

RICE UNIVERSITY

**The Effect of Synaptic Plasticity on Spatial Representation
and Navigation**

by

Kathryn Ruth Ward

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Master of Arts

APPROVED, THESIS COMMITTEE:

Steven Cox, Chairman
Professor of Computational and Applied Mathematics

Danny Sorensen
Professor of Computational and Applied Mathematics

Mark Embree
Associate Professor of Computational and Applied Mathematics

Harel Shouval
Assistant Professor of Neurobiology and Anatomy
University of Texas Health Science Center at Houston

HOUSTON, TEXAS

APRIL, 2009

Abstract

The Effect of Synaptic Plasticity on Spatial Representation and Navigation

by

Kathryn Ruth Ward

Synaptic plasticity, or the change in weight of the connections between cells, is a key mechanism underlying the brain's spatial representation and navigation functions. Experimentalists have shown that grid cells in the medial entorhinal cortex fire in hexagonal patterns within an environment, or set of visual cues. Grid cells provide the input for place cells, which fire primarily at one location in the environment and are found in the hippocampus, a region essential for both learning and memory. I have built a computational model to examine how synaptic plasticity affects the interactions among grid cells and place cells. This work demonstrates that a rate-based plasticity model drives the weights from grid cells to place cells to such a distribution that place cells form single firing fields. Furthermore, a spike-timing-dependent plasticity model applied to the connections among place cells causes place fields to shift backward as observed experimentally.

Acknowledgements

I would first like to thank the many professors who have invested in me over the last few years. I would particularly like to thank my advisor, Dr. Cox, for his guidance and encouragement. I would also like to thank Dr. Sorensen and Dr. Hewitt for their assistance in helping me write this thesis and prepare for my defense, Dr. Shouval for the guidance and instruction he has given me this semester, and Dr. Embree for the support and contagious enthusiasm he has shown.

I would like to give a special thanks to Dr. Knierim and Dr. Savelli, who have given me much guidance and support in helping me to understand the biology of the system. Thanks also to Horatio Voicu and Katherine Johnston, who both contributed largely to the initial stages of this work.

Finally, I would like to thank my family and friends whose loving support has lifted me to where am I today, and I thank my Father for the talents He has given me and for the peace and strength He provides for me.

Contents

Abstract	ii
Acknowledgements	iii
List of Figures	vii
List of Tables	ix
1 Introduction	1
1.1 Spatial Representation	3
1.1.1 Neuronal Network for Spatial Representation	3
1.1.2 The Double Rotation Experiment	5
1.2 Synaptic Plasticity	8
1.2.1 Postsynaptic-Gated Rate-Based Plasticity	9
1.2.2 Spike-Timing-Dependent Synaptic Plasticity	11
2 Methods	15
2.1 Architecture of the Model	16

2.2	Single-Cell Dynamics	18
2.2.1	Grid Cell Model	19
2.2.2	CA3 Cell Model: Integrate and Fire	23
2.3	Network Dynamics: Synaptic Plasticity	31
2.3.1	Implementation of Postynaptic-Gated Rate-Based Plasticity	31
2.3.2	Implementation of Spike-Timing-Dependent Plasticity	32
3	Results	37
3.1	Formation of Place Fields	38
3.1.1	Learning Rate	40
3.1.2	Other Factors Affecting Place Field Formation	43
3.2	Backward Shift	45
3.2.1	Simple Example	46
3.2.2	Backward Shift in the Full Model	47
4	Conclusion	52
4.1	Applications	53
4.1.1	Mismatch Stage of the Double Rotation Experiment	53
4.1.2	Various Other Applications	55
4.2	Future Work	57
4.2.1	Improvements to the Model's Speed and Capability	58
4.2.2	Improvements to the Model's Accuracy	59

4.3 Summary 60

Bibliography **62**

List of Figures

1.1	Firing fields of grid cells and place cells	3
1.2	Schematic of the hippocampus and entorhinal cortex	6
1.3	Connections among regions in the network	6
1.4	Experimental results supporting rate-based plasticity	11
1.5	Experimental results supporting STDP	13
2.1	Architecture of the model	16
2.2	Graphical demonstration of grid cell parameters	21
2.3	Comparison of the effect of γ values	24
2.4	Circuit diagram of an integrate-and-fire neuron	25
2.5	Weight change due to STDP	33
3.1	Place field formation	41
3.2	Effect of the learning rate	44
3.3	Simple example of the backward shift	48
3.4	Backward shift in the full network model	51

4.1	Cell response to the double rotation of cues	56
4.2	Speed-up in PETSc	59

List of Tables

2.1	Grid Cell Parameters	21
2.2	Integrate and Fire Parameters	29
2.3	Weight Matrix Parameters	30
2.4	Rate-Based Plasticity Parameters	33
2.5	STDP Parameters	36

Chapter 1

Introduction

Synaptic plasticity, or the process in which the strength of the connection between neurons changes in time, enables the brain to both learn and store memories. Everyday functions, such as remembering the events of the day or learning to navigate from work to home, depend on this process. Because it plays such a key role in many functions, a reliable computational model of synaptic plasticity would enable researchers to gain better insight into the mechanisms underlying these functions. In this work I have built a computational model of synaptic plasticity and have examined its effect on the brain's spatial representation and navigation functions.

The brain provides spatial information through the activity of grid cells and place cells, which have strong spatial tuning. In 2005 Moser and his lab discovered that grid cells primarily spike, or send current to other cells, in hexagonal patterns within an environment (Hafting et al. [12]). O'Keefe and Nadel discovered place cells in 1978.

Place cells, whose primary source of input is the network of grid cells, primarily spike at one location within an environment (O'Keefe and Nadel [22]). In the spatial representation setting, an environment refers to the set of visual cues that surround a given organism. Figure 1.1 shows the firing fields of both grid cells and place cells.

To ensure that my computational model can predict biological results, I have simulated a spatial representation experiment currently being performed by James Knierim and his lab at the University of Texas Medical School. In this experiment a rat runs around a circular track while electrodes lowered into its brain record the spiking behavior of the cells (Lee et al. [19]). I have worked closely with Knierim to ensure that my model is consistent with the biology of the system. In turn, the computational model I have developed provides insight into the mechanisms underlying observed results from experiments such as his.

This work focuses on two phenomena that occur as the rat circles the track. First, each place cell integrates input from grid cells, which spike at several locations along the track, to form a place field at one location. Second, many place fields shift backward, in the opposite direction that the rat is moving. This thesis demonstrates that rate-based synaptic plasticity applied to the connections from grid cells to place cells can enable place fields to form. Furthermore, spike-timing-dependent plasticity (STDP) applied to the connections among place cells can cause place fields to shift backward.




Figure 1.1: Firing fields of grid cells and place cells. The experimental data presented here shows the firing fields of a grid cell (a) and a place cell (b) as a rat explores a rectangular enclosure. The gray lines show the trajectory of the rat. (a) Each red dot represents a location at which the grid cell has spiked. The firing locations of grid cells form hexagonal patterns. (b) Each red dot represents a location at which the place cell has spiked. Place cells spike at one location within a given environment. Taken from Witter and Moser [29].

1.1 Spatial Representation

After moving to a new city, one must consult a map to navigate from one place to another. In time, however, the brain builds an internal map, enabling one to navigate without external aid. Grid cells and place cells belong to the network that performs this function. In this section I describe the network's anatomy, focusing on the characteristics that lend it to spatial representation. I then describe the double rotation experiment the model simulates. This experiment provides further insight into the interactions among the cells in the network.

1.1.1 Neuronal Network for Spatial Representation

Place cells are located in the hippocampus, an essential region for learning and memory (O'Keefe and Nadel [22]). Without the hippocampus, the brain would be unable to store memories for longer than a few minutes. One famous example of this memory

loss is a man known as H.M. In 1953, H.M. sustained damage to his hippocampus during brain surgery. From that day on, he was unable to retain new memories from day to day, though he could clearly recall events from his childhood (Cohen and Eichenbaum [7]). Patients such as H.M. have shown that one vital function of the hippocampus is to act as a relay between short-term and long-term memory.

The hippocampus is also an important region for the generation of neurological diseases such as Alzheimer's disease and epilepsy. The early and extensive changes characteristic of Alzheimer's disease occur in the hippocampus. During an epileptic seizure, cells spike synchronously at a much higher rate than normally, and this synchronous over-activity among cells often originates in the hippocampus (Traub et al. [28]).

Figure 1.1.1 shows a portion of the rat brain. The regions important for spatial representation are the entorhinal cortex (EC) and the hippocampal formation (HF). The EC contains two subregions, the lateral entorhinal cortex (LEC) and the medial entorhinal cortex (MEC). Cells in the LEC have weak spatial tuning, showing little preference for one location in an environment over another. The MEC, however, contains grid cells, which have strong spatial tuning. The hippocampus is composed of the dentate gyrus (DG), CA1, and CA3 regions. Place cells are contained in the CA1 and CA3 regions of the hippocampus.

The regions form several connections onto each other, as shown in Figure 1.1.1. Two pathways exist from the EC to the CA3, one direct pathway and one pathway

that passes through the DG (Traub et al. [28]). The direct pathway provides the majority of input to the CA3 because cells in the DG have a low spike rate (Knierim, personal communication, March 3, 2009). The CA3 then projects onto itself and onto the CA1, which also receives input from the EC. The CA1 projects back onto the EC, providing the output of the hippocampus (Traub et al. [28]).

The CA3 subregion of the hippocampus is particularly important because of its recurrent collateral connections. These connections are strong enough that a burst of spikes from a single CA3 cell can evoke a burst of spikes in another CA3 cell (Traub et al. [28]). These strong connections cause the CA3 to have a low threshold for the generation of an epileptic seizure. The recurrent collateral connections also affect place cells in the CA3, as demonstrated in the double rotation experiment (Lee et al. [19]).

I model the behavior of place cells in the CA3 region of the hippocampus during the double rotation experiment, examining their interactions with other cells and their dependence on synaptic plasticity. Because of the importance of synaptic plasticity among cells in the CA3 region, the understanding gained from the model has implications for many important areas in addition to spatial representation and navigation.

1.1.2 The Double Rotation Experiment

Certain aspects of the network described in the previous section are not yet fully understood. To further investigate the interactions among cells in the network, the

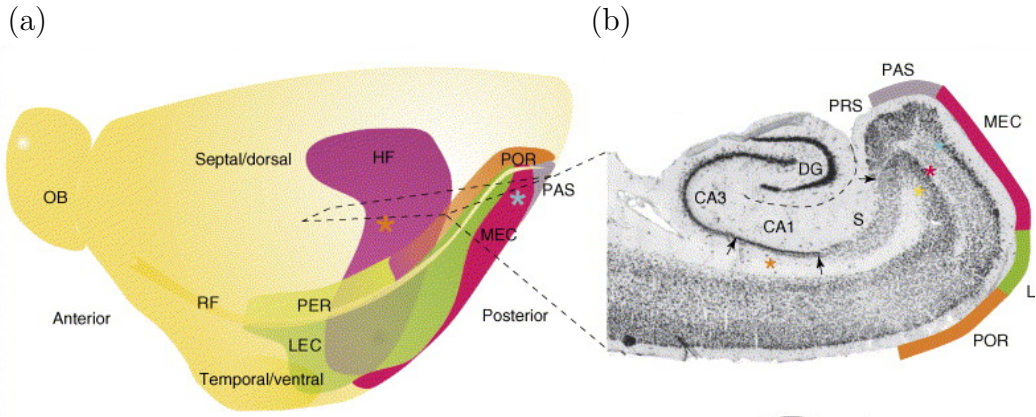


Figure 1.2: Schematic of the hippocampus and entorhinal cortex. (a) This schematic shows the posterolateral view of the left hemisphere of a rat brain. Grid cells are contained in the medial entorhinal cortex (MEC). Both regions of the entorhinal cortex, the MEC and the lateral entorhinal cortex (LEC), drive cells in the hippocampal formation (HF). (b) This schematic shows a horizontal section of the rat brain. Place cells are contained in the CA1 and CA3 subregions of the hippocampus. Taken from Witter and Moser [29].




Figure 1.3: Connections among the network regions in the network. The entorhinal cortex, which is divided into the MEC and LEC subregions, connects onto the dentate gyrus (DG), CA1, and CA3 subregions of the hippocampus. The DG connects onto itself and onto the CA3, and the CA3 connects onto itself and the CA1. The CA1 connects back onto the entorhinal cortex, providing the output of the network. The other regions shown are the presubiculum (PrS), parasubiculum (PaS), and subiculum (S). These regions do not affect the firing fields of place cells in the hippocampus. Taken from Amaral and Witter [1].

Knierim lab has performed the double rotation experiment in which a rat runs clockwise around a circular track with various local and distal cues. The radius of the track is 33 cm, and the rat takes an average of 18 seconds to complete one lap. Thus the rat's average velocity as it circles the track is 11.5 cm/sec.

The experiment is divided into two stages, the learning stage and the mismatch stage. During the learning stage, the Knierim lab records the spiking behavior of grid cells and place cells as the rat initially circles the track. After the rat has learned the cues over a period of days, the lab performs the mismatch stage, in which they rotate the distal cues clockwise and the local cues counter-clockwise. They then record the response of grid cells and place cells to this change of environment.

The computational model of this work simulates two phenomena that occur during the learning stage of the experiment. First, place cells form single firing fields as the rat initially explores the track. Second, once the place fields have formed, they shift backward, in the counterclockwise direction (Yu et al. [31]). Synaptic plasticity accounts for both phenomena in the model.

As Blum and Abbott have proposed, the backward shift of place fields may have important implications for sequence learning and navigation. Place fields shift backward in situations where a rat repeatedly follows a particular route, such as a circular track. Because of their spatial tuning, each place cell in the CA3 primarily spikes at one location along the route. When the firing fields shift backward, the place cells spike earlier along the route, enabling the rat to anticipate its next step (Blum and

Abbot [5]).

The double rotation experiment is one of the few experiments that examine the difference between place cells in the CA1 and in the CA3. During the learning stage, most place cells in the CA3 shift backward, while place cells in the CA1 do not. The differences are even more apparent in the mismatch stage of the experiment. Grid cells in the MEC provide the majority of spatial information to the hippocampus, and most place cells in the CA1 adjust their firing fields when the cues are rotated according to the response of grid cells. Place cells in the CA3, however, do not respond according to the grid cell response, implying that some unknown source or mechanism controls place cells in the CA3 when the environment changes. I describe the results from this stage more fully in Section 4.1.1.

1.2 Synaptic Plasticity

Synaptic plasticity is necessary for any function concerning learning and memory, including spatial representation and navigation. Experiments such as the Morris water maze have shown that rats lose the ability to navigate a given route when synaptic plasticity is prevented from occurring (Morris [21]). With respect to the double rotation experiment, Ekstrom has shown that the backward shift of place fields only occurs if synaptic plasticity is present (Ekstrom et al. [8]).

Synaptic plasticity has impacted the study of neuroscience since Donald Hebb first proposed his plasticity theory in 1949. This theory, known as Hebbian learning,

asserts that if two neurons repeatedly spike at the same time, the weight between them increases (Hebb [14]). Although the basic Hebbian learning model fails to explain many biological phenomena, it has become the cornerstone for the numerous models of synaptic plasticity that now exist. In this work I demonstrate that two types of Hebbian plasticity, rate-based plasticity and spike-timing-dependent plasticity, can account for the formation and backward shift of place fields.

1.2.1 Postsynaptic-Gated Rate-Based Plasticity

The first phenomenon the model simulates is that place cells driven by input from grid cells form single firing fields. Researchers such as Solstad and Moser in 2006 and Franzius in 2007 have shown that there exists a set of weights specifying the strength of connections between neurons such that the current from grid cells in the MEC causes place cells to form single firing fields (Solstad and Moser[26] and Franzius [9]). In both cases, however, the weights were calculated beforehand, and plasticity was not implemented.

Francesco Savelli, a member of the Knierim lab, has shown that, given small initial weights, synaptic plasticity can drive the weights to a stable distribution in which grid cells beget place cells (Savelli and Knierim [24]). However, he models place cells in the CA1 region of the hippocampus, a region that contains no collateral connections among its cells. As a result, the backward shift in the place cells' firing fields cannot be observed. I have implemented the rate-based plasticity model examined by Savelli

and confirmed that this type of plasticity can enable place fields to form in the CA3.

The rate-based plasticity model is based on experimental data, such as that shown in Figure 1.2.1. The connection between two cells is called a synapse, where spikes from the presynaptic cell affect the postsynaptic cell. In this experiment Heynen et al. stimulated cells in the hippocampus to spike at a given rate and measured the average response of postsynaptic cells due to each presynaptic spike. When they stimulated the presynaptic cells to spike at a rate of 10 Hz, the postsynaptic responses due to individual spikes decreased. In other words the weights from presynaptic cells to postsynaptic cells decreased, a process known as long-term depression (LTD). When they stimulated the presynaptic cells to spike at a rate of 100 Hz, the weights increased, a process known as long-term potentiation (LTP).

In their book *Spiking Neuron Models*, Gerstner and Kistler describe the basic rate-based plasticity model, which has the form

$$\Delta W_{ij}(t) = k(R_j(t) - \bar{R})R_i(t), \quad (1.2.1)$$

where W_{ij} is the weight from presynaptic cell j to postsynaptic cell i , R_j and R_i are the instantaneous spike rates of the presynaptic and postsynaptic cells, respectively, \bar{R} is the threshold rate, and k is a parameter that specifies the learning rate. For this model, the weight only changes if the postsynaptic cell has a nonzero spike rate, and the presynaptic spike rate determines the direction of weight change (Gerstner and Kistler [11], pp. 356-362). In Section 2.3.1, I provide a detailed explanation of the implementation of this plasticity model.




Figure 1.4: Experimental results supporting rate-based plasticity. (a) Cells in the hippocampus are stimulated to spike at a given rate. (b) When presynaptic cells spike at a rate of 10 Hz, the weights decrease. When the presynaptic cells spike at a rate of 100 Hz, the weights increase, recovering to a greater value than the baseline weight. Taken from Heynen et al. [15].

1.2.2 Spike-Timing-Dependent Synaptic Plasticity

As the rat circles the track in the double rotation experiment, the firing fields of place cells in the CA3 shift backward due to synaptic plasticity at work among the region's recurrent collateral connections. The backward shift implies that synaptic plasticity has driven the recurrent CA3 connections to a feedforward state in which a place cell that spikes at one location is not only driven by grid cells, but is also driven by place cells that spike just before it on the track. To simulate this effect, I have implemented a spike-timing-dependent plasticity (STDP) model because of its effectiveness in driving a network to a feedforward state.

True to its name, STDP models depend on the exact timing of individual pairs of pre- and postsynaptic spikes rather than spike rates. Researchers such as Markram in 1997 and Bi and Poo in 1998 have shown that this type of plasticity is present in certain biological settings (Markram et al. spikes before the presynaptic cell, LTD

occurs. The magnitude of LTP or LTD is determined by the exact time difference between the two spikes.

A variety of STDP models have been developed that differ both in their assumptions and resulting weight distributions. One important distinction among STDP models is the way in which each spike pair contributes to the weight change. Song and Abbott developed an STDP model in 2000 in which each spike pair contributes equally to the weight change (Song and Abbott [27]). In 2002 Froemke and Dan claimed that this model fails in high-spiking regimes, which occur naturally in the brain. They developed a model in which the first spike pair within a burst of spikes dominates over other pairs in determining the weight change (Froemke and Dan [10]). Figure 1.2.2 shows that the Froemke and Dan model is more accurate than the Song and Abbott model in bursting regimes. I implement the Froemke and Dan model of STDP for this work.

The Knierim lab has shown that both the STDP model developed by Song and Abbott and a biophysical synaptic plasticity model developed by Shouval cause a backward shift in a simplified setting. They simulated a network of 1000 input place cells each connected to a single output place cell. They modelled each input place cell to have an identical Gaussian place field, and they distributed the firing fields evenly around the track. They then showed that the firing field of the output place cell shifted backward due to synaptic plasticity applied to the weights from the input place cells (Yu et al. [31] and Yu et al. [32]).




Figure 1.5: Experimental results supporting STDP. (a) The circles and triangles represent the percentage weight change due to an individual spike pair. The x-axis represents Δt , the difference between the postsynaptic spike time and the presynaptic spike time, and the y-axis represents the percentage of weight change. Taken from Bi and Poo [4]. Plots (b) and (c) compare the accuracy of the Song and Abbot model to the Froemke and Dan model in predicting the weight change induced by natural spike-train segments. (b) For each spike pair, the predicted percentage weight change from the Song and Abbot STDP model is plotted against the measured percentage weight change from experimental data. Ideally, all circles would lie on the diagonal. (c) For each spike pair, the predicted weight change from the Froemke and Dan model is plotted against the measured percentage weight change from experimental data. Taken from Froemke and Dan [10].

I have built on the ideas used by Shouval and Knierim to show a similar backward shift in a more biologically realistic setting. In my model, each place cell receives input from both excitatory and inhibitory cells in the CA3 as well as from grid cells in the MEC. The place cells form firing fields due to rate-based plasticity applied to the connections from a random group of grid cells. The resulting place fields vary in strength and are not evenly distributed around the track, making it more difficult to demonstrate the backward shift. I then show that the STDP model developed by Froemke and Dan applied to the random connections among CA3 cells causes a backward shift in the place fields. I provide a detailed description of the network I have simulated in Section 2.1.

Chapter 2

Methods

Due to the complexity of the brain, numerous simplifications and assumptions must be made to build a computational model of any brain function. The brain's spatial representation function is no exception. The network that performs this function spans several regions, and the spiking behavior of cells within each region is not yet fully understood. Due to these unknown factors and the computational cost that accompanies incorporating each region, I have approximated the architecture of the network in the brain with the simplified architecture described in Section 2.1.

After setting the network architecture for the model, I simulate the dynamics of the system on two levels. First, I simulate the single-cell dynamics of each individual cell in the network. The single-cell dynamics are fairly well understood, but modelling these dynamics can be very computationally expensive. Second, I simulate the network dynamics by implementing two types of synaptic plasticity, each based on a




Figure 2.1: Architecture of the model. The model includes the MEC and the hippocampal CA3 regions of the brain. Grid cells in the MEC generate spike trains based on the rat’s position on the track. These spike trains drive excitatory CA3 cells, which in turn drive both inhibitory and other excitatory CA3 cells. Inhibitory CA3 cells provide feedback inhibition to excitatory CA3 cells. The connections among cells in each region are defined by the weight matrices, W^i , $1 \leq i \leq 4$. A rate-based plasticity model controls $W^1(t)$, and an STDP model controls $W^2(t)$.

different induction protocol, that drive the network to a state in which the network dynamics of the model follow the network dynamics observed experimentally.

2.1 Architecture of the Model

Figure 2.1 shows the architecture of the model. The first level of input is the location of the rat on the track. This information determines the firing rate of grid cells in the MEC. As discussed in the introduction, grid cells are spatially selective, primarily spiking in hexagonal patterns within an environment. Thus a grid cell has a higher probability of spiking when the rat is near one of its grid points. I provide a detailed explanation of the grid cell model in Section 2.2.1.

When a grid cell spikes, it sends current to excitatory cells in the hippocampal CA3 region. These excitatory cells are the potential place cells of the system. A cell is characterized as being excitatory or inhibitory based on its effect on other cells. When

an excitatory cell spikes, it tends to depolarize the cell, or increase its voltage. When an inhibitory cell spikes, it tends to hyperpolarize the cell, or decrease its voltage.

Place cells are found in both the CA1 and CA3 regions of the hippocampus, but this work focuses on place cells in the CA3 because of this region's importance and the interesting behavior of its place cells. The connections from the MEC to the CA3 are the important connections for the formation of place fields. The recurrent collateral connections among excitatory CA3 cells are the important connections for the backward shift of place fields.

I have incorporated feedback inhibition into the model for two reasons. First, feedback inhibition is biologically present in the CA3. Second, it provides a mechanism for spreading out the firing fields of place cells. When a place field forms at one location, it drives inhibitory cells to spike at that same location. The inhibitory cells then inhibit other place cells, lowering the probability that they will form place fields at the same location.

The architecture of the model is an approximation of the true anatomy of the system. As discussed in Section ??, two pathways from the MEC to the CA3 exist, one that passes through the DG and one that directly connects the MEC to the CA3. In this work I implement only the direct pathway because cells in the DG spike sparsely. A second anatomical pathway not incorporated in the model is the connection from the LEC to the CA3. Because cells in the LEC have very weak spatial tuning, the LEC provides little information about the rat's location on the

track (Hargreaves et al. [13]). This model can easily be extended to incorporate other anatomical regions.

I construct four weight matrices to define the interactions among cells. Each nonzero element in the initial weight matrix corresponds to a synapse, or connection, between two cells. I choose the synapses randomly according to the density of connections. Rate-based plasticity affects W^1 , which stores the weights from MEC cells to excitatory CA3 cells, and STDP affects W^2 , which stores the weights among excitatory CA3 cells. I provide a detailed description of each plasticity model in Section 2.3.

I set the constant weight for each synapse in W^3 and W^4 , which specify the weights between excitatory and inhibitory CA3 cells, randomly between zero and the upper weight bound. I set the initial weights for all synapses in W^1 and W^2 to the same value, which is small enough that CA3 cells initially have a low spike rate.

2.2 Single-Cell Dynamics

Before determining network dynamics, the model must first determine the dynamics of each individual cell within the network. I use two different methods to model the dynamics of individual cells, one for modelling grid cells in the MEC and one for modelling cells in the CA3. This is due to the architecture of the model. Because grid cells provide the first level of input, I simulate their firing patterns based on

experimental data. The situation is different, however, for cells in the CA3, which integrate input from other cells. I use the integrate and fire model to simulate the dynamics of individual CA3 cells.

2.2.1 Grid Cell Model

As Hafting suggested, I define a grid cell by three parameters, the tilt θ , where

$$0 \leq \theta < \pi/3;$$

the base b , where

$$0 < b < \infty;$$

and the offset $\delta = (r, \phi)$, where

$$0 \leq r < b \quad \text{and} \quad 0 \leq \phi < 2\pi.$$

The tilt specifies the angle from the Cartesian x-axis to the grid axis, the base specifies the distance from one grid point to another along a grid lateral, and the offset specifies both the magnitude, r , and direction, ϕ , of the difference between the grid center and the origin (Hafting et al. [12]). The grid center is given by

$$\mathbf{c} = (r \cos(\phi), r \sin(\phi)). \tag{2.2.1}$$

The set of grid points, $G(\theta, b, \delta)$, form hexagonal patterns. This set is the union of two sets, $G_1(\theta, b, \delta)$ and $G_2(\theta, b, \delta)$, that are staggered with respect to each other.

The set is defined by

$$G(\theta, b, \delta) = G_1(\theta, b, \delta) \cup G_2(\theta, b, \delta), \quad (2.2.2)$$

where

$$G_1(\theta, b, \delta) = \{\mathbf{c} + kb(\cos \theta, \sin \theta) + 2jh(-\sin \theta, \cos \theta) : j, k \in \mathbb{Z}\}, \quad (2.2.3)$$

$$G_2(\theta, b, \delta) = \{\mathbf{c} + (k + (1/2))b(\cos \theta, \sin \theta) + (2j - 1)h(-\sin \theta, \cos \theta) : j, k \in \mathbb{Z}\}, \quad (2.2.4)$$

where h is the grid height, given by

$$h = b \tan(\pi/3)/2. \quad (2.2.5)$$

Figure 2.2 shows an example of G , where the elements of G_1 are marked with black circles, and the elements of G_2 are marked with red diamonds. The grid's tilt, base, offset, and center are also demonstrated.

It is important to vary the grid cell parameters so that place cells receive enough variety in their grid cell input. I generate the parameters for each grid cell randomly within the appropriate ranges given in Table 2.2.1. I vary the tilt θ between 0 and $\pi/3$ because a tilt of 0 radians is equivalent to a tilt of $\pi/3$ radians. Biologically, the base of grid cells lies between 28 and 73 cm (Solstad et al. [26]). However, I confine the base of the grid cells in the model to lie between 28 and 50 cm because the radius of the track is only 33 cm wide. I have chosen the bounds for the offset, δ , so that the center of the grid always lies within the track.




Figure 2.2: Graphical demonstration of grid cell parameters. This plot shows an example of a grid, $G(\theta, b, \delta)$, where $\theta = \pi/4$ rad, $b = 1.146$, and $\delta = (0.496, \theta)$. The black circles and red diamonds represent the elements of G_1 and G_2 , respectively. The large circle specifies the grid center.

Table 2.1: Grid Cell Parameters

Parameter	Value	Description
θ	0-pi/3	tilt
b	28-50 cm	base
r	0-33 cm	offset magnitude
ϕ	0-2 π	offset direction
γ	0.018	spread of grid
t_{ref}	3 ms	refractory period
f_{max}	20 Hz	maximum spike rate

The firing rate of a grid cell depends on the rat's distance to the nearest grid point. This distance is defined by the metric

$$d(\mathbf{x}, G) = \min\{|\mathbf{x} - \mathbf{y}| : \mathbf{y} \in G\}, \quad (2.2.6)$$

where \mathbf{x} denotes the position of the rat.

Each grid cell generates a spike train that depends on $d(\mathbf{x}, G)$. A spike train is represented by a vector containing the times at which the cell spikes. This vector is set by computing the interspike interval (ISI) between each pair of spikes. Rodieck has shown that the ISIs from experimental data follow an exponential distribution, where the probability is high for a short ISI and decays exponentially as the interval grows (Rodieck [23]). Because the interval between events in a Poisson process also follows the exponential distribution, I model each grid cell's spike train to be a non-homogeneous Poisson process, where the firing rate depends on the rat's distance to the grid.

I follow the algorithm outlined below to generate a thinned, nonhomogeneous Poisson spike train for each grid cell. This approach was taken by Savelli in his grid cell model (Savelli and Knierim [24]). The algorithm incorporates the cell's refractory period, t_{ref} , which represents the interval of time after a cell spikes in which it cannot generate another spike. Biologically, the refractory period is due to the depletion of resources that a cell needs to generate a spike. The algorithm also incorporates the maximum firing rate of the cell, f_{max} . If the rat were to remain fixed on one of the cell's grid points, the cell's average firing rate would be f_{max} .

Algorithm: Generate a thinned, nonhomogeneous Poisson spike train

1. Initialize $t_0 = 0$.
2. Propose an ISI from an exponential distribution with mean $1/f_{\max}$, given by

$$\Delta t = \max(\text{exprnd}(1/f_{\max}), t_{\text{ref}}), \quad (2.2.7)$$

where f_{\max} is the maximum spike rate, and t_{ref} is the cell's refractory period.

3. At $t = t_0 + \Delta t$, accept a spike with probability

$$P = \exp\left(\frac{-d^2(\mathbf{x}, G)}{\gamma b^2}\right), \quad (2.2.8)$$

where \mathbf{x} denotes the position of the rat, and γ determines the spread of the grid.

4. Set $t_0 \rightarrow t_0 + \Delta t$; return to step 2.

Figure 2.3 demonstrates the effect of γ on the spread of the firing fields of a grid cell as a rat explores a rectangular enclosure.

2.2.2 CA3 Cell Model: Integrate and Fire

Simulating the dynamics of cells in the CA3 requires a much different method than that used for grid cells in the MEC. Unlike the grid cells of the model, which provide




Figure 2.3: Comparison of the effect of γ values on the firing fields of the cell. As γ increases, the width of the Gaussian increases, causing the spread of the firing fields to increase. For all plots, $\theta = \frac{\pi}{4}$ radians, $b = 15$ cm, and $\delta = (2 \text{ cm}, \theta \text{ rad})$. The rat runs the same trajectory for all values of γ .




Figure 2.4: Circuit diagram of an integrate-and-fire neuron. The cell membrane acts as a leaky capacitor permeable to a leak current, capacitive current, excitatory current, and inhibitory current. Each channel is characterized by the conductance and reversal potential. The synaptic conductances, g_E and g_I , are gated by neurotransmitter from a presynaptic cell.

the first level of neuronal input, CA3 cells integrate neuronal input from both grid cells and other CA3 cells. Thus to simulate the dynamics of an individual CA3 cell, I must determine the cell's response to spikes from other cells within the network.

Several methods have been developed to model single-cell dynamics, and there is a give and take between accuracy and computational cost for each method. In 1952, Hodgkin and Huxley developed an impressive phenomenological model for single-cell dynamics (Hodgkin and Huxley [16]). While this model is accurate, it is too computationally expensive to be useful for my network containing over a thousand cells. Kellems is currently working to accurately model spike times for individual cells using model reduction techniques (Kellems et al. [18]). In the future I will use his model for the CA3 cells because of its accuracy at little computational cost.

Currently, I am using the integrate and fire (IAF) model for single-cell dynamics. The voltage of a cell generally has slight changes until it reaches some threshold value, at which point the voltage drastically increases and decreases, and the cell is said to have spiked, or fired. IAF approximates the voltage of the cell in the subthreshold regime, when the voltage is less than the threshold voltage, and approximates the times at which the cell spikes.

The dynamics of the integrate-and-fire cell are governed by three coupled ODEs describing the cell's voltage, excitatory conductance, and inhibitory conductance. Figure 2.2.2 shows the circuit representation of the cell. Kirchoff's Current Law

applied to this circuit determines the change in voltage of cell i , given by

$$c_m \frac{dV_i}{dt}(t) = g_L(V_L - V_i(t)) + g_i^E(t)(V_E - V_i(t)) + g_i^I(t)(V_I - V_i(t)) \quad (2.2.9)$$

for $V_i(t) \leq V_{th}$.

The cell membrane acts as a leaky capacitor permeable to the flow of chloride ions, resulting in a constant membrane conductance per unit area, g_L , and membrane capacitance per unit area, c_m . In the absence of driving current, the cell's voltage decays to its resting voltage, V_L , at a rate given by $\tau_m = c_m/g_L$.

In addition to the leak conductance, the cell also has transient synaptic conductances caused by spikes from other cells that form synapses, or connections, onto the cell. When a presynaptic cell spikes, it releases chemicals known as neurotransmitter that briefly bind to receptors on the postsynaptic cell, inducing a transient synaptic conductance biased by an associated reversal potential. The excitatory reversal potential, V_E , is much larger than the threshold voltage of a cell, and the inhibitory reversal potential, V_I , is close to the resting voltage of a cell.

Excitatory and inhibitory conductances are driven by spikes from excitatory and inhibitory cells, respectively. In the model an excitatory CA3 cell receives excitatory input from grid cells in the MEC and from other excitatory CA3 cells, and it receives inhibitory input from inhibitory CA3 cells. Thus the excitatory and inhibitory conductances for excitatory cell i are governed by

$$\tau_E \frac{dg_i^E}{dt}(t) = -g_i^E + \sum_{j=1}^{N^{\text{MEC}}} W_{ij}^1(t) S_j^1(t) + \sum_{j=1}^{N^{\text{CA3(E)}}} W_{ij}^2(t) S_j^2(t) \quad (2.2.10)$$

and

$$\tau_I \frac{dg_i^I}{dt}(t) = -g_i^I + \sum_{j=1}^{N^{\text{CA3(I)}}} W_{ij}^3 S_j^3(t), \quad (2.2.11)$$

where N^{MEC} denotes the number of MEC cells, $N^{\text{CA3(E)}}$ denotes the number of excitatory CA3 cells, and $N^{\text{CA3(I)}}$ denotes the number of inhibitory CA3 cells.

$S_j(t)$ denotes the spike train for cell j , given by

$$S_j(t) = \sum_{k=1}^{\#\mathbf{T}_j} \delta(t - \mathbf{T}_j(k)), \quad (2.2.12)$$

where $\mathbf{T}_j(k)$ is a vector containing the times at which cell j spikes, $\#\mathbf{T}_j$ is the number of spikes contained in \mathbf{T}_j , and δ denotes the Dirac-delta function, which is only nonzero when $t = \mathbf{T}_j(k)$.

The conductances for inhibitory CA3 cell i are very similar, governed by

$$\tau_E \frac{dg_i^E}{dt}(t) = -g_i^E(t) + \sum_{j=1}^{N^{\text{CA3(E)}}} W_{ij}^4 S_j^4(t) \quad (2.2.13)$$

and

$$g_i^I = 0. \quad (2.2.14)$$

Figure 2.1 demonstrates the connections used in the above equations.

These equations describe the dynamics of each cell in the subthreshold regime. When the voltage reaches its threshold value, the cell generates a spike, and the current time is added to the vector \mathbf{T}_j . The voltage is then reset to its reset value and held constant until the cell's refractory period has transpired.

Table 2.2: Integrate and Fire Parameters

Parameter	Value	Description
C_m	$1 \mu\text{F}/\text{cm}^2$	membrane capacitance
g_L	$0.3 \text{ mS}/\text{cm}^2$	leak conductance
V_L	-70 mV	resting voltage
V_I	-70 mV	inhibitory reversal potential
V_E	0	excitatory reversal potential
V_{reset}	-65 mV	reset voltage
t_{ref}	3 ms	refractory period
τ_E	5 ms	excitatory decay constant
τ_I	5 ms	inhibitory decay constant

Table 2.3: Weight Matrix Parameters

Parameter	Value	Description
N^{MEC}	540	number of MEC cells
$N^{\text{CA3(E)}}$	1200	number of excitatory CA3 cells
$N^{\text{CA3(I)}}$	96	number of inhibitory CA3 cells
W_{max}^1	0.2-0.3	maximum weight for W^1
δ^1	0.2	density of W^1
W_{max}^2	0.2-0.3	maximum weight for W^2
δ^2	0.25	density of W^2
W_{max}^3, W_{max}^4	0.8	maximum weight for W^3, W^4
δ^1, δ^2	0.2	density of W^3, W^4

2.3 Network Dynamics: Synaptic Plasticity

The weight matrices used to approximate the single-cell dynamics change in time due to synaptic plasticity. Both the formation and the backward shift of place fields depend on the manner in which the weights change. I have implemented a rate-based plasticity model to determine the dynamics of W^1 , or the weights from MEC cells to excitatory CA3 cells, and an STDP model to determine the dynamics of W^2 , or the weights among excitatory CA3 cells. Both models are phenomenological, and the characteristics of each cause the network dynamics in the model to resemble the network dynamics observed experimentally.

2.3.1 Implementation of Postsynaptic-Gated Rate-Based Plasticity

I have implemented the postsynaptic-gated rate-based plasticity model to simulate the formation of place fields. This plasticity model determines the change in weight of the connections from MEC grid cells to excitatory CA3 cells. The model is called postsynaptic-gated because no change in weight occurs if the instantaneous firing rate of the postsynaptic cell is zero. According to the plasticity model, the change in weight at time t from presynaptic cell j to postsynaptic cell i is given by

$$\Delta W_{ij}^1(t) = k(R_j(t) - \bar{R})R_i(t), \quad (2.3.1)$$

where k is a parameter that determines the learning rate, $R_j(t)$ is the instantaneous rate of presynaptic cell j , $R_i(t)$ is the instantaneous rate of postsynaptic cell i , and \bar{R} is a constant threshold rate. The rate of the presynaptic cell, $R_j(t)$, determines the direction of weight change. If $R_j(t) < \bar{R}$, LTD occurs; if $R_j(t) > \bar{R}$, LTP occurs.

To compute the cell's instantaneous spike rate, I convolve the cell's spike train with an exponential kernel. Given a spike train $S_j(t)$ as defined in equation 2.2.12, the instantaneous spike rate for cell j is given by

$$R_j(t) = \beta \int_0^\infty \exp\left(-\frac{t-s}{\tau_{Ca}}\right) S_j(s) ds. \quad (2.3.2)$$

The decay constant τ_{Ca} corresponds to the diffusion rate of calcium. This parameter choice is appropriate because synaptic plasticity depends on the density of calcium at the synapse (Yang et al. [30]). I set the constant β such that if the rat were to stay fixed on a grid point for a sufficiently long time, the mean of the instantaneous spike rate of the cell would be f_{\max} , the parameter used in Equation 2.2.7 to generate the spike trains for grid cells.

2.3.2 Implementation of Spike-Timing-Dependent Plasticity

The STDP model applied to the recurrent connections within the network of excitatory CA3 cells causes the backward shift in place cells' firing fields as seen experimentally. This backward shift suggests that the network is driven to a feedforward state in which place cells that spike at one location on the track drive place cells that spike at the next location.

Table 2.4: Rate-Based Plasticity Parameters

Paramter	Value	Description
k	1-5 ms	learning rate
β	10 Hz	rate constant
τ_{Ca}	100 ms	Calcium decay rate
\bar{R}	5 Hz	threshold rate

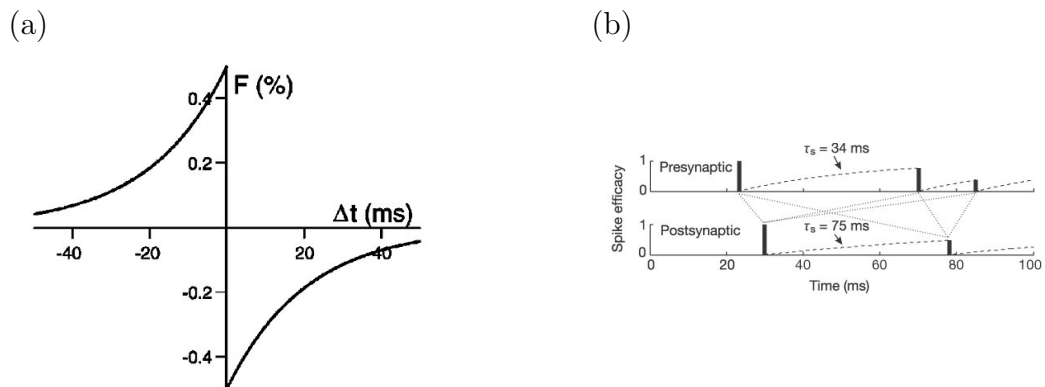


Figure 2.5: Weight change due to STDP. (a) This standard STDP curve shows the percentage weight change, F , due to Δt , the difference between the presynaptic spike time and the postsynaptic spike time. The curve is set to match experimental data as best as possible. Taken from Song and Abbott [27]. (b) In the Froemke and Dan model the weight change induced by each spike pair depends on the pre- and postsynaptic efficacy values in addition to $F(\Delta t)$. The efficacy value for each cell is set to 0 when the cell spikes, then recovers exponentially to 1. Taken from Froemke and Dan [10].

The basic STDP model proposed by Song and Abbott states that if the presynaptic cell spikes before the postsynaptic cell, LTP occurs. If the postsynaptic cell spikes before the presynaptic cell, LTD occurs. The percentage weight change due to the presynaptic spike at time t_{pre} and the postsynaptic spike at time t_{post} is given by

$$F(\Delta t) = \begin{cases} A_+ \exp(\Delta t/\tau_+) & \text{if } \Delta t < 0, \\ A_- \exp(\Delta t/\tau_-) & \text{if } \Delta t \geq 0, \end{cases} \quad (2.3.3)$$

where $\Delta t = t_{pre} - t_{post}$ (Song and Abbott [27]). The percentage weight change is plotted in Figure 2.3.2.

To avoid situations in which the weights become arbitrarily large, I restrict the weight of each synapse to lie within the interval $[0, W_{max}^2]$. In the Song and Abbott STDP model, the weight change due to each pair of a pre- and postsynaptic spikes is given by

$$\Delta W^2 = F(\Delta t)W_{max}^2. \quad (2.3.4)$$

Froemke and Dan proposed a modified version of the STDP model in which the first spike in a burst of spikes provides the largest contribution to the weight change. In their model each spike has an efficacy value associated with it that depends on the amount of time that has passed since the cell's previous spike. The pre- and postsynaptic efficacy values are given by

$$\varepsilon_{pre}(\Delta t_s) = 1 - \exp(-\Delta t_s/\tau_{pre}), \quad (2.3.5)$$

$$\varepsilon_{post}(\Delta t_s) = 1 - \exp(-\Delta t_s/\tau_{post}), \quad (2.3.6)$$

where Δt_s denotes the amount of time that has passed since the cell's previous spike, τ_{pre} denotes the presynaptic efficacy decay rate, and τ_{post} denotes the postsynaptic efficacy decay rate. These efficacy values are demonstrated in Figure 2.3.2.

According to the Froemke and Dan model, the weight change due to each spike pair is given by

$$\Delta W_{ij}^2 = \varepsilon_{pre}(\Delta t_s)\varepsilon_{post}(\Delta t_s)F(\Delta t)W_{max}^2, \quad (2.3.7)$$

where $F(\Delta t)$ denotes the percentage weight change from the Song and Abbot model, given by equation 2.3.3, and W_{max}^2 again denotes the maximum weight allowed at each synapse (Froemke and Dan [10]). Note that if the efficacy values of the pre- and postsynaptic spikes are set to one, the Froemke and Dan model is identical to the Song and Abbott model.

Because a place cell spikes in bursts each time the rat crosses its firing field, it is very important to properly sum the weight change due to each spike pair. To stay as true to biology as possible, I have implemented the STDP model of Froemke and Dan. This model is effective in causing a backward shift in the place fields.

Table 2.5: STDP Parameters

Parameter	Value	Description
A_+	20	maximum percentage of weight increase
A_-	21	maximum percentage of weight decrease
τ_+	20 ms	decay constant for weight increase
τ_-	20 ms	decay constant for weight decrease
τ_{pre}	20 ms	presynaptic efficacy decay constant
τ_{post}	50 ms	postsynaptic efficacy decay constant

Chapter 3

Results

The computational model I have built is a useful tool only if it agrees with experimental data. To determine the reliability of the model, I have tested it against the experimental data obtained during the learning stage of the double rotation experiment. Two phenomena occur during this stage. First, place cells in the CA3 integrate neuronal input from grid cells in the MEC to form single firing fields. Second, these firing fields shift backward as the rat continues to circle the track. My computational model predicts both phenomena, demonstrating that rate-based synaptic plasticity can enable place cells to form single firing fields, and STDP can cause the place fields to shift backward.

For all simulations presented in this thesis, I use 540 grid cells, 1200 excitatory CA3 cells, and 96 inhibitory CA3 cells. As Ascoli has estimated, the CA3 contains an estimated 200,000 cells, but it is computationally unrealistic to simulate a network of

this size. To ensure that cells receive enough variety in their input, I use a density of 0.2 for W^1 , which stores the weights from MEC cells to excitatory CA3 cells, and a density of 0.25 for W^2 , which stores the weights among excitatory CA3 cells. These densities are much larger than 2 to 10 percent densities found experimentally (Ascoli et al. [2]).

To handle the stochastic input from grid cells and the nonlinearities of the differential equations governing CA3 cells, I use the Backward Euler timestepping method to update the dynamic variables of the system. Even given the relatively small network size of the model, W^1 alone contains 129,600 dynamic variables that must be updated at each timestep. W^2 adds another 360,000 dynamic variables, giving the system almost 0.5 million dynamic variables. The rat takes 18 seconds to complete one lap, and I use a timestep of 1 ms. Thus I must update almost 0.5 million variables 18,000 times to simulate a single lap around the track. In Section 4.2.1, I propose methods to handle the enormous computational cost that accompanies solving this dynamical system.

3.1 Formation of Place Fields

The first experimental behavior the model simulates is that place cells form single firing fields along the track. This occurs despite the fact that place cells receive the majority of their spatial input from grid cells, which spike at several locations forming hexagonal patterns. This work demonstrates that the postsynaptic-gated rate-based

plasticity model applied to W^1 enables place fields to form. Figure 3.1 shows an example of the resulting firing fields of CA3 cells in the model. For each location on the track, I calculate the spike rate of a cell as described in Section 2.3.1. Unless otherwise noted, all parameter values for the results presented here are taken from Table 2.2.2 and Table 2.3.1.

Rate-based plasticity polarizes the weights in W^1 , as shown in Figure 3.1d. Because the plasticity model is postsynaptically gated, the weights can potentially change each time the postsynaptic cell, or the CA3 cell in this case, has a positive spike rate. The presynaptic cell, or MEC cell in this case, then determines the direction of weight change. Thus if there is a high correlation between the spike trains of MEC cell j and excitatory CA3 cell i , rate-based plasticity drives the weight from cell j to cell i , W_{ij}^1 , to its upper bound. If there is a low correlation, rate-based plasticity drives W_{ij}^1 to zero.

In order to evaluate the results of the model, I use two criteria to define a place cell. First, a place cell must attain a spike rate of at least 3 Hz at some point on the track. Second, the cell must have only one firing field. I define a firing field to be an interval on the track over which the cell has a nonzero spike rate such that the cell spikes at a rate greater than one-third its maximum rate at some point within the interval.

Using these two criteria, I have found that place cells consistently comprise over 75 percent of the excitatory CA3 cells in my simulations given a range of parameter

values. In reality, place cells comprise only about 20 percent of the CA3 cells in any given environment (Knierim, personal communication, March 2009). However, it is necessary to generate such a large percentage of place cells in the model because the model simulates only 1200 excitatory CA3 cells, which is less than 1 percent of the estimated 180,000 excitatory CA3 cells contained in the CA3 (Ascoli et al. [2]). Three parameters significantly impact the formation of place fields: the learning rate of the rate-based plasticity model, the upper bound for the weights from the MEC to the CA3, and the variation in the base values for the network of grid cells.

3.1.1 Learning Rate

The rate-based plasticity model is effective only if the learning rate, specified by k in Equation 2.3.1, is sufficiently large. A fast learning rate is necessary to handle competition among different locations on the track. If the learning rate is high, when the CA3 cell spikes, the weights from all grid cells that have a high spike rate at that location are increased by a large enough amount that the cell has a much higher probability of spiking at that same location on the track during the subsequent laps. More importantly, the weights from the grid cells that do not have a high spike rate at that location are decreased by a large enough amount that the cell has a small probability of spiking at other locations on the track. If the learning rate is slow, the weights oscillate back and forth, and the cell forms multiple firing fields.

To demonstrate the effect of the learning rate, I present a simplified example in

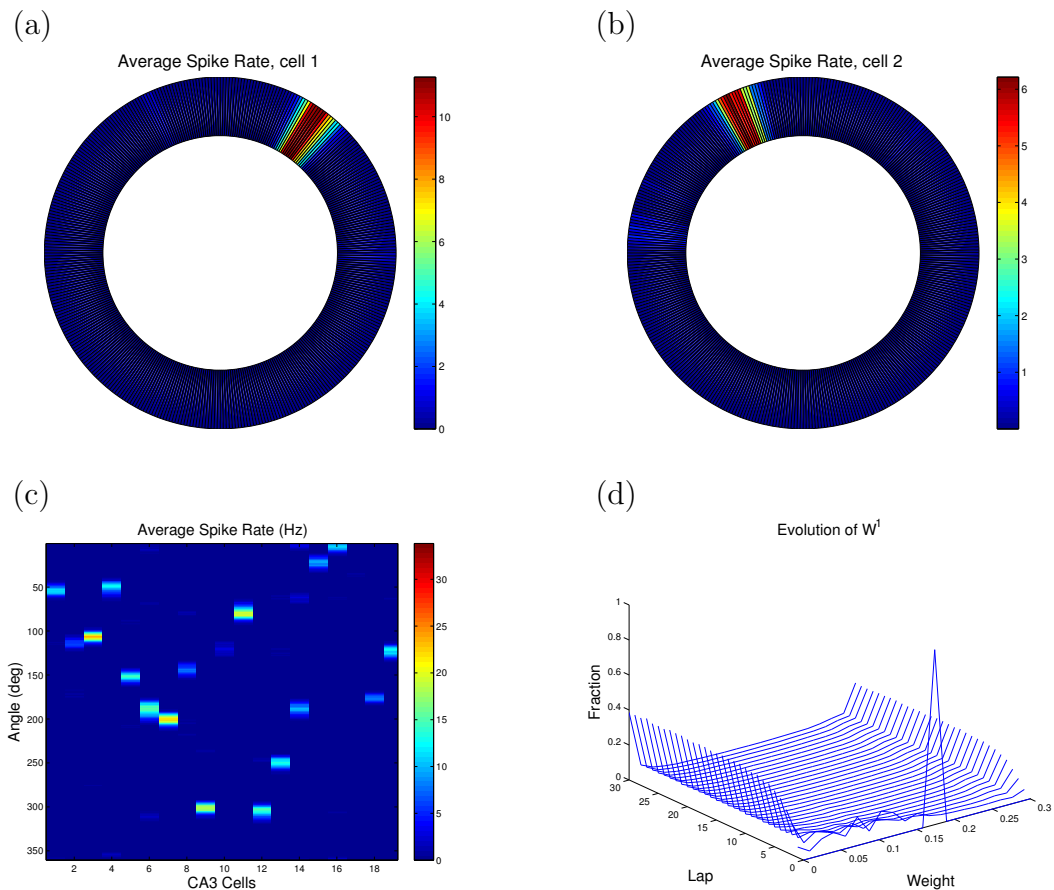


Figure 3.1: Place field formation. Rate-based plasticity coupled with grid cell input enables place cells to form in the CA3. (a,b) These two examples demonstrate the place fields that form along the track. For each cell, I have plotted the average spike rate (Hz) for each degree on the track. (c) The y-axis represents the track, where each position is defined by its angle (deg). I show the firing fields of the first nineteen CA3 cells. (d) Rate-based plasticity polarizes the weights from the MEC to the CA3, driving all weights to zero or to their upper bound, which is 0.3 for this example. I initially set all weights to 0.18.

which a CA3 cell receives input from twenty grid cells. I have set the parameters of these grid cells so that each cell in group A, which consists of the first ten grid cells, forms a single firing field centered at 120 degrees. Each cell in group B, which consists of the last ten grid cells, forms a single firing field centered at 295 degrees. All weights are initially set to 0.18. Holding all other parameters constant, I vary the learning rate and examine how it affects the MEC input and the firing fields of the CA3 cell. The MEC input at each degree on the track is given by

$$\text{input}(\text{deg}) = \sum_{\text{deg}=1}^{360} W_j(\text{deg})R_j(\text{deg}), \quad (3.1.1)$$

where W_j is the weight from MEC cell j to the CA3 cell, and R_j is the instantaneous spike rate of MEC cell j .

Figure 3.1.1 shows the results of this experiment. When $k = 0.05$, the CA3 cell forms two firing fields, one centered at 120 degrees and one centered at 295 degrees. Due to the slow learning rate, the weights oscillate back and forth, and neither group A nor group B obtains full control of the CA3 cell. If the learning rate is fast, however, the CA3 cell forms a single firing field. When the rat passes 120 degrees on the track, the weights from group A significantly increase while the weights from group B significantly decrease. This causes the CA3 cell to spike very little, if at all, when the rat passes 295 degrees. When the rat again passes 120 degrees, the CA3 cell has a larger spike rate, further polarizing the weights between group A and group B.

A large range of values for the learning rate permits place fields to form. In the simple example described above, I varied only the learning rate and held all

other parameters constant. The CA3 cell consistently formed a single firing field for $k \in [0.5, 10]$.

Note that because the learning rule is fast, the locations of the place fields depend on the rat's initial position on the track. In the situation above, if the place cell were to spike at 295 degrees before it spiked at 120 degrees, the place field would form at 295 degrees. This is consistent with experimental data, in which place cells form different firing fields in different environments.

3.1.2 Other Factors Affecting Place Field Formation

In addition to a fast learning rate, the variety in base values and the upper weight bound also significantly impact the formation of place fields. Rate-based plasticity isolates the location at which the place cell is driven by a large number of grid cells. The success of the plasticity model requires this group of grid cells to have only one firing location in common. If the group shares multiple firing locations along the track, the place cell has no mechanism for choosing one location over another, and it will form multiple firing fields. A variety of base values for the grid cells ensures that a random group of grid cells generally does not have multiple firing fields in common. In the model I randomly choose the bases of the grid cells from 28 cm and 50 cm, which falls within the biologically realistic range (Witter and Moser [29]).

The third important factor is the choice of the upper bound for the weight from MEC cells to excitatory CA3 cells. If this value is too small, CA3 cells never receive




Figure 3.2: Effect of the learning rate. (a) MEC cells in group A form firing fields centered at 120 degrees (left). MEC cells in group B form firing fields centered at 295 degrees. The colorbar represents the sum of the spike rates (Hz) of each cell in the respective group. (b-d) The CA3 cell forms a single place field given a sufficiently fast learning rate, k . The two left plots show the input from all MEC cells to the CA3 cell. I calculate the MEC input according to Equation 3.1.1. The center plot shows the average spike rate (Hz) of the CA3 cell taken over all 15 laps. The right plot shows the final weight distribution.

enough input from the MEC to spike consistently. Place cells in this situation form a weak place field over which the cell has a very low spike rate. If the upper bound is too large, a single grid cell has too large of an impact on the CA3 cell. Because grid cells spike at several locations, this causes the CA3 cell to spike at several locations as well. The success of the rate-based plasticity model in forming place fields requires the upper bound to be small enough such that the CA3 cell only spikes consistently when a group of its presynaptic MEC cells spike together.

In my simulations, I set the density of connections between the MEC and the CA3 to be 20 percent. Because the model simulates far fewer grid cells than exist in the MEC, this high density value is necessary to ensure that place cells receive enough variety in their grid cell input. Using this density value, I have found that place fields form given an upper weight bound between 0.2 and 0.3.

3.2 Backward Shift

As the rat circles the track in the double rotation experiment, place cells in the CA3 not only form single firing fields but also experience a backward shift in their firing fields. This backward shift is important for navigation and may have important implications on the mismatch stage of the double rotation experiment, discussed in Section 1.1.2. The place cells in the model also show this backward shift in their firing fields as STDP changes the weights among excitatory CA3 cells.

3.2.1 Simple Example

To provide some intuition for how STDP is successful in causing place fields to shift backward, I begin with a simple example of a ring of twelve place cells. The cells are arranged in a ring architecture, where each cell connects with its two neighbors, as shown in Figure 3.2.1. Each place cell receives input from the MEC only when the rat is at the location on the track corresponding to the cell number. Cell 1, for example, only receives input when the rat is at the position on the track corresponding to the number 1 on a clock. In this simple example, the track is small enough that the rat completes a lap in 120 ms. This fast lap time allows a network of only twelve place cells to represent the track.

As the rat runs clockwise around the track, STDP drives the network to a feed-forward state in which each cell only receives CA3 input from the CA3 cell directly preceding it. As described in Section 2.3.2, the basic principle of STDP is that if the presynaptic cell spikes before the postsynaptic cell, the weight increases; if the postsynaptic cell spikes before the presynaptic cell, the weight decreases.

Consider the weights between Cell 1 and Cell 2. Each time the rat circles the track, Cell 1 spikes before Cell 2. With respect to the connection from Cell 1 to Cell 2, Cell 1 is the presynaptic cell, and Cell 2 is the postsynaptic cell. Thus when Cell 1 spikes before Cell 2, the weight from Cell 1 to Cell 2 increases. On the other hand, the weight from Cell 2 to Cell 1 decreases because in this case, the postsynaptic cell spikes before the presynaptic cell. In this manner STDP drives all feedforward

weights to their upper bound and drives all feedback weights to zero.

The resulting feedforward architecture causes a backward shift in the firing fields of the place cells, demonstrated in Figure 3.2.1. Consider again the connection from Cell 1 to Cell 2. As the weight from Cell 1 to Cell 2 increases, each spike from Cell 1 has a greater effect on Cell 2. This additional excitatory input causes the voltage of Cell 2 to reach its threshold value sooner in time, resulting in a backward shift in its firing field.

3.2.2 Backward Shift in the Full Model

Similar to the effect of STDP on the ring of place cells, STDP drives the weights among CA3 place cells in the full model to a largely feedforward state in which place fields shift backward. I initially set all weights among excitatory CA3 cells to 0.05, a low enough value that CA3 interaction does not interfere with the formation of firing fields due to grid cell input and rate-based plasticity. STDP then changes the weights among excitatory CA3 cells at the same time that rate-based plasticity changes the weights from MEC cells to excitatory CA3 cells. Figure 3.2.2 shows an example of the resulting backward shift observed in the network of place cells.

Two different types of backward shift occur. Some place fields shift backward but retain their shape, as shown in plot (a). Others become more skewed in the negative direction, as shown in plot (b). The latter case is due to input from the MEC, which does not shift backward.




Figure 3.3: Simple example of the backward shift. This simple network consists of twelve place cells in the CA3, each receiving input from the MEC at one location on the track. (a) Each place cell initially connects with its two neighbors in the ring. (b) STDP drives the network to a feedforward state in which each cell sends current to the cell following it in the ring. (c) Each place cell shifts backward due to increased input from the place cell preceding it. Each dot represents the average firing field for a place cell during laps 1-4 (red), laps 5-12 (blue), laps 13-23 (green), and laps 24-30 (black). The firing fields become stable after thirty laps because all weights have been driven to their upper or lower bound.

The resulting weight distribution due to STDP is very different from the resulting weight distribution due to rate-based plasticity. According to the rate-based plasticity model I have implemented, weights can decrease even if the presynaptic cell has not spiked. This causes every weight to be driven to its upper or lower bound. According to the STDP model, weights change very little due to pairs of pre- and postsynaptic spikes separated by more than 50 ms. In the model the rat takes 50 ms to pass each degree on the track, which means that only weights between cells with overlapping place fields change.

I calculate the shift of each cell by taking the difference of the center of mass (COM) of the place field during the last five laps and the COM of the place field during laps 3-7. I discount the first two laps to give place cells sufficient time to form their place fields. The COM of a place field is given by

$$\text{COM} = \frac{\sum_{x=1}^{360} xR(x)}{\sum_{x=1}^{360} R(x)}, \quad (3.2.1)$$

where $R(x)$ is the average spike rate of the cell at x degrees on the track. I then classify each cell as having a backward shift if $\text{COM} \leq -1$, a forward shift if $\text{COM} \geq 1$, or no shift if $-1 < \text{COM} < 1$.

As the rat circles the track, some place fields shift backward, some shift forward, and some remain fixed. This variance in the place cell response agrees with experimental data, in which all three responses are observed (Knierim, personal communication, March 2009). Despite this, the dominant effect of STDP is a backward shift. For the example shown in Figure 3.2.2, the number of place fields that shift backward is

almost double the number of place fields that shift forward. The mean of the shift among all place cells is -0.88 degrees. I find similar results for other simulations using a range of parameter values given in Table 2.3.2.

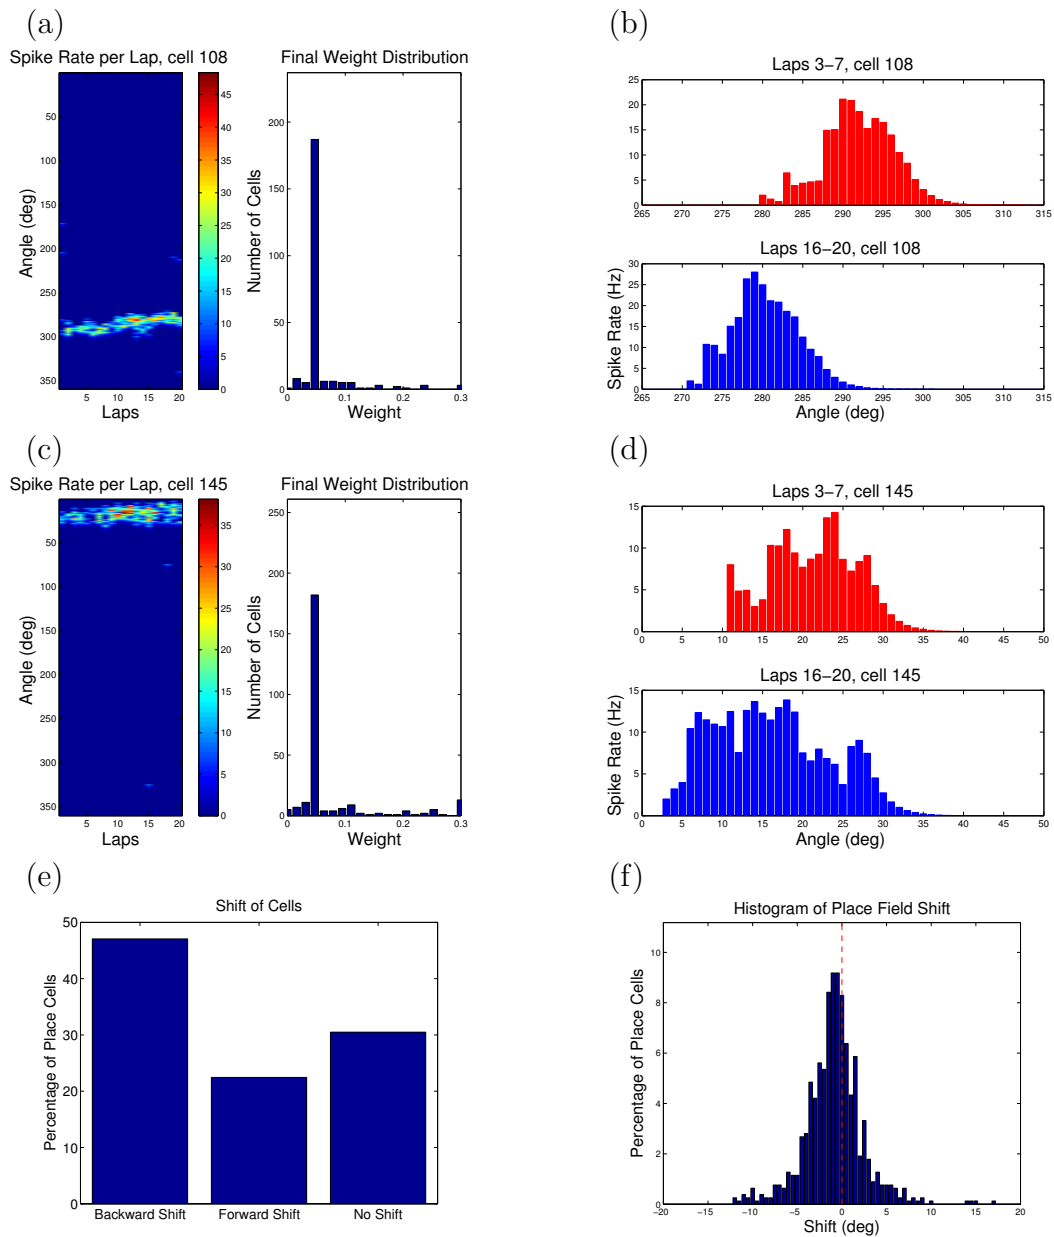


Figure 3.4: Backward shift in the full network model. (a-d) These examples show the backward shift in two individual CA3 cells. I again define each location on the track by its angle (deg). The COM of the firing field for cell 108 is 291.9 degrees for laps 3-7 and 280.3 degrees for laps 16-20, giving cell 108 a backward shift of 11.6 degrees. Cell 145 has a similar backward shift, but its place field becomes more skewed in the negative direction. The COM of the firing field for cell 145 is 21.2 degrees for laps 3-7 and 16.1 degrees for laps 16-20, giving it a backward shift of 5.1 degrees. (e) For this simulation 47.1 percent of the place cells shift backward, 22.4 percent shift forward, and 30.5 percent do not shift. (f) This histogram shows the shift of each place cell. The mean shift is -0.88 degrees.

Chapter 4

Conclusion

I have built a reliable computational model for the learning stage of the double rotation experiment that predicts the behavior of place cells observed experimentally. As the rat initially circles the track, rate-based plasticity changes the weights from the MEC to the CA3 such that place cells integrate input from grid cells to form single firing fields. STDP then changes the weights among the recurrent collateral CA3 connections to cause the firing fields to shift backward along the track.

As discussed in Chapter 1, synaptic plasticity at work in the CA3 region of the hippocampus is important for learning and memory as well as diseases such as Alzheimer's disease and epilepsy. Thus a better understanding of the effect of synaptic plasticity on the network of cells in the CA3 region would impact important areas of research both inside and outside the context of spatial representation and navigation. Because it predicts biological phenomena, the model is a solid platform that I

will expand and improve to shed light on the mechanisms underlying these important functions.

4.1 Applications

Because I designed the model to simulate the learning stage of the double rotation experiment, one direct application for the model is to simulate the mismatch stage of the experiment as well. This second stage shows that place cells in the CA3 are influenced by some unknown source or mechanism, resulting in a difference between place cell behavior in the CA1 and in the CA3. Additionally, the model can be expanded and applied to a variety of other areas outside of the spatial representation setting.

4.1.1 Mismatch Stage of the Double Rotation Experiment

The mismatch stage of the double rotation experiment is performed after the rat has learned the cues along the track and the place cells have formed their place fields. In this stage the distal cues are rotated clockwise, and the local cues are rotated counter-clockwise. As the rat again runs clockwise around the track, experimenters record the response of grid cells in the MEC and place cells in the CA1 and CA3. Figure 4.1.1 demonstrates the response of each group of cells.

The Knierim lab has found that most grid cells rotate clockwise with the distal cues and shift (Knierim, personal communication, March 6, 2009). The shift is important

because it drastically changes the locations on the track at which the grid cell spikes. Consider a grid cell that has two grid points lying directly on the track, shown in red in Figure 4.1.1c. When the grid rotates, the two red grid points still lie directly on the track. Thus without a shift, the grid cell's firing fields would simply rotate along with the distal cues. When the grid shifts, however, the two grid points marked in red fall off the track, no longer affecting the firing fields of the grid cell. Two new grid points, however, shift onto the track, causing the grid cell to spike at locations unrelated to the rotation of the cues.

Because the grid cells have shifted, place cells receive very different input from the grid cells after the cues have been rotated. Most place cells in the CA1 respond by either rotating with the distal cues or remapping, forming a new firing field unrelated to the rotation of the cues (Lee et al. [19]). The former response is due to the rotation of grid cells with the distal cues, and the latter response is due to the shift of grid cells, changing the locations at which a place cell receives its input.

Place cells in the CA3, however, respond very differently. While some remap, most place cells rotate counter-clockwise, following the local cues (Lee et al. [19]). This implies that some source or mechanism other than grid cell input controls the firing fields of these CA3 place cells when the environment changes. As discussed in Section 1.1, although place cells receive the majority of their spatial input from grid cells in the MEC, they also receive weak spatial input from cells in the LEC. Knierim has hypothesized that if the weak input from the LEC, which is difficult to

analyze, were to shift with the local cues, the tendency of place cells to already shift backward may combine with this weak LEC input to cause place cells to follow the local cues (Knierim, personal communication, March 6, 2009). Because the computational model accurately predicts behavior of cells during the learning stage of the experiment, one direct application for the model is to simulate the mismatch stage of the experiment to test this hypothesis.

4.1.2 Various Other Applications

Synaptic plasticity is essential for learning and memory to take place in the brain. Thus the model can be adapted and applied to a variety of settings concerning learning and memory in general. One application related to the formation of place cells in the CA3 is the formation of grid cells in the MEC. Grid cells receive input from head direction cells, which primarily spike when the rat's head is facing a certain direction. Grid cells integrate this directional input along with input concerning the rat's velocity to form firing fields in hexagonal patterns. Incorporating head direction cells into the model would allow me to investigate the formation of grid cells and provide a more complete platform for studying spatial representation and navigation in the brain.

The study of epilepsy is an important application of the model outside the context of spatial representation. At the onset of an epileptic seizure, a group of cells in the brain spikes synchronously at a much higher rate than normally. Because a burst

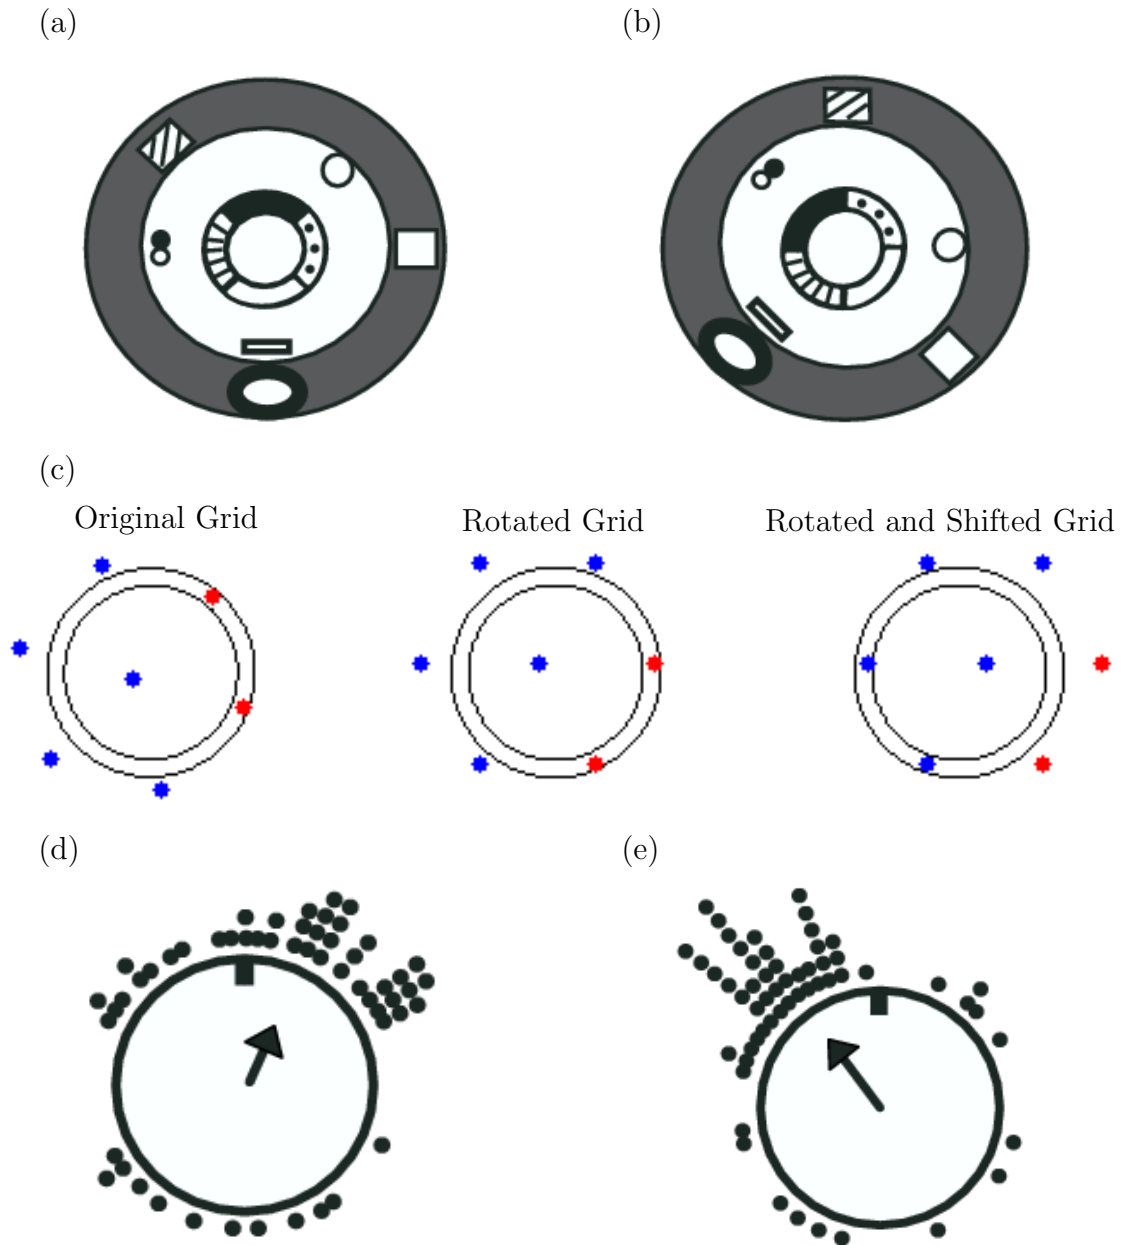


Figure 4.1: Cell response to the double rotation of cues. (a) This picture shows an overhead view of the local and distal cues along the track in their original positions. (b) During the mismatch stage, the distal cues are rotated clockwise, and the local cues are rotated counter-clockwise. The mismatch is 90 degrees for this example. (c) Grid cells rotate with the distal cues, then shift. The dots represent grid points of a grid cell. Before the double rotation, the grid cell spikes when the rat nears one of its two red grid points. After the double rotation, however, these red grid points have shifted off the track and no longer affect the cell. (c) Most place cells in the CA1 rotate with the distal cues or remap. (d) Most place cells in the CA3 rotate with the local cues, in the opposite direction of the rotation of the grid cells. Taken from Lee et al. [19].

of spikes from a single presynaptic cell in the CA3 can cause a burst of spikes in a postsynaptic cell, the CA3 region of the hippocampus has a low threshold for the generation of an epileptic seizure (Traub et al. [28]). This thesis provides the foundation for studying epilepsy because it focuses on synaptic plasticity in the CA3 region of the hippocampus.

4.2 Future Work

Although the model in its current state predicts the formation and backward shift of place fields in the CA3, several improvements must be made to the model to make it useful for applications such as those described in the previous section. The model is computationally limited in the number of cells it can simulate, as described in Chapter 3. Given the network size that I used for this work, the system has almost 0.5 million dynamic variables that must be updated at each timestep. Consequently, the simulated rat of the model currently takes about ten minutes to complete a single lap, a feat that takes an actual rat about eighteen seconds to complete. Motivated by these computational limitations, I will work to improve the speed and capability of the model by implementing the code in PETSc and applying model reduction techniques to the weight matrices.

I will also improve the model by increasing the accuracy of approximating the dynamics of both individual cells and of the weights influenced by synaptic plasticity. The accuracy of the integrate and fire model in predicting spike times of a cell is

sufficient for the double rotation experiment, but other applications may require a more accurate model. Furthermore, while experimenters such as Markram, Bi, and Poo have shown that STDP occurs in the brain, STDP cannot explain the dependence of the weight change on rate, which is also observed experimentally (Sjöström et al. [25]). One alternative to the STDP model is a biophysical plasticity model developed by Shouval. I will implement other plasticity models such as Shouval's to provide the user a greater choice in the plasticity model used.

4.2.1 Improvements to the Model's Speed and Capability

I will improve the speed and capability of the model in two ways. First, I will parallelize the code using the Portable, Extensible Toolkit for Scientific Computation (PETSc) software. This software is an ongoing project among a collaboration of computer scientists that allows the user to implement code in parallel with as much ease as implementing code in C. PETSc was first made available to users in 1995, and the creators continue to make improvements on the software to this day (Balay et al. [3]).

In 2007 Kellems implemented an integrate and fire network of cells in the PETSc environment. Figure 4.2.1 shows the speed-up of his model as he incorporated more processors. I will build on the work that Kellems has done to add synaptic plasticity to his integrate and fire network, taking advantage of the fast matrix-vector products in PETSc.




Figure 4.2: Speed-up in PETSc. Kellems has developed a model for a network of cells using the integrate and fire model. Both the initialization time and the simulation time significantly decreased as he incorporated more processors. I will be adding synaptic plasticity to his model. Taken from [17].

In addition to implementing the code in PETSc, I will also apply model reduction techniques to the weight matrices. Using these techniques I will increase the speed and capability of the model by identifying the weights that significantly affect the system and approximating the dynamics using a lower dimensional weight matrix.

4.2.2 Improvements to the Model's Accuracy

I have implemented the integrate and fire model to approximate the single-cell dynamics. This model is computationally inexpensive, but does not take into account many biological factors, such as the distance between the synapse and the cell body. Kellems

is currently applying model reduction techniques to develop a model for single-cell dynamics that is both accurate and computationally inexpensive (Kellems et al. [18]). In the future I will replace the integrate and fire model I am currently using with his accurate model for single-cell dynamics.

I approximate the dynamics of the weights among CA3 cells using the STDP model developed by Froemke and Dan. While this model explains the weight change in certain situations, it cannot explain all experimental results. For example, Markrum, one of the initial researchers who experimentally observed STDP, found that LTP could only be induced when the frequency of the spike pairs was at least 10 Hz (Markrum et al. [20]). STDP cannot explain this rate dependence. The STDP model is also dependent on the choice of the upper bound for the weight, which has no biological significance.

One alternative to this phenomenological model is a biophysical model for synaptic plasticity based on the biological properties of the synapse. Shouval has developed such a model that can explain the weight change due to both rate-based plasticity and spike-timing-dependent plasticity. His model is based largely on the dynamics of calcium, which is necessary for synaptic plasticity to occur (Castellani et al. [6]).

4.3 Summary

I have built a computational model examining the effect of synaptic plasticity on spatial representation and navigation. This model brings together established paradigms

concerning the firing behavior of grid cells, rate-based plasticity, and spike-timing-dependent plasticity. The model predicts that synaptic plasticity enables place cells in the CA3 to integrate input from grid cells in the MEC to form single firing fields that shift backward along a learned route, as seen in experiments such as the double rotation experiment.

Because the model agrees with experimental results, it has become a reliable platform that I can adapt and apply to a variety of settings concerning synaptic plasticity within a network of neurons. One such application is the study of the onset of epileptic seizures, a network phenomena that largely depends on synaptic plasticity. As I build on the platform developed in this work by implementing the model in PETSc, applying model reduction techniques to the weight matrices, and incorporating other methods for approximating single-cell dynamics and synaptic plasticity, the model will become a useful, reliable tool for investigating network dynamics.

Bibliography

- [1] D.G. Amaral and M.P. Witter. The three-dimensional organization of the hippocampal formation: A preview of anatomical data. *Neuroscience*, 31:571–591, 1989.
- [2] G.A. Ascoli, L. Hunter, J.L. Krichmar, J.L. Olds, and S.L. Senft. Computational neuroanatomy of the hippocampus, April 2009. <http://www.krasnow.gmu.edu/L-Neuron/index.html>.
- [3] S. Balay, K. Buschelman, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L.C. McInnes, B.F. Smith, and H. Zhang. PETSc users manual. Technical Report ANL-95/11-Revision2.1.5, Argonne National Laboratory, 2004.
- [4] G. Bi and M. Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of Neuroscience*, 18:10464–10472, 1998.
- [5] K.I. Blum and L.F. Abbott. A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, 8:85–93, 1996.
- [6] G.C. Castellani, E.M. Quinlan, L.N. Cooper, and H.Z. Shouval. A biophysical model of bidirectional synaptic plasticity: dependence on AMPA and NMDA receptors. *Proceedings of the National Academy of Sciences*, 98:12772–12777, 2001.
- [7] N. Cohen and H. eichenbaum. *Memory, Amnesia, and the Hippocampal System*. The MIT Press, Cambridge, Massachusetts, 1993.
- [8] A.D. ekstrom, J. Meltzer, and B.L. McNaughton. NMDA receptor antagonism blocks experience-dependent expansion of hippocampal place fields. *Neuron*, 31:631–638, 2001.
- [9] M. Franzius, R. Vollgraf, and L. Wiskott. From grids to places. *Journal of Computational Neuroscience*, 22:297–299, 2007.
- [10] R.C. Froemke and Y. Dan. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, 416:433–438, 2002.

- [11] W. Gerstner and W. Kistler. *Spiking Neuron Models*. Cambridge University Press, New York, NY, 2002.
- [12] T. Hafting, M. Fyhn, S. Molden, M. Moser, and E. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436:801–806, 2005.
- [13] E.L. Hargreaves, G. Rao, I. Lee, and J.J. Knierim. Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. *Science*, 308:1792–1794, 2005.
- [14] D.O. Hebb. *The Organization of Behavior; A Neuropsychological Theory*. Wiley, New York, NY, 1949.
- [15] A. Heynen, W. Abraham, and M. Bear. Bidirectional modification of CA1 synapses in the adult hippocampus in vivo. *Nature*, 381:163–166, 1996.
- [16] A.L. Hodgkin and A.F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117:500–544, 1952.
- [17] A. Kellems. The parallel integrate-and-fire simulator, April 2009. <http://www.caam.rice.edu/~tkellems/hippos/PIAF/PIAF.html>.
- [18] A. Kellems, D. Roos, N. Xiao, and S. Cox. Low-dimensional, morphologically accurate models of subthreshold membrane potential. *Journal of Computational Neuroscience*, pages 1–16, 2009.
- [19] I. Lee, D. Yoganarashimha, G. Rao, and J. Knierim. Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature*, 430:456–459, 2004.
- [20] H. Markrum, J. Lübke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275:213–215, 1997.
- [21] R.G.M. Morris. Synaptic plasticity and learning: Selective impairment of learning in rats and blockade of long-term potentiation in vivo by the nmda receptor antagonist AP5. *Journal of Neuroscience*, 9:3049–3057, 1989.
- [22] J. O’Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Oxford, Clarendon, 1978.
- [23] r.W. Rodieck, N.Y.S. Kiang, and G.L. Gerstein. Some quantitative methods for the study of spontaneous activity of single neurons. *Biophysical Journal*, pages 351–368, 1962.

- [24] F. Savelli and J.J. Knierim. A Hebbian model of the formation of place fields from MEC grid inputs. *Society for Neuroscience*, 205.21/EEE14, 2007.
- [25] J. Sjöström, G. Turrigiano, and S. Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32:1149–1164, 2001.
- [26] T. Solstad, E. Moser, and G. Einevoll. From grid cells to place cells: a mathematical model. *Hippocampus*, 16:1021–1031, 2006.
- [27] S. Song, K. Miller, and L.F. Abbott. Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3:919–926, 2000.
- [28] R.D. Traub, R. Miles, R.U. Müller, and A.I. Gulyás. Functional organization of the hippocampal CA3 region: implications for epilepsy, brain waves and spatial behavior. *Network: Computation in Neural Systems*, 3:465–488, 1992.
- [29] M.P. Witter and E.I. Moser. Spatial representation and the architecture of the entorhinal cortex. *Trends in Neuroscience*, 29:671–678, 2006.
- [30] S. Yang, Y. Tang, and R.S. Zucker. Selective induction of LTP and LTD by postsynaptic calcium elevation. *Journal of Neurophysiology*, 81:781–787, 1999.
- [31] X. Yu, J.J. Knierim, I. Lee, and H.Z. Shouval. Simulating place field dynamics using spike timing-dependent plasticity. *Neurocomputing*, 69:1253–1259, 2006.
- [32] X. Yu, H.Z. Shouval, and J.J. Knierim. A biophysical model of synaptic plasticity and metaplasticity can account for the dynamics of the backward shift of the hippocampal place fields. *Journal of Neurophysiology*, 100:983–992, 2008.