

**Network Consistency and Hippocampal Dynamics: Using
the properties of cell assemblies to probe the hippocampal
representation of space.**

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Jadin Cole Jackson

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

A. David Redish, Adviser

September 2006

©Jadin Cole Jackson 2006

Abstract

Neural representations are distributed. This means that information about sensory, motor, or cognitive variables is spread across a population of neurons. To get at the dynamics of information processing in the brain requires an understanding of the neural code not only at the single cell level, but also at the population level. While a specific variable or distribution may be reflected by a population code, measures are needed that reflect the quality of such a distributed representation. This dissertation focuses on this problem through two avenues of research.

Using attractor network models, we developed and characterized two measures of representational quality. One, *coherency*, uses known response properties of individual neurons to compare the current activity pattern of an ensemble with the expected activity pattern. The result is a statistical measure of representational quality that can detect dynamic anomalies in a population's representational state. Another measure, *ensemble consistency*, has the advantage that it requires no explicit assumptions regarding the tuning of neurons in the ensemble, making it ideal for exploring the

dynamics of deep brain structures where neuronal response parameters are controversial or unknown.

With these measures, we then examined the temporal dynamics of the hippocampal representation while rats performed tasks of increasing spatial and cognitive complexity. We examined single cell fluctuations and ensemble (up to 100 simultaneously recorded neurons) modulation to clarify the role of the hippocampus in spatial and temporal processing. Finally, we explored the effects of behavior on the network dynamics of awake sharp waves, hippocampal states that have only recently been characterized yet have long been known to be a source of variability in the spatial firing of hippocampal neurons.

Acknowledgements

This work would not have been possible without the generous help and commitment from many people. The experiments, surgeries, and histology presented here were part of a concerted effort from the technicians and graduate students of the Redish Lab. I would particularly like to thank Chris Boldt, Giuseppe Cortese, and Kelsey Seeland for building hyperdrives in addition to their help running experiments, training rats, and setting up and assisting during surgery. Mallika Arudi, Dan Bernal, and Morgan Little deserve special recognition for their help with managing the massive amounts of data collected during these experiments in addition to their help running experiments, training rats, and setting up and assisting during surgery. I would like to thank Sarah Jutila for preparing the histological slides. I also would like to thank Deborah Bang, Monica Kumar, Susan Nwoke, Steven Sullivan, and Jon Waataja for their help with data collection as well as for helpful discussions.

Certain individuals contribute invaluable to the development of one's ideas through their discussions, questions, and suggestions. Adam John-

son, Beth Masimore, and Neil C. Schmitzer-Torbert are three of these such individuals to whom I am greatly indebted for their influence on me intellectually, for their commiseration, and for their friendship.

I would also like to thank Robert delMas for statistical help, Duane Nykamp for helpful mathematical discussion, and Josh Gordon for helpful suggestions with data analysis.

Finally, I would like to thank my thesis committee for all their time and effort: Dick Poppele, John F. Soechting, James Kakalios, and A. David Redish. My advisor David Redish deserves special recognition for enduring the wiley antics of the author of this dissertation and for his supportive guidance through the various challenges that family life imposes on the careers of his graduate students.

This work was supported by funding from NIH grant R01-MH06829 and by an NSF-IGERT training grant #9870633.

Table of Contents

1	Introduction	1
1.1	Overview	1
2	Distributed Representations	7
2.1	Distributed Representations	7
2.1.1	Parallel Distributed Processing	7
2.1.2	Content Addressable Memories	10
2.1.3	Cell Assemblies	11
2.2	Reconstruction Techniques	12
2.3	Coherent Representations	15
2.3.1	What is a <i>Coherent representation</i> ?	15
3	Hippocampus	19
3.1	Hippocampus	19
3.1.1	Memory	20
3.2	Overdispersion	30
3.2.1	Inhomogeneous Poisson Process	31
3.2.2	Doubly Stochastic Inhomogeneous Poisson Process	33
3.2.3	Olypher Network Model	34
3.3	Goal Tasks	35
3.3.1	Hippocampal Lesions	35
3.3.2	Over-Representation of Goal	36
3.3.3	Effects on Place Cell Stability	37

3.3.4	Effects on Place Cell Temporal Variability	38
3.4	Local Field Potentials	38
3.4.1	Generation	39
3.4.2	Hippocampal Anatomy	40
3.4.3	Hippocampal LFP Types	46
3.4.4	Hippocampal LFPs Depend Strongly on Electrode Position	53
3.5	Reactivation	54
3.5.1	Slow Wave Sleep	56
3.5.2	Paradoxical Sleep (REM)	60
3.5.3	Awake Reactivation	61
3.5.4	Theories	62
4	Ensemble Measures	71
4.1	Introduction	72
4.2	The simulations	74
4.3	Simulation Methods	77
4.3.1	The Network Model	78
4.3.2	Constructing the Training Set	83
4.3.3	Reconstruction	84
4.4	Coherency	85
4.4.1	The mathematics of the coherency measurement	86
4.4.2	Coherency differentiates representational quality (Is- sue 1)	91
4.4.3	Coherency can detect dynamic changes in network activity (Issue 2)	92
4.4.4	Coherency can be used to detect the resolution of ambiguity (Issue 3)	93
4.5	Ensemble Consistency	95
4.5.1	The Mathematics of Ensemble Consistency	96
4.5.2	Statistical Justification of EC	105
4.5.3	EC differentiates representational quality (Issue 1)	107

4.5.4	EC can detect dynamic changes in network activity (Issue 2)	108
4.5.5	EC can be used to detect the resolution of ambiguity (Issue 3)	110
4.6	Discussion	112
5	Methods: Experimental	117
5.1	Introduction	117
5.2	Subjects	118
5.3	Behavioral training	118
5.4	Surgery	122
5.5	Electrode Positioning	123
5.6	Recordings	123
5.6.1	Neurophysiology	124
5.6.2	Behavioral tracking and Behavioral Control	125
5.7	LFP analysis	127
5.7.1	SW detection	127
5.7.2	Theta Detection	128
5.8	Unit Recording	130
5.8.1	Clustering	131
6	Overdispersion	134
6.1	Introduction	135
6.2	Tuning Curve Construction	137
6.3	Dispersion measurement	138
6.4	Replication of overdispersion result	139
6.5	Task dependence: LT vs. OF vs. OFG	141
6.6	Cell Assemblies	146
6.6.1	Cell pair correlations	147
6.6.2	Testing Ensemble Modulation: Overdispersion vs. Coherency	152

6.6.3	Correlation of dispersion z-score with the number of passes through a place-field	157
6.6.4	Correlations between Dispersion and Behavioral En- tropy	158
6.6.5	Correlations with Speed	161
6.7	Splitting Representational Maps	162
6.7.1	Map Splitting Method	163
6.7.2	Results of Map Splitting	168
6.8	Conclusions	173
7	Sharp-wave emission	183
7.1	Introduction	184
7.2	New Analyses	187
7.2.1	Co-firing Reactivation Analysis	188
7.2.2	Pattern Identification	189
7.2.3	Randomized Ensemble Controls	192
7.2.4	Behavioral entropy	193
7.3	Results	194
7.3.1	Linear Track.	196
7.3.2	Two-Dimensional Tasks.	198
7.3.3	Comparisons across tasks.	201
7.4	Discussion	209
8	Discussion	212
8.1	Coherency	212
8.1.1	Reconstruction versus Coherency: Implications for understanding complex network states.	212
8.1.2	Representational versus Non-Representational Co- herency	215
8.1.3	Bayesian Methods and Uncertainty	220
8.1.4	Coherency in the Brain	222
8.2	Overdispersion	224

8.2.1	Reference Frames	224
8.3	Experience Dependent Sharp Waves Ripples and Awake Reactivation	230
8.3.1	Sharp waves in CA1 are highly regulated network events	230
8.3.2	Brain State and Sharp Wave Ripple Emission	231
8.3.3	Replay generation	233
8.3.4	Forward versus Reverse Replay	237
9	Conclusion	245

Chapter 1

Introduction

“...the intuitions needed for understanding biological information processing are not easily available. Only by wresting them from actual experience does one gain a feel for what questions need to be asked, and develop a language in which to ask them.” (Marr, 1975)

1.1 Overview

How does the brain work? The question can be addressed at different levels. For example, at a gross level, what brain structures are required for memory; or, on the molecular level, how do individual neurons change their biochemical properties to encode memories? In this dissertation, the function of the brain is studied using electrophysiological (recording the electrical activity emitted by neurons), and computational techniques to

examine specific aspects of how ensembles of neurons within a brain structure collaborate to encode memories, behaviors, and external variables. Specifically, this work seeks to differentiate between network modulation of neural activity and noise observed in the activity of single cells. This involves recording the electrical activity of many neurons (up to 100 neurons recorded simultaneously, yet separately) while an animal is awake and operating under (near) normal conditions. The work in this dissertation falls into two major categories: (1) recording neural ensembles in the hippocampus of awake, behaving animals while concurrently studying their navigational behavior and (2) development of new statistical tools for quantifying neural activity at the multi-neuron level. Together, these two aspects of my research address new questions about the dynamics of information processing in the brain and how behavior affects changes in the brain's networks.

First, a network model with strong interactions between units is used to explore concepts of the consistency of neuronal activation within a network. We characterize and discuss two types of ensemble measure: representation based and pattern based measures. Both types are effective at detecting a lack of self consistency in ensemble firing patterns. We then use these concepts to explore dynamic phenomena in the hippocampus.

The hippocampus is a brain structure considered to be crucial for learning and memory in mammals (for review see Redish, 1999). Humans with hippocampal damage show marked short-term memory loss especially

after temporary distractions. Additionally, these hippocampal patients have difficulty storing new memories. The rodent hippocampus is most noted for the strong spatial selectivity of its pyramidal cells, or *place cells* (O'Keefe and Dostrovsky, 1971). Each neuron displays electrical activity (or *fires spikes*) only in particular regions in an environment. Past work has revealed many aspects of hippocampal network dynamics including the modulation of place cell activity by well-known brain states such as the theta rhythm (a strong 6-10 Hz electrical oscillation prominent during movement) and sleep (Vanderwolf, 1971; O'Keefe and Nadel, 1978). There are also explicit theories about how the manner in which information is processed by the hippocampus should modulate place cell firing (O'Keefe and Nadel, 1978; Buzsáki, 1989; Redish, 1999). In sum, the hippocampal representation of space allows us to probe specific aspects of the processing of spatial information in the brain.

For example, the spatial specificity of the hippocampus allows us to examine the effects that varying the cognitive demands in navigational tasks have on hippocampal processing. For instance, it is known that the hippocampus is required for navigation to un-cued or hidden goal locations as well as solving other spatial problems (Milner, 1970; O'Keefe and Nadel, 1978; Kesner and Novak, 1982; Morris et al., 1982; Squire, 1992; Redish, 1999). In fact, hidden goal locations that are constant from day to day are over represented by hippocampal neurons (Hollup et al., 2001a). Additionally, increasing the spatial requirements of a task and enriching

the variety of an animal's spatial experiences increases the stability of the hippocampal spatial representation (Kentros et al., 2004). Finally, adding a spatial requirement to a foraging task improves the reliability of the spatial firing of hippocampal cells (Olypher et al., 2002).

If the hippocampus represents an environment in multiple ways depending on a task's demands, then switching between these spatial maps, or *frames of reference*, should result in detectable changes in the dynamics of the hippocampal network. Modifying the experimental environment of a task can have a variety of effects on the spatial map the hippocampus uses to represent a task depending on the magnitude and type of modification (Quirk et al., 1990; Markus et al., 1995; Anderson and Jeffery, 2003; Knierim, 2002; Knierim and Rao, 2003; Lee et al., 2004b; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004; Leutgeb et al., 2005; Wills et al., 2005). Even changing the behavioral requirements of a task within a session affects the spatial map the hippocampus uses to represent a task (Markus et al., 1995; Redish et al., 2000; Rosenzweig et al., 2003). There is some evidence that this reference frame switching may be a regular phenomenon even without experimental behavioral manipulation (McNaughton et al., 1983; Wood et al., 2000; Fenton and Muller, 1998; Láncsý et al., 2001; Olypher et al., 2002). For example, this has been observed in linear environments where hippocampal map depends on the direction of travel (McNaughton et al., 1983; O'Keefe and Recce, 1993). Little is understood how this relates to two-dimensional tasks where little

directional firing is observed (Muller et al., 1994; Redish, 1999).

Fenton and Muller (1998) found that while place cells show spatially reliable place fields, they exhibit much more temporal variability than would be predicted by a simple Poisson process. This variability, or *overdispersion*, is task dependent (Olypher et al., 2002). It has been suggested that this overdispersion may be the result of the rat switching reference frames at a mean rate of 1 to 2 times per second (Lánksy et al., 2001; Olypher et al., 2002). We replicated these findings of excessive variance in place cell discharge and its task dependence. We found that on a task with known reference-frame switching (the linear track), splitting by representational state (direction of travel) results in greatly reduced variability. We then present evidence that, contrary to previous research, there are significant local interactions between some neurons with overlapping place-fields, but that overall this effect is weak. We show that these effects are correlated with local behavioral variability and the level of repetition. We then demonstrate that it is possible to extract different spatial firing maps based on a separation of ensemble firing patterns. These firing maps are indicative of separate reference frames being used in an alternating pattern on the task. The variability of firing patterns within these reference frames is greatly reduced, approaching the variance expected by a single point-process stochastic model.

Finally, we explore the dynamics of a previously ignored brain state in the awake rat that has long been a known behaviorally dependent source

of neuronal firing variability in the awake rat: the *sharp wave ripple*. We demonstrate that neuronal firing during sharp wave ripples is indeed well organized and depends on the nature and repetition of an animal's spatial experiences. Together, these results shed new light on dynamics of distributed information processing in the hippocampus.

Chapter 2

Background:

Distributed Representations

2.1 Distributed Representations

Neural representations are distributed. This means that information in the brain's networks is spread across a population of neurons. The significance of this massively parallel processing architecture has important repercussions not only for the way information is handled by an organism, but also for the way we approach the study of the brain.

2.1.1 Parallel Distributed Processing

Parallel distributed processing (PDP) architectures have a variety of advantages over current computing platforms. For instance, let us consider

the question posed in the overview: How does the brain work? One reason for asking this question is because even some of the simplest organisms have incredible abilities to interact with their environment and accomplish tasks that are far beyond the most state-of-the-art machine learning systems in the artificial intelligence (AI) and robotics fields.

How does the brain do this with units that are that far slower than modern transistors (by at least six orders of magnitude)? First of all, in contrast to the primarily serial processing architecture of modern computers, the brain can process incoming information by distributing information over vast networks of neurons that can process the information in parallel. Computations are performed both within a neuron and across the network as neurons interact within and across networks to process inputs.

To illustrate the power of the brain, consider an every-day task like retrieving a food item such as a cucumber from the back of a cluttered refrigerator. For most people, this is accomplished with apparent ease, but simply describing the process is difficult and programming a machine to accomplish this task would currently be intractable. First, the item is identified visually against all odds (by computer vision standards) in that: 1) the entirety of the cucumber is not immediately visible, being blocked by numerous items in the way such as the milk, juice, ketchup, etc.; 2) the last time you saw a cucumber was perhaps days ago; 3) the individual variability in the shape and size of cucumbers is enormous; and 4) the color variation which depends on the incident light is incredible. Second,

ignoring the 3 to 4 degrees of freedom at each finger, you will probably use at least 6 degrees of freedom in your arm alone to position your hand at the cucumber. In performing this reach to the back of the refrigerator, you must find the following computationally non-trivial solutions: 1) a path and orientation of the entirety of the arm to minimize interference with the objects cluttering the refrigerator compartment so as to refrain from spilling the milk or breaking a jar of pickles; 2) a path and orientation of the cucumber-arm assembly for retrieval (a new addition to your art's weight, inertia, and cumbersome geometry). The simplest of these tasks could take supercomputers days to solve, even for a well defined set of initial conditions; yet you accomplish this with little difficulty in less than 10 s.

The distributed networks of the brain allow visual input, tactile input, proprioceptive input, and motor output to be processed simultaneously with each influencing the other in real-time. In the context of our cucumber example, brushing an object that was hidden behind the milk (a tactile experience) will lead to a rapid compensation in limb trajectory and geometry, with simultaneous visual and proprioceptive guidance. This "real time", processing of information-rich input from multiple sources is only one of the numerous advantages of parallel distributed processing. The highly interconnected nature of distributed systems leads to many emergent properties such as content addressability.

2.1.2 Content Addressable Memories

Another powerful attribute of distributed networks is that they are “content addressable”. This means that activation of a selective part of the network can lead to activation of elements associated with that part. More specifically, a particular input from one sensory modality may lead to the activation of units related to other sensory modalities. For instance, given the sight of the cucumber it is possible to retrieve the taste of cucumbers, the expected feeling of grasping the cucumber, perhaps recipes that call for a cucumber, even the name “cucumber”. Thus, these memories of attributes associated with a given item are linked by the distributed nature of the brain’s networks.

The important feature of content addressability in the context of this dissertation, is that the memory retrieval process activates broad patterns of units associated during the encoding of stored values. Thus, recall is in essence a co-activation of an assembly of neurons. The fidelity of the memory lies in the binding and successful recruitment of all neurons in the encoding assembly as well as the preservation of the relative activation pattern of these neurons. This property allows us to make specific analytical statements of what a “coherent” representation of a given state should look like.

2.1.3 Cell Assemblies

Cell assemblies as first hypothesized by Hebb in his landmark book, are neurons that have been repeatedly co-active and therefore have by mechanisms of plasticity become “associated” so that they are more likely to fire together (Hebb, 1949). If neurons are part of the same cell assembly they should be bound together temporally and possibly result in a transient “reverberation” (Hebb, 1949). Hebb (1949) went so far as to suggest that this may be the mechanism of consciousness or perception. These ideas continue to re-emerge due to assembly-like phenomena in the hippocampus (Wilson and McNaughton, 1994; Buzsáki, 1996; Kudrimoti et al., 1999; Harris et al., 2003; Harris, 2005). Harris et al. (2003) presented data suggesting the existence of cell-assembly dynamics in the awake hippocampus. The authors demonstrated that while the spatial location of an animal predicted the firing of hippocampal neurons, using the firing of other neurons in the ensemble greatly improved prediction of the firing of the neuron of interest. Their analysis revealed grouping of neuronal assemblies that tended to fire together. The firing of these assemblies could interrupt spatial responses. This was taken by the authors as evidence for the internal cognitive processes conceptualized by Hebb (1949) in his theory of cell assemblies (Harris et al., 2003; Harris, 2005).

These data also fit with recent observations of non-spatial network coding in the hippocampus reported by Lin et al. (2005). Using a multiple components analysis they found that specific patterns of activation across

the hippocampal network could be used to reliably discriminate between salient task events such as air-puffs, shaking the environment, and sudden downward displacement of the animal's cage. They termed these "neuronal cliques". (Lin et al., 2005)

2.2 Reconstruction Techniques

Given the properties of distributed systems, one may ask how we extract meaning from the activation of units within the system. In the case of sensory systems, this may amount to asking what stimulus is most likely generated by a particular network activation pattern. In a sense, the general approach amounts to understanding the way individual units within the network code for stimuli and invert this encoding computation to decode the joint likelihood of observing a stimulus given the population response (Rieke et al., 1997) *.

One of the first methods used to extract a behavioral variable from the activity of neurons was in motor cortex with the application of vector reconstruction (Georgopoulos et al., 1983). Assigning a vector to each neuron based on its mean preferred stimulus and its current firing rate, vector reconstruction is simply the vector mean of these firing rate weighted neuronal tuning vectors (Mardia, 1972; Batschelet, 1981). This is also known

* Actually, in their discussion of reconstruction techniques, Rieke et al. (1997) focused primarily on the reconstruction of a stimulus from the temporal pattern of spikes emitted from a single neuron. However, these concepts can be generalized and extended to populations of neurons as will be seen in the methods discussed in the following paragraphs.

as the *population vector* (Georgopoulos et al., 1983, 1988).

Assuming a directional representation over circular, spherical, or hyperspherical angles, let \vec{v}_k be a unit vector pointing in the preferred direction ϕ_k of neuron k . In the two-dimensional plane we would write: $\vec{v}_k = \langle \cos(\phi_k), \sin(\phi_k) \rangle$. The reconstructed vector \vec{R} is then the vector mean of the preferred vectors of all units with lengths proportional to their firing rates F_k :

$$\vec{R} = \sum_k F_k \cdot \vec{v}_k \quad (2.1)$$

Since this method is sensitive to the uniformity of the distribution of preferred directions ϕ_k , one variation on the vector reconstruction method is the optimal linear estimator method employed by Salinas and Abbott (1994). This method determines a mapping of the vectors \vec{v}_k to a new set of vectors $\vec{\lambda}_k$ that minimize the reconstruction error (the difference between the actual and reconstructed directions). The result is similar to population vector reconstruction:

$$\vec{R} = \sum_k F_k \cdot \vec{\lambda}_k \quad (2.2)$$

Another commonly used reconstruction technique is that of template matching (e.g. Wilson and McNaughton, 1994). Define a firing rate vector (unfortunately also referred to as a *population vector*) $\mathbf{f}(t) = \{f_1(t), f_2(t), f_3(t), \dots, f_N(t)\}$, where $f_i(t)$ is the firing rate of neuron i at time t , and N is the number of neurons in the ensemble. In template matching, this firing rate vector is matched against the mean vectors of ensemble firing calculated for

each value of the behavioral variable. This “matching” can be based on maximum correlation, minimum error, etc. The best match is taken as the reconstructed variable.

Recently, there has been a boom in the application of statistical inference techniques based on Bayes rule to reconstruct behavioral and stimulus variables from neuronal activity. The core premise of Bayesian reconstruction is that if one knows the probability of observing neuronal activity for given a behavioral variable of interest, it is possible to use Bayes rule to derive the probability of seeing a particular behavioral variable given current neuronal activity. We write this as:

$$P(X|S) = \frac{P(S|X)P(X)}{P(S)} \quad (2.3)$$

where $P(X|S)$ is the probability of seeing a particular behavioral variable X given current neuronal activity S , $P(S|X)$ is the probability of observing neuronal activity S given behavioral variable X was observed, $P(X)$ is the probability of observing behavioral variable X , and $P(S)$ is the probability of seeing activity S . The reconstructed value X is then taken as the value that maximizes $P(X|S)$. An example of an application of this method to extract an animal’s spatial location from hippocampal ensembles is provided by Zhang et al. (1998).

2.3 Coherent Representations

While the development of ensemble based reconstruction methods such as those described above has allowed us to probe more deeply into the brain's processing of behavioral information, we run the risk of assuming that an animal's brain rigidly adheres to representing the present behavioral status of the animal. In doing so, reconstruction errors are viewed as "noise in the system", and we forget the cognitive questions that are fundamental to our understanding of the brain's inner workings. For instance, what is recall or confusion and how does the brain represent competing values in ambiguous situations?

To answer these questions, we need to consider how units within a network function together to form a coherent representation i.e. one that is internally consistent across all units.

2.3.1 What is a *Coherent representation*?

A *coherent* or self-consistent representation is one in which the firing of all neurons in a network conforms to some model. For instance, if one records from an ensemble of motor cortical cells, one possible model of the network would be to assume that the firing of each neuron is tuned to the direction of movement. This tuning, if it exists, should dictate the interactions of the neurons in the network. If the network is representing a particular direction, all neurons with any tuning to that direction

should be firing to some degree specified by their respective tuning curves and neurons that are tuned to very different directions should be responding very little. In other words, neurons with similar preferred directions should respond similarly if the network is responding in a manner consistent with the data set used to construct the neuronal tuning curves. If this is not true, there is a fundamental difference between your model and the current status of the network. This principle allows for the formulation of a measure of the *coherency* or *self-consistency* of a neural ensemble. (See Figure 2.1.)

Figure 2.1 B, C, and D show three hypothetical states for a network made up of neurons with tuning curves shaped like the one depicted in Figure 2.1 A, but centered at even intervals along x . The behavioral variables \hat{x}_1 and \hat{x}_2 are shown for reference. The pattern in B is consistent with behavioral variable \hat{x}_2 but not with \hat{x}_1 . A reconstruction algorithm would yield value \hat{x}_2 . If the actual value was \hat{x}_1 , then reconstruction error $|\hat{x}_2 - \hat{x}_1|$ would be high even as the network state is internally consistent. The left mode of the pattern in C is consistent with behavioral variable \hat{x}_1 but neither mode is consistent with \hat{x}_2 . A vector based reconstruction algorithm would yield value \hat{x}_2 , while template matching or Bayesian methods would yield \hat{x}_1 or the right peak depending on the noise in the system. If the actual value was \hat{x}_1 , then reconstruction error $|\hat{x}_2 - \hat{x}_1|$ would either be low or high depending on the reconstruction method and the noise in neuronal activity. However, neither reconstruction measure would reveal

the underlying representational ambiguity. The state in D is not consistent with either behavioral variable \hat{x}_1 or \hat{x}_2 . However, each reconstruction method would yield a value such as \hat{x}_1 or \hat{x}_2 even though the underlying state is complete confusion. Each of these scenarios suggests very different cognitive processes are occurring in this brain network, accessing these through an appropriate measure of ensemble internal consistency is

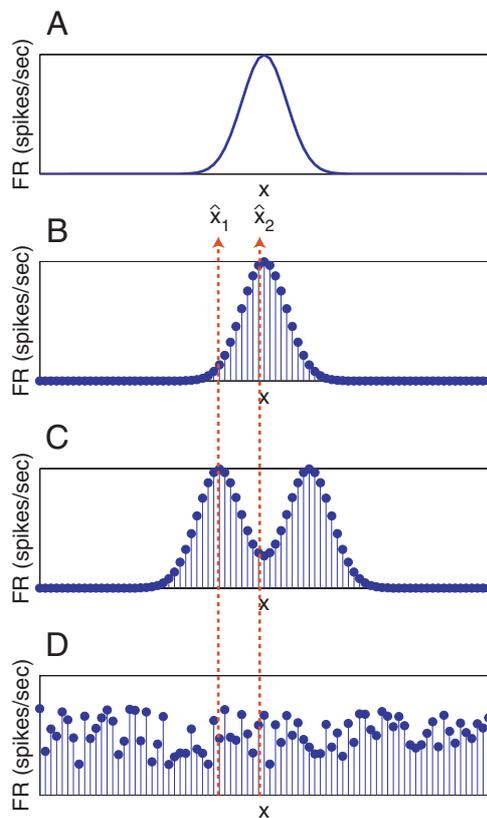


Figure 2.1: Self consistency. (A) An example unimodal tuning curve. The stimulus or behavioral variable is on the x -axis with firing rate represented along the y -axis. (B) A “coherent” network firing pattern. The stimulus or behavioral variable is on the x -axis with firing rate of each neuron represented along the y -axis. Each line represents the location of a neuron’s preferred stimulus, with height equal to the neuron’s firing rate. If each neuron in a network had unimodal tuning curves identical to the neuron represented in A but with the peak firing occurring at a different preferred stimulus x , then when the preferred stimulus of the neuron in A is presented, this is the expected network firing pattern. This pattern is consistent with behavioral variable \hat{x}_2 but not with \hat{x}_1 . (C) A bimodal representation would represent an ambiguous or incoherent state of the network described in B, since the unimodal tuning curves would predict only one mode of activity should be possible of a single stimulus x and that outside this mode neurons should be silent. One mode is consistent with behavioral variable \hat{x}_1 but neither mode is consistent with \hat{x}_2 . (D) As in C this representation would represent a confused or incoherent state of the network described in B, since the unimodal tuning curves would predict a prominent mode of activity and that outside this mode neurons should be silent. This state is not consistent with either behavioral variable \hat{x}_1 or \hat{x}_2 .

one primary aim of this dissertation.

Chapter 3

The Hippocampus as Our Model System

3.1 Hippocampus

In order to explore the biological validity of the concepts of ensemble measurement that we have introduced, a biological system that is both well-understood and yields itself to large scale experimental recordings is needed. As we will discuss in this chapter and the next, the hippocampus is a brain structure that fits this description very well. The hippocampus has been the subject of numerous anatomical studies and is arguably one of the most thoroughly investigated regions of the brain due to the incredible organization of its structural features and connectivity. Additionally, with the availability of advanced chronic electrophysiological tech-

niques the laminar structure allows for precise localization of electrodes for high-yield recordings of up to 100 neurons or more with as few as 10-12 tetrodes.

Beyond these issues, the hippocampus is a brain structure critical for the formation of specific types of memory, therefore damage to this structure is particularly debilitating. So, characterizing the function of this region is an important step toward understanding the adverse effects of many neurological diseases including Alzheimer's Disease, schizophrenia, autism, and many others.

3.1.1 Memory

While the unique structure of the hippocampus has drawn the attention of many neuroanatomists throughout the years including Ramon y Cajal and Lorente de No, one of the first major indications of its functional significance came after treating a patient for intractable epilepsy (Scoville and Milner, 1957).

On September 1st, 1953 at the age of 29, a patient referred to as H. M. underwent surgery for bilateral resection of his medial-temporal lobes as a dramatic effort to ameliorate his pharmacologically intractable and increasingly debilitating epileptic attacks. Upon recovery, it was found that this young man of average intelligence had lost the ability to form new memories. Besides displaying this profound *anterograde amnesia*, H. M. was unable to recall many events from the months preceding the surgery

and was even unable to remember the death of an uncle that had occurred three years earlier. In spite of these gross memory disturbances, the patient was able to remember many remote childhood events. Post-recovery neurological testing revealed little else of concern and in fact demonstrated a slight increase in H. M.'s I.Q. (from 103 on the Wechler I.Q. scale pre-operative, to 118 nine years post-operative). (Scoville and Milner, 1957; Milner, 1970).

The dramatic amnesic effects of medial-temporal lobe removal on H. M. triggered an immediate re-examination of other psychiatric patients that underwent similar surgical procedures (Scoville and Milner, 1957). Since then, comparing the results of a variety of lesion studies has revealed much about the specific effects of hippocampal damage and the relationship of the hippocampus to the surrounding cortices adding much insight into the function of the normal hippocampus (Scoville and Milner, 1957; Milner, 1970; Squire, 1992).

Place

The Place Cell. The studies showing that damage to the hippocampal formation and its inputs generated intense interest in the hippocampal formation as a locus of memory. Yet it was the discovery of the *place cell* in 1971 that gave the first glimpses into the nature of hippocampal information processing (O'Keefe and Dostrovsky, 1971). After chronically implanting rats with an electrode microdrive, a simple neurological work-

up revealed that some rarely firing neurons in the hippocampus demonstrated spatial responses and sometimes directional responses as the animal was passively moved around the testing platform (O'Keefe and Dostrovsky, 1971). Although this paper hardly presented convincing evidence for the navigational and spatial functions now attributed to the hippocampus, it was later followed by another paper and a landmark book by the same first author that conclusively demonstrated the spatial information present in the firing of hippocampal neurons (O'Keefe, 1976; O'Keefe and Nadel, 1978). Since then, the spatial tuning properties of hippocampal pyramidal neurons have been thoroughly investigated (See Redish, 1999 for review).

Place Maps and Reference Frames

Since an ensemble of hippocampal place cells will generally tend to have a stable pattern of place-fields distributed uniformly throughout an environment, this is taken as the neural instantiation of a spatial map or reference frame (See O'Keefe and Nadel, 1978; Redish, 1999 for review). The spatial map as represented by CA1 pyramidal cells degrades gracefully in response to environmental manipulations such that small changes in the environment result in small changes to the pattern of cells activated and their preferred locations of activation (See Redish, 1999 for review).

The concept of *partial remapping* emerged in a study by Quirk et al. (1990). The authors used two environments, a circular and a square arena, and a light-dark-light followed by a dark-light session to study the effect

of initial darkness upon entering an environment versus darkness falling after light exposure. In the light-dark-light portion of the session, cells maintained a high spatial correlation in their firing responses across the 3 conditions. However, in the dark-light portion of the session, many cells remapped when animals entered an environment in the dark: some had new place fields, some lost place fields, some retained similar place fields, and some gained a place field after previous silence. This partial remapping was shown for a pair of cells where one maintained its place field and a previously silent cell gained a place field when started in the dark. These remappings often persisted even after the lights were turned on. Remapped fields could be reinstated if the animal was subsequently reintroduced to the environment in the light. This experiment demonstrates that place cells are not simply controlled by sensory cues, yet they are stable in the dark over long periods. The authors suggested that place-cells may be controlled by a path-integrator, that this persistent dependence on initial conditions during only the trial is similar to episodic memory or working memory. They coined the term *partial remapping* and suggested that the network will remap to the extent that two environments are different (Quirk et al., 1990).

It was later demonstrated that partial remapping can result from modification of task behavioral parameters (Markus et al., 1995). A simple cue environment and multiple cue environment was used to test the effect of sensory “richness” on place cell firing as well as multiple tasks: rats

ran on an open field (either randomly foraging for pellets or in a directed manner) or on a plus maze. Markus et al. (1995) found that within the open field, simply changing the task from foraging to directed running (tapping on 1 of 4 locations in sequence) for food caused place fields to partially remap. The remapping effect for foraging versus directed search on a plus maze was much weaker than in the open field, but some remapping did occur. It was shown that the remapping gradually formed across and ensemble and tended to correlate with the animal acquiring the new directed behavior. Changing the apparatus with a session to an 8-arm radial maze also had a total remapping effect, even though the animal was in the same room. It was also shown that place-fields were less directional on the open-field than on the 8-arm radial maze (Markus et al., 1995).

This partial remapping effect has recently been demonstrated in ensembles of CA1 neurons suggesting that the CA1 representation is modified proportionally to magnitude of change in the environment (Anderson and Jeffery, 2003; Knierim, 2002; Knierim and Rao, 2003; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004). It was demonstrated that in similar environments, changing the wall color or scent (e.g. lemon or vanilla) can cause a partial remapping of the hippocampal network (Anderson and Jeffery, 2003). Likewise, in experiments where distal and local landmarks (e.g. cues on the wall and features of the surface of an annular track, respectively) are moved in opposite directions to create a mismatch, some neurons follow the local cues, some follow the distal cues, and some place

cells split their place fields (Knierim, 2002). In some neurons, this behavior was dynamically dependent upon the experience of the animal during the mismatch trial such that spatial firing on early laps was different from firing on late laps (Knierim, 2002). In a subsequent experiment, it was shown that translations in the track's location can result in both reversible and non-reversible alterations in a place-cell's spatial firing field (Knierim and Rao, 2003).

More recently, a similar mismatch paradigm to that used by Knierim (2002) yielded a dissociation between CA3 and CA1 responses to local and distal cue manipulations. As in the study by Knierim (2002), CA1 neurons demonstrated a variety of responses including dropping place fields, adding place fields, following local cues, following distal cues, or providing an intermediate firing pattern such as splitting a firing field to follow both local and distal cues. CA3 neurons, however, favored the local cue configuration while rarely following distal cues, though ambiguous firing patterns and added or deleted place fields were also observed (Lee et al., 2004b). As a result, the correlation of the CA1 map between cue and mismatch trials was very poor, while the correlation of the CA3 map between cue and mismatch trials was merely shifted spatially (Lee et al., 2004b). This was independently confirmed by two other laboratories (Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004) *. Leutgeb

*Vazdarjanova and Guzowski (2004) and Leutgeb et al. (2004) each also independently demonstrated quite conclusively that the CA3 population represents an environment with a much sparser code than does CA1, in that fewer neurons are activated in CA3

et al. (2004) exposed animals to similar (same arena or smaller arena with the same geometry) or different local environments in either similar or different experimental rooms and found that the spatial pattern of activation in CA1 neurons increasingly differentiated between the different combinations of environmental manipulations (Leutgeb et al., 2004). The CA3 map was only the same when the animal was exposed to the same local environment in the same room, any other combination yielded an independent remapping such that there was nearly zero overlap between maps and the correlation between maps was zero (Leutgeb et al., 2004). Vazdarjanova and Guzowski (2004) used immediate-early gene activation to differentiate between the effects of temporally distinct environmental experiences on the CA1 and CA3 neuronal populations. They compared groups of animals that either experienced the same environment twice, experienced two completely different environments in different rooms, or experienced some gradation in between. Depending on where along this continuum of changes in the local and distal environmental cues and animal's experience was, the difference in the CA1 neuronal population activated on the first task from the population activated on the second task was consistent with a graded partial remapping of ensembles. Not until there was a total change in the local and global environment was there a fully independent set of neurons activated by the second environment (i.e. total remapping). The CA3 population, on the other hand, demonstrated an independent set of neurons activated by the second environment than in CA1 for a given environment.

all or none mapping of an environment such that similar environmental configurations activated the same population of CA3 neurons, and a total change in the local and global environment triggered an independent population (Vazdarjanova and Guzowski, 2004).

Taken together these studies suggest that the CA3 representation is a stable all-or-none mapping of a pattern associated with an environment. Considering that CA1 gets input from CA3 as well as entorhinal cortex (see Section 3.4.2) and can form a place representation in the absence of CA3 input (Mizumori et al., 1989; Brun et al., 2002), this suggests that the CA1 representation represents the similarity between the entorhinal cortical spatial input and the CA3 localization of the animal to a particular environment. The partial remapping reported earlier (Quirk et al., 1990; Markus et al., 1995; Anderson and Jeffery, 2003; Knierim, 2002) must therefore reflect identified similarities between environmental and task parameters.

Leutgeb et al. (2005) clarified these results by demonstrating that modifying local features (the shape and color of the behavioral enclosure) while keeping place constant primarily generated changes in the maximal firing rates of CA3 neurons without changing the actual spatial map. This modification of spatial firing rates appears to decorrelate the map used to represent one enclosure from the map used to represent another enclosure even though the neurons are firing in the same spatial locations; this is termed *rate remapping* (Leutgeb et al., 2005). Leutgeb et al. (2005) suggest that rate

remapping is an independent ensemble code that can be used to differentiate altered cue configurations while preserving the coding for a particular spatial location. The CA1 ensembles demonstrated a similar phenomenon that appeared to be coupled to the partial remapping discussed earlier (Leutgeb et al., 2005). In his commentary on this study, Buzsáki (2005a) likens this phenomenon to the human experience of simultaneous representation of nested environments such as representing our location in a city reliably whether we are in a bus or a sports car. This process may underlie the task dependent remapping within an identical environment as observed by Markus et al. (1995).

Switching Reference Frames

These data on reference frame formation (i.e. the generation of a spatial firing map across a hippocampal ensemble) suggest that the hippocampus can switch the reference frame used to solve a task depending on the task requirements. This is consistent with experimental data that required such a switch of reference frames for proper task performance (Redish et al., 2000; Rosenzweig et al., 2003). In these studies animals were contained in a box fixed to a movable linear track. At the beginning of a trial, the box was opened and animals ran down the track towards a barrier that was fixed in a location with respect to the room. For an animal to receive reward, it had to pause in a location that was always a fixed distance from the barrier and therefore also consistent with respect to the room. Between trials, the box-

track was moved with respect to the room such that the goal/reward location was always a different distance from the release point at the start of the trial. Animals tended to use one ensemble firing map for representing their location with respect to their distance from the box (the *box-reference frame*) and another for representing their location with respect to the room (the *room-reference frame*). When animals correctly found the reward location, they had switched to the room-reference frame. On trials when this switch did not happen, the animal ran past the goal. This ensemble reference frame switch was detected using a coherency ratio which will be discussed later (Redish et al., 2000; See also Chapter 4). The switch in coherency ratio was more closely related to the time the animal had spent out of the box than to any other task parameter explored (Redish et al., 2000). Interestingly, aged animals had to run farther along the track before switching reference frames (Rosenzweig et al., 2003). These data suggest that the hippocampus can switch reference frame while performing a task that requires such a switch, and that this switch may be necessary for successful spatial behavior.

Finally, Touretzky and Redish (1996) presented a computational hippocampal model that addressed how the hippocampus may switch between representational reference frames. They suggested that changing reference frames results in directionality on the linear track, that the reference frames may be initially tied to each of the track ends, and that the directionality develops as an animal becomes “indifferent” to one of the ref-

erence frames when traveling toward the other end of a track (Touretzky and Redish, 1996). They state that the hippocampal place cells are therefore not direction-sensitive, but reference-frame-sensitive (Touretzky and Redish, 1996). Similarly, McNaughton et al. (1996) suggested that “each behaviorally significant location in the environment [may become] a reference center” resulting in the a directional differences in the hippocampal map of an environment. In linear environments, the uniqueness of the directionally dependent view of these reference frames is reinforced with each experience; while during foraging behaviors where the view of multiple reference centers can be accessed from multiple directions, this has modeled using competitive learning (Sharp et al., 1990). Thus, on open-environments no particular reference frame will have significance and the environment will be represented with a single map, while on a linear track where the animal shuttles back and forth, switching reference centers will switch the hippocampal reference frame (Touretzky and Redish, 1996; McNaughton et al., 1996; Redish, 1999).

3.2 Overdispersion

While network state is a source of known variability in the firing of neurons, another source is as yet unknown. The firing of place cells is reliable enough spatially to infer the rodent’s position to within 1 cm given only the current firing pattern in a hippocampal ensemble and the spa-

tial tuning of each neuron (Wilson and McNaughton, 1993). Fenton and Muller (1998) found that place cells exhibit extreme temporal variability or *overdispersion* in that their firing patterns are much less reliable temporally than would be predicted by an inhomogeneous Poisson process based on the neuron's own tuning curve. Láncsý et al. (2001) characterized this overdispersion in terms of a doubly-stochastic Poisson process that switches between two mean spike emission rates at a mean interval of between 1 and 2 seconds. They propose that this switching may result from the animal switching reference frames about once every second or two (Láncsý et al., 2001; Olypher et al., 2002). Note, this is at a much larger time scale than can be simply explained by the phenomenon of theta phase precession (Skaggs et al., 1996), where place cell firing correlates strongly with the phase of the theta rhythm. One possible explanation for this phenomenon is that rats switch the cues with which they reference their position (McNaughton et al., 1994).

3.2.1 Inhomogeneous Poisson Process

One possible model of the behavioral or stimulus dependence of a neuron's firing is that of an inhomogeneous Poisson point process where the intensity of the spike emission rate at any given time is dependent on the stimulus or behavioral value at that moment. This is the approach used by Fenton and Muller (1998) to examine the variability of hippocampal neurons as the animal passes through their place field.

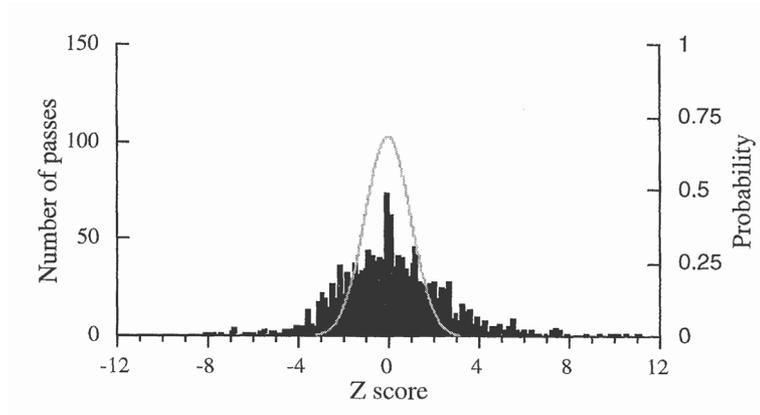


Figure 3.1: **Overdispersion of hippocampal place cells on the Open Field task.** A histogram of z-scores for the number of spikes emitted on a pass through a place field given the expected number of spikes predicted by a Poisson point process model. Fenton and Muller (1998) reported that neurons displayed excess variability on the open field foraging task (Figure from Fenton and Muller, 1998: z-scores for 1440 passes).

In an inhomogeneous Poisson point process, the expected number N of emissions between time t_0 and time t_1 is given by

$$N = \int_{t_0}^{t_1} \lambda(t) dt \quad (3.1)$$

where $\lambda(t)$ is the time-dependent intensity function of the point process.

Using this model, the spiking of a neuron can be compared with the expected firing S given the mean rate of firing R_i for each spatial position i taken from the tuning curve for each pass through the neuron's place field. Since the tuning curve is a discrete estimate of the spike emission intensity function at each location (i.e. $\lambda(t) \simeq R(x(t))$), the expected number of spikes N emitted by a neuron on a pass depends on the portion of the

tuning curve traversed as follows:

$$N = \sum_i R_i \Delta t \quad (3.2)$$

where R_i is the firing rate predicted by the tuning curve at each position sample of the pass and Δt is the video sampling rate.

As described in Fenton and Muller (1998), for $N > 4$, the Poisson distribution can be approximated by a normal distribution with mean $\mu = N$ and variance $\sigma^2 = N$. Thus, the Z-transformed distribution of N for all passes through a place field can be calculated as follows:

$$Z = \begin{cases} \frac{S - N - \frac{1}{2}}{\sqrt{N}} & \text{if } S \geq N, \\ \frac{S - N + \frac{1}{2}}{\sqrt{N}} & \text{if } S < N. \end{cases} \quad (3.3)$$

where S is the number of spikes actually emitted. The factor of $1/2$ is a correction for the discrete distribution.

3.2.2 Doubly Stochastic Inhomogeneous Poisson Process

The doubly stochastic inhomogeneous Poisson point process model presented by Láńksý et al. (2001) to explain the overdispersion process is similar to equation 3.1. However, $\lambda(t)$ is replaced with a $\Lambda(t)$, a two state *intensity process* with intensities $\lambda_1(t)$ and $\lambda_2(t)$, corresponding to states 1 and 2, respectively. The instantaneous intensity at time t_0 therefore de-

depends on which state the system happens to be in (state 1 or state 2) given random fluctuations. Thus, the instantaneous intensity at time t_0 is either $\lambda_1(t_0)$ or $\lambda_2(t_0)$.

Lánksý et al. (2001) provided a derivation of the *Fano factor* for such a doubly stochastic point process. Using the Fano factor of the data from Fenton and Muller (1998) and assuming equal dwell time, they estimated the dwell period in either state to be about 1 to 2 s.

3.2.3 Olypher Network Model

Olypher et al. (2002) expanded on the reference-frame switching discussion begun by Lánksý et al. (2001), to speculate on the anatomical origins of the place cell variability. They framed their ideas in a computational model to generate quantitative comparisons with the experimental data. Their model is depicted in Figure 3.2

In their model a subset of inputs randomly alternates between high and low intensity states. They vary the depth of modulation α and the average switching period T of these inputs and measure the variability in the output of their integrate and fire neuron. They suggest that the best fit to the data reported by Fenton and Muller (1998) is a 10% modulation depth with a mean switching period of 1 s.

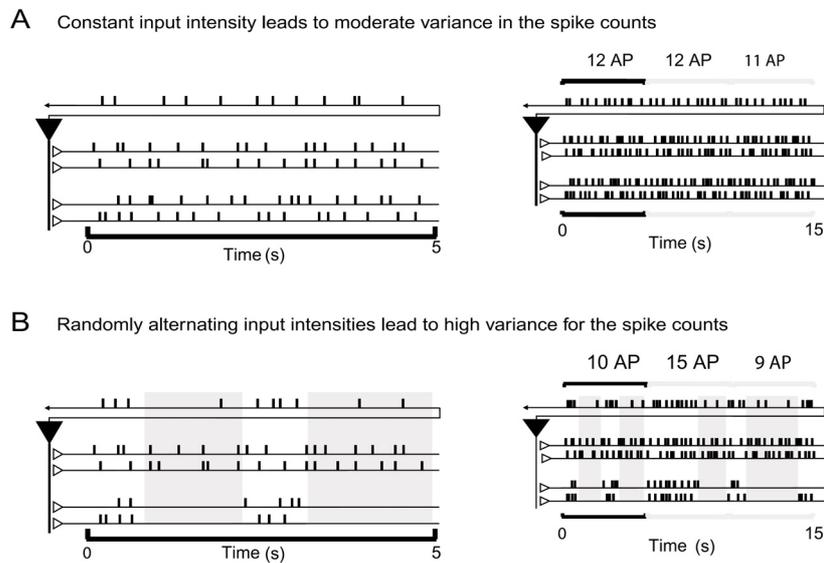


Figure 3.2: **Integrate and Fire Model of Overdispersion: Switching of a subset of the inputs to a neuron.** **A** Constant intensity inputs (bottom four rasters) to a neuron result in constant output intensity (top raster). **B** Switching of a subset of the inputs to a neuron between high and low firing rates (bottom two rasters) results in large variability in the output (top raster). The upper two inputs remain unmodulated. (From Olypher et al., 2002.)

3.3 Goal Tasks

3.3.1 Hippocampal Lesions

Both animals and humans exhibit spatial learning deficits following hippocampal damage (Scoville and Milner, 1957; Milner, 1970; O'Keefe and Nadel, 1978; Kesner and Novak, 1982; Morris et al., 1982; Squire, 1992; Redish, 1999; Clark et al., 2000). For instance, the patient H. M. discussed earlier had difficulty learning a stylus maze which required learning of a specific spatial path or sequence across an experimental board (Milner,

1970). The board had an array of bolts which, when touched, would count as an error if the bolt was not part of the path from start to finish of the maze (Milner, 1970). Likewise, rat trained on a serial position task, performed at chance levels after hippocampal lesions (Kesner and Novak, 1982). More direct links between deficits in goal directed spatial memory and hippocampal function have also been explored (Morris et al., 1982; see Redish, 1999 for review). In order to explore the link between the hippocampus and goal memory in rodents, Morris et al. (1982) used a water maze paradigm where rats were required to navigate to a hidden platform in order to escape the water. Control rats (sham operated and cortical damage controls) quickly learned the location of the hidden platform reducing their escape latency and favoring the platform's location on probe trials where the platform was not present. Animals with bilateral hippocampal lesions are dramatically impaired on both learning measures, although they eventually would find it each session (Morris et al., 1982).

3.3.2 Over-Representation of Goal

Since tasks that require goal-directed navigation require an intact hippocampus, it is important to understand the effects that spatial goals exert on the hippocampal network. For instance, it was recently reported that requiring an animal to navigate to a fixed goal can lead to recruitment of place cells that respond to the goal region resulting in an over-representation of the goal (Hollup et al., 2001a). This phenomenon was observed on an

annular water maze task. Constant and variable platform locations were used to test the goal dependence of place field locations. It was found that place fields clustered more densely in the platform region and in the region leading up to the platform for the case of uni-directional swimming. Training with variable platform locations abolished this clustering. (Hollup et al., 2001a) Thus, static goal locations can have a profound long-term effect on the hippocampal representation, and this effect can be minimized by varying the goal location day-to-day.

3.3.3 Effects on Place Cell Stability

Goal directed navigational requirements have other long-term effects on the hippocampus. Kentros et al. (2004) recently demonstrated that increasing the cognitive requirements of an experience increases the long-term stability of place fields in mice. The place fields of mice that were randomly wandering in an environment were generally quite different when the mice were re-exposed to the same environment six or more hours later. Scattering food in the arena resulted in a slight but non-significant increase in this stability; exposing animals to a novel environment significantly increased stability above the wandering group; and requiring animals to navigate to a goal region to terminate stress-inducing cues (noise and bright light) dramatically increased place-field stability above the other conditions. These effects were modulated by D1 and D5 dopamine receptor pharmacology. (Kentros et al., 2004) These data sug-

gest that motivational influences in general and goal-directed navigation in particular affect the stability of the hippocampal spatial representation.

3.3.4 Effects on Place Cell Temporal Variability

Finally, the temporal variability of place-cell firing, or overdispersion, depends on the task requirements. Specifically, place cells fired more reliably during an approach to an invisible goal than during random foraging. Simply introducing a goal-directed navigational requirement reduced temporal variability in place-cell firing, but place cell firing in the five seconds prior to reaching the goal was more reliable than either the foraging task or the goal-directed task as a whole. (Olypher et al., 2002)

These studies suggest that introducing a goal-directed navigational requirement to a task critically engages hippocampal circuitry such that the hippocampal representation is stabilized and perhaps biased towards a single representational reference frame. These changes are most likely due to the spatially-dependent cognitive requirements of the task that demand navigational calculations which depend on the precision of hippocampal information processing.

3.4 Local Field Potentials

Besides being viewed at the cellular level as action potentials, the electrical signatures of information processing within the brain can be measured as

a filtered summation of synaptic inputs across the network in the form of local field potentials (LFP). While action potential propagation is an electrical event, the LFP has more to do with synaptic transmission within a region than the propagation of action potentials in the local neuronal population. Thus the LFP is the low-frequency component of a large-scale superposition of changes in the electric field within a structural region resulting from network-wide synaptic activity.

3.4.1 Generation

The primary mode of information transfer between neurons is through the chemical synapse. Although there are many types of synapses known to exist within the nervous system, the general model of a synapse consists of a pre-synaptic axon terminal on the input-side and a post-synaptic aggregation of receptors on output-side. When an action potential from the pre-synaptic neuron reaches the axon-terminal, a chain reaction of events takes place beginning with the activation of voltage sensitive Ca^{2+} channels, the influx of calcium, and activation of vesicle docking protein complexes which ultimately leads to the vesicular release of neurotransmitter into the synaptic cleft (the space between the pre- and post-synaptic neuron). The binding of neurotransmitter to receptors in the post synaptic membrane opens these channels resulting in a depolarizing or hyperpolarizing current depending on the selective permeability of the ion channels. The transfer of ions across the cell membrane is coupled with compen-

satory currents at the cell body. The summation of the resulting neuronal currents distributed across a network constitutes a considerable change in the charge distribution within the structure. This fluctuation in the distribution of charge alters the local electric field and is measured as a fluctuation in the electrical potential at electrodes placed nearby. Local field potentials are therefore the compound effect of synaptic transmission in the brain and depend on the anatomy and cytoarchitecture of the region of interest. (For review see Niedermeyer and Lopes da Silva, 1999.)

3.4.2 Hippocampal Anatomy

The strength and type of local field potential oscillation observed within a region depends heavily on the local structural geometry and cytoarchitectural organization. The hippocampal formation consists of the entorhinal cortex, the dentate gyrus, the hippocampus proper, and the subicular complex (Amaral, 1987; Amaral and Witter, 1989).

Hippocampal Anatomy and Circuitry

The hippocampus proper is divided into 4 major regions: the entorhinal cortex (EC), the dentate gyrus (DG), Ammon's horn (cornu ammonis; or CA region), and the subiculum. The CA region is often subdivided into 3 or 4 sub-regions: CA1, CA2, CA3, and sometimes CA4[†]. The CA subre-

[†]The obscure divisions of the CA4 sub-field corresponds most closely to the polymorphic zone of the dentate gyrus (Amaral, 1987)

gions make up what is generally referred to as the hippocampus (Amaral, 1987). The longitudinal axis of the hippocampus stretches along the lateral ventricle from the septal portion to the amygdala in the temporal lobe forming a banana or 'C' shape.

The major cortical input to the hippocampus comes from entorhinal cortex. The granule cells of the dentate gyrus receive divergent input from layers II and III of the entorhinal cortex (Amaral, 1987). The input is divergent in that the proportion of DG innervated by a given length of EC is at least 2.5. This means, a region corresponding to 10% of EC, will project to a region corresponding to 25% of the dentate gyrus (Amaral and Witter, 1989). The EC projects topographically to DG such that the caudo-lateral EC is associated with septal regions of DG, and rostro-medial EC is associated with the temporal regions of DG (Dolorfo and Amaral, 1998; Burwell, 2000). The entorhinal cortex also projects to the CA fields and to the subiculum (Amaral, 1987; Amaral and Witter, 1989).

The DG granule cells send very narrow mossy fiber projections to CA3 such that a region along the longitudinal axis of a given size in DG is connected to a region of similar size in CA3 (Amaral and Witter, 1989). The pyramidal cells of the CA3 region, in turn, send out axons which bifurcate to form broad associational projections and the Schaffer collateral projections (Amaral, 1987; Amaral and Witter, 1989). The associational projections distribute widely along the longitudinal (or long) axis of CA3, while the Schaffer collateral projections go to CA1 and distribute in a broad

but systematically organized 3-dimensional pattern (Amaral and Witter, 1989). For example, a pyramidal cell in CA3 located midway along the septo-temporal longitudinal axis and near DG along the transverse axis would have its axon terminals distributed close to the CA1-CA2/3 border synapsing on the distal dendrites of CA1 pyramidal cells. Progressing toward the temporal pole would reveal that this CA3 pyramidal cell's axons terminate closer to the cell bodies of CA1 pyramidal cells near the CA1-subiculum border. (Amaral and Witter, 1989)

CA2 pyramidal neurons are much like those in CA3 but they do not receive mossy fiber input from granule cells of DG, yet they are much larger than the pyramidal cells of CA1 (Amaral, 1987). CA1 projects to the subicular complex and back to the entorhinal cortex (Amaral, 1987; Amaral and Witter, 1989). The subicular complex can be subdivided into the subiculum, the pre-subiculum, and the para-subiculum (in order of their divisions progressing from CA1 to EC) (Amaral, 1987). The regions of the subicular complex project to different layers of the entorhinal cortex. The subiculum projection terminates in EC layer IV, the pre-subiculum projection terminates primarily in layer III, and the para-subiculum projects to EC layer II (Amaral, 1987). The simplified loop of EC → DG → CA3 → CA1 is referred to as the tri-synaptic loop although the shortest path from EC through the hippocampus back to EC, a true *loop*, is one synapse (EC → CA1 → EC) and the longest path is five synapses (excluding recurrent connections: EC → DG → CA3 → CA1 → Subiculum → EC). (See (Ama-

ral, 1987) for review.)

Cytoarchitecture of the CA1 Pyramidal cell Layer

Some of our most detailed knowledge of the cellular basis for local-field-potentials comes from the hippocampus. The regularity of the CA1 sub-field of the hippocampus is especially striking. It consists of 5 layers: the stratum oriens, the stratum pyramidale, the stratum radiatum, the stratum lucidum, and the stratum lacunosum-moleculare. These layers are populated with a variety of interneurons that release the inhibitory neurotransmitter γ -aminobutyric acid (GABA). There are a wide variety of GABAergic interneurons found in CA1 sub-region of the hippocampus, however the contribution of only four primary types of interneuron to hippocampal local field potentials have been well studied: the basket cell, the axo-axonic cell, the bistratified cell, and the O-LM cell. (See top of Figure 3.3 for a diagram of projections)

The Basket Cell. Of the three types of inhibitory basket cell found in CA1, I will discuss only the GABAergic parvalbumin staining (PV^+) basket cells of the stratum pyramidale since the most is known about their participation in hippocampal LFP oscillations. PV^+ basket cells synapse primarily in the stratum radiatum on the proximal dendrites and cell bodies of CA1 pyramidal neurons, and probably receive glutamatergic input from entorhinal cortex, thalamus, and from pyramidal cells in CA1

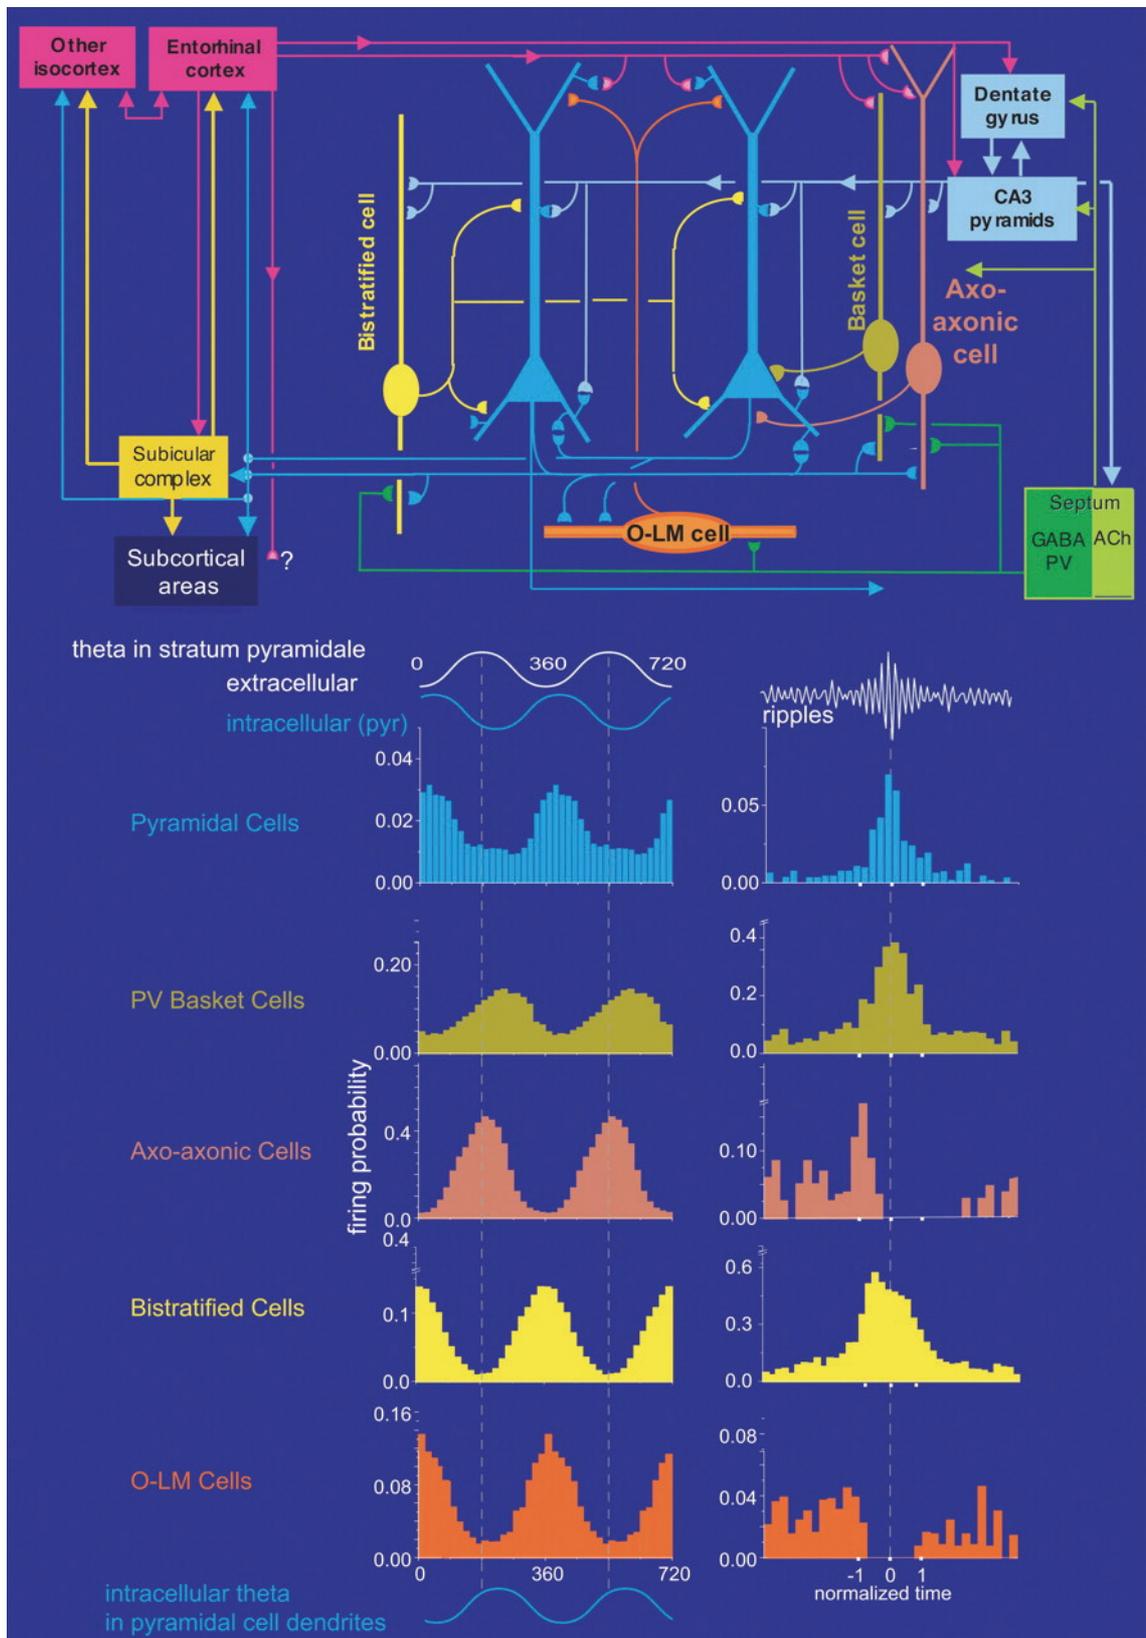


Figure 3.3: The hippocampal circuitry and spiking in response to LFP oscillations. (From Somogyi and Klausberger, 2005)

and CA3. Their dendritic arbor can reach from the stratum lacunosum-molecular to the stratum oriens. These PV⁺ basket cells may be electrically coupled through gap junctions and also form synapses on other interneurons including themselves (so called *autapses*; Pawelzik et al., 2003) and on other PV⁺ basket cells. (Somogyi and Klausberger, 2005)

The Axo-Axonic Cell. Parvalbumin staining axo-axonic cells reside in the stratum pyramidale, and synapse in the stratum oriens on the dendrites, somata, and axon initial segments of as many as 1200 CA1 pyramidal neurons. They presumably receive glutamatergic input from entorhinal cortex, from thalamus, and from pyramidal cells in CA1 and CA3. Like the basket cells, their dendritic arbor can reach from the stratum lacunosum-molecular to the stratum oriens, however the arborization in stratum lacunosum-molecular is much more extensive. (Somogyi and Klausberger, 2005)

The Bistratified Cell. PV⁺ bistratified cells also stain positive for somatostatin (SM) and neuropeptide-Y. They also reside in the stratum pyramidale and probably receive glutamatergic input from entorhinal cortex and pyramidal cells in CA1 and CA3. The axons of bistratified cells distribute in the stratum oriens and stratum radiatum and synapse on CA1 pyramidal neurons, basket cells, and other interneurons. Their dendritic arbor can range from the stratum radiatum to the stratum oriens. (Somo-

gyi and Klausberger, 2005)

The O-LM Cell. O-LM cells are found in the stratum oriens, stain positive for parvalbumin and SM, and synapse in the stratum lacunosum-moleculare on the distal dendrites of CA1 pyramidal cells and on other interneurons. They presumably and receive projections from entorhinal cortex and pyramidal cells in CA1 and CA3 as their dendrites are confined to the stratum oriens. (Somogyi and Klausberger, 2005)

We will discuss the contribution of these interneurons to hippocampal LFPs as we explore each LFP state. (See top of Figure 3.3 for a diagram of projections.)

3.4.3 Hippocampal LFP Types

The orderly structure of the hippocampus results in a number of high amplitude oscillations in local field potential recordings, the most notable of which are the sharp-wave associated fast ripples and the hippocampal theta rhythm.

Theta

The hippocampal theta rhythm or rhythmic slow activity (RSA) is a highly synchronous LFP state resulting from a complex interplay between the resonant frequencies of the local circuit of pyramidal neurons in the hippocampus and local inhibitory feedback, GABAergic input from pace-

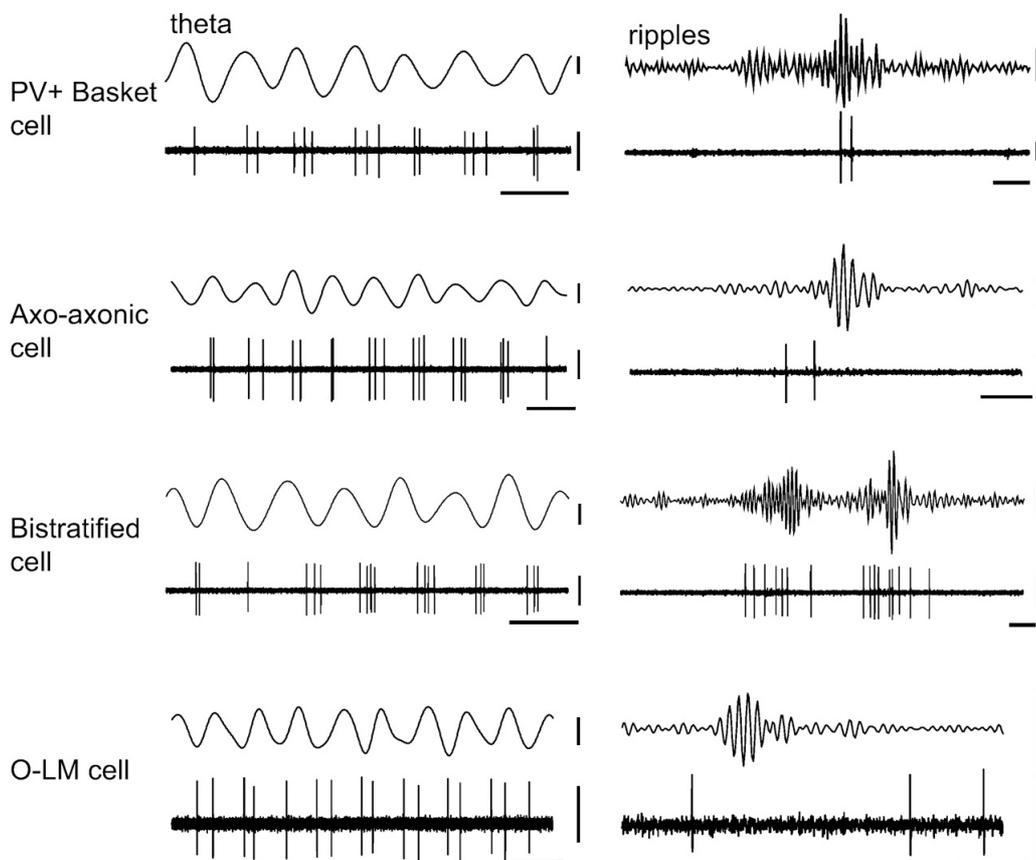


Figure 3.4: **The response of four interneuron types to prominent hippocampal LFP oscillations.** (From Somogyi and Klausberger, 2005)

maker neurons in the medial septum, and cholinergic modulation from medial septal inputs (Green and Arduini, 1954; O'Keefe and Nadel, 1978; Stewart and Fox, 1989; Nerad and McNaughton, 2006). The first thorough analysis of theta in the hippocampus reported that this large amplitude nearly sinusoidal 6-12 Hz oscillation was reciprocally related to synchronization of the cortical EEG, was abolished by septal and fornical lesions, and was related to the animal's (rabbits, cats, and monkeys) state

of arousal (Green and Arduini, 1954).

The primary components of theta rhythm generation involve three components: a septal muscarinic cholinergic input which depolarizes hippocampal interneurons and pyramidal neurons, a rhythmic GABAergic septal input to local interneurons, and a local oscillatory circuit comprised of hippocampal pyramidal neurons and GABAergic interneurons (Stewart and Fox, 1989). The cholinergic input from the septum depolarizes the local circuit to enable oscillation, the GABAergic input provides rhythmic input to the local circuit inhibitory interneurons to spatially synchronize the hippocampal theta rhythm (Stewart and Fox, 1989). The septal input appears to be distributed across the lateral and medial septal nuclei (Nerad and McNaughton, 2006).

This model is supported by several lines of data. Hippocampal theta is disrupted by lesions of the fornix and septum and systemic and septal atropine treatment (O'Keefe and Nadel, 1978; Buzsáki et al., 1983; Stewart and Fox, 1989). Furthermore, recent evidence suggests that distinct interneuron types participate in the theta rhythm at different phases (Klausberger et al., 2003; Somogyi and Klausberger, 2005). PV⁺ basket cells fire maximally at falling edge of the theta LFP recorded extracellularly at the pyramidal cell layer while axo-axonic neurons fire maximally at the peak. Bistratified cells and O-LM cells fire maximally at the trough of the theta wave. The distribution of these various inhibitory inputs along the dendrosomatic axis of CA1 pyramidal cells is consistent with 180° phase shift in

the intracellularly recorded membrane potential from the distal dendrites to the soma (Klausberger et al., 2003; Somogyi and Klausberger, 2005). (See Figure 3.3) As a result of this distribution of inputs to the CA1 field, the theta wave reverses polarity in the stratum radiatum. (See O'Keefe and Nadel, 1978.)

Behavioral Significance: Theta Phase Precession The theta rhythm is present during ambulatory movement, stimulus evoked attentive behavior, and during paradoxical (REM) sleep. The highly organized theta rhythm leads to well ordered spiking of CA1 pyramidal neurons in the spatio-temporal domain (O'Keefe and Recce, 1993; Skaggs et al., 1996). This highly ordered spiking compresses behaviorally relevant temporal sequences into the optimal LTP induction window of approximately 40 ms (see (Bi and Poo, 2001) for review) while preserving firing order. Such repetitive stimulation of intrinsic and extrinsic circuits fits the requirements needed for synaptic modification of hippocampal and cortical target regions such as the entorhinal cortex (Bliss and Lømo, 1973; Bliss and Gardner-Medwin, 1973; Chrobak and Buzáki, 1998).

LIA

Large-irregular activity (LIA) is observed as a desynchronization of the hippocampal EEG in comparison with the highly regular theta rhythm. The primary frequency is slower than theta, with periods of quiescence

punctuated by sharp deflections in the LFP termed *sharp waves* (see below for more information on sharp waves). LIA has often thought of as release of subcortical inhibition that is present during theta (O'Keefe and Nadel, 1978; Buzsáki et al., 1983).

Behavioral Significance LIA is normally observed in awake animals as they eat, drink, groom, and rest as well as in sleeping animals in the early stages of slow-wave sleep.

Sharp Waves A distinctive feature of LIA, sharp-waves (SW) are large amplitude LFP events that result from the depolarization of CA1 pyramidal neurons by CA3 Schaffer collateral inputs from the CA3 field. This massive depolarization is often accompanied by a fast ripple oscillation of approximately 200 Hz (Buzsáki et al., 1983; Csicsvari et al., 2000). At least 10% of the CA3 population must fire within about 100 ms to trigger a fast-ripple event in CA1 (Csicsvari et al., 2000). This fast ripple oscillation is then driven locally by interactions between the PV⁺ basket cells, bistratified cells, and the pyramidal cell population (Hirase et al., 2001; Csicsvari et al., 2000; Klausberger et al., 2003, 2004; Somogyi and Klausberger, 2005). The PV⁺ basket cells and bistratified cells fire phase locked to the ripple oscillation with peak firing rates occurring just after negative trough of the SW ripple (Klausberger et al., 2003, 2004). This can be seen in Figures 3.3 and 3.4 reproduced from Somogyi and Klausberger (2005).

This results in a temporally alternating current source-sink pattern flanking the stratum pyramidale (See Figure 3.5 reproduced from Ylinen et al., 1995) during the ripple event (Ylinen et al., 1995). The circuitry activated during sharp waves is same as that activated during theta (Ylinen et al., 1995).

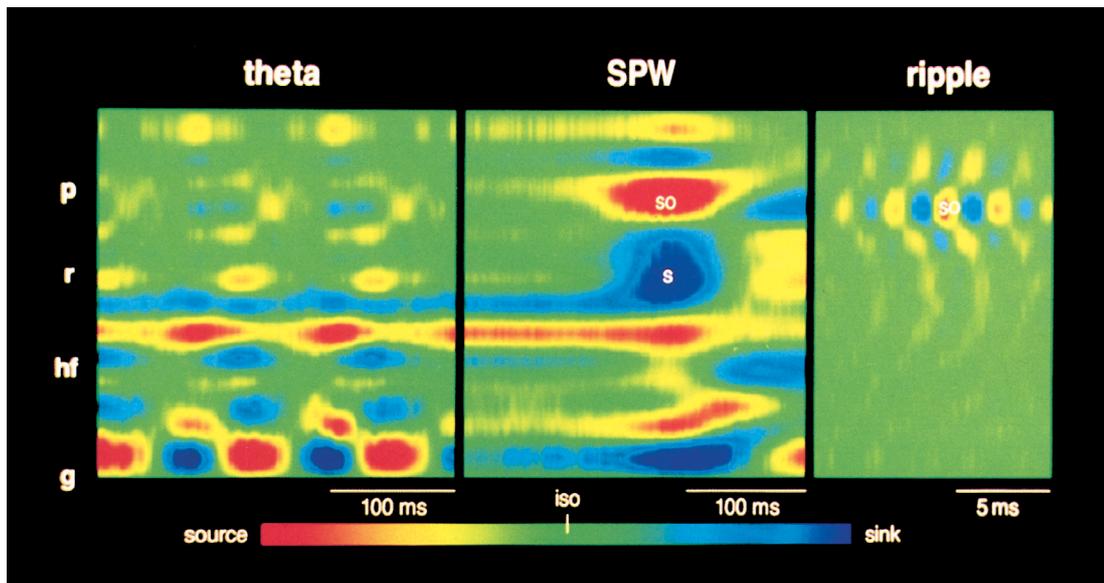


Figure 3.5: **Current Source Density Dynamics of theta, sharp wave, and sharp wave associated ripple.** (From (Ylinen et al., 1995))

SW are transferred from CA1 through the subicular complex and into entorhinal cortex. EC layer V-VI neurons spike with SW (Chrobak and Buzsáki, 1994). SW ripples occur in EC and subiculum shortly after CA1 SW ripples suggesting hippocampus drives cortical outputs during SW events (Chrobak and Buzsáki, 1996). EC ripples reverse polarity at layer II-III border possibly due to stimulation of layers V-VI by CA1 and subicular

SW (Chrobak and Buzsáki, 1996). Thus, the sharp-wave ripple event can be propagated from the hippocampus out to cortical regions.

Behavioral Significance As mentioned earlier, SW are a hallmark of LIA which means that they are associated with behaviors such as quiet resting, grooming, drinking, eating, and slow wave sleep. However, it has been recently reported that SW also occur during ambulatory movement outside of LIA (O'Neill et al., 2006). Neuronal firing during these so-called *exploratory sharp wave ripples* (eSWR, to differentiate them from SW occurring during immobility, iSWR) favors a neuron's place fields such that a neuron is more active during an eSWR if the animal is in that neuron's place field (O'Neill et al., 2006).

SIA

Small Irregular Activity (SIA) is a less well known network state of the hippocampus characterized by a short-duration (approximately 1 s), low-amplitude desynchronization of the hippocampal EEG. SIA can be triggered by electrical stimulation of the lateral pathway (nucleus reticularis pontis caudalis and the lateral raphe nucleus) or the medial pathway (the septum, or fornix). SIA is eliminated by lesions of the lateral pathway while it is not affected by septal or fornical lesions. It is therefore thought that there may be two separate types or modes of SIA: one dependent upon serotonergic input from the lateral pathway, the other due to hyper-

activation of the theta generating mechanisms (possibly resulting from over stimulation in experimental conditions). In the first case, the serotonergic input reduces the theta-rhythmic activity of hippocampal interneurons, while the latter case may be a result of broad neuronal depolarization. Thus, the induction of SIA by serotonergic input is probably the most physiologically relevant. (O'Keefe and Nadel, 1978)

Behavioral and Neurophysiological Correlates SIA can be elicited behaviorally through startling an animal from sleep (or near sleep), during approach to food, during freezing behavior, or during a heightened level of arousal during sleep (Jarosiewicz and Skaggs, 2004b; O'Keefe and Nadel, 1978). Recent ensemble recordings suggest that during SIA, the hippocampal CA1 population represents a static memory of the animal's current position by a nearly tonic activation of neurons with place fields at the animal's location prior to entering SIA (Jarosiewicz et al., 2002; Jarosiewicz and Skaggs, 2004a). Thus, manipulating the animal's position after it enters SIA has no effect on the hippocampal representation; it remains fixed representing the original location (Jarosiewicz and Skaggs, 2004a).

3.4.4 Hippocampal LFPs Depend Strongly on Electrode Position

The localization of the source-sink pairs resulting from inhibitory-excitatory push-pull coupling of pyramidal neurons and interneurons at the pyrami-

dal cell layer means that sharp waves and their associated ripples reverse as an electrode is advanced past the hippocampal cell layer. This reversal can be exploited to precisely position an electrode in the hippocampal pyramidal cell layer. Likewise, as was mentioned earlier, the theta wave reverses in the stratum radiatum. As electrodes are advanced ventrally past the CA1 pyramidal cell layer relatively large theta suddenly decreases and then inverts its phase giving way to a very high amplitude theta oscillation near or just below the hippocampal fissure (O'Keefe and Nadel, 1978). Thus, in our experiments a second reference electrode was lowered into the hippocampal fissure whenever possible to obtain this high signal-to-noise theta rhythm.

3.5 Reactivation

Due to the orderly structure of the hippocampus, it has been used as one of the major circuit models for exploring the effects of LTP (Bliss and Lomo, 1973; Bliss and Gardner-Medwin, 1973). The strong spatial response of neurons in the hippocampus also allows us to probe the effects of spatial experience on the network dynamics of the hippocampus (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978). When coupled with the clear delineation between local field potentials during the various stages of sleep and wakefulness (Green and Arduini, 1954; O'Keefe and Nadel, 1978), these features of the hippocampus have allowed neuroscientists to

probe deeply into the effect of learning during awake behavior on post-processing during sleep (Pavlides and Winson, 1989; Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Kudrimoti et al., 1999; Nádasdy et al., 1999; Louie and Wilson, 2001; Lee and Wilson, 2002) and during on-line processing during the task (Jackson et al., 2005a,b; O'Neill et al., 2006; Foster and Wilson, 2006).

Pavlides and Winson (1989) first presented evidence suggesting that plasticity induced by behaviorally evoked hippocampal neuronal activity in the waking animal affects the hippocampal activity patterns present during sleep. The authors recorded 2 or more place cells simultaneously from rats while they were confined to a portion of the environment containing a place field from only one of the neurons. Then, animals were allowed to sleep in an area that did not contain a place field from either neuron (only neurons that did not have a place field in the home cage were included) while recording from both neurons. The neurons that had place fields in the location the animal was confined to were more active and fired more multiple spike bursts during slow-wave sleep (SWS), pre-REM (PREM), and REM sleep than the neurons whose place fields were not visited. There were also slight increases in firing during quiet waking (QA) and still alert (SAL) behavioral states, but these were not significant. Switching which neuron's place field the animal was confined to, switched the offset in firing during sleep such that the neuron that had a place field in the location the animal was confined to during the second

test was more active during QA, SWS, PREM, and REM sleep than the neuron whose place field was not visited. Finally, the authors found that there was a specific increase in the number of inter-spike-intervals in the 2-4 ms range during QA, SWS, PREM, and REM sleep in the neurons that had a place fields in the location the animal was confined to than in the neurons whose place fields were not visited. They argue that this 400 Hz spike frequency during bursts is optimal for inducing LTP in target structures (Pavlides and Winson, 1989).

3.5.1 Slow Wave Sleep

Following the observations of Pavlides and Winson (1989), it was not until much later that the first ensemble-level interactions between behavioral co-activation and reactivation during sleep were examined (Wilson and McNaughton, 1994). Hippocampal ensembles were recorded before, during, and after a behavioral session of either an open field environment or a linear track environment. The authors demonstrated that neuron pairs with overlapping place fields had weak cross-correlations before an animal experienced the task, but had strongly correlated firing during the task and had increased cross-correlations during slow-wave sleep (SWS) following the task. Neurons pairs with non-overlapping place fields had significantly lower cross-correlations during the task and in the sleep afterwards. They showed that the temporal connectivity during behavior was therefore preserved during sleep after a task (Wilson and McNaughton,

1994). These data indicate that the firing of neurons during SWS is strongly related to the firing of neurons during behavior. They also suggest that performing a task results in spatially dependent co-activation of neurons which enhances the network-level coupling during sleep such that the spatial representation present during behavior is recreated in the firing of neurons during sleep. Thus, the spatially dependent cross-correlations within an ensemble are reactivated in SWS.

Later, Skaggs and McNaughton (1996) reported that the temporal bias between neuron pairs during behavior (an indicator of the temporal sequence of spiking) was strongly correlated with the bias during SWS after a task but not during SWS before the task. Furthermore, comparing the number of neurons with both positive correlations during the task and during sleep (N_{++}) with the number of neurons with positive correlations during the task and negative correlations during sleep (N_{+-}) revealed increases in N_{++} over N_{+-} for sleep following task behavior but not for sleep before the task (Skaggs and McNaughton, 1996). Thus, Skaggs and McNaughton (1996) demonstrated that the temporal ordering of neuron pairs during a task is preserved during sleep following a task, but is not present during sleep prior to the task. This suggests that the temporal sequences of ensemble activation generated during behavior are replayed in the slow wave sleep immediately following an experience.

Kudrimoti et al. (1999) followed up these reactivation studies by investigating the influence of specific LFP states, behavioral experience, task fa-

miliarity, and reactivation measurements. The authors examined the firing rate measures of Pavlides and Winson (1989), a correlations measure, and explained variance. They used explained variance to account for the contribution of pre-existing correlations in sleep prior to task performance, which was particularly important for their comparison of the effects of novel versus familiar tasks. The authors explored the specific effects of reactivation during sharp wave associated ripple complexes in slow-wave sleep, demonstrating that reactivation occurs during these ripple events. Their data show a decay in reactivation with increasing time spent sleeping following the task. Novel experiences are also reactivated following sequential exposure to familiar and then novel environments in the same session. However, the reactivation of these novel patterns was less than half the reactivation of the familiar environment even though the animals were exposed to each environment for the same length of time. There was no significant reactivation of behavioral firing patterns observed during REM sleep and REM seemed to have no observable effect on the correlation structure within subsequent SWS epochs. In summary, novel and familiar environments are reactivated to differing degrees within the sharp wave ripple events that occur during SWS, and this reactivation decays during SWS, presumably until it reaches baseline levels prior to the next session.

Up to this point, only pair-wise evidence for the replay and reactivation of behavioral cell assemblies had been shown. Likewise, the spik-

ing depended strongly on the spatial progression of the animal's behavior. Nádasdy et al. (1999) therefore used template-matching of spike sequences and joint-spiking measures to investigate the higher-order temporal ordering of replay in firing patterns reactivated after wheel-running. The authors reported that the number of repeated spike-sequences during wheel running was highly improbable given chance activation (as defined by multiple shuffling controls including across-spike-train shuffling and within-spike-train shuffling). This suggests that the hippocampal network generates temporal spiking patterns with an exquisite level of precision and repeatability during awake behavior even for relatively non-spatial tasks. This is one primary requirement for inducing plasticity for long-term storage of information in cell assemblies. Using joint-probability maps to measure the probability of observing a particular temporal sequence of spikes from three neurons revealed that there were a significant number of spike triplets common to both the run and post-sleep, while the number of triplets common to run and pre-sleep were not significantly above chance. Finally, it was found that the most compressed sequences were correlated with increased power in the sharp-wave ripple band (at approximately 160 Hz). (Nádasdy et al., 1999) Together, their data suggest that the precise temporal sequences present during awake behavior are stored and replayed in a temporally compressed manner during sharp-waves in the sleep following a behavioral episode.

The temporally compressed replay of behavioral ensemble spike se-

quences was conclusively demonstrated for long strings by Lee and Wilson (2002). Ordering neurons by the spatial sequence of their place field centers on the linear track allowed combinatorial testing for the presences of these sequences during sharp-wave events during slow-wave sleep. Highly significant ordered strings of ensemble firing patterns were observed in the sleep following a behavioral session, but rarely in the sleep prior to the behavioral session. The median compression ratio of the replayed firing sequences was 19.7 times faster than the behavioral sequences produced by the animal traversing the track (Lee and Wilson, 2002).

In summary, the temporal ensemble spiking patterns generated during behavior are replayed during sharp-wave events in slow-wave sleep in well-ordered temporally compressed bursts (Nádasdy et al., 1999; Lee and Wilson, 2002). This reactivation decays during SWS (Kudrimoti et al., 1999) until very little replay is present during the sleep prior to the next behavioral session (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Kudrimoti et al., 1999; Nádasdy et al., 1999; Lee and Wilson, 2002). However, these sequences, while rarely reactivated in pre-task sleep, are still present resulting in stronger reactivation for familiar experiences than for novel experiences (Kudrimoti et al., 1999).

3.5.2 Paradoxical Sleep (REM)

While reactivation during SWS has been well studied, REM reactivation has not received nearly the same attention. After Pavlides and Winson

(1989) reported increased firing of neurons whose place fields were activated during behavior, Kudrimoti et al. (1999) did not find significant reactivation of behavioral firing patterns during REM sleep. They did report, however, that REM firing patterns were similar to the flanking SWS epochs which demonstrated reactivation of the behavioral firing patterns (Kudrimoti et al., 1999). Two years later, Louie and Wilson (2001) demonstrated temporally compressed replay of ensemble behavioral firing sequences during REM sleep. The modulation dynamics of theta rhythm power observed during the task sequence were also preserved during REM sleep. It is possible that the approximately two-fold compression of the replay during REM episodes could have interfered with the explained-variance analysis used by Kudrimoti et al. (1999).

3.5.3 Awake Reactivation

It has long been known that SW events occur during awake states such as grooming, feeding, and quietly resting (See O'Keefe and Nadel, 1978). However, only very recently has the phenomenon of reactivation been investigated during awake SW.

In two conferences, we presented evidence that reactivation occurs during awake sharp waves in a behaviorally dependent manner (Jackson et al., 2005a,b). The following year, O'Neill et al. (2006) reported that two types of SW were emitted by awake animals: immobile sharp wave ripples (iSWR), and exploratory sharp wave ripples (eSWR). The iSWR is found

during states such as grooming, feeding, and quietly resting, while the eSWR is emitted while the animal is moving but in a low theta power state (O'Neill et al., 2006). Both ripples have very similar properties in that the firing rate of a pyramidal neuron within these ripples depends on whether the animal is in the neuron's place field, and that these ripples therefore carry sufficient information to constitute a reactivation of the pattern present during high theta power exploration. O'Neill et al. (2006) also demonstrated that reactivation increased from the first 10 min to the last 10 min of their task. Shortly after this result was published, Foster and Wilson (2006) reported that some awake sharp-waves carry a reverse replay signal such that the temporal sequence of place-cell firing during behavior was reversed during SW. While these studies demonstrate that awake SW ripples are accompanied by the reactivation of firing patterns occurring during theta, the behavioral dependence of these SW ripples and the associated reactivation remains unclear. There are many theories that predict that SW ripple emission should depend strongly on an animal's level of experience.

3.5.4 Theories

Theories of hippocampal function (Buzsáki, 1989; Buzsáki et al., 1994; McNaughton et al., 1996; Shen and McNaughton, 1996; Redish and Touretzky, 1998; Redish, 1999) predict that asymmetric plasticity (Levy and Steward, 1983; Bi and Poo, 2001) applied to recurrent connections within CA3

through experience of repeated spatial sequences during theta will lead to storage of sequences within the recurrent connectivity matrix (Levy and Steward, 1983; Muller et al., 1991; Blum and Abbott, 1996; Redish and Touretzky, 1998). During states in which the network was uncoupled from its entorhinal inputs (e.g. slow wave sleep and LIA, Chrobak and Buzsáki, 1994, 1996; Chrobak et al., 2000), uncorrelated noise in the system would then cascade across these strengthened synapses producing a replay of this stored information during sharp-waves (Buzsáki, 1989; Buzsáki et al., 1994; Buzsáki, 1996; Buzsáki et al., 1994; Ylinen et al., 1995; Shen and McNaughton, 1996; McNaughton et al., 1996; Redish and Touretzky, 1998; Redish, 1999; Csicsvari et al., 2000).

Consolidation

The observation that H. M. could remember temporally remote episodes and environments but not the events leading up to his surgery or anything new since, has inspired a number of the theories about the consolidation of memory (Marr, 1970, 1971; McNaughton, 1983; Buzsáki, 1989; McClelland et al., 1995; Sejnowski and Destexhe, 2000; Hoffmann and McNaughton, 2002). At the core of these theories, activation of the hippocampus during behavior modifies the recurrent connectivity matrix within CA3. During sleep the information stored in this connectivity matrix is played out to the cortex transferring the memory from a labile medium in to a more permanent store. This reactivation of stored sequences during sleep is

thought to underlie the consolidation of memory traces to other brain regions (Buzsáki, 1989; McClelland et al., 1995; Sejnowski and Destexhe, 2000; Hoffmann and McNaughton, 2002).

The Two-Stage Memory System

Early theories of cortical function were quick to point out the need for a two-stage learning system (Marr, 1970, 1971; McNaughton, 1983; Buzsáki, 1989; Buzsáki et al., 1994; McClelland et al., 1995; McNaughton et al., 1996; Buzsáki, 1996; Redish, 1999). In his model of neocortex, Marr (1970) realized that requirements for training cortical classificatory units imposed constraints that could only be solved by access to a simple associative memory system (Marr, 1971) that could present collections of data to units for refinement of their connections. Marr (1970) also saw the need to separate cortical plasticity mechanisms for training from the normal functioning state of his units. These constraints basically led to a two step proposal for the construction of class units in his cortical model: 1) on-line collection of data in an associative memory store; 2) retrieval of data from the memory store for construction and refinement of classificatory units (Marr, 1970). Nearly thirty-five years later, one of the most successful, comprehensive, and physiologically-justified theories of memory trace formation is based on a very similar framework.

A more complete conceptualization emerged in The Two-Stage Model presented by Buzsáki (1989) which is based on reciprocal interactions be-

tween the hippocampus and neocortex. Founded on an extensive review of hippocampal neurophysiology and experiments of his own, Buzsáki (1989) suggested the following dynamic process:

1. During theta-associated behaviors, the dentate gyrus transfers processed entorhinal cortical activation to the CA3 pyramidal cells. The resulting combination of recurrent CA3 synaptic circuitry, direct entorhinal input, and mossy fiber input from the dentate granule cells results in a weak heterosynaptic potentiation of behaviorally relevant activation vectors. (Buzsáki et al., 1983; Buzsáki, 1989; Buzsáki et al., 1994; Buzsáki, 1996; Buzsáki, 2005b)
2. The resulting potentiation of the CA3 recurrent connectivity matrix lends the pyramidal cells of the CA3 field to synchronous population bursts when sub-cortical inhibition is released at the termination of exploratory behaviors. These bursts in CA3 induce long-term potentiation in the CA3 recurrent network as well as in the Schaffer collateral input to CA1. The depolarization of the CA1 network by CA3 increases the synchrony of CA1 pyramidal cells resulting in a population bursts that have a powerful effect on neocortical targets. (Buzsáki et al., 1983; Buzsáki, 1989; Buzsáki et al., 1994; Buzsáki, 1996; Buzsáki, 2005b; Ylinen et al., 1995; Chrobak and Buzsáki, 1996)

Additional important features of this model are that the neocortical input during theta selects the CA3 population that will trigger the sharp

wave burst (Buzsáki, 1989). This model differs from a pure Hebbian mechanism however in that there are no classical teaching or “detonator” synapses; cortical, mossy fiber, and recurrent inputs are all modified and play a part in shaping the connectivity of the CA3 network (Buzsáki, 1989). The perforant path inputs to the hippocampus from EC are silent during SW allowing stored information to be replayed without interference from incoming information (Buzsáki et al., 1994). The temporally graded potentiation of CA3 recurrences should result in the activation of the most recently and most highly potentiated synapses first transferring activation to the least potentiated synapses (assumed to be also the most temporally distant in this model) resulting in a compressed reversed-order replay of stored inputs at the termination of exploration (Buzsáki et al., 1994; Buzsáki, 1996). As discussed in an earlier section, this prediction of compression (Nádasdy et al., 1999; Louie and Wilson, 2001; Lee and Wilson, 2002) and reverse replay (Foster and Wilson, 2006) have been recently confirmed. ‡ It should be noted, however, that the reverse replay has only been observed in awake sharp waves (Foster and Wilson, 2006), while all other analyses of SW occurring during SWS reveal forward replay of behavioral sequences (Nádasdy et al., 1999; Louie and Wilson, 2001; Lee and Wilson, 2002). This actually fits well with the processes described in the series of papers on the the Two-Stage Memory System (Buzsáki et al., 1983;

‡This awake “reverse replay” could also be evidence for the *self-localization* process of the Redish and Touretzky (1998) model (see below).

Buzsáki, 1989; Buzsáki et al., 1994; Buzsáki, 1996), however given that the asymmetric nature of LTP mechanisms expected to underlie hippocampal plasticity (Bi and Poo, 2001) cannot readily explain this phenomenon, more research is required to understand the differences between these replay processes.

According to the Two-Stage Memory System model, disrupting sharp waves should disrupt memory trace formation (Buzsáki et al., 1994; Buzsáki, 1996). More recent additions to the model include the addition of *recall* to the *storage and replay* model discussed above (Redish, 1999). The recall process proposed by Redish (1999), like the self-localization process (Redish et al., 1998), involves direct local view input to the CA3 and CA1 subregions to reset path integration. Another addition to the model includes viewing the theta rhythm as “an essential temporal organizer” that allows for navigation in both Euclidean space and in neuronal space (Buzsáki, 2005b). This is because compression of input sequences coupled with the rules of synaptic plasticity result in an orderly evolution of sequences of cell assemblies leading to higher order associations across spatial and non-spatial (i.e. neuronal) input vectors (Buzsáki, 2005b). Thus, the “semantic”-like groupings of associated non-spatial input sequences and “map”-like groupings of associated spatial input sequences that were presented in a temporally discontinuous manner during theta are assembled, packaged, and transferred back to cortex by population bursts during hippocampal sharp waves (Buzsáki, 2005b).

While the The Two-Stage Memory System discussed above (Buzsáki et al., 1983; Buzsáki, 1989; Buzsáki et al., 1994; Buzsáki, 1996) is an implicit, or word-based, model in the sense that it is a thorough description of a phenomenon based on solid neurophysiological evidence, Redish and Touretzky (1998) presented an explicit model (i.e. a computational instantiation) of the interaction between pattern storage during exploration and replay during sharp waves. The model consisted of multiple networks corresponding to cortical inputs (including entorhinal and polysensory areas) coupled to a CA3-CA1 model. In their model, SW occurring upon entry into an environment were associated with a self-localization process that centered the animal's hippocampal activity on its current location. During goal directed navigation, a Hebbian learning rule was applied to the synapses within hippocampal network. They demonstrated explicitly that a vector field of route information was stored in the CA3 recurrent connections of their model. Following this "exploratory" phase, in a simulation of LIA, random input elicited a replay of the animal's previous goal directed routes. The output of these replays were used to train the cortical networks. Untrained cortical networks (assumed to be similar to hippocampal lesion conditions) were unable to spontaneously navigate to the goal in subsequent simulations. Trained cortical networks, however, were able to navigate to the goal with ease. This propagation of path information stored in CA3 out to cortex is taken as a model of the phenomenon of consolidation (see below; Redish and Touretzky, 1998).

It is interesting to note, that Redish and Touretzky (1998) did not use the entire reactivation-sequence but only a sub-sampled one to train the cortical networks during consolidation in order to reduce computer time. This actually mimics the phenomenon of temporal compression now known to exist in replayed sequences (Nádasdy et al., 1999; Louie and Wilson, 2001; Lee and Wilson, 2002).

In summary, Redish and Touretzky (1998) successfully implemented an explicit model of the The Two-Stage Memory System proposed originally by Buzsáki (1989). They demonstrated quantitatively the feasibility of this conceptual framework through a simplified abstraction of a complex group of networks known to exist in the brain.

McNaughton et al. (1996) and Shen and McNaughton (1996) have also suggested that Hebbian plasticity applied to the specific hippocampal maps that were activated during behavior would result in a reactivation of those same maps during sleep. Their interest was in the reactivation of the hippocampal maps that were used on a task, as opposed to more general hippocampal activation. In their simulations, Shen and McNaughton (1996) demonstrated that such sleep-related reactivation occurs in an attractor model of the hippocampus if one assumes potentiation of overlapping place cells within a behaviorally activated hippocampal map of an environment (an associative mechanism) or a specific reduction of inhibition of cells within a behaviorally activated map (a non-associative mechanism) (Shen and McNaughton, 1996).

In conclusion, we can extrapolate from this two-stage model that if these theories are correct, then the theories discussed above predict that the emission of awake sharp wave ripple events should increase in number with experience within a session and that the organization of ensemble firing during those awake sharp-waves should improve with experience. These increases in sharp-wave emission and reactivation should depend on the level of repetition of spatial sequences. Since experimental evidence suggests SW activity in CA3 can initiate CA1 SWs *in vivo* and *in vitro* (Buzsáki et al., 1983; Ylinen et al., 1995; Csicsvari et al., 1999a; Behrens et al., 2005), any changes in SW activity in CA3 should be observable in CA1 ripples as well.

Chapter 4

Simulation Studies: Characterizing and Developing Ensemble Measures

Abstract

Developments in neurophysiology and data acquisition have enabled the recording of large ensembles of neurons. These advances provide an opportunity to probe the function of networks in the brain. While current analytical methods focus on extracting the values of variables thought to be encoded in these networks, little attention has been given to measures that assess the quality of this encoding or to assessing the network-level interactions between neurons. This chapter presents two ensemble analyses

that address these issues by testing the consequences of our assumptions about information processing at the network level. We characterize these analyses using a simulated neural network and demonstrate that the information provided by these ensemble measures allows us to detect network firing patterns that violate our assumptions about the network.

4.1 Introduction

With the advancement of recording technology, the ability to record large ensembles of neurons simultaneously has allowed us to probe deeper into information processing in the brain (Brown et al., 2004; Buzsáki, 2004). These developments in data acquisition have stimulated the invention of new data analysis tools for extracting information from distributed neural representations of behavioral variables (Brown et al., 2004). Among these tools, *reconstruction* techniques which extract a task variable from the ensemble firing patterns have been the most widely applied ensemble analyses in neuroscience (Salinas and Abbott, 1994; Brown et al., 2004, 1998; Georgopoulos et al., 1986; Wilson and McNaughton, 1993; Zhang et al., 1998). Generally speaking, the usefulness of these reconstruction algorithms has been assessed by the error between the reconstructed value and the actual task variable of interest (Salinas and Abbott, 1994; Wilson and McNaughton, 1993; Brown et al., 1998; Zhang et al., 1998). However, when considering the possibility that an animal's cognitive processes may

at times be independent of the current stimuli or task parameters (Harris et al., 2003; Hebb, 1949), it is important to note that this reconstruction error may not be due to a failure of the reconstruction algorithm or due to “random” noise in the system. The case may be that the neuronal ensemble could actually be processing that variable in a coherent manner. Therefore, it is of great importance that we measure the self consistency of a neural ensemble.

For example, Redish et al. (2000) presented a measure of representational consistency which they used to determine the reference frame with which a hippocampal ensemble’s firing pattern was most consistent; they termed this measure *coherency*. Redish et al. (2000) used a three-step analysis: calculate the activity packet (a weighted sum of tuning functions based on the actual activity), calculate the expected activity packet (a weighted sum of tuning functions based on the expected neural activity), and then coherency was defined as the dot-product between the two packets. Coherency was measured as hippocampal ensembles realigned between two coordinate systems. They demonstrated that this method was indeed able to detect the time of transition between representational reference frames in hippocampal ensembles. Furthermore, this transition was correlated with the animal’s behavior on the task.

This is an example where considering the “coherency” of an ensemble’s firing patterns provided deeper insight into the network dynamics. In this chapter, we will explore variations on two ensemble quality measures

through simulation studies and will then turn to other concepts for assessing representational quality. The first measure, an improved *coherency*, provides information about how consistent the ensemble's firing pattern is with the known spatial tuning of each neuron. The next measure, *ensemble consistency* (EC), quantifies how consistent a firing pattern is with a distribution of previously measured firing patterns. These measures will be characterized using neural network simulations of three examples of dynamic phenomena: 1) random network firing vs. a stable activity mode; 2) smooth rotation of the represented value vs. a jump in the representation to a distant value; and 3) ambiguous (or bimodal) representation vs. a single value. In our characterization of these measures, we will also discuss the output of a population vector reconstruction algorithm and the limited information it provides. Later we will discuss the application of these methods to neural data.

4.2 The simulations

Simulations provide a fast, efficient, and, most importantly, controlled means of generating data for the purposes of characterizing ensemble measures. The attractor network presented here is related to networks used to model the rodent head direction system (Skaggs et al., 1995; Redish et al., 1996; Zhang, 1996; Goodridge and Touretzky, 2000; Sharp et al., 2001). Briefly, this network employed symmetric local excitatory connections be-

tween neurons with similar preferred directions and global inhibition with periodic boundary conditions. Thus, this network can be thought of as a circular ring of neurons with a stable attractor state consisting of a single mode of active neurons. This mode can be located anywhere on the ring. Figure 4.1 provides a schematic visualization of this attractor network. Three issues were examined:

Issue 1: Random Network Firing vs. Stable Activity Mode When started from random noise, neurons in a ring attractor will compete until a group of neighbors wins and the network settles to a stable mode of activity at that location (i.e. representing one direction). Neurons with preferred directions near this direction will have higher firing rates than those distant from this direction. Thus, the final mode will be randomly selected given a random input (Wilson and Cowan, 1973; Kohonen, 1977).

Issue 2: Rotation vs. Jump When this system is in a stable state (i.e. representing one direction), and network inputs drive neurons with preferred directions near the represented direction (within 60° in our network), the represented direction will shift toward the input (Redish et al., 1996; Zhang, 1996; Samsonovich and McNaughton, 1997; Redish, 1999). In contrast, when the network inputs drive neurons with preferred directions far from the represented direction (greater than 60° in our network), the system will non-linearly jump to a new direction if the strength of the

drive is large enough to overcome the global inhibition (Zhang, 1996; Samsonovich and McNaughton, 1997; Redish, 1999).

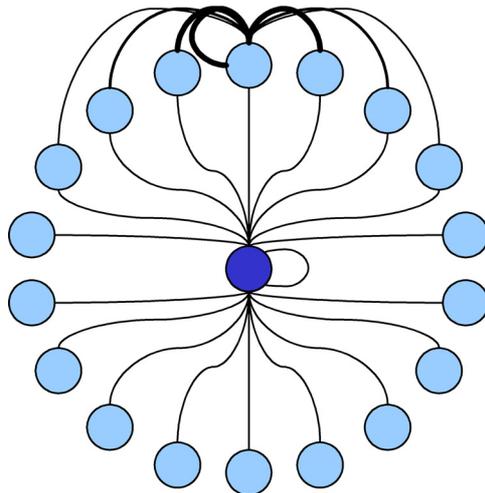
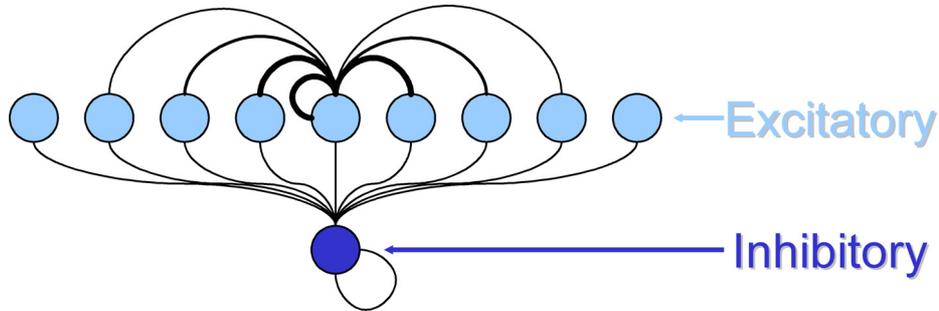


Figure 4.1: (*top*) The attractor is formed by connecting all neighboring excitatory neurons with weights based on their representational distance (e.g. neurons with similar preferred directions have stronger excitatory connections). All excitatory neurons project with equal weights to an inhibitory interneuron which projects back to all neurons with globally uniform inhibition. All neurons have recurrent connections as well. (*bottom*) Applying periodic boundary conditions effectively creates a ring attractor where each neuron can be thought of as representing a particular angular position on a circle.

Issue 3: Ambiguous vs. single valued representations If the network is started from a bimodal state (i.e. with inputs at two different directions), the population of neurons representing each input location will compete until one group (or direction) wins. The result will depend on the noise in the network.

4.3 Simulation Methods

We used this well-studied attractor network (Wilson and Cowan, 1973; Amari, 1977; Kohonen, 1984; Redish et al., 1996; Zhang, 1996; Redish, 1999) in order to facilitate the exploration of known dynamic, transient states and compare them to known stable states. This allowed a fuller study of the ensemble measures discussed in this chapter than would have been possible if we had used experimental data.

This type of network has been used to model a variety of neural structures, including the head direction system of the rodent (Skaggs et al., 1995; Redish et al., 1996; Zhang, 1996; Goodridge and Touretzky, 2000; Sharp et al., 2001), place cells within the hippocampus (McNaughton et al., 1996; Shen and McNaughton, 1996; Zhang, 1996; Samsonovich and McNaughton, 1997; Redish, 1999; Káli and Dayan, 2000; Guazzelli et al., 2001), the formation of ocular-dominance columns (Obermayer et al., 1992; Miller, 1995), control of saccadic error in the superior colliculus (Sparks, 1986; Munoz et al., 1991; van Opstal and Kappen, 1993; Arai et al., 1994) and in

the basal ganglia (Arbib, 1995), and memory storage within cortex (Wilson and Cowan, 1973; Kohonen, 1982, 1984).

4.3.1 The Network Model

Simulations were based on those presented in Redish (1999). Firing rate of each unit was continuous and normalized to be between 0 and 1. Continuous, random, Gaussian noise was added in order to better simulate physiological conditions. 75 excitatory neurons (voltage, V_k^E ; firing rate, F_k^E , and synaptic drive, S_k^E) and 1 inhibitory neuron (voltage, V^I ; firing rate, F^I ; and synaptic drive, S^I) were used in these simulations. The excitatory neurons were arranged with uniform spacing of the preferred directions along a ring topology. This simplified boundary conditions and other calculations. Specific parameters are given in Table 4.1.

The neural analogs of voltage and firing rate are straight-forward. The neural analog of synaptic drive is derived by normalizing the synaptic α -function (here taken to be an exponential decay) by the synaptic weight (Pinto et al., 1996). Thus, rather than calculating firing rate, multiplying that by the synaptic weight and then producing a synaptic effect; we calculate firing rate, produce the synaptic effect, and then multiply that by the synaptic weight.

The voltage for each excitatory neuron was calculated from the synaptic drive of each excitatory neuron and from the synaptic drive for the inhibitory neuron. The weight matrix $W_{E \leftarrow E}$ was a Gaussian kernel of stan-

Table 4.1: Neural network simulation parameters

V_k^E	membrane voltage of excitatory neuron k	var
$F_k^E(t)$	normalized firing rate of excitatory neuron k	var $\in [0,1]$
$S_k^E(t)$	synaptic drive from excitatory neuron k	var
$IN_k(t)$	extra-network excitatory input to excitatory neuron k	var $\in [0,2]$
ξ_G	Gaussian noise added to S_k^E	$\mu = 0, \sigma = 0.1$
γ^E	tonic inhibition to each excitatory neuron	-1.5
τ^E	decay time constant for presynaptic effect of E-neuron	10 ms
V^I	membrane voltage of inhibitory neuron	var
$F^I(t)$	normalized firing rate of inhibitory neuron	var $\in [0,1]$
$S^I(t)$	synaptic drive from inhibitory neuron	var
γ^I	tonic inhibition to inhibitory neuron	-7.5
τ^I	decay time constant for presynaptic effect of inhibitory neuron	2 ms
$W_{E \leftarrow E}$	synaptic weight kernel for E-to-E connections	see text
$W_{E \leftarrow I}$	synaptic weight for I-to-E connections	-8.0
$W_{I \leftarrow E}$	synaptic weight for E-to-I connections	0.88
$W_{I \leftarrow I}$	synaptic weight for I-to-I connections	-4
Δt	time-step	1 ms

dard deviation 72° , providing a local excitation function. The weight $W_{E \leftarrow I}$ was a constant providing global inhibition.

$$V_k^E(t) = \sum_j W_{E \leftarrow E}(\phi_k - \phi_j) \cdot S_j^E(t) - W_{E \leftarrow I} \cdot S^I(t) + \gamma^E + IN_k(t) \quad (4.1)$$

where ϕ_k and ϕ_j represent the preferred directions of neuron k and neuron j , respectively. Firing rate was taken as a simple sigmoidal function of

voltage.

$$F_k^E(t) = \frac{1 + \tanh V_k^E(t)}{2} \quad (4.2)$$

Noise was added to the excitatory neurons in the synaptic drive equation.

$$S_k^E(t + \Delta t) = S_k^E(t) + \frac{\Delta t}{\tau^E} \cdot (F_k^E(t) - S_k^E(t)) + \xi_G(t) \cdot S_k^E(t) \quad (4.3)$$

where ξ_G was drawn randomly at each time-step from a normal distribution with variance 0.1 and mean 0. Thus, the actual variance of the injected noise was signal dependent. Synaptic efficacy was limited to the range $[0, 1]$.

Functions for the inhibitory interneuron were similar. However, for simplicity, noise was not included in the interneuron's synaptic drive.

$$V^I(t) = \sum_j W_{I \leftarrow E} \cdot S_j^E(t) - W_{I \leftarrow I} \cdot S^I(t) + \gamma^I \quad (4.4)$$

$$F^I(t) = \frac{1 + \tanh V^I(t)}{2} \quad (4.5)$$

$$S^I(t + \Delta t) = S^I(t) + \frac{\Delta t}{\tau^I} \cdot (-S^I(t) + F^I(t)) \quad (4.6)$$

As noted above, this network has been well studied. In particular, the effects of starting conditions and extra-network input are well known (Wilson and Cowan, 1973; Samsonovich and McNaughton, 1997; Redish, 1999).

We will focus on three issues:

Issue 1: Random Noise vs. Stable Activity Mode. When random noise

is in the network, neurons will compete with each other until a group of neurons with neighboring preferred directions begins to dominate. This group wins and the network settles to a stable mode of activity representing one direction. Cells with preferred directions near this direction will have much higher firing rates than those far away. The direction represented will be a consequence of the random input (Wilson and Cowan, 1973; Kohonen, 1977). We performed this simulation by injecting random noise such that $IN_k(t)$ was uniformly distributed between 0 and 0.75 for any neuron k at any time t for 500 time-steps.

Issue 2: Rotation vs. Jump. Starting with a stable state (representing a single direction), an external synaptic current is provided to a group of neurons off-set from the represented direction, the represented direction will then shift toward the input, passing through intermediate values (Redish et al., 1996; Zhang, 1996; Samsonovich and McNaughton, 1997; Redish, 1999). Gaussian input with a standard deviation of 21.5° ($\sigma^2 = 20$ neurons) was provided with an amplitude of 1.5. This input was provided for 200 time-steps fixed at -86° (-1.5 rad), then shifted at 0.3° (0.005 rad) per time-step for 600 steps (for a total of 3 rad), and then maintained at $+86^\circ$ ($+1.5$ rad) for the last 200 time-steps of the simulation. This input-train resulted in a smooth 172° (3 rad) *rotation*.

In contrast, when the system is in a stable state and sufficient extra-network excitatory drive is provided far away from the represented direction (beyond 60° in our simulations), the firing patterns of the system will change to encode a new value without encoding intermediate values in the interim (Zhang, 1996; Samsonovich and McNaughton, 1997; Redish, 1999). Gaussian input with a standard deviation of 21.5° ($\sigma^2 = 20$ neurons) was provided with an amplitude of 1.5. This input was provided for 400 time-steps fixed at -86° (-1.5 rad), then shifted to $+86^\circ$ ($+1.5$ rad) for the last 600 time-steps of the simulation. This input-train resulted in a 172° (3 rad) *jump*.

Issue 3: Ambiguity. When a simulation is started with extra-network drive provided at two inputs, the system will settle to represent a unique direction depending on the random noise in the units (Redish, 1999). Two Gaussian inputs with an amplitude of 0.75 and a standard deviation of 21.5° ($\sigma^2 = 20$ neurons) were provided for 100 time-steps fixed at -86° (-1.5 rad) and $+86^\circ$ ($+1.5$ rad), to prime the network into an ambiguous, bimodal state. These inputs were then terminated and the two modes of activity competed for the remaining 400 time-steps of the simulation. This input-train resulted in two competing modes of activity separated by 172° (3 rad), and ended with a single winning mode at either -86° (-1.5 rad) or $+86^\circ$ ($+1.5$ rad).

4.3.2 Constructing the Training Set

In order to generate the tuning curves for neurons in the attractor network and probability distributions for our coherency and consistency measures, an attractor network simulation was run for 10000 time-steps. This stable training set is a smooth rotation and represents the null hypothesis that no jumps or ambiguous states are present in the network (see below).

Tuning curves

The activity $F_k(t)$ of each neuron k was averaged for each position of the input $x(t)$ to yield the neuron's tuning to all input directions:

$$\mathbb{T}_k = \text{mean}(F_k(t)|x(t)) \quad (4.7)$$

Empirical Cumulative Distribution

The *probability density function* f of a continuous random variable z follows the rule that the probability $P(z - \Delta z < z < z + \Delta z)$ of observing a value z within some range $\pm\Delta z$ is given by the formula:

$$P(z - \Delta z < z < z + \Delta z) = \int_{z-\Delta z}^{z+\Delta z} f(Z)dZ \quad (4.8)$$

The *cumulative distribution function* F of a continuous random variable z follows the rule that the probability $P(z < Z)$ of observing a value of z

smaller than the value Z is given by the formula:

$$P(z < Z) = F(z) = \int_{-\infty}^z f(z)dz \quad (4.9)$$

We approximated the cumulative distribution function of the data in the training set by finding the proportion of the values in the training set \hat{z} that were less than the value of interest z . This “empirical cdf” was determined by the following equation:

$$\text{cdf}_{\hat{z}}(z) = \text{Prop}(\hat{z} < z) \quad (4.10)$$

where $\text{cdf}_{\hat{z}}(z)$ is the *empirical cumulative distribution function* of variable z given the values in sample \hat{z} , and $\text{Prop}(\hat{z} < z)$ is the proportion of values in sample \hat{z} that are less than the value z .

4.3.3 Reconstruction

Many methods have been used for reconstruction (Georgopoulos et al., 1983; Wilson and McNaughton, 1993; Salinas and Abbott, 1994; Rieke et al., 1997; Zhang et al., 1998). For a uniform ring-topology such as the one used here, the simplest is the weighted vector mean (Mardia, 1972; Batschelet, 1981) also known as the *population vector* (Georgopoulos et al., 1983, 1988). Let \vec{v}_k be the preferred vector for excitatory neuron k , defined as the unit vector in the direction of the preferred direction ϕ_k of neuron k : $\vec{v}_k =$

$\langle \cos(\phi_k), \sin(\phi_k) \rangle$). The reconstructed vector \vec{R} is the weighted mean of the preferred vectors:

$$\vec{R} = \sum_k F_k^E \cdot \vec{v}_k \quad (4.11)$$

The reconstructed direction $\hat{\phi}$ is taken from the orientation of resultant vector \vec{R} from this vector sum. While we use this method for simplicity, it is important to note that coherency (defined below) works with any reconstruction method. It can also be used as a comparison between neural activity and expected neural activity given the real behavioral values.

4.4 Coherency

The first measure we will discuss is the coherency measure of Jackson and Redish (2003). This measure improves upon the one presented by Redish et al. (2000). While the coherency measure used by Redish et al. (2000) was capable of detecting a transition in an ensemble's representation, the measure itself was not statistically justified nor was it well characterized. The coherency measure needed rethinking and needed to be tested under standardized network conditions. Therefore, using a standard ring attractor neural network, an improved method was characterized to show that indeed network dynamics could be measured.

The primary objective of the “new and improved” coherency is to first measure how different the actual representational state is from the expected representational state and then set a statistical threshold beyond

which the observer believes the actual state is “significantly” different from the expected state.

4.4.1 The mathematics of the coherency measurement

As specified by Redish et al. (2000), an important step in measuring the population properties is to quantify the contribution of all neurons to our knowledge of the representational space. This was done by constructing an *activity packet*: a weighted sum of tuning curves. Jackson and Redish (2003) presented a modified computation of the activity packet as follows. If $T_k(x)$ is the tuning curve of neuron k over the representational space x and $F_k(t)$ is its firing rate at time t , then the ensemble’s activity packet $A(x, t)$ is given below:

$$A(x, t) = \frac{\sum_k T_k(x) \cdot F_k(t)}{\sum_k T_k(x)} \quad (4.12)$$

where \sum_k is the sum over all neurons in the ensemble.

Likewise, the redundant information contained in the tuning curves allows us to reconstruct the expected ensemble firing pattern given either the current value of x or the value of x inferred from the current neural firing pattern (designated \hat{x}). Thus, our best estimate given \hat{x} of what the firing rate of cell k at time t should be is the previously measured average firing rate of cell k when the animal experiences x . This is precisely the definition of the tuning curve: $T_k(x) = \mathbf{E}[F_k(t)|x(t)]$. So, the expected

activity packet $\hat{A}(x, t)$ is simply

$$\hat{A}(x, t) = \frac{\sum_k \mathbf{T}_k(x) \cdot \mathbf{E}[F_k(t) | \hat{x}(t)]}{\sum_k \mathbf{T}_k(x)} \quad (4.13)$$

or

$$\hat{A}(x, t) = \frac{\sum_k \mathbf{T}_k(x) \cdot \mathbf{T}(\hat{x}(t))}{\sum_k \mathbf{T}_k(x)} \quad (4.14)$$

The actual and expected activity packets can be compared using any of a variety of comparisons. As mentioned previously, Redish et al. (2000) used a dot product to measure the similarity between the actual and expected activity packets:

$$C_{DP}(t) = \hat{A}(x, t) \cdot A(x, t) \quad (4.15)$$

Here, we use a C to denote that the measure C_{DP} measures the similarity, or consistency, between the actual and expected representations. Subsequently, Jackson and Redish (2003) used a root-mean-squared-error (RMSE) measure of the difference between the actual and expected activity packets:

$$I_{RMS}(t) = \frac{\sqrt{\int_x (A(x, t) - \hat{A}(x, t))^2 dx}}{\int_x \hat{A}(x, t) dx} \quad (4.16)$$

where the integration is done over the entire representational space. We use I to emphasize that I_{RMS} measures how *incoherent*, or inconsistent, the actual representation is with the expected representation. Later, Johnson et al. (2005), used a measure of the variance of the difference between the

two activity packets:

$$I_{VAR}(t) = \frac{\text{var}_x(A(x,t) - \hat{A}(x,t))}{\int_x \hat{A}(x,t) dx} \quad (4.17)$$

Note, that I_{VAR} can be translated into the same units as I_{RMS} by taking the square-root of the numerator; this is equivalent to calculating the standard deviation of the difference between the two representations. We will denote this I_{STD} :

$$I_{STD}(t) = \frac{\text{stdev}_x(A(x,t) - \hat{A}(x,t))}{\int_x \hat{A}(x,t) dx} \quad (4.18)$$

It should also be noted that alternative measurements can be used to compare the two activity packets depending on the objectives of the analysis. In the case of C_{DP} , one might be more interested in testing how similar an ensemble activity pattern is to a hypothesized representation. I_{RMS} might be used when interested in identifying absolute differences across the population firing pattern, whereas I_{STD} (or I_{VAR}) might be invoked to measure relative differences between actual and expected ensemble activations since these measures subtract off the contribution of the mean difference between the two activity packets.

One of the primary contributions of Jackson and Redish (2003) was to provide a statistical interpretation of *coherency*. The coherency of the ensemble was defined as the probability of accepting the null hypothesis

that the actual and expected activity packets are the same, denoted

$$H_0 : A(x, t) = \hat{A}(x, t) \quad (4.19)$$

The probability of accepting the null hypothesis was found by empirically determining the probability distribution of the measurement of choice and calculating the probability of match between the expected and observed activity packets. *Coherency* is then defined as this probability of match. If the measure implemented detects differences between the activity packets, this probability is equal to the probability of seeing a larger difference between the actual and expected representation given the data in a carefully constructed training set I_{tr} : **Coherency** = $P(H_0) = 1 - \text{cdf}_{I_{tr}}(I(t))$. If this probability is sufficiently small, the actual and expected activity packets are more different than a large majority of the samples in our training set and we can reject the null hypothesis that the actual representation is the same as the expected representation.

Likewise, if the measure implemented detects similarities between the activity packets (e.g. the dot product measure), the probability of a match is equal to the probability of seeing a smaller amount of similarity between the actual and expected representation: **Coherency** = $P(H_0) = \text{cdf}_{C_{tr}}(C(t))$. If this probability is sufficiently small, the actual and expected activity packets are less similar than a large majority of the samples in our training set and we can reject the null hypothesis that the actual represen-

tation is the same as the expected representation.

The CDF can be calculated by running a simple experiment where the population is known to be representing the null hypothesis (i.e. that no dynamic anomalies are occurring in the network). This was done in the simulations below by forcing a slow rotation through five turns in both directions (See section 4.3.2). This stable training set represents the null hypothesis that no jumps or ambiguous states are present in the network.

Given this null hypothesis, we can set a threshold below which the expected and actual activity packets are *significantly* different. If the CDF is constructed correctly, system activity consistent with the tuning curves should evenly explore the range from 0 to 1. Thus, a significance threshold of 0.05 could be defined such that the system should spend 5% of the time with $I(t)$ values that give a p -value < 0.05 . For example, in a simulation with a continuous directional input at 180° , 4.9% of the values out of 5000 time-steps had RMSE values below the $\alpha = 0.05$ level.

For ensemble firing patterns that differ from the state predicted by the tuning-curves, the p -value will be very close to zero, indicating a system state that is incompatible with the activity predicted by the tuning curves, even after accounting for the variability due to the stochasticity of neuronal firing rates. Due to the non-linearity of the various $I(t)$ measurements discussed, a very small significance value is preferred to reduce false alarms. When the firing patterns are inconsistent with each other, they fall well below even reasonably small significance levels. Thus, we

will use a significance level of $\alpha = 0.001$ for the rest of the simulations.

In our simulations, all 3 cases could be cleanly dissociated using any of the coherency measures discussed above.

4.4.2 Coherency differentiates representational quality (Issue 1)

As noted above (Issue 1), when an attractor network is started from random, noisy values, it settles to a stable state such that only cells with preferred directions near a specific orientation are active. Figure 4.3 shows an example of a network settling. Reconstruction always provided an orientation and did not differentiate the random and settled states (Figure 4.3).

In contrast, coherency differentiated the random and settled states (Figure 4.3). Because of the non-linearities of the measure, coherency detected the time of settling accurately, displaying a stark difference between the two states. While in the random state, the coherency measurement showed that the random state was significantly different from the expected “bump” of activity ($p < 0.005$). After the network transitioned to a stable representational state at approximately time-step 342, coherency showed a higher probability of match.

4.4.3 Coherency can detect dynamic changes in network activity (Issue 2)

When an attractor network is in a stable state, providing synaptic input to neurons with preferred directions near the direction being represented by the network forces a rotation in the represented direction. Chaining this extra-network excitation to the represented direction forces the network to rotate continuously. In contrast, if sufficient extra-network excitation is provided to a group of neurons with preferred directions far from the encoded direction, the firing patterns of the system will change to encode the value consistent with the excitation without passing through intermediate values in the interim (Redish et al., 1996; Zhang, 1996; Samsonovich and McNaughton, 1997; Redish, 1999; Issue 2, above).

Reconstruction showed a smooth transition through intermediate orientations in both the rotation condition (Figure 4.4) and the jump condition (Figure 4.5). Reconstruction thus suggested that both of these transitions were simple rotations, yet the dynamics of these two transitions were fundamentally different. Coherency, however, detected the difference. In the jump condition, I_{RMS} and I_{STD} showed a strong transient increase at the time of transition (I_{RMS} : time-steps 562–609, $p < 0.005$; I_{STD} : time-steps 561–609, $p < 0.005$; see Figure 4.5B,C), but no corresponding increase during the rotation (time-steps 200–800, $p > 0.005$, see Figure 4.4B,C). Likewise, similar results were found using the dot-product measure of activity

packet similarity. In the jump condition, C_{DP} showed a strong transient decrease at the time of transition (time-steps 580–604, $p < 0.005$, see Figure 4.5D), but no corresponding increase during the rotation (time-steps 200–800, $p > 0.005$, see Figure 4.4D).

4.4.4 Coherency can be used to detect the resolution of ambiguity (Issue 3)

When the attractor network simulation begins with synaptic input to one group of neurons with nearby preferred directions, the network will settle to a bump of activity centered on that group. In other words, there will be one value represented by the network: the value consistent with the preferred direction of the most active neurons (Wilson and Cowan, 1973; Kohonen, 1977). However, when the attractor network simulation begins with synaptic input to two separate groups of neurons, the network will transiently display bimodal activity. In other words, there will temporarily be two bumps of activity, one at the center of each input, and therefore two values represented. The system will quickly settle to a state where only one value is represented; the result is dependent on the separation of the inputs (Redish, 1999).

When the two inputs are separated by a large enough distance (when the difference in preferred direction is greater than 60° in our network), the two values compete: the final represented value is equal to one or the

other of the two input values (See Figure 4.6).

The network can be seen as resolving ambiguity by forcing a multimodal input to settle to a unimodal firing pattern. The representation of the eventual steady-state of the system can be measured using a standard reconstruction algorithm. However, as in the earlier examples, reconstruction provided no information as to the time at which the system resolved the ambiguity. Reconstruction provided incorrect results at time-steps 0–110 in the distantly-separated case (Figure 4.6).

Coherency, on the other hand, did provide that timing information. When the representation was still ambiguous, the system was incoherent ($p < 0.005$, Figure 4.6). In the well-separated case (in which the inputs compete, Figure 4.6), coherency transitioned from a highly significant, low probability of match to within the non-significant range at time-step 113 (I_{RMS} and I_{STD}) and 107 (C_{DP}), indicating that the network had successfully resolved the ambiguity (Figure 4.6). The low probability of match ($p < 0.005$) at early time-steps indicated that reconstruction provided invalid results, while the high probability of match afterwards indicated that the reconstructed directions after approximately time-step 110 were valid. Thus, coherency successfully detected the time of transition to a stable state.

In summary, coherency is a robust measure of the state of a neural ensemble. It distinguished between consistent and inconsistent representational states by differentiating a random noise state from a coherent repre-

sentational state. Coherency detected transient changes in the state of the neural ensemble when it distinguished between smooth rotations, where intermediate values were encoded, and jumps, where intermediate values were not encoded. These two transitions were ambiguous using traditional reconstruction techniques. Coherency also differentiated between ambiguous bimodal states and stable unimodal states.

4.5 Ensemble Consistency

Like the coherency measure described above, Ensemble Consistency (EC) can be used to measure the quality of a neural representation within an ensemble. However, the EC method has two specific advantages: it does not require explicit knowledge of the neuronal response parameters (the tuning curves), nor does it require an explicit hypothesis of the encoded value. In other words, EC makes fewer assumptions about what neurons are encoding. Thus, it is possible to measure the consistency of the dynamic relationship between neurons in an ensemble with little or no knowledge about what they encode.

EC relies on the assumption that if neurons in an ensemble are working together to represent some behaviorally-relevant parameters, their firing rates will maintain the same relationships under similar conditions. This assumption, therefore, stems from the cell assembly hypothesis (Hebb, 1949).

4.5.1 The Mathematics of Ensemble Consistency

EC was assessed by defining an N -dimensional firing rate vector $f(t)$:

$$f(t) = \{f_1(t), f_2(t), f_3(t), \dots, f_N(t)\} \quad (4.20)$$

where $f_i(t)$ is the firing rate of neuron i at time t , and N is the number of neurons in the ensemble. If the ensemble firing patterns maintain the same relationship across neurons for similar conditions, these points should be grouped in the same region of firing-rate-space.

To test the similarity of an ensemble firing pattern to firing patterns previously recorded, one has only to find the probability of observing that point given the previously recorded firing patterns. In this way, one can measure how consistently an ensemble responds to a stimulus and behavioral conditions.

In order to measure the likelihood of observing a sample point relative to an expected distribution, we estimate the density of the distribution at the sample point. Density estimation is commonly done by dividing a space into bins and estimating the average density in each bin. This method has two major problems: its memory usage can be enormous, and the resolution is limited by bin size. For example, a data set consisting of just 15 simultaneously recorded neurons at a 15-bin resolution would require 3.5 exabytes of memory (3.5 billion GB).

To overcome this limitation, we adopted the method of Kernel Den-

sity Estimation (KDE) (Silverman, 1986). In this method, each point in the group is assigned Gaussian parameters to spread out its contribution to the overall density. Normalizing this distribution gives an estimate of the joint probability density distribution. The local density of the sample point can be calculated by evaluating the density contributed by each individual Gaussian in the group and summing the result. Thus, the group of training set points is transformed into a continuous estimate of the local density at the sample point. Figure 4.7 shows an example for a two cell ensemble taken from our training set.

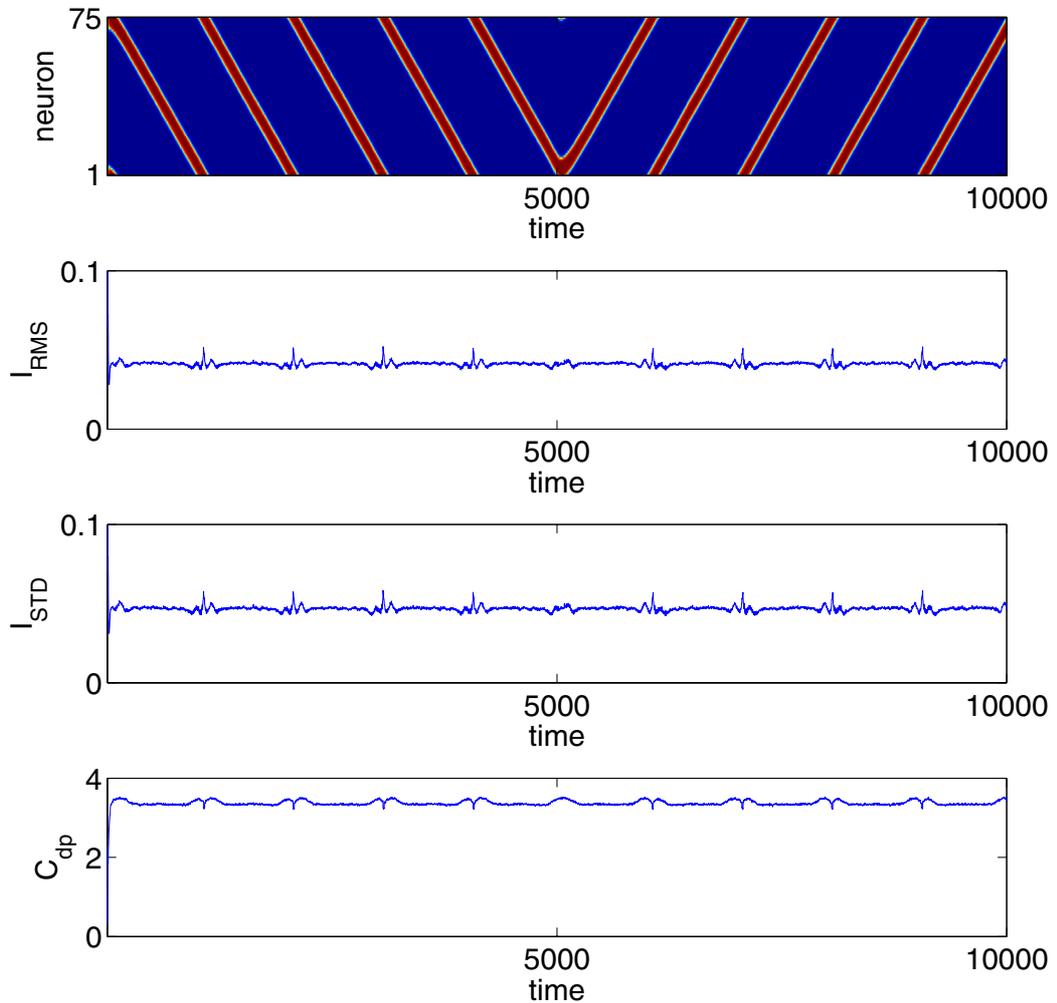


Figure 4.2: A smooth rotation induced in the network yields a stable training set. (A) The neural activity. Time is shown in milliseconds on the x -axis. Neurons ordered by their preferred direction ($0^\circ - 360^\circ$) along the y -axis, shaded according to their firing rate. (B) The I_{RMS} measure of inconsistency between actual and expected activity packets. (C) The I_{STD} measure of inconsistency between actual and expected activity packets. (D) The C_{DP} measure of consistency between actual and expected activity packets.

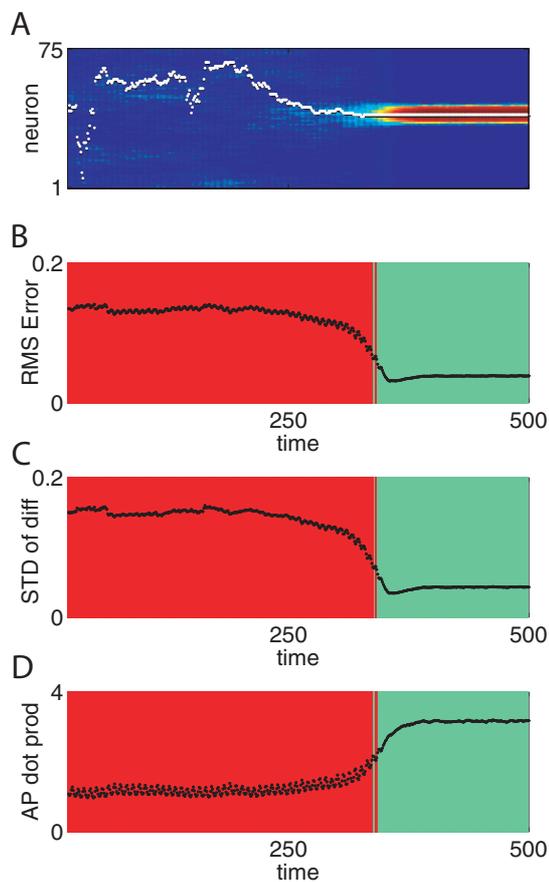


Figure 4.3: A simulation started with random input to the network settles to a stable state. (A) The neural activity. Time is shown in time-steps on the x -axis. Neurons ordered by their preferred direction ($0^\circ - 360^\circ$) along the y -axis, shaded according to their firing rate. White dots indicate the direction extracted from the population activity using population-vector reconstruction. Note, that the reconstruction algorithm yields a position whether or not there is an actual mode of activity present at that location. (B) The I_{RMS} measure of inconsistency between actual and expected activity packets. During the random state, the discrepancy between the actual and expected activity packets is high ($p < 0.005$, red zone). Upon reaching the stable state at time-step 342, the difference drops ($p > 0.005$, green). (C) The I_{STD} measure of inconsistency between actual and expected activity packets. As in B, during the random state, the discrepancy between the actual and expected activity packets is high ($p < 0.005$, red zone). Upon reaching the stable state at time-step 342, the difference drops ($p > 0.005$, green). (D) The C_{DP} measure of consistency between actual and expected activity packets. During the random state, the similarity between the actual and expected activity packets is low ($p < 0.005$, red zone). Upon reaching the stable state at time-step 343, the similarity increases ($p > 0.005$, green).

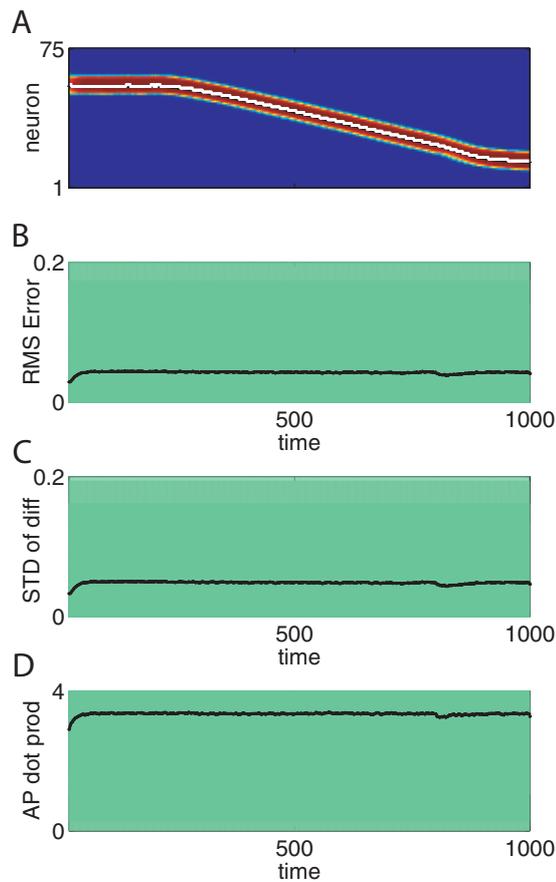


Figure 4.4: A smooth rotation induced in the network yields stable results. (A) The neural activity. Time is shown in time-steps on the x -axis. Neurons ordered by their preferred direction ($0^\circ - 360^\circ$) along the y -axis, shaded according to their firing rate. White dots indicate the direction extracted from the population activity using population-vector reconstruction. The reconstructed position follows the activity of the network faithfully. (B) The I_{RMS} measure of inconsistency between actual and expected activity packets. Throughout the rotation, the network maintains a stable state with a small difference between the actual and expected activity packets ($p > 0.005$, green). (C) The I_{STD} measure of inconsistency between actual and expected activity packets. As in B, throughout the rotation the network maintains a stable state with a small difference between the actual and expected activity packets ($p > 0.005$, green). (D) The C_{DP} measure of consistency between actual and expected activity packets. Throughout the rotation, the network maintains a stable state with a strong similarity between the actual and expected activity packets ($p > 0.005$, green).

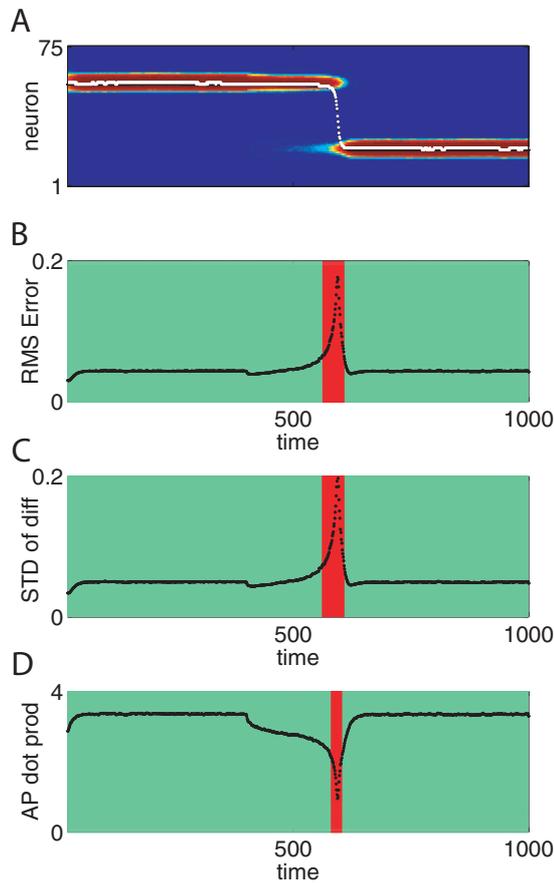


Figure 4.5: Coherency detects a discontinuity. (A) The neural activity. Time is shown in time-steps on the x -axis. Neurons ordered by their preferred direction ($0^\circ - 360^\circ$) along the y -axis, shaded according to their firing rate. White dots indicate the direction extracted from the population activity using population-vector reconstruction. Note that the reconstructed position shows a smooth rotation from the initial position of activity before the jump, through positions where there is no network activity, to the final location of activity after the jump. (B) The I_{RMS} measure of inconsistency between actual and expected activity packets. The discrepancy between the actual and expected activity packets is low during the stable state, before and after the jump ($p > 0.005$, green), but high during the transient bimodal activity state at the moment of the jump from time-steps 562–609 ($p < 0.005$, red zone). (C) The I_{STD} measure of inconsistency between actual and expected activity packets. As in B, the discrepancy between the actual and expected activity packets is low during the stable state, before and after the jump ($p > 0.005$, green), but high during the transient bimodal activity state at the moment of the jump from time-steps 561–609 ($p < 0.005$, red zone). (D) The C_{DP} measure of consistency between actual and expected activity packets. The similarity between the actual and expected activity packets is high during the stable state, before and after the jump ($p > 0.005$, green), but low during the transient bimodal activity state at the moment of the jump from time-steps 580–604 ($p < 0.005$, red zone).

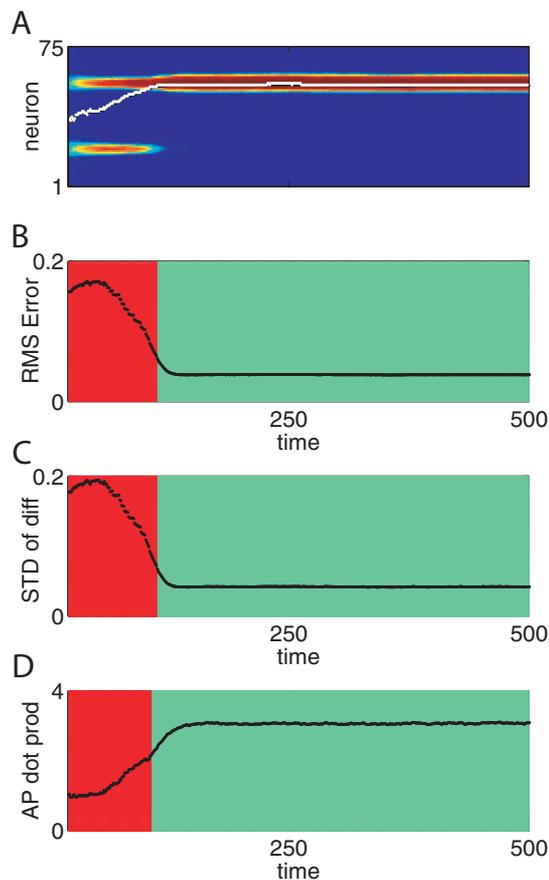


Figure 4.6: A simulation started with competing inputs settles to a single mode of activity. (A) The neural activity. Time is shown in time-steps on the x -axis. Neurons ordered by their preferred direction ($0^\circ - 360^\circ$) along the y -axis, shaded according to their firing rate. White dots indicate the direction extracted from the population activity using population-vector reconstruction. Note that the reconstructed position shows a smooth rotation from the mean position, where there is no network activity to the winning location. (B) The I_{RMS} measure of inconsistency between actual and expected activity packets. The discrepancy between the actual and expected activity packets is high during the initial bimodal state before the competition is resolved ($p < 0.005$, red zone), but low afterwards ($p > 0.005$, green). (C) The I_{STD} measure of inconsistency between actual and expected activity packets. As in B, the discrepancy between the actual and expected activity packets is high during the initial bimodal state before the competition is resolved ($p < 0.005$, red zone), but low afterwards ($p > 0.005$, green). (D) The C_{DP} measure of consistency between actual and expected activity packets. The similarity between the actual and expected activity packets is low during the initial bimodal state before the competition is resolved ($p < 0.005$, red zone), but high afterwards ($p > 0.005$, green).

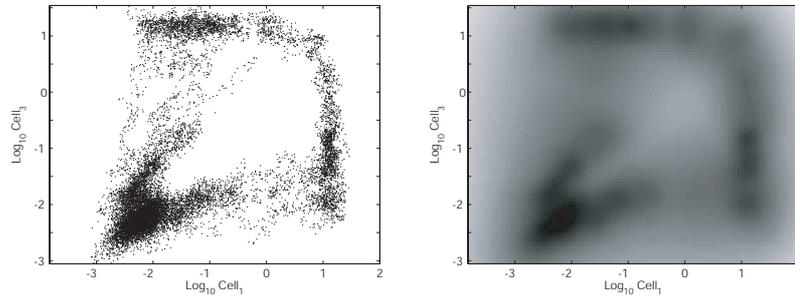


Figure 4.7: Density distribution of a two cell ensemble taken from our training set. The firing rate of cells 1 and 2 are represented by the x and y -coordinates, respectively. (left) The set of observed firing patterns \mathbf{f}_i . (right) The probability density distribution resulting from variable-width Gaussian kernel density estimation on \mathbf{f}_i . Black represents maximum density and white represents zero probability of observing a firing pattern. Note: plot is shaded by the logarithm of the density.

Mathematically, we write:

$$C_t = \sum_{i \in S} \beta_i \cdot \exp\left(-\frac{\|\mathbf{f}(t) - \mathbf{f}_i\|^2}{\sigma_i^2}\right) \quad (4.21)$$

where $\mathbf{f}(t)$ refers to the ensemble firing pattern of the point of interest at time t and \mathbf{f}_i refers to the i th firing pattern of the expected distribution S , also called the *training set*. β_i is the constant of normalization for the symmetric Gaussian associated with \mathbf{f}_i :

$$\beta_i = \frac{(\sigma_i^2)^{-(N/2)}}{\sqrt{2\pi} \cdot N_S} \quad (4.22)$$

where N is the number of neurons in the ensemble, N_S is the number of samples in the training set S , and σ_i is the width or standard deviation of the Gaussian associated with the i th sample of S .

Finally, we call C_t the *Ensemble Consistency* (EC) since it is a measure of how consistent $\mathbf{f}(t)$ is with \mathcal{S} . To improve our estimate of the density, we can let β_i depend on the n th-nearest neighbor distance of \mathbf{f}_i . First, the n th-nearest neighbor distance $D_i^{(n)}$ is found for each training set point i using the Euclidean distance measure. The standard deviation σ_i of the Gaussian associated with \mathbf{f}_i was defined as $\sigma_i = k \cdot D_i^{(n)}$. For this data set, $k = 4$ provided a good overlap between neighboring Gaussians. The order of the nearest neighbor distance, $n = 10$ was selected because the average n th order nearest neighbor distance was found to be very noisy for low n , to rise quickly as n increases, then to temporarily plateau. The beginning of this plateau represents how closely the points in the training set are packed and serves as a good reference for choosing the order, n , of the nearest neighbor. For our simulations, $n = 10$ provided a robust density estimation.

To reduce the dependence of EC on the overall ensemble firing rate, one could implement any number of normalizations. * One method would be to normalize each firing rate vector to unit length by dividing each vector by the modulus (square root of the sum of squared components). This normalization projects all points to the positive quadrant of the unit hypersphere. Another method could be to normalize all firing rate vectors to sum to one. This projects all points to the hyper-plane with vertices inter-

*The formulation and physical meaning of these normalizations arose from ideas generated in discussions with Adam Johnson, who first attempted to apply EC after normalizing all firing rate vectors to sum to 1.

secting all axes at 1. Both, methods will result in a bias toward the center of the positive quadrant of the space if firing rates are uniformly distributed prior to normalization. However, depending on the type of network inhibition, it is possible that these normalizations could result in uniform projections. For instance, an architecture that relies on a divisive inhibition that is proportional to the sum of the activities of the excitatory neurons would be best suited for using the later normalization (i.e. normalizing all firing rate vectors to sum to 1). In our simulations, we characterized EC applied to the raw firing rate vectors as well as EC applied to the firing rate vectors normalized to sum to 1.

In summary, Ensemble Consistency is the local probability density of the N -dimensional firing rate vector at a point in time t . Low probability densities represent firing patterns unlikely to occur under the conditions of the training set, and high probability densities represent firing patterns that were often seen in the ensemble under training-set conditions.

4.5.2 Statistical Justification of EC

As stated before, EC measures density not probabilities. Transforming these density measurements into a probability of seeing a density less than or equal to the observed density is complicated by the fact that these densities are unlikely to be unimodal, as can be seen in Figure 4.7. This multimodality occurs because neurons have overlapping tuning curves: the firing of one cell may occur either with or without the firing of the other cell.

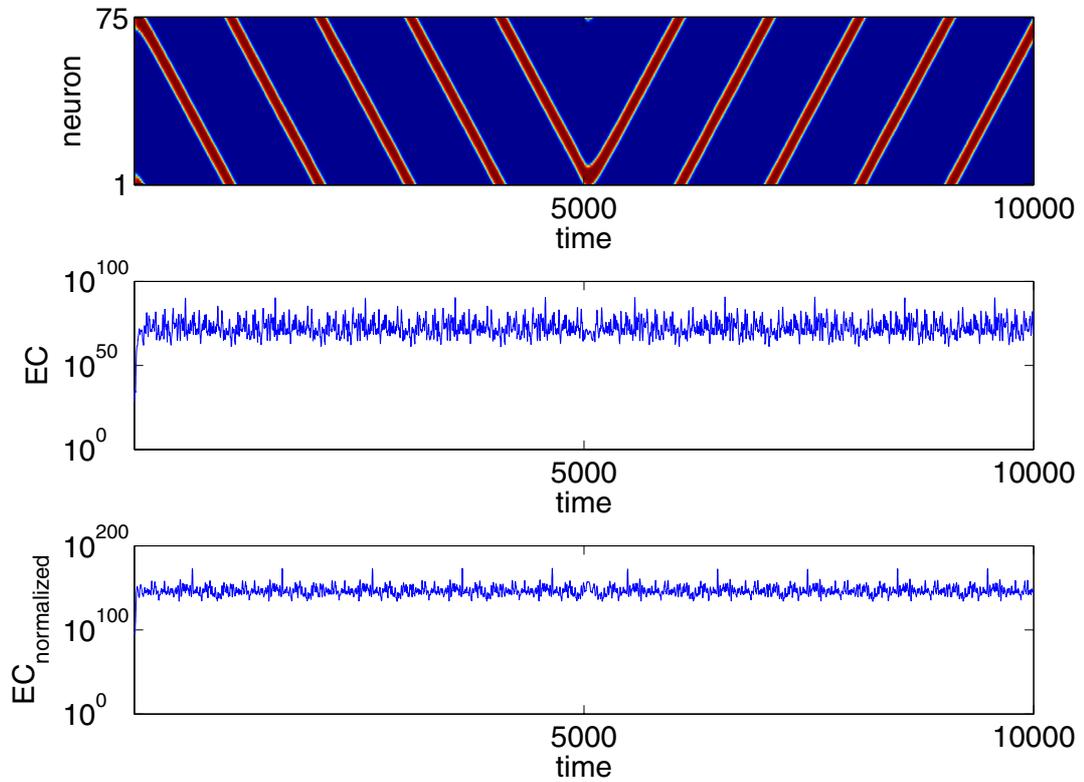


Figure 4.8: A smooth rotation induced in the network yields a stable training set. (A) The neural activity. Time is shown in milliseconds on the x -axis. Neurons ordered by their preferred direction ($0^\circ - 360^\circ$) along the y -axis, shaded according to their firing rate. (B) The EC measure of local probability density. (C) The EC measure calculated from normalizing each firing rate vector to sum to 1. Note: the same set was used as in figure 4.2

The likelihood of observing a particular density was estimated through a “leave-one-out” approach. For each sample \mathbf{f}_i in the training set \mathcal{S} , the density at \mathbf{f}_i was measured using Equation 4.21 summing over all samples in \mathcal{S} except for \mathbf{f}_i . This provided a set of densities found in the ensemble

under normal (training) conditions. The training set used here was the same set used in the Coherency section. As stated earlier, was created by forcing a rotation of the network through the full range of directions for five revolutions and then a reverse rotation through the same number of revolutions. This sampled the parameter space evenly. These densities were used to construct a $\text{cdf}_{EC_{tr}}$ of expected log-density values for one-sided significance testing as in the Coherency section (See Eq 4.10). The null hypothesis $H_0 = \mathbf{f}_i \in \mathcal{S}$ was rejected if the probability of observing a smaller density than the local density at \mathbf{f}_i was small. A stringent significance threshold of $\alpha = 0.001$ was used to reduce the false-alarm rate.

4.5.3 EC differentiates representational quality (Issue 1)

Figures 4.10–4.13 show the results of a 15-neuron ensemble taken from the population of 75 excitatory neurons in the attractor network simulation. The cells are ordered by their preferred direction with the preferred direction of each neuron being 4.8° from its neighbor. Similar results were obtained with randomly sub-sampled ensembles as long as the component tuning curves spanned the parameter space. Smaller ensembles were particularly dependent on fortuitous choices of the component neurons in the ensemble to span the parameter space.

Figure 4.10 shows a simulation started with random activity. As reviewed above, this network favors a single “bump” of activity, and soon

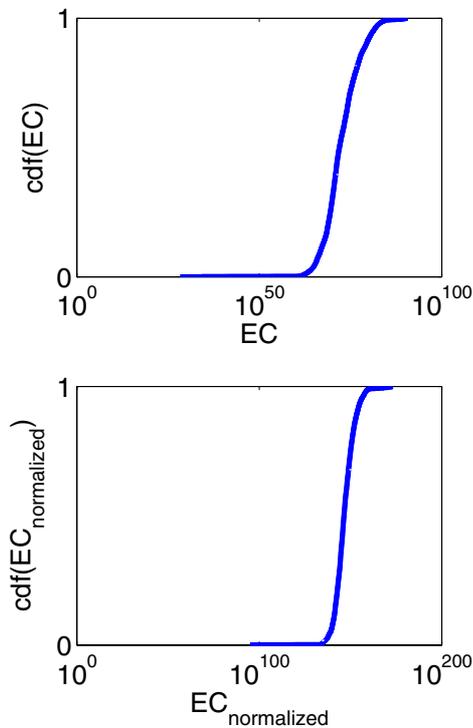


Figure 4.9: The empirically derived $\text{cdf}_{EC_{tr}}$ for both EC measures given the training set. (A) The $\text{cdf}_{EC_{tr}}$ of the EC measure of local probability density. (B) The $\text{cdf}_{EC_{tr}}$ of the EC measure calculated from normalizing each firing rate vector to sum to 1.

settled to a stable state with only a few neurons in an excited state. While the network was in the random state, EC was very near zero ($p < 0.001$), but transitioning to the stable state resulted in higher densities. Thus, EC differentiated between random and stable activity.

4.5.4 EC can detect dynamic changes in network activity

(Issue 2)

EC was stable during a smooth rotation of the network activity (see Figure 4.11). However, when the representational state jumped discontinuously,

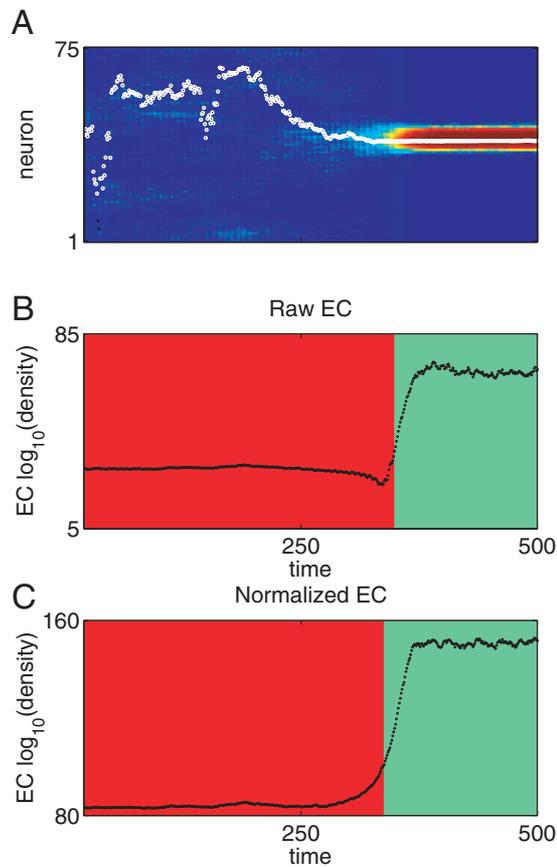


Figure 4.10: A simulation started with random input to the network settles to a stable state. (A) The neural activity. Time is shown in milliseconds on the x -axis. Neurons ordered by their preferred direction along the y -axis, shaded according to their firing rate. (B) Ensemble Consistency (EC). During the random state, the probability density, or EC, is low ($p < 0.001$, red zone). Upon reaching the stable state at time-step 348, EC rises ($p > 0.001$). (C) The same is true if the firing rate vectors are normalized. During the random state, the probability density is low ($p < 0.001$, red zone). Upon reaching the stable state at time-step 337, the density rises ($p > 0.001$).

EC detected the period of discontinuity (see Figure 4.12). Thus, EC differentiated between a jump in the representation and a stable rotation of the network firing pattern. Ensemble Consistency remained stable throughout the rotation (Figure 4.11). During the jump, EC was very low ($p < 0.001$, red zone, Figure 4.12), signifying a group of firing patterns far from the distribution of training set values (i.e. a dynamic instability that did not occur during training).

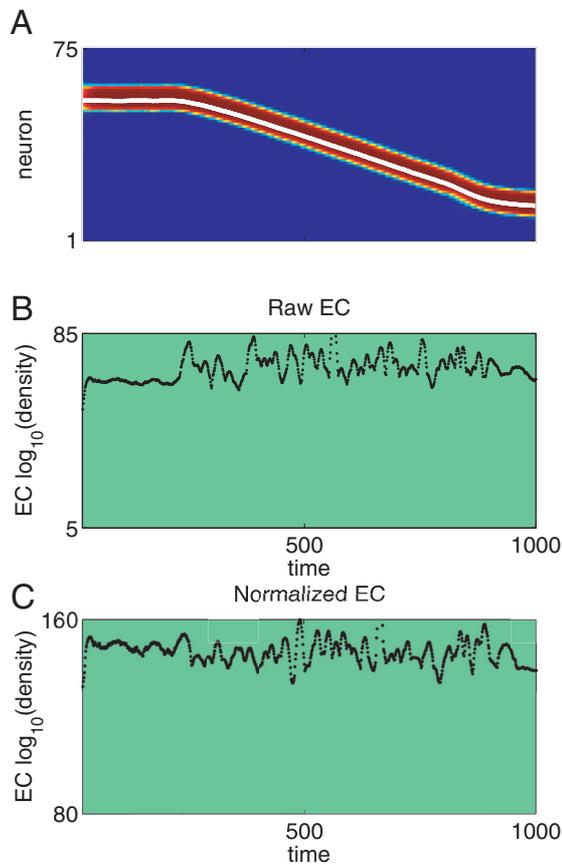


Figure 4.11: EC is stable throughout a rotation. (A) The neural activity rotates smoothly. Time is shown in milliseconds on the x -axis. Neurons are ordered by their preferred direction along the y -axis and shaded according to their firing rate. Whether using the raw firing rate vector (B) or the normalized firing rate vector (C), Ensemble Consistency (EC) shows a high probability density throughout the rotation ($p > 0.001$).

4.5.5 EC can be used to detect the resolution of ambiguity

(Issue 3)

Figure 4.13 shows a simulation started with two inputs separated by 170° . Since this network favors a single “bump” of activity, the two modes competed until one won and the network settled to a stable state with only a few neurons active. While the network was in this ambiguous state, EC was very low ($p < 0.001$), but transitioning to the stable state resulted in higher densities. Thus, EC differentiated between ambiguous and coher-

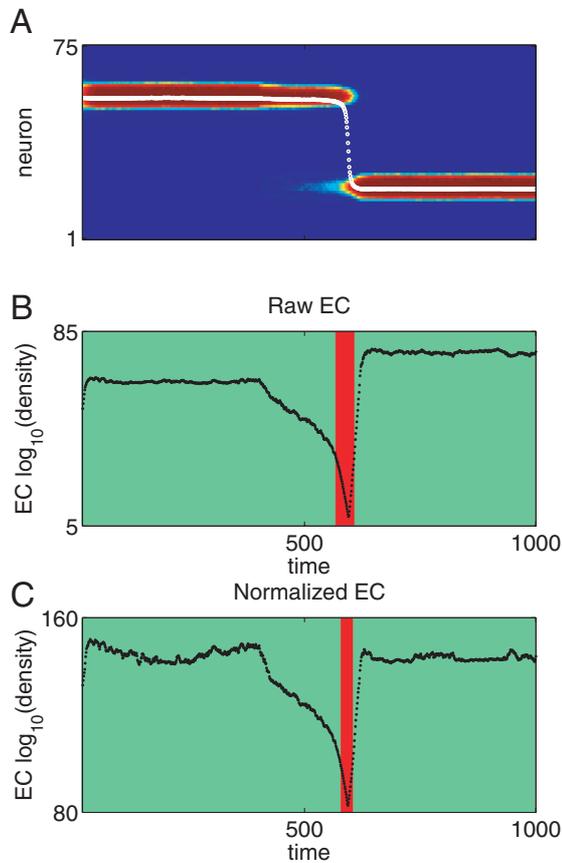


Figure 4.12: EC detects a discontinuity. (A) The neural activity jumps during a dynamic instability. Time is shown in milliseconds on the x -axis. Neurons are ordered by their preferred direction along the y -axis and shaded according to their firing rate. Whether using the raw firing rate vector (B) or the normalized firing rate vector (C), Ensemble Consistency (EC) shows a low probability density during the jump (Raw: time-steps 567–608, $p < 0.001$; Normalized: time-steps 578–605, $p < 0.001$; red zone). EC is high during the stable state, before and after the jump ($p > 0.001$, green).

ent activity.

In summary, the EC method can assess the consistency of an ensemble with little or no knowledge of the neural encoding. This is a powerful method for examining learning in structures that have complex representations of cognitive function and may be especially useful for examining deep brain structures where the behavioral parameters signalled by the neural activity are controversial or unknown.

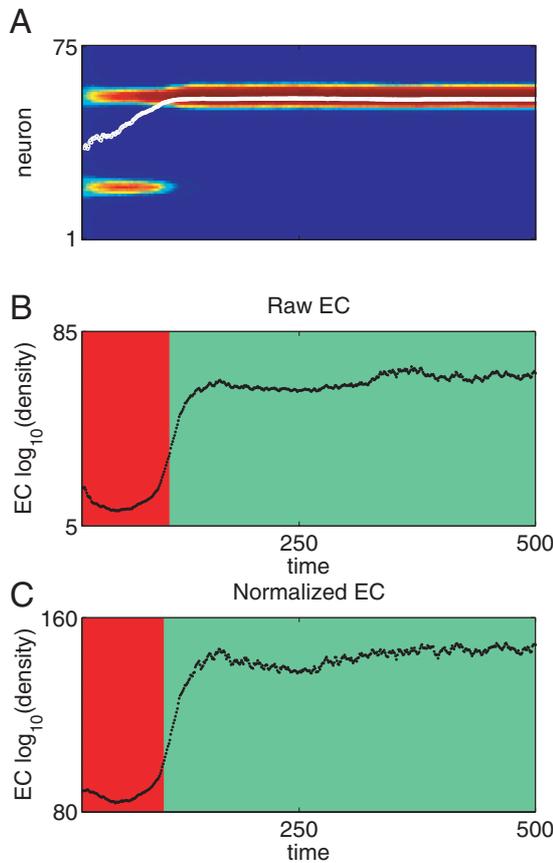


Figure 4.13: A simulation started with competing inputs settles to a single mode of activity. (A) The neural activity jumps during a dynamic instability. Time is shown in milliseconds on the x -axis. Neurons are ordered by their preferred direction along the y -axis and shaded according to their firing rate. Whether using the raw firing rate vector (B) or the normalized firing rate vector (C), the probability density is low during the competition (Raw: time-steps 0–113, $p < 0.001$; Normalized: time-steps 0–107, $p < 0.001$; red zone). Upon resolving the ambiguity and settling to a single value, EC rises ($p > 0.001$, green).

4.6 Discussion

The simulations in this chapter have shown that ensemble-level measures of network activity such as EC and Coherency provide useful information about network dynamics. This information is not readily available from reconstruction techniques such as population vector reconstruction. These ensemble quality measures can distinguish between random or ambiguous bimodal states in the network and the stable states predicted by the tuning-curves yielding statistical information that an observer can use

to test hypotheses of network function.

In order to avoid the strict assumptions made by standard parametric statistics, our methods require the assumption that the data used to construct the tuning curves and the training set be taken during a stationary period that represents the null hypothesis for an experiment. This is not a new or special assumption; it is a fundamental statistical concept and an assumption intrinsic to the concept of the tuning-curve. After this measurement is made, coherency is capable of detecting non-stationary events and other deviations from the control, or training, set.

It is possible, however, to apply these ensemble measures without constructing a training set. For example, in the case where fluctuations in the network's state are expected to coincide with other variables, one may find correlations between coherency or EC and the variables of interest. Johnson et al. (2005) used this approach in analyzing recordings from the rodent head-direction system. They found that low coherency values were strongly correlated with larger errors in reconstructing the animal's head-direction from the firing of neural ensembles in the post-subiculum. In a later chapter, this approach will be applied to examine the source of variability in hippocampal place-cells.

We are not the first to measure the reliability of a reconstruction estimate. There have been statistically based methods that use the length of the reconstructed population vector (Moore, 1980; Smyrnis et al., 1992; Ashe et al., 1993) but these methods require nearly complete revision when

attempting to apply them to other systems with different symmetries. For example, applying such a framework based on unimodal tuning curves to systems with bimodal, multimodal, or spatial tuning curves (O’Keefe and Conway, 1978; O’Keefe and Speakman, 1987; Sharp, 1996; Blair et al., 1997) would require a complete re-derivation of the statistics. In contrast, the methods presented in this chapter are general and will work for any type of tuning without modification.

While this chapter has focused on using coherency and EC to analyze an attractor network with a uniform ring topology, it should be noted that coherency is applicable to any neural system. Thus, multi-dimensional systems with bimodal, multimodal, or even heterogeneous tuning-curves are accessible to the coherency method. For example, some hippocampal place cells exhibit multiple place fields. The coherency method presented here requires no modifications to accommodate an ensemble with a mixture of multimodal and unimodal responses; the only requirement is the ability to calculate a tuning curve for each individual cell. Knowing the tuning curve allows the construction of an expected activity packet, to which the actual activity packet can be compared. Likewise, EC is also a general method, yet it circumvents the need to construct tuning curves. Instead of making assumptions about the encoding, Ensemble Consistency measures the dynamic relationship between neurons by using a density estimate to measure how close the current firing pattern is to other firing patterns observed in the training set. In this way, EC is like an abstracted

method of generalized template matching. This freedom from explicit assumptions of neuronal tuning allow EC to be applied to structures that have complex representations of cognitive function and may be especially useful for examining deep brain structures where the behavioral parameters signalled by the neural activity are controversial or unknown.

EC is limited, however, by its sensitivity to the training set. Because there is no normalization for the number of training set points collected for each stimulus, EC requires an even sampling of parameter spaces. Otherwise, densities of some firing patterns could be over- or underestimated. But, it is important to note that many behavioral tasks can be constructed such that the animal evenly samples the entire parameter space on every trial. Schmitzer-Torbert and Redish (2002) use a continuous T maze which requires the animal to make a series of T choices before receiving food reward along a return ramp without turning around. This required the animal to sample each portion of the task equally on each lap. Averbeck et al. (2002) recorded ensembles from monkeys copying geometric shapes. Each component of the shape is sampled equally on each trial. Note that this restriction only applies to the training set. Thus, a rat running a circular track will sample all head directions equally. Using the circular track as a training set, EC could be used with ensembles of head-direction cells to examine questions in an open-field in which directions are not sampled equally.

Any simultaneously recorded neural ensemble that can be used to re-

construct a behavioral variable can also be used to measure coherency. Reconstruction with simultaneously-recorded neural ensembles has been used in place cells (Wilson and McNaughton, 1993; Brown et al., 1998; Zhang et al., 1998) and in motor cortex (Salinas and Abbott, 1994; Averbeck et al., 2000). Accurate coherency measures can be taken from as few cells as are needed to provide adequate coverage of the parameter space by their tuning curves (i.e. enough to have reliable reconstruction). Our examples in the results section were based on 75-cell ensembles. The EC examples presented in Jackson and Redish (2004) used only 15-cell ensembles from similar simulations. The key to smaller ensembles is how well the neuronal responses span the parameter space. Note that EC is not dependent on uniformity of the neuronal responses. Thus, a 15-cell ensemble with tuning curves that span the space of potential stimuli, but are not uniformly distributed will provide a sufficient basis for the EC method. In our simulations, we found accurate detection of the three dynamic transitions discussed above with as few as 10 cells (data not shown). With modern recording technologies, simultaneous recordings of 20–100 cells are common. Coherency and EC can therefore be used with any of these recordings to detect differences in representational quality as well as transient representational events.

Chapter 5

Methods:

Experimental Procedures and Analyses

5.1 Introduction

In order to explore the biological relevance of concepts of network coherency developed in our simulations, we decided to apply these ideas to two poorly understood phenomena that result in highly variable firing in the hippocampus: sharp-wave associated population bursts and the excess-variability observed in the firing patterns of single units. An important factor in choosing the hippocampus as our model system was that the morphology of this structure allows large-scale chronic record-

ings of up to 100 or more neurons simultaneously in awake-behaving animals. Since the spatial firing of hippocampal units is task dependent, we recorded from rats as they ran on each of three task that varied in spatial and behavioral complexity: a linear shuttle task, a foraging task in a cylindrical arena, and a goal-directed spatial navigation task in the same cylindrical arena.

5.2 Subjects

Male Brown-Norway Fisher-344 hybrid cross rats were housed individually in a specific pathogen-free (SPF) vivarium maintained on a synchronous Day/Night cycle. Animals were handled daily for 15 min for at least one week prior to beginning behavioral training. One day prior to commencement of behavioral training, animals were denied access to food in their home cage while water access remained *ad libitum*. Subsequently animals received their full daily complement of food on the tasks based on their behavioral performance. All procedures were approved by the University of Minnesota IACUC and met all NIH guidelines for animal use in research.

5.3 Behavioral training

Food deprived rats were trained to run on a series of multiple tasks, including shuttling back and forth along a 137 cm by 15 cm linear track (LT),

foraging for scattered pellets in an 92 cm-diameter cylindrical arena (OF), or navigating to a small, invisible goal for food reward in the same cylindrical arena (OFG).

The linear track was constructed of gray carpeted ply-wood elevated above the experimental table top by 2 × 4 lumber blocks (Figure 5.1). Animals shuttled back and forth on the track to trigger food reward upon entry into the goal zone surrounding a feeder; this goal zone was not re-armed until the animal had triggered the opposite goal zone.

The cylindrical arena consisted of a flat, 94 cm diameter circular table top inserted into a tree tub with cylindrical walls and a cue card spanning approximately deg 90 (Figure 5.2). In the OF random foraging task, the rat retrieved food pellets as they scattered randomly at Poisson intervals. OFG used the same cylindrical arena as the foraging task, but delivered food only after the rat entered a 7 cm diameter goal region, randomly selected each day. Once the rat entered the goal region, a tone was played and the feeders were triggered. The goal was randomly placed each session, but remained constant within a session. Like the foraging task, the food scattered randomly upon dispensing. The goal was not re-armed until the rat had been outside of the goal and its 14 cm surrounding region for 4 seconds. Since OFG and OF used the same arena, the arena was wiped down with 70% isopropyl alcohol before each OF and OFG task to reduce the carry-over of odor cues from task to task.

Naive rats were trained on each task individually until proficient. Train-

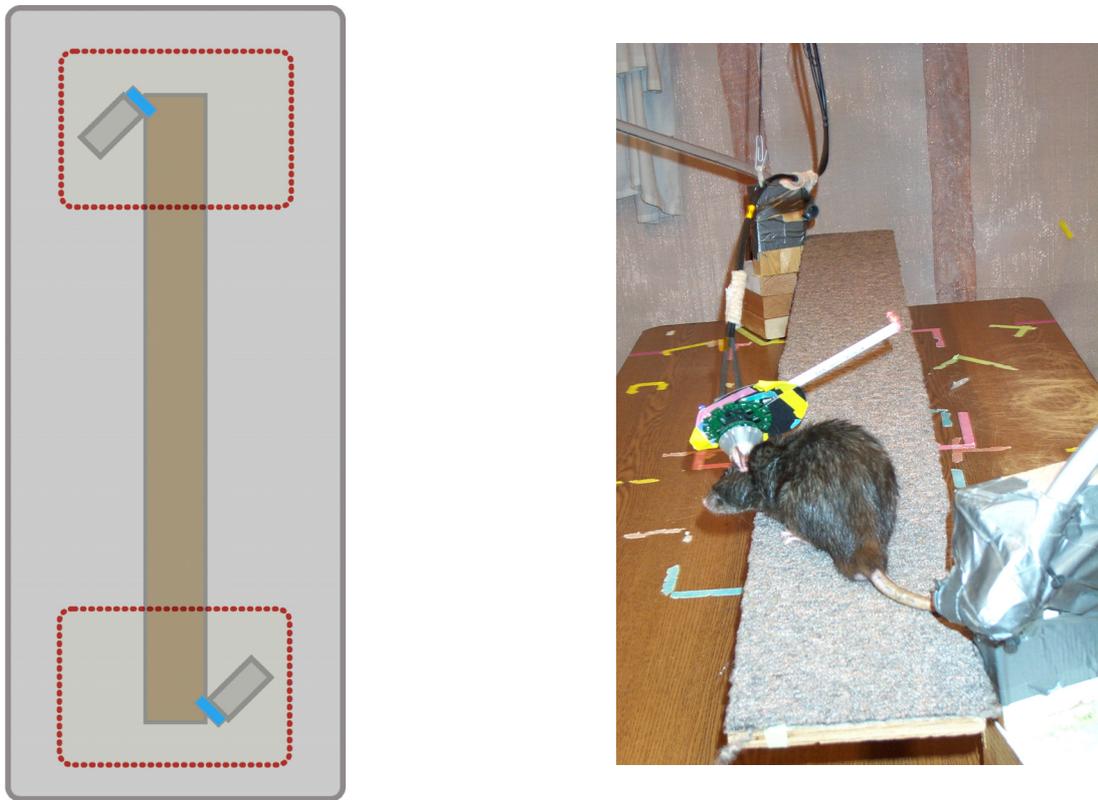


Figure 5.1: **The Linear Track Task.** (*left*) Diagram of the linear track. Red dotted lines indicate goal zone surrounding feeder locations. (*right*) Photograph of an animal on the linear track. Note feeder at near end on right side, and at far end on left side of the track. Animal in photograph has a boom for tracking head direction; the boom was not used on most animals in our data set.

ing began with single 30 – 40 min sessions on a single task (task training order counter-balanced across rats) until an animal was proficient on that task: full coverage of arena on OF; at least 30 successful goal entries on OFG; at least 50 laps on LT. Animals were then trained on the next task to proficiency. This continued until animals had been trained on each task individually. This usually took about one week per task.

Next, animals were trained for at least four days on the three-task protocol such that they encountered each ordering at least once: LT-OF-OFG; LT-OFG-OF; OF-OFG-LT; OFG-OF-LT. Thus, final-training and the post-

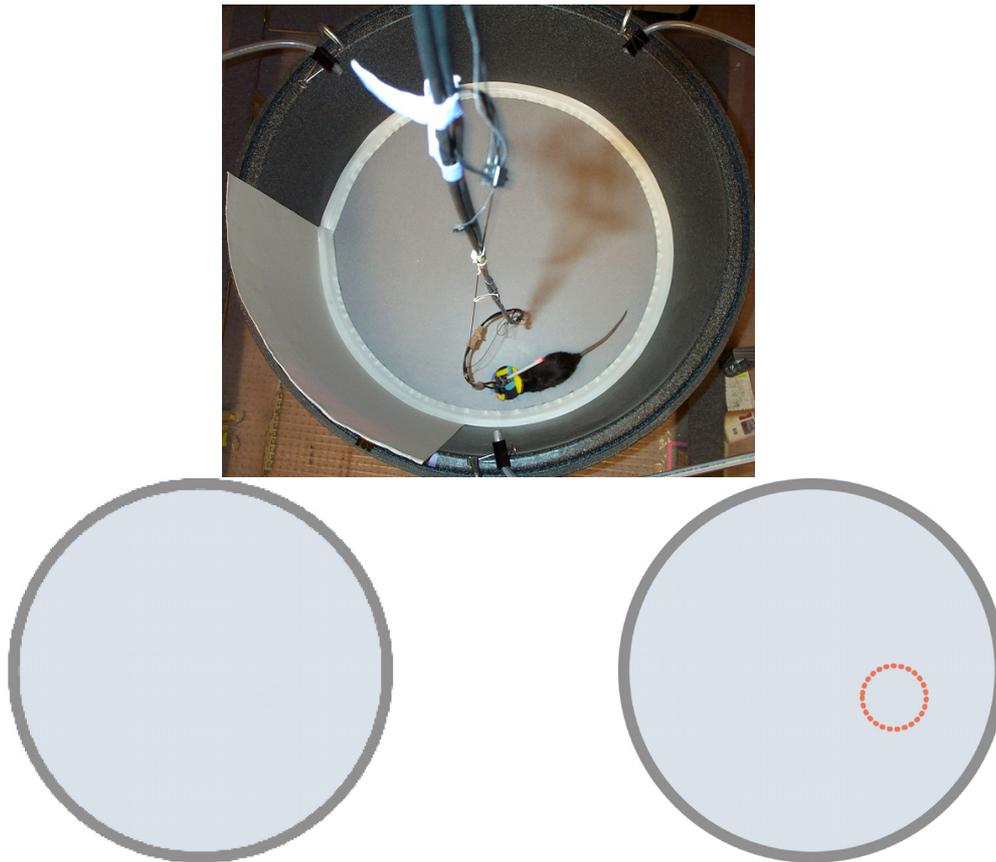


Figure 5.2: **The Open Field and Open Field Goal Tasks.** (*top*) Photograph of an animal in the arena. Note feeder tubes enter at the top of the arena wall near the bottom-center of the photo and on the top right and left sides. This configuration in combination with the pellet velocity and scattering after bouncing on the arena floor resulted in a highly uniform pellet landing distribution. Animal in photograph has a boom for tracking head direction; the boom was not used on most animals in our data set. (*left*) Diagram of the open field. (*right*) Diagram of the open field goal task. Red dotted lines indicate the invisible goal zone which varied from session to session.

implantation recording sessions consisted of 20 minute exposure to each of the three tasks pseudo-randomly ordered each day (goal location also varied pseudo-randomly each day) with a 5 minute rest period before and after each task.

5.4 Surgery

Once a rat was running proficiently on all three tasks, it was implanted with a 14-tetrode microdrive (Kopf Neuro-Hyperdrive, David Kopf Instruments, Tujunga, CA; 12 tetrodes and 2 references) at Bregma (-3.8 mm A/P, $2.0 - 2.5$ mm M/L). Rats were deeply anesthetized with an intraperitoneal injection of Nembutal (sodium pentobarbital, 40-50 mg/kg, Abbot Laboratories, North Chicago, IL), shaved on the scalp, and placed on a stereotax. A 0.5-2.0% isofluorane-oxygen mixture was then provided to maintain general anesthesia. 0.1cc Dualcillin (Phoenix Pharmaceutical Inc., Saint Joseph, MI) injections were administered to each hind limb. The scalp was disinfected first with alcohol then with Betadine (Purdue Frederick, Norwalk, CT). Skin and fascia were removed from the skull around the implantation site, and the wound was cauterized. Holes were drilled for the 8-9 jeweler's screws and 1-ground screw which were distributed around the implant to anchor it to the skull. Once the screws were in place, a craniotomy was opened above the target (A/P Bregma -3.8 mm, M/L $+2.0$ mm; Paxinos and Watson, 1998) using a surgical trephine (Fine

Science Tools, Foster City, CA), and the hyperdrive was lowered into place. A steel wire from the hyperdrive ground terminal was connected to the steel wire soldered onto the ground screw using Amphenol pins. A Silastic (Dow Corning 3140) barrier filled the space between the hyperdrive bundle and the skull. Dental acrylic (Perm Reline and Repair Resin, The Hygenic Corp., Akron, OH) was used to fix the hyperdrive to the bone screws and seal the wound. After removal from the stereotax, 3 cc saline was administered subcutaneously. Some rats received regimens of 0.1 cc Baytril (2.27% enrofloxacin, Bayer Corp., Shawnee Mission, KS) injected subcutaneously each day following surgery for three days. Animals received 0.8 cc Children's Tylenol orally immediately upon waking and in their water supply (25 mL mixed in 0.275 L water) during recovery.

5.5 Electrode Positioning

Recordings were then taken from the pyramidal layer of the CA1 region of hippocampus. The pyramidal layer was identified by the presence of strong high-frequency (100–200 Hz) ripples (Ylinen et al., 1995).

5.6 Recordings

Recordings were carried out in 10 foot × 10 foot room enclosed in copper screen. All electrophysiological and video tracker recordings were digi-

of 0.3-1.0M Ω . Each tetrode was individually drivable.

Extracellular action potentials were recorded at 32kHz in 1 ms when the voltage crossed a threshold set by the experimenter on any of the four channels on a tetrode. The one millisecond window of data was taken for each action potential. The signals were first amplified at the headstage with unity gain amplifiers, then passed through multistrand cables and a commutator before reaching variable gain amplifiers (1-50,000x). There, they were band pass filtered from 600-9000Hz for spike recordings using 48 channels of a Neuralynx 64 channel Cheetah system, or filtered from 1-475Hz and sampled at 2 kHz for local field potential recordings (LFP) using 16 channels of a Neuralynx 64 channel Cheetah system (Neuralynx, Tucson, AZ). Binding of recording cables due to rotation of the rat was minimized by a torque-sensing, motorized 72-channel commutator (Neuralynx, Tucson, AZ; Dragonfly, Ridgeley, WV; AirFlyte, Bayonne, NJ).

5.6.2 Behavioral tracking and Behavioral Control

The positions of LEDs mounted on the animal's head stage were recorded by a camera mounted in the center of the recording room's ceiling. The video frame data was sampled at 60 Hz and digitized and time-stamped by a Cheetah data acquisition system (Neuralynx, Tucson, AZ); pixels that crossed thresholds set by the experimenter were recorded to disk. Real-time position data was accessed by in-house behavioral control software implemented in Matlab (The Mathworks, Natick, MA). This soft-

ware used the serial ports to communicate with an experimental control box (constructed by JCJ) to trigger food delivery (45 mg pellets: Research Diets, New Brunswick, NJ; food dispensers: Med-Associates, St. Albans, VT) and simultaneously signal the Cheetah recording system for a synchronous food delivery time-stamp (each feeder had a unique digital identification).

Position data were then pre-processed for post-hoc analysis by extracting the center of mass of all supra-threshold pixels. Video interlacing effects were removed from the data through linear interpolation of odd and even position samples (two 30 Hz time-series) to produce two 60 Hz time series, which were then averaged to yield a single, stable 60 Hz time series.

Lap Times: Linear Track Since the linear track was aligned along the x -direction of the video data, the x -position was taken as the linearized 1-dimensional projection of the animal's behavior. The x -velocity was calculated using a 64th order low pass FIR filter with a 2 Hz high frequency cut off by filtering forwards and backwards (to eliminate phase shifts). Plots of the x -position versus x -velocity displayed a clustering of low velocity at the track ends. These clusters were selected manually for each session and lap times were either defined as the time from departure from one end of the track to the next departure from the same end. Inter-lap-intervals were defined as the time spent in these low-velocity clusters at the track ends

(1.7 cm/s \pm 2.5 cm/s; mean speed \pm SD).

Goal Entry Simulations To quantify goal-directed behavior, the animal's behavioral output with respect to the goal could be compared to the same animal's behavior with respect to the same goal on a different task (e.g. foraging in the open-field) or to a pseudo-goal region of the same size randomly selected within the arena. These pseudo-goal-entry simulations were constructed the same as the task-control code such that after the animal's position entered the pseudo-goal, the entry time was saved and a 4-second time-out was imposed before the next pseudo-goal-entry was recorded. Thus, actual goal-entry behavior could be compared to behavior with respect to locations with no task-related significance.

5.7 LFP analysis

Generally noisy channels were removed from the analysis. Any event of max/min voltage (during which the amplifiers reached the + or - rails) on any channel was removed from all channels including 0.5 s before and after.

5.7.1 SW detection

SW events were extracted by down-sampling the LFP traces by a factor of 2 (using an anti-aliasing low-pass filter), and bandpass filtering from 100-

250Hz. Amplitude for each trace was found via Hilbert-transform and then averaged across traces. The distribution of log-transformed average amplitude was used to find samples more than 2.5σ from the mean power. Higher and lower values of σ (e.g. 2σ , 3σ , 4σ) yielded qualitatively similar results. Visual inspection of a subset of the data revealed ripple events synchronous across LFP channels. Threshold crossings shorter than 20ms were removed, the remaining events were concatenated if less than 100ms apart. 20 ms was added to the beginning and end of each SW to capture the tails of the SW. Raw threshold detections were also analyzed and yielded qualitatively similar results, likewise merging threshold crossings less than 100ms apart before discarding short (< 20 ms) events yielded qualitatively similar results. To reduce the possibility of non-LIA, high-frequency events contaminating the analysis we removed all SW events detected during high-theta/low-delta periods. Thus, the SW represented in our analyses would be most analogous to the immobile sharp-wave ripple (iSWR) events of O'Neill et al. (2006). The results in this paper are from this conservative set, however there was little qualitative difference introduced by including all SW events. (See Figure 5.4)

5.7.2 Theta Detection

A similar method was used for detecting theta epochs. Theta times were extracted by down-sampling the LFP traces by a factor of 5 (using an anti-aliasing low-pass filter), and bandpass filtering from 6-10Hz to ob-

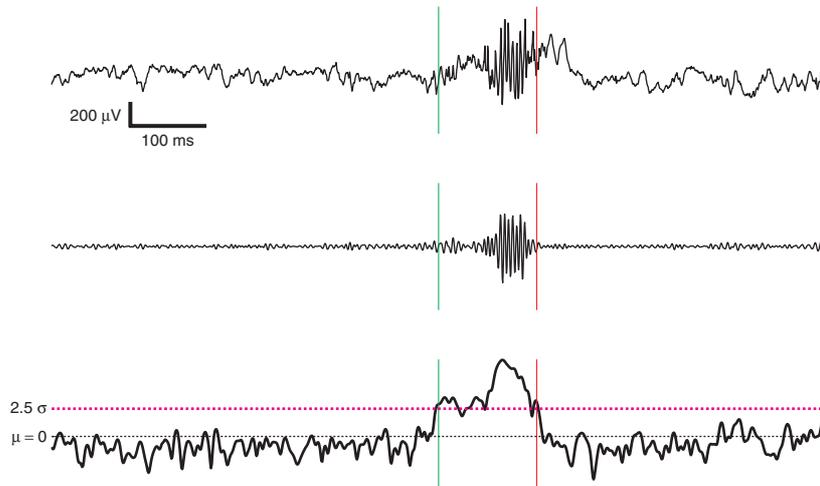


Figure 5.4: **Sharp wave ripple event detection.** Time is on the x -axis, voltage is plotted on the y -axis; scale bar indicates 100 ms and 200 μV , respectively. (*top*) The raw voltage trace. (*middle*) Same voltage trace as above band pass filtered from 100 Hz to 250 Hz. (*bottom*) Z-score of instantaneous power averaged across all EEG channels. Thick, red dashed line indicates the detection threshold of 2.5 standard deviations above mean (session mean of the average signal). Note that power trace drops below threshold; threshold crossing times were found then merged if they occurred within 100 ms of another event (solid green and red lines)

tain theta-band signals, and from 2-4Hz to obtain delta-band signals. Amplitude for each trace's band was found via Hilbert-transform and then averaged across traces to obtain two averaged signals: an average theta-band amplitude and an average delta-band amplitude. The distribution of the log-transformed ratio (theta/delta) of average amplitudes was used to identify samples with a low power-ratio more than 1σ from the session mean, these were taken to be *non-theta* brain states. Higher and lower values of σ (e.g. 0.5σ , 1.5σ , 2σ) yielded qualitatively similar results. Visual inspection of a subset of the data revealed low theta amplitude epochs

that clustered at locations of immobility (i.e. the linear track ends). These non-theta epochs were concatenated if less than 500 ms (the low frequency cutoff for the delta-band) apart, and events smaller than 100 ms (the high frequency cutoff for the theta-band) were removed. Raw threshold detections were also analyzed and yielded qualitatively similar results. *Theta* epochs were taken as the complementary set of times; these high theta band power and low delta band power epochs tended to coincide with times when the animal was moving.

5.8 Unit Recording

Tetrodes allow the discrimination of extracellularly recorded spikes from multiple different units by comparing spike waveform properties across channels. This is because the voltage traces resulting from spikes from a neuron will be distributed differentially across the four channels depending on the spatial relationship to the tetrode (See Figure 5.5). Likewise, spikes from other neurons will form different patterns of voltage distributions across the tetrode since two neurons cannot occupy the same physical location. These patterns can be exploited through clustering techniques to isolate the spikes of putative individual units, thereby allowing the recording of up to 20 or more neurons per tetrode depending on the packing of the local neuronal population.

5.8.1 Clustering

Waveform features were calculated from the 1 ms, 32-sample spike waveforms and clustered using a combination of automatic and manual clustering algorithms (MClust, A.D. Redish et al, <http://mclust.sourceforge.net>, Klustakwik, K.D. Harris et al, <http://klustakwik.sourceforge.net>).

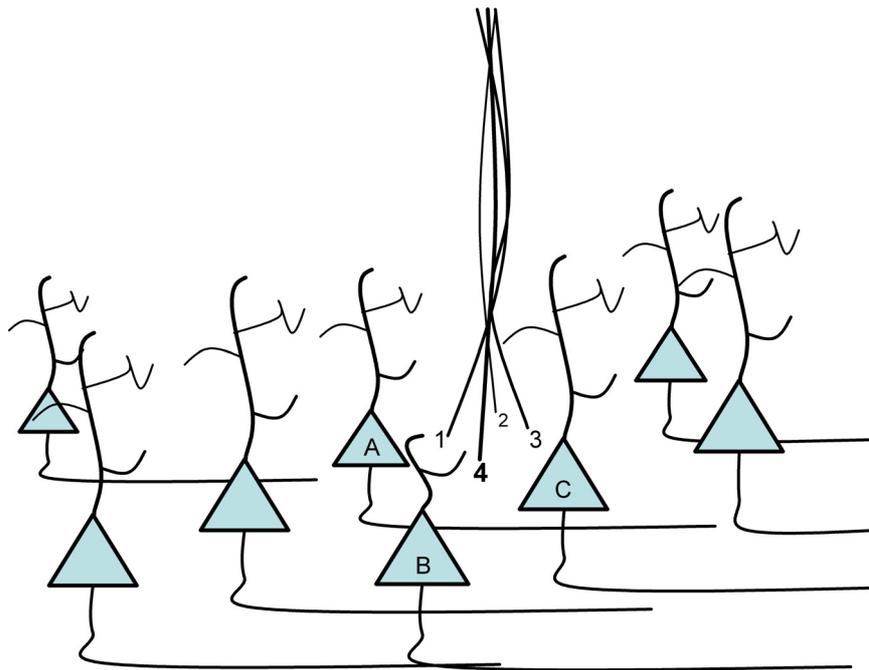


Figure 5.5: **tetraode diagram**. The tetraode is depicted as 4 black strands intertwined. Voltage traces resulting from spikes from neuron A will be larger on channel 1 than on channels 2 and 4 and much larger than channel 3. Likewise, spikes from neuron C will be larger on channel 3 and 4 than on channels 1 and 2. Simply plotting the peak voltage of waveforms recorded on this tetraode on the Channel 1 - peak voltage \times Channel 3 - peak voltage axes will result in a clustering of spikes from neuron C at large values for voltage on channel 3 and small values for channel 1 and the clustering of spikes from neuron A at large values for voltage on channel 1 and small values for channel 3. Thus, the spikes from neuron A will be separated from C on this projection of waveform features. Continuing this process will allow the separation of spikes from neuron B from those emitted by A and C.

net). Only clusters with firing rates below 2 Hz were used; these tended to have the characteristic bursting inter-spike interval typical of hippocampal pyramidal neurons. All analyses were also run only including clusters with high quality (Schmitzer-Torbert et al., 2005) ($L_{ratio} < 0.2$ and $ID > 15$). For example, cluster 10 on Figure 5.6 would be excluded from further analysis based on low cluster quality since it has $L_{ratio} > 0.2$ (actual $L_{ratio} = 0.76$) and $ID < 15$ (actual $ID = 10.5$). Qualitatively similar results were obtained when the full data set was used. Neurons included in our analyses had few interspike intervals less than 3 ms and exhibited a strong bursting phase. *Neurons whose waveforms were not stable

***Cluster Isolation Quality:** The separation of a cluster from other signals recorded on the same tetrode was quantified using the L_{ratio} and isolation distance (ID) metrics previously described (Schmitzer-Torbert et al., 2005). Briefly, L_{ratio} and ID are measures of the number of points surrounding a cluster and the distance of the cluster from all other points (i.e. all other recorded signals), respectively. Both measures are calculated using the 8-dimensional space of waveform energy (on each of 4 channels) and the first principle component of the waveform (on each of 4 channels). First, calculate the Mahalanobis distance $D_{i,C}^2$ of point i from the center C of the cluster of interest as

$$D_{i,C}^2 = (x_i - \mu_C)^T \Sigma_C^{-1} (x_i - \mu_C) \quad (5.1)$$

where μ_C is the vector pointing to the mean of cluster C in the 8-D feature space, x_i is the location of point i in the feature space, and Σ_C is the covariance matrix of cluster C . L_{ratio} is then calculated using the following formula:

$$L_{ratio} = \frac{\sum_{i \notin C} 1 - \text{CDF}_{\chi_{df}^2}(D_{i,C}^2)}{n_C} \quad (5.2)$$

where the sum is over all points i not contained in cluster C , $\text{CDF}_{\chi_{df}^2}$ is the cumulative distribution function of the χ^2 distribution, df is the number of degrees of freedom (the number of dimensions in our feature space), and n_C is the number of spikes in cluster C . ID is simply the Mahalanobis distance $D_{i,C}^2$ of the n_C^{th} nearest point not included in the cluster. L_{ratio} and ID are highly correlated with the Type I and Type II error rates associated with clustering spikes. ID is most correlated with Type I clustering errors (incorrect inclusions or false alarms), and L_{ratio} is most highly correlated with Type II

enough across the entire session to reliably cluster or whose waveforms drifted towards the spike detection threshold were not clustered and were excluded from further analysis.

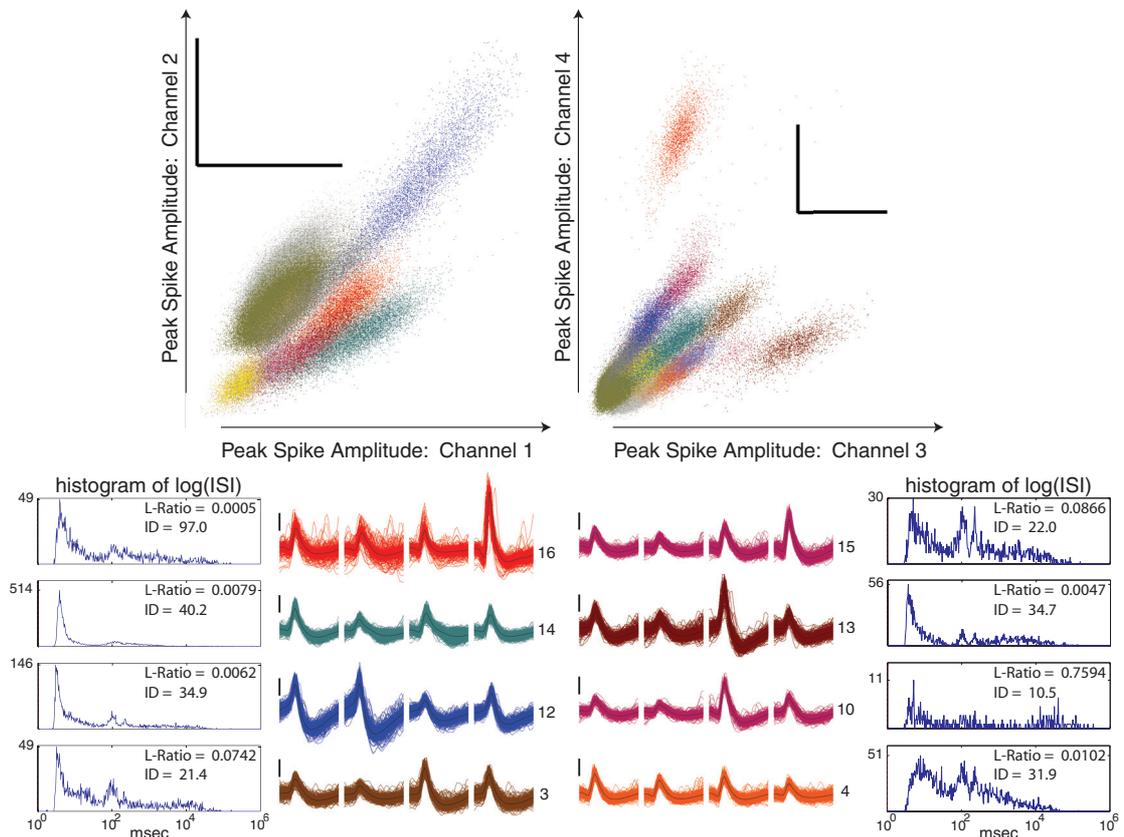


Figure 5.6: Example tetraode recordings: clustering, waveforms, and firing patterns. Spikes were clustered according to multiple waveform features including peak spike amplitude, energy, and principle components. Clustered spikes are shown with different colors for each cluster on two projections: peak spike amplitude on channel 1 versus channel 2, and peak amplitude on channel 3 versus channel 4. One millisecond waveforms for 8 of the 16 separable clusters are shown below color coded by cluster color. Interspike interval (ISI) histograms from each neuron are typical of hippocampal pyramidal neurons. Cluster quality values are inset within each cluster's ISI histogram. All scale bars are 100 μ V. Data from R048-2004-07-14 TT7.

clustering errors (incorrect exclusions or misses).

Chapter 6

Overdispersion: Behavioral and Network Sources of Excess Variability in Hippocampal Place Cell Discharge

Abstract

While the firing of hippocampal pyramidal neurons is highly specific spatially, the temporal variability of single unit firing patterns is actually quite large. In this chapter, we replicate previous studies by demonstrating a

similar degree of variability in our hippocampal recordings taken from rats as they forage for food and find that goal-directed behavior in the same environment reduces this variability in a manner similar to previous reports. We then explore the hypothesis that this variability observed at the single-unit level is related to modulation at the network-level resulting from the animal using multiple spatial reference frames. We demonstrate that hippocampal responses on the linear track, a task that induces reference-frame switching in the hippocampus, can emulate the highly variable patterns observed during our foraging task. Finally, we apply the concept of the cell-assembly to derive multiple spatial maps and demonstrate that the network-level switching between these maps is consistent throughout the task and depends on the task-related goals of the animal.

6.1 Introduction

The rodent hippocampus is most noted for the strong spatial selectivity of its pyramidal cells, or *place cells* (O'Keefe and Dostrovsky, 1971). Much work has revealed many aspects of hippocampal network dynamics including the modulation of place cell activity by well-known brain states such as the theta rhythm and sleep (Vanderwolf, 1971; O'Keefe and Nadel, 1978). There are also different theories about the how the manner in which information is processed by the hippocampus should modulate place cell firing (O'Keefe and Nadel, 1978; Buzsáki et al., 1983; Redish, 1999). For

example, if the hippocampus uses multiple reference frames to represent the environment while an animal is performing a navigational task, then switching reference frames constitutes switching the cues to which each pyramidal cell is responsive and will, therefore, switch the distribution of cells that are active (Touretzky and Redish, 1996; Redish and Touretzky, 1997; Redish et al., 2000). There is some evidence that this reference frame switching may be a regular phenomenon. Fenton and Muller (1998) found that while place cells are reliable spatially, they exhibit much more temporal variability than would be predicted by a Poisson process. This excess variability, or *overdispersion*, is task dependent (Olypher et al., 2002). It has been suggested that this overdispersion may be the result of the rat switching reference frames at a mean rate of 1 to 2 times per second (Lánksý et al., 2001; Olypher et al., 2002).

Little is known about the hippocampal dynamics of reference frame switching. The primary goal of this chapter will be to examine the dynamics and variability of hippocampal place cell firing patterns in the context of the rest of the network. To explore these issues, we begin by replicating the the original findings of Fenton and Muller (1998) in the open-field foraging task (OF). We then attempt to replicate the task-depedent decrease in overdispersion observed by Olypher et al. (2002). To explore the possibility that reference-frame switching may contribute to overdispersed firing, we introduce the linear track, a task with directionally-dependent hippocampal firing that resembles a switch between reference frames. Next,

we consider behaviorally-related modulations in place-cell firing such as velocity dependence, path variability, and path repetition. Finally, we apply the concept of the cell-assembly (Hebb, 1949) to split ensemble spatial firing-patterns into two separate maps. We show that while coherency is unable to detect a network-wide fluctuation given whole-session average firing maps, the times when network firing patterns were clustered to one or the other map represent significant switches in the coherency of one or the other map. These switching times are strongly modulated by the behavioral parameters of our goal-directed tasks and yield lower variability than is observed in session-average maps.

6.2 Tuning Curve Construction

The tuning of a cell is the average or expected firing rate of a neuron measured over a given behavioral variable; in this case, this variable is the animal's spatial location. Tuning curves were constructed by binning the task area into $11 \text{ pixel} \times 11 \text{ pixel}$ bins and creating two 2-dimensional histograms: a histogram of the number of spikes emitted in each bin and a histogram of the number of video-tracker samples in each bin. The occupancy time for each bin was determined by dividing the number of positions samples per bin by the video sampling rate. The firing rate per bin was determined by dividing the spike count in each bin by the occupancy time in that bin.

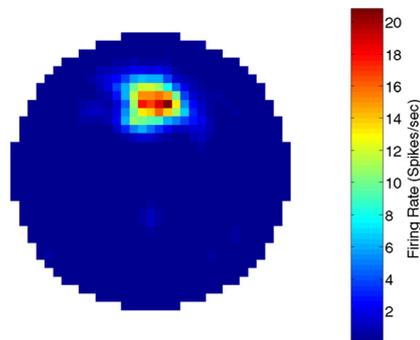


Figure 6.1: **Tuning Curve.** An example of a tuning curve of a place cell recorded on openfield. R031-2003-05-14 TT04-01

6.3 Dispersion measurement

One possible model of the behavioral or stimulus dependence of a neuron's firing is that of an inhomogeneous Poisson point process where the intensity of the spike emission rate at any given time is dependent on the stimulus or behavioral value at that moment. This is the approach used by Fenton and Muller (1998) to examine the variability of hippocampal neurons as the animal passes through their place fields. In order to compare the results of our experiment with previous reports of overdispersion, we implemented the analysis described in their paper (Fenton and Muller, 1998).

First, a neuron's place field was taken as the largest contiguous body of tuning curve pixels with non-zero firing rate. Therefore, every pixel in a place field shared a border with at least one other pixel in the place field. The center of a place field was taken as the 3×3 group of pixels with the highest mean firing rate.

An animal's pass through a place field was only considered if it satisfied all of the following conditions: (1) the pass went through the center of the place-field (2) the pass lasted longer than 1 s, (3) the tracking of the animal was continuous throughout the pass with high theta rhythm and not interrupted by sharp-waves. The spiking of a neuron was compared with the firing expected given its tuning curve during each pass through its place field that met the above conditions.

6.4 Replication of overdispersion result

Fenton and Muller (1998) reported that hippocampal place cells have surprisingly highly variable firing rates. They compared the number of spikes fired by place cells to the number of spikes predicted by an inhomogeneous Poisson point process with rate parameters based the spatial tuning of the cells. If neurons were as variable as this maximally variable model, then a z-score based on the expected mean and variance should follow a Gaussian distribution with zero mean and unit variance. The number of spikes actually emitted by the neurons were much more variable than their model predicted, exhibiting a non-zero mean ($\mu = 0.18$) and a variance almost 6 times that predicted by the model ($\sigma^2 = 5.9$). This excess variance is termed *overdispersion* (see Fig 6.2).

On our open-arena foraging task (OF), we observed a similar phenomenon of high variance ($\sigma^2 = 6.0$, $N_{passes} = 4034$, $n_{rats} = 6$). Figure 6.2

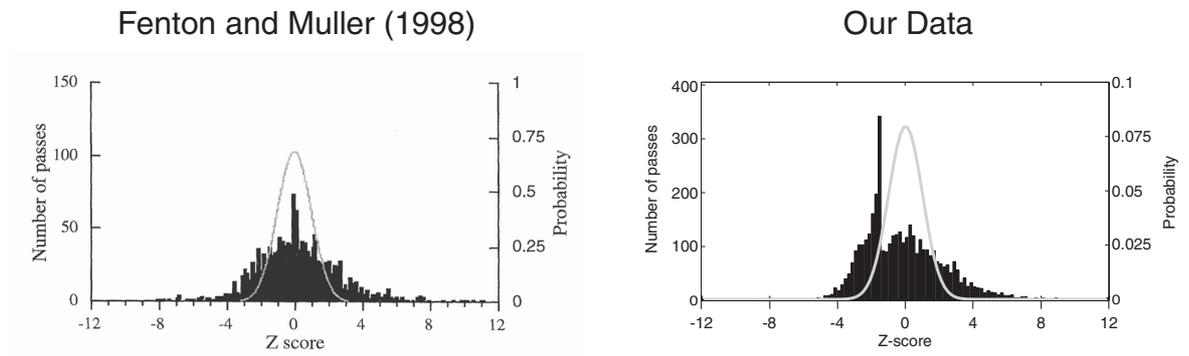


Figure 6.2: **Overdispersion of hippocampal place cells on the Open Field task.** Both figures contain histograms of z-scores for the number of spikes emitted on a pass through a place field given the expected number of spikes predicted by a Poisson point process model. (*left*) Fenton and Muller (1998) reported that neurons displayed excess variability on their open field foraging task (figure from Fenton and Muller (1998): z-scores for 1440 passes). (*textitright*) Place cells on our open field foraging task also exhibited a comparable amount of excess variability (4034 passes). Note the peak in our distribution also appears on the Fenton and Muller (1998) distribution, but is shifted left.

shows the distribution of z-scores comparing the number of spikes emitted on each pass through a place field to the expected distribution of spikes give an inhomogeneous Poisson process mode. The thin gray line shows the expected distribution with zero-mean and unit variance. Note the peak in our distribution also appears on the Fenton and Muller (1998) distribution, but is shifted left. This peak is the result of passes through place fields when no spikes were fired. Removing no-fire passes results in a smooth, unimodal distribution of the same width (data not shown).

6.5 Task dependence: LT vs. OF vs. OFG

Three years after Fenton and Muller (1998) described the overdispersion phenomenon, Láncský et al. (2001) presented a doubly stochastic Poisson process model where, in addition to the average spike output expected from the firing field, they assume a temporal rate switching process such that the average spike output is given by the temporally weighted mean of two rates. Láncský et al. (2001) conjectured that if overdispersion is related to a switching or modulation resulting from interaction between CA1's two main inputs, entorhinal cortex and CA3, then this excess variance may be due to a switching between cognitive states such as using rat-centered and room-centered coordinates for navigation.

Later, Olypher et al. (2002) extended the mathematical model of Láncský et al. (2001) into an artificial neural network model that included two input sources (e.g. CA3 and entorhinal cortex). They varied the modulation amplitude and duration of one of the inputs and compared the results of this parameterization to the data originally presented by Fenton and Muller (1998). A modulation depth of 10% with a period of approximately 1 s best fit the original data. Olypher et al. (2002) then used an experimental paradigm that conditioned animals to use primarily room-based coordinate frames and compared the results to those of control animals running the same task as in Fenton and Muller (1998) (expected to be switching coordinate frames). In their paradigm, an animal was either

allowed to forage for randomly scattered pellets (control), or required to navigate to a hidden goal, stable in room coordinates, to receive reward (navigator). They found that as animals approached a goal, the variability in firing rates of CA1 pyramidal neurons converged towards the variance predicted by unmodulated inhomogeneous Poisson process (as low as $\sigma^2 = 1.74$ during the 5 s prior to reaching the goal). However, for the foraging animals the variance was high ($\sigma^2 = 4.87$).

Our three-task experiment was designed to replicate and extend the experimental conditions in Olypher et al. (2002), adding a further condition (the linear track). Since it is known that hippocampal place fields exhibit directional firing on the linear track (McNaughton et al., 1983; O'Keefe and Recce, 1993), the hippocampal representation of linear track effectively consists of two separate maps, or reference frames. It is thought that this results from the use of two reference frames (corresponding to the two directions of travel; Redish and Touretzky, 1997; Redish, 1999). If this is true, then we can predict when an animal will be in a particular reference frame on the linear track. Ignoring directionality should approximate a doubly stochastic process where spike emission rates switch between the rates computed for each direction. Thus, we predicted that on the linear track, a non-directional analysis would yield overdispersion comparable to that observed previously (Fenton and Muller, 1998; Olypher et al., 2002) and a directional analysis would yield a dispersion z-score distribution approaching unit variance. A comparison across tasks therefore

should yield insight into the viability of the map-switching hypothesis: since the claim is that map switching can generate overdispersed firing, hippocampal neurons on a task that switches maps should demonstrate overdispersed firing distributions. As we shall see, further comparisons across tasks will also be instructive.

As mentioned earlier, our neurons on the OF task were highly variable ($\sigma^2 = 6.0$, $N_{passes} = 4034$, $n_{rats} = 6$), which compares favorably with previous work ($\sigma^2 = 5.9$ in Fenton and Muller, 1998 and $\sigma^2 = 4.87$ in Olypher et al., 2002). On OFG, however, we did not observe the same magnitude of reduction in variance as Olypher et al. (2002) (5 s pre goal: $\sigma^2 = 5.1$, $N_{passes} = 896$, $n_{rats} = 6$). There was however, a reduction in the variance as the animals approached the goal: $\sigma_{4s}^2 = 4.9$, $\sigma_{3s}^2 = 5.1$, $\sigma_{2s}^2 = 4.6$ (See Figure 6.3). In sum, we were unable to replicate the over-all decrease in variance on our OFG task that was observed by Olypher et al. (2002), but we were able to show a decrease in variance as the animal approached the goal. This decrease was not as striking as the $\sigma^2 = 1.74$. These differences may be due to differences in our behavioral paradigm. Olypher et al. (2002) required rats to pause for 1 s in a 20 cm diameter goal region that was stable across sessions. In our experiment, we required no pause after entry* into our randomly placed 7 cm goal region which varied from session to session. Eliminating the pause requirement may have allowed unplanned goal entry to trigger food delivery, mixing planned and un-

*This reduced the likelihood of the animal slipping into LIA after entering the goal.

planned data into our analysis and increasing variability. We attempted to reduce these unplanned goal entries by reducing the goal region's diameter. Indeed, observing the animals' behavior revealed apparent goal-directed movements with occasional goal triggers that appeared to be accidental.

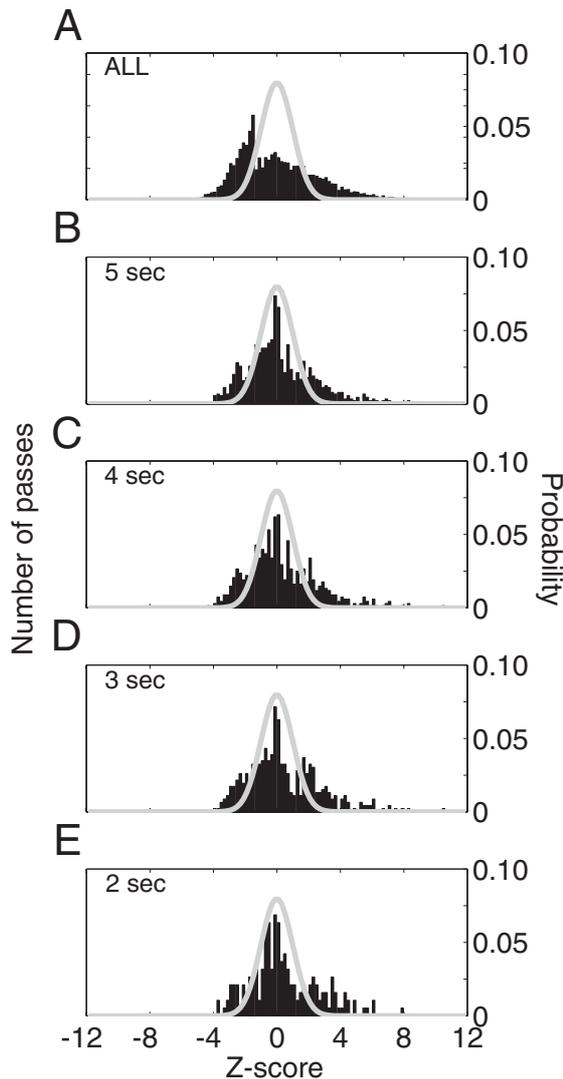


Figure 6.3: **Overdispersion on OFG.** (A) Overall variability on the OFG task ($\sigma^2 = 6.6972$, $N_{passes} = 4675$). (B) 5 seconds prior to goal entry ($\sigma_{5s}^2 = 5.0936$, $N_{passes} = 896$). (C) 4 seconds prior to goal entry ($\sigma_{4s}^2 = 4.9051$, $N_{passes} = 678$). (D) 3 seconds prior to goal entry ($\sigma_{3s}^2 = 5.0829$, $N_{passes} = 462$). (E) 2 seconds prior to goal entry ($\sigma_{2s}^2 = 4.6395$, $N_{passes} = 189$). Note that the primary change to the distribution of z-scores observed as the animal approaches the goal, is the reduction in the low-z-score peak associated with no-fire passes and an increase in the positive tail of the distribution. This indicates that the chance that at least some spikes are fired is increasing along with the overall firing rate.

Table 6.1: Task dependent variability

Task	σ^2	N_{passes}
LT	7.0	4139
LT $A \rightarrow B$	2.1	1005
LT $B \rightarrow A$	2.8	1505
OF	6.0	4034
OFG	6.7	4675
OFG 5s	5.1	896
OFG 4s	4.9	678
OFG 3s	5.1	462
OFG 2s	4.6	189

Our predictions of the effects of switching behavior on the neuronal response variability in linear track were confirmed. Ignoring directionality yielded a highly variable distribution of firing rates ($\sigma^2 = 7.0$, $N_{passes} = 4139$, $n_{rats} = 6$). However, splitting by direction resulted in a trend towards convergence of the actual firing rates with the Poisson-process based model (See figure 6.4; LT $A \rightarrow B$: $\sigma^2 = 2.1$, $N_{passes} = 1005$; LT $B \rightarrow A$: $\sigma^2 = 2.8$, $N_{passes} = 1505$; $n_{rats} = 6$). In our data, we can substantiate that there is a dramatic difference between the reference frames represented in the hippocampus on the two directions of the linear track since there is strong neurophysiological evidence to suggest that this is true: the pattern of place fields in one direction is independent of the pattern in the opposite direction. Thus, at least at the hippocampal level, there are effec-

tively two unique maps observed on the linear track. This dramatic shift in reference frame on LT results in a highly variable discharge of place cells when firing is compared with a tuning curve that was averaged over both reference frames ($\sigma^2 = 7.0$, see Figure 6.4A). Splitting cell firing by the reference frame in use and comparing to the appropriate reference frame reduces (or explains) a large portion of the variability of these firing rates ($\sigma^2 = 2.1$, Figure 6.4B; $\sigma^2 = 2.8$, Figure 6.4C). These data suggest that network-wide modulation such as the reference frame switching on linear track can indeed produce variability on the order of that observed in the open-arena foraging task. However, this does not confirm that reference frame switching is responsible for the excess variability observed on the open-arena foraging and goal tasks.

6.6 Cell Assemblies

Harris et al. (2003) and Harris (2005) presented data suggesting the existence of cell-assembly dynamics in the hippocampus. Cell assemblies as first hypothesized by Hebb in his landmark book (Hebb, 1949), are neurons bound together by mutual inputs and local interactions. If cell-assembly dynamics are taking place in the hippocampus we should see the following effects on dispersion: activation of neurons should be related to

1. the activation of cells with neighboring place fields;

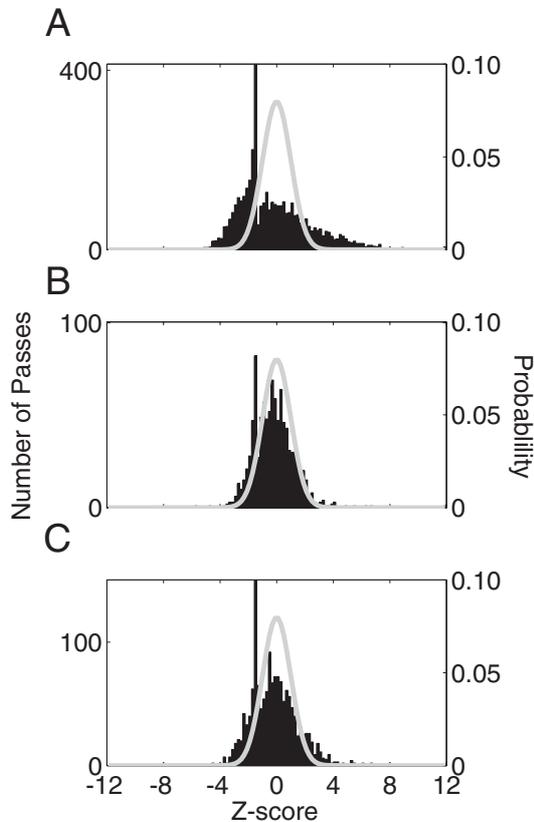


Figure 6.4: **Overdispersion on LT.** (A) Ignoring directionality a highly variable distribution of firing rates ($\sigma^2 = 7.0$, $N_{passes} = 4139$, $n_{rats} = 6$). However, splitting by direction resulted in a trend towards convergence of the actual firing rates with the Poisson-process based model: (B) LT $A \rightarrow B$ — $\sigma^2 = 2.1$, $N_{passes} = 1005$; (C) LT $B \rightarrow A$ — $\sigma^2 = 2.8$, $N_{passes} = 1505$.

2. the sequence of activation of cells with neighboring place fields (e.g. B differs if $A \rightarrow B \rightarrow C$ or $C \rightarrow B \rightarrow A$); and
3. the amount of repetition of activation sequences.

6.6.1 Cell pair correlations

One of the most interesting questions surrounding the phenomenon of overdispersion is whether this is a network-wide phenomenon or not. Could overdispersion be solely due to a noise process intrinsic to the neu-

ron and independent of network level processing? This question is tied to the cell assembly hypothesis mentioned above. The first attempt at answering this question was by Fenton and Muller (1998) in their original paper reporting overdispersion. Fenton and Muller (1998) performed pair-wise correlations of z-scores between two pairs of cells with overlapping place fields. They found no significant correlation for either pair of simultaneously recorded cells on their open arena foraging task.

Since these analyses, Láncsy et al. (2001); Olypher et al. (2002) tried to theoretically explain the source of overdispersion by fitting the data with models that assumed upstream modulation of the hippocampus. Olypher et al. (2002) also went one step further, implementing a behavioral control to bias animals into one hypothetical cognitive state. They required animals to navigate to a hidden goal to receive food reward expecting that this would force animals to only a room-based reference frame during navigation to the goal. The weakness in these data is that they do not and cannot substantiate that the animals are indeed using different mixtures of coordinate systems during the goal and foraging tasks. Both studies also propose mechanisms (viz. reference frame switching) that should result in pair-wise correlations contrary to the findings of Fenton and Muller (1998).

We, therefore, looked for correlations in the dispersion z-score between pairs of simultaneously recorded neurons in our data. The overlap of tuning curves for all neuron pairs in an ensemble was measured as the num-

ber of pixels above zero firing rate in both neurons' tuning curves. The distribution of overlap values was divided such that pairs in the upper 75% of overlap scores for each task were chosen for analysis. The pass times for each neuron in a pair were then compared to find all passes that overlapped temporally by any amount. The dispersion z-score values for these matched passes were then added to the pool of pair-wise data. Correlations were then performed on this pooled data. Higher and lower overlap cutoffs (50% and 95%, when possible) were also assessed and qualitatively similar results were obtained. Randomizations consisted of randomizing one neuron's z-score values across passes prior to matching the pass times for the pair.

Since our experiments generated large ensembles of simultaneously recorded neurons, it was possible to replicate these pair-wise analyses. For each task, the dispersion of all cell pairs with the most overlap was compared (the pairs with overlap in top 75% of overlap among all pairs on a task). To control for non-path-specific interactions between neurons, we randomized the order of the z-scores for the mutual passes (passes that went through both neurons' place-fields) for one neuron in the pair. On the open field task we found that there were indeed weak but highly significant positive correlation between dispersion z-score for neurons with overlapping place fields ($\rho = 0.047$, $P(\rho = 0) = 0.0048$, $N_{pairs} = 966$). This correlation was absent in randomized controls ($\rho = 0.0092$, $P(\rho = 0) = 0.59$) suggesting that sheer increases in the number of points were

not resulting in a false appearance of coupling between cells. On the OFG task, there was a similar level of correlation ($\rho = 0.049$, $P(\rho = 0) = 0.00095$, $N_{pairs} = 1105$) which was absent in the randomized controls ($\rho = 0.016$, $P(\rho = 0) = 0.27$). This correlation almost doubled when considering the five seconds prior to goal entry ($\rho = 0.095309$, $P(\rho = 0) = 0.036457$, $N_{pairs} = 323$).

The linear track with its two-reference-frame nature was, however, not correlated significantly under any condition (LT non-directional: $\rho = -0.010616$, $P(\rho = 0) = 0.48112$, $N_{pairs} = 897$; LT $A \rightarrow B$: $\rho = -0.018438$, $P(\rho = 0) = 0.59647$, $N_{pairs} = 222$; LT $B \rightarrow A$: $\rho = 0.036874$, $P(\rho = 0) = 0.11463$, $N_{pairs} = 359$). The randomized controls were also not significant (LT: $\rho = -0.018$, $P(\rho = 0) = 0.24$; LTAB: $\rho = -0.028$, $P(\rho = 0) = 0.42$; LTBA: $\rho = -0.015$, $P(\rho = 0) = 0.53$).

The significant correlation of z-score across cell pairs on OF and OFG suggests that there is indeed a process coupled across cells, perhaps at the network level, that is influencing their variability. The fact that we do not see this on LT suggests that either the expected anti-correlation between place-cells that prefer opposite directions is swamping the expected correlation between cells that prefer the same directions or that hippocampal dynamics are fundamentally different on LT compared to tasks in the open arena. If the former is true (that the anti-correlated neuron pairs are swamping the contributions of the correlated neuron pairs), it remains possible that the overdispersion on OF and OFG is due to a network-wide

process such as reference frame switching. The evidence in favor of this is that there is a dramatic increase in correlation between cell pairs on LT when we exclude all passes when either neuron in a pair did not fire, however this increased correlation does not reach significance ($\rho = 0.037$, $P(\rho = 0) = 0.078$, $N_{pairs} = 897$). Furthermore, it may be that if there exists a switching process, on OF and OFG the switch is between correlated maps such that the cell pair correlations would not have a strong anti-correlated component between maps. Thus, the depth of modulation between reference frames could underlie the difference between LT and the open-arena tasks (OF and OFG). On LT, the two-directions have nearly orthogonal maps such that the correlation between maps for each direction are small (data not shown) and the modulation between maps is therefore nearly 100%. However, if the hippocampal dynamics truly are fundamentally different on LT compared to tasks in the open arena, then these data may argue against a clean switching process on OF and OFG. This is because our example of a switching process (LT), has very different pair-wise interactions than the highly significant interaction seen on OF and OFG.

Thus, we have partial evidence that prediction #1 (from our cell assembly discussion, Section 6.6) may be true: there is significant correlation between z-scores of cell pairs. These data suggest that the neuronal firing variability may be due to transient interaction or coupling between subsets of neurons with overlapping place fields. This type of interaction

fits Hebb's (1949) conceptualization of the *cell assembly* (Hebb, 1949; McNaughton and Morris, 1987; Harris, 2005). The fact that this correlation is weak ($\rho \simeq 0.05$ and is not found in some cell pairs (Fenton and Muller, 1998) indicates that #2 and #3 may be important.

6.6.2 Testing Ensemble Modulation: Overdispersion vs. Coherency

Since, a weak but highly significant positive correlation between the z-scores of cell pairs suggested that the phenomenon of overdispersion may be a network-wide modulation of the hippocampal system on OF and OFG, we predicted that the dispersion z-score should be correlated with the ensemble measures discussed earlier. To test this, we compared the dispersion z-score to the I_{RMS} and I_{STD} measures of ensemble deviations from expected tuning. The I_{RMS} measure was used since it is specifically sensitive to deviations in the firing levels of the ensemble. The I_{STD} measure was used since it is insensitive to ensemble wide increases in activity, but it measures relative differences in the activity across the tuning space.

Coherency analysis method

As stated in the Simulation Methods, coherency is a comparison of the actual and expected activity of an ensemble, represented by the *actual* and *expected activity packet* ($A(x, t)$ and $\hat{A}(x, t)$, respectively). The activity packet

is well defined for any given state of any system. Thus, comparing actual and expected activity packets can provide a measure of the state of the system at any given moment in time. We therefore applied measures derived from the simulation studies to our neurophysiological data to test questions regarding fluctuations in the state of the hippocampal network.

Actual and expected activity packets were calculated from the tuning curves of each neuron and from the firing rates of each neuron at a time t given the spatial tuning $T(x, y)$ of each neuron in the ensemble. The actual and expected activity packets were compared using the I_{RMS} and I_{STD} measures (See Equations 4.16 and 4.18, respectively).

As mentioned earlier, the I_{RMS} measure is sensitive to absolute differences in ensemble firing across the population while the I_{STD} measure is sensitive to relative differences in ensemble firing across the population. In the *leave-one-out* coherency analyses discussed next, the neuron of interest was excluded from the construction of both activity packets. Since the experimental coherency analyses were intended for comparison with other variables, there was no need to construct a training set. Instead, correlations were performed to measure the strength of any relationship between the I measures and the variable of interest.

leave-one-out analysis

We calculated the I_{RMS} and I_{STD} measures for the entire ensemble excluding the neuron for which we were computing the dispersion z-score (see

Figure 6.5). The mean correlation between the absolute z-score and I_{RMS} across sessions on OF was -0.0015 ± 0.069 (mean \pm SE) and for I_{STD} was -0.026 ± 0.042 (mean \pm SE). These correlations were not significantly different from zero (I_{RMS} : $P(\text{mean} = 0) = 0.76$, I_{STD} : $P(\text{mean} = 0) = 0.62$; two-sided Wilcoxon sign rank test). A similar result was found for OFG. The mean correlation between absolute z-score and I_{RMS} across sessions on OF was 0.0012 ± 0.026 (mean \pm SE) and for I_{STD} the mean correlation was -0.048 ± 0.042 (mean \pm SE). As in OF, these correlations were also indistinguishable from zero (I_{RMS} : $P(\text{mean} = 0) = 0.68$, I_{STD} : $P(\text{mean} = 0) = 0.23$; two-sided Wilcoxon sign rank test).

The mean correlation between z-score and I_{RMS} across sessions on LT was 0.017 ± 0.039 (mean \pm SE) and for I_{STD} was 0.010 ± 0.044 (mean \pm SE). Neither correlation was above zero (I_{RMS} : $P(\text{mean} = 0) = 0.78$; I_{STD} : $P(\text{mean} = 0) = 0.55$; two-sided Wilcoxon sign rank test). Splitting LT by direction, there was no significant positive correlation between either measure with the dispersion z-score (I_{RMS} $A \rightarrow B$: -0.070 ± 0.030 , $P(\text{mean} = 0) = 0.027$, $B \rightarrow A$: -0.0016 ± 0.047 , $P(\text{mean} = 0) = 1.00$; I_{STD} $A \rightarrow B$: 0.047 ± 0.032 , $P(\text{mean} = 0) = 0.19$, $B \rightarrow A$: 0.0010 ± 0.042 , $P(\text{mean} = 0) = 0.52$).

If this generalized lack of significance were due to a lack of overlapping cells with which to estimate ensemble firing (i.e. if the z-scored cell had the only place-field that the animal was passing through) then we would expect to see a positive dependence of the correlations on ensemble size.

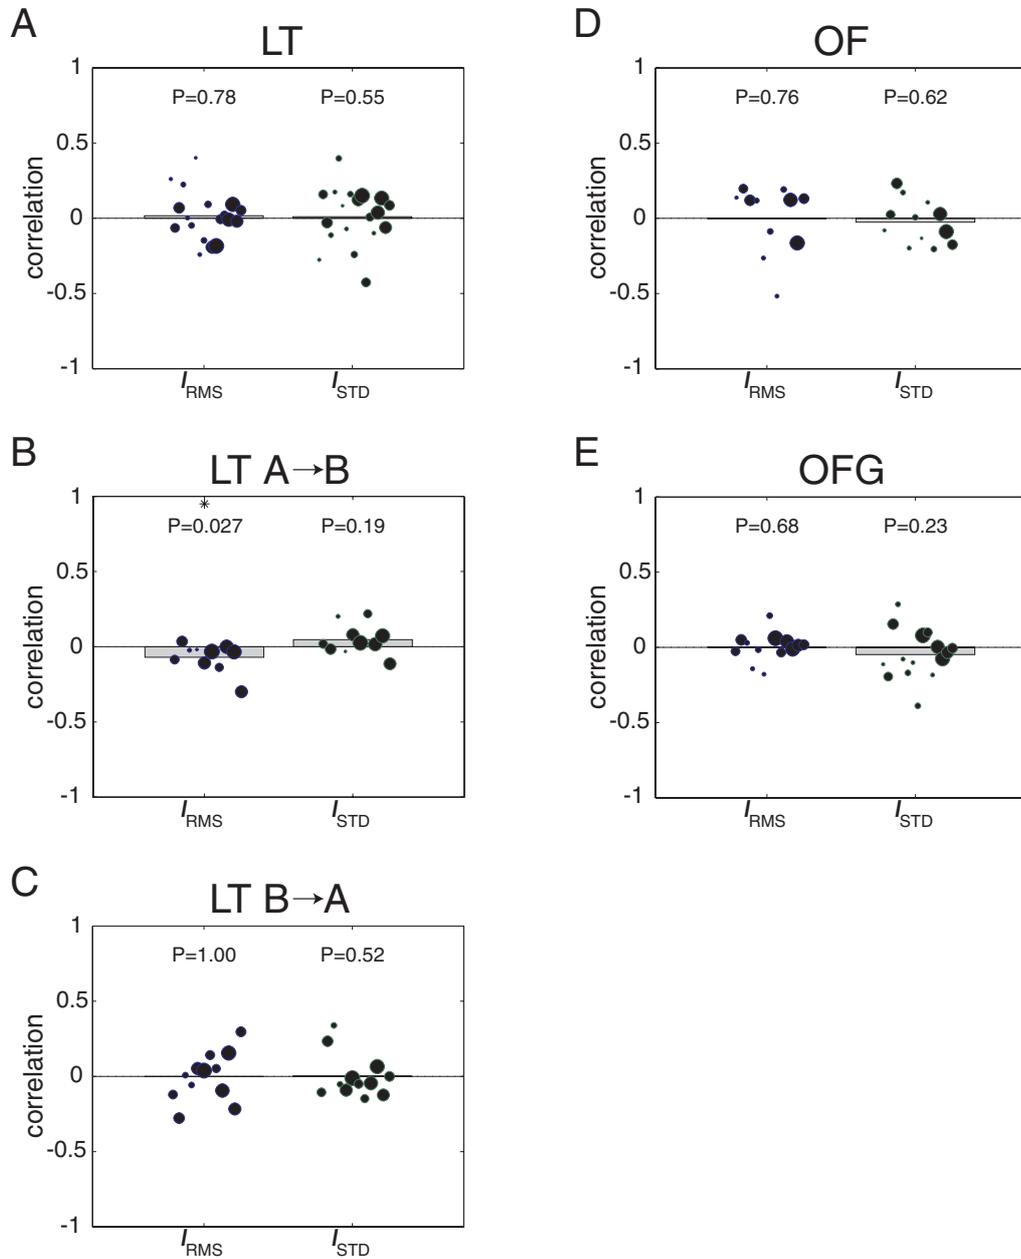


Figure 6.5: **Leave-One-Out: Coherency vs. Dispersion.** Each neuron's z-score is compared with the coherency of the ensemble excluding that cell. The correlation between dispersion z-score and each ensemble measure (I_{RMS} on the left, I_{STD} on the right) is plotted for each task: (A) LT; (B) LT A → B; (C) LT B → A; (D) OF; and (E) OFG. The bars show the mean correlation across sessions, and the circles show the session correlations. The circle size is proportional to the ensemble size. P-values for the null hypothesis that the mean correlation is zero are above each bar. An asterisk above the P-value indicates significance at the level of $P < 0.05$.

However, there was no significant effect of ensemble size on the correlation for the I_{RMS} measure with dispersion z-score (OF: $P(\text{mean} = 0) = 0.60$; OFG: $P(\text{mean} = 0) = 0.34$; LT: $P(\text{mean} = 0) = 0.16$; LT $A \rightarrow B$: $P(\text{mean} = 0) = 0.93$; LT $B \rightarrow A$: $P(\text{mean} = 0) = 0.90$), nor was there a significant correlation between I_{STD} and z-score (OF: $P(\text{mean} = 0) = 0.70$; OFG: $P(\text{mean} = 0) = 0.27$; LT: $P(\text{mean} = 0) = 0.42$; LT $A \rightarrow B$: $P(\text{mean} = 0) = 0.43$; LT $B \rightarrow A$: $P(\text{mean} = 0) = 0.62$). Finally, the analyses in this section required cells to share at least 20 pixels of overlap with at least one other cell. Increasing the stringency of the analysis by requiring cells to have a higher degree of overlap (even up to 100 pixels) with at least one neighbor did not improve the correlation between z-score and I_{RMS} or z-score and I_{STD} (data not shown).

In summary, since LT did not have a significant correlation between coherency and absolute z-score when this task is known to demonstrate switching, this implementation of the coherency measurement is unable to detect ensemble-wide modulation that is related to the fluctuations in a single neuron's firing pattern as observed in the dispersion z-score measurement. If switching between multiple maps is responsible for overdispersion, this failure is likely due to our inability to know *a priori* what maps to expect on OF and OFG, in other words we were using tuning curves constructed by averaging over firing from all maps used throughout the session. Thus, if the hippocampus spends roughly equal time using each map (i.e. using each reference frame), then each map is equally incoherent

compared with the average of these maps and we would see no modulation effect.

6.6.3 Correlation of dispersion z-score with the number of passes through a place-field

If the cell assembly hypothesis is true, it may be possible that some groups of cells are more tightly bound temporally than others, regardless of spatial firing. Since there appeared to be very little network-wide modulation that correlated with dispersion of a neuron's firing, one possible explanation for the significant pair-wise dispersion correlation result could be that small groups of neurons are modulated together, but that this modulation is separate from the rest of the ensemble. One possible mechanism for this is the experience-dependent modification of place-fields observed by Mehta et al. (1997) whose authors report that place-fields expand backwards (firing earlier with each approach) and increase their firing rate as an animal repetitively passes through the place-field. To examine the effect of this phenomenon in our data, we asked whether the dispersion z-score was correlated with the number of times the animal passed through a place field. On the open field, there was no significant correlation between a neuron's z-score on a particular pass and the cumulative number of passes through that neuron's place field up to that pass ($\rho = 0.024$, $P(\rho = 0) = 0.13$). However, there was a significant positive correlation on

OFG ($\rho = 0.043$, $P(\rho = 0) = 0.0030$). Thus, neurons fired more (higher z-scores) as animals passed through their place fields more. On the linear track, the dispersion z-score was also correlated with the number of passes whether ignoring directionality or splitting by direction (LT: $\rho = 0.047$, $P(\rho = 0) = 0.0024$; LT $A \rightarrow B$: $\rho = 0.086$, $P(\rho = 0) = 0.0064$; LT $B \rightarrow A$: $\rho = 0.056$, $P(\rho = 0) = 0.029$).

6.6.4 Correlations between Dispersion and Behavioral Entropy

To further test this correlation of z-score with spatial behavior, the probability of transitioning from one pixel to any other pixel was determined for each pass through the pixel. The average entropy of the set of transitions for each pass through a place field was then calculated. Thus, for each movement, an average cumulative entropy could be determined for that movement. Therefore, for each pass through a place-field resulting in a dispersion z-score, we could compare that z-score of the cumulative entropy of all behavior in that same region. This tests item #2 of our predictions from the cell assembly hypothesis (see above): is the activation of a neuron dependent upon the sequentiality of the animal's behavior?

Cumulative Behavioral Entropy Calculation Per Pass Through Place Field.

The x, y -position data was binned into identical bins as the tuning curves: 11 pixel \times 11 pixel blocks (3 cm \times 3 cm) and the transition probability from each bin into every other bin was updated as the animal transitioned from one bin to another as above. With each transition (e.g. moving from bin j to bin k), the Shannon entropy H_j of the transition probability from bin j to all other bins was calculated for the previously occupied bin:

$$H_j = \sum_i^N -p_{i,j} \log_2 p_{i,j} \quad (6.1)$$

Thus, a cumulative record of the entropy of every location traversed in the animal's path throughout the session was compiled. For comparison with the dispersion z-score of each pass through a place fields, the average cumulative entropy of the path during that pass was used.

Dispersion and Behavioral Entropy

As before, we used the absolute z-score since we expect that more variable spatial behavior should produce more variable firing. Since the variability in behavior is measured in entropy which increases monotonically with increasing disorder, the absolute z-score provides the appropriate comparison, increasing monotonically with larger deviations from the mean expected firing. There were significant positive correlations between an animal's behavioral entropy and the magnitude of the dispersion observed

on the linear track and the open-field goal task (LT: $\rho = 0.078 \pm 0.020$ $P(\rho = 0) = 0.0022$; OFG: $\rho = 0.081 \pm 0.029$ $P(\rho = 0) = 0.0085$). This correlation was even stronger on OFG when considering the five seconds leading up to goal entry ($\rho = 0.18908 \pm 0.048353$ $P(\rho = 0) = 0.0015$). Splitting by direction on LT yielded significant negative correlation in the $A \rightarrow B$ direction ($\rho = 0.12 \pm 0.051$ $p(\rho = 0) = 0.049$), and near zero correlation in the $B \rightarrow A$ direction ($\rho = -0.0069 \pm 0.059$ $P(\rho = 0) = 0.91$). There was no significant correlation on OF ($\rho = 0.013 \pm 0.045$ $P(\rho = 0) = 0.57$). These data are shown in Figure 6.6

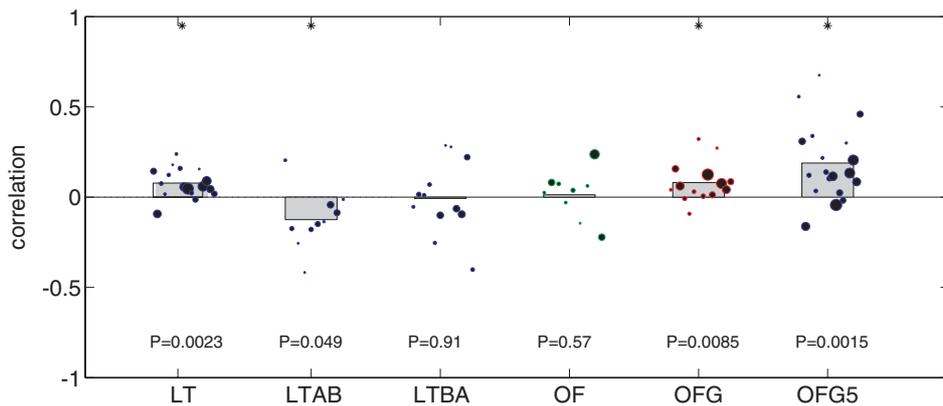


Figure 6.6: **Magnitude of dispersion is correlated with behavioral entropy.** The correlation between the absolute dispersion z-score and the average cumulative transition entropy of the locations used to calculate the z-score. The bars show the mean correlation across sessions, and the circles show the session correlations. The circle size is proportional to the ensemble size. P-values for the null hypothesis that the mean correlation is zero are below each bar. An asterisk above the bar indicates significance at the level of $P < 0.05$. OFG5 is for the 5 s preceding goal entry.

These data indicate that the less ordered the behavior is in a specific region of space, the less reliable the activation of any neuron that fires prefer-

entially within this region. These results indeed fit with the prediction #2. However, the task dependence of this difference (e.g. high correlation on OFG and low correlation on OF) is interesting; this suggests that another variable may be influencing the overdispersion effect.

6.6.5 Correlations with Speed

The small but significant pair-wise correlations and correlations with number of passes and behavioral entropy would not fully explain the large variance observed in the over-dispersion phenomenon. Likewise, the lack of significance of the coherency in the leave-one-out analysis argues against a clear network-wide modulation. However, there have been reports of speed and direction modulation of place cell activity (McNaughton et al., 1983; Markus et al., 1995; Huxter et al., 2003). Therefore, we examined the relationship between speed and the dispersion z-score. On all tasks, the correlations between z-score and velocity were positive and highly significant (see Table 6.2).

We tested this trend by performing a regression on each cell's dispersion as a function of velocity. The average slope was significantly above zero on OF and OFG (OF: slope= 0.019 ± 0.0044 , $P(\text{slope} = 0) = 0.000048$; OFG: slope= 0.018 ± 0.0047 , $P(\text{slope} = 0) = 0.00022$; slopes are mean \pm SE). This relationship was also found on LT and LT $A \rightarrow B$, but was not significant on LT $B \rightarrow A$ (LT: slope= 0.010 ± 0.0049 , $P(\text{slope} = 0) = 0.037$; LT $A \rightarrow B$: slope= 0.0077 ± 0.0028 , $P(\text{slope} = 0) = 0.010$; LT $B \rightarrow A$:

Table 6.2: Velocity vs. Z-score Correlation

Task	ρ	$P(\rho = 0)$
LT	0.13	$6.3 * 10^{-16}$
LTAB	0.11	0.0014
LTBA	0.099	0.00043
OF	0.14	$3.9 * 10^{-18}$
OFG	0.12	$6.8 * 10^{-15}$
OFG 5s	0.17	$3.1 * 10^{-7}$

slope= 0.0036 ± 0.0033 , $P(\text{slope} = 0) = 0.29$; slopes are mean \pm SE). The intercepts were significantly negative on all task conditions except on LT $B \rightarrow A$ (negative, but not significant; data not shown).

6.7 Splitting Representational Maps

The evidence thus far has been pointing to a sub-ensemble modulation of neuronal firing, however, it remains to be seen whether there is indeed multiple spatial maps that underlie this partial-network or cell-assembly-like binding of pair-wise firing. To examine the spatial properties of sub-ensemble interactions, we clustered the firing patterns that occurred within each pixel of an environment to construct multiple whole-environment spatial firing maps. Briefly, for each pixel, the firing patterns consisting of a collection of all firing rate vectors for each video tracker sample where the animal was detected in that pixel were clustered using a k-means algo-

rithm[†]. Figure 6.7 shows two patterns isolated from one 11 cm × 11 cm pixel. Based on these two ensemble firing-pattern clusters for each pixel, we constructed two 2-dimensional firing rate maps by sorting each pixel's clusters into one or the other map by maximizing correlations with neighboring pixels. Finally, the times when the animal's firing patterns were detected in each map were extracted and used to partition the behavior into either of the two representational states.

6.7.1 Map Splitting Method

To examine the spatial properties of sub-ensemble interactions, we clustered the firing patterns that occurred within each pixel of an environment to construct multiple whole-environment spatial firing maps.

Map Splitting by Clustering:

For each 11 pixel × 11 pixel bin (3 cm × 3 cm) bin, all ensemble firing patterns observed when the animal was in that bin were clustered using a k-means algorithm using correlation as the distance metric. Since previous research has suggested that the overdispersion phenomenon is best described using a two-state model (Lánksý et al., 2001; Olypher et al., 2002), we set the k-means algorithm to output 2 clusters, however larger

[†]Since previous research has suggested that the overdispersion phenomenon is best described using a two-state model (Lánksý et al., 2001; Olypher et al., 2002), we set the k-means algorithm to output 2 clusters. Larger numbers of clusters were tried, but there was insufficient data for such analyses.

numbers of clusters were also tried. Extracting more than 2 clusters fragmented the data for most analyses, thus there was insufficient data for dispersion analyses with more than 2 clusters. Figure 6.7 shows two patterns isolated from one 11 cm \times 11 cm pixel. Based on the clusters for each pixel, we constructed two 2-dimensional firing rate maps by sorting each pixel's clusters into one or the other map by maximizing correlations of the cluster's mean firing pattern with the mean firing pattern of clusters in the neighboring pixels. Finally, the times when the animal's firing patterns were detected in each map were extracted and used to partition the behavior into either of the two representational states. These times were referred to as the switching times.

Dispersion (Leave-One-Out):

In order to examine the dispersion of neuronal firing rates, a leave-one-out approach was implemented. For each neuron in an ensemble, the switching analysis was performed on the ensemble excluding that neuron to obtain switching times for the ensemble. These ensemble switching times were applied to the left-out neuron to partition the position and spike data into separate states for the construction of tuning-curves for each state. The dispersion analysis was then run for the state times using the tuning-curves associated with that state, just as was done for the linear-track directional task. To control for the partitioning of states, the tasks were divided in half and the dispersion was run on those halves.

Map Switching Analyses:

Whole-ensemble switching times were used for the rest of the analyses described below. Dwell time in each state was calculated as the product of the position sampling period and the number of position samples detected in each respective state, divided by the number of transitions into each state. Similarly, the switching rate was calculated as the number of transitions into each state divided by the product of the position sampling period and the total number of position samples detected on the task. These values were averaged over the 10 sessions used in this analysis.

For the peri-event time histogram (PETH) switching analysis, the time of switching was measured relative to the onset of key task events such as food delivery (on OF and LT) or to the qualifying tone (on OFG). The number of switching times at each time lag were binned into 0.1 s bins for the 6 s before and after each event, summed for all event times, and normalized by the number of positions samples detected in each bin to yield the transition rate at each time lag from the event. The rates for each session were averaged and the standard error of the mean was calculated for each bin. The transition rate after an event was compared to the transition rate before the event using an unpaired t-test given the mean across sessions. The 95% confidence intervals corrected for multiple comparisons were found using a bootstrap. For each session, 50 randomly selected event times were drawn from a uniform distribution and PETH of switching times were created for each of these pseudo-sessions.

The distribution of 50 runs \times 120 bins followed a normal distribution very closely. The mean and standard deviation of this distribution were used to calculate the 95% confidence intervals corrected for multiple comparisons ($.025/N_{times}, .975/N_{times}$), where $N_{times} = 120$ (the number of bins per PETH).

Coherency Calculation:

The whole-ensemble switching times were also used for the coherency calculation. The actual and expected activity packets were compared using three measures: $I_{RMS}(t)$ (Equation 4.16); $I_{STD}(t)$ (Equation 4.18); and $C_{DP}(t)$ (Equation 4.15).

As stated earlier, the I_{RMS} measure is sensitive to absolute differences in ensemble firing across the population while the I_{STD} measure is sensitive to relative differences in ensemble firing across the population. C_{DP} measures the similarity between the actual and expected activity packets and is sensitive to absolute firing differences. To calculate the coherency ratio CR_{method}^{1-2} to examine the maps derived from in the Map-Switching Analysis (See 6.7.1), the data set was split in half by interleaving minutes. Thus, one half of the data set consisted of all *odd* one-minute blocks of data, and the other consisted of all *even* one-minute blocks of data. One half was used to derive the state-dependent tuning curves for the other given the switching times between maps and visa-versa. This eliminates issues of tautology. The ensemble measure (I_{RMS} , I_{STD} , or C_{DP}) of one half was then calculated

given the tuning curves of the other half and visa versa. The coherency ratio was used to detect state-switching (see State-Switching Analysis):

$$CR_{method}^{1-2} = (\mathbf{Coherency}_1 - \mathbf{Coherency}_2) / (\mathbf{Coherency}_1 + \mathbf{Coherency}_2)$$

where *method* was either *RMS*, *STD*, or *DP*, $\mathbf{Coherency}_1$ and $\mathbf{Coherency}_2$ are the coherency values for maps 1 and 2, respectively. **Coherency** was defined as the proportion of times in the same data set that the actual and expected activity packets matched as well or better than the sample of interest:

$$\mathbf{Coherency} = 1 - \text{cdf}_{I(t)}(I(t)) \quad (6.2)$$

OR

$$\mathbf{Coherency} = \text{cdf}_{C(t)}(C(t)) \quad (6.3)$$

where $\text{cdf}_{measure}(measure)$ means that the I_{RMS} , I_{STD} , or C_{DP} values from each half were concatenated and used to calculate the cumulative distributions cdf. ‡

This index $CR_{measure}^{1-2}$ of how coherent one state was with respect to the other will be above zero if the ensemble firing pattern is more similar to map 1 than to map 2, and less than zero if ensemble firing is more similar to map 2 than to map 1. The coherency of firing patterns occurring in one half of the session was calculated using tuning curves derived from ensemble firing given the switching times in the other half of that session, thereby circumventing any tautological issues. The coherency index was

‡This does not result in a tautology since this used for a relative comparison. At most, this concatenation only lowers our sensitivity working against a significant result.

then aligned to switching times from one state to the other for each task for the half of the data that was not used to construct the tuning curves. It was then binned into 1 s bins. Similar results were found when the tautology was used. The coherency ratio in the one second after an event was compared to the coherency ratio in the one second before the event using an unpaired t-test given the session means.

6.7.2 Results of Map Splitting

Figure 6.8 shows that this method can isolate both directional representational states observed on LT. Times when the firing patterns were clustered into any pattern that was associated with *map#1* are in red and times when firing patterns were clustered into *map#2* are in green; all position samples are shown in black. Since our method could separate putative representational states of the hippocampus, we constructed spatial tuning curves for each state. On the linear track, these state-derived tuning curves closely matched the tuning curves based on the animal's direction of movement (See Figure 6.8). This same process was applied to OF and OFG data.

Applying this analysis to split the OF and OFG maps resulted in similar splits on both tasks. Within the same task, there were clear instances where multiple-place-field neurons had one place field split between maps, where place-fields were entirely assigned to one map and not the other, or where a place field in one map had a higher firing rate than in the other map. These are all classic examples of partial remapping and rate remapping.

Both types of remapping have been observed to occur when animals are placed in different enclosures within the same environment (Leutgeb et al., 2005; Anderson and Jeffery, 2003) and when changing tasks within the same environment (Markus et al., 1995). Buzsáki (2005a) hypothesized that this may be the way that animals keep track of nested environments (e.g. driving your car – you have an environment within and the environment without). For example, a hippocampal representation of the animal's location with respect to the room may be similar to the representation of its location in the arena in that the majority of neurons may fire in the same locations in both maps with differences between these two maps coded by a differences in the firing rate distributed across the ensemble. Thus, if the animal is switching between two (or more) maps while using the task, the tuning-curves used to calculate the dispersion z-score will be a time average of these maps and splitting these maps should reveal tuning curves with partial rate remapping across the states.

Next, we examined the variability of neuronal firing within these two states. Table 6.3 shows that there was much greater dispersion for the whole-task maps than for the task split maps. In fact, the task splits were often less than half the variance of the original data [§]. The larger variance (as compared to the expected unit-variance distribution) was due almost exclusively to a skew towards higher firing rates, rather than a symmetric

[§]Task original values are different from the data in Table 6.1, because they are from the subset of sessions with more than 20 neurons recorded simultaneously.

Table 6.3: Dispersion of firing patterns on each task before and after splitting by representational map: a Leave-One-Out analysis.

Task	Original	Map ₁	Map ₂
LT	6.7	2.9	2.4
OF	5.2	2.5	4.6
OFG	5.8	3.5	2.8

distribution around zero as in the original maps. To test whether this decrease in the variance was simply an artifact of dividing the task into two states, the dispersion of each task divided in half was examined. There was no comparable reduction in variance (data not shown).

An interesting question that follows from this splitting of hippocampal representational states is whether these artificially derived states are of any relevance to the animal. We examined this in two ways: 1) testing whether these states were consistent enough to be useful to higher brain regions; and 2) testing whether the switching times are in anyway related to the task requirements.

We tested the consistency of these states using *coherency* to compare network states derived from a subset of the data to the switching dynamics of the rest of the data from that session. To do this, we used a method similar to that of Redish et al. (2000). As expected, there was a significant increase in the coherency index as the ensemble switched from state 2 to state 1, and a significant decrease in the coherency index when the

Table 6.4: **Significant Changes in Coherency at Transitions.** The probability that the coherency index $CR_{method}^{1-2} = (\text{Coherency}_1 - \text{Coherency}_2) / (\text{Coherency}_1 + \text{Coherency}_2)$ of both states was not increased for transition ($S_2 \rightarrow S_1$) or decreased for transition ($S_1 \rightarrow S_2$) was calculated for each transition between states data shown in Figure 6.10.

Task	Transition	$P^{RMS}(before > after)$	$P^{STD}(before > after)$	$P^{DP}(before > after)$
LT	($S_1 \rightarrow S_2$)	$3.2 * 10^{-5}$	0.010	$6.7 * 10^{-8}$
OF	($S_1 \rightarrow S_2$)	$2.5 * 10^{-8}$	$3.6 * 10^{-10}$	$4.1 * 10^{-14}$
OFG	($S_1 \rightarrow S_2$)	$3.3 * 10^{-7}$	$3.3 * 10^{-7}$	$2.5 * 10^{-13}$
Task	Transition	$P^{RMS}(before < after)$	$P^{STD}(before < after)$	$P^{DP}(before < after)$
LT	($S_2 \rightarrow S_1$)	$1.4 * 10^{-8}$	$1.1 * 10^{-12}$	$8.6 * 10^{-13}$
OF	($S_2 \rightarrow S_1$)	$4.1 * 10^{-17}$	0.034	$8.0 * 10^{-24}$
OFG	($S_2 \rightarrow S_1$)	$9.4 * 10^{-16}$	$1.5 * 10^{-5}$	$1.0 * 10^{-16}$

ensemble switched from state 1 to state 2 (for statistics, see Table 6.4). This indicates that these states are stable and represent real differences in the ensemble firing pattern, suggesting a real network oscillation.

The temporal dynamics of state switching were next explored. Table 6.5 shows the transition statistics for the state switching processes described above. The average switching rate $\mu_{switch}^{\#}$ was approximately 3.2 Hz, with the average cycle rate $\mu_{switch}^{\#\rightarrow\#}$ of approximately 1.6 Hz. The in-state dwell was approximately 380 ms; thus, these states often persisted through multiple theta cycles on average.

To examine the relation of these temporal dynamics to the task behavioral parameters, a peri-event time histogram (PETH) was constructed

Table 6.5: State transition Statistics.

Task	μ_{switch} (Hz)	$dwell$ (sec)
LT	1.5	0.35
OF	1.6	0.31
OFG	1.6	0.33

from all state transition times derived from the ensemble state to yield the transition rate ratio leading up to each task event (viz. food delivery on LT or OF, and the qualifying tone on OFG). These normalized PETHs are shown in Figure 6.11. There were significant increases in switching rates following reward related events on all tasks (LT: $mean(post - pre) = 0.019$, $P(pre = post) = 0.0045$, $t = 2.9$; OF: $mean(post - pre) = 0.014$, $P(pre = post) = 0.0014$, $t = 3.3$; OFG: $mean(post - pre) = 0.021$, $P(pre = post) = 0.0023$, $t = 3.1$, 2 sample unpaired t-test). These results suggest that the state transition analysis above is pulling out ensemble dynamics relevant to the animal's behavior. If one also considers the coherency results that suggest a network-wide modulation occurring at these transition times, these data indicate that key task events such as food reward are associated with shifts in the representational state of the network. Given that these shifts occur at a greater frequency than the task behavioral cycle, it is possible that these shifts in network state correspond to internal cognitive shifts in motivation or behavioral planning such as spatial target selection

(e.g. switching from pellet search behavior to targeted navigation on OF).

6.8 Conclusions

In this chapter we explored evidence for reference frame switching in the hippocampus using three different tasks: the linear track (LT), foraging in the open-arena (OF), and goal-directed navigation in the open-arena (OFG). We were able to replicate findings of excessive variance in place cell discharge on OF and the findings of reduced variability as animals approached a goal. We demonstrated that on a task with known reference-frame switching (LT), splitting by representational state (LT-directional) resulted in greatly reduced variability. We presented evidence that, contrary to previous research, there were significant local interactions between some neurons with overlapping place-fields, but that overall this effect was weak. We showed that in accordance with the cell-assembly hypothesis, these effects are correlated with the local behavioral variability and level of repetition. Finally, we were able to extract different spatial firing maps based on a separation of ensemble firing patterns. These firing maps are indicative of separate reference frames being used in an alternating pattern on the task. The variability of firing patterns within these reference frames is greatly reduced, approaching the variance expected by a single point-process stochastic model.

One question that comes to mind when considering the results of the

coherency analysis is why the leave-one-out analysis yielded no significant relationship between the ensemble measures and the dispersion measure if there truly is a switching process. The answer lies in the null result for the leave-one-out analysis on the non-directional linear track. This tells us that if we have two maps used on one task, that the tuning-curves will be an average of both maps. Assuming each map is active approximately the same amount of time as the other map, then when either map is active, the error in the ensemble activation is equally high. In other words, a switch from one map to the other results in similar I_{RMS} and I_{STD} measures.

One possible interpretation for this reduced variance of the split times, is that it is simply a phenomenon of explaining the variance by adding another parameter. This may be the case, however it would have to be argued that it was a judicious parameter choice because splitting the map in other ways does not result in lowered variance (e.g. splitting the tasks in half — first 10 min versus second 10 minute; data not shown). Regardless of whether or not the smoothing that results from maximizing correlations with neighboring pixels within a map is responsible for the reduction in variance after splitting, there remains the temporal significance of the transition times between maps. The maps correlate well with the maps corresponding to the directions of travel on linear track and tend to switch right after reaching the opposite feeder. Likewise, on the OFG, the transition between maps corresponds closely with the reward, with an

increase in transition rate between maps after the qualifying tone.

What do these data tell us about cell assemblies? Cell assemblies are by definition neurons bound together by common inputs and reciprocal connections such that their spiking is more interrelated within an assembly than across an assembly. If we apply this definition to our hippocampal data, we see that there are multiple continuous cell assemblies. *Multiple* in the sense that distinct spatial tuning maps result from clustering the firing patterns observed on the task. *Continuous* in the sense that each map is extended in time and space such that the pattern of neurons active within one map gradually evolves as the animal moves through space until the hippocampus switches to the next map. Testing our continuity assumption is beyond the scope of this chapter. It is the assumption of spatial continuity that allowed us to define the switching times between maps.

The evidence for these cell assemblies is in the pair-wise correlations. If neurons are part of the same cell assembly they could be bound together spatially and temporally, while neurons from a different cell assembly may be bound together spatially and remain temporally independent. This is a recently reported phenomenon (Harris et al., 2003; Harris, 2005). These data also fit with recent observations of non-spatial network coding in the hippocampus reported by Lin et al. (2005). Using a multiple components analysis they found that specific patterns of activation across the hippocampal network could be used to reliably discriminate between salient task events such as air-puffs, shaking the environment, and sudden down-

ward displacement of the animal's cage. The ensemble activation patterns in our data overlapped to some extent in that there was a generalized similarity between the two maps extracted from OF and OFG with variations in firing rate depending on which state the network was in and occasional remapping of a few neuronal responses to new preferred locations in the environment. This level of overlap is reminiscent of the the overlap in activation patterns (*neural cliques*) observed by Lin et al. (2005).

Our results argue for a faster switching rate with deeper modulation than suggested by Olypher et al. (2002). As mentioned earlier, the prominent, left-shifted peak in our data (see Fig 6.2) is due to no-fire passes. Figure 6.9 depicts numerous state-remappings where a place field exists in one state and is absent in the other state. If an animal runs through this region in one state, spikes will be fired, while in the other state spikes will not be fired. Markus et al. (1994) reported the remapping of spatial responses in an environment to changes in the task reward contingencies within same session. It may be that training our animals on two behaviorally different tasks has produced hippocampal spatial representations that contrast these behaviorally different tasks ultimately resulting in deeper modulation between maps used by our animals. This would increase the incidence of no-fire passes on both tasks and contribute a large left-shifted component to the distribution of z-scores.

One further question remains: we can account for half the variance in the overdispersion of place-cell firing by splitting the ensemble map

into two temporally independent spatial reference frames, but there is still two to three times the variance expected given a Poisson point process. While we have thus provided more direct evidence for the reference frame switching hypothesis, there remains as yet unexplained variability in the hippocampal neuronal firing rates. This variability is likely generated from a number of sources such as speed modulation, variations in plasticity as a function of experience and the regularity of spatial behavior. There also remains the distinct possibility that other internal cognitive processes such as navigational planning also influence the temporal variability observed in hippocampal neurons.

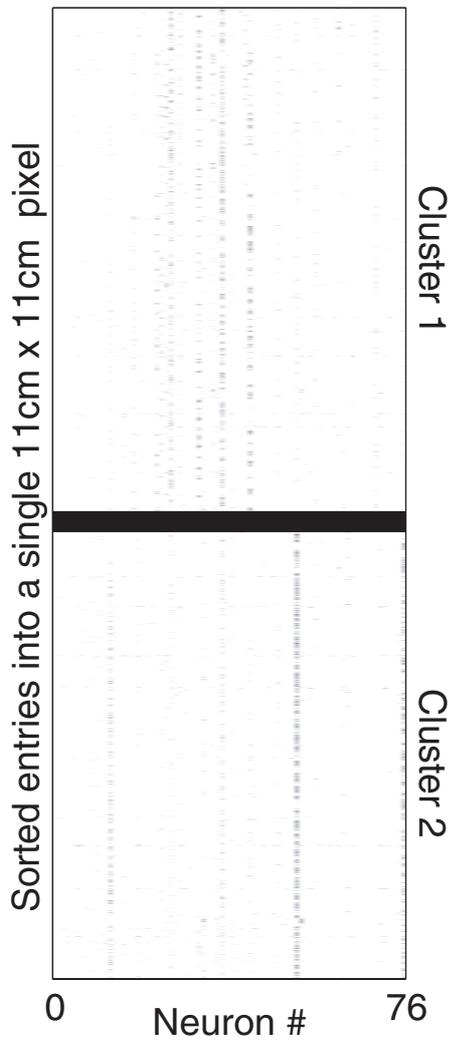


Figure 6.7: **Two distinct firing patterns isolated within a single pixel.** Firing rate vectors from each video tracker sample found within an $11\text{ cm} \times 11\text{ cm}$ pixel were clustered using a k-means algorithm. Shown at left is the firing rate for each neuron (shading by firing rate) for each time detected in the pixel of interest. Times have been sorted into one of two clusters. Times when no firing occurred within the ensemble are not shown.

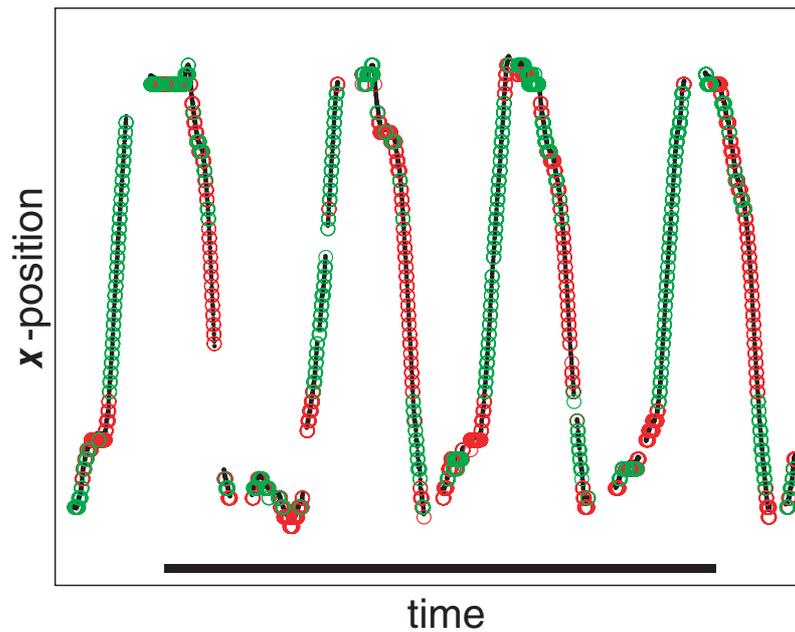


Figure 6.8: **Two distinct maps correspond to two directions of travel on LT.** Firing rate vectors were clustered in two clusters using a k-means algorithm and sorted into one of two maps by maximizing correlations with neighboring pixels based on cluster means. Times corresponding to which firing rates were clustered to which map are plotted in red (map 1) or green (map 2). All position samples are plotted beneath in black. Note that the red and green times segregate to movement in one direction or the other. Scale bar is 1 minute, movement along the x-direction is approximately 1 m. Gaps in x-position data are either due times removed when the animal was in a non-theta LFP state, sharp-wave times, or poor video tracking.

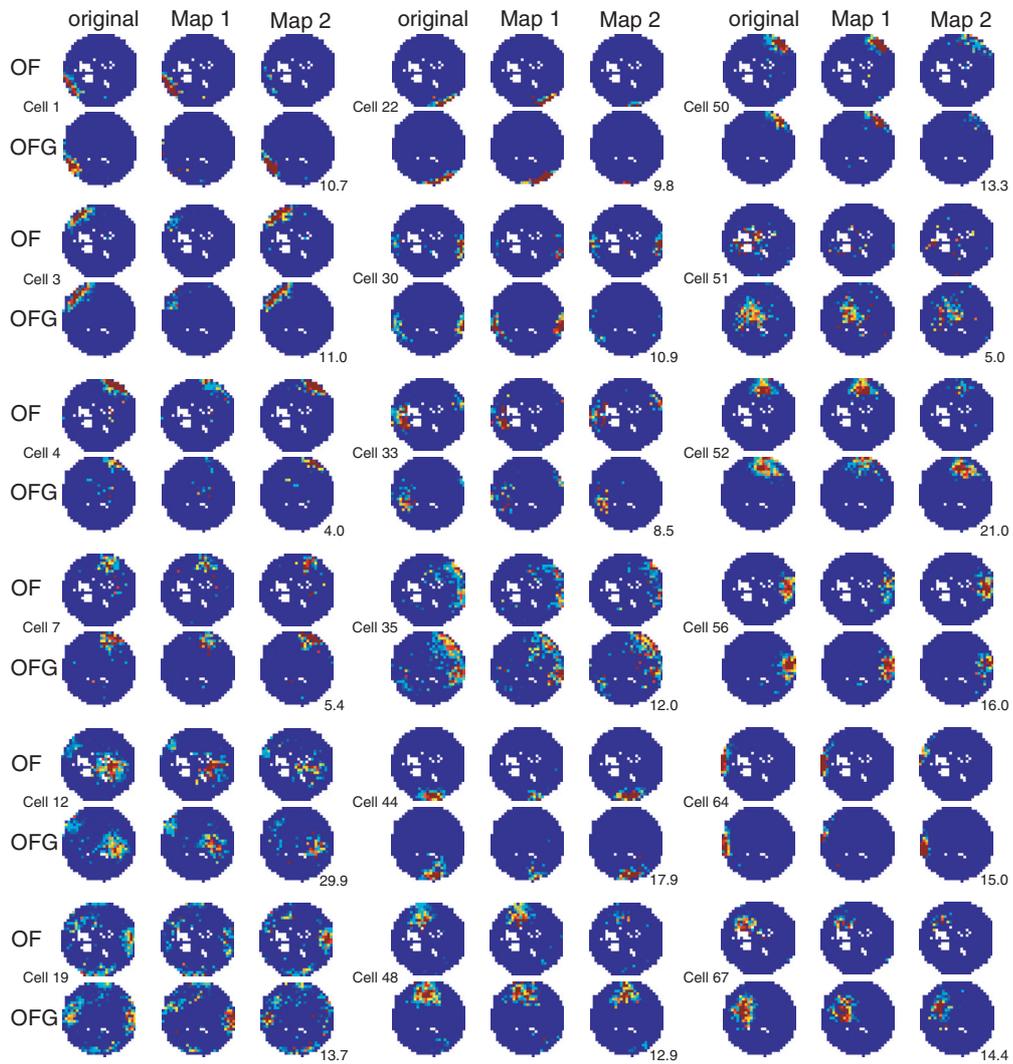


Figure 6.9: **Two distinct maps are found on OF and OFG.** Firing rate vectors were clustered in two clusters using a k-means algorithm and sorted into one of two maps by maximizing correlations with neighboring pixels based on cluster means. Small number in the lower right-hand corner of each group is the firing rate associated with dark red on th map. Not all neurons are represented. The assignment of Map 1 and Map 2 for OF and OFG was based on which map (or stack of tuning curves) on OF correlated best with Map 1 vs Map 2 on OFG. Note there are instances of switched maps from OF to OFG, instances of dropped place-fields from Map 1 to Map 2, and even instances of rate remapping between Maps where one field has a higher firing rate in one map than in the other.

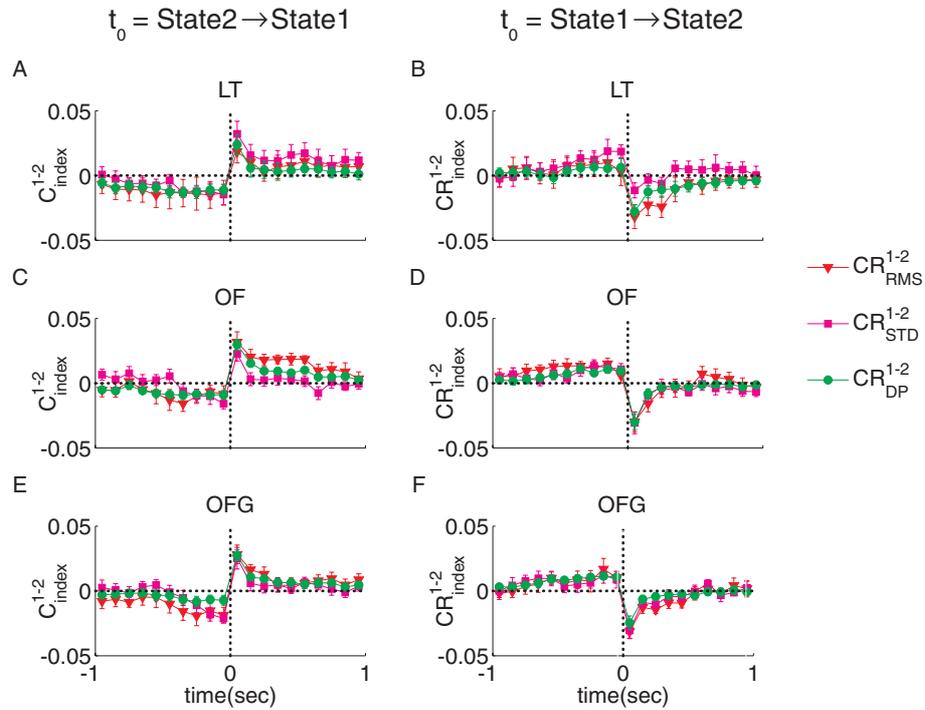


Figure 6.10: **States are stable features of ensemble information processing.** The coherence index $CR_{\text{measure}}^{1-2} = (\text{Coherency}_1 - \text{Coherency}_2) / (\text{Coherency}_1 + \text{Coherency}_2)$ given each ensemble measure was calculated from non-overlapping sets of data to avoid tautological issues. (A) On LT, coherence ratio increases when transitioning from state 2 to state 1. Results are shown for each ensemble measure. (B) LT transitions from state 1 to state 2 show a significant drop in coherence ratio. (C) On OF, coherence ratio increases when transitioning from state 2 to state 1. (D) OF transitions from state 1 to state 2 show a significant drop in coherence ratio. (E) On OFG, coherence ratio increases when transitioning from state 2 to state 1. (F) OFG transitions from state 1 to state 2 show a significant drop in coherence ratio. Error bars show standard error of the mean over sessions. Statistics are provided in Table 6.4.

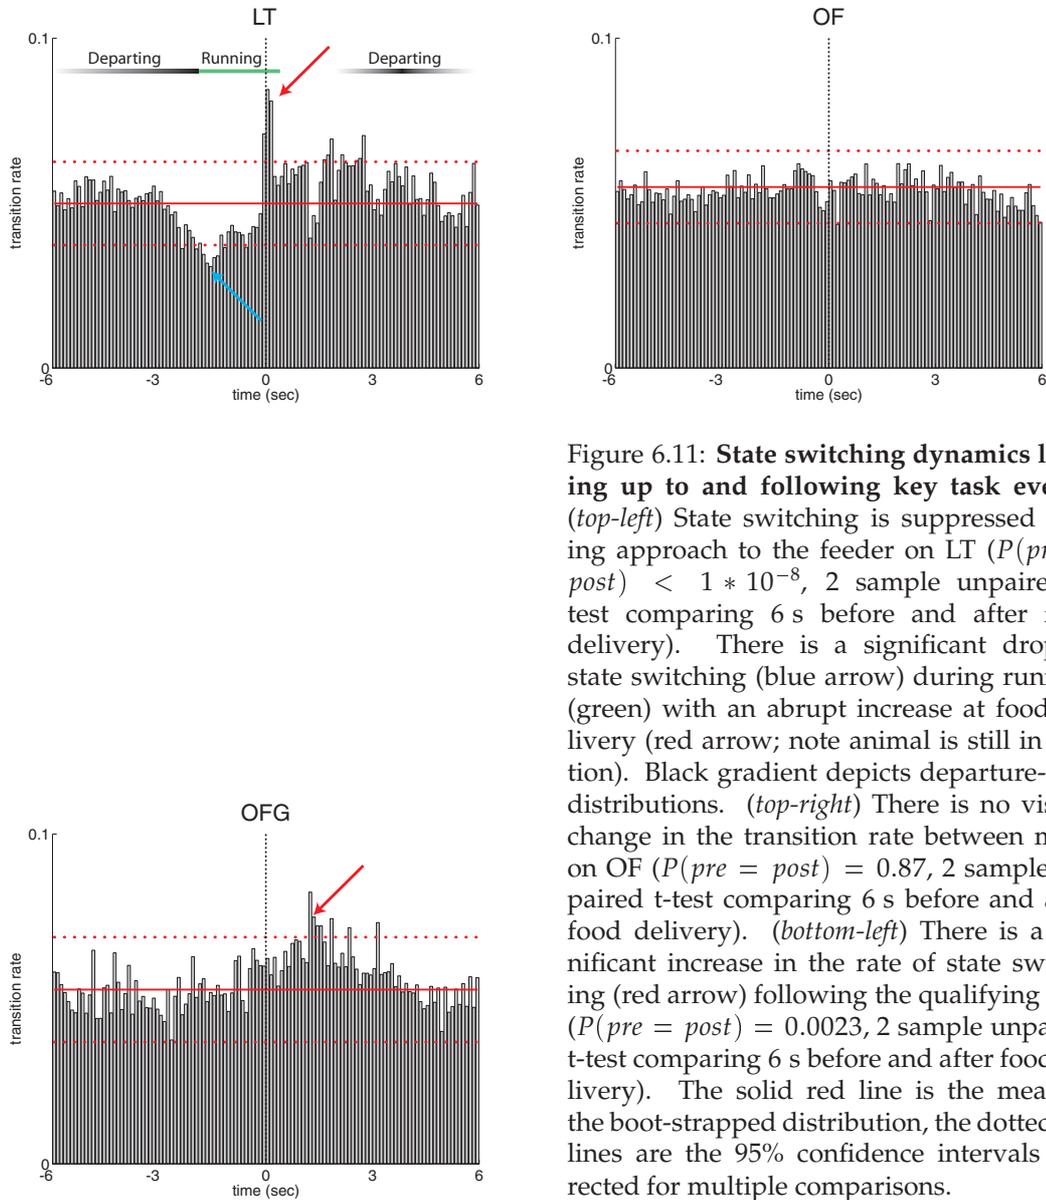


Figure 6.11: **State switching dynamics leading up to and following key task events.** (*top-left*) State switching is suppressed during approach to the feeder on LT ($P(\text{pre} = \text{post}) < 1 * 10^{-8}$, 2 sample unpaired t-test comparing 6 s before and after food delivery). There is a significant drop in state switching (blue arrow) during running (green) with an abrupt increase at food delivery (red arrow; note animal is still in motion). Black gradient depicts departure-time distributions. (*top-right*) There is no visible change in the transition rate between maps on OF ($P(\text{pre} = \text{post}) = 0.87$, 2 sample unpaired t-test comparing 6 s before and after food delivery). (*bottom-left*) There is a significant increase in the rate of state switching (red arrow) following the qualifying tone ($P(\text{pre} = \text{post}) = 0.0023$, 2 sample unpaired t-test comparing 6 s before and after food delivery). The solid red line is the mean of the boot-strapped distribution, the dotted red lines are the 95% confidence intervals corrected for multiple comparisons.

Chapter 7

Sharp-wave Emission: Experience Dependent Hippocampal Dynamics.

Abstract

Network-level modulation that contributes to the variability in the firing of hippocampal units is evident in the difference between theta-related activity and the activity present during sharp-waves. Hippocampal firing patterns present during behavior when 6-10 Hz theta oscillations dominate the local field potential (LFP) record are reactivated during sharp waves, transient LFP events, present in rest and subsequent slow-wave sleep. Theories of hippocampal processing suggest that sharp-waves arise

from strengthened plasticity during theta-related behaviors. We tested these predictions by recording neural ensembles and LFPs from rats running tasks requiring different levels of behavioral repetition. The number of sharp-waves emitted increased during sessions with more regular behaviors. Reactivation became more similar to behavioral firing patterns across the session. This enhanced reactivation also depended on the regularity of the behavior. These results confirm the predictions made by current theories and support the hypothesis that sharp-waves arise from potentiated synapses within the hippocampal network.

7.1 Introduction

As discussed earlier, damage to the hippocampal formation causes memory deficits in the storage of facts, temporal episodes, and new spatial environments. Analogous deficits are observed in both humans and animals following hippocampal lesions (Scoville and Milner, 1957; O'Keefe and Nadel, 1978; Kesner and Novak, 1982; Morris et al., 1982; Squire, 1992; Redish, 1999; Clark et al., 2000; Fortin et al., 2002). In rodents, neural correlates of spatial aspects of this memory are found in the strong spatial tuning of the hippocampal pyramidal neurons (*place cells*, O'Keefe and Dostrovsky, 1971; Redish, 1999). Tasks with repetitive spatial components have been used to study reactivation of spatially-dependent firing patterns during slow-wave sleep (Wilson and McNaughton, 1994; Skaggs and

McNaughton, 1996; Kudrimoti et al., 1999; Nádasdy et al., 1999; Lee and Wilson, 2002). It is thought that this reactivation of stored sequences during sleep underlies the consolidation of memory traces to other brain regions and that the loss of the hippocampus therefore leads to anterograde memory deficits (McNaughton et al., 1983; Marr, 1970, 1971; Buzsáki, 1989; Squire, 1992; McClelland et al., 1995; Sejnowski and Destexhe, 2000).

In the hippocampus, network states are characterized by distinct oscillatory patterns in the local field potentials (LFP). In rats, hippocampal LFPs show two clearly identifiable oscillatory patterns, (1) a 7–10 Hz regular oscillation (*theta*), seen during attentive behaviors, such as running, as well as REM sleep, and (2) a more broad-spectrum pattern (*LIA*), seen during other behaviors, such as grooming, eating, and slow-wave sleep (Vanderwolf, 1971; O'Keefe and Nadel, 1978). *LIA* is punctuated by transient LFP events termed *sharp waves* (SW), identified by high-frequency (100–250) Hz *ripple* oscillations (O'Keefe and Nadel, 1978; Buzsáki et al., 1983; Ylinen et al., 1995). During wakefulness, *LIA* and SW brain states similar to those activated during sleep are observed (Vanderwolf, 1971; O'Keefe and Nadel, 1978; See Figure 7.2). During these awake sharp-waves, ensemble firing patterns are re-activated and this activity appears to grow with time (O'Neill et al., 2006).

Theories of hippocampal function (Shen and McNaughton, 1996; Buzsáki, 1989; Redish and Touretzky, 1998; Redish, 1999) predict that asymmetric plasticity (Levy and Steward, 1983; Bi and Poo, 2001) applied to re-

current connections within CA3 through experience of repeated spatial sequences during theta will lead to storage of sequences within the recurrent connectivity matrix (Levy and Steward, 1983; Muller et al., 1991; Blum and Abbott, 1996; Redish and Touretzky, 1998). During states in which the network was uncoupled from its entorhinal inputs (e.g. slow wave sleep and LIA, Chrobak and Buzsáki, 1994, 1996; Chrobak et al., 2000), uncorrelated noise in the system would then cascade across these strengthened synapses producing a replay of this stored information during sharp-waves (Buzsáki, 1989; Ylinen et al., 1995; Redish and Touretzky, 1998; Redish, 1999; Csicsvari et al., 2000). These theories predict that the emission of awake sharp wave ripple events should increase in number with experience within a session and that the organization of ensemble firing during those awake sharp-waves should improve with experience. These increases in sharp-wave emission and reactivation should depend on the level of repetition of spatial sequences. Since experimental evidence suggests SW activity in CA3 can initiate CA1 SWs *in vivo* and *in vitro* (Behrens et al., 2005; Buzsáki et al., 1983; Csicsvari et al., 1999a; Ylinen et al., 1995), any changes in SW activity in CA3 should be observable in CA1 as well.

We explicitly examined the task-dependence of sharp wave emissions and hippocampal ensemble reactivation in well-trained rats running three tasks of varying complexity. Hippocampal neural ensembles and local field potentials were recorded from the CA1 region of six rats as they ran

three behavioral tasks daily. The tasks included shuttling back and forth along a linear track (LT), foraging for scattered pellets in a cylindrical arena (“open field”, OF), and navigating to a goal for food reward in the same cylindrical arena (OFG). We demonstrate that both sharp wave emission and ensemble reactivation during sharp waves increased with experience. These increases depended on both the behavioral repetition and the regularity (measured as path entropy). These data were also presented in the previous chapter (See Chapter 6).

7.2 New Analyses

Since the past experimental results have demonstrated forward replay of reactivated ensembles (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Nádasdy et al., 1999; Lee and Wilson, 2002) with recent reports of backward replay (Foster and Wilson, 2006) and since both forward and backward replay is predicted by various theories (Levy and Steward, 1983; Buzsáki et al., 1994; Buzsáki, 1996; Redish and Touretzky, 1998), we developed a non-directional ensemble reactivation analysis, that takes into account the ensemble activation patterns present in theta and during sharp-waves. Also, since we infer from these theories that specific aspects of the spatial behavior (e.g. the regularity and repetition) are responsible for engaging the mechanisms of sharp wave generation and associated ensemble reactivation, we therefore developed a new behavioral

analysis: behavioral entropy.

7.2.1 Co-firing Reactivation Analysis

Using only spikes that occurred in strictly high theta states (see Methods: Theta Detection), all neuron pairs that fired more often within \pm half the median theta cycle given a spike from each neuron of interest (i.e. all other neurons in the ensemble not on the same tetrode) were considered to co-fire. This was assessed by finding the maximum bin of the cross-correlation. If the maximum of the cross-correlation was at the zero bin, that pair of neurons “co-fired” during theta.

The same co-firing criterion was applied to the spikes emitted during SWs except that the median SW duration was used for the bin-width of the cross-correlations between non-tetrode cell pairs. Since the mean SW duration did not significantly change over laps (see Results), this measurement of co-firing was unlikely to be artificially inflated by any changes in SW period. Using the median SW duration partially compensates for the compression of replay shown during slow-wave sleep (Lee and Wilson, 2002; Nádasdy et al., 1999) since the SW length was approximately half the length of theta (45 ± 5.1 %; approximately 2 times compressed), which is already up to 10 times compressed relative to behavior (Skaggs et al., 1996). If the maximum of the cross-correlation was at the zero bin, a pair of neurons “co-fired” during SWs. Any noise in determining this measure would only detract from the signal, therefore this definition of co-firing

enforces the conservative nature of the analysis below.

7.2.2 Pattern Identification

Computing the temporal co-firing for all cell pairs from different tetrodes yielded a pattern of 1's and 0's for the theta state s_θ and for the SW state s_{SW} (1 = co-firing, 0 = not co-firing; see example 7.1 and 7.2).

$$s_{SW} = 10111 \dots 100 \quad (7.1)$$

$$s_\theta = 00011 \dots 100 \quad (7.2)$$

These two ensemble co-firing patterns were compared as follows: (1) s_c was computed as the exclusive OR (XOR) of the the two binary patterns s_{SW} and s_θ (see example Eq. 7.3) ;

$$s_c = s_{SW} \oplus s_\theta = 10100 \dots 000 \quad (7.3)$$

(2) given the null hypothesis H_0 that s_{SW} and s_θ are independent with respect to each other and randomly related, the probability $\text{Prob}(k \leq n_c | N, p_c)$ of observing a particular number of 1's, where $n_c = \sum s_c$ or less (1's are mismatches between s_{SW} and s_θ) in the pattern s_c can be computed from the binomial cumulative distribution (Eq. 7.4) with parameters p_c (the proba-

bility of observing a 1; Eq. 7.7), and the length N of pattern s_c .

$$\text{Prob}(k \leq n_c | N, p_c) = \sum_{k=1}^{n_c} \frac{N!}{k!(N-k)!} p_c^k (1-p_c)^{N-k} \quad (7.4)$$

p_c , the probability of observing a 1 in s_c , is determined from the proportion of ones p_{SW} in the SW co-firing pattern and the proportion of ones p_θ in the theta co-firing pattern (7.7).

$$p_{SW} = \frac{\sum s_{SW}}{N} \quad (7.5)$$

$$p_\theta = \frac{\sum s_\theta}{N} \quad (7.6)$$

$$p_c = p_{SW}(1 - p_\theta) + p_\theta(1 - p_{SW}) \quad (7.7)$$

This is a one-tailed measure of how significantly different from random the similarity between two patterns is. If s_{SW} and s_θ are very similar, s_c has very few ones, and n_c is smaller than expected by random chance; therefore, $\text{Prob}(k \leq n_c | N, p_c)$ is low. Because the probability of a mismatch p_c is calculated from the proportion of 1's and 0's in the SW and theta co-firing patterns, $\text{Prob}(k \leq n_c | N, p_c)$ is robust to differences in overall proportions of ones and zeros between the two patterns. Furthermore, since the null hypothesis assumes the patterns are randomly related, dependencies between elements within either pattern are therefore expected to be randomly related. The result is that the individual elements of the output of the XOR operation are expected to be independent and identi-

cally distributed. Violations arising out of dependencies that occur in both sets argue for a rejection of the null hypothesis. Our analytical calculations and empirical simulations found that the binomial approximation is conservative and will underestimate significance levels. This is because, of lower and upper limits to the number of possible matches, which reduces the possibility of seeing low probability events (i.e. events that argue for a rejection of the null hypothesis). These limits come increasingly into play as the proportion of 1's and 0's in the two patterns deviates from 50%, as is often the case in our data. Thus, as the proportion of 1's and 0's gets farther from even in either or both patterns, the measure becomes more conservative with n_c approaching either $N * p_c$ or $N * (1 - p_c)$.

Correlations between the s_{SW} and s_{θ} patterns yielded temporal trends and differences between the experimental group and randomized controls qualitatively similar to the above reactivation analysis for each task. Thus, results presented in the paper hold for other methods of comparing the ensemble co-firing patterns.

Finally, it is important to note that this analysis uses the entire ensemble co-firing pattern, exploiting information about neurons that are not active on the task or pairs that did not have overlapping place-fields. Thus, we are not measuring whether neuron pairs correlated in theta are increasing their correlation in SW; we are testing whether the same *cell assembly* is present in SW that was present in theta. In other words, if neuron pairs were co-activated during SW that were not co-active during theta

(or visa versa), our measure would show a lower value than for matched co-activations in both states. Our application of this method here is not temporally biased, and is therefore insensitive to the ordering of the reactivated patterns. Furthermore, as stated earlier, any noise in our estimation of co-activations would reduce the similarity of the s_{SW} and s_{θ} patterns and argue in favor of the null hypothesis. However, to reduce the likelihood of clustering artifacts influencing the analysis (Quirk and Wilson, 1999), only neuron pairs recorded across tetrodes were included in the analysis.

7.2.3 Randomized Ensemble Controls

Two randomized controls were used for comparison with experimental reactivation. The first randomized control (SWAP) randomizes spike identity across the ensemble while maintaining the over-all ensemble state dependent firing. The second randomized control (SHUFF) shuffles the inter-spike-intervals (ISI) of each neuron preserving neuronal firing statistics but disrupting state-dependent firing and controlling for the contribution of silent cells to the reactivation analysis. (See Fig 7.1). Both randomizations were run 8 times for each session and the data for each session were averaged for within-session comparisons.

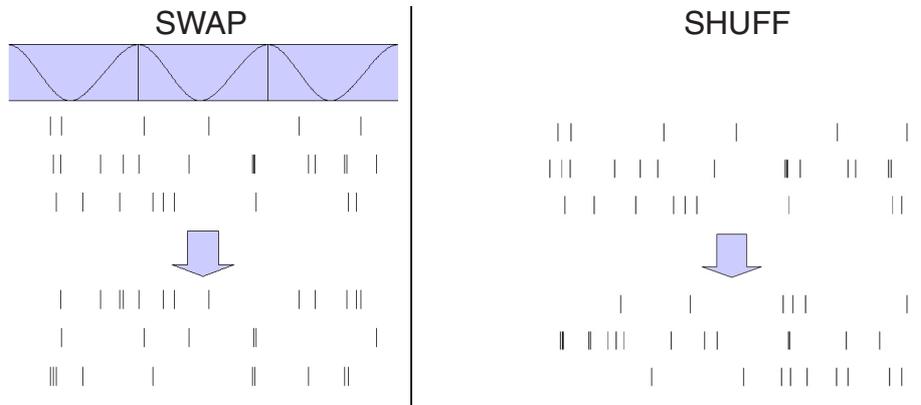


Figure 7.1: **Methods: Randomized controls for cell-assembly similarity.** Diagrammatic. The top three rasters represent spikes from three simulated neurons. Time runs along the x -axis. **SWAP**: The spikes are shuffled across neurons preserving their timing but changing the neuron they are assigned to. This preserves the overall ensemble firing patterns with respect to the oscillatory state shown at top. **SHUFF**: The intervals between spikes are shuffled within each neuron's spike train in the bottom rasters, preserving each neuron's firing statistics but disrupting ensemble state-dependent temporal firing patterns.

7.2.4 Behavioral entropy

For each task on each day, the x, y -position data was binned into $40 \text{ pixel} \times 40 \text{ pixel}$ blocks ($11 \text{ cm} \times 11 \text{ cm}$) and the transition probability from each bin into every other bin was calculated for each block of time. Thus, our 640×480 video capture yielded a 16×12 bin array and a 192×192 transition matrix, containing the transitions between the bins. The transition matrix's Shannon entropy was calculated using all non-zero transitions as

$$H = \sum_j^N \sum_i^N -p_{i,j} \log_2 p_{i,j} \quad (7.8)$$

where $p_{i,j}$ is the time-independent probability of transition from bin j to bin i and $N = 16 * 12 = 192$ is the total number of pixels.

7.3 Results

696 spike-trains were recorded over 24 sessions in ensembles of up to 96 neurons/session (30 ± 31 neurons/session; mean \pm SD); LFP data from 17 additional sessions were included for which spike-trains were not available. Sharp waves (SW) were identified by a threshold applied to the average amplitude across tetrodes of the Hilbert-transformed local field potentials, band-pass filtered from 100 to 250 Hz. As observed in previous experiments (O'Keefe and Nadel, 1978; Buzsáki et al., 1983; Christian and Deadwyler, 1986), SW occurred during the tasks (LT, OF, OFG) concentrated when the rat paused in running (See Figure 7.2). Overall, the number of SW events per time spent in non-theta states increased with laps on each task (slope > 0 , $P(\text{slope} = 0) < 0.0003$). See Figure 7.3.

To determine whether these sharp wave events were also associated with experience dependent changes in neuronal firing, we examined the cross-neuron co-firing probabilities. When averaged over the entire 20 min session, over all tasks, the pattern of neuron pairs that were co-active during awake sharp waves were significantly more similar to the pattern that were co-active during theta than would be expected given random neuronal activity (ANOVA, $P < 10^{-7}$, $F = 19.41$). These results were com-

pared to two randomized controls: SWAP, in which the spike identity was swapped within the ensemble, preserving timing and ensemble firing properties, and SHUFF, in which the spike times were shuffled within spike train, preserving each neuron's overall firing rate. Both randomiza-

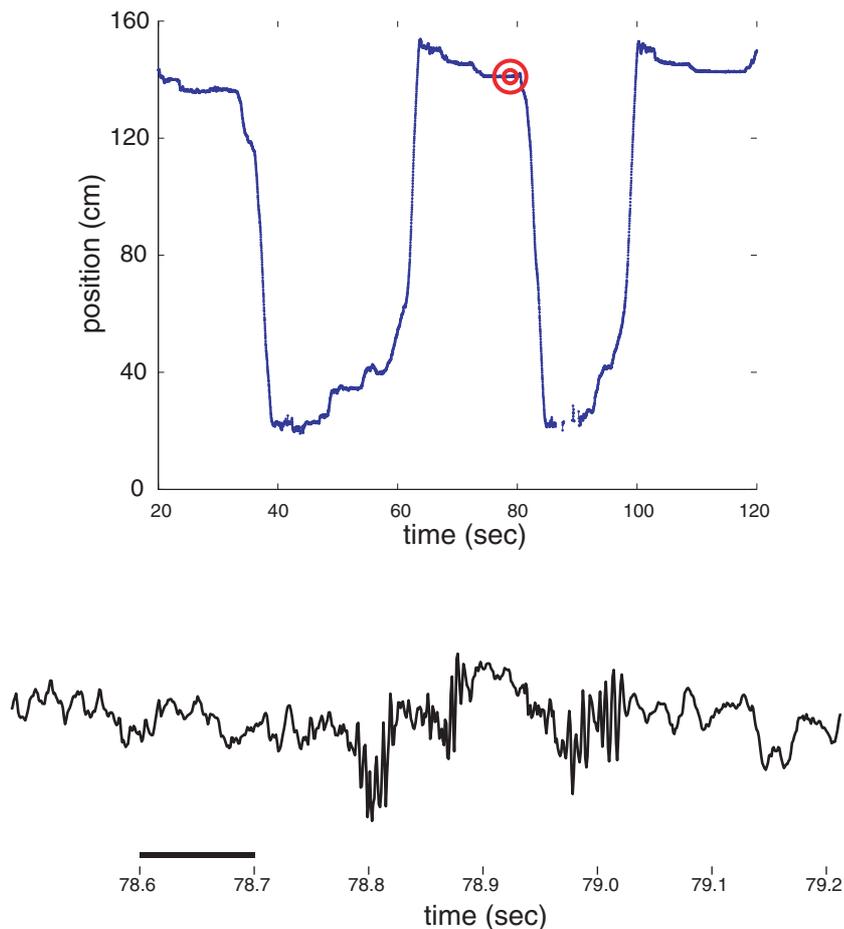


Figure 7.2: **Awake Sharp waves were detected.** (*top*) The rat's location when the sharp-wave (below) was detected. Time is on the x -axis, position along the length of the track is plotted on the y -axis. Bullseye indicates location of rat when SW occurred, just prior to beginning the journey to the next feeder. (*bottom*) Sharp wave ripple event. Time is on the x -axis, voltage is plotted on the y -axis. Scale bar indicates 100 ms

tions removed the effect. See Figure 7.4.

The similarity between SW firing and theta firing also increased across the session. For each session, this was measured by dividing the number of SWs in the session in half (first half, second half), and comparing the similarity of co-active firing in each half with the co-active firing in theta. Overall, the similarity tended to increase (one-sided Wilcoxon paired signed rank test comparing reactivation during the first $\frac{n}{2}$ SW to reactivation during the second $\frac{n}{2}$ SW, $p < 0.02$). See Figure 7.5. While there was a strong overall effect, the strength of the effect on each task differed. We next consider each task individually.

7.3.1 Linear Track.

Sharp waves occurred during the linear track task, concentrated at the ends of the track where the rat received food and paused between each lap (0.19 ± 0.11 SW/sec in non-theta; mean \pm SE). The number of SW events per lap increased throughout the session (slope > 0 , $P(\text{slope} = 0) < 0.00002$, see Fig. 7.3). This increase in SW could not be attributed to a change in performance, behavior at the track ends, changes in time spent in non-theta states or changes in rate of transition between non-theta and theta (see controls, below). The duration of the SW showed no detectable change across the session ($P(\text{slope} = 0) = 0.30$, ns), however there was a significant increase in the amplitude of the SW ($P(\text{slope} = 0) < 0.00005$). Awake SWs are associated with reactivation of the ensembles active dur-

ing theta (O'Neill et al., 2006; Foster and Wilson, 2006) as are SWs that occur in slow-wave sleep after a session (Wilson and McNaughton, 1994; Kudrimoti et al., 1999; Hoffmann and McNaughton, 2002; Nádasdy et al., 1999). The SW emitted during waking states were also associated with reactivation of firing patterns observed during the theta-associated components of behavior (ANOVA, $P < 10^{-10}$, Fig. 7.4). The proportion of neurons included in each sharp wave did not significantly change within each session ($P(\text{slope} = 0) = 0.84$, ns). Nor did the average firing rates occurring within a SW change within a session ($P(\text{slope} = 0) = 0.94$, ns). The reactivation itself, however, did increase in similarity to co-firing patterns observed during theta across the task (one-sided Wilcoxon paired signed rank test comparing reactivation during the first $\frac{n}{2}$ SW to reactivation during the second $\frac{n}{2}$ SW, $P < 0.05$, see Fig. 7.5).

Linear track controls. The increase in SW emission rate on the linear track could have been caused by an increase in the time spent resting between laps at the track ends, however there was no detectable increase in lap duration ($P(\text{slope} = 0) > 0.88$, ns) nor in the resting time between laps ($P(\text{slope} = 0) > 0.79$, ns). Neither was there a corresponding increase in the rate of transition out of theta ($P(\text{slope} = 0) > 0.28$, ns), nor in the time spent in non-theta brain states ($P(\text{slope} = 0) > 0.14$, ns). Since we did not have video image data (only LED coordinates), an analysis of specific behaviors was not possible (e.g. grooming, resting, chewing, etc.). However,

to test for changes in the activity level at the track ends, we compared how the mean and standard deviation of the animal's speed while at the track ends changed across laps. There was no change in either of these measures of activity level (mean movement speed $P(\text{slope} = 0) > 0.55$, ns; standard deviation of movement speed $P(\text{slope} = 0) > 0.36$, ns). As a final control, to check for the possibility that increases in overall LFP power with experience could affect SW detection, we ran the equivalent analyses with the band pass filter set to the theta band power. No detectable increases in threshold crossings, neuronal recruitment, or firing rate were observed at the theta band, suggesting that the increase in SW event detection is due to a specific enhancement in the SW frequency band. Thus, the robust increase in SW emission on LT cannot be explained by overt changes in the animal's state of arousal or intensity of behavior over the first 30 laps.

7.3.2 Two-Dimensional Tasks.

Since the increase in SW emission during awake states was obtained on a linear track where the animal repetitively traversed the same path, which presumably would allow for repeated strengthening of the network's connectivity, we examined whether these increases occurred in more complex environments and tasks with less spatial behavioral regularity. Each day, in addition to the linear track session, each rat ran two additional 20 min sessions: one session in which the rat foraged for food randomly distributed in a 92 cm diameter cylindrical arena ("open-field", OF) as well

as one session in a goal-oriented foraging task (“open-field with goal”, OFG). In the goal-oriented task, whenever the rat crossed the hidden goal, 2–5 pellets were delivered. Once the goal was triggered, the rat had to remain outside of the goal for 4 seconds in order to re-arm the goal. The goal-location changed pseudo-randomly from day to day, but remained constant within each day. The landing location of pellets was randomly distributed throughout the environment on both OF and OFG.

Sharp waves also occurred on OF and OFG concentrated where the rat paused to receive food or rest (OF: 0.11 ± 0.10 SW/sec in non-theta, mean \pm SE; OFG: 0.042 ± 0.061 SW/sec in non-theta, mean \pm SE). The rate of SW emission did not increase significantly on the open-field task ($P(\text{slope} = 0) > 0.12$), nor was there a detectable change in SW amplitude ($P(\text{slope} = 0) > 0.16$, ns). As in LT, there was no detectable change in the duration of the sharp wave ($P(\text{slope} = 0) > 0.10$, ns). When averaged over the entire 20 min session, the pattern of neuron pairs that were co-active during awake sharp waves were significantly more similar to the pattern that were co-active during theta than would be expected given random neuronal activity for both tasks ($P < 10^{-9}$, $F = 31.23$, Fig. 7.4). The similarity between cell assemblies active during SW and theta showed no significant increase in the OF task, (one-sided Wilcoxon paired signed rank test, $P = 0.28$, ns, Fig. 7.5).

On the open-field goal (OFG) task, the occurrence of SW showed a strong trend (approaching, but not reaching significance when corrected

for multiple comparisons, $P(\text{slope} = 0) < 0.002$, Fig. 7.3; ns due to multiple comparisons). As in OF, the duration of SWs did not change significantly ($P(\text{slope} = 0) > 0.57$, ns), nor did the amplitude of SWs ($P(\text{slope} = 0) > 0.44$, ns). When averaged over the entire 20 min session, the set of neuronal pairs that were co-active during awake sharp waves were significantly more similar to the set that were co-active during theta than would be expected given random neuronal activity for both tasks (ANOVA, $P < 10^{-9}$, $F = 23.39$, Fig. 7.4). The similarity between cell assemblies active during SW and theta on OFG did not show a detectable increase (one-sided Wilcoxon paired signed rank test, $P = 0.10$, ns, Fig. 7.5).

Two-dimensional task controls. For the two-dimensional tasks, there was no significant increase in lap duration (OFG, $P(\text{slope} = 0) > 0.70$, ns), neither was there a corresponding increase in the rate of transition out of theta (OFG, $P(\text{slope} = 0) > 0.74$, ns) nor in the time spent in non-theta brain states (OFG, $P(\text{slope} = 0) > 0.10$ ns). Because the lap times used in the OF task were taken from the matched OFG task occurring on the same day, these times were not useful in terms of behavioral controls for the open-field; however, the brain state controls exhibited trends towards a slight increase in both the rate of transition out of theta (OF, $P(\text{slope} = 0) = 0.02$) and the time spent in non-theta (OF, $P(\text{slope} = 0) = 0.01$), but neither was significant when multiple comparisons were taken into account. This may be due to differences in food

delivery contingencies, particularly that food delivery was not linked to behavioral output in the OF task.

7.3.3 Comparisons across tasks.

A two-way ANOVA comparing SW emission across laps and tasks revealed a strong effect of task ($F = 19.1, P < 1 * 10^{-8}$). SW emission rates were significantly higher on LT than on either OF or OFG. ($P < 0.05$, Tukey HSD criterion for multiple comparisons); SW emission rates on OFG were higher on OF than on OFG ($P < 0.05$, Tukey HSD criterion for multiple comparisons).

Since SW emission rates and reactivation time courses differed across tasks, we tested whether those differences in SW emission rates correlated with differences in spatial sequence behavior. To test this, each spatial task was binned into 40 pixel \times 40 pixel blocks (11 cm \times 11 cm) and the transition probability from each bin into every other bin was calculated for each block of time. (See Methods.) As expected, an ANOVA comparing the entropy of these transition probabilities across tasks and time revealed a strong effect of task-type ($F = 101.3, P < 10^{-12}$) with LT having significantly lower transition entropy compared with either OF or OFG. Consistent with the linear track being a 1-dimensional task and both open-field tasks being 2-dimensional, LT had approximately half the spatial transition entropy of the behavior on OF and OFG (LT: 27.4 ± 15.2 bits ; OF: 53.7 ± 28.9 bits; OFG: 56.2 ± 27.0 bits, mean \pm SD).

Because sequential behavior is thought to engage hippocampal plasticity mechanisms (Mehta et al., 1997; Ekstrom et al., 2001; Shen et al., 1997), and both SW emission and reactivation increased with time (Figs. 7.3, 7.5), we tested the extent to which SW emission and reactivation were dependent on the interaction of the two factors of behavioral regularity. Because a sufficient number of SWs were available to measure emission as a function of lap, the emission rate of SWs in non-theta states was measured as a function of lap and the entropy of all behavior leading up to that lap (“cumulative entropy”). SWs were more likely to be emitted earlier with more regular behavior (stepwise multiple linear regression: significant effect of lap number $P < 0.00001$, and of entropy $P < 0.00001$, with a strong interaction $P < 0.00001, F = 52.9$; See Fig. 7.6).

This same trend could be seen in each task. On LT sessions, SWs were more likely to be emitted earlier with more regular behavior. (stepwise multiple linear regression: significant effect of lap number $P < 0.00001$, and of entropy $P < 0.0005$, with a strong interaction between the two $P < 0.00001, F = 15.0$; Fig. 7.6). On OF sessions, a similar interaction effect was observed (stepwise multiple linear regression: weak effect of lap number $P < 0.002$ ns, by multiple comparisons; an effect of entropy $P < 0.001$; and an interaction between the two $P < 0.00005, F = 10.9$; Fig. 7.6). OFG sessions demonstrated a similar interaction (stepwise multiple linear regression: weak effect of lap number $P < 0.002$ ns, by multiple comparisons; no significant effect of entropy $P > 0.31$ (ns); however,

the interaction between the two was significant $P < 0.0005$, $F = 7.88$; Fig. 7.6).

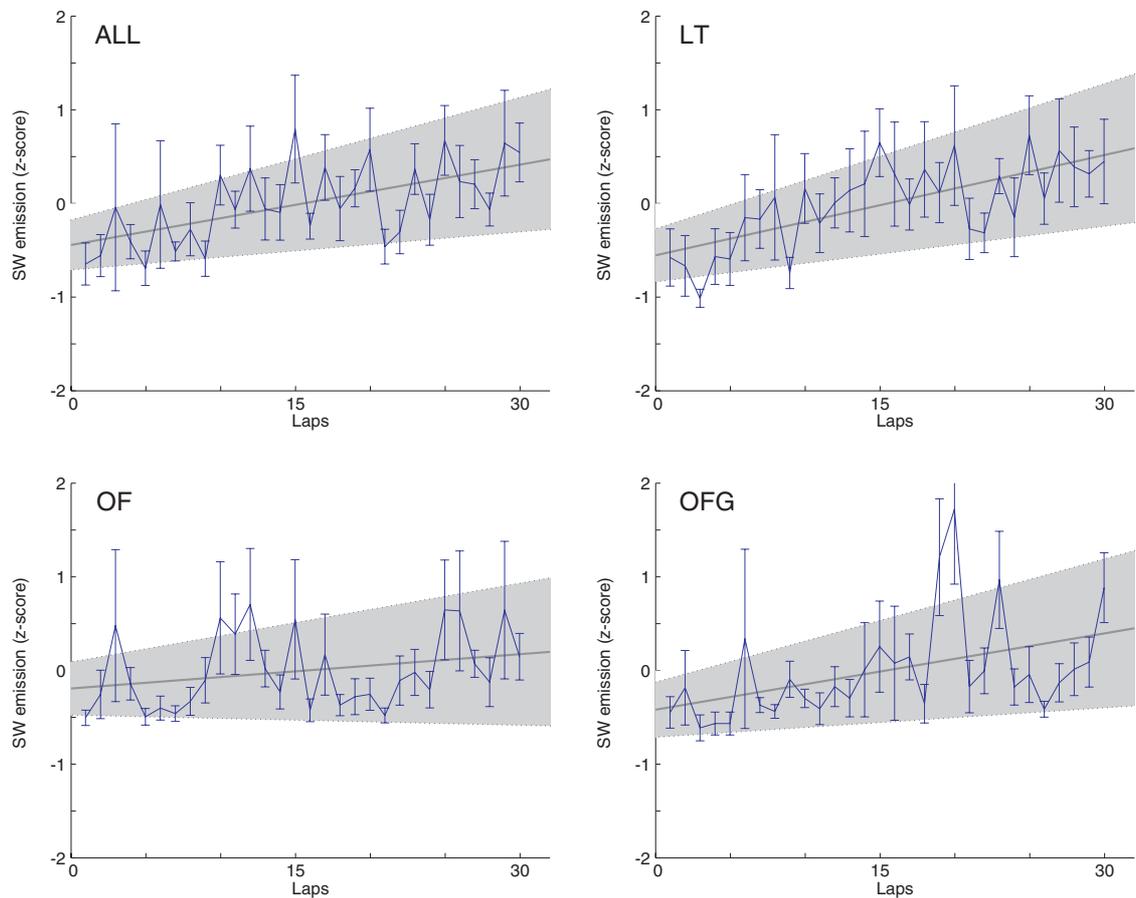


Figure 7.3: **Experience dependent changes in SW ripple events.** Mean and SE (blue solid line) of SW ripple events emission rate normalized by the time spent in non-theta was calculated from individual averages across animals. Linear regression line (dark gray) and 95% regression confidence intervals (light gray) show an increase in the SW emission rate for all four conditions. (ALL) Overall, including all three tasks: $R^2 = 0.063$, $F = 14$, $P(\text{slope} = 0) < 0.0002$. (LT) Linear track, $R^2 = 0.098$, $F = 19$, $P(\text{slope} = 0) < 0.00002$. (OF) Open field, $R^2 = 0.011$, $F = 2.3$, $P(\text{slope} = 0) > 0.12$. (OFG) Open field with goal, $R^2 = 0.056$, $F = 10.3$, $P(\text{slope} = 0) < 0.002$.

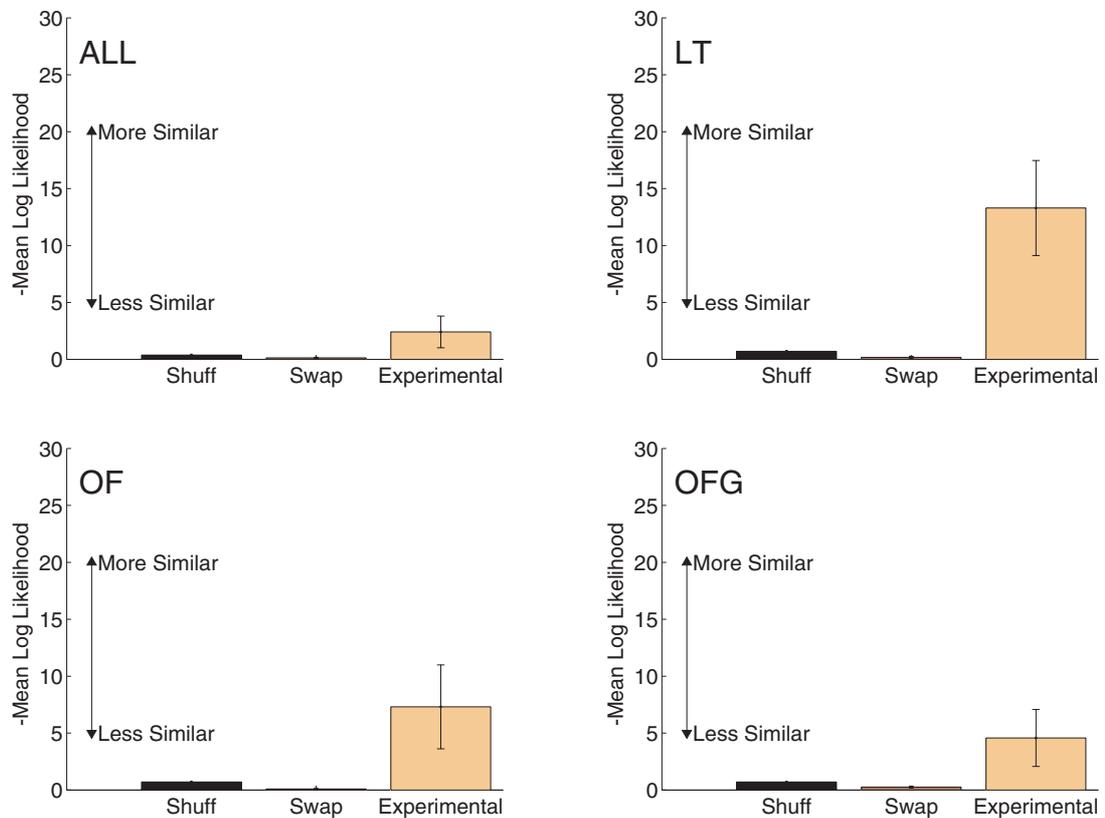


Figure 7.4: **Sharp waves during awake behavior include the same cell assemblies as occur during theta.** The similarity between cell co-firing during SWs and during theta are shown for each condition (see Methods). A one-factor ANOVA was used to compare each condition. The difference between the true similarity and the randomized controls was checked with Tukey post-hoc statistics (Zar, 1999). For all four conditions, the cell assemblies active during SW were more similar to those seen during behavior (theta) than would be expected from either random control, including SWAP (preserving timing and ensemble firing properties) and SHUFF (preserving each neuron's overall firing rate). Note that the ALL condition is an analysis over all sessions, not an average of the other three conditions. (ALL) Overall, including all three tasks: $P < 10^{-10}$, $F = 80.4$. (LT) Linear track, $P < 10^{-10}$, $F = 31.23$. (OF) Open field, $P < 10^{-9}$, $F = 23.39$. (OFG) Open field with goal, $P < 10^{-8}$, $F = 19.41$.

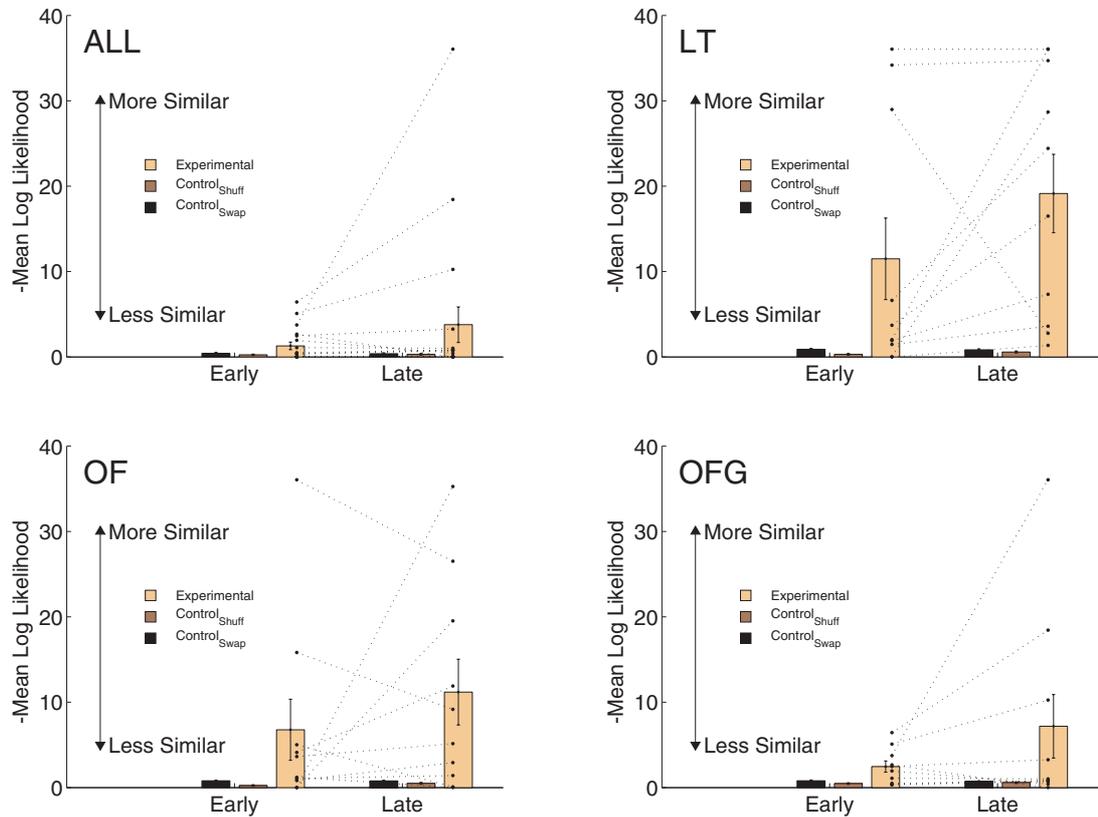


Figure 7.5: **The assemblies became more coherent through the session.** The sharp waves in each session were divided into two halves by the median occurring sharp waves, providing the same number of sharp waves in two blocks (an *early* block, and a *late* block). If the cell assemblies co-firing in the sharp waves become more similar to the cell assemblies occurring during theta, we would expect the similarity to increase between the two blocks. The similarity did increase for the linear track, and for the overall condition. But the increase was not significant for the two dimensional conditions. One-sided nonparametric Wilcoxon signed rank tests were used (Zar, 1999). (ALL) Overall, including all three tasks: $P < 0.02$. (LT) Linear track, $P < 0.05$. (OF) Open field, $P < 0.28$. (OFG) Open field with goal, $P < 0.10$.

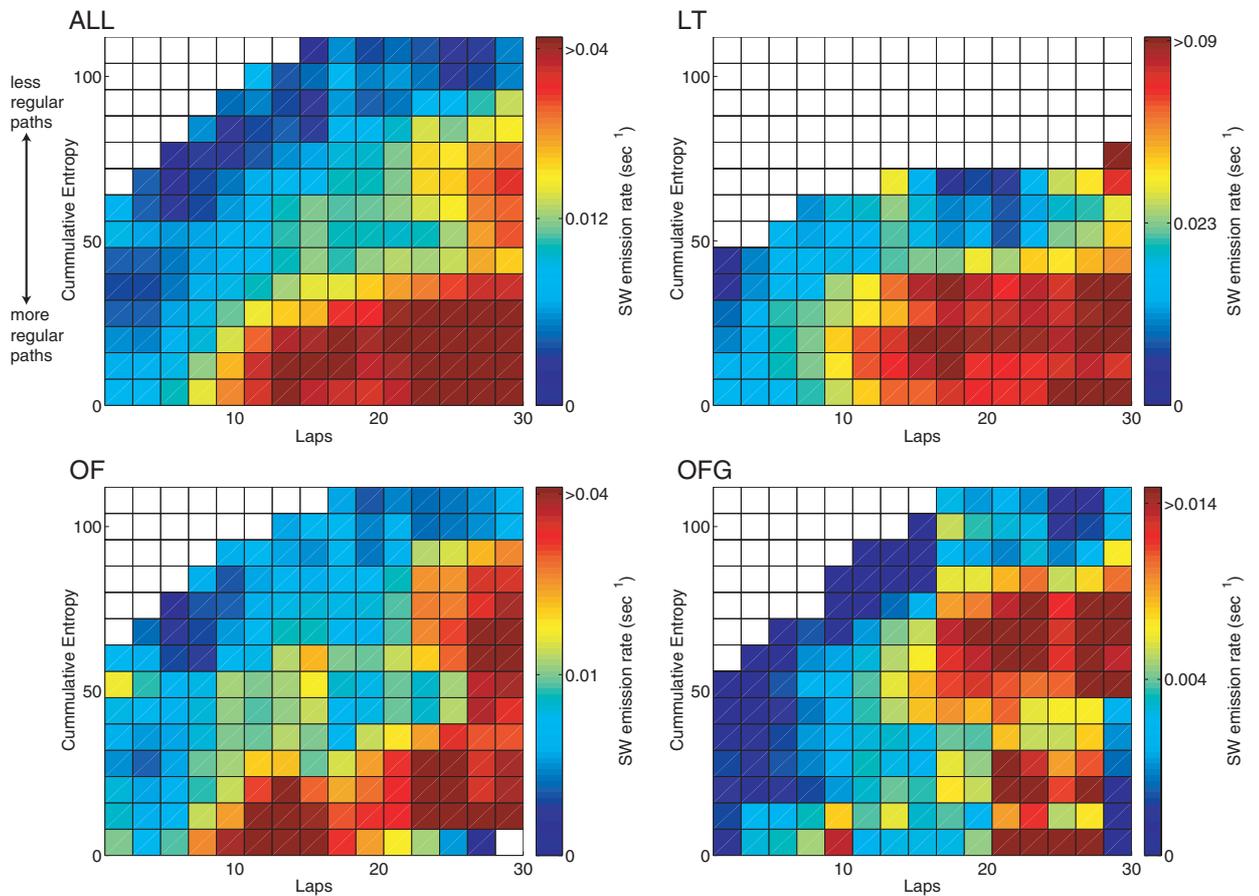


Figure 7.6: **Dependence of sharp wave emission on the sequential repetitiveness of the behavior.** # SW events normalized by time spent in non-theta states as a function of lap number and behavioral entropy. SW emission increased with lower entropy (more regular paths) and on later laps (with more experience). **(ALL)** Each lap for each session on each task (LT, OF, or OFG) contributed one 3-dimensional point to the analysis. For each bin, points were radially averaged to determine average SW emission given the cumulative regularity and experience. Stats: stepwise regression on raw (i.e. unaveraged) data showed an effect of lap number $P < 0.00001$, an effect of entropy $P < 0.00001$, and an interaction between the two $P < 0.00001$. **(LT)** Same as ALL except only LT sessions were used. Stats: stepwise regression on raw data showed an effect of lap number $P < 0.00001$, an effect of entropy $P < 0.0005$, and an interaction between the two $P < 0.00001$. **(OF)** Same as ALL except only OF sessions were used. Stats: stepwise regression on raw data showed an effect of lap number $P < 0.002$ (ns, by multiple comparisons), an effect of entropy $P < 0.001$, and an interaction between the two $P < 0.00005$. **(OFG)** Same as ALL except only OFG sessions were used. Stats: stepwise regression on raw data showed an effect of lap number $P < 0.002$ (ns, by multiple comparisons), an effect of entropy $P > 0.31$ (ns), and an interaction between the two $P < 0.0005$.

Because reactivation could not be measured for each lap, the total time spent in theta was used as the temporal measure. Theories suggest that information storage in hippocampus occurs during theta (Buzsáki, 1989; Hasselmo and Bower, 1993; Redish, 1999). Reactivation similarity also increased with more regular behaviors and time spent in theta as is evidenced by a gradient from high-entropy/low laps (upper left) towards low-entropy at higher laps (lower right). This effect was significant (stepwise multiple linear regression: significant effect of entropy $P < 0.005$, and of time spent in theta $P < 0.05$, with a strong interaction $P < 0.001$). The low number of data points for this analysis was sufficient only for the pooled data from the three tasks, not for individual task comparisons of reactivation as a function of spatial regularity and experience. See Fig. 7.7.

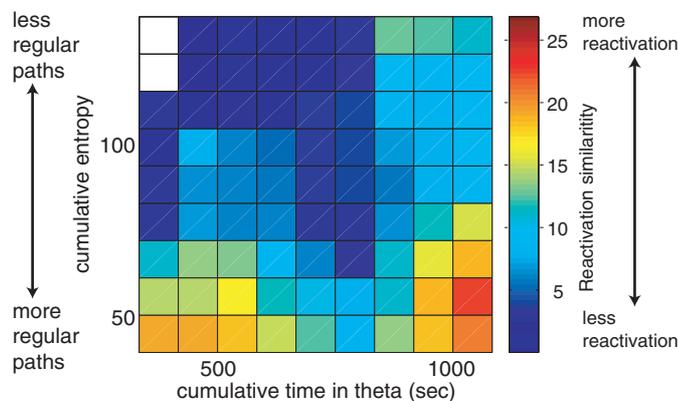


Figure 7.7: Dependence of reactivation on the sequential repetitiveness of the behavior. Reactivation similarity (measured as negative log likelihood of similarity relative to randomness) as a function of behavioral entropy and total time spent in theta. Note gradient from high-entropy/low laps (upper left corner) towards low-entropy at higher laps (lower right corner). Stats: stepwise regression showed an effect of entropy $P < 0.005$, for time in theta $P < 0.05$, and an interaction between the two $P < 0.001$.

7.4 Discussion

Taken together, these data confirm predictions that both sharp waves and reactivation develop with experience across multiple behavioral tasks. This increased probability of sharp wave emission depended on repetition and the regularity of the behavior. Ensemble firing patterns during task performance were reliably reactivated during sharp waves emitted on task. The similarity between the ensemble firing patterns in sharp waves and theta increased across the task and also depended on the regularity of the behavior.

On the linear track and multiple-T tasks, the animal's trajectories were highly repeated with specific reward delivery sites (back and forth along a thin track, around a loop on a track, respectively). In contrast, on the open-field tasks, the trajectories were highly variable and the average reward distribution was uniform. These differences in behavior and reward distribution may account for the higher rate of reactivation on LT. On the goal-oriented task (OFG), the repeated approach to the goal may account for the increase in SW emission rate with experience, which was not observed in the foraging task (OF). It is possible that instrumental activation of reward may also be important. Quantitative analysis of these variations in behavioral structure between LT, OF, and OFG revealed a gradation in the repetition of spatially sequential behavior. As shown in Fig. 7.6, combining the lap-to-lap variability of individual sessions across all tasks

showed that an increase in the occurrence of SW was strongly related to the cumulative sequential repetitiveness of behavior.

Since the density of SW emission during non-theta states depended on the behavioral structure of the task, our results are consistent with SWs being generated by the sequential firing of neurons as a result of plasticity induced by sequential behavior (Skaggs and McNaughton, 1996; Nádasdy et al., 1999; Csicsvari et al., 1999a; Redish, 1999). Our data are consistent with the SW itself being a consequence of a noise-driven firing cascade across potentiated synapses within CA3 (Ylinen et al., 1995; Shen and McNaughton, 1996; Redish and Touretzky, 1998; Csicsvari et al., 1999a, 2000; Behrens et al., 2005). Experimental evidence suggests SW activity in CA3 initiates CA1 SWs (Behrens et al., 2005; Buzsáki et al., 1983; Csicsvari et al., 1999a; Ylinen et al., 1995). Likewise, Jackson et al. (submitted) demonstrated an experience-dependent increase in SW emission in both CA1 and CA3 regions of the hippocampus. Theories of hippocampal function suggest that information is stored in the hippocampus during the theta state and replayed out via ripple activity during LIA (Marr, 1971; McNaughton et al., 1983; Buzsáki et al., 1983; Redish and Touretzky, 1998; Redish, 1999). Theories predict that asymmetric plasticity (Levy and Steward, 1983; Bi and Poo, 2001) applied to recurrent connections within CA3 through experience of repeated spatial sequences (Levy and Steward, 1983; Blum and Abbott, 1996; Redish and Touretzky, 1998) results in replay of this stored information. Our data strongly support this hypothesis — increases in

SW emission and reactivation were dependent on the level of repetition of regular spatial sequences.

Finally, it cannot be completely ruled out that gross physiological changes may underlie the neurophysiology reported here. For example, the increased rate of SW emission may depend on other variables correlated with experience such as increased body/brain temperature associated with prolonged behavioral output (Moser et al., 1994). However, these effects can be argued against in part due to persistence of the experience dependent SW emission increase on LT whether it was the first task in the series or the last task (data not shown) and to the stability of the non-theta state and the behavior on the LT task.

Chapter 8

Discussion

8.1 Coherency

8.1.1 Reconstruction versus Coherency:

Implications for understanding complex network states.

In the introduction, I asked “what is recall or confusion and how does the brain represent competing values in ambiguous situations?” In chapter 4, possible network interpretations of these situations were implemented: ambiguity that results from random activity and ambiguity that results from multimodal states (e.g. the transient states of the jump and competition simulations). While the traditional application of reconstruction techniques tend to avoid these issues, we have demonstrated that simple assumptions about neuronal tuning or about the binding of cell assem-

blies allow the construction of ensemble measurements that can detect the transient network anomalies mentioned above.

In our simulations studies (Jackson and Redish, 2003, 2004; see chapter 4), we compared the output of a standard reconstruction algorithm (Georgopoulos et al., 1983) to two classes of ensemble coherency measures. The population vector reconstruction of the direction represented by the network always yielded an answer even when random or bimodal activity patterns were present. This simplistic comparison addresses a more general trend in the application of reconstruction techniques. Reconstruction alone can not be used to infer internal states of an animal's sensory and cognitive networks such as the difference between random firing and well-represented variables. This is particularly important when considering issues of memory and recall. One function of memory is to appropriately link a current experience to a past experience; in the case of the hippocampus, this may mean using the same spatial map as was previously used in an environment. However, a primary usefulness of a memory is in its ability to influence disconnected experiences through recall of past events or episodes (see O'Keefe and Nadel, 1978 and Redish, 1999 for review). In this case of recall, one would expect that neuronal firing would, by definition, be disconnected from the current behavioral state of the animal. Recall may be detected by reconstruction methods as a reconstructed value that is very different from the current behavioral value. Usually, these values are considered noise to be removed from a reconstruction al-

gorithm (for example, see Zhang et al., 1998). Using a coherency method like those presented here, will allow an investigator to judge whether these aberrant reconstructions are truly valid representational events.

Johnson et al. (2005) demonstrated the power of these properties in the rodent head direction system. They recorded neuronal ensembles from the rodent post-subiculum and compared the coherency measure of equation 4.17 with the error in the reconstructed head direction signal computed using three methods: population vector reconstruction (Eq. 2.1; Georgopoulos et al., 1983, 1988), the optimal linear estimator method (Eq. 2.2; Salinas and Abbott, 1994), and a Bayesian method based on a Poisson firing assumption (See Eq. 2.3; Zhang et al., 1998). They found a high correlation between reconstruction error and coherency such that high coherency was associated with low reconstruction error while low coherency values co-occurred with high reconstruction errors. This indicates that in the head direction system, reconstruction error is not associated with the representation of head orientations different from the current physical orientation in the same environment. Coherency of the head-direction system may therefore be an important part of an animal's internal sense of directional certainty.

In short, coherency and reconstruction can be used as complementary tools to investigate the implications of our assumptions about neuronal tuning, information processing, and network states in the brain. Furthermore, since assumptions of neuronal tuning are required to imple-

ment any reconstruction methods, coherency allows the formulation of hypotheses about network function to statistically test the network-level implications of these assumptions.

8.1.2 Representational versus Non-Representational Coherency

One important aspect of ensemble coherence measures addressed in this dissertation is that assumptions about network function lead to testable hypotheses in the ensemble domain. To explore this assertion, we demonstrated that making assumptions about the tuning properties of neurons allows the construction of a representational quality measure such as *coherency*. To complement this analysis, we also demonstrated how simple assumptions of distributed processing and the formation of cell assemblies lead to hypotheses that can be tested with measures like *ensemble consistency*. These two measures constitute members of two classes of ensemble self consistency measures.

How Does One Choose Whether Representational or Non-Representational Measures Should Be Used? Simply put, representational measures are useful when representational questions are being asked, while non-representational measures can avoid the bias of assuming a particular network function and may allow the testing of the effects of experimental manipulations on networks of unknown function. It may seem like representational and non-representational measures are interchangeable, but the power of non-

representational measures arises when exploring brain structures of unknown or controversial function.

For example, representations in deep brain structures such as the nuclei of the basal ganglia are not well understood. The neuronal responses to behavioral variables in these structures are complex and there is much controversy in the literature as to the exact function and tuning of neurons in these networks. In such instances, questions related to learning and refinement in these networks can still be assessed using a non-representational measure such as ensemble consistency. This may avoid the complications involved in interpreting and communicating the results that might be obtained using a representational method.

On the other hand, a representational method may prove useful in testing specific hypotheses regarding the tuning of neurons in these structures. One could imagine assuming that striatal projection neurons encode limb movement direction, constructing a training set and tuning curves in one experimental paradigm and testing the coherency of the same ensemble under a different behavioral paradigm. If the behavioral context is important, the neuronal responses in the second paradigm would not fit the tuning derived from the first paradigm and the coherency would be very poor. It may be, however, that this network of neurons is still working in the same way and that cell assemblies are preserved. In such a case, a non-representational measurement would be relatively unchanged between the two paradigms.

Pitfalls of using Representational or Non-Representational Measures

The freedom from tuning-assumptions provided by our non-representational ensemble consistency measure comes at a cost. Since we are measuring the probability density of previously observed firing patterns at the location of a test firing pattern, this density will depend on the time the animal/network spent exploring that region of firing-rate space. Ultimately this is affected by how evenly a behavioral task requires the network to explore the range of possible factors that cause these firing patterns. Therefore, there is an implicit requirement placed on the experimenter to understand the possible tuning properties of the neurons in the network of interest and to either design experiments that evenly explore these variables or to shape the questions and analyses to avoid the adverse influence of this behavioral dependence. Thus, some representational knowledge is required.

Using the representational formulation of our coherency measurement, by contrast, allows us to normalize for behavioral effects from the beginning. This is done in the construction of the tuning curve, where the number of spikes observed at a given value of our behavioral variable is divided by the time spent at that value. The greatest challenge lies in the proper construction of the tuning curves for the question of interest. If one averages over events of interest, they will be undetectable since standard events and interesting events will both be seen as fluctuations around the mean. For instance, applying our coherency analysis to LT, OF, and OFG

using tuning curves constructed from the entire session did not reveal any state-switching effects that correlated with fluctuations in the dispersion z-score in our overdispersion analyses. This is because the tuning curves represent the average of both reference frames and both reference frames constituted consistent deviations from the expected network activity predicted by the tuning curves. Thus, one reference frame was not more or less “coherent” than the other. Once we had split the reference frames, we did see a difference in coherency associated with a map switch.

Our Use of Representational and Non-Representational Measures In our exploration of the overdispersion phenomenon, we employed a representational coherency measure to distinguish between the representation of one reference frame from another. To apply a non-representational measure such as the ensemble consistency method, we could have constructed training sets based on firing patterns from each map and measured how similar firing patterns before or after a switch were to one or the other training set. This would certainly have been acceptable, however the results would not have been nearly as simple to relate back to our representational assertions since we ultimately wanted to test the *representation* of reference frames.

In contrast, our exploration of reactivation during awake sharp wave ripples employed a non-representational measure to avoid the bias of comparing reactivation to spatial measures of behavior and to capture the tem-

poral aspects of encoding and reactivation laid out in the theoretical literature. For example, if map-switching can result in two neurons with similar place fields on opposite maps being active on a task, then those two neurons may never be simultaneously active. The Two-Stage Model (See Section 3.5.4) suggests that it is the simultaneous activation of neurons during theta that leads to the reactivated ensemble pattern during sharp waves. A naive representational measure would therefore, ignore the map-switching effects of our overdispersion result and predict reactivations that are not consistent with the theories we wanted to test.

Rate versus Timing Finally, we should note that the coherency measures in our simulations and overdispersion analyses are rate based. This has limitations when considering the precise timing of spikes and especially for dealing with bursting neurons. For instance, recent modeling had demonstrated that bursts are robust, quantal events, that allow the encoding of stimuli through spike count with low temporal variability between spikes (Kepecs and Lisman, 2003). This indicates that the highest information throughput in networks of bursting neurons may be contained in the number of spikes per burst, not in the firing rate (Kepecs et al., 2002; Kepecs and Lisman, 2003). The reactivation measure used in our sharp-wave ripple analyses, however, was based on the cross correlation between neurons. This is therefore a timing based measure that would be enhanced by mutual bursting of neurons.

8.1.3 Bayesian Methods and Uncertainty

There is much more that can be done to understand the consequences of our assumptions about the distributed encoding of information in neuronal populations. For instance, if one assumes that networks can encode probability distributions over behavioral variables, then ensemble measures can be formulated in probabilistic terms. For example, as briefly introduced in the introduction, it is possible to use Bayes rule to derive the probability of seeing a particular behavioral variable given the neuronal activity recorded in the ensemble. We write this as:

$$P(X|S) = \frac{P(S|X)P(X)}{P(S)} \quad (8.1)$$

where $P(X|S)$ is the probability of seeing a particular behavioral variable X given current neuronal activity S , $P(S|X)$ is the probability of observing neuronal activity S given behavioral variable X was observed, $P(X)$ is the probability of observing behavioral variable X , and $P(S)$ is the probability of seeing activity S . The reconstructed value X is then taken as the value that maximizes $P(X|S)$.

In the context of the results presented in the Overdispersion chapter, possible applications of this approach could include a Bayesian decoding paradigm with the explicit specification of a state modulation parameter related to the spatial reference frame being represented within the network. Instead of the simple spatial encoding model above, state variables

could be introduced such that equation 8.1 can be rewritten to include a state modulation parameter:

$$P(X|S, \Lambda) = \frac{P(S|\Lambda, X)P(\Lambda|X)P(X)}{P(\Lambda)P(S|\Lambda)} \quad (8.2)$$

where the added variable Λ is a state variable that affects the way an ensemble responds to some internal state. We may want to try to infer the probability of observing a particular state given the ensemble firing pattern and the behavioral variable X :

$$P(\Lambda|X, S) = \frac{P(S|\Lambda, X)P(X|\Lambda)P(\Lambda)}{P(X)P(S|X)} \quad (8.3)$$

Since $P(\Lambda)$ is unknown, a maximum likelihood method would be needed to reverse infer these parameters that we think influence Λ . This would amount to a Bayesian inference problem similar to that of a Hidden Markov Model. $P(S|\Lambda, X)$ would then contain our state dependent maps.

Other attractive possibilities associated with probabilistic methods of ensemble analysis include the use of entropy to measure the amount of disorder in the distribution of the represented variable (e.g. equation 8.1). High entropy distributions would be analogous to disordered firing of neurons in the ensemble, while low entropy distributions would correspond to all neurons agreeing more or less on the same variable. A possible interpretation of the entropy of $P(X|S)$, therefore may be the network's uncertainty with respect to variable X . Higher entropy would correspond

to greater uncertainty and visa versa. This may serve as quantitative measure of confusion and uncertainty and may represent an estimate of the amount of information needed to “convince” the network to choose on a particular value of X .

8.1.4 Coherency in the Brain

Our discussion of coherency so far, has revolved primarily around what coherency means to the researcher. Perhaps a more interesting question is, “what does coherency mean in the context of the brain?” In other words:

1. Does coherency matter to the network?
2. Does coherency matter to downstream structures?

Does coherency matter to the network? The answer to this question may depend substantially on the architecture of the network in question. But, first let us consider some sources or factors influencing coherency: input coherency and connectional coherency.

One major influence on a network’s coherency is the information impinging on a network. It should be fairly obvious that uniform random noise on the network inputs will be less helpful in organizing a network’s activation pattern than a single unambiguous network pattern. Furthermore, if the network is plastic such that it is learning from inputs, such as an associative network, an input that is similar to one the network has

seen before should result in more coherent activation than an input pattern very different from all previous inputs. This leads us to the other main factor, connectional coherency.

Whether a network can detect (i.e. is affected by) its own coherency depends on its internal connectivity. For example, if there is broad inhibition that is driven by broad input from excitatory neurons, coherent activation may mean that only a few cells are highly activated. In this case, incoherent activation could include the firing of a large number of neurons, which would then shut down the network via inhibitory feedback resulting in generally low firing rates. In such a case, sparse activity within a window of time is favored and broad network activation is extinguished. This, however, says nothing about which groups of cells are activated, if indeed more than one neuron *is* activated. Thus, the only constraint formed by a broad inhibition on what is considered “coherent” activity would be that the activity must somehow be sparse. Another example of the interaction of coherency and network connectivity would be the in the case of recurrent connectivity. In this case, the network forms attractor states such that particular patterns of activity are favored over others (Hopfield, 1982). Coherency would then be related to the depth of a pattern in the energy space. The most *coherent* patterns would be defined as those with the *lowest energy* *.

*Incoherent activity would be molded into a coherent output pattern by the attractor dynamics of the network in this case. The length of time this takes should depend on the energy space of the attractor such that an incoherent pattern that lies in a portion of the

Does coherency matter to downstream structures? The temporal effects of coherent activity on output structures is especially important in spiking networks where the relative timing of spikes across inputs is important. This is especially evident in biological systems where spike-timing dependent plasticity (STDP) is a common feature (Bi and Poo, 2001). In this case, one may imagine that a coherent activation in one network would organize the members of that pattern with of a higher probability of being active within some window of time. Such a coordination would increase the probability of spikes reaching neurons in a downstream structure more closely spaced in time, thereby increasing the potential for depolarizing a neuron and invoking plasticity mechanisms within the critical temporal window. If an asymmetric plasticity rule is in effect, the relative coherency within and between multiple input networks becomes increasingly important, such that coherent activation across multiple network will lead to maximal induction of plasticity.

8.2 Overdispersion

8.2.1 Reference Frames

Touretzky and Redish (1996) suggested that changing reference frames based on the track ends results in directionality on the linear track. Mc-space that has a small gradient and is far from any basin will take longer to converge to a coherent representation than a pattern near a basin with a steep energy gradient.

Naughton et al. (1996) suggested a similar framework where the map that an animal uses may be related to which navigational goal or landmark is most significant for the animal's immediate behavior. In models reviewed by Koene et al. (2003), the entorhinal cortex specifies potential goal locations given goal related input from pre-frontal cortex, CA3 specifies the animal's current location, and CA1 decides the appropriate next position to pursue in light of this goal information. Experimental data from Wood et al. (2000) and Ferbinteanu and Shapiro (2003) show that some CA1 pyramidal neurons may fire only when a specific goal location is intended by an animal and may also depend on the path taken.

The CA3 representation is a stable all-or-none mapping of a pattern associated with an environment (Lee et al., 2004b; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004). CA3 has been shown to demonstrate at least two independent methods for orthogonalizing its representation: complete remapping for changes in the global environment (Lee et al., 2004b; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004; Wills et al., 2005) and rate remapping to differentiate local environmental changes while preserving the coding for a particular spatial location with respect to distal cues (Leutgeb et al., 2005). Considering that CA1 gets input from CA3 as well as entorhinal cortex and can form a place representation in the absence of CA3 input (Mizumori et al., 1989; Brun et al., 2002), this suggests that the CA1 representation is related to the similarity between the entorhinal cortical spatial input and the CA3 localization of the animal to

a particular environment. Thus, the degree of partial remapping (Quirk et al., 1990; Markus et al., 1995; Anderson and Jeffery, 2003; Knierim, 2002) might therefore reflect the difference between environmental and task parameters.

Given these data (Mizumori et al., 1989; Wood et al., 2000; Brun et al., 2002; Ferbinteanu and Shapiro, 2003; Lee et al., 2004b; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004; Wills et al., 2005; Leutgeb et al., 2005) and the theoretical framework presented above (Touretzky and Redish, 1996; McNaughton et al., 1996; Koene et al., 2003), it makes sense to differentiate between entorhinal cortical input and CA3 input. Olypher et al. (2002) speculated that different inputs to the CA1 place cell may have different temporal dynamics; these data are therefore compatible with this idea. We suggest that it is in this sense of goal directed navigation that a reference frame is useful for describing an animal's location within a configuration of local and distal cues.

The model presented by Olypher et al. (2002) accounted for place cell variability by modulating a subset of inputs, randomly alternating between high and low intensity states. They varied the depth of modulation and the average switching period of these inputs and suggested that the best fit to the data reported by Fenton and Muller (1998) is a 10% modulation depth with a mean switching period of 1 s. The mean switching periods found in our data were about one third of this duration, approximately 350 ms or 3 theta cycles. To obtain such a large overdispersion

would require a deeper modulation than was suggested by Olypher et al. (2002). As can be seen in our data, this modulation can be up to 100% with entire place fields being turned on and off with each switch. It is possible that the partial remapping represented by this deep modulation is due to an interference between entorhinal input representing various goal locations and CA3 input representing the animal's current location. The model presented by Olypher et al. (2002) is actually similar to recent goal-directed navigation models (See Koene et al., 2003 for review). Our map-switching data also fit with data from Ferbinteanu and Shapiro (2003) and Wood et al. (2000) where the firing of some CA1 pyramidal neurons depended on the specific goal location intended by an animal and on the path taken. It is still not clear, however, whether the entorhinal cortex carries such goal-related information, in that the most extensive studies of entorhinal cortical neurons reveal grid-like firing resembling a semi-Cartesian representation of space (Fyhn et al., 2004; Hafting et al., 2005) that may be modulated by the animal's direction and velocity depending on the cortical layer of the neuron (Sargolini et al., 2006).

If this goal related input exists in entorhinal cortex, switching between goals may result in a switch to a different entorhinal cortical input pattern to CA1. This switch would have a partial remapping effect. Behaviorally, this switch could be as simple as a switch from wandering to chasing a pellet, or from chasing pellets to navigating to a hidden goal. In our data it is clear that there is a consistent switch happening on all three tasks we

examined. While it is not clear what causes a switch on the open field, the switching on the linear track and on open field goal occurred at goal locations consistent with the hypothesis that the switching of goals is communicated to the CA1 layer of hippocampus. This is evident in the significant suppression of reference frame switching prior to arrival at the goal on OFG or LT, and the significant increase in reference frame switching after goal entry.

The fact that the two reference frames and switching times produced by our analysis did not reduce the variability in neuronal firing to the unit variability expected from a Poisson point process model may suggest that more than two goal states are used by the animal. In OFG, there are at least three different motivational states that can be named: searching for pellets, chasing pellets that have just been delivered, and navigating to the goal. Others may include navigating to the wall (a preferred location of refuge for our rats) and navigating to the cue-card that spans a section of the wall (for some reason all of our rats tend to favor this location above all others). Applying these alternative motivational “goals” to OF we arrive at four potential motivational states that may effect CA1: searching for pellets, chasing pellets that have just been delivered, navigating to the wall, and navigating to the cue-card. The variability on LT was very low after accounting for state-switching, but it did not reach unity. This level may indicate the lower limit of our ability to predict place-cell firing with tuning curves based only on space. It has been suggested that velocity is also

encoded in the firing rate of hippocampal neurons (McNaughton et al., 1983; Huxter et al., 2003). We explored the possibility that this variability may be a summation of multiple influences including velocity, experience dependent plasticity, and the regularity of the animal's path through the place-field. None of these demonstrated a sufficient correlation to account for the excess variance observed, even after splitting representational reference frames. It was also noted that the majority of this remaining excess variability was due to a long tail of higher firing than would be predicted by the tuning curves derived for each state. Thus, variability may be due to the above influences as well as attentional or neuromodulatory levels. Given the overall shape of the firing distributions after accounting for the possible effects of reference-frame switching, however, it appears that the majority of the dispersion salient to the questions of this thesis has been accounted for.

Could CA3 show overdispersion? The model discussed so far for generation of overdispersion in CA1 is as follows: goal information filtered by entorhinal cortex directs the partial remapping of CA1 to reflect a reference frame with respect to this goal. CA3 receives direct entorhinal input as well as dentate granule input. The recurrent connections within CA3 are thought to generate an attractor network (McNaughton and Morris, 1987; McNaughton et al., 1996; Shen and McNaughton, 1996; Zhang, 1996; Samsonovich and McNaughton, 1997; Redish and Touretzky, 1997,

1998; Tsodyks, 1999; Redish, 1999; Káli and Dayan, 2000; Guazzelli et al., 2001). Recent experimental evidence has bolstered this notion (Lee et al., 2004b; Leutgeb et al., 2004; Vazdarjanova and Guzowski, 2004; Leutgeb et al., 2005). If it is possible for CA3 to rapidly switch between rate maps and/or spatial maps then this could also generate overdispersion. This switching would add further variability to the CA1 representation.

8.3 Experience Dependent Sharp Waves Ripples and Awake Reactivation

8.3.1 Sharp waves in CA1 are highly regulated network events

In our data, we observed substantial changes in sharp-wave emission rates and sharp-wave amplitude with experience. This increased amplitude is related to the level of population participation in CA3 and population synchrony (Csicsvari et al., 1999b). As discussed in section 3.4.2, there are a number of neuron types that regulate the sharp wave ripple oscillation. After a population burst from CA3 depolarizes CA1 pyramidal neurons and interneurons in the CA1 field, the ripple is initiated by axo-axonic and/or O-LM cell firing which hyperpolarizes the membrane of the pyra-

midal cells and de-inactivates the voltage sensitive ionic-channels preparing them for rebound bursting (Klausberger et al., 2003, 2004). The oscillation is then driven locally by interactions between the PV⁺ basket cells, bistratified cells, and the pyramidal cell population (Hirase et al., 2001; Csicsvari et al., 2000; Klausberger et al., 2003, 2004; Somogyi and Klausberger, 2005). The PV⁺ basket cells and bistratified cells fire phase locked to the ripple oscillation with peak firing rates occurring just after negative trough of the SW ripple (Klausberger et al., 2003, 2004). The termination of the sharp wave may occur after axo-axonic and/or O-LM cell firing. In our data, we do not see an increase in the ripple duration or in place cell firing rates. However, we did see an increase in sharp-wave ripple amplitude which is related to the strength of input from CA3 (Buzsáki et al., 1983; Csicsvari et al., 1999b). This suggests that the ripple duration is tightly regulated by the CA1 network of inhibitory interneurons, each with its own regulatory function within the sharp wave: the maintenance of firing rates through within-ripple inhibition by PV⁺ basket cells and bistratified cells and the maintenance of ripple duration by axo-axonic and/or O-LM cell firing (Klausberger et al., 2003, 2004; Somogyi and Klausberger, 2005).

8.3.2 Brain State and Sharp Wave Ripple Emission

There are many possible reasons why sharp wave ripple emission rates might grow with experience on a task. One reason for the growth in this emission rate could be that animals are simply more tired. It is known that

sharp waves are a feature of LIA, an EEG defined brain state that is common when animals rest (See O'Keefe and Nadel, 1978 for review). If animals spend more time in LIA, then one would expect more sharp-waves to be emitted. To control for this, we reported the rate of SWR emission as the number of sharp waves per second spent in LIA. We also measured the time in LIA on LT and found that there was no significant change over the window of behavior we were examining. Likewise, in case the SWR emission depended on the actual transition into LIA or on crossing some threshold of time spent in LIA, we measured the rate of transition into LIA and found no significant change on LT. Thus, changes in gross brain state were not likely to be the cause of increased SWR emission. To measure the outward behavioral indications of exhaustion, we looked at the time it took to run each lap and the time spent resting between laps, these did not change on LT. We also examined the sessions when LT was first versus the sessions when animals ran LT last, there was no noticeable difference in the SWR emission slope on these. Together, these data indicate that the increases in SWR emission are due to behaviorally induced changes within the brain.

One possible explanation for this increase in SWR emission also related to exhaustion would be sleep deprivation issues. Since rats are nocturnal and our rats were maintained on a synchronous day/night cycle, they were running during their sleep period. Recently it was shown that sleep deprived animals show a rebound in sharp wave ripple emission (Pono-

marenko et al., 2003). Ponomarenko et al. (2003) measured rebound of sharp-wave ripples after sleep-deprivation and stimulant treatment (amphetamine and two dosages of modafinil). Sharp wave ripple emission was dramatically higher after sleep-deprivation and stimulant treatment. This rebound in ripple emission was significantly correlated with the time the animal was awake (Ponomarenko et al., 2003). One could imagine that as animals run a task longer, they become more likely to nod off into early stages of slow wave sleep and then this rebound mechanism kicks in. Such a scenario may explain our increase in sharp-waves; however, this would argue for increased sharp-wave emission through out the entire session, which as discussed above was not observed comparing early and late linear track tasks. Furthermore, Ponomarenko et al. (2003) suggest that the increased theta activity resulting from modafinil and amphetamine treatment is the driving factor behind the sharp-wave ripple emission rebound. If this was the case, then the increased time in theta for their animals means more time for storage of information and argues in favor of our results, not against our results.

8.3.3 Replay generation

An important issue repeated throughout the reactivation literature is that there is no reactivation in slow-wave sleep prior to a task but enhanced reactivation in the sleep following a task (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996; Kudrimoti et al., 1999; Nádasdy et al.,

1999; Lee and Wilson, 2002). However, there has been no direct connection between the behavioral activity during a task with the development of this reactivation. Our results connect the behavioral dependence of previous slow-wave-sleep studies to specific task parameters, namely that the repetition of behavioral sequences is specifically responsible for the increased reactivation following a behavioral session. Furthermore, we explicitly examined the assumptions that theta-dependent ensemble-activity leads to the plasticity mechanisms that generate sharp waves. If Hebbian-like mechanisms underlie this increase in sharp waves, we also would expect that the regularity of the animal's path and hence the repeatability of neuronal activation sequences would also be important. In support of this hypothesis, sharp wave emission depended on both the repetition (the number of laps) and the regularity of this repetition (the path entropy). Since these LTP-like mechanisms are expected to underlie the formation of cell assemblies that are reactivated during sharp waves, we developed a measure of co-firing that specifically tested for the similarity in the cell assembly patterns between the theta and sharp wave states. The similarity in cell assemblies increased with both the time spent in theta states and the regularity of the behavior. These increases in SW emission and reactivation were robust on the linear track without concomitant increases in behavioral states or LFP states that are naturally associated with higher sharp wave emission. The other tasks provided a means of comparing the results of highly regular behavior to goal oriented two-dimensional explo-

ration and to random foraging. Taken together, these results confirm the predictions that arise from the Two-State Model: that theta-related plasticity is fundamental to the generation of sharp waves. These results are therefore consistent with the Two-State Model.

NMDA-dependence of Sharp Wave Emission Finally, plasticity mechanisms that could serve the general theoretical roles mentioned earlier (See section 3.5.4) have been observed experimentally (Bliss and Lømo, 1973; Bliss and Gardner-Medwin, 1973; Bi and Poo, 2001; Behrens et al., 2005; Kentros et al., 1998; Ekstrom et al., 2001; Shen et al., 1997). Temporally-asymmetric plasticity (Bi and Poo, 2001), SW observations in slice (Behrens et al., 2005), and place field modulation with experience (place-field expansion, Mehta et al., 1997) are dependent on NMDA-receptor integrity (Collingridge et al., 1983; Bi and Poo, 2001; Behrens et al., 2005; Kentros et al., 1998; Ekstrom et al., 2001; Shen et al., 1997). These data imply that experience-dependent effects on sharp-wave emission should also depend on NMDA-receptor integrity. While waking SW have been included in previous analyses (Kudrimoti et al., 1999; Foster and Wilson, 2006; O'Neill et al., 2006), neither the task-dependent properties nor the NMDA-R-dependence of sharp waves emitted during waking states have been explicitly studied. Collaborating with Adam Johnson, preliminary evidence suggests that sharp wave emission is also dependent on the integrity of NMDA-receptor mediated synaptic transmission (See Figure 8.1).

Together with the results in this thesis, these data would support learning theories which suggest that NMDA-dependent mechanisms strengthen cell-assemblies during tasks and that sharp waves reactivate those cell-assemblies in subsequent rest states. Cell assemblies have been directly observed in theta during behavior (Wilson and McNaughton, 1993; Harris et al., 2003; Leutgeb et al., 2005; Wills et al., 2005). Stabilization of a cell-assembly in a novel environment is dependent on intact LTP-mechanisms (Austin et al., 1990; Barnes et al., 1997; Kentros et al., 1998). Even in familiar environments, place fields expand along the direction of travel with experience within a session (Mehta et al., 1997; Lee et al., 2004a) and this expansion is NMDA-receptor dependent (Ekstrom et al., 2001; Shen et al., 1997), which implies that NMDA-receptor dependent LTP-like effects are occurring during behavior. As has previously been shown *in vitro*, sharp waves can be generated from CA3 (dissociated from its entorhinal cortex inputs), and NMDA-receptor dependent LTP induction increases the probability of sharp wave occurrence (Behrens et al., 2005). While extra-hippocampal effects of NMDA-receptor blockade cannot be excluded, our data nonetheless suggest that a mechanism necessary for LTP is also important for sharp wave emission *in vivo* and that the behavioral regularity thought to induce NMDA-receptor dependent plasticity in the hippocampus can also produce increases in sharp wave events as well as the reactivation of cell-assemblies.

8.3.4 Forward versus Reverse Replay

As discussed earlier (See sections 3.5.3 and 3.5.4), reverse replay has been observed in awake sharp waves (Foster and Wilson, 2006), while other studies of SW occurring during SWS reveal forward replay of behavioral

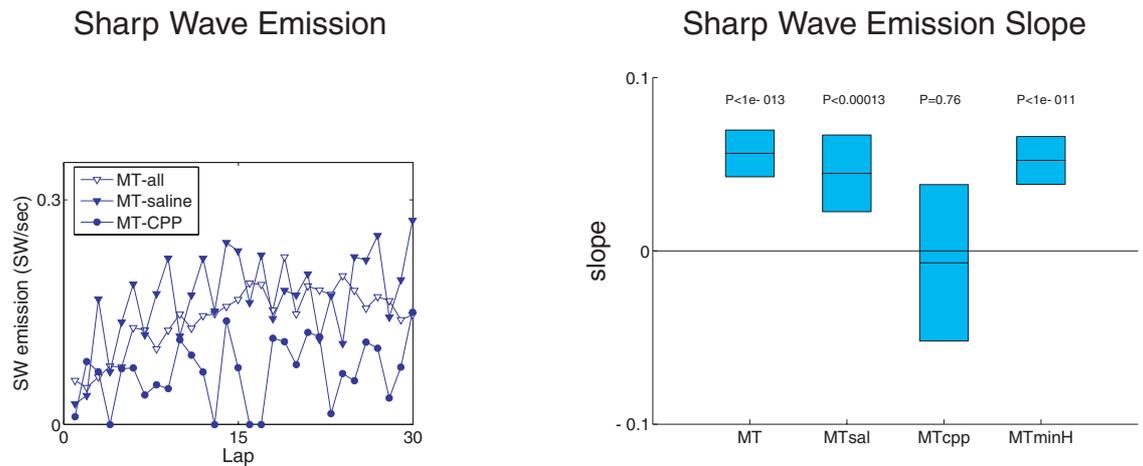


Figure 8.1: Dependence of sharp wave emission on NMDA receptor integrity. (left) Mean SW emission. # SW events per unit time spent in non-theta states as a function of lap number on their multiple-T (MT) task after IP injection of the NMDA-receptor antagonist (\pm)-3-(2-Carboxypiperazin-4-yl)propyl-1-phosphonic acid (CPP), MT after saline injection, and all MT sessions (all sessions prior to Saline or CPP administration sessions). Comparing SW emission across laps and pharmacological groups revealed a strong effect of pharmacological group ($P < 2 * 10^{-6}$, $F = 26.2$, two-way ANOVA: saline vs. CPP). In a separate ANOVA comparing a MT-untreated experiment with the saline and CPP sessions, SW emission rates on MT and MT-saline were not significantly different ($P > 0.05$, Tukey HSD criterion for multiple comparisons). SW emission rates after treatment with the NMDA receptor antagonist CPP (MT-CPP) were significantly lower than emission rates on either MT or after treatment with vehicle (MT-saline; $P < 0.05$, Tukey HSD criterion for multiple comparisons). (right) SW emission slope. Regression slope of SW emission per lap (z-score) and 95% confidence intervals of regression slope are shown for each treatment. SW emission increased in untreated (MT sessions prior to Saline or CPP, $P(\text{slope} = 0) < 0.00001$, $n = 7$) and saline injected animals (MT-Saline, $P(\text{slope} = 0) < 0.0002$, $F = 16.3$, $n = 3$). This increase in slope was not observed after CPP (MT-CPP, $P(\text{slope} = 0) = 0.76$, $F = 0.09$, $n = 3$, ns). Using sessions with at least as irregular paths as the best MT-CPP session did not abolish the increase in SW-emission (MT-high entropy, $P(\text{slope} = 0) < 0.00001$).

sequences (Skaggs and McNaughton, 1996; Nádasdy et al., 1999; Louie and Wilson, 2001; Lee and Wilson, 2002). Based solely on the NMDA-receptor depend mechanisms discussed above (Collingridge et al., 1983; Shen et al., 1997; Kentros et al., 1998; Bi and Poo, 2001; Ekstrom et al., 2001; Behrens et al., 2005), the theories we have discussed would predict that replay should be forward due to the asymmetric post-synaptic long-term potentiation of behaviorally activated synapses (Buzsáki, 1989; Buzsáki et al., 1994; McNaughton et al., 1996; Shen and McNaughton, 1996; Redish and Touretzky, 1998; Redish, 1999). However, one of these theories has explicitly stated that they predict reverse replay (Buzsáki, 1989; Buzsáki et al., 1994; Buzsáki, 1996). Buzsáki (1989) suggest that temporally graded potentiation of CA3 recurrents should result in the activation of the most recently and most highly potentiated synapses first transferring activation to the least potentiated synapses (assumed to be also the most temporally distant in this model) resulting in a compressed reversed-order replay of stored inputs at the termination of exploration (Buzsáki, 1989; Buzsáki et al., 1994; Buzsáki, 1996). This prediction of reverse replay is based on a labile potentiation of synapses in the CA3 recurrent matrix that is solidified by LTP mechanisms invoked during sharp-wave bursts (Buzsáki, 1989; Buzsáki et al., 1994; Buzsáki, 1996). This reverse replay mechanism would require a rapid decay in the potentiation of synapses. Such a rapid decay is observed shortly after LTP induction and is referred to as short-

term synaptic enhancement (STE[†]; see Fisher et al., 1997 for review). The STE phase is presynaptic, is comprised of at least four components, and can last milliseconds to minutes (Fisher et al., 1997). If this STE mechanism is behind the generation of reverse replay, we would predict that it should become forward-ordered as STE decays to baseline and LTP dominates.

In the behavioral context, a trip from one end of a linear track to the other would result in a rapid progression through theta-modulated firing of place-cells with place-fields along the way. The most recently activated place fields would have the highest STE, while the others would have already begun to decay. The mechanism described by Buzsáki (1989) would therefore lead to a reverse replay of the previously activated sequence as observed in Foster and Wilson (2006). However, as the STE phase decays to a stable potentiated level (viz. LTP), there should be a dis-ordering and then a reordering in the forward direction within minutes of ceasing behavior and commencing sleep.

We did not test these effects in our data, because we were interested in the more general cell assembly related reactivation effects that depend on experience. The measure that we employed was decidedly non-directional, such that ordering effects would not add noise to our results. However, it would be interesting to test these ordering effects directly by (1) blocking STE without effecting LTP and conversely (2) blocking LTP without

[†]One component of STE is post-tetanic potentiation (PTP)

effecting STE. One would expect that (1) should block reverse replay leaving forward replay intact, while (2) should block or reduce forward replay with little or no effect on reverse replay immediately after a lap on the linear track. Hippocampal LTP can be blocked through NMDA-receptor antagonists, and through interference with intracellular signalling mechanisms such as the Ca^{2+} /calmodulin-dependent kinase pathway in the post-synaptic structure, while STE is dependent upon multiple different Ca^{2+} -dependent pathways in the presynaptic terminal that can be differentiated by various pharmacological and genetic manipulations such as interfering with mitochondrial Ca^{2+} release (Fisher et al., 1997).

Self-localization versus Reverse Replay An alternative explanation of the reverse replay discovered by Foster and Wilson (2006), is that it could be construed to constitute the self-localization signal predicted by and demonstrated in the simulations of Redish and Touretzky (1998). In the case of their simulations, ambiguous sensory input upon entering an environment is filtered via the attractor dynamics of CA3 to localize the animal in the environment. This settling process would look like a “reverse replay” in that the hippocampal activity would converge on the animal’s location. If this were a linear track with the animal placed at one end, the activity in one reference frame would appear to move backward while the activity in the other reference frame would appear to move forward. Depending on whether or not one reference frame dominates the attrac-

tor dynamics, either or both possibilities may be observed. They assume that the animal has had previous experience on the environment such that the local views and path-integration coordinates have been learned and a “cognitive graph” (Muller et al., 1991; basically the spatial representation of this environmental reference frame) has been stored in the CA3 recurrent connections. One way to distinguish between a reverse replay and the self localization process would be the noisiness of the initial state which should be low in the case of a reverse replay, but should start high in the self-localization state and converge on the animal’s location. In fact, the data presented by Foster and Wilson (2006) show both types of phenomena (supplemental online material, Foster and Wilson, 2006). This difference in noise should be testable using the coherency methods presented here. The self-localization hypothesis is an interesting possibility that would conflict with the STE interpretation of the mechanism behind reverse replay since Redish and Touretzky (1998) envisioned this as an important part of memory recall when an animal enters an environment.

Given the theoretical underpinning of the Redish and Touretzky (1998) self-localization process, it would not be expected to be modified by manipulations of LTP or STE within the same session, since it should depend on already established connections. A further way of differentiating between self-localization and reverse replay is that reverse replay should be coherent at fine time scales (20-50 ms), while self-localization should start out as incoherent random activation and converge to a coherent represen-

tation at the animal's location. The coherency measures presented here should provide such information. A critical issue that may help resolve this question is whether or not sensory information can influence the hippocampus during awake sharp waves.

Does sensory information influence awake sharp waves? It is known that removing CA3 input to CA1 does not abolish the strong place selectivity in CA1 (Mizumori et al., 1989; Brun et al., 2002). Furthermore, the anatomical data indicates that CA1 neurons receive direct EC inputs and project back to the same cortical columns from which those inputs came (See Buzsáki, 1996 for review). Considering the anatomical data in conjunction with the experimental data, it is most likely that CA1 reactivates the patterns where the conjunction between entorhinal and CA3 input is most consistent. This is because the Schaffer collateral inputs from CA3 that synapse on CA1 cells receiving strong activation from EC will have relatively more advantage over synapses on CA1 cells without strong EC activation during a behavioral experience. The awake sharp waves reported by O'Neill et al. (2006) suggest that sensory information does influence neuronal firing during awake sharp waves ripples. In their data, the firing rates of neurons during a SWR depended on the location that the sharp-wave was emitted. If the SWR was emitted inside a neuron's place field, that neuron would have higher firing rate than if the SWR was emitted outside its place field (O'Neill et al., 2006). The effect is strong

enough that spatial tuning during SWRs looks very similar to spatial tuning during theta-related exploration (O'Neill et al., 2006). In some of our preliminary analyses (data not shown), it appears that on average the reconstruction is centered on the animal's current location. These data suggest that entorhinal input may be biasing the CA3 sharp-wave initiation sight and/or interacting through CA1 processing with the Schaffer collateral inputs that are carrying the sharp wave to ultimately carry information about the animal's current location.

If sensory information is indeed influencing hippocampal processing during sharp waves, this argues in favor of the Redish and Touretzky (1998) model of self localization, suggesting that the reverse replay observed by Foster and Wilson (2006) may be guided by sensory input to at least some extent. Whether or not the animal uses this information to localize itself in an environment remains an open question. If the awake SWR is self localization, then disrupting this process should result in increased instability in place fields according to the Redish and Touretzky (1998) model. In their model, the self-localization process is a path-integration reset. Disrupting path integration would disrupt the spatial reference necessary for maintaining register between local view inputs and self motion information.

In our data, we saw very low sharp-wave ripple emission at the beginning of a task which increased with experience. Since our animals were highly trained, the behavioral need for self-localization processes was un-

likely. Although we have not tested this, it is quite possible that the reactivation on our task may contain a mixture of reverse, forward, and disordered replays. This is because these highly-trained animals presumably start with well-organized connectivity between CA3 pyramidal neurons for these tasks, but further behavioral induction of potentiation and STE mechanisms will temporarily change the hierarchy of short-term synaptic potentiation in these synapses. As this short-term potentiation due to STE mechanisms swamps the stored connectivity in CA3, we would expect increased disorder of the reactivated patterns and initiation of more reverse-directed replays.

Chapter 9

Conclusion

In this thesis, we have explored the consequences of some basic concepts of ensemble coherency and neural information processing and through these ideas confirmed decades old predictions and discovered new insights into hippocampal function.

First, ensemble measures such as the various reconstruction methods discussed in the introduction have proven a useful means of probing the brain's networks to examine the ability of an ensemble to process behavioral variables (Georgopoulos et al., 1983; Wilson and McNaughton, 1993; Johnson et al., 2005). What is new is our expanded understanding of the ensemble: the importance of representational coherency and its relationship to network function. We demonstrated through simulation studies that applying concepts based on assumptions fundamental to the concepts of distributed representation and the cell assembly confers the ability to

probe the dynamics of neuronal information processing. After characterizing these ensemble consistency and representational coherency measures in a simulation environment, we demonstrated their applicability to physiological data. The concepts conveyed in this thesis are merely a drop in an ocean of possibilities. There still remains much work to be done to examine the theoretical, methodological, and biological implications of distributed representations.

Our application of concepts of distributed representations then allowed us to examine network related sources of place cell firing variability. It had already been known that the hippocampus can switch between reference frames on special tasks that require this cognitive switch for proper task performance (Redish et al., 2000). It was also known that goals exert a special influence on the hippocampus (Hollup et al., 2001b) and that removal of the hippocampus disrupts goal-directed navigation on uncued tasks (Morris et al., 1982; O'Keefe and Nadel, 1978; Redish, 1999). Furthermore, evidence existed that goal-directed tasks stabilized the variability in the firing of single cells in CA1 (Olypher et al., 2002; Kentros et al., 2004), but there was no clear network-level data that suggested reference frame-switching was related to this phenomenon. The novel contributions of the data presented in this dissertation are many in these regards. First, we have demonstrated that reference-frame or map switching is a *standard* operational process within the hippocampus on tasks as "simple" as random foraging. Reference frame switching is strongly related to or in-

fluenced by goal-directed behavior. This was made evident by the significant modulation of reference-frame switching in precise temporal relation to reward-related task cues on the linear-track and on our goal-directed open-field task. Interestingly, even though we have long known that individual directions on the linear track are often represented as separate maps, we were able to determine the precise time the animal switched between these maps: a the time of food delivery, before reaching the end of the track. This switching phenomenon explains, in part, the excess variability of place-cell firing reported in the hippocampus (Fenton and Muller, 1998) and supports hypotheses of the source of this variability (Lánksý et al., 2001; Olypher et al., 2002). However, the analyses presented here were unable to fully explain the excess variability in place-cells indicating that this model of reference-frame switching is not complete. More work needs to be done to understand the construction of reference frames within the networks of the hippocampus and surrounding cortices. It is also possible, that while our map-splitting analysis found two maps, the analysis may not have optimally stitched these maps together since our hill-climbing algorithm may settle at local maxima. Furthermore, it is certainly possible to apply the map-splitting analysis to extract more than two maps. While we did not have sufficient data to thoroughly examine more than two maps, experiments specifically targeting the details of this map-switching phenomenon could address this issue. For instance, longer recordings with probe trials or shifts in task parameters could be

employed to collect more data and specifically examine the dependence of the switch on reward parameters.

Finally, we set out to examine a long known source of spatial and temporal firing variability in the hippocampus: the awake sharp wave. We were able to confirm the detection of awake sharp waves in our task. We demonstrated that indeed, the same sub-ensembles activated during theta (the storage state) were reactivated during these awake sharp waves. Most importantly, it has been long hypothesized that the emission of sharp-waves should be dependent upon the level of potentiation between active neuronal assemblies during theta-related behavioral exploration. While recent in-vitro evidence confirmed the influence of hippocampal long-term-potentiation on the emission of sharp-waves (Behrens et al., 2005), we were able to confirm the behavioral consequences of repetition and regularity on the sharp wave emission rate. We demonstrated that increased behavioral regularity and repetition interacted to increase both the sharp-wave emission rate and the completeness of the reactivated cell assemblies. What remains to be understood is how these sharp wave events affect cognitive performance and the network coherency within the hippocampus and in down-stream cortical output structures. Specifically, it has been hypothesized that disrupting these sharp wave ripple events should disrupt memory trace formation. It is expected that the reactivation of cell-assemblies during these sharp wave bursts is critical for high-order associations that span discontinuous temporal episodes (Buzsáki,

1989; Buzsáki et al., 1994; Buzsáki, 1996; McNaughton et al., 1996; Buzsáki, 2005b; Kudrimoti et al., 1999; Redish, 1999). Therefore, experimentally intervening in the generation or transfer of sharp wave-related information should have specific effects on an animal's ability to form complex linkages between the cause-and-effect contingencies within tasks with temporal delays and interruptions. Indeed, there is strong evidence that the hippocampus is necessary for learning sequences of this type from various modalities (Scoville and Milner, 1957; Milner et al., 1968; Smith and Milner, 1981; Morris et al., 1982; Kesner and Novak, 1982; Reed and Squire, 1998; Redish, 1999; Clark et al., 2000; Fortin et al., 2002; Burman et al., 2006). In conclusion, assuming that information is processed in a distributed manner, applying the concept of the cell-assembly, and looking at within-ensemble dynamics has allowed us to probe more deeply into the subtleties of information processing dynamics in biological neural systems *in vivo*. This has enlightened our understanding of variability at both the single cell and network levels. We were able to use this understanding to discover the network-level modulation that results from various requirements placed on behavioral activity during awake, alert states. We examined how this activity and what kind of activity affects hippocampal on-line and off-line processing through the analysis of state-switching and sharp wave associated reactivation, respectively.

Bibliography

- Amaral DG (1987) Memory: anatomical organization of candidate brain regions. In: Handbook of Physiology; The Nervous System (Plum F, ed.), pp. 211–294. Progress in Brain Research, Bethesda: American Physiological Society.
- Amaral DG, Witter MP (1989) The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience* 31(3):571–591.
- Amari SI (1977) Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics* 27:77–87.
- Anderson MI, Jeffery KJ (2003) Heterogeneous modulation of place cell firing by changes in context. *Journal of Neuroscience* 23(26):8827–8835.
- Arai K, Keller EL, Edelman JA (1994) Two-dimensional neural network model of the primate saccadic system. *Neural Networks* 7(6/7):1115–1135.
- Arbib M, ed. (1995) *The Handbook of Brain Theory and Neural Networks*. MIT Press.
- Ashe J, Taira M, Smyrnis N, Pellizer G, Georgakopoulos T, Lurito JT, Georgopoulos AP (1993) Motor cortical activity preceding a memorized movement trajectory with an orthogonal bend. *Experimental Brain Research* 95:118–130.
- Austin KB, Fortin WF, Shapiro ML (1990) Place fields are altered by NMDA antagonist MK-801 during spatial learning. *Society for Neuroscience Abstracts* 16:263.

- Averbeck BB, Chafee MV, Crowe DA, Georgopoulos AP (2002) Parallel processing of serial movements in prefrontal cortex. *Proceedings of the National Academy of Sciences, USA* 99:13172–13177.
- Averbeck BB, Crowe DA, Chafee MV, Georgopoulos AP (2000) Neural activity during copying of geometrical shapes. *Society for Neuroscience Abstracts* .
- Barnes CA, Suster MS, Shen J, McNaughton BL (1997) Multistability of cognitive maps in the hippocampus of old rats. *Nature* 388(6639):272–275.
- Batschelet E (1981) *Circular statistics in biology*. New York: Academic Press.
- Behrens CJ, van den Boom LP, de Hoz L, Friedman A, Heinemann U (2005) Induction of sharp wave-ripple complexes in vitro and reorganization of hippocampal networks. *Nature Neuroscience* 8(11):1560–1567.
- Bi GQ, Poo MM (2001) Synaptic modification by correlated activity: Hebb's postulate revisited. *Annual Review of Neuroscience* 24(1):139–166.
- Blair HT, Lipscomb BW, Sharp PE (1997) Anticipatory time intervals of head-direction cells in the anterior thalamus of the rat, implications for path integration in the head-direction circuit. *Journal of Neurophysiology* 78(1):145–159.
- Bliss TVP, Gardner-Medwin AR (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology* 232:357–374.
- Bliss TVP, Lømo T (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology* 232:331–356.
- Blum KI, Abbott LF (1996) A model of spatial map formation in the hippocampus of the rat. *Neural Computation* 8(1):85–93.

- Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA (1998) A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience* 18(18):7411–7425.
- Brown EN, Kass RE, Mitra PP (2004) Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature Neuroscience* 7(5):456–461.
- Brun VH, Otnæss MK, Molden S, Steffenach HA, Witter MP, Moser MB, Moser EI (2002) Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science* 296:2243–2246.
- Burman MA, Starr MJ, Gewirtz JC (2006) Dissociable effects of hippocampus lesions on expression of fear and trace fear conditioning memories in rats. *Hippocampus* 16:103–113.
- Burwell RD (2000) The parahippocampal region: Corticocortical connectivity. *Annals New York Academy Of Sciences* 911:25–42.
- Buzsáki G (1989) Two-stage model of memory trace formation: A role for “noisy” brain states. *Neuroscience* 31(3):551–570.
- Buzsáki G (1996) The hippocampo-neocortical dialogue. *Cerebral Cortex* 6(2):81–92.
- Buzsáki G (2004) Large-scale recording of neuronal ensembles. *Nature Neuroscience* 7(5):446–451.
- Buzsáki G (2005a) Similar is different in hippocampal networks. *Science* 309:568–569.
- Buzsáki G (2005b) Theta rhythm of navigation: Link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus* 15:827–840.
- Buzsáki G, Bragin A, Chrobak JJ, Nadasdy Z, Sik A, Hsu M, A Y (1994) Oscillatory and intermittent synchrony in the hippocampus: relevance to memory trace formation. In: *Temporal Coding in the Brain* (Buzsáki G, Llinas R, Singer A, Berthoz W, Christen Y, eds.), pp. 145–172. Heidelberg: Springer.

- Buzsáki G, Leung LW, Vanderwolf CH (1983) Cellular bases of hippocampal EEG in the behaving rat. *Brain Research* 287(2):139–171.
- Christian EP, Deadwyler SA (1986) Behavioral function and hippocampal cell types: Evidence for two non-overlapping populations in the rat. *Journal of Neurophysiology* 55:331–348.
- Chrobak JJ, Buzsáki G (1998) Operational dynamics in the hippocampal-entorhinal axis. *Neuroscience and Biobehavioral Reviews* 22(2):303–310.
- Chrobak JJ, Buzsáki G (1994) Selective activation of deep layer (V-VI) retrohippocampal neurons during hippocampal sharp waves in the behaving rat. *Journal of Neuroscience* 14(10):6160–6170.
- Chrobak JJ, Buzsáki G (1996) High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat. *Journal of Neuroscience* 16(9):3056–3066.
- Chrobak JJ, Lörincz A, Buzsáki G (2000) Physiological patterns in the hippocampo-entorhinal cortex system. *Hippocampus* 10:457–465.
- Clark RE, Zola SM, Squire LR (2000) Impaired recognition memory in rats after damage to the hippocampus. *Journal of Neuroscience* 20(23):8853–8860.
- Collingridge G, Kehl S, McLennan H (1983) Excitatory amino acids in synaptic transmission in the schaffer collateral-commissural pathway of the rat hippocampus. *Journal of Physiology* 334:33–46.
- Csicsvari J, Hirase H, Czurkó A, Buzsáki G (1999a) Fast network oscillations in the hippocampal CA1 region of the behaving rat. *Journal of Neuroscience* 19(RC20):1–4.
- Csicsvari J, Hirase H, Czurkó A, Buzsáki G (1999b) Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *Journal of Neuroscience* 19(274–287):274–287.
- Csicsvari J, Hirase H, Mamiya A, Buzsáki G (2000) Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events. *Neuron* 28:585–594.

- Dolorfo CL, Amaral DG (1998) Entorhinal cortex of the rat: Organization of intrinsic connections. *The Journal Of Comparative Neurology* 398:49–82.
- Ekstrom AD, Meltzer J, McNaughton BL, Barnes CA (2001) NMDA receptor antagonism blocks experience-dependent expansion of hippocampal “place fields”. *Neuron* 31:631–638.
- Fenton AA, Muller RU (1998) Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proceedings of the National Academy of Sciences, USA* 95:3182–3187.
- Ferbinteanu J, Shapiro ML (2003) Prospective and retrospective memory coding in the hippocampus. *Neuron* 40(6):1227–1239.
- Fisher SA, Fischer TM, Carew TJ (1997) Multiple overlapping processes underlying short-term synaptic enhancement. *Trends In Neurosciences* 20(4):170–177.
- Fortin NJ, Agster KL, Eichenbaum HB (2002) Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience* 5(5):458–462.
- Foster DJ, Wilson MA (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* .
- Fyhn M, Molden S, Witter MP, Moser EI, Moser MB (2004) Spatial representation in the entorhinal cortex. *Science* 305(5688):1258–1264.
- Georgopoulos AP, Caminiti R, Kalaska JF, Massey JT (1983) Spatial coding of movement: A hypothesis concerning the coding of movement direction by motor cortical populations. *Experimental Brain Research Suppl.*(7):327–336.
- Georgopoulos AP, Kettner RE, Schwartz AB (1988) Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience* 8(8):2928–2937.
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419.

- Goodridge JP, Touretzky DS (2000) Modeling attractor deformation in the rodent head-direction system. *Journal of Neurophysiology* 83:3402–4310.
- Green JD, Arduini AA (1954) Hippocampal electrical activity in arousal. *Journal of Neurophysiology* 17:533–557.
- Guazzelli A, Bota M, Arbib MA (2001) Competitive hebbian learning and the hippocampal place cell system: Modeling the interaction of visual and path integration cues. *Hippocampus* 11(3):216–239.
- Hafting T, Fyhn M, Molden S, Moser MB, Moser E (2005) Microstructure of a spatial map in the entorhinal cortex. *Nature* 436:801–806.
- Harris KD (2005) Neural signatures of cell assembly organization. *Nature Reviews Neuroscience* 6:399–407.
- Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsáki G (2003) Organization of cell assemblies in the hippocampus. *Nature* 424(6948):552–556.
- Hasselmo ME, Bower JM (1993) Acetylcholine and memory. *Trends in Neurosciences* 16(6):218–222.
- Hebb DO (1949) *The Organization of Behavior*. New York: Wiley. Reissued 2002 LEA.
- Hirase H, Leinekugel X, Csicsvari J, Czurkó A, Buzsáki G (2001) Behavior-dependent states of the hippocampal network affect functional clustering of neurons. *Journal of Neuroscience* 21(RC145):1–4.
- Hoffmann KL, McNaughton BL (2002) Coordinated Reactivation of Distributed Memory Traces in Primate Neocortex. *Science* 297(5589):2070–2073.
- Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI (2001a) Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *Journal of Neuroscience* 21(5):1635–1644.

- Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI (2001b) Place fields of rat hippocampal pyramidal cells and spatial learning in the watermaze. *European Journal of Neuroscience* 13(6):1197–1197.
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences, USA* 79:2554–2558.
- Huxter J, Burgess N, O'Keefe J (2003) Independent rate and temporal coding in hippocampal pyramidal cells. *Nature* 425(6960):828–832.
- Jackson JC, Johnson A, Redish AD (2005a) Hippocampal sharp wave events increase during behavior with experience within session. *Computational Neuroscience* .
- Jackson JC, Johnson A, Redish AD (2005b) Sharp-wave events and correlated neuronal activity increase during behavior. *Society for Neuroscience Abstracts* .
- Jackson JC, Johnson A, Redish AD (submitted) Hippocampal sharp waves during awake states depend on repeated sequential experience and nmda-receptor integrity .
- Jackson JC, Redish AD (2003) Detecting dynamical changes within a simulated neural ensemble using a measure of representational quality. *Network: Computation in Neural Systems* 14:629–645.
- Jackson JC, Redish AD (2004) Measuring ensemble consistency without measuring tuning curves. *Neurocomputing* 58-60C:91–99.
- Jarosiewicz B, McNaughton BL, Skaggs WE (2002) Hippocampal population activity during the small-amplitude irregular activity state in the rat. *Journal of Neuroscience* 22(4):1373–1384.
- Jarosiewicz B, Skaggs WE (2004a) Hippocampal place cells are not controlled by visual input during the small irregular activity state in the rat. *Journal of Neuroscience* 24(21):5070–5077.
- Jarosiewicz B, Skaggs WE (2004b) Level of arousal during the small irregular activity state in the rat hippocampal EEG. *J Neurophysiol* 91:2649–2657.

- Johnson A, Seeland KD, Redish AD (2005) Reconstruction of the postsubiculum head direction signal from neural ensembles. *Hippocampus* 15:86–96.
- Káli S, Dayan P (2000) The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model. *Journal of Neuroscience* 20(19):7463–7477.
- Kentros C, E H, Hawkins RD, Kandel ER, Shapiro M, Muller RV (1998) Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280:2121–2126.
- Kentros CG, Agnihotri NT, Streater S, Hawkins RD, Kandel ER (2004) Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* 42:283–295.
- Kepecs A, Lisman J (2003) Information encoding and computation with spikes and bursts. *Network: Computation in Neural Systems* 14(1):103–118.
- Kepecs A, Wang XJ, Lisman J (2002) Bursting neurons signal input slope. *The Journal of Neuroscience* 22(20):9053–9062.
- Kesner RP, Novak JM (1982) Serial position curve in rats: Role of the dorsal hippocampus. *Science* 218(4568):173–175.
- Klausberger T, Magill PJ, Márton LF, Roberts JDB, Cobden PM, Buzsáki G, Somogyi P (2003) Brain-state and cell-type-specific firing of hippocampal interneurons *in vivo*. *Nature* 421:844–848.
- Klausberger T, Marton LF, Baude A, Roberts JD, Magill PJ, Somogyi P (2004) Spike timing of dendrite-targeting bistratified cells during hippocampal network oscillations *in vivo*. *Nature Neuroscience* 7(1):41–47.
- Knierim JJ (2002) Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. *Journal of Neuroscience* 22(14):6254–6264.
- Knierim JJ, Rao G (2003) Distal landmarks and hippocampal place cells: Effects of relative translation versus rotation. *Hippocampus* 13(5):604–617.

- Koene RA, Gorchetchnikov A, Cannon RC, Hasselmo ME (2003) Modeling goal-directed spatial navigation in the rat based on physiological data from the hippocampal formation. *Neural Networks* 16(5-6):577–584.
- Kohonen T (1977) *Associative Memory: A System-Theoretical Approach*. New York: Springer.
- Kohonen T (1982) Self-organized formation of topologically correct feature maps. *Biological Cybernetics* 43:59–69.
- Kohonen T (1984) *Self-Organization and Associative Memory*. New York: Springer-Verlag.
- Kudrimoti HS, Barnes CA, McNaughton BL (1999) Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *Journal of Neuroscience* 19(10):4090–4101.
- Lánksý P, Fenton AA, Vaillant J (2001) The overdispersion in activity of place cells. *Neurocomputing* 38-40:1393–1399.
- Lee AK, Wilson MA (2002) Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36:1183–1194.
- Lee I, Rao G, Knierim J (2004a) A double dissociation between hippocampal subfields: Differential time course of ca3 and ca1 place cells for processing changed environments. *Neuron* 42:803–815.
- Lee I, Yoganarasimha D, Rao G, Knierim J (2004b) Comparison of population coherence of place cells in hippocampal subfields ca1 and ca3. *Nature* 430:456–459.
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005) Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles. *Science* 309(5734):619–623.
- Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI (2004) Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1. *Science* 305(5688):1295–1298.

- Levy WB, Steward O (1983) Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* 8(4):791–7.
- Lin L, Osan R, Shoham S, Jin W, Zuo W, Tsien JZ (2005) Identification of network-level coding units for real-time representation of episodic experiences in the hippocampus. *PNAS* 102(17):6125–6130.
- Louie K, Wilson MA (2001) Temporally structured replay of awake hippocampal ensemble activity during Rapid Eye Movement sleep. *Neuron* 29:145–156.
- Mardia KV (1972) *Statistics of Directional Data*. New York: Academic Press.
- Markus EJ, Barnes CA, McNaughton BL, Gladden VL, Skaggs WE (1994) Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus* 4:410–421.
- Markus EJ, Qin Y, Leonard B, Skaggs WE, McNaughton BL, Barnes CA (1995) Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience* 15:7079–7094.
- Marr D (1970) A theory for cerebral neocortex. *Proceedings of the Royal Society* 176:161–234.
- Marr D (1971) Simple memory: A theory of archicortex. *Philosophical Transactions of the Royal Society of London* 262(841):23–81.
- Marr D (1975) Approaches to biological information processing. *Science* 190:875–876.
- McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review* 102(3):419–457.
- McNaughton BL (1983) Associative properties of hippocampal long term potentiation. In: *Neurobiology of the Hippocampus* (Seifert W, ed.), pp. 433–447. New York: Academic Press.

- McNaughton BL, Barnes CA, Gerrard JL, Gothard K, Jung MW, Knierim JJ, Kudrimoti H, Qin Y, Skaggs WE, Suster M, Weaver KL (1996) Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *Journal of Experimental Biology* 199(1):173–186.
- McNaughton BL, Barnes CA, O'Keefe J (1983) The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research* 52:41–49.
- McNaughton BL, Knierim JJ, Wilson MA (1994) Vector encoding and the vestibular foundations of spatial cognition: Neurophysiological and computational mechanisms. In: *The Cognitive Neurosciences* (Gazzaniga M, ed.), pp. 585–595. Cambridge MA: MIT Press.
- McNaughton BL, Morris RGM (1987) Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences* 10(10):408–415.
- Mehta MR, Barnes CA, McNaughton BL (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences, USA* 94:8918–8921.
- Miller KD (1995) Receptive fields and maps in the visual cortex: Models of ocular dominance and orientation columns. In: *Models of neural Networks III* (Domany E, van Hemmen JL, Schulten K, eds.). New York: Springer-Verlag.
- Milner B (1970) Memory and the medial temporal regions of the brain. In: *Biology of Memory* (Pribram KH, Broadbent DE, eds.), pp. 29–50. New York: Academic Press.
- Milner B, Corkin S, Teuber H (1968) Further analysis of the hippocampal amnesia syndrome: 14-year follow-up study of H. M. *Neuropsychologia* 6:215–234.
- Mizumori SJY, McNaughton BL, Barnes CA, Fox KB (1989) Preserved spatial coding hippocampus CA1 pyramidal cells during reversible suppression in CA3c output: Evidence for pattern completion in hippocampus. *Journal of Neuroscience* 9(11):3915–3928.

- Moore BR (1980) A modification of the Rayleigh test for vector data. *Biometrika* 67(1):175–180.
- Morris RGM, Garrud P, Rawlins JNP, O'Keefe J (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* 297:681–683.
- Moser EI, Moser MB, Andersen P (1994) Potentiation of dentate synapses initiated by exploratory learning in rats: dissociation from brain temperature, motor activity, and arousal. *Learning and Memory* 1(1):55–73.
- Muller RU, Bostock E, Taube JS, Kubie JL (1994) On the directional firing properties of hippocampal place cells. *Journal of Neuroscience* 14(12):7235–7251.
- Muller RU, Kubie JL, Saypoff R (1991) The hippocampus as a cognitive graph. *Hippocampus* 1(3):243–246.
- Munoz DP, Pélisson D, Guitton D (1991) Movement of neural activity on the superior colliculus motor map during gaze shifts. *Science* 251:1358–1360.
- Nádasdy Z, Hirase H, Czurkó A, Csicsvari J, Buzsáki G (1999) Replay and time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience* 19(2):9497–9507.
- Nerad L, McNaughton N (2006) The septal eeg suggests a distributed organization of the pacemaker of hippocampal theta in the rat. *European Journal of Neuroscience* 24:155–166.
- Niedermeyer E, Lopes da Silva F, eds. (1999) *Electroencephalography: Basic Principles, Clinical Applications, and Related Fields*. Baltimore MD: Williams and Wilkins, fourth edn.
- Obermayer K, Blasdel GG, Schulten K (1992) Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Physical Review A* 45(10):7568–7588.
- O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Experimental Neurology* 51:78–109.

- O'Keefe J, Conway DH (1978) Hippocampal place units in the freely moving rat: Why they fire where they fire. *Experimental Brain Research* 31:573–590.
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Research* 34:171–175.
- O'Keefe J, Nadel L (1978) *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press.
- O'Keefe J, Recce M (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.
- O'Keefe J, Speakman A (1987) Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research* 68:1–27.
- Olypher AV, Lánský P, Fenton AA (2002) Properties of the extra-positional signal in hippocampal place cell discharge derived from the overdispersion in location-specific firing. *Neuroscience* 111(3):553–566.
- O'Neill J, Senior T, Csicsvari J (2006) Place-selective firing of ca1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron* 49:143–155.
- Pavlidis C, Winson J (1989) Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *Journal of Neuroscience* 9(8):2907–2918.
- Pawelzik H, Hughes DI, Thomson AM (2003) Modulation of inhibitory autapses and synapses on rat ca1 interneurons by gaba-a receptor ligands. *Journal of Physiology* 546(3):701–716.
- Paxinos G, Watson C (1998) *The Rat Brain in Stereotaxic Coordinates*. New York: Academic Press, fourth edn.
- Pinto DJ, Brumberg JC, Simons DJ, Ermentrout GB (1996) A quantitative population model of whisker barrels: Re-examining the Wilson-Cowan equations. *Journal of Computational Neuroscience* 3(3):247–264.

- Ponomarenko AA, Lin JS, Selbach O, Haas HL (2003) Temporal pattern of hippocampal high-frequency oscillations during sleep after stimulant-evoked waking. *Neuroscience* 121(3):759–769.
- Quirk GJ, Muller RU, Kubie JL (1990) The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience* 10(6):2008–2017.
- Quirk MC, Wilson MA (1999) Interaction between spike waveform classification and temporal sequence detection. *Journal of Neuroscience Methods* 94:41–52.
- Redish AD (1999) *Beyond the Cognitive Map: From Place Cells to Episodic Memory*. Cambridge MA: MIT Press.
- Redish AD, Elga AN, Touretzky DS (1996) A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems* 7(4):671–685.
- Redish AD, McNaughton BL, Barnes CA (1998) Reconciling Barnes et al. (1997) and Tanila et al. (1997a, 1997b). *Hippocampus* 8:438–443.
- Redish AD, Rosenzweig ES, Bohanick JD, McNaughton BL, Barnes CA (2000) Dynamics of hippocampal ensemble realignment: Time vs. space. *Journal of Neuroscience* 20(24):9289–9309.
- Redish AD, Touretzky DS (1997) Cognitive maps beyond the hippocampus. *Hippocampus* 7(1):15–35.
- Redish AD, Touretzky DS (1998) The role of the hippocampus in solving the Morris water maze. *Neural Computation* 10(1):73–111.
- Reed JM, Squire LR (1998) Retrograde amnesia for facts and events: Findings from four new cases. *Journal of Neuroscience* 18(10):3943–3954.
- Rieke F, Warland D, de Ruyter van Steveninck R, Bialek W (1997) *Spikes*. Cambridge MA: MIT Press.
- Rosenzweig ES, Redish AD, McNaughton BL, Barnes CA (2003) Hippocampal map realignment and spatial learning. *Nature Neuroscience* 6(6):609–615.

- Salinas E, Abbott L (1994) Vector reconstruction from firing rates. *Journal of Computational Neuroscience* 1(1-2):89–107.
- Samsonovich AV, McNaughton BL (1997) Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience* 17(15):5900–5920.
- Sargolini F, Fyhn M, Hafting T, McNaughton BL, Witter MP, Moser MB, Moser EI (2006) Conjunctive Representation of Position, Direction, and Velocity in Entorhinal Cortex. *Science* 312(5774):758–762.
- Schmitzer-Torbert N, Jackson J, Henze D, Harris K, Redish A (2005) Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* 131(1):1–11.
- Schmitzer-Torbert NC, Redish AD (2002) Development of path stereotypy in a single day in rats on a multiple-T maze. *Archives Italiennes de Biologie* 140:295–301.
- Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry* 20:11–21.
- Sejnowski TJ, Destexhe A (2000) Why do we sleep? *Brain Research* 886(2):208–223.
- Sharp PE (1996) Multiple spatial/behavioral correlates for cells in the rat postsubiculum: Multiple regression analysis and comparison to other hippocampal areas. *Cerebral Cortex* 6(2):238–259.
- Sharp PE, Blair HT, Cho J (2001) The anatomical and computational basis of the rat head-direction cell signal. *Trends in Neurosciences* 24(5):289–294.
- Sharp PE, Kubie JL, Muller RU (1990) Firing properties of hippocampal neurons in a visually symmetrical environment: Contributions of multiple sensory cues and mnemonic processes. *Journal of Neuroscience* 10(9):3093–3105.
- Shen B, McNaughton BL (1996) Modeling the spontaneous reactivation of experience-specific hippocampal cell assemblies during sleep. *Hippocampus* 6(6):685–693.

- Shen J, Barnes CA, McNaughton BL, Skaggs WE, Weaver KL (1997) The effect of aging on experience-dependent plasticity of hippocampal place cells. *Journal of Neuroscience* 17(17):6769–6782.
- Silverman BW (1986) *Density Estimation for Statistics and Data Analysis* (Monographs on Statistics and Applied Probability). London: Chapman and Hall.
- Skaggs WE, Knierim JJ, Kudrimoti HS, McNaughton BL (1995) A model of the neural basis of the rat's sense of direction. In: *Advances in Neural Information Processing Systems 7* (Tesauro G, Touretzky DS, Leen TK, eds.), pp. 173–180. Cambridge MA: MIT Press.
- Skaggs WE, McNaughton BL (1996) Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271:1870–1873.
- Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6(2):149–173.
- Smith ML, Milner B (1981) The role of the right hippocampus in the recall of spatial location. *Neuropsychologia* 19(6):781–793.
- Smyrnis M, Taira M, Ashe J, Georgopoulos AP (1992) Motor cortical activity in a memorized delay task. *Experimental Brain Research* 92:139–151.
- Somogyi P, Klausberger T (2005) Defined types of cortical interneurone structure space and spike timing in the hippocampus. *J Physiol (Lond)* 562:9–26.
- Sparks DL (1986) Translation of sensory signals into commands for control of saccadic eye movements: Role of primate superior colliculus. *Physiological Review* 66(1):118–171.
- Squire LR (1992) Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychology Review* 99(2):195–231.

- Stewart M, Fox SE (1989) Two populations of rhythmically bursting neurons in rat medial septum are revealed by atropine. *Journal of Neurophysiology* 61(5):982–993.
- Touretzky DS, Redish AD (1996) A theory of rodent navigation based on interacting representations of space. *Hippocampus* 6(3):247–270.
- Tsodyks M (1999) Attractor network models of spatial maps in hippocampus. *Hippocampus* 9(4):481–489.
- van Opstal AJ, Kappen H (1993) A two-dimensional ensemble coding model for spatial-temporal transformation of saccades in monkey superior colliculus. *Network: Computation in Neural Systems* 4:19–38.
- Vanderwolf CH (1971) Limbic-diencephalic mechanisms of voluntary movement. *Psychological Review* 78(2):83–113.
- Vazdarjanova A, Guzowski JF (2004) Differences in Hippocampal Neuronal Population Responses to Modifications of an Environmental Context: Evidence for Distinct, Yet Complementary, Functions of CA3 and CA1 Ensembles. *J Neurosci* 24(29):6489–6496.
- Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J (2005) Attractor Dynamics in the Hippocampal Representation of the Local Environment. *Science* 308(5723):873–876.
- Wilson HR, Cowan JD (1973) A mathematical theory of the functional dynamics of cortical and thalamic tissue. *Kybernetik* 13:55–80.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.
- Wilson MA, McNaughton BL (1994) Reactivation of hippocampal ensemble memories during sleep. *Science* 265:676–679.
- Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H (2000) Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* 27:623–633.
- Ylinen A, Bragin A, Nadasdy Z, Jando G, Szabo I, Sik A, Buzsáki G (1995) Sharp wave-associated high-frequency oscillation (200 Hz) in the intact

- hippocampus: Network and intracellular mechanisms. *Journal of Neuroscience* 15(1):30–46.
- Zar JH (1999) *Biostatistical Analysis*. Upper Saddle River, NJ: Prentice-Hall, fourth edn.
- Zhang K (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience* 16(6):2112–2126.
- Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ (1998) Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology* 79:1017–1044.