

Fig 4.9 shows the setup we used to implement the lateral bump-layer ring interconnectivity. A dsPIC was configured to read output spikes from the bump neuron chip and route the spikes recurrently to multiple synapses at its input. The synaptic weights were set to implement a “difference-of-exponentials”-shaped connectivity pattern for each neuron with its nearest neighbors receiving the strongest

excitatory connection and distant neighbors receiving inhibitory connections that decay in strength with distance.




Fig 4.9. Testing setup showing the Bump chip, the routing microcontroller, and the serial connection to the PC.

In addition to implementing the connectivity pattern between neurons, the microcontroller can “weave” externally-provided spikes into the stream of recurrent activity going into the chip. This feature is used to implement the connectivity between the rotation chips and the bump chip.

4.4.2 Angular Velocity Integration

The mathematical integration of angular velocity (i.e., moving the bump of activity) is implemented using two parts, the disinhibition circuit and the rotation circuits.

4.4.2.1 Disinhibition

The disinhibition neuron layer is a copy of the neuron chip which receives “one-to-one” inhibitory connections from the bump chip. On this chip, the global

inhibitory neuron is suppressed and no lateral connectivity is implemented. A dsPIC microcontroller reads the activity of the bump chip and drives the inhibitory synapses at matching locations on the disinhibition chip. All disinhibition neurons are biased to have a tonic spiking response *except* when the bump neurons suppress this activity. This connection results in a spiking pattern on the chip that appears as the negative of the bump neurons.

4.4.2.2 Left and Right Rotation

The left and right rotation chips are copies of the neuron chip that receive a global excitatory input current proportional to the rotation rate; the left-rotation chip is driven only for leftward rotations and the right-rotation chip is driven only for rightward rotations. The global inhibitory neuron is suppressed on each chip and no lateral connections are implemented. A dsPIC microcontroller reads the activity from the disinhibition neuron chip and drives the inhibitory synapse for the matching location on both chips. In leftward rotation, for example, the neurons on the left-rotation chip at the current bump location fire with a rate proportional to the leftward speed input. Other neurons on the chip are suppressed by activity from the disinhibition neuron chip. A dsPIC microcontroller reads the activity from both rotation chips and routes these spikes into the bump neuron chip with a one-neuron leftward shift for spikes originating from the left-rotation chip and a one-neuron rightward shift for spikes originating from the right-rotation chip.

Fig 4.10 shows the connectivity of the entire HD system used to collect the results presented in the following section.




Fig 4.10. HD System Connectivity: The neuron chips are labeled in the figure and the other boards are the dsPIC microcontroller implementing the AER routing in the system.

4.5 Testing Results

The neuron chip was fabricated using a commercially-available $0.5\mu\text{m}$ 2-poly, 3-metal process using the MOSIS Service. The four neuron chips of the system together consume 1.5mW with a 5V power supply during normal operation.

4.5.1 Creating a Stable Bump of Activity

One of the main functions of the HD model is the formation and maintenance of a bump of neural activity that does not decay, spread uncontrollably, or drifts laterally around the ring. This bump of sustained activity acts as a working memory for the current estimate of spatial orientation in the HD system and moves around the ring of neurons as the orientation of the head in space changes. Most models of the

HD system have assumed that the bump of active neurons have a Gaussian-shaped spatial activity distribution. In this work, we demonstrate the operation of our system in this mean rate regime as well as in a novel synchronous mode regime and report our results in this section.

4.5.2 Mean Rate Mode of Operation

In the mean-rate mode of operation, the aim is to operate the network such that the activity pattern has a Gaussian-like spatial distribution of firing rates around the ring of neurons. To achieve analog-valued firing rates and avoid synchronization effects, all synapses are set to provide low-amplitude, long-duration currents to prevent the incoming spikes from imposing any significant temporal structure on the post-synaptic neuron firing times. All neurons in the ring also receive an identical tonic input current that drives them (in the absence of other inputs) to fire at a very low rate. In Fig 4.11 we show raster plots of the activity recorded from the ring of neurons when the global inhibitory neuron was: 1) prevented from firing, and 2) permitted to fire. In the absence of the global inhibitory neuron feedback, we see the formation of multiple activity bumps due mainly to the lateral interconnections between the neurons. In this case, the presence of multiple bumps on the ring helps to stabilize the entire pattern. Due to the projection kernel, each bump actively sustains itself and prevents the formation of another bump immediately next to it on either side through long range lateral inhibition. This effectively “repels” other bumps. Due to the ring structure of the array, a steady-state pattern of *bumps* is quickly achieved that can counteract the drifting of individual bumps, resulting in an overall activity pattern that is stable.




Fig 4.11. Left panel, neural response spike rasters of the ring neurons to a constant input current with the global inhibitory neuron suppressed. Multiple bumps of activity are present. Right panel; response of the ring neurons to a constant input current with the global inhibitory neuron providing feedback. The bump is initiated at an unstable location but eventually drifts to a more stable location. With the global inhibition activated, the activity bump occurs with lower firing rates than the case shown in the left panel.

When the global inhibitory neuron provides feedback, it creates a competition between the various bumps and the bump with the strongest drive (due to mismatch in this case) will remain active while the other ones will be suppressed. The operation of the global inhibitor reduces the firing rates of the neurons in the active bump. The width of the bump is controlled by the width of the excitatory portion of the projection kernel for each neuron. In the example shown in Fig 4.11, a kernel that produces a four neuron-wide bump is used. More examples of width control are shown in the next section. In the case of one active bump, the stabilizing repulsion “force” from the other bumps around the ring is absent, hence, the bump is prone to

mismatch and can drift to stable locations on the ring. An example of such drifting is shown in the right panel of Fig 4.11.

4.5.2.1 Bump Stability in the Mean Rate Mode of Operation

Mismatch between circuit elements is one of the major problems that we face in analog VLSI circuits. Fig 4.12 shows the response of the neurons in the bump chip for the same tonic input, ideally driving them all to fire at the same rate. The right panel of Fig 4.12 shows the frequency response of the neurons normalized by the average firing frequency of the group (80 Hz). The figure shows that there is a wide range of deviation from the average which shows the presence of mismatch in the response of the neurons. This mismatch in the response of the neurons can disturb the operation of our system especially in the mean rate mode of operation. Renart et al. [91] showed that in a simulated network of heterogeneous spiking neurons, bumps of activity are prone to drifting to global attractor states due to the mismatch between the neurons and synapses. Although they demonstrated that stabilization of the drift was possible with synaptic plasticity or dramatic increases in the network size, these approaches were not practical for our system and thus we did not investigate them.

With only a small number of neurons in the ring layer (only 32 neurons) compared to the ring systems analyzed by Renart et al. (4096 neurons) [91], the effect of mismatch on the operation of our system can be dramatic. To study the extent of drifting in this system, we initiated a bump of activity (four neurons-wide) at each location in the ring and monitored the position of the activity bump. Fig 4.13 shows the evolution of the activity centroid (activity-weighted position computed using a 25 msec integration window) initiated at each of the 32 possible locations around the

ring. For illustration purposes, we show only the time intervals where significant drifts in the bump location occurred.




Fig 4.12. Left panel, rasters showing response of bump neurons all of them receiving the same tonic input. Right panel, deviation from mean firing freq by individual neurons.




Fig 4.13. Each trace represents the centroid of the activity of a bump started at each of the 32 possible locations around the ring. The plot shows the presence of some global attractors; for example, the bumps with initial center of mass near neurons 6, 7, 8, 9, and 10 eventually migrate to a centroid near neuron 9. The plot shows the two time intervals where the drift in the bump occurred; the first is the first 0.5 sec of operation; then another window after 2.5 seconds until 5 seconds, after 5 seconds we did not see any other drift in the bump location. A neuron # wrap-around was used for better visualization of the attractor basins. f

the bump is crucial for a reliable memory of orientation. Movement of the bump should only occur with actual head movement. As seen in Fig 4.13, when operating in the mean-rate mode of operation, we encountered instabilities at some locations. We attribute these instabilities to mismatched circuit properties and now describe the conditions under which this can occur.

In the mean-rate mode of operation, the shape of the local interconnection pattern strongly influences the pattern of firing rates within the bump. Because the neuron firing rates are a steady-state solution of the recurrent projection pattern (which

includes the spike generation nonlinearities) they cannot be simply studied analytically; however, we can gain some intuitive insights by examining typical solutions (e.g., Fig 4.11, right). By using a projection pattern like that shown in Fig 4.1 (with an excitatory neighborhood of 3 neurons and inhibition beyond that), neurons near the middle of a bump (that is four neurons wide) receive excitation from both sides and receive no inhibition. Neurons on the edge of the bump receive excitation only from one side and no inhibition. The neurons just *past* the edge (which are not active, or “off-edge” neurons) receive strong inhibition from the opposite edge neuron and excitation from three bump neurons. Because this off-edge neuron would inhibit the opposite edge neuron *if it were active*, these two neurons are in competition with each other and can exhibit the desired bistable behavior. To ensure stability of the bump location, the off-edge neurons need to stay inactive when the bump neurons are active. To achieve this, the net input current to the off-edge neurons on both sides must be zero or negative. Because the difference in input current is relatively small for these two competing neurons, mismatch can upset the required relationship for stability. If the off-edge neuron is receiving even a small net positive input current, the neuron will eventually produce a spike and inhibit the opposite edge neuron, shifting the centroid. The right panel of Fig 4.11 shows an example where the bump was initiated at neurons 28, 29, 30, and 31 but ultimately drifts to a more stable location (neurons 27, 28, 29, and 30). The instability at neurons (28, 29, 30, and 31) can be related to the slow firing rate of neuron 31 (that can be attributed to mismatch) that is not strong enough to inhibit all of the excitatory input to neuron 27, leaving this neuron with a net excitatory input that is integrated over

time to ultimately generate a spike and shift the bump. Since the global inhibition (which uses long current pulses) inhibits all neurons “equally”, it reduces firing rates for all neurons and thus provides the winner-take-all function, but it does not improve the margin for stability. Although increasing the strength of the local inhibitory projection would increase stability, this will come at the cost of making it difficult to move the bump intentionally.

Although we have demonstrated the ability to initiate and maintain bumps in a mean rate mode, this mode of operation suffered significantly from the problem of attractors and thus we did not test the system any further in this mode. The mean rate mode of operation may be suitable for large arrays of neurons where small drifts would not result in big errors of estimation or where learning could be applied efficiently to reduce drifting. In small systems like ours, however, any small drift results in significant errors in the estimation of spatial orientation.

4.5.3 Synchronized (Bursting) Mode of Operation

A “synchronized” or “bursting” mode of operation, natural to models that use spiking neurons, resulted in the best control over the bump movement and minimal drift towards attractors. In this configuration, the strengths of the lateral excitatory connections were tuned to be relatively short and strong, encouraging neighboring neurons to fire soon afterwards. We explore this mode and its operation as a memory structure for modeling the HD system in the following sections.




Fig 4.14. Left panel, neural response spike rasters of the ring neurons to a constant input current with the global inhibitory neuron suppressed. Multiple bumps of activity are present. Right panel, response of the ring neurons to a constant input current with the global inhibitory neuron providing feedback. A single, synchronously firing group of neurons emerges.

The first demonstration is the formation of a bump of activity. The neurons in the bump chip were all biased with the same DC bias to have a low tonic firing rate (this bias condition is maintained throughout the operation of the system). We monitored the activity of the bump neurons in two conditions: first with the global inhibition suppressed and then with the global inhibition allowed to fire. The results are shown in Fig 4.14. Without global inhibition, multiple bumps of activity appear on the network as discussed in the section (*Mean Rate Mode of Operation*) and their interaction helps to stabilize them all (as in Fig 4.11, left). When the global inhibitory neuron is allowed to fire, it imposes a global competition on the active neuron groups and allows only one group of neurons (i.e., bump) to remain active.

Fig 4.15 shows the time-evolution of the activity centroid (activity-weighted position computed using a 25 msec integration window) for bumps started at each of the 32 possible locations around the ring. Throughout the course of operation, we did not see major drifting in the location of the centroid and the bump remained stable at all 32 possible locations around the ring.




Fig 4.15. Each trace represents the centroid of the activity of a bump started at each of the 32 possible locations around the ring in the synchronized mode of operation. The plot shows that the bumps initiated at all locations are stable (i.e., no drift was detected during the course of operation). The same neuron # wrap-around as in Fig. 13 was used.

4.5.3.1 Bump Stability in the Synchronized Mode of Operation

As discussed in the section (*Bump Stability in the Mean Rate Mode of Operation*), the stability of the bump location around the ring is a key issue for our system. In the synchronous mode of operation (unlike the mean-rate mode), we did not encounter

drifting in the bump location with time. We attribute this stability to two factors: 1) the strong, pulsatile, excitation imposed by the lateral interconnections and 2) the strong, pulsatile inhibition imposed by the global inhibitory neuron. Below we discuss details as they relate to bump stability.

Consider the case where all neurons are initially at the same zero state, receiving a constant excitatory bias input. If no other inputs are provided and the neurons and interconnections are identical, all neurons would integrate up to the threshold at the same time and fire synchronously. If one group of neurons receives extra input current (e.g., from long-duration recurrent synaptic input or from external signals), the neuron with the largest input will spike first (“the winning neuron”) and it will provide a strong excitatory pulse of current to its nearest neighbors and a strong inhibitory pulse of current to its more-distant neighbors. The strong lateral excitatory projection dramatically accelerates the winning neuron’s excitatory neighborhood, producing a group of near-synchronous spikes while strongly inhibiting (i.e., slowing down or resetting) the off-edge neurons. When the synchronous group fires spikes, they excite the global inhibitory neuron to fire its spike very shortly after the active bump neurons. Because the inhibition arrives while the active neurons are still in their refractory period, this inhibition does not affect their firing rate. For non-firing neurons, however, this strong inhibitory pulse resets their membrane potential to zero.

A byproduct of the synchronous mode is that the neurons will have identically high firing rates. In the synchronous mode, the first neuron to fire excites its six nearest neighbors. The second neuron to fire (ideally one of the first neuron’s direct neighbors) further excites the neighborhood, but suppresses neurons four locations

away. Ultimately a group of four neurons will fire that do not receive any inhibition. By providing sufficiently strong local inhibition, we can ensure that the off-edge neurons never fire, to ensure stability. Because this inhibition is short, it is unlikely to interfere with other external inputs that direct the intentional movement of the bump. This feature is important because it creates a larger parameter space over which stability is created without negative consequences for movement. Furthermore, the four active neurons excite the global inhibitory neuron with a short latency such that the global inhibitory pulse can provide additional inhibition that also benefits stability without interfering with movements of the bump. Once the refractory period is over and the global inhibition has passed, all neurons begin integrating again towards threshold to start a new cycle.

In this scheme, the global inhibitory neuron plays two roles: it discourages multiple bumps around the ring and it resets all neurons to their zero state after the bump neurons have fired a spike, erasing any residual charge in the “non-bump” neurons due to mismatch and preventing these “non-bump” neurons from firing over time.

4.5.4 Controlling the Bump Width

In the previous experiment we showed the emergence of a four-neuron-wide bump as a result of the local connectivity between the neurons and global inhibitor. The width is controlled by adjusting the span of the lateral excitatory-inhibitory projection pattern between the neurons in the ring; wider projections produce a wider bump and narrow projections produce narrow bumps. In Fig 4.16 we show three different examples of widths obtained by changing the recurrent projection pattern width.

One interesting property of the synchronous case is that the firing rate of the group

is largely locked to the most active neuron in the group. Wider bump widths, therefore, result in fewer neurons determining the bump firing rates thus reducing the variability seen around the ring.




Fig 4.16. Activity bumps with variable widths of three, four, and five neurons wide. For the three neuron-wide bump, each neuron was connected to 10 others (four excitatory and six inhibitory). For the four neuron-wide bump each neuron was connected to 12 others (six excitatory and six inhibitory), and for the five neuron-wide bump each neuron was connected to 14 others (eight excitatory and six inhibitory).




Fig 4.17. System data, for rotating the bump to the left with variable speed.

4.5.5 System Data

In Fig 4.17 we show the activity from all of the neuron chips in the system for the case of head rotation to the left that starts slowly then speeds up with time. The first (top) panel shows the spike activity recorded from the bump neuron chip, the second panel shows the spike activity from the disinhibition neuron chip, the third panel shows the left-rotation neuron chip, and the fourth (bottom) panel shows the spike activity from the right neuron chip. The bump is initially stationary at neurons 8, 9, 10, and 11, while neurons on the disinhibition neuron chip show the inverse activity. When the global left-rotation signal is applied to the left-rotation chip, only neurons corresponding to the current bump location become active, resulting in the bump moving to the left. As the velocity signal increases the spiking rate on the left-rotation chip increases leading to an increase in the rotation speed of the bump. Throughout the experiment no right rotation signal was supplied to the system and thus there is no activity on the right rotation chip.

4.5.6 Moving the Bump




Fig 4.18. Moving the bump around the ring in both directions with variable rotation speed. The bump starts at rest. At 6.2 sec, the system is fed with an increasing leftward rotation velocity until it reaches a maximum of 153.2 deg / sec at 10.4 sec. Starting from 14.8 sec, it is slowly brought back to a full stop at 16.6 sec. In the second part starting from 21.3 sec, an increasing right-ward rotation velocity is fed to the system until it reaches a maximum speed of 123.3 deg / sec at 22.3 sec. At 29.3 sec, the rightward velocity signal is decreased slowly until the system reaches a rest at 33.3 sec.

To demonstrate the ability to move the activity bump around the ring in both directions, a variable head velocity signal is introduced to the system. Fig 4.18 shows an example experiment where the bump was initially at neurons 1, 2, 3, and 4, and then moved around the ring with variable (manually-controlled) left-rotation speeds. When the left rotation signal was taken away, the bump stopped at its last location and remained stable there. After three seconds, the right-rotation was applied and the bump rotated around the ring in the other direction before coming smoothly to a stop when the speed signal was removed.

4.5.7 Drift in position estimation

The HD system is an angular odometry system that maintains and updates the current estimate of head orientation by integrating angular velocity of the head. Two potential problems facing this system are inaccurate velocity integration and drift of the bump while the head is stationary. The brain can partially overcome this drift by using information from other sensory cues (e.g. vision) to correct or recalibrate the HD system. If the animal finds itself in a “familiar” (previously visited and memorized) location and the HD system reading is in conflict with the recalled angle, the system will be strongly driven by the memorized orientation to force the activity of the HD system neurons to point to this orientation [41].

In Fig 4.19 we compare the evolution of head position as estimated by our system with the head position as estimated by a perfect integrator when driven with three different velocities. The data show that the bump velocity is not constant as it travels around the ring. It should be noted that because the speed variation pattern is repeated on each cycle, these errors are primarily integration nonlinearities and not integrated noise. This type of error, can lead to a drift in the estimated orientation when integrating multiple small turns as seen in the biological HD system. Fig 4.20 shows the error between the estimated head orientation from our system and the corresponding perfect integrator for the three examples described in Fig 4.19. To determine the expected drift error for any given movement, we can use Fig 4.20 by calculating the difference in drift errors between the start and stop orientations. *For example, the worst-case movement at 29.7 deg/sec occurs when moving from 48 degrees to 155 degrees, resulting in an expected error of -11.5 degrees.*




Fig 4.19. The black, blue and green lines show the evolution head position as represented by the centroid of the bump of the activity as the system is driven to perform a 360° rotation around the ring using three different rotation speeds and the dotted gray lines show the expected head position if the same velocities were presented to a perfect velocity integrator system.

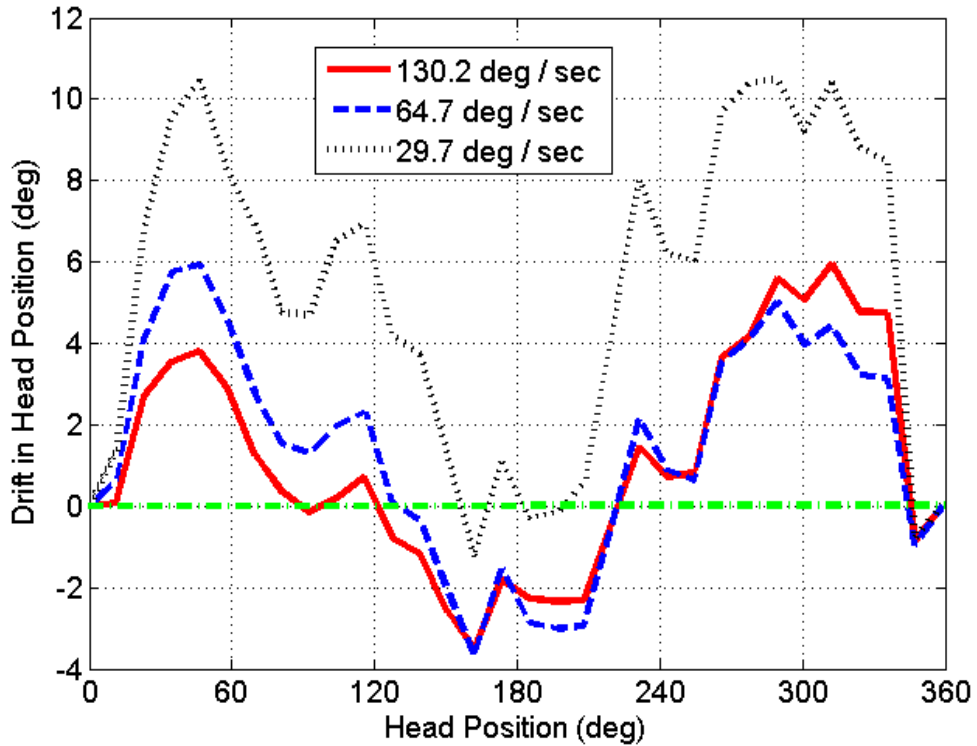


Fig 4.20. The error in estimating the head position in degrees as coded for using our system vs. using a perfect integrator for the same three examples shown in Fig. 19.

4.5.8 Resetting the bump

A noisy or weakly nonlinear integrator used in a spatial memory system (similar to what happens in the brain) should be capable of resetting the position of activity to correct for the drift. In this section we demonstrate that position-resetting is possible, moving the activity bump of the HD chip directly to different locations on the array of neurons. Throughout this experiment, the global inhibitory neuron is switched OFF and we rely on the individual neurons' inhibitory synapses to provide the necessary inhibition for the experiment by directly stimulating them using AER pulses. The measurement was begun with no activity on the array. At 8 seconds, a short burst of spikes were sent to neurons 10, 11, 12, and 13 that created a stable bump of activity at these neurons that persisted beyond the stimulation. At 32 seconds, the bump was relocated to neurons 20, 21, 22, and 23 by projecting inhibition to all of the neurons in the array just prior to stimulation of the neurons of the new location. The test was repeated at 48 sec to move the bump to neurons 5, 6, 7, and 8 using the same strategy as shown in Fig 4.21. Although the global inhibitory neuron was suppressed during this experiment, it only needs to be suppressed transiently at the time of transition.




Fig 4.21. Resetting the Bump Location: Externally-provided stimulation input is shown in big black dots, spikes from the ring neurons are in gray dots and the inhibitory neuron spikes are in red along the top of the graph.

4.5.9 Characterizing the HD System

The ability to create, sustain, and control the movement of a bump of activity on a ring of neurons provides the desired infrastructure for creating an HD system. This system acts by integrating the head velocity to represent and maintain a memory of the head orientation in space. For this system to act as an integrator of head velocity, it is important to control the speed of the bump around the ring linearly with the head rotation velocity as described in the following equations. Let $\theta(t)$ (degrees) be the head's orientation in space, with $\theta \in (0, 360]$.

We can write the equation

$$\frac{d\theta(t)}{dt} = V_r(t) \quad (1)$$

for the change of $\theta(t)$ as a function of the head rotation velocity $V_r(t)$ (deg/sec). By integrating (1) we can compute the head's orientation.

$$\theta(t) = \theta_0 + \int_{t_0}^t V_r(t) dt \quad (2)$$

θ_0 is the orientation of the head at the time the rotation started ($t = t_0$).

In the neuromorphic VLSI system, the head angle θ is represented by the location of the activity bump on the ring of (N) neurons. The angle is therefore discretized into 32 possible bump locations around the ring to represent the entire 360° range of angles for the HD system. Hence, the resolution of the system is $360/32 = 11.25^\circ$ where the angle 0° , for example, is coded by activity at neurons 0, 1, 2, and 3, i.e. at $N = 0$. Similarly, the angle $360 \cdot (31/32) = 348.75^\circ$ is coded by activity at neurons 31, 0, 1, and 2, i.e. at $N = 31$. In general, the head orientation θ (degrees) can be computed from the bump location N using

$$\theta = (360 \cdot N)/32, \quad (3)$$

and the location $N = n$ is defined as the activity being at neurons ($n, n+1, n+2$, and $n+3$), where $n \in [0, 31]$. Using difference equations to represent rates of change we could write

$$\frac{\Delta\theta}{\Delta t} = \frac{360}{32} \left(\frac{\Delta N}{\Delta t} \right). \quad (4)$$

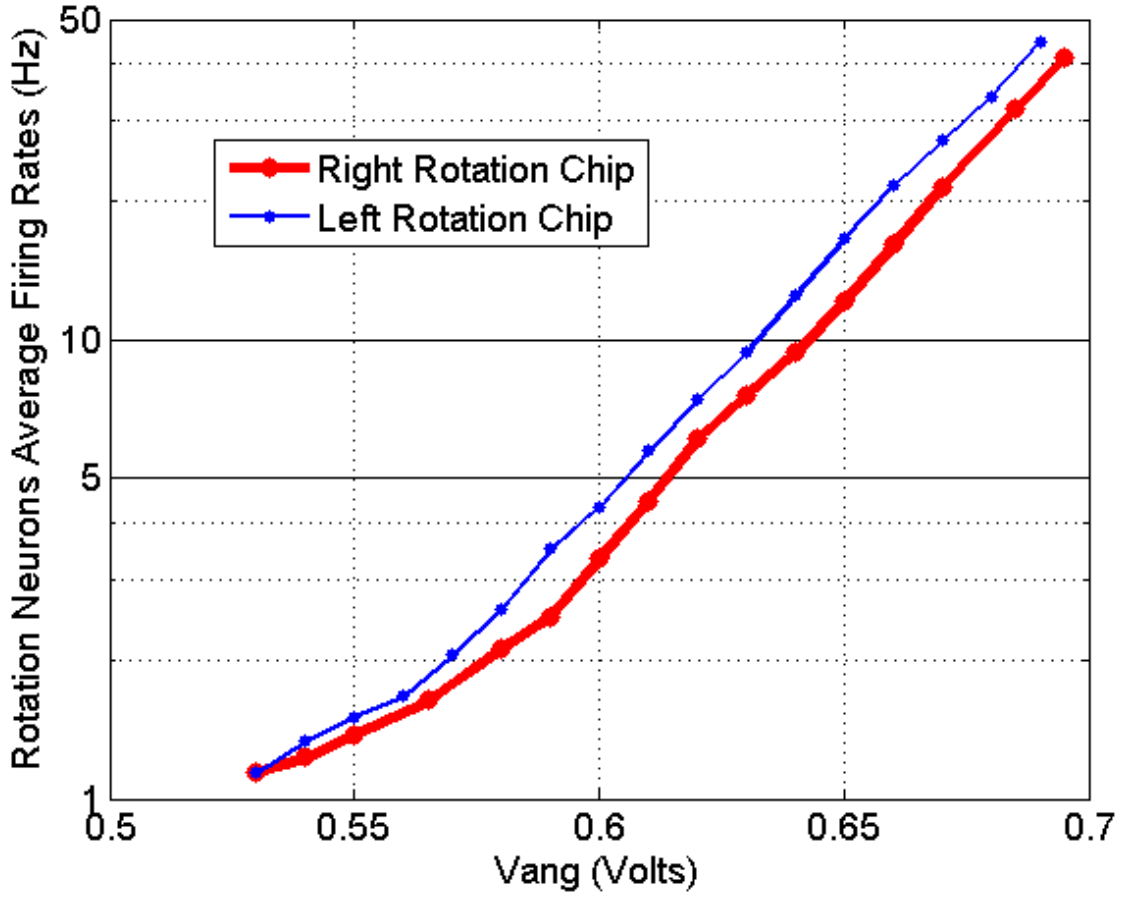


Fig 4.22. Average spiking rates on the rotation chips and the corresponding speed signal V_{ang} .

Fig 4.22 shows the measured relationship between the average firing rate of the left- and right-rotation neurons on each chip and the bias voltage V_{ang} (volts) used to control the V_{inj} bias of the left and right rotation chips. The angular velocity input is logarithmically represented by the applied bias voltage V_{ang} that controls a charging current of a rotation neuron, driving it to fire spikes. Because the current is an exponential function of the applied voltage, the spiking rate of the neuron linearly reflects the magnitude of the charging current and thus the angular velocity input.

$$f_{R,L}(t) = C \cdot \exp(\beta_{R,L} V_{ang(R,L)}(t)) \quad (5)$$

β and C are constants, and $f_L(t)$ and $f_R(t)$ are the left and right rotation chips mean

firing rate, respectively. While the measurements show a systematic mismatch between the response of the left and right rotation chips, it is possible to pre-compensate for this mismatch at the input.

In Fig 4.23, we show the relationship between the firing rates of the rotation chips and the bump rotation velocity we recorded from our system. This relationship is fairly linear throughout the entire range of operation. The rate of change of the bump location, $(\Delta N/\Delta t)$, can be computed as a linear function of the left and right rotation speeds, encoded by the firing rate of the left and right rotation chip neurons

$$\frac{\Delta N}{\Delta t} = \alpha_L f_L(t) - \alpha_R f_R(t) \quad (6)$$

where α_L, α_R are constants that can be determined by fitting straight lines to the data shown in Fig 4.23. Although signals for opposing directions should never be active simultaneously, the system is configured such that only one of the rotation chips is allowed to be active at a given time.

From equations (4) and (6) we can also write,

$$\frac{\Delta \theta}{\Delta t} = \frac{360}{32} \cdot (\alpha_L f_L(t) - \alpha_R f_R(t)) \quad (7)$$

In other words, the angular rate of change of the bump location is linearly dependent on the firing rate of the neurons in the left and right rotation chips in our system.

One important specification of the system is the usable range of head rotation velocities. In Fig 4.23 we show that the linearity of operation is maintained over a wide range of head rotation speeds from as slow as one neuron per second (11.25 deg/sec) to as fast as 50 neurons per second (562.5 deg/sec). The system is relatively well-matched for the two rotation directions.

Equations (6) and (7) show how the system uses the two head rotation velocity

inputs to drive the bump of activity around the ring, acting as a mathematical integrator.




Fig 4.23. The average bump velocity measured in (Neurons per Sec.) as a function of average spiking rates on the rotation chips. The vertical bars represented one standard deviation from the average.

4.6 Conclusions

In this chapter we present a neuromorphic VLSI system based on a recurrent, spiking neuron network that can create and maintain a stable, self-sustaining, activity pattern that can be used to mathematically integrate global rotation signals. The dsPIC® microcontrollers are programmable and allow the implementation of the AER routing in the system to have the flexibility to investigate different connectivity schemes. A fixed routing scheme could be easily implemented using simple digital

logic which will result in significant reduction in the power consumption and the entire system could be integrated onto a single chip.

An important observation of this work is that the mean-rate mode of operation is vulnerable to the problem of drifting towards attractors produced by the mismatch found in the neuron and synapse circuits, making it difficult to move the bump of activity in a controlled fashion. Drifting *is* avoidable with minimal mismatch and careful parameter tuning. On the other hand, in the “synchronized” mode of operation, the cyclical hard reset of the bump neurons by the inhibitory neuron prevents the small mismatches from building up over time, reducing drift [23]. Movement of the bump, however, now requires a stronger “kick”, which can easily be provided by the spiking left and right movement neurons.

Chapter 5 : Online Error Correction in HD System

5.1 Introduction

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...)[47, 72]. 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 [41, 92, 93]. 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.

In the previous chapter we presented a neuromorphic HD system, that worked nicely, however, it suffered from noise which leads to noisy integration and error in estimating the orientation of the head in the environment. In this chapter, 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 to keep the noisy estimate of the orientation aligned with the environment.

5.2 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 5.1 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 orientation. The head is equipped with a sonar system that can detect objects and report their range.

A detailed block diagram of the system is shown in Fig 5.2.




Fig 5.1. 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.

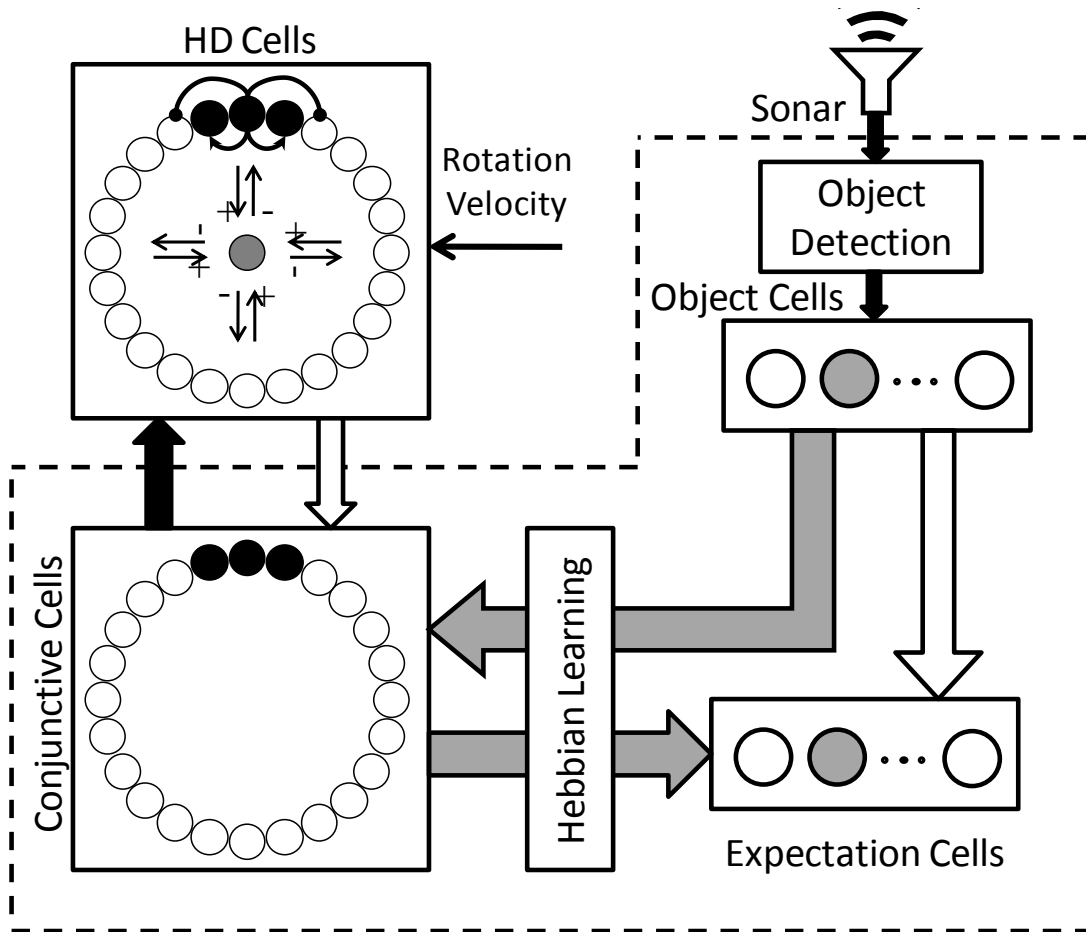


Fig 5.2. System block diagram. The blocks enclosed in the dashed-line box are implemented in software and the other blocks are in hardware. The black arrows indicate predefined non-plastic synaptic connections, the arrows in grey show the plastic synaptic connections, and the white arrows indicate non-plastic connections used as teacher signals to guide the learning process.

5.2.1 Sonar Object Detection and Object Cells

The sensory input to this system is an air-coupled, ultrasonic sonar mounted on a rotary platform that acts as the “head” in our system. The sonar sends short ultrasonic pings (~ 0.5 ms) in a relatively narrow beam and uses the echo time to determine the distances of different objects. For simplicity in our example, we used 2.5 inch diameter polyvinyl chloride pipes (a.k.a., “poles”) as objects and used the distance as a cue for identity. Following a sonar ping, the echo envelope waveform from the receiver is provided to a simple perceptron [94] array whose outputs indicate which object has been detected or whether no object has been detected. The “no object cell” is a special cell that is inhibited by all of the other object cells and only becomes active when no targets are detected.

5.2.2 Head Direction Cell System

We use a neuromorphic VLSI-based HD system (that is fully described in [95]) that integrates angular velocity signals (rate-encoded digital pulses, or “spikes”) to maintain an estimate of the head orientation. In the present system, angular velocity signals are generated from the angle encoder of the sonar platform. Due to transistor fabrication mismatch and practical limitations of calibration (drift in parameters, nonlinearities, etc.), however, this system suffers from errors in the integration process which accumulate over time.

5.2.3 Conjunctive Cells

Conjunctive cells were first discovered in the entorhinal cortex of rats [96]. These cells are thought to integrate information about location, direction and distance

to provide the necessary information to move the activity in the grid-cell network. In our working model, we extend the definition slightly to include information about objects in the environment (i.e. sonar objects) to form a combined representation for the environment that can be used to correct for any drift in the HD representation similar to the model presented in [67] where visual objects were used to perform the correction for the drift. The conjunctive cells receive excitatory connections from the HD cells to mirror their activity but also project excitation back to the HD cells to move their activity to correct for any drift when more reliable information about orientation (i.e., sensory-triggered memory) is available. The conjunctive cells also receive plastic excitatory connections from the object cells to learn the association between detected objects and orientations. The plastic connections are modified using a simple Hebbian learning rule. Moreover, the “no object cell” does not send any projection to the conjunctive cell layer to avoid activating any conjunctive cells when no target is present.

5.2.4 Expectation Cells

The expectation cells are a set of neurons that learn the association between orientation and sensory input like the conjunctive cells, but they instead reflect the activity of the object cells (i.e., sensory input). These cells learn to recreate the object cell pattern present at each orientation through weak excitatory projections from the object cells (teacher signal) and plastic connections from the conjunctive neurons. If, at a particular orientation, the sensory system does not detect an object that was previously seen there, the expectation cells reactivate the expected pattern.

The expectation cell layer has a recurrent set of connections that implement a winner-take-all function that allows only one neuron to be active at a time.

5.2.5 System Operation

Our system is designed to continuously learn the environment in which it is operated, but when *a priori* knowledge about the space exists it compares its expectations about the environment with the information coming from its sensors to determine whether to learn a new association, to fix its internal estimate for its location, or understand that it is lost and should not learn new associations until it can realign itself with the environment.

The system is initialized with all plastic synapses set to zero, indicating that no objects are associated with any direction. This includes the “no object” state. Rotating the sonar platform generates rotational velocity signals that are provided to the HD system and subsequently integrated to an estimate of position. During the first rotation of the sonar around the environment, the conjunctive cells are only driven by the HD cells since no objects have been associated with any position. As it rotates, the object-to-conjunctive (OC) synapses rapidly learn to activate the conjunctive cells (i.e., orientation) associated with the observed objects. At the same time, the conjunctive-to-expectation (CE) synapses rapidly learn to activate the expectation cells (i.e., object identification) associated with each orientation. Note that the “no object” cell does not have an OC synapse, but does excite its “twin” in the expectation cell group, thus learning to expect the “no object” situation.

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.

5.3 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°).

5.3.1 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 5.3 shows an example of the activity in the system.

5.3.2 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 5.4 shows the activity of the cells in this case.

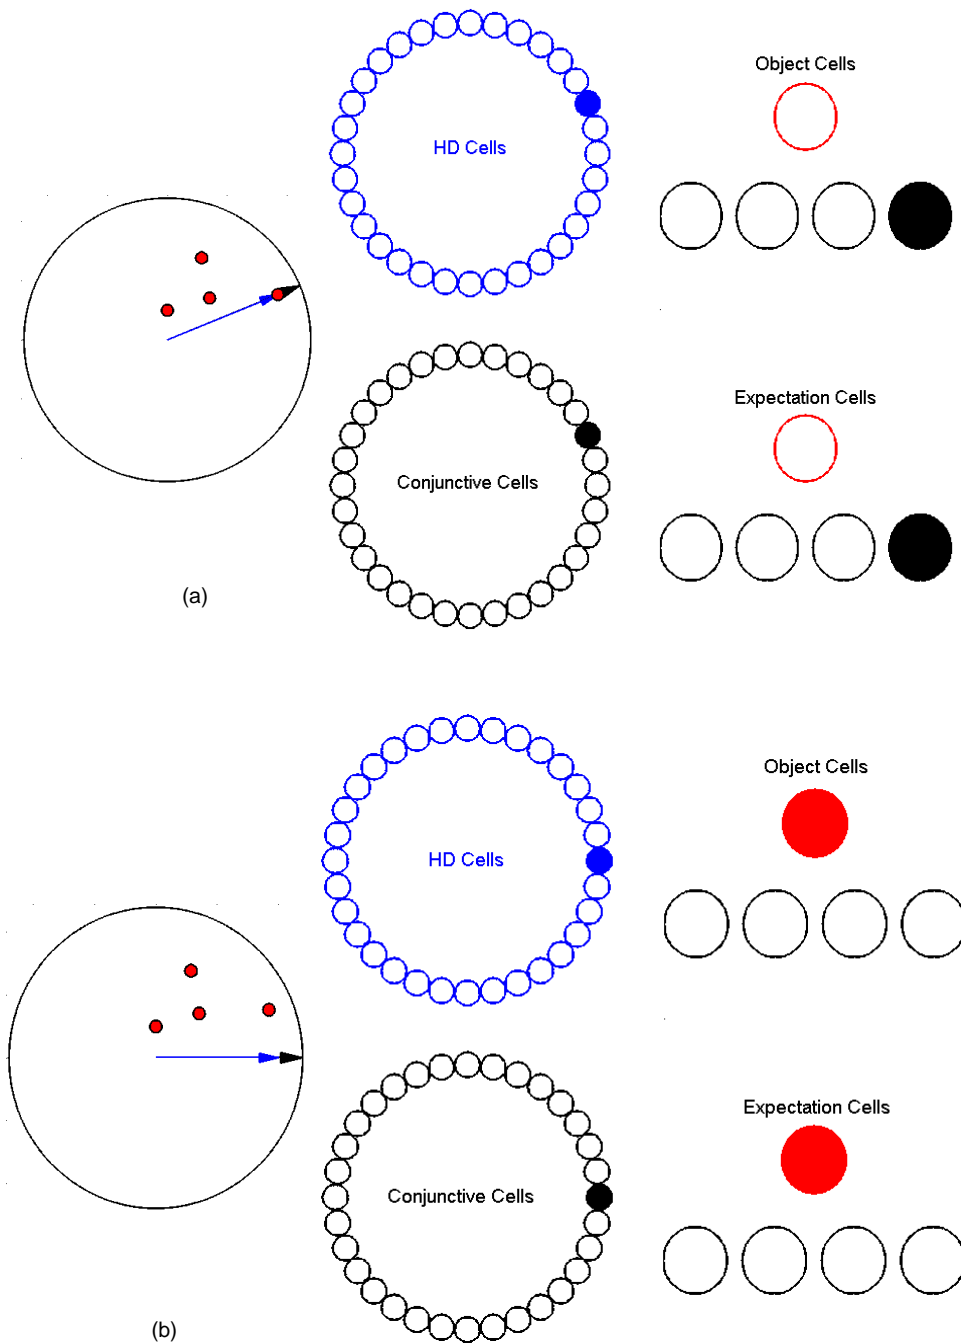


Fig 5.3. This figure shows the activity in the system when the HD estimate is aligned with the actual orientation in space. The left panel shows a schematic for the arena with the targets as red circles. The black arrow represents the actual head position and the blue arrow is the position as estimated 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.

5.3.3 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 5.5 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.

5.3.4 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 5.6 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.




Fig 5.4. 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.

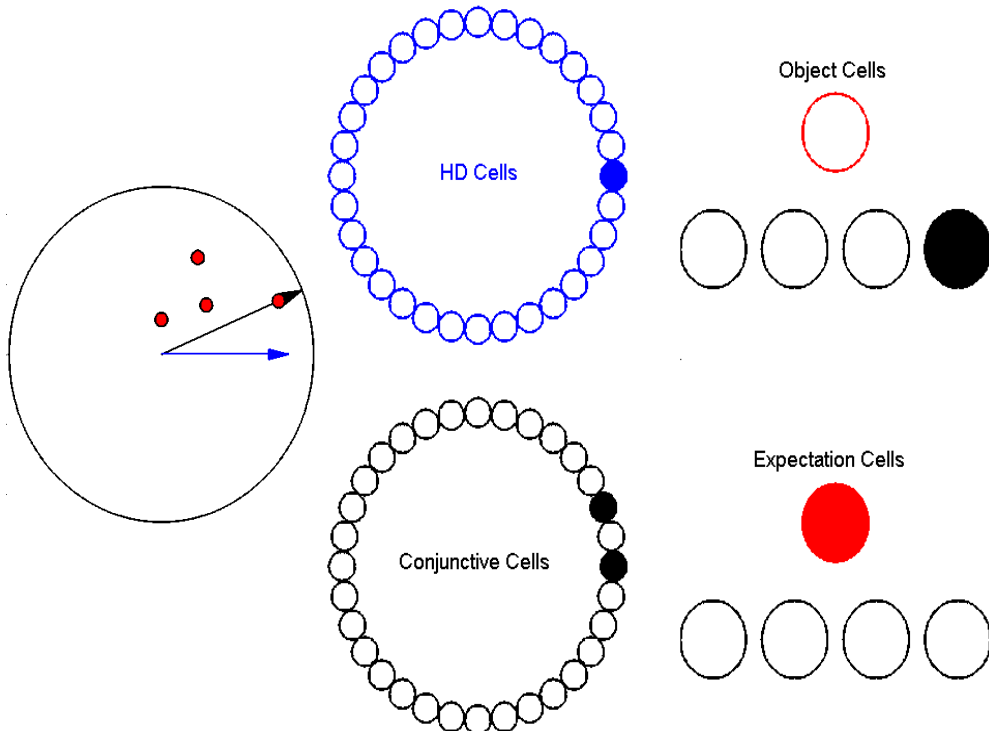


Fig 5.5. The system will reset. Based on the HD estimate of orientation, the system was not expecting to see a target, however, the sensory data show the presence of object #4. The system will reset the HD system to point toward the position of object #4. Note that the activity in the conjunctive cells reflect both the HDS' estimate and the stored orientation.




Fig 5.6. 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.

5.4 Conclusion

In this chapter we have presented an extension to the neuromorphic head-orientation odometry system presented in chapter 4 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.

Chapter 6 : Grid Cell System

6.1 Introduction

Discovered in 2005, grid cells revolutionized the way we think about the space encoding in the mammalian brain [19]. These cells respond periodically as the animal walks in an environment such that the response fields form a hexagonal grid anchored to the environment without being influenced by environmental cues. Neural recordings have revealed that cells in different layers exhibit different spatial frequencies in their firing fields [19]. The discovery of these cells changed the way we think about place cells and led to novel theories on how the place cells of the brain can be formed. One of the leading theories about how grid cells are used is to sum the activity from different spatial frequencies like a Fourier series to produce a pulse-like place cell response [72]. Fig 6.1 shows the response of 4 grid cells from the entorhinal cortex of a rat, modified from [19]. In chapters 2 and 3 we went over the neurobiology of these cells and over the models proposed for how they can be formed. In our quest to build neuromorphic VLSI system that can be used for spatial navigation, we built head-direction cells [97] and online, sensor-based realignment of orientation estimates [98]. In this chapter, we describe our efforts to develop a compact VLSI circuit for modeling a population of grid-cells.

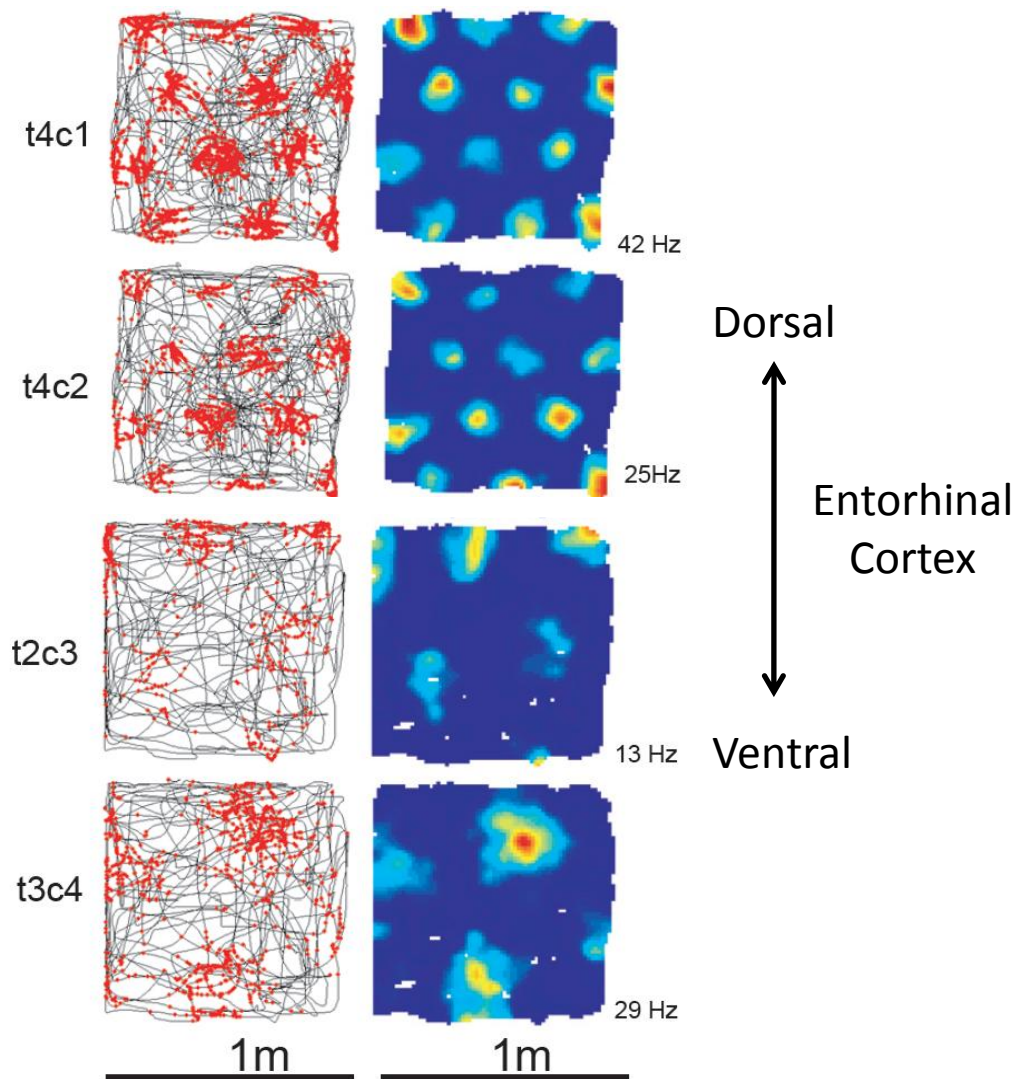


Fig 6.1. Response of 4 grid cells from the entorhinal cortex of a rat, modified from [19].

6.2 Modeling the Grid Cells

In chapter 3 discussed the available models for the grid cell operation. There are two main classes of models for grid pattern formation; one suggests that the grid pattern is created by attractor dynamics due to feedback interconnections in the MEC formation and the other suggests that the grid pattern is a result of constructive interference of multiple subthreshold oscillations injected into the network's neurons [72, 73]. In this chapter, our implementation of the grid cells is based on the model using attractor dynamics. In the continuous attractor models, a single position in space is represented by an attractor state in the system. As the animal moves, the network activity shifts to “nearby” attractors.

6.3 Simulating Grid Formation

We used Matlab to simulate this model; we used a sheet of (25 x 25) neurons, with toroidal boundary conditions and random initial interconnection weights. For this sheet of neurons to exhibit grid like activity, the neurons need to be interconnected using a difference of Gaussians-like weight pattern. We have tested the hypothesis described in [72], for the development of such weight pattern in the entorhinal cortex starting from random initial connections; the hypothesis assumes that spontaneous waves of activity (similar to the ones forming the retinotopic map [72]) occur at the entorhinal cortex during development and simple Hebbian learning is applied to modify the connections. In our simulations we have obtained a difference of Gaussians connection pattern as required for the grid formation. We monitored the evolution of the activity in the network from random initial activation levels till it reached a steady state level which showed a grid like activity over the

entire sheet. The biological grid cells have multiple spatial frequencies; this can be obtained by varying the width of the interneuron excitatory connections, by varying the frequency of the spontaneous waves of activity used to generate the weights. Fig 6.2 shows two examples for the weight pattern generated for the neuron at location (0, 0) on the sheet. Fig 6.3 shows the response field of a neuron in a virtual arena for the weight patterns from Fig 6.2.




Fig 6.2. Two examples of interconnection weights for neuron at location (0, 0) on a 2D sheet of neurons used to simulate grid cells activity.




Fig 6.3. Firing rates (in Hz) of one simulated grid neuron over a virtual arena using the connection weights schemes shown in Fig 6.2, showing that the spatial frequency of the response grid changes by changing the width of the connection matrix for the neurons on the sheet. The neuron in the left panel is connected to its neighbors using the weight pattern shown in Fig 6.2 (left) and the one in the right panel uses the weights from Fig 6.2 (right panel).

6.4 *The Grid Circuit*

To simulate the grid cell network, we designed and fabricated a two-dimensional, current-mode, pixel-array circuit that implements the attractor system using mutual excitatory and inhibitory projections between each pixel and its neighbors. The pixels are arranged on a Cartesian grid with continuous (i.e., toroidal) boundary conditions to eliminate edge effects in the operation of the system. Although the output from each pixel is a digital spike, the core pattern generation circuits are operated as analog circuits. To save significant chip area, however, we have implemented a clocked pixel-readout for the spiking neurons that conform to the address-event protocol (AER) [89] for integration into other AER-based systems. When a pixel generates a spike, its address and a communication request signal are sent off the chip and when the receiving system has captured the information it sends back an acknowledge signal. We used minimum size transistors for the digital circuits, for the analog circuits we used transistors with ($W/L = 9/9 \lambda$).

6.4.1 Chip Organization

Fig 6.5 shows the connectivity diagram of one pixel of the 16 x 16 array. To implement the desired attractor network, each pixel sends and receives excitation locally and inhibition more distally. By controlling the relative strengths and widths of the excitatory and inhibitory kernels, the shape of the activity in the array can be controlled. Each pixel directly interacts with its four nearest neighbors (above, below, right, and left) in the grid (we used the letters N, S, E, and W to label these connections). A set of current diffusor [99] circuits implement and control the spread of the inhibitory and excitatory kernels.




Fig 6.4. Micrograph of the Grid Chip. The chip has an array of 16 x 16 analog pixels. Scanner circuits are used to select pixels, allowing the activity of each pixel to be read out in the form of AER pulses

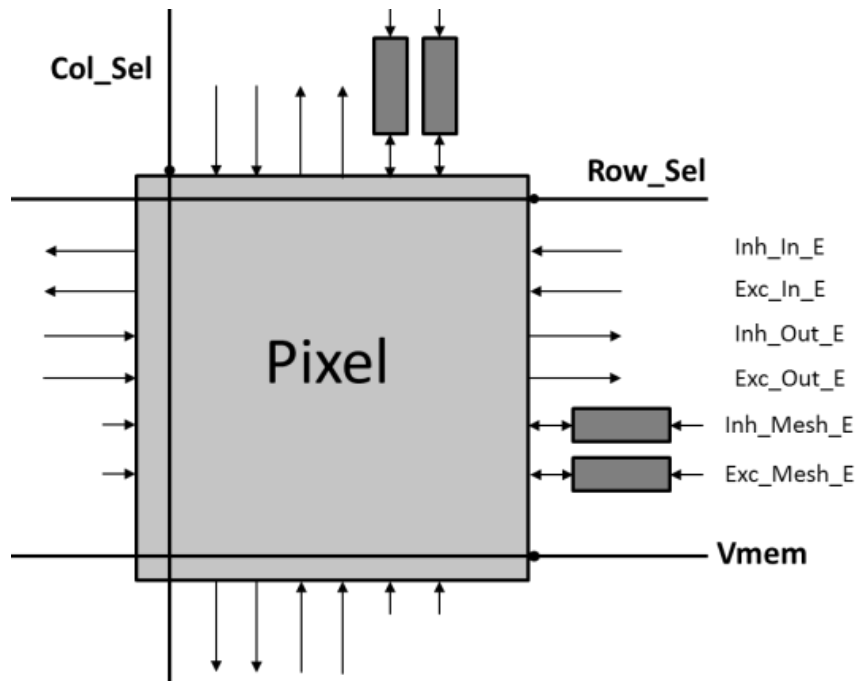


Fig 6.5. Diagram showing the connectivity of one pixel. The 16 x 16 array is formed by tiling this pixel. Each pixel receives row and column select lines, is connected to each of its neighbors by six local signal lines two of which go through pFET diffusers (small shaded box), and shares a global output voltage line (Vmem).

The output current level at each pixel is integrated onto a local capacitor to mimic a neural membrane capacitance. To implement the synchronous spiking readout, a 2D scanner that uses a *clk* signal generates *Col_Sel* and *Row_sel* signals to address the pixels to send the capacitor voltage to a common spike-generation mechanism. If the local capacitor voltage exceeds a threshold, a spike is generated and the membrane voltage is reset.

6.4.2 The Pixel

Fig 6.6 shows the schematic of the pixel circuit. Each pixel consists of a core circuit (M13 and M14 near the center) and four main circuit blocks (the ‘projection network’, ‘inhibitory input’, ‘excitatory input’, ‘readout’). The drain current of M14 represents a constant excitatory input current to the pixel which is controlled by the global voltage signal *V_{in_pixel}* for all pixels on the chip. The current flowing in M13 is the output current of the pixel which differs between pixels due to the extensive interconnectivity.

6.4.3 Projection Network

This network mirrors the pixel’s output activity (the current in M13) to both the excitatory kernel (using pFETs M2, M4, M6, and M8) and the inhibitory kernel (using pFETs M1, M3, M5, and M7) which project directly to the four nearest neighbors. The current going to the excitatory network additionally goes through a differential pair (not shown) that is used to control the center of the excitatory projection kernel. In this paper we only address the case where the excitatory component of the kernel is centered.

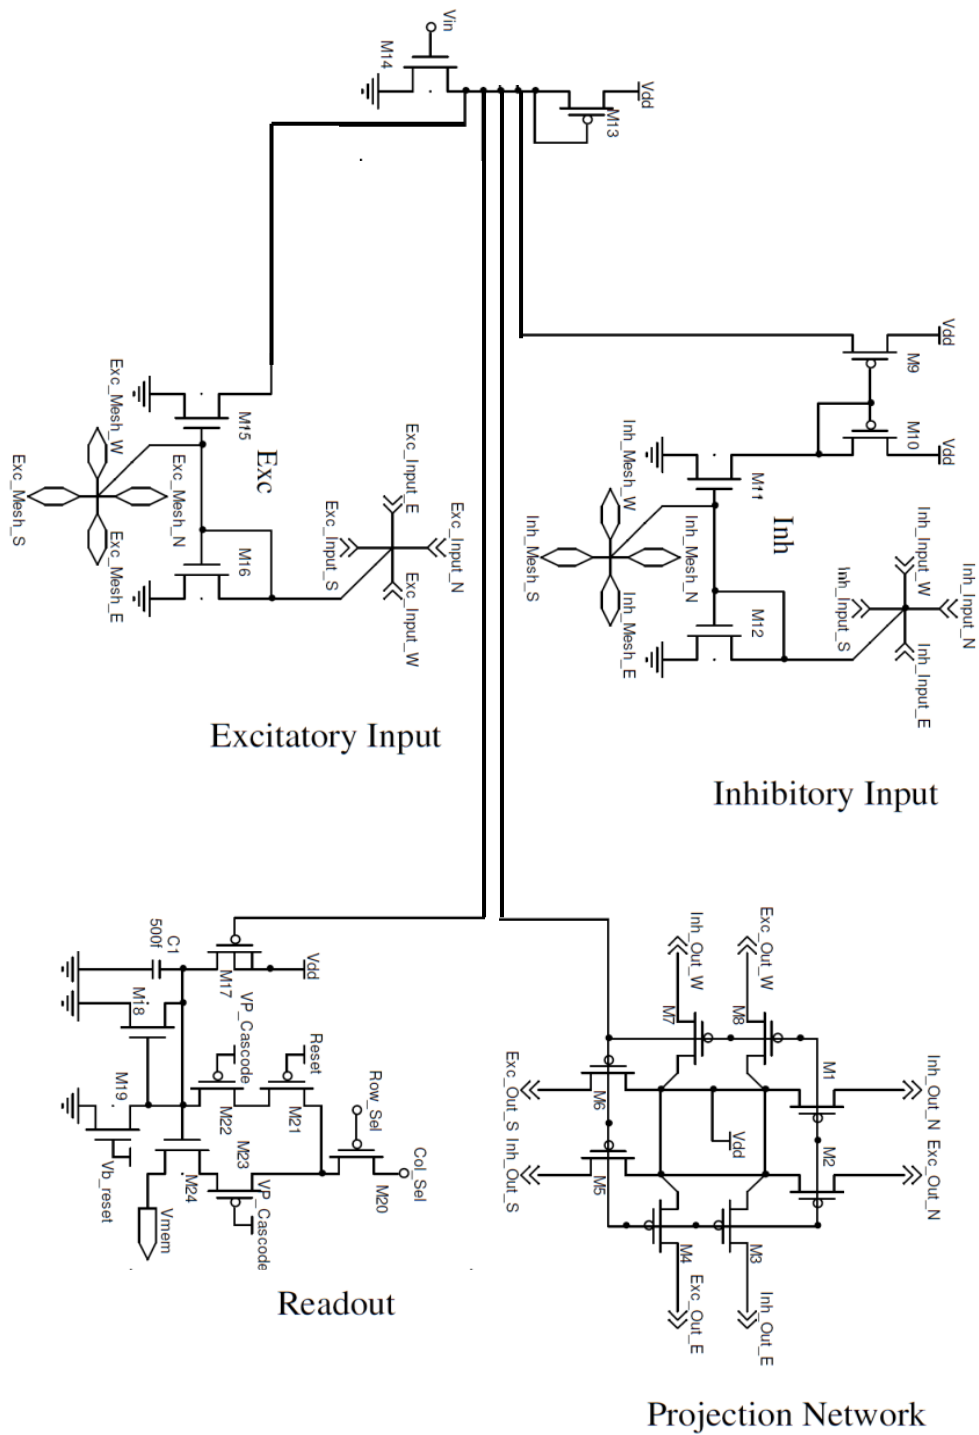


Fig 6.6. Pixel Circuit. The pixel circuit consists of the central node (between M13 and M14) and four circuit blocks that implement the inhibitory and excitatory projection kernels and a readout circuit.

6.4.4 Inhibitory Input

The inhibitory component of the interconnection kernel is created by summing the four nearest-neighbor inhibition currents (“*Inh_Input_x*”, where x varies with direction) and adding them to the current diffuser network (at node “*Inh*”) to compute the local average inhibition (via the connections “*Inh_Mesh_x*”). A copy of this local average inhibitory current is then used to inhibit the pixel locally. The local average inhibitory current is found at the drain of M11 and is mirrored to the central summing node via the pFET mirror (M9 and M10). The same computation is applied to the currents arriving from the neighbors in other directions.

6.4.5 Excitatory Input

This network is similar in operation to the inhibitory network. Like the inhibitory network, neighboring excitatory currents are summed and a local average excitatory current is computed at the node labeled “*Exc*” using the diffuser network. Similarly, the diffuser transistor gate voltage controls the width of the excitatory kernel. In this case, however, the current is not mirrored so that the excitatory current adds in same direction as the constant bias current transistor M14.

6.4.6 Readout

To observe the spatial pattern of current in the array, a two-dimensional scanner circuit is used to scan the array pixel by pixel. A pixel is selected when *Col_Sel* is logic high and *Row_Sel* is logic low, activating the readout circuit of that pixel. The output current of the pattern formation network flowing in M13 is mirrored to M17 which charges a capacitor C1. M24 operates as a source-follower when the

pixel is selected and a voltage-shifted version of the capacitor voltage appears at the node V_{mem} , a common readout node to all pixels of the array. The biasing circuit for M24 is discussed in the next section. If a pixel drives V_{mem} above a threshold, an output spike is sent off-chip and a reset signal will be generated to reset the pixel's capacitor. Details of the generation of these signals are discussed in a later subsection. The V_{b_reset} signal creates a refractory period to control how long the pixel has to wait before it is allowed to begin integrating again.

6.4.7 Request Generation

Fig 6.7 shows the circuit that receives the output voltage of the pixel (V_{mem}) and decides whether to generate an output spike or not. M25 is the bias transistor for the source-follower circuit (M24) in each pixel. The input voltage V_{mem} is buffered by a follower circuit (M26-M30) and compared to a threshold voltage by a wide-range transconductance amplifier (M31-M39). This comparator's output pulls Req_bar logic low only if the output voltage is higher than the threshold voltage, triggering a request to the AER system.

6.4.8 Reset Generation

Fig 6.8 shows the circuit that generates the *Reset* signal after a pixel has triggered a request for an output spike. After the AER system sends out the address of the pixel that generated the spike, the AER system pulls the Ack_bar signal logic low, which in turn will pull the *Reset* line to logic low. *Reset* is a global signal (connected to all pixels) but only the pixel selected by Col_Sel and Row_Sel in Fig 5.2 will reset and all other pixels are not affected.

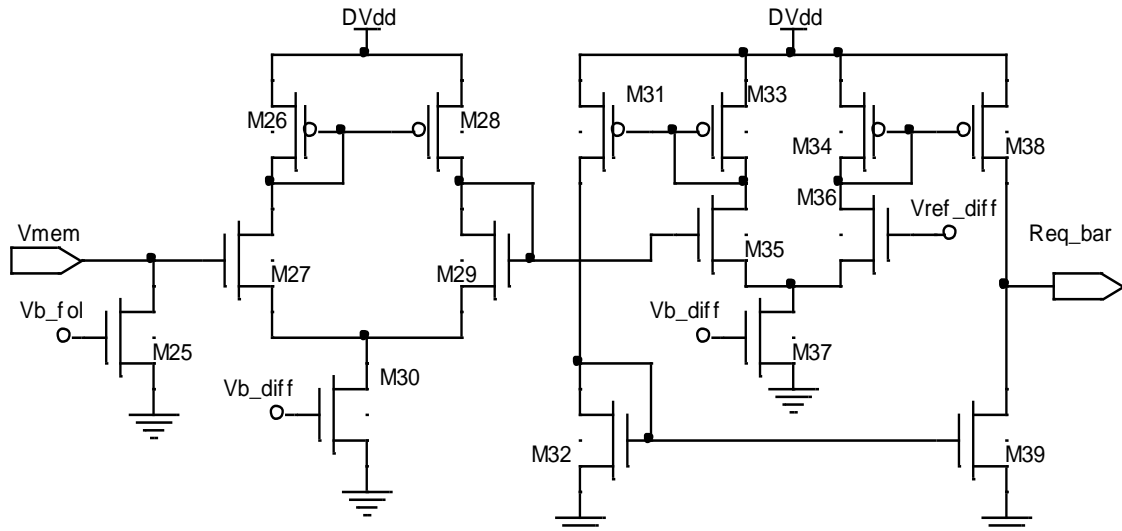


Fig 6.7. Request generation. M25 provides the bias for the source-follower transistor in every pixel (M24 in Fig. 1). The voltage output from the selected pixel is buffered and compared to V_{ref_diff} , to generate a Req_bar signal when the pixel voltage is higher than the spiking threshold.

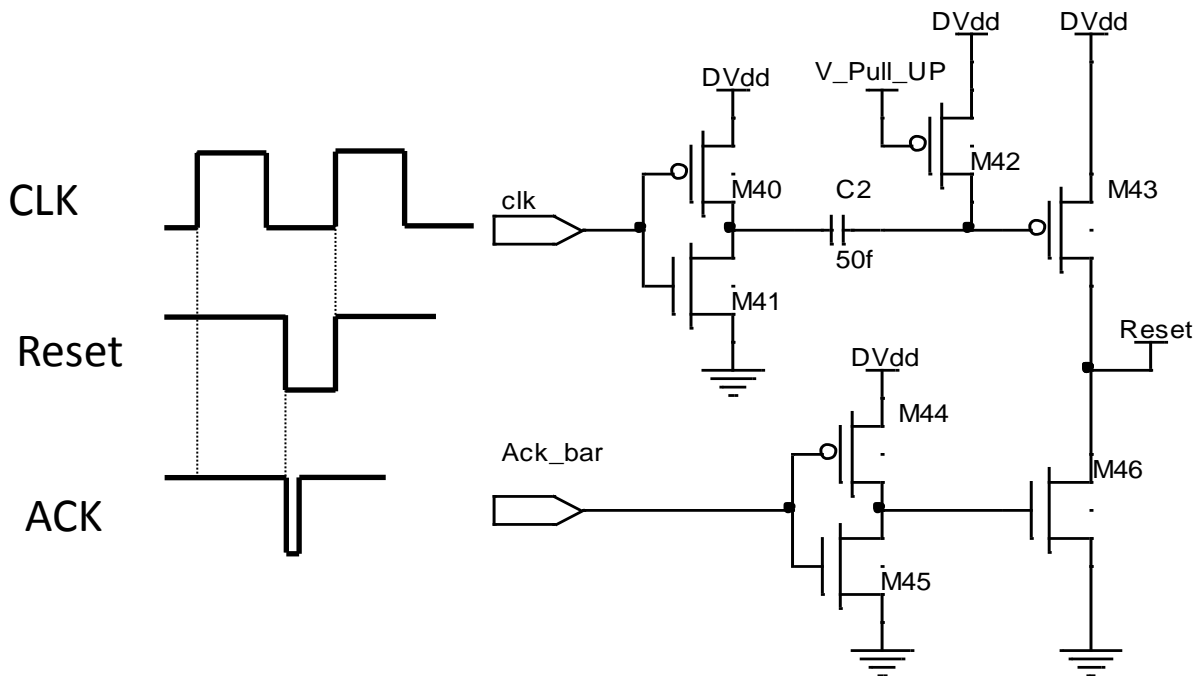


Fig 6.8. Reset generation. The Ack_bar signal coming from the AER system pulls the $Reset$ signal low. The $Reset$ signal can only be raised high by the rising edge of a new clk cycle as the scanner selects the next pixel.

To raise the *Reset* signal, the positive edge of the scanner *clk* signal is to create a logic-low voltage pulse through *C2* that raises *Reset* using the pFET M43. The bias voltage *V_pull_up* controls the time required for the gate of M43 to be pulled back to *DVdd*. As pixels are selected and fire output spikes, the *Reset* line is only raised at the end of the cycle, preventing multiple output spikes per clock cycle.

6.5 Chip Results

A 1.5mm x 1.5mm chip was fabricated using a commercially-available 0.5 μm CMOS process. We have tested the functionality of the chip to generate grid-like patterns, by providing a constant current to all of the pixels and selecting the control voltages for the diffusor networks to create narrow excitatory kernels and wider inhibitory projection kernels. We used a *clk* frequency of 100 KHz to scan the activity of the pixels.




Fig 6.9. Raster plot of 0.2 sec of activity recorded from the chip, this raster translates to the patterns shown in Fig 6.10. The rasters show pixels firing at different frequencies which create analog bumps of activity on the chip.




Fig 6.10. Reconstructed grid cell circuit activity pattern on the chip at three different spatial frequencies. The brightness of the pixel represents the rate of spikes generated by the pixel (spikes/sec).

The activity of the chip was recorded as AER spikes with the 16 x 16 array of pixels resulting in addresses from 0 to 255. Fig 6.9 shows example raster plots of the spiking activity as recorded from the chip. Since the chip is organized in rows and columns it is simple to construct 2D images using the spike rate as intensity. The spike generation is linear with the local pattern current, thus the image reflects the spatial pattern linearly. Fig 6.10 shows example grid patterns of activity on the chip for three different spatial frequencies. These three plots were recorded from the same chip by changing only the two parameters controlling the width of the excitatory and inhibitory projection kernels.

6.6 Modeling the Chip Connectivity

In order to qualitatively describe the operation of the chip, we present in this section, a mathematical analysis of the connectivity pattern on the chip.

The chip has two resistive networks (one for excitation and the other for inhibition) connecting the pixels, in the analysis we are using superposition to compute the operation of the circuit.

6.6.1 1-Dimensional Analysis

To simplify the analysis we will start by a 1-Dimensional array of pixels and study the effects of the interconnections on the activity of the network. For n pixels, each pixel i has an input current x_i and an output current y_i . The current y_i is directly projected to the neighboring pixels $i-1$ and $i+1$, each one of these neighbors diffuses back a part of that current to the pixel i through the diffuser networks, an excitatory part ay_i and an inhibitory part by_i . Moreover, each neighboring pixel $i-1$ and $i+1$

sends a copy of its current y_{i-1} and y_{i+1} to the pixel i , a part of that current diffuses back to the original pixels and another part act on the pixel i itself cy_{i-1} and cy_{i+1} excite the pixel i and dy_{i-1} and dy_{i+1} inhibit the pixel. Putting all these currents together and applying Kirchhoff's current law at the node i we get

$$x_i = (1 + 2ab)y_i + (d - b)y_{i-1} + (d - b)y_{i+1} \quad (1)$$

When we apply Kirchhoff's current law at each pixel we get

$$b = 1 - 2*a, d = 1 - 2*c \quad (2)$$

This is a system of linear equations that can be solved for y_i 's

$$\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} 1 + 2ab & d - b & 0 & \dots & d - b \\ d - b & 1 + 2ab & d - b & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ d - b & 0 & \dots & d - b & 1 + 2ab \end{bmatrix} \times \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} \quad (3)$$

Assuming a 16 x 1 array of pixels with circular boundary conditions to mimic the connectivity on the chip and running a numerical simulation for this 1-D version of the system we show the unit step function of the connectivity and the effect of changing the α and β factors which represent the diffusion factor in the excitatory and inhibitory networks. Fig 6.11 shows the response of a 16 x 1 array of pixels interconnected as described in equation (2) for an impulse input at pixel 8. The figure shows that the response of the circuit exhibits the difference of Gaussians pattern required for creating the grid-pattern.

6.6.2 2-Dimensional Analysis

Using the same concept we applied for the 1-D analysis in the previous section, we present here a 2-D system of equation describing the connectivity in the

network as well as some numerical simulations from Matlab that replicate some the data we recorded from the chip.

Using the same terminology as the previous section with an addition of an index variable j for the other dimension, we can write the steady state equation linking the output of the pixel y_{ij} with its input x_{ij} as follows.

$$x_{ij} = (1 + 4ab)y_{ij} + (d - c)y_{i-1j} + (d - c)y_{i+1j} + (d - c)y_{ij-1} + (d - c)y_{ij+1} \quad (4)$$

The constraints defined in (2) become

$$b = 1 - 4*a, d = 1 - 4*c \quad (5)$$

Fig 6.12 shows Matlab simulation for a 16 x 16 array where a grid-like response emerges on the array, similar to what we record from the chip.

To study the stability of the system, we evaluated the condition number for the system of equations described by (4) over the entire parameter range of a and c the results of that analysis is shown in Fig 6.13. The analysis shows that the system is stable over the blue strip in the figure which is a narrow band where a and c values are close and is mostly unstable otherwise.




Fig 6.11. Matlab generated response of a 16 x 1 array of pixels using the model of the interconnection on the chip for an impulse input at pixel 8. ($a = 0.17$, $c = 0.2$)




Fig 6.12. Matlab generated response of a 16 x 16 array of pixels using the model of the interconnection on the chip showing the emergence of grid-like patterns on the pixels ($a = 0.17$, $b = 0.2$)




Fig 6.13. Condition Number results for the system of Equations described in (4) for the entire range of parameters of a and c .

6.7 Moving the Grid Pattern

The chip that we designed for emulating the grid cells, we had two objectives, the first was to create grid-like patterns of activity over the pixels of the chip, and the second was to move the patterns of activity over the chip in a controlled way using only analog displacement velocity signals. In the chip presented in this chapter we had circuits that were designed to realize this functionality, and we did the proper simulation of the functionality of these circuits before sending it out to foundry. Unfortunately, when testing the chip, those circuits did not perform as well as we expected and we were not able to move the grid pattern over the chip in a controlled way. The mismatch appeared to be significant enough that the movement controls did not move the pattern, but appeared to grow the bump size in the desired direction

until the pattern catastrophically bloomed into a stripe spanning the entire chip. To characterize for the level of mismatch on the chip, we ran Monte Carlo simulations on the operation of the circuits we have on the chip. Varying the threshold voltage of transistors, we noticed that when the mismatch in the transistors' threshold exceeds 15%, the movement of the grid pattern becomes distorted and the movement becomes irregular or stuck at certain locations.

Since the operation of the grid cells was crucial for the creation of place cells, we decided to move to software based grid cells and use that as the basis of creating the place cells as we discuss in the next chapter.

6.8 Conclusion

In this chapter we have presented a neuromorphic chip that can produce the spatially-periodic, hexagonal activity patterns believed to exist in the grid cell layers of entorhinal cortex. Multiple spatial frequencies were demonstrated. The grid patterns were stable and did not drift with time. Although the pixels exhibit significant mismatch, the large-scale pattern is quite evident.

Chapter 7 : Place Cell System

7.1 Introduction

Spatial navigation is a critical survival skill for many animals. Neurophysiological studies in mammals of the mechanisms underlying this skill have revealed a variety of neurons that support internal models of position that integrate both internally-generated (e.g., vestibular and proprioceptive signals) and externally-generated (e.g., vision or sonar) cues. As we discussed earlier, a subset of neurons in the hippocampus are neurons called “place cells” which become active when the animal is in a particular spot in a given environment [11]. Those cells are thought to be activated by the convergence of activity from both a translational and rotational internal estimate of the animal’s movements [72]. In addition, sensory cues in the environment are used to “anchor” the response fields to reference objects in the world to correct for drifts in the estimate. The internal *orientation* estimate is performed by the “head-direction cell” system [42] and the internal *translational* estimate is performed by the “grid-cell” system [96]. Together they provide the information needed to produce the place-cell responses. Using inputs from multiple spatial frequencies to create place cells is similar to a Fourier construction of a bump of activity [72].

The biological data show that the formation of stable place cell responses in a certain environment can develop during the first encounter of the animal in this space [47]. This fast emergence of place cell responses without extensive experience implies that the animal did not slowly create the place fields based solely on the

sensory information in the space. The odometry data is always available, however, since it comes from measurements of its own motion and motor signals, and can be used to create the place cell responses. On the other hand, the place cell system can quickly learn to link sensory cues (landmarks) in the space to the place fields; rotating salient cues in an environment with stable place fields result in a corresponding rotation in the place cell responses.

In previous chapters, we presented our implementation for a neuromorphic head direction cell system and grid cell system. In this chapter, we discuss how the place cells can be created using the activity from grid cells and head direction cells.

The results we present in this chapter are mainly from simulations using MATLAB® to test our ideas and plan its implementation in future hardware.

7.2 Head Direction Cells to Place Cells

Head direction cells serve to keep track of an animal's orientation in a given space. This information is extremely valuable in spatial navigation to guide the direction of movement of the grid activity pattern across the grid cell layers [47]. In the next section we discuss how grid cells can be used to create place cells.

7.3 Grid Cells to Place Cells

As discussed earlier, response fields of Grid Cells have been found to have different spatial frequencies. It has been proposed that the grids of activity have been move over the sheet of neuron to follow the animal's displacement in the environment [100]. The grid patterns associated with neurons in the entorhinal cortex

are thought to be a result of the connectivity between the neurons as described in the previous chapter.

There are two main theories regarding the velocity input to the activity of grid cells [71, 72]. The first theory states that each layer of Entorhinal cortex has a different connectivity pattern, creating the different spatial frequencies of the grids and that the velocity signal drives all the layers equally to produce the motion. The second theory states that all the layers of the entorhinal cortex that exhibit grid cells have similar connectivity patterns and that the different spatial frequencies emerge from the drive of the velocity to the grid cells. It is still not known which one of these two theories is true if any. Since our aim is to create place cells, we have investigated both theories in our simulations and in this section we present the results.

7.3.1 Grid Cells with Similar Velocity Drive and Different Connectivity

In this section we test the first assumption that states that the displacement velocity signal drives the different layers of the entorhinal cortex equally, thus the difference in the spatial frequencies between the grid layers is due only to the connectivity patterns between the neurons of each layer.

In our simulations we address the reverse of the problem, by considering the following question, if we have a group of grid cells and an agent moving randomly in an environment, can a place cell be created at a specific location in the environment using grid cell activity alone? In the section to follow, we use Hebbian-learning to find the projection weights from neurons in the grid cell layer to a place cell such that a place field can be formed in the selected location.

In this case, we have grid cells with three different spatial frequencies. The spatial frequencies were picked arbitrarily with one constraint that they are multiple harmonics of the same base frequency. The initial phases of the grids are picked at random for the initial position of the agent in the space. The agent moves around randomly in the space; whenever it enters into the desired place field for the place cell, a teacher signal forces the place cell to fire and Hebbian-learning strengthens the projections from the active grid cells to the place cell. If the agent is at a location outside the place field but the place cell is activated, an anti-Hebbian learning rule “unlearns” that location. The size of the arena, the location of the place field, and the initial position of the agent in the arena are arbitrarily chosen.

Fig 7.1 shows the three grid cells we used in this simulation in which the three cells have different spatial frequencies. The cells’ activity is discrete with three levels of activity. Fig 7.2 shows two examples of place fields that were trained in our system. We are able to train cells with place fields as small as one pixel and as large as 10 x 10 pixels over the training arena. In the bottom panel of Fig 7.2 an extra region of sensitivity was created in the place field. Such anomalies occurred in our system as the desired place fields were made larger. Fig 7.3 shows the weight matrices that produce the place fields shown in Fig 7.2. The left column of Fig 7.3 corresponds to the place cell on the top in Fig 7.2



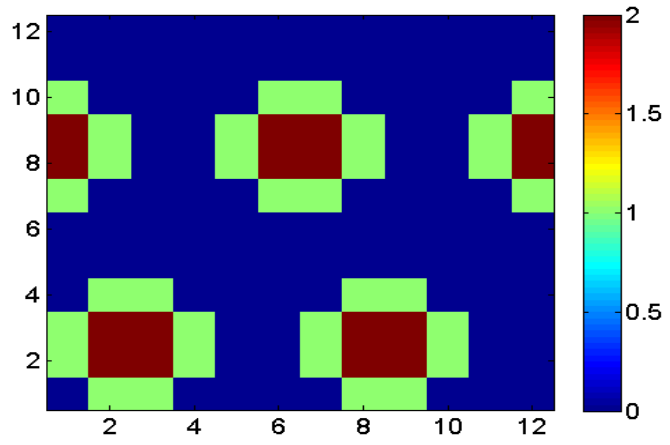


Fig 7.1. The response of the three grid cells used to create the place cell. We have three different spatial frequencies for the grid cells' responses.





Fig 7.2. Two examples for place fields that system was able to learn, the place field can be as small as one pixel or as big as 10 x 10 pixels.




Fig 7.3. Weight matrices for the two place cells shown in Fig 7.2

7.3.2 Grid Cells with Similar Connectivity and Different Velocity Drive

In this case, we reversed the variables from the previous case. We used only one connectivity pattern for all of the grid layers, however, the velocity signals had different weights on the activity of the grid cells. The grid cell layer consisted of three cells similar to the case shown in the previous section. We kept the setup of the experiments the same.




Fig 7.4. The grid cell response used to create the place cells in this experiment.




Fig 7.5. Two examples for place fields that system was able to learn, the place field can be as small as one pixel or as big as 10 x 10 pixels.




Fig 7.6. Weight matrices for the two place cells shown in Fig 7.5

Fig 7.4 shows the response of the grid cell we used in this experiment. Using three different initial phases (randomly chosen) we generated the grid cells activity for the experiment. Fig 7.5 shows two examples of place fields we were able to train in the system. We could train cells with place fields as small as one pixel and as large as 10 x 10 pixels over the training arena. As the place field of the place cells gets bigger, we noticed that an ambiguity in the place field started to emerge. Fig 7.6 shows the weight matrices that result in the place fields shown in Fig 7.5, the left column of Fig 7.6 corresponds to the place cell on the top in Fig 7.5.

7.4 Conclusions

In this chapter we presented a simulation of how the grid cells can be used to create place cells. From our experience, the construction of place cells from grid cells is computationally simple and does not require a lot of training; we were able to obtain stable place fields within 1000 iterations into the experiments. The fact that we were able to create place cells from grid cells implies that the formation of place cells in the brain could be explained based on the activity of the grid cells.

Chapter 8 : Conclusions and Future Directions

This dissertation was focused on the study of the neurobiological structures involved in performing spatial navigation in mammals. The animal of interest was the big brown bat, however, due to the lack of neurobiological data from that specific species, we adopted models from the rat literature with the assumption that both animals have similar brain structures. Over the course of the thesis, new neurobiological findings in the echolocating bat assured our initial assumptions regarding the similarities between the two species.

The neural structures involved in representing spatial layout of the environment in the brain of mammals can be broken down into the head direction cell system, the grid cell system and the place cell system. The head direction cells keep track of the orientation in the environment, the grid cells keep track of the displacement in the environment and the place cells keep track of the absolute location in the environment.

Our goal was to use our understanding of the neural structures to design neuromorphic VLSI circuits that mimic the computation of these circuits. We were able to design a neuromorphic head direction cell system. The system we developed was driven only by analog (asynchronous spikes) rotation velocity signals and was able to integrate the velocity signals into position, mimicking the biological head direction cell system. The neuromorphic VLSI head direction cell system we designed was noisy (both spatially and temporally) and similar to its biological counterpart, the position estimate would drift with time. To solve that problem, we combined information from a sonar system (sensory) with the head direction system

in a hybrid (hardware / software) system that was able to associate the location of cues in the environment and use that knowledge to correct for the drift in the head direction cell system estimate.

Following the head direction cell system we designed a grid cell system. To reduce the system complexity, we went from spike-based computation to a completely analog core with a spiking representation at the output that was generated by a comparator while scanning all of the neurons sequentially. The neuromorphic VLSI grid cell system successfully generated grid-like patterns of activity whose spatial frequency is controlled by parameters on the chip, however, the patterns could not be moved over the chip. Our built-in mechanism for controlled pattern movement suffered from big levels of mismatch on the chip which resulted in the grid patterns being very stable and not movable.

Since our aim was to understand the neural circuitry involved in spatial navigation we went ahead and moved our grid cells to the software domain to test our ability to create place cells based on grid cell activity. Our simulation showed that place cells can indeed be formed using as few as three grid cells. We tested two hypotheses for place cell formation from grid cells, and showed that both of them did produce place cells. With the formation of place cells, we have been able to model the most relevant parts of the neural circuits that, to our knowledge, contribute to the encoding of space in the brain of the rats.

For future directions we can suggest that the place cell system be moved from the software domain back to the hardware domain to have a completely hardware based system. Also we think that the system should be tested as a whole using a robotic

platform which will be challenging from the perspective of tuning all the pieces of hardware and have them work together at the same time. Our overall approach, for the head-direction system when applied to hardware is low power (internal to the chips) and would be suitable for micro-aerial vehicles, which is one of the next platforms on which our system will be tested.

Finally, our knowledge about the computational algorithms performed by the brain changes by the day with new stunning discoveries. To benefit from the biological superiority of computation, we must keep an eye out for discoveries coming from the neuroscience which might lead to revolutionary ways to our engineering techniques.

Bibliography

- [1] A. S. Etienne, "Mammalian Navigation, Neural Models and Biorobotics," *Connection Science*, vol. 10, pp. 271-289, 1998.
- [2] R. Chavarriaga, T. Strosslin, D. Sheynikhovich, and W. Gerstner, "A Computational Model Of Parallel Navigation Systems In Rodents," *Neuroinformatics*, vol. 3, pp. 223-241, 2005.
- [3] A. D. Redish, *Beyond The Cognitive Map: From Place Cells To Episodic Memory*. Cambridge: The MIT Press, 1999.
- [4] A. Arleo, "Spatial Learning And Navigation In Neuro-Mimetic Systems," PhD Dissertation/Thesis, Ecole Polytechnique Federal de Lausanne, 2000.
- [5] J. O'Keefe and L. Nadel, *The Hippocampus As A Cognitive Map*: Oxford University Press, 1978.
- [6] J. J. Leonard and H. F. Durrant-Whyte, "Simultaneous Map Building And Localization For An Autonomous Mobile Robot," in *Intelligent Robots and Systems '91. Intelligence for Mechanical Systems, Proceedings IROS '91. IEEE/RSJ International Workshop on*, 1991, pp. 1442-1447.
- [7] M. Muller and R. Wehner, "Path Integration In Desert Ants, *Cataglyphis Fortis*," *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, vol. 85, pp. 5287-5290, 1988.
- [8] K. v. Frisch, *The Dance Language And Orientation Of Bees*: Belknap Press, 1967.
- [9] F. L. W. Ratnieks. (2000) How Far Do Honey Bees Forage? *Bee Improvement*. 10-11.
- [10] M. J. Milford, *Robot Navigation From Nature* vol. 41: Springer, 2008.
- [11] J. O'Keefe and J. Dostrovsky, "The Hippocampus As A Spatial Map. Preliminary Evidence From Unit Activity In The Freely-Moving Rat," *Brain Research*, vol. 34, pp. 171-175, 1971.
- [12] J. S. Taube, R. U. Muller, and J. B. Ranck, Jr., "Head-Direction Cells Recorded From The Postsubiculum In Freely Moving Rats. I. Description And Quantitative Analysis," *The Journal Of Neuroscience*, vol. 10, pp. 420-435, 1990.
- [13] J. S. Taube, R. U. Muller, and J. B. Ranck, Jr., "Head-Direction Cells Recorded From The Postsubiculum In Freely Moving Rats. II. Effects Of

- Environmental Manipulations," *The Journal Of Neuroscience*, vol. 10, pp. 436-447, 1990.
- [14] P. E. Sharp, H. T. Blair, and J. Cho, "The Anatomical And Computational Basis Of The Rat Head-Direction Cell Signal," *Trends in neurosciences*, vol. 24, pp. 289-294, 2001.
- [15] R. G. Morris, "Spatial Localization Does Not Require The Presence Of Local Cues," *Learning And Motivation*, vol. 12, pp. 239-260, 1981.
- [16] J. O'Keefe, "Place units in the hippocampus of the freely moving rat," *Experimental neurology*, vol. 51, pp. 78-109, Apr 1976.
- [17] J. O'Keefe and M. L. Recce, "Phase Relationship Between Hippocampal Place Units And The EEG Theta Rhythm," *Hippocampus*, vol. 3, pp. 317-330, 1993.
- [18] B. J. Young, G. D. Fox, and H. Eichenbaum, "Correlates Of Hippocampal Complex-Spike Cell Activity In Rats Performing A Nonspatial Radial Maze Task," *The Journal Of Neuroscience*, vol. 14, pp. 6553-6563, 1994.
- [19] 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, pp. 801-806, 2005.
- [20] M. M. Yartsev, M. P. Witter, and N. Ulanovsky, "Grid cells without theta oscillations in the entorhinal cortex of bats," *Nature*, vol. 479, pp. 103-7, Nov 3 2011.
- [21] N. Ulanovsky and C. F. Moss, "Hippocampal Cellular And Network Activity In Freely Moving Echolocating Bats," *Nature Neuroscience*, vol. 10, pp. 224-233, 2007.
- [22] N. Ulanovsky and C. F. Moss, "What The Bat's Voice Tells The Bat's Brain," *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, vol. 105, pp. 8491-8498, 2008.
- [23] J. J. Knierim and B. L. McNaughton, "Hippocampal Place-Cell Firing During Movement In Three-Dimensional Space," *Journal Of Neurophysiology*, vol. 85, pp. 105-116, 2001.
- [24] N. Ulanovsky and C. F. Moss, "Dynamics of hippocampal spatial representation in echolocating bats," *Hippocampus*, vol. 21, pp. 150-61, Feb 2011.
- [25] R. A. Brooks, "Intelligence Without Representation," *Artificial Intelligence*, vol. 47, pp. 139-159, 1991.

- [26] D. Cliff, J. A. Meyer, and S. W. Wilson, "Computational Neuroethology: A Provisional Manifesto," in *The International Conference on Simulation of Adaptive Behavior From Animals to Animats*., Paris, France, 1991.
- [27] H. Moravec and A. Elfes, "High Resolution Maps From Wide Angle Sonar," presented at the IEEE International Conference on Robotics and Automation (ICRA'85), 1985.
- [28] A. Elfes. (1989) Using Occupancy Grids For Mobile Robot Perception And Navigation. 46-57.
- [29] B. Kuipers and Y. T. Byun, "A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations," *Journal of Robotics and Autonomous Systems*, vol. 8, pp. 47-63, 1991.
- [30] B. Kuipers and T. S. Levitt. (1988) Navigation and Mapping in Large-Scale Space. *AI Magazine*. 25-43.
- [31] S. Thrun and A. Bucken, "Integrating Grid-Based And Topological Maps For Mobile Robot Navigation," in *AAAI/IAAI*, 1996, pp. 944-950.
- [32] S. Thrun and A. Bucken, "Learning Maps For Indoor Mobile Robot Navigation," *Artificial Intelligence*, vol. 99, pp. 21-71, 1998.
- [33] E. C. Tolman, "Cognitive Maps In Rats And Men," *Psychological Review*, vol. 55, pp. 189-208, 1948.
- [34] R. G. Robertson, E. T. Rolls, P. Georges-Francois, and S. Panzeri, "Head Direction Cells In The Primate Pre-Subiculum," *Hippocampus*, vol. 9, pp. 206-219, 1999.
- [35] S. J. Mizumori and J. D. Williams, "Directionally Selective Mnemonic Properties Of Neurons In The Lateral Dorsal Nucleus Of The Thalamus Of Rats," *The Journal Of Neuroscience*, vol. 13, pp. 4015-4028, 1993.
- [36] S. I. Wiener, "Spatial And Behavioral Correlates Of Striatal Neurons In Rats Performing A Self-Initiated Navigation Task," *The Journal Of Neuroscience*, vol. 13, pp. 3802-3817, 1993.
- [37] J. Cho and P. E. Sharp, "Head Direction, Place, And Movement Correlates For Cells In The Rat Retrosplenial Cortex," *Behavioral Neuroscience*, vol. 115, pp. 3-25, 2001.
- [38] J. S. Taube, "Head Direction Cells Recorded In The Anterior Thalamic Nuclei Of Freely Moving Rats," *The Journal Of Neuroscience*, vol. 15, pp. 70-86, 1995.

- [39] H. T. Blair, J. Cho, and P. E. Sharp, "Role Of The Lateral Mammillary Nucleus In The Rat Head Direction Circuit: A Combined Single Unit Recording And Lesion Study," *Neuron*, vol. 21, pp. 1387-1397, 1998.
- [40] H. T. Blair and P. E. Sharp, "Functional Organization Of The Rat Head Direction Circuit," in *The Neural Basis of Navigation*, P. E. Sharp, Ed., ed: Kluwer Academic Publishers, 2002.
- [41] J. S. Taube and J. P. Bassett, "Persistent Neural Activity In Head Direction Cells," *Cerebral Cortex*, vol. 13, pp. 1162-1172, 2003.
- [42] J. S. Taube, "The Head Direction Signal: Origins And Sensory-Motor Integration," *Annual Review Of Neuroscience*, vol. 30, pp. 181-207, 2007.
- [43] B. L. McNaughton, L. L. Chen, and E. J. Markus, "'Dead Reckoning,' Landmark Learning, And The Sense Of Direction: A Neurophysiological And Computational Hypothesis," *Journal of Cognitive Neuroscience*, vol. 3, pp. 190-203, 1991.
- [44] R. W. Stackman and J. S. Taube, "Firing Properties Of Head Direction Cells In The Rat Anterior Thalamic Nucleus: Dependence On Vestibular Input," *The Journal Of Neuroscience*, vol. 17, pp. 4349-4358, 1997.
- [45] 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," *Journal of Neuroscience*, vol. 15, pp. 6260-6270, 1995.
- [46] J. J. Knierim, H. S. Kudrimoti, and B. L. McNaughton, "Place Cells, Head Direction Cells, And The Learning Of Landmark Stability," *The Journal Of Neuroscience*, vol. 15, pp. 1648-1659, 1995.
- [47] E. I. Moser, E. Kropff, and M. B. Moser, "Place Cells, Grid Cells, And The Brain's Spatial Representation System," *Annual Review Of Neuroscience*, vol. 31, pp. 69-89, 2008.
- [48] M. W. Jung and B. L. McNaughton, "Spatial Selectivity Of Unit Activity In The Hippocampal Granular Layer," *Hippocampus*, vol. 3, pp. 165-182, 1993.
- [49] G. J. Quirk, R. U. Muller, J. L. Kubie, and J. B. Ranck, Jr., "The Positional Firing Properties Of Medial Entorhinal Neurons: Description And Comparison With Hippocampal Place Cells," *The Journal Of Neuroscience*, vol. 12, pp. 1945-1963, 1992.
- [50] P. E. Sharp and C. Green, "Spatial Correlates Of Firing Patterns Of Single Cells In The Subiculum Of The Freely Moving Rat," *The Journal Of Neuroscience*, vol. 14, pp. 2339-2356, 1994.

- [51] D. G. Amaral and M. P. Witter, "The Three-Dimensional Organization Of The Hippocampal Formation: A Review Of Anatomical Data," *Neuroscience*, vol. 31, pp. 571-591, 1989.
- [52] D. G. Amaral, "Emerging Principles Of Intrinsic Hippocampal Organization," *Current opinion in neurobiology*, vol. 3, pp. 225-229, 1993.
- [53] E. T. Rolls, Y. Miyashita, P. M. Cahusac, R. P. Kesner, H. Niki, J. D. Feigenbaum, *et al.*, "Hippocampal Neurons In The Monkey With Activity Related To The Place In Which A Stimulus Is Shown," *The Journal Of Neuroscience*, vol. 9, pp. 1835-1845, 1989.
- [54] R. Muller, "A Quarter Of A Century Of Place Cells," *Neuron*, vol. 17, pp. 813-822, 1996.
- [55] N. Burgess, M. Reece, and J. O'Keefe, "A Model Of Hippocampal Function," *Neural Networks*, vol. 7, pp. 1065-1081, 1994.
- [56] M. A. Wilson and B. L. McNaughton, "Dynamics Of The Hippocampal Ensemble Code For Space," *Science Mag.*, vol. 261, pp. 1055-1058, 1993.
- [57] M. R. Mehta, M. C. Quirk, and M. A. Wilson, "Experience-Dependent Asymmetric Shape Of Hippocampal Receptive Fields," *Neuron*, vol. 25, pp. 707-715, 2000.
- [58] C. Pavlides and J. Winson, "Influences Of Hippocampal Place Cell Firing In The Awake State On The Activity Of These Cells During Subsequent Sleep Episodes," *The Journal Of Neuroscience*, vol. 9, pp. 2907-2918, 1989.
- [59] P. E. Sharp, "Subicular Cells Generate Similar Spatial Firing Patterns In Two Geometrically And Visually Distinctive Environments : Comparison With Hippocampal Place Cells," *Behavioral Brain Research*, vol. 85, pp. 71-92, 1997.
- [60] J. O'Keefe and N. Burgess, "Geometric determinants of the place fields of hippocampal neurons," *Nature*, vol. 381, pp. 425-428, May 30 1996.
- [61] G. Buzsaki, "Feed-Forward Inhibition In The Hippocampal Formation," *Progress in Neurobiology*, vol. 22, pp. 131-153, 1984.
- [62] R. G. Morris and U. Frey, "Hippocampal Synaptic Plasticity: Role In Spatial Learning Or The Automatic Recording Of Attended Experience?," *Philosophical Transactions Of The Royal Society Of London, Series B. Biological Sciences*, vol. 352, pp. 1489-1503, 1997.
- [63] W. B. Levy and O. Steward, "Temporal Contiguity Requirements For Long-Term Associative Potentiation/Depression In The Hippocampus," *Neuroscience*, vol. 8, pp. 791-797, 1983.

- [64] L. F. Abbott and K. I. Blum, "Functional Significance Of Long-Term Potentiation For Sequence Learning And Prediction," *Cerebral cortex*, vol. 6, pp. 406-416, May-Jun 1996.
- [65] H. Eichenbaum, P. Dudchenko, E. Wood, M. Shapiro, and H. Tanila, "The Hippocampus, Memory, And Place Cells: Is It Spatial Memory Or A Memory Space?," *Neuron*, vol. 23, pp. 209-226, 1999.
- [66] J. O'Keefe, "Do hippocampal pyramidal cells signal non-spatial as well as spatial information?," *Hippocampus*, vol. 9, pp. 352-64, 1999.
- [67] 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," *Advances In Neural Information Processing Systems*, vol. 7, pp. 173-180, 1995.
- [68] A. D. 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, pp. 671-685, 1996.
- [69] J. S. Taube, J. P. Goodridge, E. J. Golob, P. A. Dudchenko, and R. W. Stackman, "Processing The Head Direction Cell Signal: A Review And Commentary," *Brain Research Bulletin*, vol. 40, pp. 477-84, 1996.
- [70] J. P. Goodridge and D. S. Touretzky, "Modeling Attractor Deformation In The Rodent Head-Direction System," *Journal of Neurophysiology*, vol. 83, pp. 3402-3410, 2000.
- [71] M. C. Fuhs and D. S. Touretzky, "A Spin Glass Model Of Path Integration In Rat Medial Entorhinal Cortex," *The Journal Of Neuroscience*, vol. 26, pp. 4266-4276, 2006.
- [72] 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, pp. 663-678, 2006.
- [73] N. Burgess, C. Barry, and J. O'Keefe, "An Oscillatory Interference Model Of Grid Cell Firing," *Hippocampus*, vol. 17, pp. 801-812, 2007.
- [74] M. E. Hasselmo, "Grid cell mechanisms and function: contributions of entorhinal persistent spiking and phase resetting," *Hippocampus*, vol. 18, pp. 1213-29, 2008.
- [75] E. T. Rolls, S. M. Stringer, and T. Elliot, "Entorhinal Cortex Grid Cells Can Map To Hippocampal Place Cells By Competitive Learning," *Network: Computation in Neural Systems*, vol. 17, pp. 447-465, 2006.

- [76] A. Arleo and L. Rondi-Reig, "Multimodal Sensory Integration And Concurrent Navigation Strategies For Spatial Cognition In Real And Artificial Organisms," *Journal of Integrative Neuroscience*, vol. 6, pp. 327-366, 2007.
- [77] U. M. Erdem and M. Hasselmo, "A goal-directed spatial navigation model using forward trajectory planning based on grid cells," *Eur J Neurosci*, vol. 35, pp. 916-31, Mar 2012.
- [78] J. S. Taube, "Sensory Determinants Of Head Direction Cell Activity," in *The Neural Basis of Navigation: Evidence from Single Cell Recording*, P. E. Sharp, Ed., ed: Kluwer Academic Publishers, 2001.
- [79] X. Xie, R. H. Hahnloser, and H. S. Seung, "Double-Ring Network Model Of The Head-Direction System," *Physical Review E, Statistical, Nonlinear, And Soft Matter Physics*, vol. 66, 2002.
- [80] K. Zhang, "Representation Of Spatial Orientation By The Intrinsic Dynamics Of The Head-Direction Cell Ensemble: A Theory," *The Journal Of Neuroscience*, vol. 16, pp. 2112-2126, 1996.
- [81] 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 : the official journal of the Society for Neuroscience*, vol. 25, pp. 1002-1014, 2005.
- [82] T. Degris, L. Lachèze, C. Boucheny, and A. Arleo, "A spiking neuron model of head-direction cells for robot orientation," in *The Eighth International Conference on Simulation of Adaptive Behavior, From Animals to Animats*, 2004, pp. 255-263.
- [83] P. Stratton, M. Milford, J. Wiles, and G. Wyeth, "Automatic Calibration of a Spiking Head-Direction Network for Representing Robot Orientation," presented at the Australian Conference on Robotics and Automation, Sydney, Australia, 2009.
- [84] C. Laing and C. Chow, "Stationary bumps in networks of spiking neurons.," *Neural Comput*, vol. 13, pp. 1473-94, Jul 2001.
- [85] T. P. Vogels, K. Rajan, and L. F. Abbott, "Neural network dynamics," *Annu Rev Neurosci*, vol. 28, pp. 357-76, 2005.
- [86] T. P. Trappenberg and D. I. Standage, "Multi-packet regions in stabilized continuous attractor networks," *Neurocomputing*, vol. 65-66, pp. 617-622, 2004.
- [87] M. A. Basso and R. H. Wurtz, "Neuronal activity in substantia nigra pars reticulata during target selection," *Journal of Neuroscience*, vol. 22, pp. 1883-1894, 2002.

- [88] G. Indiveri, E. Chicca, and R. J. Douglas, "A VLSI Array Of Low-Power Spiking Neurons And Bistable Synapses With Spike-Timing Dependent Plasticity," *Neural Networks, IEEE Transactions on*, vol. 17, pp. 211-221, 2006.
- [89] K. A. Boahen, "Point-To-Point Connectivity Between Neuromorphic Chips Using Address Events," *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on*, vol. 47, pp. 416-434, 2000.
- [90] J. V. Arthur and K. Boahen, "Silicon Neurons That Inhibit To Synchronize," in *International Symposium on Circuits and Systems (ISCAS 2006)*, Island of Kos, Greece, 2006.
- [91] A. Renart, P. C. Song, and X. J. Wang, "Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks," *Neuron*, vol. 38, pp. 473-485, May 2003.
- [92] D. P. Goodridge JP, Worboys KA, Golob EJ, Taube JS., "Cue control and head direction cells.," *Behav Neurosci*, vol. 112, pp. 741-69, 1998.
- [93] R. H. Hahnloser, "Emergence Of Neural Integration In The Head-Direction System By Visual Supervision," *Neuroscience*, vol. 120, pp. 877-891, 2003.
- [94] K. A. S. Hertz J.A., Palmer R.G., *Introduction To The Theory Of Neural Computation, Volume I*: Westview Press, 1991.
- [95] 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.
- [96] F. Sargolini, M. Fyhn, T. Hafting, B. L. McNaughton, M. P. Witter, M.-B. Moser, *et al.*, "Conjunctive Representation of Position, Direction, and Velocity in Entorhinal Cortex," *Science*, vol. 312, pp. 758-762, May 5, 2006 2006.
- [97] T. M. Massoud and T. K. Horiuchi, "A Neuromorphic VLSI Head Direction Cell System," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 58, pp. 150 - 163, 2011.
- [98] T. M. Massoud and T. K. Horiuchi, "Online correction of orientation estimates using spatial memory in a neuromorphic head direction system," presented at the International Symposium for Circuits and Systems (ISCAS 2011), Rio, Brazil, 2011.
- [99] A. G. Andreou and K. A. Boahen, "Translinear circuits in subthreshold MOS," *Analog Integrated Circuits and Signal Processing*, vol. 9, pp. 141-166, 1996.

- [100] Y. Burak and I. R. Fiete, "Accurate path integration in continuous attractor network models of grid cells," *PLoS Comput Biol*, vol. 5, p. e1000291, Feb 2009.

List of Publications

Journal Publications

Massoud T.M., and Horiuchi T.K., "A Neuromorphic VLSI Head Direction Cell System". *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 58, Issue 1, 2011, pp. 150 - 163. (DOI: 10.1109/TCSI.2010.2055310).

Articles

Massoud T.M., and Horiuchi T.K., "Learning to correct orientation estimates using spatial memory", *The Neuromorphic Engineer Newsletter*, September 2011. (DOI: 10.2417/1201108.003828).

Conference Publications

Massoud, T. M., and Horiuchi, T. K., "A Neuromorphic VLSI Grid Cell System". *IEEE International Symposium on Circuits and Systems (Systems (ISCAS 2012)*, May 2012, Seoul, Korea.

Massoud, T. M., and Horiuchi, T. K., "Online Correction of Orientation Estimates Using Spatial Memory in a Neuromorphic Head Direction System". *IEEE International Symposium on Circuits and Systems (ISCAS 2011)*, May 2011, Rio, Brazil. pp. 2429 - 2432, 2011. (DOI: 10.1109/ISCAS.2011.5938094)

Massoud, T. M., and Horiuchi, T. K., "A Neuromorphic Head Direction Cell System". *IEEE International Symposium on Circuits and Systems (ISCAS 2009)*, May 2009, , Taipei, Taiwan. pp. 565 - 568, 2009. (DOI: 10.1109/ISCAS.2009.5117811)

Horiuchi, T. K., Bansal, C., and **Massoud, T. M.**, "Binaural Intensity Comparison in the Echolocating Bat Using Synaptic Conductance". *IEEE International Symposium on Circuits and Systems (ISCAS 2009)*, May 2009, Taipei, Taiwan. pp. 2153 - 2156, 2009. (DOI: 10.1109/ISCAS.2009.5118222)