Physiological mechanisms of hippocampal memory processing: experiments and applied adaptive filtering

by

David P. Nguyen

B.S. Electrical Engineering
Boston University, 1999

M.S. Electrical and Computer Engineering
Georgia Institute of Technology, 2001

Submitted to the Department of Brain and Cognitive Sciences on July 8, 2008
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Neuroscience
at the
Massachusetts Institute of Technology

September 2008

Author: _____
Department of Brain and Cognitive Sciences
July 8, 2008

Certified by: _____
Matthew Wilson
Professor of Neurobiology
Thesis Supervisor

Accepted by: _____
Matthew Wilson
Professor of Neurobiology
Chairman, Committee for Graduate Students

Physiological mechanisms of hippocampal memory maintenance: experiments and applied adaptive filtering

by

David P. Nguyen

The hippocampus is necessary for the formation and storage of episodic memory, however, the computations within and between hippocampal subregions (CA1, CA3, and dentate gyrus) that mediate these memory processing functions are not completely understood. We investigate by recording in the hippocampal subregions as rats execute an augmented linear track task. From these recordings, we construct ensemble rate representations using a point process adaptive filter to characterize single-unit activity from each subregion. We compared the dynamics of these rate representations by computing average max rate and average rate modulation during different experimental epochs and on different segments of the track. We found that the representations in CA3 were modulated most when compared to CA1 and DG during the first 5 minutes of experience. With more experience, we found the average rate modulation decreased gradually across all areas and converged to values that were not statistically different. These results suggest a specialized role for CA3 during initial context acquisition, and suggest that rate modulation becomes coherent across HPC subregions after familiarization.

Information transfer between the hippocampus and neocortex is important for the consolidation of spatial and episodic memory. This process of information transfer is referred to as memory consolidation and may be mediated by a phenomena called "replay." We know that the process of replay is associated with a rise in multi-unit activity and the presence of ripples (100-250 Hz oscillations lasting from 75ms to 100ms) in CA1. Because ripples result from the same circuits as replay activity, the features of the ripple may allow us to deduce the mechanisms for replay induction and the nature of information transmitted during replay events. Because ripples are relatively short events, analytical methods with limited temporal-spectral resolution are unable to fully characterize all the structure of ripples. In the thesis, we develop a framework for characterizing, classifying, and detecting ripples based on instantaneous frequency and instantaneous frequency modulation. The framework uses an autoregressive model for spectral-temporal analysis in combination with a Kalman filter for sample-to-sample estimates of frequency parameters. We show that the filter is flexible in the degree of smoothing as well as robust in the estimation of frequency. We demonstrate that under the proposed framework ripples can be classified based on high or low frequency, and positive or negative frequency modulation; can combine amplitude and frequency information for selective detection of ripple events; and can be used to determine the number of ripples participating in "long ripple" events.

Thesis Supervisor: Matthew Wilson
Title: Professor of Neurobiology

ACKNOWLEDGEMENTS

I would like to take this opportunity to acknowledge everyone who has, in large part, contributed to making graduate school not only possible but amazing in every way.

First, to my family near and far, through all the ups and downs, you have always put me first and encouraged me to do the best that I can. For this, I extend my greatest respect and endless appreciation.

To my closest friends, Dave Cianciolo, Femi "the thunder" Dosunmi, Hau Le, Anh Bungcayao, Katie Milton, The Tavitians (Dave, Diane, and Dillon), and Yvonne Cao, thank you for continually keeping me grounded and moving forward.

I am fortunate to have so many wonderful colleagues in the Wilson and Neurostatistics Research Labs. I thank all lab members, past and present, for being thoughtful and engaging. In particular, I thank Uri Eden and Patrick Purdon for sharing all their insights from sports to neuroscience, Matt Jones for teaching me the ins and outs of Newbury Street fashion and in vivo tetrode-ology, Davie Foster for building the Rat4 patch box, Tom McHugh for teaching me to be compassionate and for being a great source feedback, Yuki Koyama-McHugh for her advice and constant encouragment, Julie Scott for being super caring, Fabian Kloosterman for his collaborative spirit, Valerie Ego-Stengel + Daoyun Ji + Jun Yamamoto + Hector Penagos + Tom Davidson + Miguel Remondes + Kevin Kneville + Eric Jonas + Sujith Vijayan + Stephen Gomperts + Greg Hale + Stuart Layton + Dan Bendor + Gabriela Czanner + Anne Smith + Anne Dreyer + Antonio Molins + Camilo Lamus + Iahn Cahigas + Dennis Dean + Sri Sharma + Gordon Pipa + Victor Solo for their support and effort to help, and James H. Williams for making the day bright with his smile.

To all the donors and administrative personnel affiliated with the BCS department, I extend my gratitude for your kindness and attention to graduate student welfare. In particular, Denise Heintze has been simply one of the most caring, outstanding, and wonderful people I've ever met. In addition, Brandy Baker has proven time and again to be super helpful and full of chocolate. Thank you to the Singleton, Picower, and McGovern families for sharing your vision and helping to make neuroscience at MIT a reality.

I extend my deepest gratitude to my long time mentor, Professor Emery Brown for his constant, unwavering support. As an undergraduate working in his lab up till this very day, I have learned directly and indirectly from Emery the value of hard work and collaboration. It is through our interactions that I have come to see the bigger picture in science and life.

I feel exceptionally fortunate to have the support and guidance of Professor Riccardo Barbieri over the years. Riccardo's insights into everything between life, science, and engineering have truly made a significant impact on my development. In regard to this disseration, Riccardo has gone above and beyond by reading every revision with patience. Thank you.

I thank Professor Christopher Moore for helping me to find common sense in a life not so different from those in phdcomics, and, even more, for helping me to find just the right balance between exploration and focus.

I wholeheartedly thank my advisor, Professor Matthew Wilson, for looking out for me all these years and providing me with all the resources I could possibly need to explore my interests in neuroscience. I truly value what Matt has patiently instilled in me: the process of due diligence in science, that there is nothing wrong with following your own convictions, and sleeveless workout shirts never go out of style.

Thesis Committee Members

Matthew A. Wilson, Advisor
Professor of Neurobiology, Department of Brain and Cognitive Sciences and Biology, MIT

Emery N. Brown, Committee Chair
Professor of Computational Neuroscience and Health Sciences and Technology,
Department of Brain and Cognitive Sciences, MIT-Harvard Division of Health Sciences and Technology, MIT

Christopher I. Moore, Internal Member
Assistant Professor of Neuroscience,
Department of Brain and Cognitive Sciences, MIT-Harvard Division of Health Sciences and Technology, MIT

Riccardo Barbieri, External Member
Assistant Professor of Anesthesia,
Harvard Medical School, Massachusetts General Hospital

CONTENTS

Chapter 1 : INTRODUCTION

The hippocampus is essential to the function of spatial memory encoding and consolidation. Without it, we lack the ability to store episodes of where we have been, whom we have spoken with, and the stories that define us. Such abilities are as central to the human condition as the hunger of curiosity, which compels us to understand the nature of our own existence.

This thesis explores the hippocampus from the perspectives of the electrophysiologist and biomedical engineer. Experimentally, we recorded ensemble single neurons from three sub-circuits of the hippocampus (CA1, CA3, and dentate gyrus) during a spatial task. On the signal processing side, we developed and applied analytical tools for measuring millisecond timescale activity in the brain. The combination of the two allowed us to state with better certainty the respective rate dynamics of each hippocampal sub-circuit, and to more accurately measure the structure of extracellular oscillations.

For the remainder of the introduction, we discuss a view of hippocampal function using previously published results, clearly state the questions addressed in the thesis, and provide an overview of the thesis.

BACKGROUND

The role of the hippocampus (HPC) is to aid in the maintenance of memories through the processes of information encoding and memory consolidation. These capacities of the HPC were first realized when the post-hoc analysis of a human brain lesion of the HPC led to graded retrograde amnesia [memories that were more distant in time from the surgery were more completely recalled] (Scoville & Milner, 1957). Later, in the rodent, place cells were found in the CA1 subregion of the HPC. The cells were termed "place cells" for their activity was highly correlated with the location of rat (O'Keefe, 1976). Since the discovery of place cells, there has been an exponential increase in the number of studies that examine the role of the HPC in spatial learning. Furthermore, interest in the HPC is now well represented in every concentration of neuroscience. This, in large part, is due to the observation of spatial memory related activity such as place fields, long term potentiation (LTP), and replay. In the thesis, we specifically ask what are the respective roles in the hippocampal sub-circuits in the encoding of spatial memory, and how can we better understand the nature of ripple associated replay events?

Anatomy

Examining the cross-section of the hippocampus (HPC) will reveal a three layered structure with four anatomically distinct subregions called CA1, CA3, dentate gyrus (DG), and subiculum (Witter & Wouterlood, 2002). These layers of the HPC are striking in the coronal brain slice of rats and humans, and can easily be recognized in the brains of many vertebrates - even the enchidna (spiny anteater), which evolved from a line that diverged from the human line 200 millions years ago (Manns & Eichenbaum, 2006).

Each subregion contains a simple putative excitatory cell layer containing only one type of projection (output) cell type; in CA1 and CA3 the cell type is the pyramidal cell, while in the DG it is the granule cell. Pyramidal cells have somas that are shaped like pyramids

which range from 20 to 30 microns in diameter. Granule cells have compact cylindrical somas and have an average diameter of 10 microns. In one hippocampi, the counts for the CA1 pyramidal cells are 3.0×10^5, CA3 pyramidal cells are 2.5×10^5, and DG granule cells are 1.2×10^6.

Figure 1.1: **The basic anatomy of the rodent hippocampus. A sagittal cross section of the hippocampus is shown with the respective inputs from the entorhinal cortex. Th neurons shown in the principal cell layers of the hippocampus are pyramidal neurons for CA1 and CA3, and granule cells for DG. The inputs to the hippocampus originate from the superficial layers of the entorhinal cortex, while the direct output of the hippocampus to the entorhinal cortex is the projection from CA1 to layer V of the entorhinal cortex. The modulatory input refers to a large number of projections, including the GABA-ergic input from the medial septum, which is a direct source of theta modulation in the hippocampus. Reprinted by permission from Macmillan Publishers Ltd: (Neves et al., 2008).**

Interneurons in the HPC are extremely diverse and their classification is based on features such as spike waveform, average firing rate, presence of parvalbumin, and morphhology (Han et al., 1993; Freund & Buzsaki, 1996; Klausberger et al., 2003). Although hippocampal function cannot be fully understood without accounting for interneuron

activity, relatively little about interneuron information processing has been published. This, primarily, is because interneurons are diverse and precise identification of interneuron cell types from electrophysiological signals alone is not generally reliable.

The primary inputs into the HPC originate from the superficial layers of the EC through a fiber tract called the perforant path (PP). This tract perforates at the hippocampal fissure where one branch runs along the stratum lacunosum-moleculare through CA1 and ends in CA3 all the while making synapses onto CA1 and CA3 apical dendrites; and the other branch runs through the molecular layers of the DG making synapses onto inhibitory neuron and granule cell dendrites. The axons (outputs) of the granule cell are called mossy fibers; these fibers project exclusively to the apical dendrites of the CA3 pyramidal cells. The axons of the CA3 pyramidal cells are sent to the dendrites of the CA1 pyramidal cells and, along the way, many of these CA3 axons synapse back onto CA3 pyramidal cell dendrites. The axons that provide feedback in CA3 are termed *recurrent collaterals*. The axons of CA1 project mainly to the subiculum and some to the deep layers of entorhinal cortex, as well as to the neocortex (Figure 1.1).

The perceptual system drives place cell activity

The receptive fields of hippocampal neurons are highly correlated with spatial location (O'Keefe, 1976; Muller et al., 1987). Therefore, the hippocampal neurons are referred to as "place cells" and their receptive fields as "place fields" (O'Keefe, 1976). Rotenberg & Muller (1997) proposed the theory that place cell activity in the rodent hippocampus reflects a perceptual system (Rotenberg & Muller, 1997). Central to the idea of a perception is the notion of a context. Based on the perceptual system theory we can define a stable context to be a set of observable cues with some non-zero valence that continually elicits stable responses from hippocampal cells as a function of position. In general, it is the geometric relationship between the position of the cues and the position of the animal that determine the location, shape, and rate of hippocampal response (Muller et al., 1987;

Burgess & O'Keefe, 1996; Leutgeb et al., 2004). The anatomical substrate for this theory is shown in Figure 1.2, where the coincident firing of neocortical neurons, then parahippocampal neurons, results in the activation of place cells (Marr, 1971; Squire & Zola-Morgan, 1991).
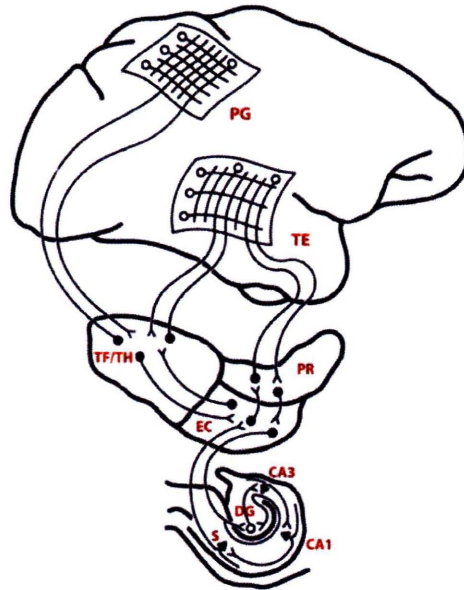
Figure 1.2 : Schematic relationship between primate neocortex and constituents of the parahippocampal region. The circuits sketched on the surface of the neocortex are the parietal (PG) and inferotemporal (TE) cortices. Activity in the parietal is highly correlated with object location while activity in the inferotemporal area is correlated with object recognition. The abbreviations are parahippocampal cortex (TF/TH), entorhinal cortex (EC), perirhinal cortex (PR), subiculum (S), dentate gyrus (DG), cornu ammonis 1 (CA1), cornu ammonis 3 (CA3). This circuit puts into context the role of the medial temporal lobe in transforming perceptual information into spatial information. Adapted from (Squire & Lindenlaub, 1990), reprinted with permission from Schattauer Verlag, Stuttgard, Germany.

Under the perceptual system theory, changes to cues with zero valence will have little effect on the position of place fields. Likewise, the perceptual system theory predicts that an alteration of the geometric relationship between animal position and salient cues will result in a change in place field shape, location, and/or rate. Muller & Kubie (1987) found that when the walls of rectangular and circular enclosures where altered to double the area of the enclosure (while the rat was not in the enclosure), one third of the place fields doubled in size without significant changes to field shape (Muller & Kubie, 1987). In contrast, when Wilson & McNaughton doubled the area of a square enclosure by lifting up a divider to allow full access to a rectangular enclosure, they found that the place cells in the familiar square area did not change dramatically (Wilson & McNaughton, 1993). The difference between these two experiments was that in the Wilson & McNaughton study the rat remained in the environment while the divider was being removed. A reasonable interpretation of this scenario is the rat was able to lower the perceptual valence of the divider as it was removed, and therefore maintain its current neuronal representation of context based on other cues. In contrast to Wilson & McNaughton, Leutgeb & Moser (2007) found that while the rat was either in a circular or square enclosure with high walls, when the walls of the enclosure were gradually morphed into the other shape, place fields in the hippocampus were found to change in rate and location (Leutgeb et al., 2007). Under the perceptual system theory, these changes in rate could be explained by the fact that the walls used in the Leutgeb & Moser experiment were too high for the rat to orient itself using distal cues and therefore the responses in the neocortex that were once coincident, naturally diverged as the geometry of the local cues diverged, leading to changes in hippocampal receptive fields. Interestingly, the perceptual system which drives the hippocampal response appears to be modular and robust in nature. For example, when the light is turned off hippocampal place cells are initially perturbed but become stable once again with some experience in the dark (Okeefe, 1976; Quirk et al., 1990; Gothard et al., 2001; Hargreaves et al., 2007). These observations are consistent with the idea that the rat can shift its perceptual system to increase the valence of vestibular information and still maintain the same place specific responses in the hippocampus. In particular, the stabilization of place fields in the dark suggests that pattern completion computations of

the hippocampal-entorhinal loop may allow for persisting hippocampal receptive fields in situations where visual cue information is not available.

Mechanisms of spatial memory encoding

The first line of evidence which implicates the hippocampus as part of the spatial memory consolidation circuit came from the observation of patient H.M., who in 1953 underwent a bilateral lesion of the medial temporal lobes in a procedure to cure his severe epilepsy. The lesions left H.M. with the capacity to recall only episodes that occurred 11 years before or earlier to the surgery, and no ability to recall new episodic experiences (Scoville & Milner, 1957; Salat et al., 2006). In addition to H.M, another patient R.B. developed marked anterograde amnesia at the age of 52. After the death of R.B., examination of his brain revealed total bilateral lesions of the CA1 subfield from rostral to caudal end (Zola-Morgan et al., 1986). The case of R.B. demonstrates that damage to the hippocampus alone can lead to severe impairment of memory processing. In particular, the fact that the amnesia was less severe in the case of R.B. than H.M., suggests that the other subregions of the hippocampus may act in parallel with CA1 to process memory (Squire & Zola-Morgan, 1991).

More detailed investigation of the function of the hippocampus was afforded by rodent *in vivo* studies of HPC activity during behavior and sleep. First, place receptive fields where observed in the hippocampus (O'Keefe, 1976). This was followed by the observation that the θ-rhythm strongly modulated the activation of place cells (Okeefe & Recce, 1993). Close examination of this modulatory activity led to the discovery of phase precession and spatial sequence compression (Skaggs et al., 1996). With the discovery of the LTP in the hippocampus (Bliss & Lomo, 1973), and the establishment of the spike-timing dependent plasticity (STDP) rule for strengthening or weakening synapses (Dan & Poo, 2004), the timescale of sequenced cell activations imposed by the θ-rhythm led to the idea that the θ-rhythm could act as a mechanism for binding information into a continuous chain of

place field activations that could serve as a memory trace for episodic information (Skaggs et al., 1996). The ability to reconstruct spatial episodes from sequential place cell activations in combination with the observation that pair-wise correlated firing during behavior and sleep was preserved (Wilson & McNaughton, 1993; Wilson & McNaughton, 1994), provided the scientific foundation for investigating a phenomenon called "replay", where episodic memory traces stored in the hippocampus would be re-expressed during sleep as part of the process of memory consolidation. Investigations in the Wilson lab reported the observation of highly unlikely expression of place cell sequences that matched those observed during behavior, supporting the view that the HPC is a locus for short term episodic memory and that the process of memory consolidation involves active engagement between the HPC and neocortex (Louie & Wilson, 2001; Lee & Wilson, 2002). A recent report by Ji & Wilson (2007) showing sequenced replay in the HPC was accompanied by sequenced replay in the visual cortex during slow wave sleep, was the first convincing line of evidence to show a sequence replay specific dialogue between the HPC and neocortex (Ji & Wilson, 2007). Perhaps the most striking observation of replay was discovered by Foster & Wilson (2006), where they observed that during awake, resting states, ripple associated activation of multi-unit place cell activity was highly structured in reverse order such that neurons fired in the reverse sequence in which they were activated during behavior. This observation of "reverse replay" is direct evidence showing that the hippocampus can recapitulate sequence reactivation of place cells without engaging in exploratory behavior itself. This brief summary ties together basic observations of hippocampal physiology to frame a larger picture of hippocampal function: the storage and recapitulation of place cell sequences as episodic memory traces.

The roles of pharmacological and genetic manipulations have made enormous contributions to the understanding of HPC function in memory encoding and consolidation. A famous study by Morris & O'keefe (1982) tested the effect of total bilateral hippocampal lesions on the performance of rats in a task where the rat was given several trials to learn a platform location that was hidden underwater. This task, termed the Morris water maze, was designed to probe the function of the hippocampus in the

processing of spatial memory. Morris & O'keefe found that lesioned rats showed marked impairment in the recall of platform position even after many trials (Morris et al., 1982). With the observation of LTP in the hippocampus (Bliss & Lomo, 1973), the Morris water maze became the standard protocol for studying the role of LTP in memory consolidation *in vivo*. In particular, infusion of NMDA receptor antagonists have been observed to block LTP *in vivo* and lead to a marked deficit in Morris water maze performance (Davis et al., 1992). A limitation of pharmacological manipulations is the placement of reagent may be inconsistent from subject to subject and the spread of the reagent is generally not well controlled. In order to address these problems, the gene for the NMDA receptor was knocked-out in a hippocampal subregion specific manner (McHugh et al., 1996; Tsien et al., 1996; Nakazawa et al., 2003; McHugh et al., 2007). Whether the NMDAR knockout was in CA1, CA3, or DG, all mice showed a marked deficit in Morris water maze performance suggesting that LTP in the hippocampus is necessary for normal memory consolidation.

As we discussed earlier in the introduction, the sequential activation of place cells within a θ-cycle may be crucial for storing the sequence through mechanisms of STDP. The ability to encode place by θ-phase suggests the spatial-temporal interactions between the θ-rhythm and place cell activity is robust (O'Keefe & Recce, 1993). Therefore, a manipulation to the HPC which leads to a degradation of the temporal-spatial properties of place cell activity should theoretically result in the impairment of spatial and episodic memory encoding and recall. This line of reasoning is supported by experiments where genetically modified mice lacking the NMDA receptor in area CA1, CA3, or DG, express large and imprecise CA1 spatial receptive fields (McHugh et al., 1996; Tsien et al., 1996; Nakazawa et al., 2001; McHugh et al., 2007). In these mice, the loss of ability to rapidly sharpen CA1 receptive fields with experience is highly correlated with impaired spatial learning. These deficits in spatial learning may be explained by the possibility that "less reliable" place fields in the knock-out mice results in poor repetition of place cell sequence activation over multiple theta cycles. Moreover, the fact that the CA1 place fields are larger in these mutant mice suggests that the output of CA1 will be less sparse due to

higher probability of overlapping fields. When compared to normal mice, the HPC of these mutant mice would likely produce fewer unique place cell sequences thereby diminishing the encoding capacity of the hippocampus. Consequently, the reinforcement of sequence memory traces by STDP will be severely disrupted. In summary, the interactions between the θ-rhythm, excitability, and spike timing are important for inducing STDP mechanisms that underlie the place cell sequence based memory traces.

Subregion specialization à la Marr et al.

Information encoding and memory recall are non-trivial processes that are attributed to the HPC. The underlying mechanisms for these functions are mediated by combining the computations of pattern completion and pattern separation (McNaughton & Morris, 1987; Rolls, 1996). Pattern completion allows for the retrieval of a complete memory trace using only a partial set of cues from the neocortex. When two or more memory traces are triggered by cue information downstream of the HPC, a process called "pattern separation" is initiated so that through a neocortical/HPC conversation only one pattern is selected. Pattern completion and separation need not be mutually exclusive and most likely co-exist temporally to make the HPC as robust as possible. David Marr's "collateral effect" provided the first theoretical framework for understanding how the architecture of the HPC might support pattern completion (Marr, 1971). Marr's theory predicted how a subset of neurons that were once part of a memory trace could be used to reactivate all neurons that constitute that memory trace in its entirety. He applied his theory to the HPC and hypothesized that the hippocampus-entorhinal loop could serve the function of pattern completion. Marr's theory was later applied to the autoassociative network of the CA3 *recurrent collaterals* in a hypothesis which positioned CA3 to be the primary pattern completion engine of the brain (McNaughton & Morris, 1987; Treves & Rolls, 1991). Many papers suggest CA3 is truly the autoassociator it is thought to be (Nakazawa et al., 2002; Moser & Moser, 2003; Lee et al., 2004; Lee et al., 2004; Leutgeb et al., 2004; Gold & Kesner, 2005; Leutgeb et al., 2006). Pattern separation, on the other hand, is a function that is attributed to the DG based on computational theory Marr initially proposed for

cerebellar function (Marr, 1969). The role of the DG would then be to orthogonalize or accentuate the differences of two competing patterns so that they are sufficiently transformed into two non-competing representations downstream. Support for orthogonalization is provided by the few-to-many-to-few projections from the (PP)-to-(granule cell)-to-(CA3 pyramidal cell) in the HPC, much like the (mossy fiber)-to-(granule cell)-to-(purkinje cell) circuit in the cerebellum. As information arrives from the EC into DG through the perforant path, small differences in EC population coding can lead to large differences in DG population coding. Therefore, memory recall under the conceptual frameworks of Marr, is a melding of pattern completion and pattern separation in proportions that are imposed by the availability of cue information, which is highly dependent on environmental conditions and task demands. Area CA1 receives input from CA3 through the Schaffer collateral projections and from EC through the temporal-ammonic pathway (Remondes & Schuman, 2003; Remondes & Schuman, 2004). The convergence of these inputs onto the apical dendrites of CA1 pyramidal cells suggests that CA1 is largely a coincidence detector. Although many experiments have sought to dissociate the functional properties of the HPC subregions within this theoretical framework, specialization of HPC subregions are still not well understood (Gilbert et al., 2001; Guzowski et al., 2004; Lee & Kesner, 2004).

Subregion cooperation

Instead of solely viewing the hippocampus as a computational unit with distinct sub-circuits that perform different functions, one may benefit from considering the possibility that the function of each sub-circuit may not be as unique as the anatomy. With the melding of hippocampal sub-circuits, a more fruitful approach might consider what functional properties are shared between the subregions and infer the function of the hippocampus as a whole based on the common properties of the hippocampal subregions. Therefore, in doing so, we are not forced to prescribe to any particular way of thinking in terms of the functionality of the subregions. We now proceed by examining the

electrophysiological evidence that is common throughout the HPC and extrapolate from there in a bottom-up like manner.

The most obvious trait found throughout the HPC, is the theta rhythm. It is a 6-12 Hz rhythm in the rodent and dominates the frequency spectrum during locomotion, consumatory behaviors, and REM sleep (Vanderwolf, 1969; Buzsáki, 2006). The theta rhythm in the HPC is driven by GABA-ergic input from the medial septum and diagonal band of broca via the fornix fiber tract (Buzsaki et al., 1983; Bland et al., 1999; Buzsaki, 2002). In the HPC, the theta rhythm has been found to be highly coherent between CA1, DG, and EC, however, detailed studies have shown this coherence to be present but not as robust in CA3 (Buzsaki, 2002).

The functional significance of the theta rhythm cannot be overstated. The activity of both excitatory and inhibitory neurons are highly regulated by the phase of the oscillation (Skaggs et al., 1996; Harris et al., 2002; Somogyi & Klausberger, 2005; Ego-Stengel & Wilson, 2007). This regulation is thought to provide three essential functions: 1) global synchronization between HPC and cortical areas (Siapas & Wilson, 1998; Jones & Wilson, 2005; Jones & Wilson, 2005) as well as structures in the basal ganglia (DeCoteau et al., 2007), and 2) a temporal compression mechanism for converting rate based representations into spike-sequence based representations (Skaggs et al., 1996; Mehta et al., 2002), and 3) a mechanism for constraining the temporal relationship of spikes for induction of LTP (Huerta & Lisman, 1995; Holscher et al., 1997; Mehta & Wilson, 2000). The continuum of theta organized activity has recently been shown to extend to grid cells where the theta rhythm has been known to exist (Buzsaki, 2002; Hafting et al., 2008).

The detailed coordination of firing throughout and beyond the HPC strongly suggests that the role of the HPC (and possibly the EC) is to convert unreliable, rate based codes from extra-hippocampal areas into reliable spike-sequence based codes in the HPC. To investigate this hypothesis on a functional level, the rat is given a stereotypes trajectory to run on. As the rat experiences this trajectory over and over again, the same sequences of place cell activations are expressed in the hippocampus. By limiting the experience of the

rat to a small number of trajectories, theoretically, the number of patterns that are then stored in the HPC will be small. Amazingly, these patterns of experience have been found to be "replayed" or "re-expressed" during REM sleep, SWS sleep, and awake-immobility (Louie & Wilson, 2001; Lee & Wilson, 2002; Foster & Wilson, 2006; Ji & Wilson, 2007). These observations support the idea that the HPC is a site which is involved in the storage and recall of short term memory (Morris, 2006), and that continued communication between the HPC and neocortex is necessary for long-term memory consolidation (Squire et al., 2004).

Recently, we have observed a shift in the view of hippocampal subregion specialization. In some respect, it is a return to the model of simple memory proposed by Marr (1971), where a memory system that contained multiple layers of coincidence detection combined with recurrent collateral feedback would provide the memory storage and recall functions expected from the hippocampus. In a report by Leutgeb & Moser (2007), where they recorded from both CA3 and DG during a morph box experiment, they concluded that the observed rate changes in the respective areas resulted in a decorrelation in pair-wise place cell activity within both CA3 and DG. This decorrelation is consistent with the idea that the subregions of the HPC act as coincidence detectors and as coincident input onto any one cell changes, so must the receptive field of the cell. If we extrapolate this idea, then changing the input pattern to the DG or the coincidence detection mechanism in the DG (by manipulation NMDARs) would affect the input patterns to CA1 and CA3. Similarly, changing the input to CA3 or altering the coincidence detection mechanism would affect the input pattern to CA1. Therefore, ultimately, changing the coincidence detection mechanism in any of the hippocampal subregions will affect the timing of action potentials needed to form sequence memory traces.

THE SCIENTIFIC QUESTIONS: HIPPOCAMPAL SUBREGIONS SPECIFIC ROLES IN MEMORY PROCESSING AND UNDERSTANDING RIPPLE ACTIVITY

Problem Statements

This thesis advances the study of two important aspects in hippocampus mediated episodic memory consolidation: 1) how is spatial information processed by the HPC, and 2) how can we better understand the structure of ripple oscillations?

1) Understanding how information is encoded in the HPC is central to understanding its function in mediating memory consolidation. In particular, there has been an interest in determining the role of the HPC in novel spatial encoding (Lee et al., 2005). Curiously, the definition of what is and is not novel can be vague. In some sense, we are always entering a novel situation if past history is taken into account. Furthermore, because the HPC is constitutively encoding spatial information, a major question is: what are the differences in HPC processing during "more novel situations" and "less novel situations"?

Novelty can be defined relative to the animal and relative to the HPC. In the first case, novelty can be defined by experience. With rats, novelty can be introduced by moving the experiment to another room or introducing objects that the rat has never seen into the environment. Novelty can be measured subjectively in rats using the time spent exploring an object or by the initial velocity of exploration in an environment. On the other hand, what is novel to the HPC cannot be controlled by the experimentalist and depends on the perceptions of the animal. There are two confounds related to this particular point. First, even in a very familiar environment, novelty can be induced by a history effect. For example, "I'm in the home cage" has no history, but "I'm in the home cage now and just came from getting chocolate ice cream" has an element of history. Second, even in a novel observation room, there may be familiar visual cues, objects, sounds, and smells that elicit

known inputs into the hippocampus. In such a case, inputs into the HPC from EC may not actually be novel to the HPC.

Taking these confounds of novelty into account, we designed a linear augmentation track task that allows the animal to become familiar with the context (the room) while introducing elements of "novelty" during the task. Novelty in this case is allowing the rat to explore areas of space that were once restricted. During this process of track running, we record multiple single-units from all three subregions of the HPC. We compare the representation of the context, within HPC subregions and across subregions, at different points of the experiment to determine how subregion specific representations of the context changes over time and with additional spatial exploration. From these comparisons, we make a statement about the dissociation and cooperation of HPC subregions in spatial memory encoding.

The process of quantifying the neuronal representations of context is a non-trivial problem. Representations in the brain are always changing and the HPC is no exception. Considering the possibility that the variation in HPC place cell firing may be a reflection of episodic encoding, it is necessary to quantify these variations as accurately as possible. Given a receptive field for a place cell, the variation that we seek to quantify may simply be in the rate to place mapping. A major problem with estimating changes in this mapping is that it requires the rat to move through the space in order to observe the action potentials. This sampling limitation is costly and, therefore, every spike that we observe is extremely valuable. In the thesis, we discuss an adaptive filter for incorporating spiking information that could operate on arbitrarily small timescales and allow the dynamics of contextual representations in the HPC to be tracked continuously.

2) The observation that episodic replay co-occurs with ripple events suggests that the ripple signatures may carry information related to the mechanism of replay generation. What is known is that ripple frequency, ripple power, and ripple duration can vary greatly (Chrobak & Buzsaki, 1996; Csicsvari et al., 1999; Csicsvari et al., 1999). However, what is not well understood is how these variations are produced and whether these variations

are strongly correlated with different computational processes in the HPC. For example, do ripples of different frequency co-occur with forward replay or reverse replay? (Diba & Buzsaki, 2007) Discovering such covariates would allow us to better understand the relationship of replay to behavioral experience and task demands.

In order to make this idea even more concrete, the theta rhythm can be used as an example. Like the ripple, the hippocampal theta rhythm is also an oscillation that can vary in frequency, amplitude, and duration. The frequency of the theta rhythm is meaningful during locomotion because it is highly correlated with the velocity of the animal (Hirase et al., 1999; Wyble et al., 2004). The amplitude of the theta rhythm can be modulated on a task-phase dependent manner (Wyble et al., 2004; Sinnamon, 2006), while the duration of the theta rhythm is highly correlated with the time the rat has been engaged in a particular behavioral state (Buzsáki, 2006). These correlates may be used to decode the state of the animal and predict the activity in areas that are anatomically connected to the HPC.

Unlike the theta rhythm, variations in ripple properties are not well-understood. This is largely due to the use of signal processing algorithms that are restricted in spectral-temporal resolution (Gillis et al., 2005; O'Neill et al., 2006). One third of this thesis is devoted to developing adaptive spectral-temporal filters for quantifying fine time-scale structure of local field potentials. Another third of the thesis is an application of the derived filter to the problem of ripple characterization and detection.

THE ANALYTICAL PROBLEM: QUANTIFYING PHYSIOLOGICAL PROCESSES ON ARBITRARY TIMESCALES

Adaptive decoding of spike trains

In a report by Bialek *et al*, they ask how much information can the spiking of a single neuron convey (Bialek et al., 1991). To address this question, the authors put forth a general paradigm called "stimulus reconstruction". The idea is that if the spikes of a neuron can be used to reconstruct the driving stimulus then it can be said that the spike train contains information about the stimulus space. In procedural terms, "the problem of reading the neural code is essentially the problem of building a (generally nonlinear) filter that operates continuously on the spike to produce a real-time estimate of the unknown stimulus waveform."

In the analysis of ensemble single-unit data from the rat hippocampus by Zhang *et al*, the authors compared Bayesian and population vector methods for position reconstruction (Zhang et al., 1998). Accurate reconstruction of position involves two aspects. First, if the spike train concurrently contains information about other physical quantities such as running velocity of theta phase modulation, then these quantities must be explicitly modeled in the decoding filter. Second, the spike train must contain position information. They found that stimulus reconstruction is best accomplished using Bayesian state-space methods when the assumption of Poisson spiking holds and independent spiking is true.

What is not explicitly discussed by Zhang *et al* is the process of stimulus reconstruction is really a two-step process that consists of an encoding stage and then a decoding stage. The encoding stage can be though of as a model selection and parameter estimation procedure, while the decoding stage can be thought of as model validation procedure. The model

which yields the lowest stimulus construction error would be the model that best accounted for the information contained in the spike train.

Significant advancements in the encoding-decoding framework have been made by the Brown Lab (Brown et al., 1998; Barbieri et al., 2000; Frank et al., 2000; Barbieri et al., 2001; Barbieri et al., 2001; Brown et al., 2001; Barbieri et al., 2002; Brown et al., 2002; Frank et al., 2002; Eden et al., 2004; Ergun et al., 2007). These reports focus on Bayesian state-space approaches to spike train decoding and model selection, and put forth general frameworks for modeling and validating time-dependent reconstruction algorithms that may be non-Poisson in spike timing and non-Gaussian in stimulus space.

For the investigation of HPC subregion dissociation and cooperation during episodic information encoding, we would like to track HPC representations of context as they change over time. Using an adaptive point process framework developed in the Brown Lab (Brown et al., 2001; Frank et al., 2002), in the thesis we track the receptive field of each cell over time. When the dynamic properties of each cell are pooled together according to HPC subregion, statements can be then be made about the relative responses of each subregion to 1) contextual familiarization, 2) novel spatial exploration, 3) and general time-dependent experience in the same environment. This approach allowed us to quantify consistent but slow changes in ensemble rate coding of the HPC.

Adaptive filters for characterization of brain rhythms

Whether recorded outside the skull (electroencephalogram, EEG) or inside the brain (local field potential, LFP), the frequency content of the brains electrical activity in the brain has been found to be crucial for the assessment of functional states, diseases, and injuries. Furthermore, many studies have shown that different areas of the brain will exhibit different patterns of rhythmic activity, and that each rhythm is highly dependent on age, mental state, stimuli, pharmacological agents, and disturbances to/from the body (Buzsáki, 2006).

Despite all that we know about brain rhythms from their generation to function, we still face many complex and unanswered questions. The role of adaptive filters could be used to help answer these questions by allowing for the fine timescale structure of brain rhythms to be computed and correlated with the activation of specific neuronal circuits. A general process for performing these measurements is to 1) select the pieces of data for analyses, 2) extract the relevant features from the data, and 3) combine the measurements to, for example, perform some classifications or make some statistical comparisons. The component that is most relevant to this thesis is step #2, the extraction of relevant features.

The methods used today fall largely into two classes: parametric and nonparametric. Generally, the difference is parametric methods are model driven, whereas, nonparametric methods are transform driven (Percival & Walden, 1993). The advantage of using parametric methods is they offer better resolution and better characterization when the underlying assumptions of the model are correct. In contrast, the advantage of using non-parametric methods is very few to zero assumptions are needed. The process of choosing one method over another largely depends on what measurements are needed for addressing the scientific question at hand.

On the most basic level, brain rhythms are characterized by modulation in power in the respective frequency bands. For this purpose, nonparametric FFT and wavelet based methods for computing power spectra are sufficient. However, in many cases the desired quantity is not easily derived from the power spectrum. For example, one would have difficulty measuring the number of peak frequencies, the location of peak frequencies, the bandwidths of brain rhythms, and relative power of brain rhythms. Moreover, with the dynamic nature of brain rhythms there is no guarantee that such measurements derived from the power spectra will be accurate. Furthermore, placing statistical boundaries on the accuracy of such measurements are extremely difficult. One way to overcome these limitations is to use parametric methods.

As I discussed earlier, the use of parametric methods requires assumptions to be made in order to justify the model. First, EEG and LFP can vary in their composition; at some

times one rhythm may be present and at other times the signal may be a superposition of multiple rhythms. Second, each rhythm can be bounded by a low and high frequency boundary. Third, within the frequency boundary, a rhythm may rapidly change in amplitude, bandwidth, center frequency of oscillation, and phase delay. In other words, each rhythm may be highly nonstationary over very short timescales (milliseconds).

A parametric method that can quantify data following these assumptions is the autoregressive (AR) model. We focus on the AR because it has a long and fruitful history as a flexible tool for spectral and time-series analysis (Zetterberg, 1977; Box et al., 1994; Oppenheim et al., 1999; Lutkepohl, 2005). Interestingly, for all that it can do, the AR has a very simple form: the univariate AR, $y_n = \sum_{p=1}^{P} a_p y_{n-p} + \varepsilon_n$, is parameterized by a model order p, a measurement error variance term σ_ε^2, and a coefficient vector term $A_p = (a_1,...,a_p)$. If the coefficients A_p are chosen correctly and ε_n is a white noise process, the AR generates a time series y_n as a superposition of $p/2$ oscillations when p is even. The peak frequency and bandwidth of each of these oscillations is also parameterized within A_p. So, with the proper definition of A_p and the ability to generate ε_n, one could generate data with similar properties to that of EEG and LFP data. More importantly, there now exists many methods for solving the inverse problem: to estimate $(p,\sigma_\varepsilon^2,A_p)$ from the data. From these parameter estimates, we can directly compute the end products of our spectral analyses: the number of rhythms in the EEG/LFP, the peak frequency of each rhythm, the bandwidth of each rhythm, and the relative power between rhythms (Oppenheim et al., 1999). Most importantly, these end products are computed from an AR model of the oscillation and not extracted from the power spectrum.

Many of the advancements made in EEG analysis, were made in the 1960's and 1970's. Zetterberg was one of the first to apply AR models to the analysis of EEG data (Zetterberg, 1969). In his 1969 report, Zetterberg acknowledges that his maximum-likelihood estimator is only appropriate for stationary parts of EEG signals. Not long after, Bohlin combined the AR linear model with the Kalman filter for adaptive estimation of

AR parameters (Kalman, 1960; Bohlin, 1977). Till this day, the same constructs proposed by Bohlin provide the basis for continued advancements in EEG/LFP analyses. For example, Foffani *et al* used partial fraction expansion to decompose each oscillation of the AR into independent analytical functions (Foffani et al., 2004); in doing so, time-varying properties of each brain rhythm could be tracked dynamically over time and with fine-temporal precision. Later, the Kalman smoother was applied to univariate and multivariate versions of AR models of EEG/LFP signals (Arnold et al., 1998; Tarvainen et al., 2004).

The concept of coherence, which receives much attention today, had been published early on by Gersch in his consideration of multidimensional EEG/LFP recordings (Gersch, 1970). Gersch used the analytic form of the auto-spectra and cross-spectra, which can easily be derived from the AR model, to analytically solve for the coherence function between two LFP signals. In parallel to EEG research, Granger used AR models in his development of Granger causality to better formulate problems in econometrics (Granger, 1969). Since the introduction of Granger's framework, every manipulation of the multivariate AR transfer function has been proposed to be a measure of causality or connectivity in the brain (Rosenberg et al., 1998; Sameshima & Baccala, 1999; Baccala & Sameshima, 2001; Hesse et al., 2003; Kus et al., 2004; Ivanov & Kilian, 2005; Chen et al., 2006; Kus et al., 2006; Schelter et al., 2006). Although these proposals do measure something, it is unclear how they can be used to make scientific inferences. It seems that more work is required to make clear the proper interpretation of proposed causality measures.

THESIS SUMMARY

Chapter 2: We seek to understand the relative contributions of the HPC subregions in the coding of episodic memory. The chapter presents *in vivo* single unit place cell data collected from 4 rats over 10 days of experimentation in an augmented linear track task, which is designed to examine the differential response of the hippocampus in "more novel" and "less novel" situations. The experiment, data, and analyses presented here are novel to the hippocampal field. The augmented linear track task as we have defined it has never been published (to the best of our knowledge). Moreover, simultaneous recordings of neurons in CA1, CA3, and DG are rarely seen in the literature. We combine, in a novel manner, an adaptive filter for measuring place cell plasticity with basic statistical analysis to the relative contributions of the HPC subregions in spatial and episodic memory processing.

Chapter 3: We examine the signal processing question of how to accurately quantify the frequency modulation and amplitude modulation of brain rhythms. We derive, implement, and characterize a Kalman filter for tracking instantaneous frequency and frequency modulation of brain rhythms based on an autoregressive model. We consider the practical aspects of adaptive filtering and ultimately find that minimizing the adaptive filter provides the most reliable measurement of frequency dynamics. We demonstrate how the filter can be applied to any brain rhythm of interest, and provide a simple example looking at θ-rhythm activity on a linear track.

Chapter 4: We apply the filter developed in Chapter 3 to the analysis of hippocampal ripple events. We demonstrate how to choose filter parameters for analyzing the ripple signal and how to interpret the frequency modulation signal in the ripple band. We build on this knowledge to address three problems in the analysis of ripples: the classification of ripples, the detection of ripples, and the deconvolution of overlapping ripple events. We propose a framework for ripple classification that is independent of ripple amplitude (a quantity that is variable with electrode position and electrode properties). We also

demonstrate how the filter may aid in selective detection of ripple events with particular frequency and amplitude properties. We feel that this framework is the first of many that will help to advance our knowledge of ripple-associated, hippocampus mediated, mechanisms for memory processing.

Chapter 5: We recap the problems faced in the thesis, how we overcame the problems, and the results.

Chapter 2 : HIPPOCAMPAL SUBREGION DISSOCIATION AND COOPERATION: AN EXAMINATION OF EXPERIENCE DEPENDENT RATE DYNAMICS

INTRODUCTION

An important and unanswered question in the study of spatial memory processing is what unique processing occurs in the hippocampus to allow for the encoding and consolidation of spatial experiences. This question can be broken down into two lines of though, one line asking do the anatomically identified subregions of the hippocampus perform some unique function which contributes to the role of the hippocampus as a whole, or do the subregions perform similar functions but are placed in different stages of processing, or a combination of both?

As we discussed in the introduction, a context is defined by elements in the environment that can be perceived. Furthermore, the act of perceiving, being aware of oneself with respect to past history and with respect to the objects in the environment, heavily influences place cell activity in the hippocampus (Jeffery, 2007). Given a stable configuration of environmental cues, an open question is how does general experience affect rate coding of the hippocampus? In particular, how does repetitive spatial experience modulate the rate, shape, and location of receptive fields in the subregions of the hippocampus? Are some subregions more sensitive to initial experiences in a context? Are some regions more sensitive to relative spatial novelty, such as the act of exploring a once restricted area of space?

In this chapter, we explore these questions using an augmented linear track task while recording *in vivo* from the principal cell layers of CA1, CA3, and dentate gyrus in awake, behaving rats. The data used for analysis were recorded after each rat was exposed to the same room for at least 4 days, and the cues used for orientation and room identification were held constant throughout every experimental session (over multiple days). We address the questions posed above using adaptive filtering techniques for tracking receptive field plasticity. Our results indicate that when a rat is reintroduced to a familiar environment, the process of re-familiarization is accompanied by markedly higher rate modulation in CA3 is over that of CA1 and DG. This difference in rate modulation is

pronounced only at the "home" end of the linear track. In addition, the difference in rate modulation between the subregions was reduced with experience such that by the end of the session (~45 minutes), the difference between subregions was no longer significant. The final rate modulation was non-zero throughout the length of the track apparatus, suggesting that with many laps back and forth, the hippocampus was still engaged in a process that alters the rate-to-space representation in the hippocampus. These findings suggest that CA3 has a unique role in the processing rapid contextual coding and provide evidence for a constitutively active process of rate modulation that engages all three subregions.

MATERIALS AND METHODS

Surgery and Recordings

Surgical procedures: Four male rats between 3 months to 6 months of age were used in the collection of data for this analysis. Each rat was placed in the sleep box between two and three hours, every day, for four or five consecutive days. After the rat was comfortable resting in the sleep box, a 20 tetrode hyperdrive was implanted and affixed to the head of the rat.

All surgical procedures met NIH guidelines and were approved by the Committee on Animal Care at MIT. The rat was anesthetized with a mixture of ketamine (0.5 cc, 100mg/ml), xylazine (0.3 cc, 20mg/kg), alchohol (0.3 cc, 200 grade), atropine (0.1 cc, 0.4mg/ml), and saline (1.0 cc). Once the rat was fully induced under anesthesia, the rat was head fixed to a stereotaxic device and was given 1% isoflurane for the remainder of the surgery. Bone screws (7 to 9) were inserted around the perimeter of the skull as anchors. A craniotomy and durotamy of 1.5mm in diameter was centered on the coordinates -3.5 mm from bregma and 2.5 mm to the right of midline. The multi-electrode array was held in place over the craniotomy and dental acrylic was used to fix the array to the skull and bone screws. A ground wire was attached to the ground pin on the connector board of the drive array and the most posterior bone screw.

Tetrode construction: Tetrodes were constructed according to (Gray et al., 1995). Polyimide coated nichrome wire (7 micron core diameter, 15 micron total diameter, Kanthal Precision Wire, Palm Coast, FL) was pulled to 50 cm in length. The wire was folded twice and twisted under constant force such that the four wires remained taught in a spiral. The wire was twisted clockwise 70 turns and then counter clockwise for 30 turns. A heated-air gun set to 500 degrees F was used to fuse the polyimide coating of the twisted wires. On one end, the wires were separated and the insulation was removed by applying a flame for a brief moment. On the other end, the tetrode wire bundle was carefully cut so that all four wires remained fused and the exposed metal tips were roughly in the same

plane. In order to increase the signal detection ability of the electrodes, electrode impedance was reduced by electroplating gold particles onto the exposed nichrome tips. The electroplating was performed by dipping the tetrode and a conductive wire into a conductive gold solution. A current was passed along the conductive wire, the gold solution, and the length of the electrode such that the conductive wire acted like an anode while the electrode wire acted like a cathode. The current was set to 3 uA and was applied in brief pulses, each 1 second long. After each pulse, the impedance of the electrode tip was measured at 1kHz. The current pulses were applied until the impedance fell between 300 and 500 kΩ. In order to aid the process of histology, each tetrode was coated with a dialkylcarbocyanine probe (Invitrogen), either DiI (red) or DiO (green), so that cells around the tetrode would contain fluorophores. The probes were applied by mixing an alcohol solution containing DiI or DiO and drawing it into a syringe. The syringe is gently pushed so a bead of solution formed at the tip of the syringe. The bead of alcohol was run back and forth over the tetrode until the color of the solution could clearly be seen as a coating over the tetrode. The two die colors were chosen to be spatially alternating (Figure 2.1).

Figure 2.1 : Tetrode Drive Construction. (A) 3-D model of tetrode drive (B) Actual 20-tetrode drive (C) Layout of tetrodes on the bottom cannula.

A total of 20 tetrodes were constructed and held in place by a hyperdrive containing 20 linear actuators. The tetrodes are connected to a linear actuator and then piped into a master cannula that was approximately 1mm in diameter. On the non-gold plated end, each of the four wires in the tetrode were individually connected to pins on a Millmax connector (14x2 configuration, 1mm pitch). The 6 tetrodes that are most medial were targeted to the granule cell layer of the upper blade of the dentate gyrus, the 6 tetrodes that are more lateral were targeted to the CA3c pyramidal cell layer, while the 6 tetrodes remaining were targeted to the CA1 pyramidal cell layer. The two reference tetrodes were targeted to the white matter tract (alveus/cingulum) between the dorsal hippocampus and cortex.

Data collection: Three types of data were digitally acquired during the experiments: rat head position, action potential waveforms, and local field potentials (LFPs) by custom data acquisition software (AD, Wilson and Frank, 1999). Infrared diodes were affixed to the multi-electrode drive array on the rat's head. The diodes were clustered into two groups: left and right.   Cameras in the ceiling were able to image these infrared diodes moving

through space at a rate 60 frames per second. The left and right diode clusters were interlaced in time such that only one would be on at any given time. Interlacing the diode frames allows for straight forward estimation of head direction and head velocity. Each frame from the camera was fed through a threshold detector, which sent the coordinates of the threshold-crossing pixels to a computer for storage. Each pixel was time-stamped using a clock running at 20 kHz. A total of 16 LFP channels were recorded simultaneously. Each LFP channel could be chosen from any of the 18x4 + 2 = 74 individual electrode channels. Each LFP signal was amplified by a gain of 1000, bandpass filtered between 1 and 475 Hz, and was digitized at 1 kHz with 12 bit precision. Each tetrode channel was pre-filtered using a bandpass filter with cutoffs of 300Hz and 6kHz. A threshold value was assigned to each channel for spike detection. If the peak amplitude of a spike crossed any of the respective thresholds, the spike waveform was recorded on all four channels. The spike waveforms were sampled at 32kHz, with 6 samples before the threshold crossing and 26 samples after the threshold crossing.

Histology: After the protocol was completed, the rat was euthanized using pentobarbital and perfused with a formalin solution. The brain was dissected, sectioned in the coronal plane with 50 μm thickness, and cover slipped. A stereoscopic, fluorescent microscope was used to visualize the electrode tracks at 10X magnification. Each brain slice was imaged three times using dark field, red filter, and green filter; the three images were combined using the Photoshop overlay function (Figure 2.2).

Data obtained over 10 days of recording between 4 animals where used in this analysis (Table 2.1). In one animal only DG activity was recorded. In another animal only CA1 activity was recorded. In two animals, activity from CA1, CA3, and DG were simultaneously recorded. The data used in this analysis was collected on days were the linear augmentation task used all three track segments. By this time, each animal had experience for at least four days.

Figure 2.2 : Coronal section of the rat dorsal hippcampus, right hemisphere, -3.5mm from bregma. The four partial tracks (DiI = orange, DiO = green) were imaged under a fluorescent microscope and overlaid over a darkfield image (10x). Above the granule cell layer of the DG, an electrolytic lesion clearly demarcates the location of the tetrode tip.

Area \	Aardvark	Camel	Emu	Foxxy	total
CA1	1 : 29	0	2 : 27	6 : 117	173
CA3	0	0	2 : 7	6 : 34	41
DG	0	1 : 14	2 : 14	6 : 36	64

Table 2.1 : Summary of isolated units from all areas of the hippocampus for four subjects including principal and interneurons. For each table entry with two numbers, the format is "number of recording days : number of cells over those days". The last column is a sum showing the total number of cells collected over all the animals and days.

Experimental Paradigm
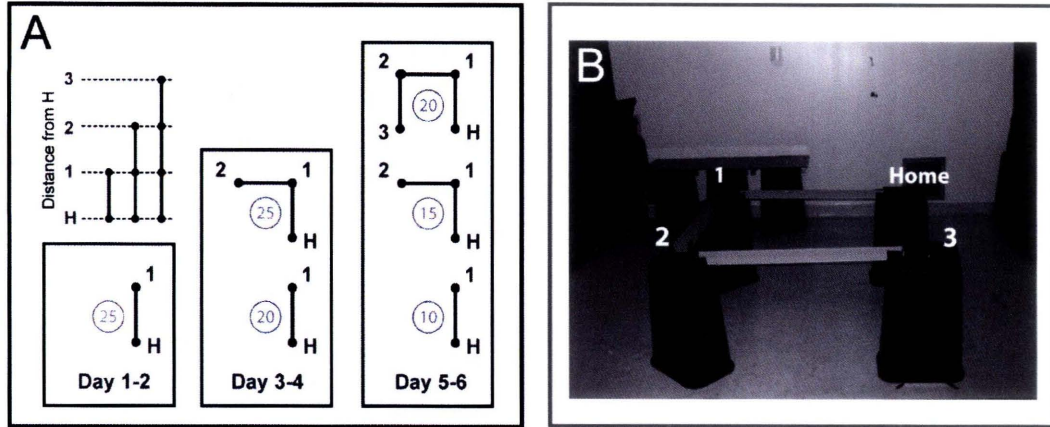
Figure 2.3 : **(A) Configuration of track augmentation experiment. Every day begins with a single segment, and depending on the day, the track may be extended up to a full U-shape. The panel in the top left is a schematic showing the conversion of the track from 2-D to 1-D coordinates. (B) A photograph of the experiment room while the track is fully extended. When the track is not fully extended, the bins are always present and the unavailable segments are placed on the ground below their locations in the photo.**

After implantation of the hyperdrive, rats were given two days of rest before adjustment of electrode position. Over the course of one to two weeks, electrodes were slowly lowered to their respective targets. In the meantime, the rats were placed on a food restricted diet. After the tetrodes were lowered and showed stable units, experimentation and data collection commenced.

Behavioral paradigm: The augmented linear track task allows for the observation of hippocampal activity over conditions of increased familiarity in a room and relatively novel spatial exploration. On the first day or recording, the rat was completely naïve and had never experienced the track before. The rats explored the track cautiously for the first 10 to 15 minutes on the first day. Normally within two days of training, the rats were comfortable running on the track. For the first two days of experimentation, the rat was exposed to only one out of three track segments (Figure 2.3). Track segments are always introduced in sequential order on each recording day. The rule of thumb was new track

segments are introduced only when the rat had at least two days of experience on each track segment.

We explicitly made certain that all cues and track manipulations were always visible to the animal. New track segments were added by lifting the new segment from the ground, removing the food well that was there, connecting the new segment at 90 degree angle, and adding a new food well to the other end of the new segment. Also, the track was not immediately surrounded by a curtain, so the room was completely open to the animal. For orientation, large objects of contrast were placed on each wall of the room. These objects remained constant throughout all days of experimentation. The track was held 60 cm above the ground. The track was 2 inches wide with 1 inch high walls. The reward locations were detachable from the track and were constructed from black foam board. Over the course of experimentation, we attempted to keep, as consistent as possible, the cues in the room, the starting location of the rat at the beginning of the experiment, and the orientation of the maze relative to the room and the cues.

Data Analysis

Clustering: The spike waveforms for each tetrode were reduced to a set of features: maximum amplitude, and peak-to-trough time. On computers running Linux, manual clustering software, Xclust2 by Matthew Wilson, was used to isolate individual units using spike features. The times of spikes for each unit were written to file. Matlab software (Natick, MA) was used to analyze spike train and position data.

Linearization and segmentation of track: The representation of rat position was converted from 2-D to 1-D using a simple distance transformation (Figure 2.3). All locations of the rat during the experiment were orthogonally projected onto a line which traced the path of the full 3-segment U-shaped track. After projecting the points onto the line, the position of the animal was represented as the distance from the end of the line representing the home location. For each day of experimentation, the 1-D representation

of the track was divided into seven sections. The spatial divisions mark the boundary between reward/corner locations and straight track segments. The purpose of these divisions is to allow comparisons to be made over time and space using all the data collected during the experiment (Figure 2.4).

Adaptive estimation of receptive field plasticity: We track receptive field plasticity using an adaptive point process filter (Frank et al., 2002; Frank et al., 2006). Briefly, the filter models the spatial receptive field using the Catmull-Rom spline, which is a simple cubic spline that is entirely defined by the location $(x_1, ..., x_K)$ and height $(\theta_1, ..., \theta_K)$ of K control points evenly spread over the linear track (Bartels et al., 1987). The formulation of the spline-based receptive field is

$$r(u) = \begin{bmatrix} u^3 & u^2 & u & 1 \end{bmatrix} \begin{bmatrix} -0.5 & 1.5 & -1.5 & 0.5 \\ 1 & -2.5 & 2 & -0.5 \\ -0.5 & 0 & 0.5 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \theta_{k-1} \\ \theta_k \\ \theta_{k+1} \\ \theta_{k+2} \end{bmatrix}$$

where for a given 1-D position x, $u = (x - x_k)/(x_{k+1} - x_k)$, for $x_k \le x \le x_{k+1}$. The objective of the filter is to update θ given past history and current information. This is accomplished using a steepest decent maximum likelihood filter (Brown et al., 2001) :

$$\hat{\theta}(n) \approx \hat{\theta}(n-1) - \varepsilon \left[\frac{\partial \lambda(n \mid H, \theta, x)/\partial \theta}{\lambda(n \mid H, \theta, x)} \left(N(n-1, n] - \lambda(n \mid H, \theta, x) \Delta \right) \right]_{\substack{H = H(n-1) \\ \theta = \hat{\theta}(n-1) \\ x = x(n)}}$$

where,

$$\frac{\partial \lambda(n \,|\, H, \theta, x)}{\partial \theta} = \begin{bmatrix} -0.5u^3 + u^2 - 0.5u \\ 1.5u^3 - 2.5u^2 + 1 \\ -1.5u^3 + 2u^2 + 0.5u \\ 0.5u^3 - 0.5u^2 \end{bmatrix},$$

$$u = \frac{x - x_k}{x_{k+1} - x_k}, \quad \text{for } x_k \le x \le x_{k+1},$$

the conditional intensity function, $\lambda(n \,|\, H, \theta, x)$, is defined to be the receptive field function, $\lambda(n \,|\, H, \theta, x) \triangleq r(x)$; the discrete time interval is Δ; and the counting function $N(n-1, n]$ indicates the number of spikes that occurred in the discrete time interval $(n-1, n]$.

The parameters used throughout the analyses were: the first control point was set to 0 cm, the last control point was set to the distance of the fully extended track, the control points were equally spaced, the number of control points was 35, the timestep was $\Delta t = 1/15$ sec, and the learning rate was set to $\varepsilon = 5$ which allowed for a maximum rate change of 3 Hz/sec.

For each neuron, the spike train was split into two spike trains based on the direction of running. Direction information was discarded at the ends of the track, such that action potentials that occurred in goal locations at the track ends were assigned to both spike trains. These two spike trains were then divided again based on track configuration, such that a total of 6 spike trains were obtained. The tracking of spatial receptive field shape was performed individually on each of these 6 spike trains. An example of the adaptive filter output can be found in Figure 2.4.
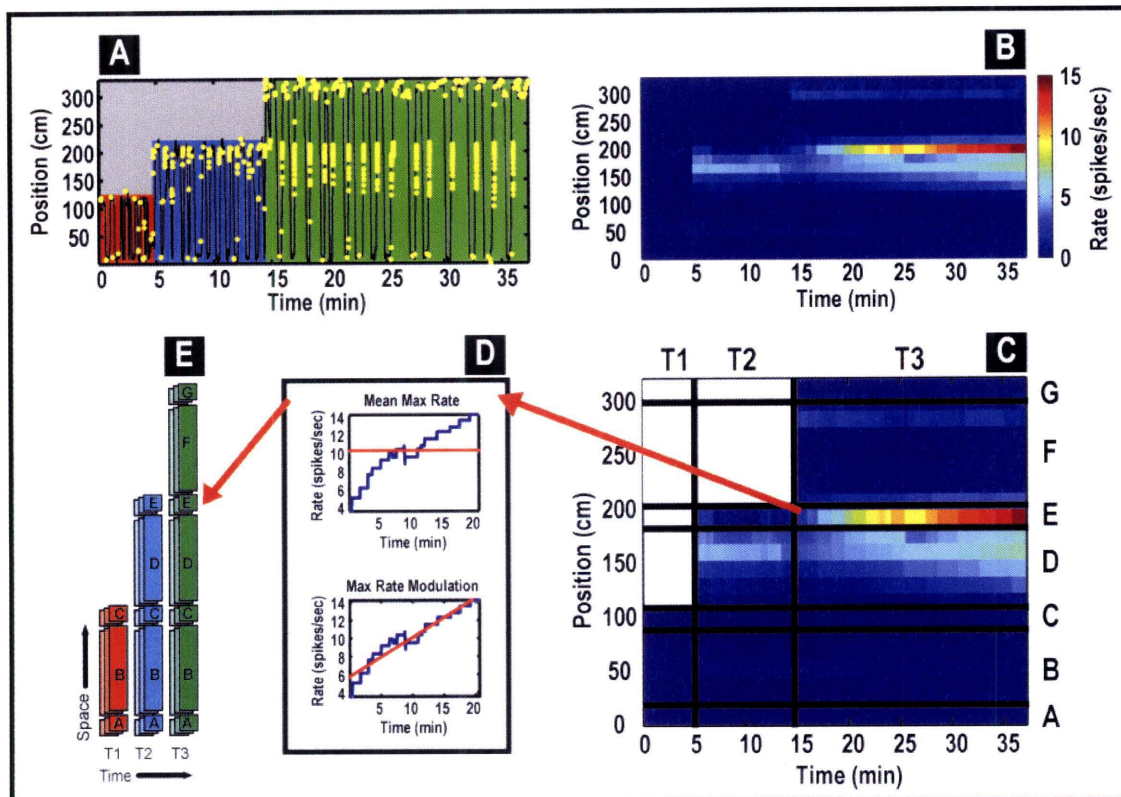
Figure 2.4 : Measuring rate dynamics over space and time. (A) An example of place cell activity. This neuron was recorded from the dentate gyrus and only spikes recorded during movement from bottom-to-top were displayed. (B) Adaptive representation of rate map over space and time for spiking activity displayed in panel A. (C) The experiment is segmented according to track configuration and different spatial segments. The spatial segments distinguish between areas of goals/corners and linear areas. (D) Extraction of rate parameters. The maximum value of the rate map is extracted from the receptive field surface within the space-time boundaries shown in Panel C at a time interval of 1 sec per sample. From the maximum rate trajectory, mean maximum rate and maximum rate modulation values are computed as shown. (E) Formation of rate templates. These two rate measures are computed for each combination of track configuration and spatial segment. Each layer in Panel E corresponds to a single cell. For each cell, the activity was divided 30 times: 15 combinations of space and track configuration x 2 directions of running. The stacking of the layers forms a template or population vector of activity for each track configuration-spatial segment condition.

Measurement of mean maximum rate and maximum rate modulation: The dynamics of receptive fields were summarized by characterizing the temporal evolution of the maximum rate-to-place mapping. The point process adaptive filter described above was used to track the receptive field dynamics. Regardless of the number of track segments available to the rat, the range of the spline spanned the fully extended track length using 35 control points. Next, each adaptive spline was sectioned according to Figure 2.4 and the maximum rate was extracted within the respective spatial boundaries at the time resolution of 1 second. A linear regression analysis was then performed on each maximum rate trajectory and from this a slope was extracted. In addition, the average value of the maximum rate was also computed. From this point we will refer to the slope of the maximum rate trajectory as the "rate modulation", the absolute value of the slope of the maximum rate trajectory as the "absolute rate modulation", and will refer to the average maximum rate as the "mean rate." These three measures were computed for each neuron for all combinations of track configuration, spatial segment, and running direction. The activity of the neurons were then pooled according to recording location (CA1, CA3, or DG) to produce population vector templates. Statistical measures were used to summarize the activity described in each population vector. In particular, the average value was computed for each template. We refer to population statistics as average mean rate, average rate modulation, and average absolute rate modulation.

Population Vector (Template) Comparisons: The U-shaped track was linearized and segmented into seven sections (Figure 2.4), such that locations of food reward and corners were considered separately from linear segments of the track. Either mean rate or absolute rate modulation for each neuron was computed in each track segment, for both directions of running, and for each track configuration. If a neuron had a firing rate below 2 Hz for all conditions, it was not considered for analysis. Interneurons were identified by setting a threshold of 5 Hz mean firing rate and removed from the pool of cells used for analysis. Table 2.1 summarizes the data used in the analysis.

We measure relative differences in average maximum firing rate or average maximum rate modulation by choosing a reference population vector (R) and another population vector (T), and performing two analyses: 1) computing Pearson's correlation coefficient, and 2) hypothesis testing using a bootstrap test on the population mean.

For the **correlation analysis**, R and T were chosen to represent the same spatial extent on the track, same direction, same anatomical recording location, but different track configurations (Figure 2.11). Pearson's correlation coefficient was computed as follows:

$$\rho_{R,T} = \frac{\text{cov}(R,T)}{\sigma_R \sigma_T} \approx \frac{\left(\frac{1}{n-1}\right)\left[\sum_i (R_i - \bar{R})(T_i - \bar{T})\right]}{\left[\frac{1}{n-1}\sum_i (R_i - \bar{R})^2\right]^{1/2}\left[\frac{1}{n-1}\sum_i (T_i - \bar{T})^2\right]^{1/2}}.$$

The **bootstrap analysis** tested for significant differences in mean rate and absolute rate modulation statistics. We performed hypothesis testing where R and T were chosen from different anatomical regions, but the spatial extent, direction of motion, and tracking configuration were the same. We also performed hypothesis testing where R and T were chosen from different track configurations, but the spatial extent, direction of motion, and anatomical recording location were the same. The bootstrap test used 10,000 iterations with population sizes of 25. A t-test performed on the sample means were used to either accept or reject the null hypothesis that R and T are similar.

RESULTS

We analyzed ensemble spike trains recorded from principal neurons of the CA1, CA3, and dentate gyrus (DG) subregions of the hippocampus in four male, adolescent rats (3 to 6 months in age) during an augmented linear track task. A total of 10 days of recording between 4 animals yielding 171, 38, and 62 principal neurons for CA1, CA3, and DG, respectively, were used in this analysis. In all the days of experimentation used in this analysis, the rats had at least four days of prior experience on the track..

Population average of mean rate is not different between subregions

Previous studies have reported large changes in the receptive field properties of HPC place cells in response to manipulations of salient cues (Muller & Kubie, 1987). These changes in receptive field properties are known as "rate remapping" or "global remapping," however, we use the term "rate modulation" to remain agnostic about the meaning of any changes in rate to place mapping. With the addition of new track segments, we examined the possibility that an individual HPC subregion might have a heightened response to the observation of relatively novel spatial experience. If this were true, we might observe a subregion specific bias in rate modulation resulting in a relatively significant increase or decrease in overall rate.

In order to explore this possibility, we computed the mean rate for each neuron as a function of track configuration and location on the track (see Methods). The mean rate for each neuron was combined according to recording location to create three sets of vectors. Table 2.3 contains the statistics of these mean rate vectors. In order to aid in the interpretation of these statistics, the data in the table is plotted in two different ways. Figure 2.5 displays the relationship between average mean rate between subregions across track configurations, and across spatial locations. We found differences in average mean rate between subregions, however, the differences did not generalize. In particular, the

average mean rate for CA3 was lower than that of CA1 and DG in spatial segment D and E but not F and G. This inconsistency neither helps to prove or disprove a unique role for CA3 during relatively novel spatial exploration. Overall, the data is consistent with previous observations that average mean rate will be high in locations of high velocity (linear segments) versus locations of low velocity (corners and reward locations) (Czurko et al., 1999).
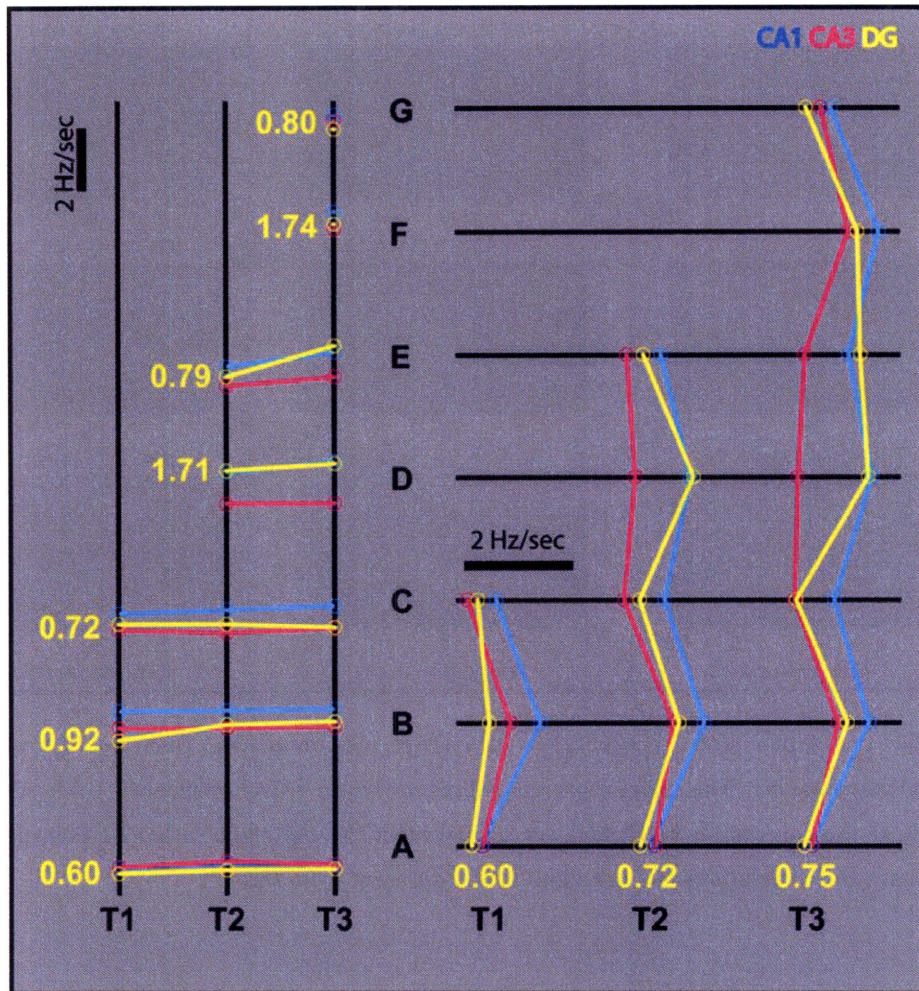
Figure 2.5 : Subregion specific dynamics of average mean rate (average of the mean maximum rate over a cell population). The data used to generate this figure is given in Table 2.3. The yellow labels mark the average mean rate in the CA1 subregion for the given condition. The data is plotted to allow for visual comparison across track configurations and within track configurations. T1, T2, and T3 mark the different track configurations, while A through G mark the different spatial segments of the track. Please refer to Figure 2.4 for more clarification.
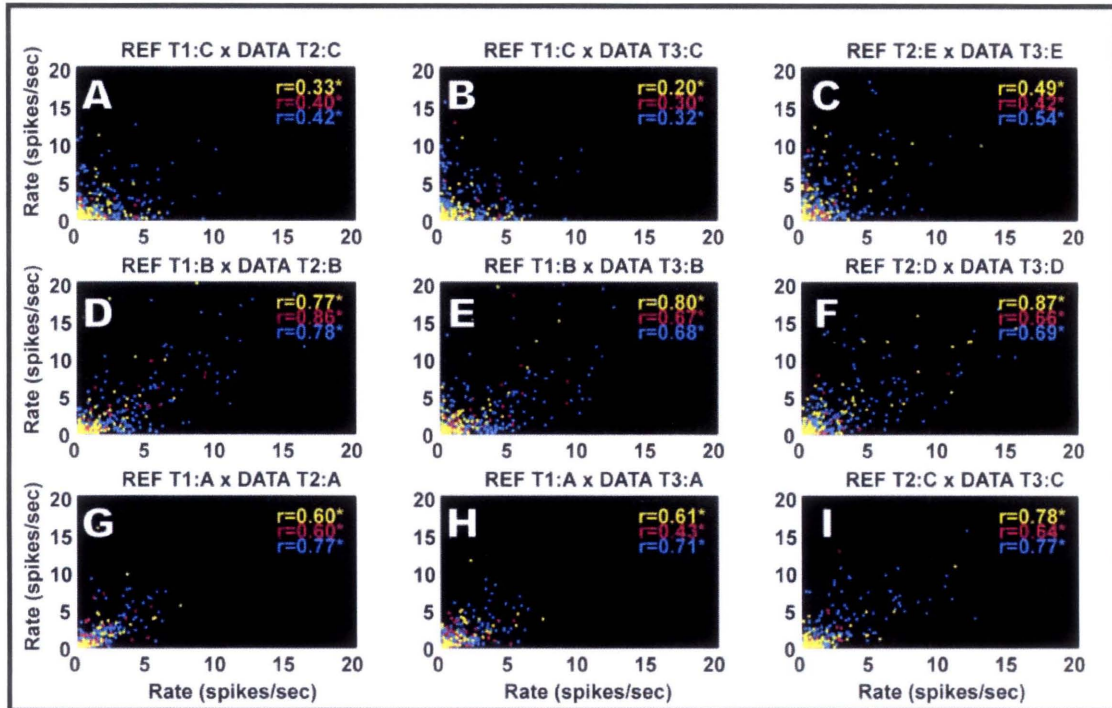
Figure 2.6 : Population vector scatter plots of average maximum rate. Refer to Figure 2.11 for overview of comparisons. The colors represent subregions: cyan = CA1, magenta = CA3, yellow = DG. The values of the correlation coefficient are displayed in the top right corner of each panel. The asterisks mark the correlations that are significantly different from zero.

Dynamics of subregion specific population coding

The function of pattern separation has been attributed to CA3 (Guzowski et al., 2004; Leutgeb & Leutgeb, 2007; McHugh et al., 2007), and more recently, experimental evidence supports the idea that pattern separation in the HPC is mediated by the cooperation of CA3 and DG through a process of coincidence detection (Leutgeb et al., 2007). Unlike these experiments, in our experiment the salient cues used for orienting in the room are stable throughout the course of the experiment. Therefore, we hypothesize

that with the addition of track segments, any changes in the mapping between salient cues, which define the context, and place field rate and location will most likely be gradual.

In order to examine the nature of the gradual changes in rate to place coding in the HPC, we compared population vector templates of average mean rate between track configurations. Figure 2.11 provides a schematic of this comparison and Figure 2.6 summarizes the results. Each analysis compares previous activity in a track section with later activity in the same track section by computing a correlation coefficient (r-value). All r-values were significant at reward and non-reward locations (Figure 2.13). For comparison between subregions, a shuffle analysis was used to determine the variability of each correlation value given a limited number of neurons (Figure 2.14).

Rate changes in the three subregions were most apparent when a reward location was converted to a corner on the track (T1:C x T2:C; T2:E x T3:E). This observation is mostly likely due to a manipulation of task phase and running velocity at the corresponding track corners (Figure 2.17), (Wyble et al., 2004; Sinnamon, 2006). In track sections where the behavior of the animal was not altered by additional track segments, changes in ensemble rate were also observed but to a lesser extent (Figure 2.6). The analysis revealed that the r-value for each subregion was statistically different in the comparison of T1:A x T3:A (Panel H). Furthermore, in the scenario T1:A x T3:A, the correlation analysis revealed that ensemble coding of the home location was least altered in DG, followed by CA1, and most altered in CA3.

It is unlikely that these differences were caused by overall increases/decreases in excitability in the subregions as a function of track configuration (Figure 2.5). Moreover, the modulation of population rate coding per subregion was similar over temporal and spatial experience except for at the home location (section A).

Figure 2.7 : **Subregion specific dynamics in average absolute rate modulation as a function of track configuration and space. The data used to generate this figure is given in Table 2.4. Asterisks mark significant decreases in rate change from one track configuration to the next (p < 0.05). T1, T2, and T3 mark the different track configurations, while A through G mark the different spatial sections of the track. Please refer to Figure 2.11 for more clarification.**

Dynamics of rate modulation as a function of track configuration and relative spatial experience

In the previous section, we examined the relative dynamics of subregion rate coding by comparing population vector activity occurring on the order of tens of minutes. With the adaptive point process filter, we observed that rate modulation is non-zero over time and the modulation varied in both direction and magnitude as a function of track configuration (Figure 2.11 and Figure 2.12). Here, we compare subregion specific average absolute rate modulation across track configurations to assess the relative change in rate dynamics within and between HPC subregions.

The average absolute rate modulation is summarized in Table 2.4 and displayed in Figure 2.7. In general, the average absolute rate modulation tended to decrease over the course of the experiment. The only condition where a significant subregion difference could be statistically confirmed was at the beginning of the experiment and in the home goal location; there, rate changes in CA3 where 65% greater than in CA1 and 175% greater than in DG.

Following the template analysis in Figure 2.11, population vectors of average absolute rate modulation were constructed and compared (Figure 2.8). Statistical analyses were computed within subregions and across track configurations. Not all correlations from the template correlation analyses were significant (Figure 2.15). We examined the variability in the r-value due to sample error by systematically varying the number of neurons used in the correlation, and randomly choosing which cells to compare (Figure 2.16, no identity shuffle). The shuffle analysis revealed that r-values were only significantly different between all three subregions under one condition, T1:A x T2:A. The corresponding r-values are CA1 = 0.32, CA3 = -0.76, DG = 0.10. Interestingly, the r-value for CA3 is large and negative, whereas, CA1 and DG r-values are low and positive.

We further assessed the tendency for absolute slope to decrease with the addition of track segments. This was accomplished by comparing the distributions of absolute rate slope

across subsequent track configurations (Figure 2.7). Right-sided empirical bootstrap tests of sample mean set to a power of $\alpha = 0.05$ showed significant decreases between T1 and T2 for spatial bins A, B, and C; and between T2 and T3 for D and E, for CA3 only (p < 0.05). For CA1, decreases in absolute slope from T1 to T2 were statistically significant for spatial bins A, B, and C; and between T2 and T3 significant decrease was observed in spatial bin C. For DG, significant decreases were only found in the transition from T1 to T2 for reward-related sites A and C.
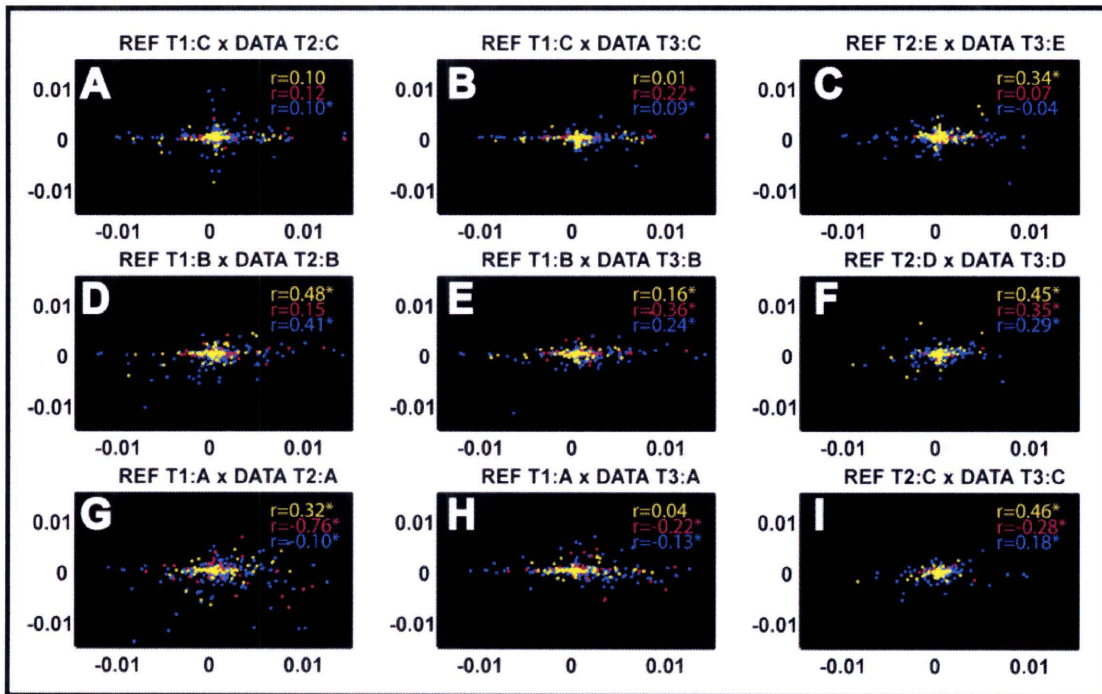
Figure 2.8 : Population vector scatter plots of maximum rate slope. Refer to Figure 2.11 for overview of comparisons. The colors represent subregions: cyan = CA1, magenta = CA3, yellow = DG.

Subregion coherence in rate modulation

Some features of ensemble rate dynamics were similar across subregions. When comparing the change in absolute rate slope between the one segment track and two

segment track configuration, all subregions showed a decrease in the amount of rate modulation, however, subsequent track manipulations did not elicit a significant and coherent decrease in absolute rate slope across all subregions (Figure 2.7). Moreover, on the three segment track configuration the ensemble rate modulation for all track locations and all subregions where not statistically different (Figure 2.7). In addition, significant differences were not found in the average firing rate at the goal location for all three track configurations (Figure 2.5).

DISCUSSION

Rapid contextual learning in CA3

We investigated relative changes in ensemble rate coding and ensemble rate dynamics between the subregions of the HPC and found unique ensemble rate dynamics in CA3 that were only present at the beginning of the experiment. In particular, within the first track configuration (5 to 10 minutes), at the home location on the track, area CA3 showed significantly greater modulations in maximum firing rate relative to CA1 and DG (Figure 2.7). Focusing on the activity in the home location at the beginning of the experiment, we found that a majority of neurons in CA3 had a preference to increase their firing rate while CA1 and DG had no preference for direction of rate modulation (Figure 2.12). Between the track configuration with one segment and the configuration with two segments, the template correlation analysis using average absolute rate modulation showed a significant negative correlation for the CA3 ensemble (r=-0.76, Figure 2.8), indicating an overall reversal in the direction of rate modulation. The differences between HPC subfields were found in the rate dynamics and not overall rate itself (Figure 2.7).

The rate dynamics observed at the goal location were probably not caused by reward related activity. In our experiments, the amount of reward did not change over the course of the experiments and the rats maintained a constant desire for reward throughout the

experiment. This desire was indicated by the constant or increasing running velocity with each extension of the track (Figure 2.17). It is possible that an internal perceptual shift of reward valence could induce rate modulations in the HPC. However, a study by Tabuchi *et al* (2003) demonstrated that changes in reward valence did not result in changes in the response of the neurons at the reward locations (Tabuchi et al., 2003). Also, reward related modulations of rate would likely be present on both ends of the track, however, we only found pronounced CA3 specific rate modulation at the home location.

In our analyses of CA3 activity, we find the rapid modulation of ensemble rates to be consistent with previous reports of rapid encoding in the CA3 subregion by NMDA receptors. Work by Nakazawa *et al* have shown that NMDA receptors (NR1) in the CA3 area are essential for rapid acquisition of novel spatial learning (Nakazawa et al., 2002; Moser & Moser, 2003; Nakazawa et al., 2003). Later, Cravens *et al* showed how the same CA3-NR1 deficient mice were able to perform contextual discrimination when tested 24 hours after contextual fear conditioning but failed if tested after 3 hours (Cravens et al., 2006). They conclude that plasticity in CA3 is essential for the rapid encoding and recall of spatial and contextual memory. Furthermore, in vitro studies have shown that NMDAR-dependent plasticity in CA3 is mossy fiber independent and recurrent collateral dependent, as well as perforant path dependent (Zalutsky & Nicoll, 1990; Nakazawa et al., 2002; Nicoll & Schmitz, 2005). These findings suggest the enhanced modulation of rate coding in CA3 that we observe in the initial stage of the experiment may reflect the induction of NMDAR-mediated learning via the perforant path. Furthermore, this NMDAR-mediated rate modulation may play an active role in the formation and recall of rapid spatial and non-spatial episodic (one trial) learning.

The distinct enhancement of CA3 rate modulation over CA1 and DG may be due, in large part, to the unique anatomy and physiology of the CA3 area. Recurrent associational and commissural connections are inputs that provide strong excitatory drive and also have the ability to induce post-synaptic LTP (Zalutsky & Nicoll, 1990). Another source of input is the sparse but powerful mossy fiber connections which are the axons of granule cells of the DG. A single mossy fiber-CA3 synapse contains many post-synaptic densities that are

completely engulfed by the presynaptic terminal (Nicoll & Schmitz, 2005). Each CA3 pyramidal cell receives input from only 50 granule cells on average, however, the strength of one mossy fiber synapse is such that one presynaptic spike is enough to elicit a post-synaptic spike (Jonas et al., 1993; Henze et al., 2002). Another unique physiological characteristic of CA3 pyramidal cells is marked frequency facilitation in the mossy fiber pathway; increasing the presynaptic spike frequency from 0.5 Hz to as little as 1 Hz can cause many-fold increases in synaptic strength (Nicoll & Schmitz, 2005). These physiological mechanisms for enhancing the integration of sparse inputs may serve to amplify changes from DG and EC input, and produce relatively larger changes in spatial rate maps at the beginning of a new experience.

Functionally, enhanced CA3 rate modulation during early experience in a context may allow the HPC to stabilize the hippocampal representation of context specific memory traces. A key question here is what does it mean to stabilize a response? We can treat CA3 as a computational unit that transforms an input pattern into an output pattern. The fact that CA3 is a circuit with powerful feedback connections implies that several cycles of feedback will be needed before stable output can be reached (Rolls & Kesner, 2006). However, once the output is stable, we can assume that plasticity in CA3 will allow the transformation to be retained and therefore allow CA3 to respond more reliably next time the same input pattern is presented. Now assume the input pattern into CA3 is unique for every spatial location and head direction. Combined with the plasticity assumption, more spatial exploration would then lead to better optimization and more robust responses from CA3 inside the context. One functional consequence of this type of optimization is with faster and more reliable CA3 output, reinforcement of synapses downstream will be more consistent. Another functional consequence is once the transformation has been optimized for input patterns afforded by the present context, CA3 would contain within itself, a memory trace of the contextual experience that is represented by a set of optimal transformation functions.

This line of reasoning is supported by several experimental results. In a study where place fields recorded in light and darkness were compared (Markus et al., 1994), Markus *et al.*

found that on a new day of experimentation, after the animal had many days of exposure to the same maze, the correlation between light and dark ensemble CA1 activity was greater when the animal was allowed to explore the environment before the lights were turned off. Markus *et al.* concluded that even with multiple days of exposure, initial exploration specific to that day was an important factor in stabilizing place field location and reliability in darkness. A study by McHugh *et al* found that in contextual fear conditioning of mice, spatial exploration and not just visual accessibility of a space, was necessary for storage of the fear memory trace (McHugh & Tonegawa, 2007). In a study by Wilson & McNaughton (1993), they found that CA1 activity became more reliable with continued exploration in a novel environment (Wilson & McNaughton, 1993). Although this activity was recorded in CA1, increasing CA3 spike time reliability would cause spiking in CA1 to become more reliable. In addition, Wilson & McNaughton observed that exploration on the novel portion of the environment did not significantly effect place field representations on the familiar half of the environment (Wilson & McNaughton, 1993).

Constitutively active pattern separation

The convergence of average absolute rate modulation between the subregions occurs over the course of the experiment. This observation suggests that initially, CA3 plays a unique role in information processing (which we have discussed above), but as CA3 transformations become optimized with spatial exploration, rate modulation of place specific neuronal activity decreases (Figure 2.7 and Figure 2.8). Interestingly, subregion specific rate modulation converges to a point where rate slopes per subregion are not statistically different in the home goal location or any other part of the track. Furthermore, rate modulation in all subregions converges to a similar, non-zero value in all areas of the track suggesting that a baseline level of plasticity may be present throughout the spatial extent of the context.

It is unclear, however, whether this baseline rate modulation in the HPC is 1) due to changes in input patterns originating outside of the HPC or 2) a result of intrinsic HPC computations. The first hypothesis suggests that constitutive HPC rate modulation may be a result of constant neocortical integration of history and continual cognitive processing. The second hypothesis suggests the synaptic weights which provide a substrate for HPC state, by virtue of the CA3 autoassociation network or EC-HPC loop, can never be constant given the same neocortical input. One could reconcile these hypotheses by extending the time of the experiment past one hour. If the first hypothesis is correct, we would expect to see the rate of modulation fall to zero as the effect of neuromodulatory mechanisms eventually turn off and/or as the salience of repetitive episodes on the track become insignificant. For the second hypothesis, we would expect to see persistent and coherent rate modulation for a theoretically limitless amount of time provided that spatial exploration is constant throughout the period.

If the mechanism for baseline rate modulation is intrinsic to the hippocampus, this mechanism may be an automatic way of incorporating time information into a conjunctive representation of event sequences. The baseline modulation we observe may reflect an ongoing process of pattern separation that would provide new neuronal codes for the same path trajectories that are separated by time. In this way, the HPC may likely serve as constitutively active "history integrator" that constantly updates current representations based on structured activity. In cooperation with the cortical "path integrator circuit" (Taube, 2007), the HPC/neocortical loop may provide a general purpose framework for efficiently and automatically encoding episodic information.

Conclusion

In conclusion, our results indicate an active role for CA3 in the initial stages of contextual familiarization. In this scenario, novel input patterns into CA3 are defined to be input patterns which elicit an unreliable output pattern from CA3. With continued exposure to a

novel input pattern, CA3 quickly finds a stable transformation and rapidly encodes this transformation such that on the next exposure the output of CA3 will be precise and reliable. Because the information storage capacity of CA3 is limited, transformations that have been learned for one context may be erased or interfered with when the animal has been in another context for a long time. This constraint on CA3's ability to store transformations is offset by its ability to rapidly make new associations. With respect to the present experiment, enhanced CA3 rate modulation is observed during the first 5 minutes of the experiment but steadily declines over the course of the next 30-40 minutes. The convergence in rate modulation between the three subregions suggests a coherent super-circuit which engages plasticity mechanisms in CA1, CA3, and DG on a constitutive basis. Such a role for plasticity is strongly supported by experimental results from the Tonegawa lab, where NMDAR specific lesions of plasticity in either CA1, CA3, or DG results in the expression of unreliable and imprecise spatial tuning of CA1 place cell activity (McHugh et al., 1996; Nakazawa et al., 2004; McHugh et al., 2007). We hypothesize that, in the subregions of the HPC, convergence of rate modulation to a non-zero value indicates a constitutively active process of pattern separation that may be driven by intrinsic properties of the HPC; and that the observed baseline rate modulation may also reflect a process of "temporal integration" so that similar event sequences that are separated in time are encoded by unique hippocampal activity patterns.

SUPPLEMENTARY INFORMATION

Figure 2.9 : Population rate statistics. A-C) Maximum rate of neurons during running (head velocity > 15 cm/sec). D-F) Maximum rate of neurons using all data. The maximum rate was computed by using a sliding window of 50ms. G-I) Average rate of neuronal activity during running. J-L) Average rate of neuronal ensemble over the whole experiment. The number of cells used for each subregion is 171 for CA1, 38 for CA3, and 62 for DG; each cell contributed two data points, one rate for each head direction. The maximum rate was computed by taking the largest rate computed from non-overlapping time bins of 500 ms.

	# Cells	Max Rate (running)	Max Rate (overall)	Mean Rate (running)	Mean Rate (overall)
CA1	171	19.9 (15.7)	29.46 (16.1)	1.63 (5.0)	1.28 (4.02)
CA3	38	15.34 (7.8)	25.5 (12.46)	0.69 (0.7)	0.68 (0.60)
DG	62	22.27 (21.6)	34.0 (22.06)	1.58 (6.9)	1.18 (5.43)

Table 2.2 : Rate statistics for each subregion of the hippocampus. For each entry the mean is followed by the standard deviation in parentheses.
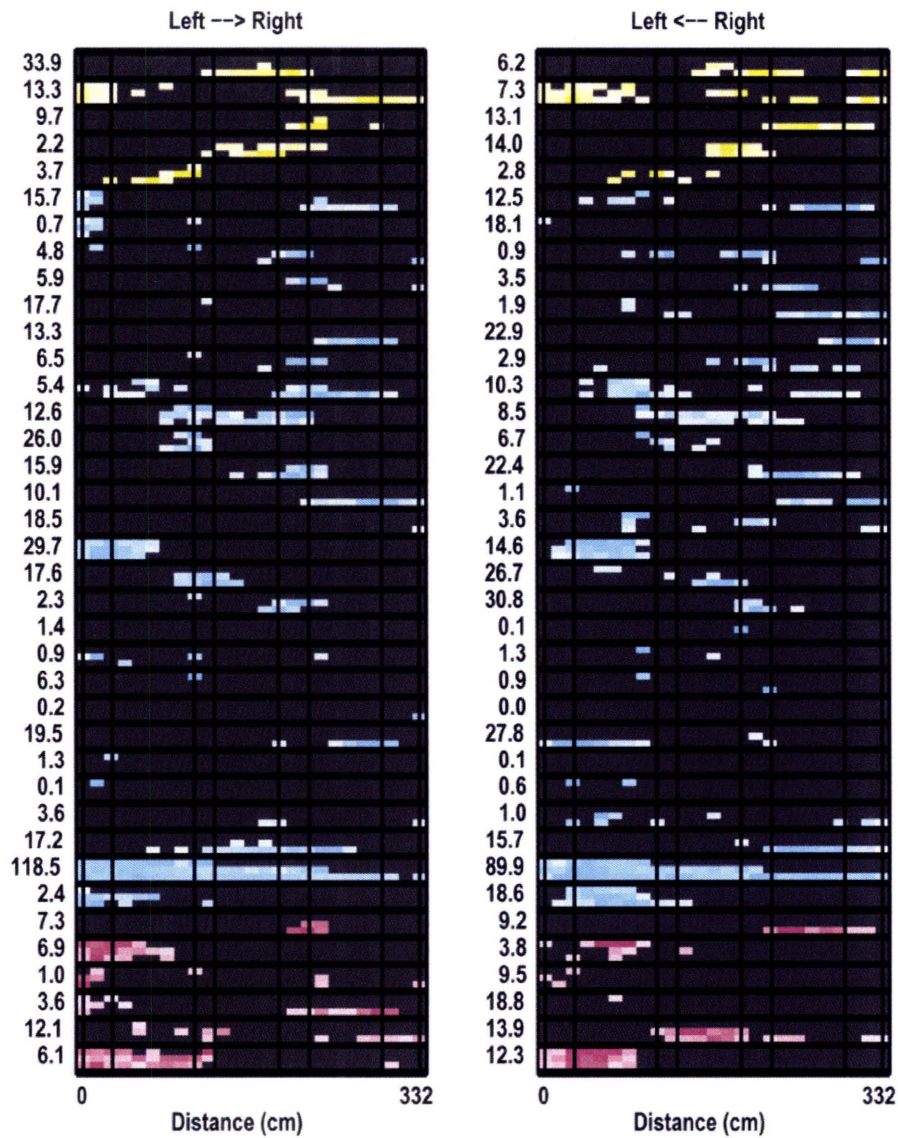
Figure 2.10 : Ensemble activity of simultaneously recorded neurons (yellow = dentate gyrus, blue = ca1, magenta = ca3). Each cell is given one color and is represented by three rows of activity for each direction of running. The three rows correspond to the three track configurations (1,2, and 3 segments). The numbers on the left correspond to the maximum firing rate for that cell over all positions and track configurations.
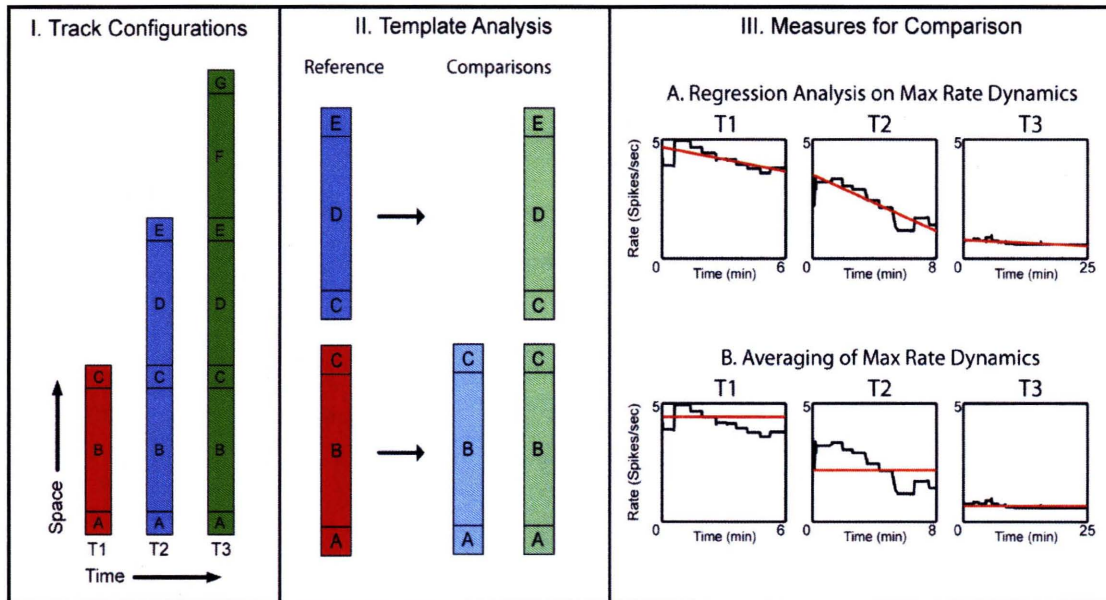
Figure 2.11 : Schematic summarizing data analysis. The Template Analysis compares activity earlier in the experiment with activity later in the experiment. Here, activity is measured in two ways: A) The maximum rate of a neuron is measured using the point process adaptive filter (see Methods), and a linear regression is used to summarize the maximum rate dynamics. B) The average of the maximum firing rate may also serve as a summary of activity in each temporal-spatial bin. These measurements are made for each neuron and then pooled over days and subjects in order to form population vectors. Neurons are pooled into different groups based on recording location (CA1, CA3, or DG). Correlation coefficients are used as a similarity measure between the template and some spatial-temporal epoch of interest (see Methods).

Figure 2.12 : Subregion specific comparison of percent cell population with increasing maximum rate. For each condition, the black line represents the 50% mark. The data used to generate this figure is given in Table 2.5. Asterisks mark significant deviation from the 50% mark (p < 0.05). T1, T2, and T3 mark the different track configurations, while A through G mark the different spatial sections of the track. Please refer to Figure 2.11 for more clarification.
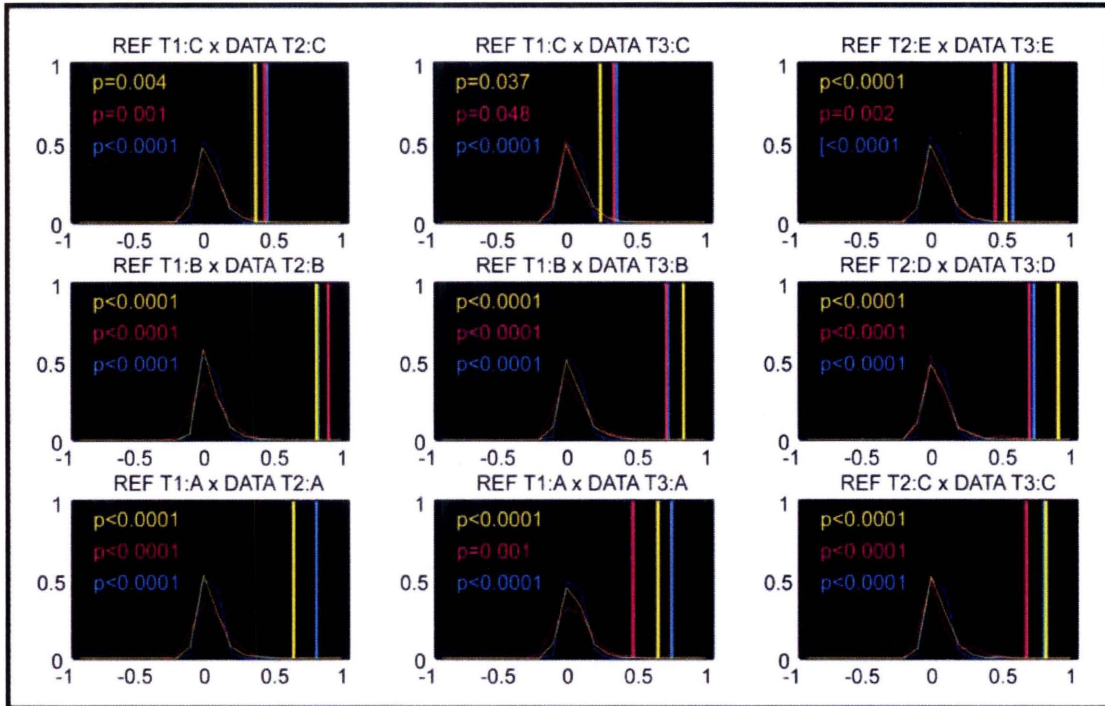
Figure 2.13 : Significance testing for correlation values. For each pair of population vector templates (average maximum rate), the order of one template was shuffled while the order in the other template remained constant, and the correlation value was computed. The distributions for the correlation value were obtained using 1000 shuffles. All correlations were significant (p < 0.05). The p-values are reported in the upper left-hand corner. The colors represent subregions: cyan = CA1, magenta = CA3, yellow = DG.
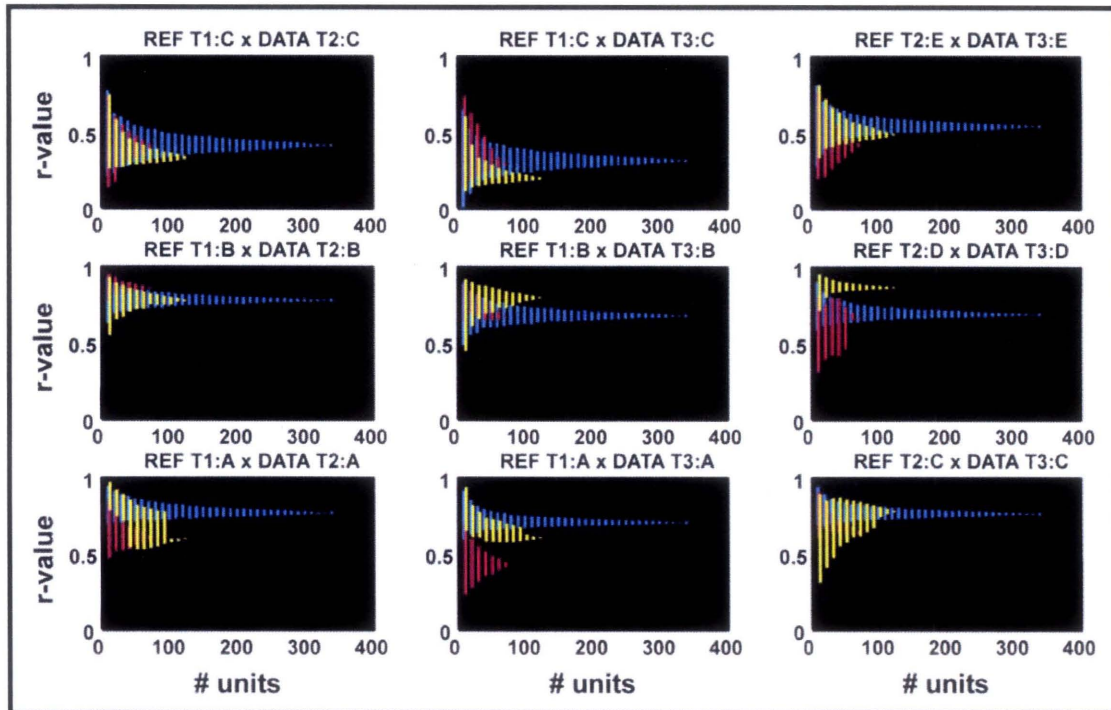
Figure 2.14 : **Correlation analysis using a partial set of neurons (average maximum rate population vectors). Cell identities remained constant between population vector templates, however, a random subset of cells were used to compute the correlation. For each population size, 1000 random samples were drawn (without replacement) from the neuron pool, and the correlation value was computed. Each vertical line represents the 25 to 75 percent confidence interval for the correlation value under each condition. The colors represent subregions: cyan = CA1, magenta = CA3, yellow = DG.**
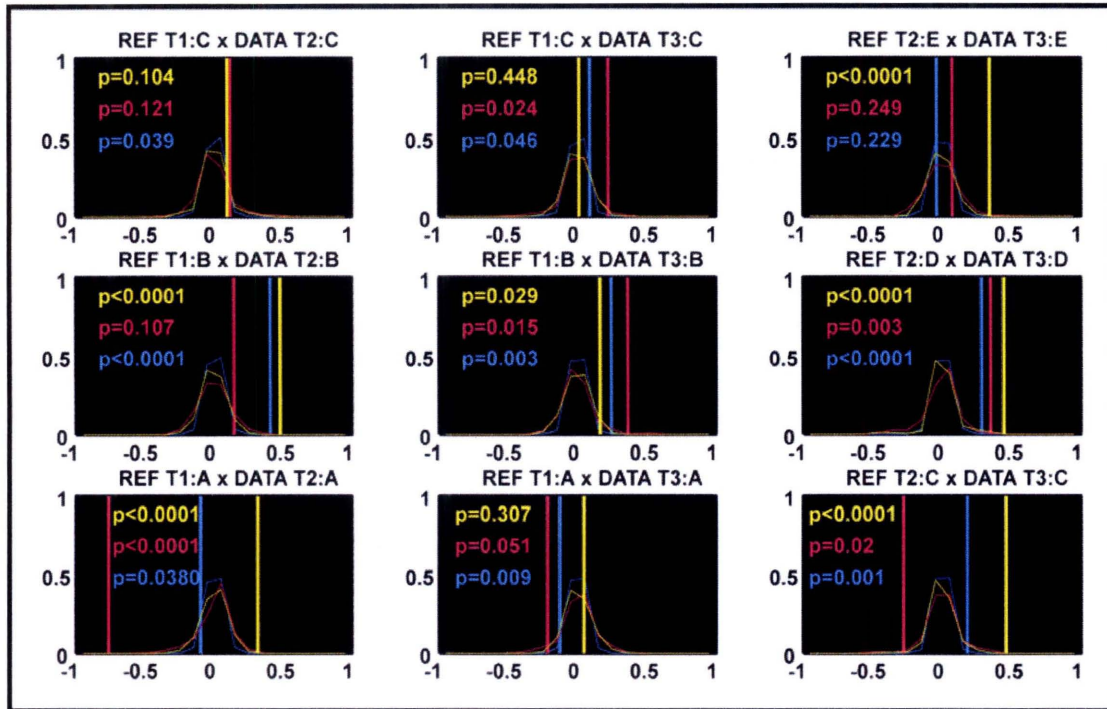
Figure 2.15 : Significance testing for correlation values. For each pair of population vector templates (absolute rate modulation), the order of one template was shuffled while the order in the other template remained constant, and the correlation value was computed. The distributions for the correlation value were obtained using 1000 shuffles. Not all correlations were significant. The p-values are reported in the upper left-hand corner. The colors represent subregions: cyan = CA1, magenta = CA3, yellow = DG.

Figure 2.16 : **Correlation analysis using a partial set of neurons (average rate modulation). Cell identities remained constant between population vector templates, however, a random subset of cells were used to compute the correlation. For each population size, 1000 random samples were drawn (without replacement) from the neuron pool, and the correlation value was computed. Each vertical line represents the 25 to 75 percent confidence interval for the correlation value under each condition. The colors represent subregions: cyan = CA1, magenta = CA3, yellow = DG.**
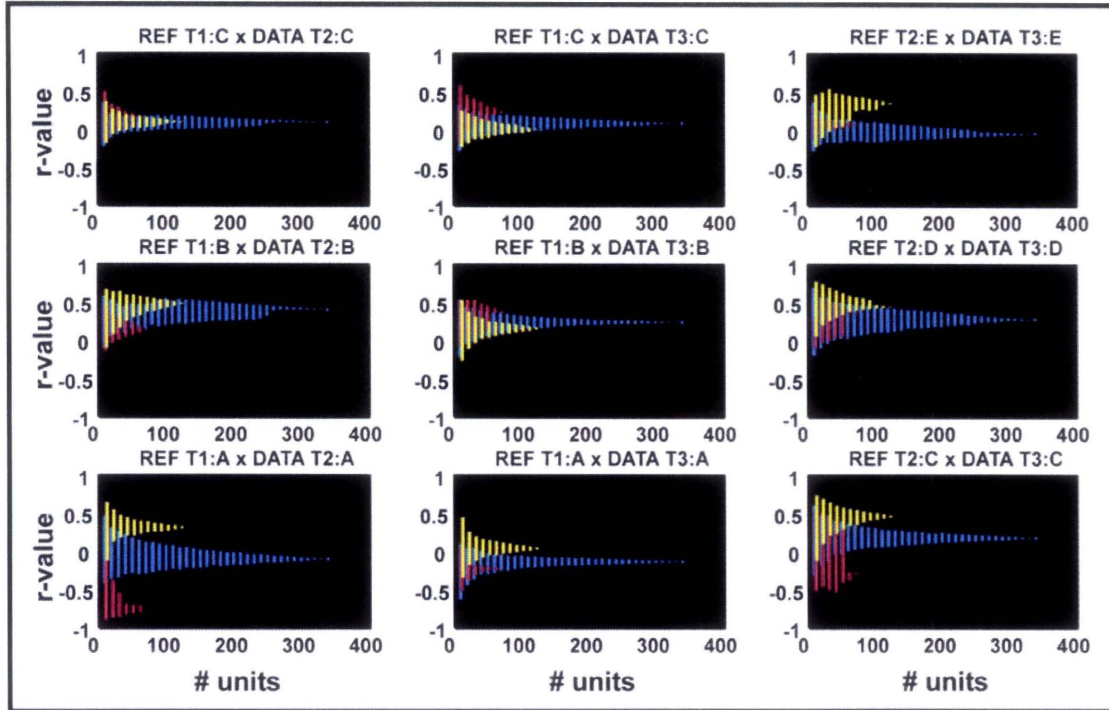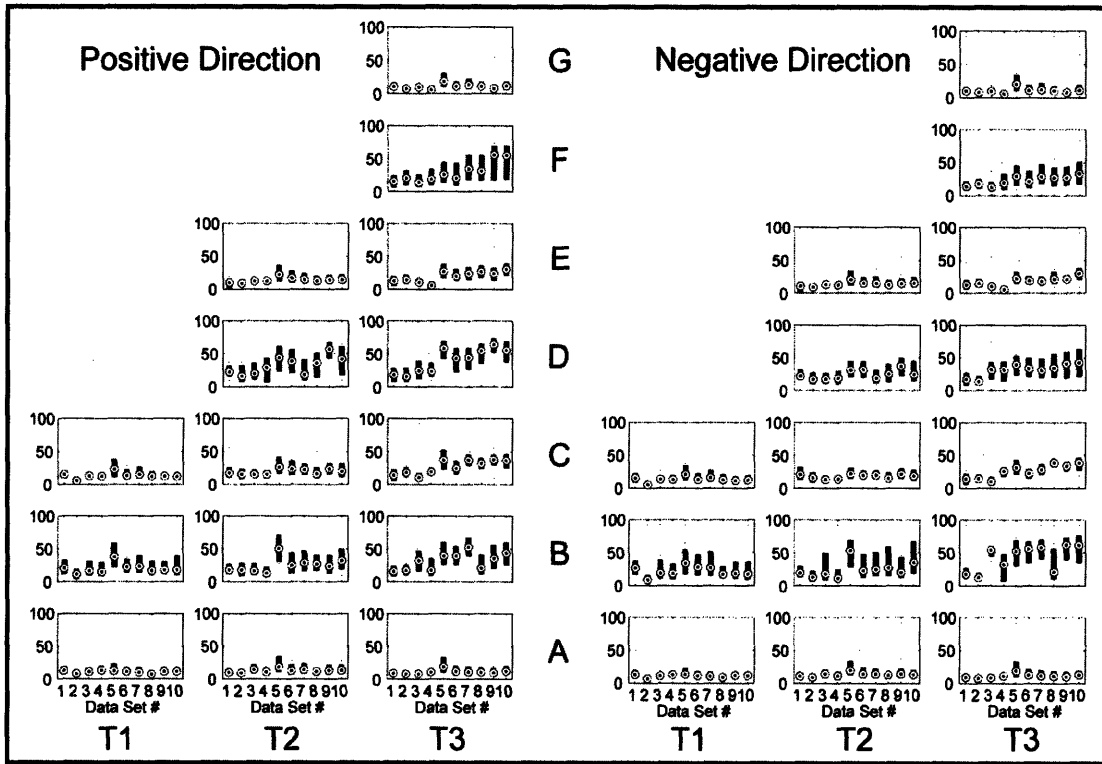
Figure 2.17 : Statistics of velocity according to spatial location and track configuration. The units are in cm/sec.

	T1	T2	T3
G			1.26 (1.92) 1.07 (1.35) 0.80 (1.40)
F			2.14 (3.45) 1.57 (2.23) 1.74 (2.82)
E		1.12 (1.79) 0.48 (0.75) 0.79 (1.71)	1.58 (2.78) 0.77 (1.50) 1.79 (3.32)
D		1.65 (2.67) 0.63 (1.39) 1.71 (3.67)	2.00 (3.35) 0.64 (1.44) 1.94 (3.50)
C	1.06 (1.77) 0.52 (1.01) 0.72 (1.16)	1.18 (2.11) 0.42 (0.75) 0.71 (1.47)	1.30 (2.32) 0.58 (1.67) 0.60 (1.25)
B	1.87 (2.87) 1.32 (1.96) 0.92 (1.67)	1.90 (3.39) 1.33 (2.07) 1.44 (4.60)	1.94 (3.69) 1.37 (2.61) 1.54 (3.60)
A	0.78 (1.24) 0.80 (1.14) 0.60 (1.16)	0.93 (1.58) 0.99 (1.51) 0.72 (1.90)	0.86 (1.45) 0.89 (1.12) 0.75 (1.49)

Table 2.3 : Statistical analysis of average maximum rate. Units are in (spikes/sec). The sample size is 171 for CA1, 38 for CA3, 62 for DG. Each table entry contains the mean and (std) for CA1, CA3, and DG, respectively. A graphical display of this table can be found in Figure 2.5.

	T1	T2	T3
G			0.50 (0.73) 0.66 (0.86) 0.38 (0.75)
F			0.35 (0.50) 0.32 (0.51) 0.43 (0.73)
E		1.08 (1.87) 0.44 (0.80) 0.56 (0.92)	0.47 (0.90) 0.19 (0.28) 0.46 (0.90)
D		0.71 (1.50) 0.30 (0.60) 0.79 (3.13)	0.40 (0.79) 0.15 (0.29) 0.48 (0.97)
C	1.34 (2.77) 0.99 (2.30) 0.98 (1.65)	0.57 (1.24) 0.27 (0.57) 0.39 (0.95)	0.35 (0.72) 0.11 (0.27) 0.27 (0.44)
B	1.40 (2.96) 1.21 (1.96) 0.76 (1.51)	0.64 (1.30) 0.43 (0.62) 0.42 (0.79)	0.38 (0.87) 0.41 (1.07) 0.26 (0.39)
A	1.42 (2.31) 2.35 (5.16) 0.85 (1.60)	0.95 (2.11) 1.25 (2.30) 0.48 (1.08)	0.47 (0.88) 0.66 (1.24) 0.30 (0.55)

Table 2.4 : Statistics analysis of average absolute rate modulation. Units are in (spikes/sec)/(1000 sec). The sample size is 171 for CA1, 38 for CA3, 62 for DG. Each table entry contains the mean and (std) for CA1, CA3, and DG, respectively. A graphical display of this table can be found in Figure 2.7.

	T1	T2	T3
G			54.11 % 51.31 % 49.21 %
F			52.94 % 61.84 % 57.81 %
E		49.41 % 60.52 % 49.21 %	50.29 % 53.94 % 57.03 %
D		53.82 % 51.31 % 51.56 %	56.47 % 60.52 % 45.3 %
C	59.11 % 55.26 % 59.37 %	49.70 % 46.05 % 43.75 %	51.76 % 52.63 % 39.84 %
B	54.11 % 61.84 % 39.84 %	47.35 % 51.31 % 50.00 %	54.41 % 55.26 % 47.65 %
A	52.35% 68.42 % 48.43 %	54.41 % 50.00 % 52.34 %	59.41 % 59.21 % 58.59 %

Table 2.5 : Percent of population with positive rate modulation. The numbers are unit-less. The sample size is 171 for CA1, 38 for CA3, 62 for DG. Each table entry contains the mean and (std) for CA1, CA3, and DG, respectively. A graphical display of this table can be found in Figure 2.12.

Chapter 3 : INSTANTANEOUS ESTIMATION OF AMPLITUDE AND FREQUENCY MODULATION OF LOCAL FIELD POTENTIALS

Abbreviations: AM = amplitude modulation, FM = frequency modulation, EEG = electroencephalogram, LFP = local field potential, AR = autoregressive model, KF = Kalman filter, AD-KF = amplitude demodulated Kalman filter, AIC = Akaike Information Criterion, OIC = Frequency based Information Criterion, PSD = Power Spectral Density

INTRODUCTION

The cellular matrix of the brain is inhomogeneous and complex, yet emerging patterns of regular oscillatory activity are reliably observed; each of these brain rhythms, occurring at frequencies between 0.02 Hz to 500 Hz, reflect the activation of unique neural circuits. It is known that each rhythm can be activated independently or in combination with other brain rhythms depending on behavioral or cognitive state. Brain rhythms can be found in both electroencephalograms (EEG) and local field potentials (LFPs). Each rhythm is specified by a frequency band and denoted by letters from the Greek alphabet $(\delta, \theta, \alpha, \beta, \gamma)$ or referenced by names such as slow oscillations and ultra fast oscillations. One approach to understanding the functional significance of brain rhythms is through the correlation of brain rhythms with behavioral state, cognitive state, experimental variables, and neurophysiology; these correlations combined provide evidence for the role of neural assemblies in brain function (Buzsáki, 2006).

The α-rhythm (8-14 Hz, in humans) is a prominent rhythm that has been associated with higher cognitive processes. Increases in α-amplitude during cognitively demanding tasks engage mental calculations, mental imagery, working memory, and visual attention. The presence of the α-rhythm may signify the long range, horizontal synchronization between cortical areas for the purpose of binding sensory information (Palva & Palva, 2007). The θ-rhythm in rats (6-12 Hz), is clearly dominant and regular in the hippocampus during locomotion and REM sleep, but is absent during other behavioral states such as immobility and slow-wave sleep (Buzsaki, 2002); one possible function for the θ-rhythm is to provide a temporal mechanism for coordinating information transfer, information encoding, and

memory consolidation by the hippocampus (Skaggs et al., 1996; Louie & Wilson, 2001; Harris et al., 2002). The γ-rhythm (30-90 Hz) is prominent in both neocortical and limbic areas, and arises from tight coupling between interneurons and excitatory neurons; the rhythm is most likely multi-functional and is currently believed to support temporal encoding and sensory binding (Bartos et al., 2007). The "ripple" event is a short, fast oscillatory event (50-100ms in duration) that is found in the hippocampus and human cortex (Chrobak & Buzsaki, 1996; Bragin et al., 1999; Clemens et al., 2007). Recently, ripple events have been implicated in memory consolidation as preserved temporal order of hippocampal place cell activity has been shown to occur on a compressed timescale within ripple events (Foster & Wilson, 2006).

Rhythmic activity is modulated in both amplitude and frequency on very fine timescales (Buzsáki, 2006). With the understanding that the brain is dynamic and that traces of important information processing can be observed in intervals as small as 10 msec (Axmacher et al., 2006; Foster & Wilson, 2006), our overall goal here is to develop a robust method for extracting neurophysiologically meaningful structure on a millisecond time basis from EEG and LFP data in order to understand the functional role of specific neural circuits. Since brain rhythms are a manifestation of ensemble neuronal activity, modulations of neural circuits should directly result in modulations in amplitude (AM) and frequency (FM) of oscillatory field potentials. The investigation of AM and FM activity in LFP's is relatively novel and nontrivial, as robust estimators for examining both AM and FM activity have not been developed. Therefore, the focus of this chapter is to characterize brain rhythms through the measurement of *instantaneous amplitude, instantaneous frequency, and frequency modulation* and consider these signals as a utility for investigating neural circuit function.

We derived and applied an adaptive algorithm for accurately tracking instantaneous frequency in EEG/LFP data. The filter models brain rhythms with an autoregressive (AR) model and estimates the AR coefficients using an adaptive Kalman filter framework together with observations of neural activity (Arnold et al., 1998; Tarvainen et al., 2004).

We show that for simple simulated sinusoids, the basic Kalman filter yields a *biased* estimate of instantaneous frequency. With a simple procedure for amplitude demodulation, we show that this bias can be corrected to obtain accurate estimates of frequency dynamics. The performance of the unbiased filter is demonstrated using simulated data. We also demonstrate how AR models with order greater than 2 can give erroneous results and that the accuracy of the frequency tracking can be improved by first band-pass filtering the LFP before estimation with the Kalman filter (KF). We developed and applied a frequency based information criterion for semi-automatic KF parameter selection called the ω-information criterion (OIC). Having outlined a recipe for accurately estimating the instantaneous frequency of brain rhythms, we analyze LFP data recorded from the CA1 subfield of the hippocampus of a rat during a reward foraging task on linear track. We show that the prominent θ-rhythm (6-12 Hz) is frequency modulated on sub-second timescales and report behavioral, task dependent structure in the amplitude and frequency modulation of the θ-rhythm.

MODELING OF A RHYTHMIC PROCESS WITH NON-STATIONARY AMPLITUDE AND FREQUENCY

A sinusoid that is modulated in both amplitude and frequency may be written in discrete time as

$$y(n) = m(n)\sin\big(\phi(n)\big), \qquad (3.1)$$

where $m(n)$ is the amplitude modulation (AM) signal and $\phi(n)$ is the time-varying phase on the unit circle. From simple principles of angular geometry, we know that the relationship between phase, angular velocity $\omega(t)$, and frequency $f(t)$ is

$$\phi'(t) = \omega(t) = 2\pi f(t),$$

$$\int \phi'(t) = \int \omega(t) = 2\pi \int f(t),$$

(3.2)

$$\phi(t) = 2\pi \int_{t=0}^{t} f(t).$$

The phase can be approximated in discrete time by

$$\phi(n) \approx 2\pi \left(\sum_{u=0}^{n} f(u)\Delta \right) + \phi(n = 0).$$

(3.3)

Equations (3.1) and (3.3) provide the most flexible description of an evolving sinusoid. In the upcoming discussion, we will show how to estimate the AM and FM components of (3.1) from observed data.

TIME-VARYING AMPLITUDE OF A SINUSOIDAL PROCESS

The definition of AM used here is the departure of the peaks of the sine wave away from 1 beyond any effects that inadequate sampling and noise may produce. The amplitude modulation signal, $m(n)$, is commonly defined using two forms: coherent or incoherent demodulation. Coherent demodulation is defined by a carrier amplitude and time-varying deviation from the carrier, $m(n) = m_0 + \delta m(n)$. Incoherent demodulation is defined by a time varying amplitude envelope. In the brain, a physiological interpretation for a carrier amplitude signal has not been established; therefore we incoherent demodulation as the physiological interpretation is more straightforward.

Estimation of the amplitude envelope is a fairly straightforward process with the use of the Hilbert transform (Oppenheim et al., 1999). The discrete-time Hilbert transform has impulse response function

$$h(n) = \begin{cases} 0, & \text{for n even,} \\ \dfrac{\pi}{2n}, & \text{for n odd,} \end{cases} \tag{3.4}$$

and frequency response function,

$$H(\omega) = \begin{cases} e^{+i\pi/2}, & 0 \le \omega < 0, \\ e^{-i\pi/2}, & 0 \le \omega < \pi. \end{cases} \tag{3.5}$$

The function $H(\omega)$ can be interpreted as a +90 degree shift in the power spectrum of the input signal, $Y(\omega)$, for negative frequency components and -90 degree shift for positive frequency components. In essence, this transformation preserves magnitude information in $y(n)$ but shifts the phase of the information 90 degrees. In the simplest case, the Hilbert transform turns a pure sine wave into a pure cosine wave. Taking the Hilbert transform twice results in a 180 degree shift, yielding $-y(n)$ or the negative of the signal of the original signal.

A Hilbert *transformer* combines $y(n)$ with $H\{y(n)\}$ to yield an analytic or complex signal, $\xi(n) = y(n) + i \cdot H\{y(n)\}$, which satisfies the property

$$\xi(n) = m(n)e^{j\varphi(n)}, \tag{3.6}$$

for

$$\varphi(n) = \arctan\left(\frac{\xi_{imag}(n)}{\xi_{real}(n)}\right). \tag{3.7}$$

Since the absolute value of $e^{j\phi(n)}$ is always 1, the amplitude envelope of $y(n)$ is the absolute value of the output of the Hilbert transformer:

$$m(n) = |\zeta(n)| = |y(n) + i \cdot H\{y(n)\}|. \tag{3.8}$$

Last, the measurement of amplitude modulation can be made by taking the sample-based derivative of the amplitude envelope

$$am(n) \approx \frac{m(n) - m(n-1)}{\Delta}. \tag{3.9}$$

TIME-VARYING FREQUENCY OF A SINUSOIDAL PROCESS

The estimation of FM requires the estimation of instantaneous frequency. Methods for frequency tracking have been proposed based on a time-varying AR model, but have not clearly demonstrated accurate measurement of frequency (Ko & Li, 1994; Mainardi et al., 1995; Arnold et al., 1998; Bencheqroune et al., 1999; Foffani et al., 2004; Tarvainen et al., 2004) . Moreover, for the analysis of LFP and EEG data, there has been little discussion on the practical aspects of applying time-varying AR models. A filter that can estimate the instantaneous frequency of brain rhythms without bias or artifacts would be invaluable to the process of understanding dynamic information processing in the brain.

In the following sections, we develop a practical, robust estimator for instantaneous frequency of a band-limited signal in the form of (3.1). The algorithm is constructed in the same fashion as (Arnold et al., 1998), using an AR process to model the EEG/LFP time series, with the addition of a Kalman smoother to estimate the AR coefficients on a sample-by-sample basis. We reiterate a KF for univariate AR models and show the filter can lead to erroneous instantaneous frequency estimates when applied naively. One source of bias is based on the construction of the KF, whereas the other source is based on the properties of the AR model. We propose modifications to the adaptive estimation procedure to address both sources of bias.

THE AUTOREGRESSIVE PROCESS

In the observation interval $t \in (0,T]$, let time be discretized such that the discrete time index is $n = t/\Delta$, $f_s = 1/\Delta$ is the sampling frequency, and $J = T/\Delta$ is the largest discrete time index. Let $y(n)$ denote the measured LFP signal, $v(n)$ denote the zero-mean Gaussian measurement noise with variance σ_v^2, and p be the order of the autoregressive (AR) process. Then the AR(p) model is

$$y(n) = \sum_{k=1}^{p} a_k y(n-k) + v(n),$$
(3.10)

with z-transform and power spectral density (PSD) functions, respectively,

$$Y(z) = V(z) / \left(1 - \sum_{k=1}^{p} a_k z^{-k}\right),$$
(3.11)

$$S(z) = \sigma_v^2 / \left[\left(1 - \sum_{k=1}^{p} a_k z^{-k}\right)\left(1 - \sum_{k=1}^{p} a_k^* z^k\right)\right].$$
(3.12)

The denominator of (3.11) is referred to as the characteristic polynomial; the magnitude of the roots must be less than 1 for stability; and the roots are referred to as a poles. The PSD of an AR(p) process, $S(z)$, is a multi-modal function that resembles a summation of peaks. The relative height, shape, and location of the peaks are entirely determined by the poles of the transfer function in (3.11), and therefore, by the AR coefficients a_k. The phase of the pole determines where the center of the peak is positioned on the frequency axis, while the height and width of the peak are determined by the magnitude of the pole; the sharpness of the peak is determined by the distance of the pole from the unit circle. Detailed properties of AR processes can be found in standard text books (Box et al., 1994; Oppenheim et al., 1999). Since the a_k's parameterize the shape of the PSD, estimation of the a_k's from the

observed data, $y(n)$, would allow for the computation of the AR poles, and, therefore, the frequencies of oscillatory activity in the data.

THE FIXED INTERVAL KALMAN SMOOTHER FOR AR MODELS

Some methods of estimation assume that a_k is fixed over time (Oppenheim et al., 1999), but the analysis of notoriously nonstationary brain rhythms requires an approach that assumes a dynamic AR process, $a_k \rightarrow a_k(n)$. The KF for univariate AR models has previously been described and is restated here in a concise manner (Arnold et al., 1998). The state-space representation for the AR process is

$$x(n) = x(n-1) + w(n),$$ (3.13)

$$y(n) = C(n)x(n) + v(n),$$ (3.14)

where $x(n) = \left[a_1(n), ..., a_p(n) \right]^T$ is the hidden or state variable corresponding the AR coefficients at discrete time n, $C(n) = \left[y(n-1), ..., y(n-p) \right]$ is the observation transformation vector, $w(n)$ is the state transition noise with covariance Σ_w, $v(n)$ is the observation noise with covariance Σ_v.

The state estimate, $x(n|J)$, is derived from all the observations, $y(n)$, in some observation interval, $n \in [1, J]$, is obtained with the fixed-interval Kalman smoother. The estimation begins by performing a forward pass, where $x(n|n)$ is computed over for $n = 1$ to $n = J$, see (3.20)-(3.22)). Then a backward pass is performed to propagate future information into state estimates at past times ((3.26)-(3.28)). The backward pass reduces

delays associated with adaptive filtering which use non-optimal learning parameters and decreases the variance of the state estimates.

The KF can be described as a sequential Bayesian estimator, therefore, we outline the algorithm using conditional notation to emphasize the history dependence in each step of the estimation process (Haykin, 1996). Eqs. (3.15) to (3.18) are conditional estimators of state of observation statistics,

$$x(n \mid j) = E\left[x(n) \mid y(1),...,y(j)\right], \qquad (3.15)$$

$$\Sigma_x(n \mid j) = \text{cov}\left[x(n) \mid y(1),...,y(j)\right], \qquad (3.16)$$

$$y(n \mid j) = E\left[y(n) \mid y(1),...,y(j)\right], \qquad (3.17)$$

$$\Sigma_y(n \mid j) = \text{cov}\left[y(n) \mid y(1),...,y(j)\right], \qquad (3.18)$$

where n need not be less than, equal to, or greater than j.

The one-step prediction equations use the state space equations, (3.13) and (3.14), to predict the state variables and their statistics.

$$x(n \mid n-1) = x(n-1 \mid n-1) \qquad (3.19)$$

$$\Sigma_x(n \mid n-1) = \Sigma_x(n \mid n-1) + \Sigma_w \qquad (3.20)$$

$$y(n \mid n-1) = C(n)x(n \mid n-1) \qquad (3.21)$$

$$\Sigma_y(n \mid n-1) = C(n)\Sigma_x(n \mid n-1)C(n)^T + \Sigma_v \qquad (3.22)$$

The fact that Σ_w and Σ_v are non-zero implies that the one-step prediction will contain errors. The following correction equations incorporate new observations with a-priori knowledge of the statistics of the state evolution to correct those errors. The Kalman gain, (3.23), scales the innovation term, $y(n) - y(n \mid n-1)$, such that $x(n \mid n)$ is the MMSE estimate of $x(n)$:

$$K(n) = \Sigma_x(n \mid n-1) C(n)^T \Sigma_y(n \mid n-1)^{-1}, \tag{3.23}$$

$$x(n \mid n) = x(n \mid n-1) + K(n)(y(n) - y(n \mid n-1)), \tag{3.24}$$

$$\Sigma_x(n \mid n) = \Sigma_x(n \mid n-1) - K(n)\Sigma_y(n \mid n-1)K(n)^T. \tag{3.25}$$

In contrast to $x(n \mid n)$, $x(n-1 \mid J)$ is an estimate which incorporates past as well as future information. The following backward, smoothing equations are applied after $x(J \mid J)$ has been evaluated.

$$x(n \mid J) = E[x(n) \mid y(1), \ldots, y(J)] \tag{3.26}$$

$$x(n-1 \mid J) = x(n-1 \mid n-1) + A(n-1)(x(n \mid J) - x(n \mid n-1)) \tag{3.27}$$

$$A(n-1) = \Sigma_x(n-1 \mid n-1)\Sigma_x^{-1}(n \mid n-1) \tag{3.28}$$

From the smoothed AR coefficient estimates, $x(n \mid J) = [a_1(n \mid J), \ldots, a_p(n \mid J)]^T$, we can easily obtain estimates of dominant frequencies in the signal at time index n, $f_k(n \mid J)$ for $1 \le k \le p/2$.

$$\left(1 - \sum_{k=1}^{p} a_k(n \mid J) z^{-k}\right) = \prod_{k=1}^{p}\left(1 - b_k(n \mid J) z^{-1}\right) \tag{3.29}$$

$$b_k(n \mid J) = r_k(n \mid J) e^{-i\omega_k(n \mid J)} \tag{3.30}$$

$$f_k(n \mid J) = \frac{f_s}{2\pi}|\omega_k(n \mid J)| \tag{3.31}$$

In condition where p is even and the poles are complex conjugates, $b_i(n \mid J) = b_k^*(n \mid J)$ for $i \ne k$ and $1 \le i,k \le p$, there KF estimator has found $p/2$ oscillations in the data. The complex phase of the poles are proportional to the dominant oscillation frequency according to (3.31). Here, $r(n \mid J)$ is the magnitude of the pole, $\omega(n \mid J)$ is the phase of the pole on the complex plane, and f_s is the sampling frequency of the observed signal. If

p is odd, then at least one pole is guaranteed to be real; real positive poles act as dampers or low pass filters, while real negative pole act as high pass filters. All poles must be less than 1 in magnitude for the AR model to be stable.

Frequency modulation (FM) is defined as the derivative of the instantaneous frequency and is computed by taking the approximated derivative of $f(n\,|\,J)$,

$$fm(n\,|\,J) \approx \frac{f(n\,|\,J) - f(n-1\,|\,J)}{\Delta}. \tag{3.32}$$

Depending on the parameter values for σ_v^2 and σ_w^2, the estimate of $fm(n\,|\,J)$ may need additional smoothing; a Hanning window corresponding to 100ms in length is suitable for this purpose.

For an AR(P) process, the KF parameters are initialized as follows $n_0 = p+1$, $C(n_0) = [\,y(p),\ldots,y(2),y(1)\,]$, $\Sigma_w = \sigma_w^2 I_{PxP}$, $\Sigma_v = \sigma_v^2$, where I_{PxP} is the $P \times P$ identity matrix. The initial state and state covariance matrix, $x(n_0\,|\,n_0)$ and $\Sigma_x(n_0\,|\,n_0)$, are initialized using the Yule-Walker equations. The selection of σ_w^2 and σ_v^2 will be discussed later in the chapter.

UNBIASED ESTIMATION OF FREQUENCY

Simulated Data: amplitude scaled chirps surrounded by white noise
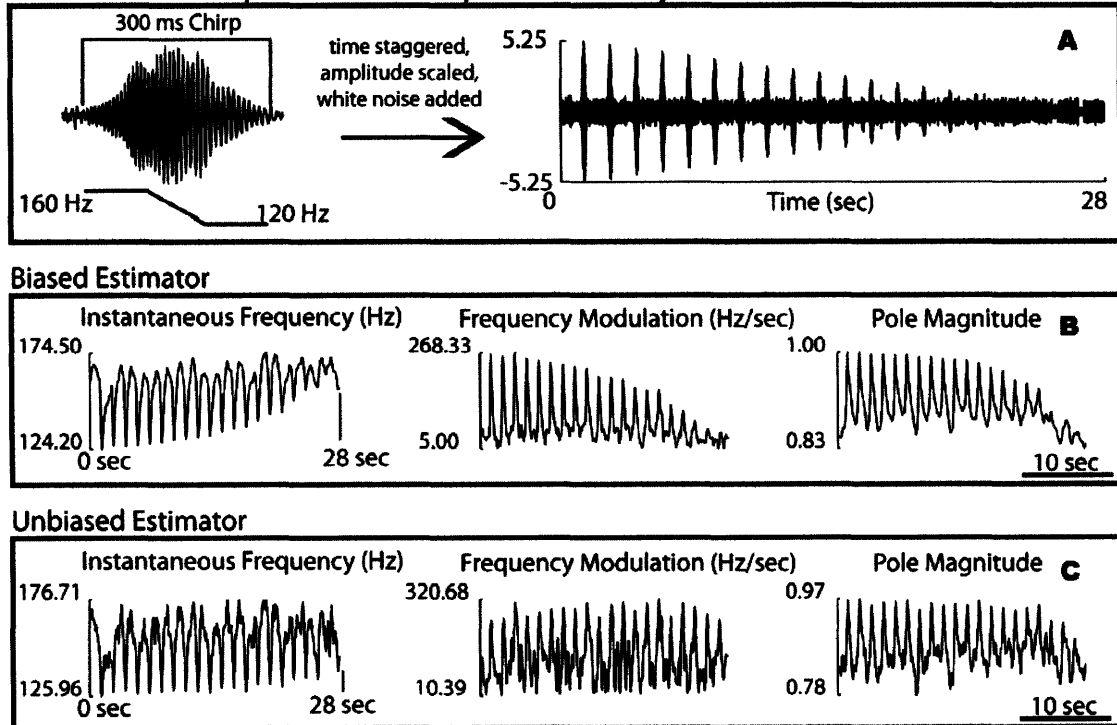
Figure 3.1 Biased instantaneous frequency estimation. (A) Simulated data was created by combining a fast, frequency and amplitude modulated signal with background noise. The signal was sampled at 800 Hz. The first chirp had an amplitude of 5 and the last chirp has an amplitude of 0.25. The background noise is white with variance 0.25. The chirps were not added to noise, but surrounded by noise in time. (B) Kalman based estimates using the biased method with parameters $p = 2, \sigma_v^2 = 0.1, \sigma_w^2 = 0.0001$. **(C) Kalman based estimates using the amplitude demodulated method with parameters** $p = 2$, $\sigma_v^2 = 0.1$, $\sigma_w^2 = 0.0001$.

The basic KF assumes the observation matrix C is constant over time; thus the steady state Kalman gain in (3.23) is guaranteed to converge to a steady state value as $n \to \infty$. On the other hand, if C is allowed to vary over time as in (3.14), the amplitude of the instantaneous Kalman gain will be a function of the amplitude of the past observation,

$\left[y(n-1),...,y(n-p)\right]$. A critical consequence is the state estimates and ultimately the pole phases will be more variable when $\left[y(n-1),...,y(n-p)\right]$ is large relative to when $\left[y(n-1),...,y(n-p)\right]$ is small. Ignoring this condition will still allow the KF to track oscillations in the data over time, but in a consistently biased manner (Figure 3.1, Panel B). Since our goal is to hone in on small time-scale changes in instantaneous frequency, artificial structure produced by biased filtering will ultimately lead to erroneous inferences.

The bias discussed here can be ameliorated by amplitude demodulating the input signal $y(n)$ using (3.8),

$$y_{AD}(n) = y(n)/m(n). \tag{3.33}$$

With $y_{AD}(n)$ the variation in peak-amplitude is extremely predictable and well modeled by the KF. This process of amplitude demodulation combined with Kalman filtering will be referred to as the *amplitude demodulated-Kalman filter (AD-KF)*. In tracking sinusoids in noise, this manipulation drastically decreases the systematic scaling of the Kalman gain with amplitude. The tradeoff here, as can clearly be seen from Figure 3.1-C, is the variance of the estimate may also increase. These variations can be systematically reduced by either increasing the value of σ_v^2 or decreasing the value of σ_w^2. Each manipulation will lead to smoother AR parameter estimates but with different qualities (Figure 3.6).
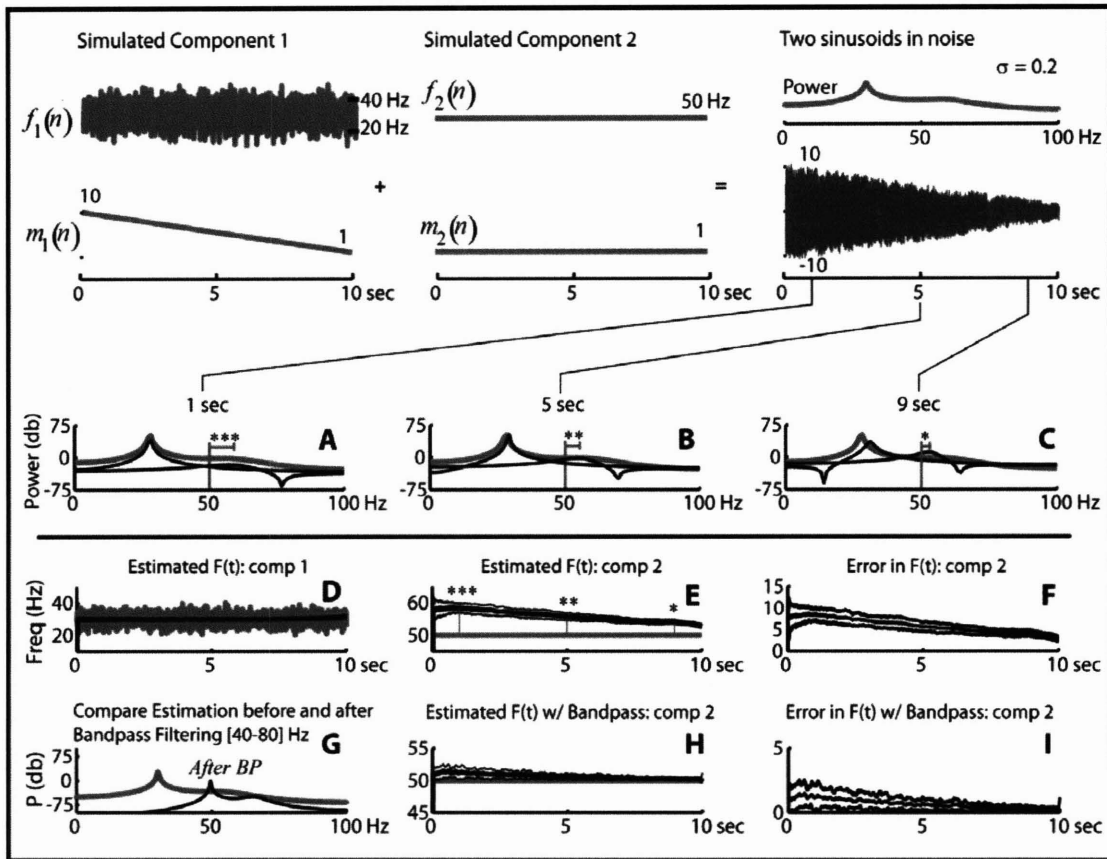
Figure 3.2 Autoregressive pole interactions. The six top-most panels illustrate how the data was simulated. The first sinusoidal component was both amplitude and frequency modulated; the frequency modulation, shown at the top, caused the instantaneous frequency to range between 20 Hz and 40 Hz; the amplitude modulation is clear from the realization of the time series. The second simulated sinusoidal component was constant in amplitude and frequency. The two components were summed and white noise ($\sigma^2 = 0.04$) was added. (A,B,C) Time-staggered AR PSD's estimated from the simulated data. The gray line is the full PSD while the dark lines represent the contribution of each pole to the PSD. The asterisks show how the representation of the 50 Hz component changes in the estimate even though it does not change in theory. (D and E) The gray lines are the true frequency while the black lines are the estimated frequencies. The estimates are shown with 1 std intervals. (F) The error in the frequency estimate of component #2, shown with 1 std intervals. (G) The original simulated signal is bandpass filtered between 40 and 80 Hz. (H and I) Bandpass filtering before estimation decreases error in the estimate. In panels A-F, the KF parameters were $p = 4$, $\sigma_v^2 = 0.1$, $\sigma_w^2 = 0.01$. In panels G-I, the AD-KF parameters were $p = 2$, $\sigma_v^2 = 0.1$, $\sigma_w^2 = 0.01$.

INTERACTIONS BETWEEN MULTIPLE POLES

AR models are advantageous in time-frequency analysis because they represent a flexible class of PSD functions that are best described as a linear summation of peaks, where the location, height, and width of each peak is dictated by the poles of the AR model. The overlap between side lobes of the peaks consequently means that the power at any given frequency is based on the interaction between multiple poles (Figure 3.2 (A-C)). The $1/f$ rule, which is well known in the analysis of brain rhythms, states that the expected power of an oscillation is inversely proportional to the frequency of the oscillation (Buzsáki, 2006). This has strong implications for frequency estimation where the sidelobes of poles at lower frequencies may dominate poles at higher frequencies; higher frequency poles can therefore vary as a function of activity at lower frequencies. For example, if the KF ($p \geq 4$) reveals changes in frequency in the γ-rhythm, these changes may be due to actual γ-rhythm modulation or may be a side effect of activity in θ, α, or β-rhythms. Figure 3.2 (A-F) illustrates the interaction between poles in the process of frequency estimation and how the interaction is amplified by the $1/f$ phenomena.

The most effective way to address this interaction is to isolate and analyze each rhythm individually. Bandpass filtering the observation signal prior to Kalman filtering helps to ensure changes in pole amplitude and location within a frequency band of interest is truly a result of activity in that band. Moreover, an AR(2) model can be assumed with reasonable certainty and model identification would no longer be necessary. In Figure 3.2, (G-I), the example demonstrates how bandpass filtering, even over a large frequency range, can help to remove artificial structure in the frequency estimates. In practice, the cutoff frequencies for bandpass filtering can be determined by visual inspection of a spectrogram.
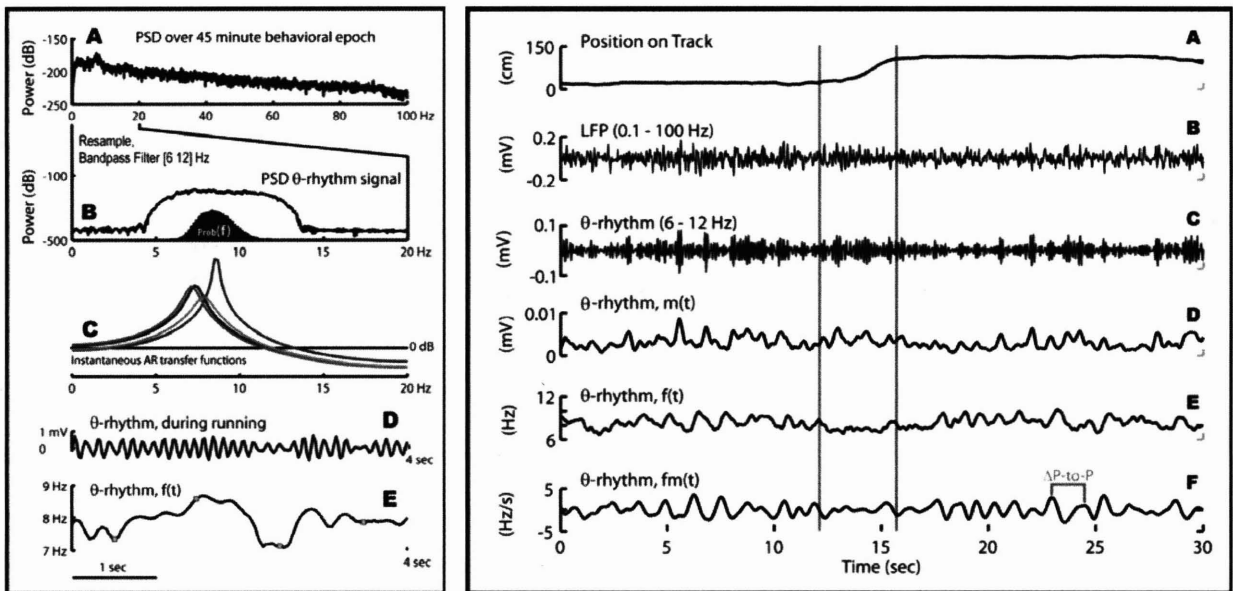
Figure 3.3 (left): Tracking θ-rhythm dynamics. (A) The PSD of the LFP estimated from 45 minutes of data. **(B)** The PSD of the LFP after downsampling by a ratio of 1/5 and bandpass filtering between 6 and 12 Hz using a 20 point FIR filter. The histogram that is shown on the same axis, $prob(f)$, is the histogram of the instantaneous frequency estimated from θ-rhythm using the AD-KF ($p = 2$, $\sigma_v^2 = 0.1$, $\sigma_w^2 = 0.07$). **(C)** AR transfer function profiles at times 0.5, 1.5, 2.5, and 3.5 seconds from panel-E. The darker lines correspond to earlier time points. **(D-E)** A 4 second example of the -rhythm and it's respective frequency estimate.

Figure 3.4 (right): State dependent θ-rhythm dynamics. An example of real data recorded from CA1 of the rat hippocampus during a linear track task. **(A)** The track is 1 meter in length and in these 30 seconds the rat runs from one end to the other. **(B)** LFP Data sampled and 200 Hz. **(C)** θ-rhythm signal sampled at 40 Hz. **(D)** The amplitude envelope of the θ-rhythm which is computed by taking the absolute value of the Hilbert transform. **(E)** The instantaneous frequency estimate. **(F)** The frequency modulation estimate. The gray vertical lines mark the interval when the rat is running.

In general, the parameters for bandpass filtering will affect the quality of the instantaneous frequency estimate. Choosing a pass band that is too large may result in the presence of more than one oscillation. In this case, the frequency represented by the AR(2) parameter estimates will be an amplitude-weighted average of all the oscillations that are present in

the pass band. This limitation is of importance when each individual oscillation represents a distinct neurophysiological process; by bandpass filtering and assuming an AR(2) process, a model-data mismatch has occurred, and individual neural processes will erroneously be lumped into one process. The consequences of this limitation cannot be stated absolutely, and will ultimately depend on the scientific question being asked and the interpretation of the adaptive filter output.

When the cutoff frequencies for the bandpass filter are known from inspection of the data, the next order of business is resampling the data to create a symmetric power distribution in the frequency domain about the point at $\pi/2$ radians. This geometry is necessary because AR components become more asymmetric in shape as their center frequency moves closer to either 0 or π radians. Therefore, the interplay between data and AR model is not uniform over all boundaries. Furthermore, bandpass filtering the data creates to sharp edges in power which add to the complexity of the data-model interactions. Since uniformity cannot be achieved the next step is to impose symmetry: the interaction between data and model will be the same at the lower and higher edges of the pass band.

A simple recipe can be used for conditioning the signal to obtain an approximately symmetric power distribution about $\pi/2$ radians. First, define the frequency range for the brain rhythm of interest $\left(\overline{f_1}, \overline{f_2}\right)$, and then evaluate a central frequency parameter $\overline{f_c} = 0.5\left(\overline{f_1} + \overline{f_2}\right)$. Second define the Nyquist and sampling frequencies to be $\overline{f_n} \approx 2\overline{f_c}$ and $\overline{f_s} \triangleq 2\overline{f_n}$. Resampling the signal to approximately $\overline{f_s}$ and then bandpass filtering between $\overline{f_1}$ and $\overline{f_2}$ will provide the largest *symmetric* range of motion for the pole in the frequency band of interest.

MODEL AND PARAMETER SELECTION

A prerequisite problem when using AR models is determining the order of the model. In this case, we have shown that the prerequisite problem is trivial since it is always possible to band limit the input signal and assume an AR(2) process. The AR(2) process in combination with the band limited signal is preferred because there are no pole-pole interactions, the tracking of instantaneous frequency is trivial since there is only one pole, and the use of frequency estimates are better defined for scientific inference.

We are then left with the problem of parameter selection. The algorithm requires four parameters to be defined $\bar{f_1}$ and $\bar{f_2}$ for bandpass filtering, and σ_w^2 and σ_v^2 for the KF. The parameters for bandpass filtering can be determined in two ways: bootstrapping from published physiology studies or by inspection of a spectrogram. In practice, the best scenario is when both methods are in agreement. If they are not in agreement, then the original scientific question should be revisited and the parameter set that best addresses the question should be chosen. In a systematic manner, the limits for bandpass filtering should be systematically varied to test if the results hold in a robust or predictable manner.

The selection of KF parameters σ_w^2 and σ_v^2 can be addressed more systematically. A multitude of parameter selection methods have been developed, each using different information criteria or different cost functions based on residuals or estimation errors of the time domain signal (Djuric & Kay, 1992; de Waele & Broersen, 2003). In relation to the task at hand, time domain optimizers are inadequate since minimizing the error in the time domain does not necessary lead to minimizing the error in instantaneous frequency estimation. In order to address this concern, we have developed a criteria based on frequency information. We call this measure OIC as the Greek letter "omega" is the roman equivalent of the letter "o".

The first step in determining the values for σ_w^2 and σ_v^2 is constraining the possible set of values. By either normalizing the data such that $2 \times std(\tilde{y}) = 1$, where $\tilde{y}(n) = y(n)/\text{constant}$, or using the amplitude demodulation procedure in (3.8), $\tilde{y}(n) = y(n)/m(n)$, we can guarantee that suitable values for both σ_w^2 and σ_v^2 lie in the range of 0.1 and 0.001. This way, once a set of parameters have been found, they can be used over and over again for the same type of data, regardless of the amplitude of the measurements. The value for σ_w^2 tends to more sensitive to sampling frequency than σ_v^2, and as the sampling frequency changes, σ_w^2 should be reduced accordingly to reflect the timescale of the relevant activity in the EEG/LFP (Figure 3.5).
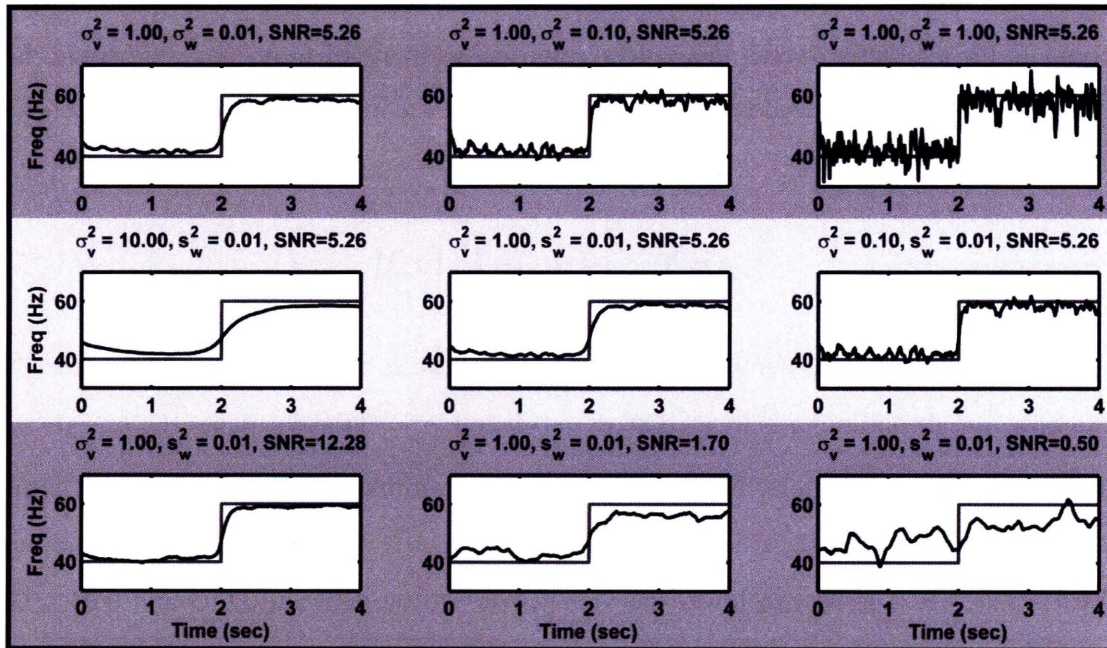

Figure 3.5 : Interactions between KF noise variance parameters and signal-to-noise ratio. (First row) Decreasing σ_w^2 improves the ability to track frequency on a sample-by-sample basis. (Second row) Increasing σ_v^2 smooths the estimate over longer timescales than increasing σ_w^2. (third row) Frequency estimates are robust up to a SNR of approximately 2.

The OIC uses local estimates of frequency over time based on the peaks in the signal. The OIC works best when the input signal to the KF is bandlimited and one oscillation clearly dominates at each point in time. After normalizing the input signal, the next step is to find all the peaks in the data that are atleast 0.5 in magnitude larger than either adjacent trough (this removes spurious noise from the estimate). Let the times of the peaks be $\tau_1, ..., \tau_M$ for M peaks and the times between the peaks be $\upsilon_m = (\tau_m + \tau_{m+1})/2$ for $m = 1, ..., M - 1$. The estimate of time-varying frequency is computed as,

$$\tilde{f}(\upsilon_m) = 1/(\tau_{m+1} - \tau_m),$$ (3.34)

where the location of the frequency estimate in time is placed between the peaks used to compute the estimate. Given this purely empirical estimate of frequency, one could then compute the following residual based on the output of the instantaneous frequency filter in (3.31):

$$\varepsilon_{WIC} = \sum_{m=1}^{M-1} \left(\tilde{f}(\upsilon_m) - \tilde{f}_{KF}(\upsilon_m) \right)^2.$$ (3.35)

An example of the FIC parameter selection procedure is show in Figure 3.6. For this example, the first minima is located at $\sigma_v^2 = 0.1$ and $\sigma_w^2 = 0.0196$. Panel (E) reveals that the most optimal parameter set yields a frequency estimate that contains abrupt jumps consistent with overfitting a noisy signal. Since noise affects the frequency estimate on a local timescale, it is best to lower the value of σ_w^2 to be between 0.005 and 0.01. The reasoning for this can be seen in Figure 3.5 where increasing σ_v^2 smoothes the estimate over long timescales, whereas decreasing σ_w^2 increasing the ability to track on small timescales.

Noisy time series will result in noisy empirical estimates of frequency. If we assume the signal is a pure sinusoidal then the power of the noise is the power in the band that does not correspond to the sinusoid. With this definition, it may be possible to compute the

SNR from PSD estimates using the multitaper wavelet method or any other PSD estimator (Percival & Walden, 1993). Having the SNR will help to determine if the optimal values under the OIC will lead to noise fitting. If the SNR is too low (< 2), the frequency estimates may not be valid at all. If the SNR is midrange, (2 < SNR < 6), then it may be sufficient to increase σ_v^2 by an order of magnitude. For large SNR, the optimal value given by the OIC is most likely sufficient for estimation.
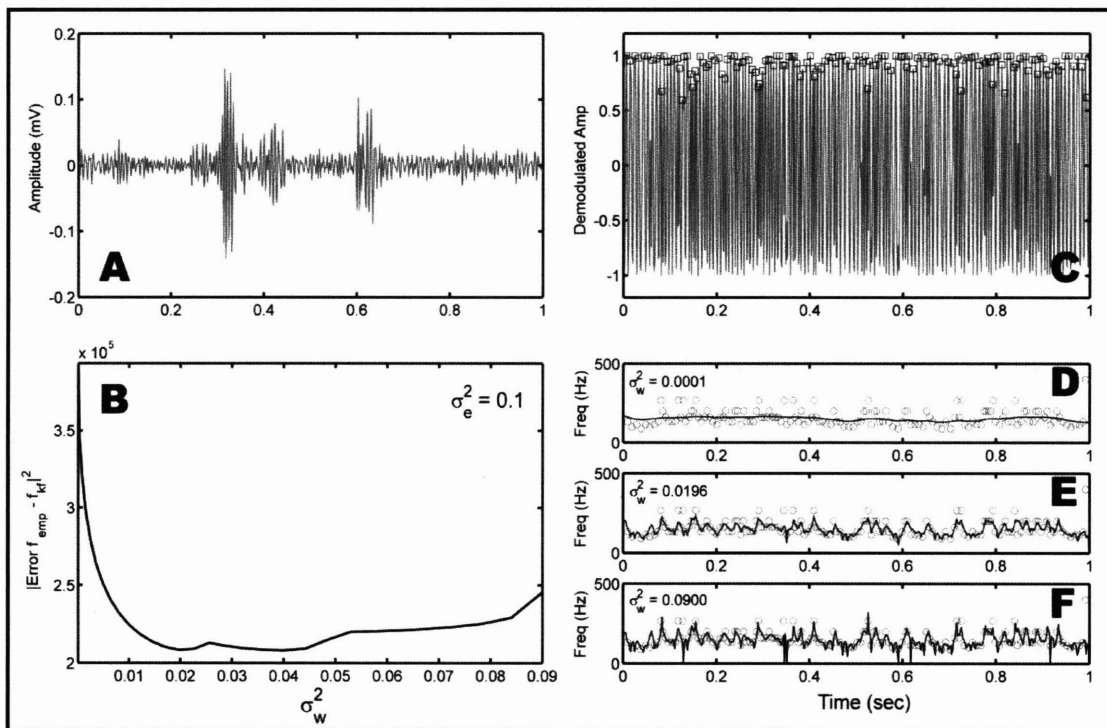
Figure 3.6: Kalman filter parameter selection. (A) 1 second trace of a ripple recorded from stratum oriens of the dorsal rat hippocampus. (B) The estimation error in frequency associated with $\sigma_v^2 = 0.1$ and σ_w^2 varying according to the x-axis. (C) The amplitude demodulated ripple signal from panel (A). The identified peaks are marked with square boxes. (D-F) Circles mark the time and value of the empirical frequency estimates. The solid line is the estimate of frequency given by the AD-KF.

TESTING GOODNESS-OF-FIT

One issue with band-pass filtering the data is that multiple sinusoids may be present in that band. Our assumption of an AR(2) model is incompatible with such scenarios. Therefore, post hoc analyses to determine were the single sinusoid assumption has been violated will be needed to verify our estimated of instantaneous frequency.

As we discussed earlier, one output of our Kalman filter is the residual time-series which contains the component of the input signal that is not captured by the adaptive AR(2) process. Ideally, if the filter captures all the rhythmic structure in the data, the residual should be uncorrelated with itself and the autocorrelation of the residual should be flat, which corresponds to a white process. It has been shown that the Ljung-Box-Pierce statistic (Box et al., 1994),

$$Q(K) = N_w (N_w + 2) \sum_{k=1}^{K} (N_w - k)^{-1} \gamma_k^2 (\hat{v}),$$

(1.36)

(where N_w is the length of the data segment used to compute the autocorrelation, $\hat{v}$ is the residual output of the Kalman filter, γ_k is the value of the autocorrelation at lag k, and K is the number of lags in the autocorrelation used to compute the statistic), is distributed according to $\chi^2(K-2)$ for an AR(2) process when γ_k is the autocorrelation of $\hat{v}$ and $\hat{v}$ is a white noise process. When $\hat{v}$ contains autocorrelated values, the distribution of γ_k should be inflated such that $Q(K)$ will outside the right-sided 95% confidence interval of the $\chi^2(K-2)$ distribution. In Figure 3.7, an example is shown where the pass-band of the pre-filter is expanded to gradually include another oscillation. As more of the second oscillation is included, the AR(2) model attempts to capture both oscillations but fails. This inability is apparent in Figure 3.7(F,G), where the autocorrelation clearly shows oscillatory structure and the Q(20) statistic is clearly outside of the expected range.
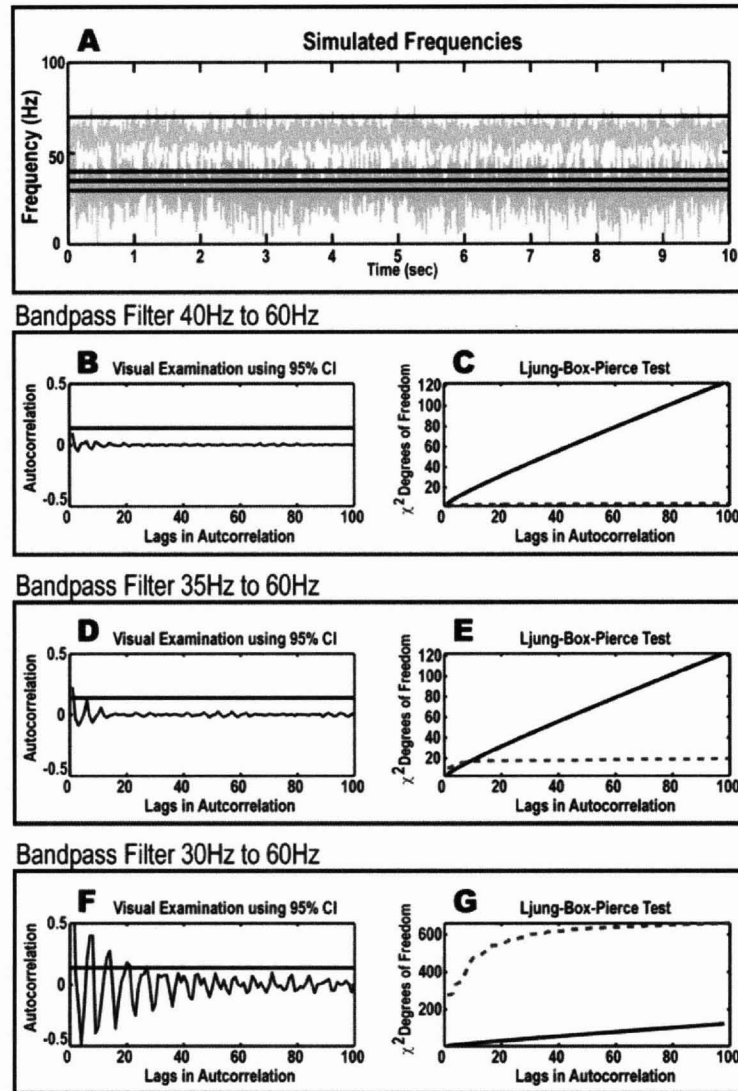
Figure 3.7 : Ljung-Box-Pierce test for lack-of-fit. (A) As the pass band becomes larger 40-60Hz, 35-60Hz, and then 30-60Hz, more of the low frequency component is included in the input to the AR(2) Kalman filter. (B,D,F) The right-half of the autocorrelation of the filter residuals. The solid horizontal line is the theoretical 95% confidence interval of the autocorrelation of a white noise process. (C,E,G) The output of the Ljung-Box-Pierce Test for lack of fit. In practice, if the dashed line is below the solid line at lag 20, this supports a good-fit of the data.

ANALYSIS OF LFP FROM THE RAT HIPPOCAMPUS

A rat implanted with a tetrode array targeted to the hippocampus was trained to run back and forth on a U-shaped track for a food reward. At the beginning of the experiment, only one out of three arms of the U-track was exposed to the rat. After 10 minutes, a second arm was exposed and so the length of the track doubled and the track resembled an L-shape. After another 15 minutes, a third track arm was exposed and the rat was allowed to run a full U-trajectory. LFP data was recorded from stratum pyramidal of the dorsal CA1 region at a rate of 1 kHz and bandpass filtered between 1 and 475 Hz. The total observation time was approximately 48 minutes. The θ-rhythm was isolated by first bandpass filtering to prevent aliasing, downsampling the LFP to a 40 Hz sampling rate, and finally bandpass filtering between 6 and 12 Hz ().

In the hippocampus, θ-oscillations are tightly bound to dominant information processing and coordination mechanisms as well as to different behaviors (Vanderwolf, 1969; Wyble et al., 2004). The goal here is to determine if there is any systematic, fine timescale structure in the amplitude and frequency modulation of θ-band activity as a function of behavioral state. Previously, reports have shown that θ-oscillations are modulated in amplitude but not frequency during running (Czurko et al., 1999; Wyble et al., 2004). During immobility, an attention related form of θ-oscillation has been reported but has not been well characterized (Vanderwolf, 1969).

We extracted instantaneous amplitude information using the magnitude of the Hilbert transform on the θ-signal. We extracted instantaneous frequency information using the AD-KF for AR(2) models with parameters $p = 2$, $\sigma_v^2 = 0.1$, $\sigma_w^2 = 0.07$. The process of analyzing θ-activity is outlined in Figure 3.3. Figure 3.4 is a typical example of θ-activity during running and non-running behaviors. First, the dynamics of θ-band activity is visibly dependent on the behavioral state of the animal. Second, when comparing amplitude and frequency dynamics, it is interesting to note that amplitude and frequency modulation can occur independent of one another. Third, while the rat is running, the

frequency of the θ-oscillation is held constant while the amplitude can vary significantly. Fourth, modulations in θ-rhythm frequency can occur at rates with an upper bound of 5 Hz/sec.
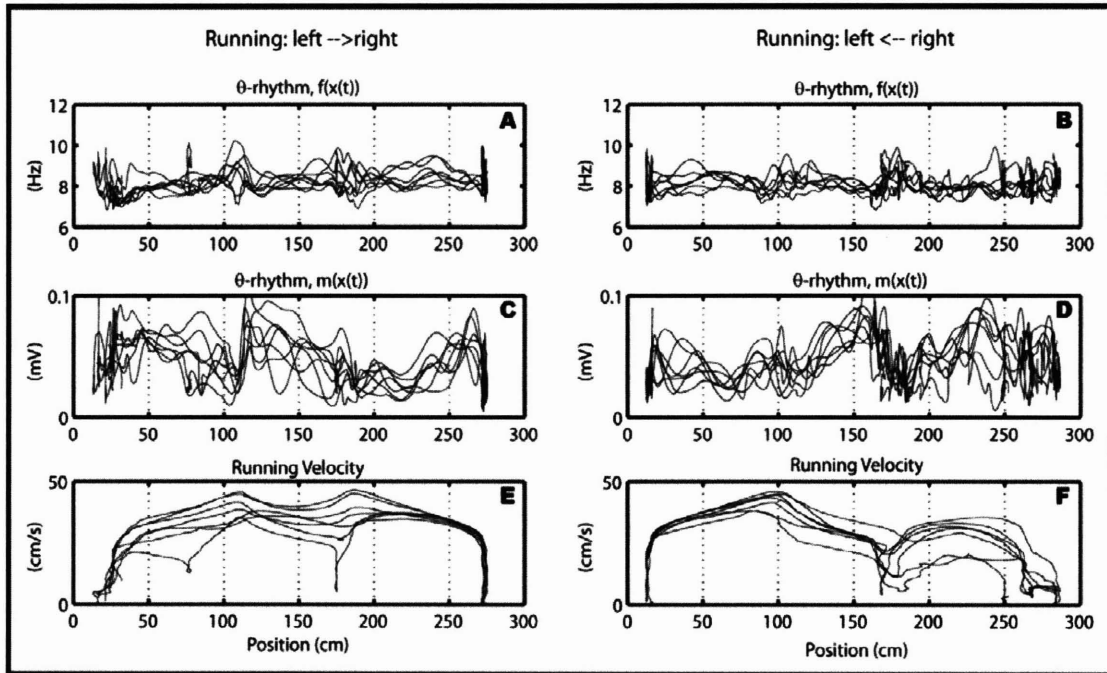
Figure 3.8: θ-rhythm dynamics are consistent on a trial-by-trial basis. The point of this figure is to show consistent and dynamic changes in θ-rhythm dynamics as a function of behavior, therefore, average responses are not shown. Here, the track is approximately 3 meters in length and is U-shaped. The distance is measured along the path of the track. Each line represents a single pass in one direction or the other and each plot contains 7 passes. In E and F, changes in velocity help to determine the location of the turns in the U-track and the track ends. (A-B) Instantaneous frequency estimates using the AD-KF ($p = 2$, $\sigma_v^2 = 0.1$, $\sigma_w^2 = 0.07$). (C-D) θ-rhythm amplitude envelope computed from the absolute value of the Hilbert transform. (E-F) The velocity of the rat as a function of position. The average time it took to run from one end to the other was 9.5 sec with a std of 2.5 sec.

Since it is hard to see the fine-timescale dynamics of θ-activity in Figure 3.4, we isolated 7 successive laps while the rat was running on the full U-shaped track. In order to make the lap-by-lap comparisons easier to visualize, we plotted each parameter as a function of position and made separate plots for each direction of running, Figure 3.8. Inspection of θ-frequency, θ-amplitude, and running velocity reveals several interesting points about θ-rhythm dynamics during locomotion. First, θ-amplitude is best predicted by distance from the most recent track end and direction may play a role as well. Second, the instantaneous frequency during run hovers around 8 Hz which is consistent with previous reports (Slawinska & Kasicki, 1998; Czurko et al., 1999; Wyble et al., 2004; Sinnamon, 2006). Third, for only the left $\rightarrow$ right trajectory, there appears to be a systematic ramping of frequency as the rat first starts to run. Fourth, at the corners of the U-track, the change in θ-amplitude is highly consistent but the change in θ-frequency is highly variable.

In general, rats will be distracted or lose motivation to run during an experiment; having the rat on the linear portion of the track does not guarantee movement. In order to investigate systematic variations in θ-dynamics as a function of locomotion behavior, we dissected the observation interval into epochs of locomotion (high velocity) and stopping (low velocity) by setting a 10 cm/sec velocity threshold. Using 48 minutes of data, some of which are shown in Figure 3.4 and Figure 3.8, we compared the distributions of θ-amplitude, θ-frequency, θ-frequency modulation, and time differences between positive peaks in θ-frequency modulation during low and high velocity epochs. Figure 3.9 shows box plots of the different parameters. For each individual epoch, the mean of each parameter was computed, and the means over all epochs were compared using a 1-way Anova. The only parameters that did not show a difference in distribution were the standard deviation of θ-amplitude, the mean of the ΔP-to-P parameters, and the standard deviation of the ΔP-to-P parameter (Figure 3.9). In general, the average θ-amplitude was higher during running although the variation was statistically the same in both conditions. The average instantaneous frequency and variation was lower during the high velocity state, which is expected (Czurko et al., 1999). The peak-to-peak time delay in

frequency modulation averaged around 0.55 sec, but was much more variable during the high velocity state.
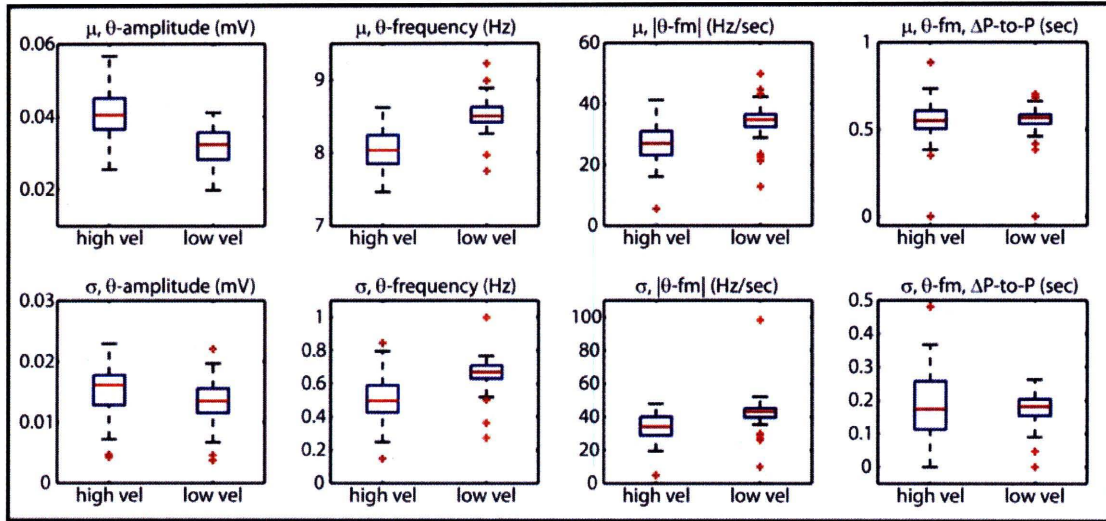
Figure 3.9. Statistical differences in the θ-rhythm dynamics between low and high velocity behaviors. The 48 minute observation interval was divided into epochs of high and low velocity (threshold 10 cm/sec). For each epoch (N=89 high vel, N=88 low vel), the mean and standard deviation of each parameter was computed. The box plots display the mean value (A-D) and standard deviation (E-H). The columns from left to right represent the amplitude, frequency, absolute value of frequency modulation, and positive-peak to positive-peak interval in frequency modulation. A visual description of θ-fm ΔP-to-P can be found **Figure 3.4(F)**. Comparisons between high and low velocity statistics were performed using a one way anova. The resulting p-values for rejecting the null hypothesis are **(A)** $p < 1.16e-11$ **(b)** $p < 3.27e-8$ **(c)** $3.33e-16$ **(d)** $p < 0.52$ **(e)** $p < 0.018$ **(f)** $p < 2.58e-6$ **(g)** $3.35e-10$ **(h)** $p < 0.49$.

These analyses demonstrate the utility of instantaneous frequency estimation in the characterization of real LFP data. We provide evidence for independent modulation of both amplitude and frequency during high and low velocity behaviors. Furthermore, using frequency modulation we were able to characterize a systematic θ-frequency modulation during high and low velocity. Overall, the adaptive filter allowed for temporally precise characterization of frequency content and was able to reveal systematic differences in θ-activity as a function of behavioral state.

DISCUSSION

We discuss three important issues when using the KF to estimate univariate autoregressive models for the purpose of estimating oscillation frequencies of multiple sinusoids with $1/f$ behavior. The first issue we discussed was the amplitude-induced bias in the estimation of instantaneous frequency. More specifically, we noted that, for the observation equation in (3.14) the Kalman gain becomes a function of the observations. This dependence prevents the Kalman gain from ever converging to steady state and, more importantly, scales the Kalman gain proportionally to the amplitude of the observations. We addressed this bias by amplitude demodulating the signal so only frequency modulation was left in the signal before Kalman filtering. The second issue of importance was the fact that interactions between poles can have undesired effects on instantaneous frequency estimates (Figure 3.2). We showed that first bandpass filtering the observation signal in some band of interest was a viable, pragmatic compromise with two advantages: reducing estimation bias due to cross-talk between oscillations and avoiding the prerequisite problem of AR model identification. Third, we developed a novel frequency based information criterion (OIC) for determining the best parameters for instantaneous frequency estimation. We applied the resulting algorithm to real data recorded from the rat hippocampus during a linear track task and found novel structure in the dynamic modulation of θ-band activity.

The frequency estimation algorithm is a relatively simple, slim-downed version of adaptive filters that have previously described (Arnold et al., 1998; Tarvainen et al., 2004). In designing the algorithm, we had two purposes in mind: unbiased instantaneous frequency estimation of brain rhythms and ease of use. We showed that bandpass filtering and modeling with an AR(2) process provides more reliable frequency estimates, and relieves the user from solving the prerequisite model identification problem. Here we have clearly shown that the increase in estimation accuracy outweighs the cost of sequentially analyzing each brain rhythm. Furthermore, because the AR(2) model is simpler, the process of determining the KF variance parameters is also more straightforward for the user. In

addition, diagnostics such as the Ljung-Box-Pierce test are readily available for verification of frequency estimates. For these reasons, software could be designed to make instantaneous frequency estimation available to the average neuroscientist.

The analysis of θ-activity in CA1 during a linear track experiment revealed interesting observations of fine time scale structure of frequency dynamics and its relation to amplitude dynamics. Briefly, those observations were instantaneous θ-frequency during locomotion was very steady at around 8Hz; during non-running behaviors θ-activity could range between 6 Hz and 11 Hz; and during both running and non-running behaviors θ-frequency was modulated at an average interval of 0.6 sec although there was less variance in this interval during non-locomotion. These results demonstrate the presence of systematic, fine-time scale structure in LFP data. These findings motivate further analyses such as investigating the activity of principal neurons and interneurons during periods of θ-frequency and/or θ-amplitude modulation in order to understand the computational significance of these modulatory events.

In summary, we developed an adaptive filter for tracking frequency of brain rhythms and clearly demonstrated how to apply this filter in order to obtain robust and interpretable measurements. The procedure is relatively simple to use with four intuitive parameters to define (frequency limits of the brain rhythm and KF variances). The OIC is highly informative to the task of choosing KF variances as makes direct use of the frequency content in the data. The output of the estimation algorithm can be interpreted directly in a physiological framework. Future directions would look into developing reliable measures of connectivity between locations in the brain using instantaneous frequency or frequency modulation; and mixed filtering that could integrate point process models of spiking activity with Kalman filtering of EEG/LFP activity.

Chapter 4 : RIPPLE DETECTION AND CHARACTERIZATION USING INSTANTANEOUS AMPLITUDE AND FREQUENCY CONTENT

Abbreviations: LFP = local field potential, MUA = multi-unit activity, SPW = sharp wave, SWS = slow wave sleep, EC = entorhinal cortex, KF = Kalman filter, AD-KF = Amplitude demodulated Kalman filter, AR = autoregressive, PSD = power spectral density, IFQ = instantaneous frequency, AI = amplitude information, AFI = amplitude and frequency information.

INTRODUCTION

In the local field potential of the rat hippocampus, high frequency (100-250Hz) and short duration (30-100ms) events, termed "ripples", can be found in slow wave sleep, awake immobility, and awake exploration (Ylinen et al., 1995; O'Neill et al., 2006). Ripples in CA1 tend to co-occur with sharp wave depolarizations that are initiated by synchronous activity in CA3 (Csicsvari et al., 2000). In the1990's, ripples were known to correlate with an increase in multi-unit activity (MUA) of principal and interneuron types (Ylinen et al., 1995). Another decade of research revealed the MUA associated with ripples contained highly structured and compressed re-expression of temporal patterns found in ensemble activity during running on a track (Lee & Wilson, 2002; Foster & Wilson, 2006; Diba & Buzsaki, 2007).

Reports from the Buzsaki Lab support the idea that synchronous activity in CA3 pyramidal cells propagates to CA1, then to subiculum, and on to entorhinal cortex; along the way strong synchrony can be observed in the LFP as ripple events (Ylinen et al., 1995; Chrobak & Buzsaki, 1996; Csicsvari et al., 2000). Csicsvari et al. (2000) recorded simultaneously from CA1 and CA3 subfields, and showed that CA1 ripples were highly correlated with the firing of individual units in CA3 and that synchrony of CA3 units were a strong predictor for SPW-Ripple complexes within a 100ms time window. Chrobak et al. (1996) showed that ripples that were simultaneously recorded from CA1 and entorhinal cortex (EC) were either synchronous or delayed by up to 30ms (CA1 preceding in time); and that cells in layer V of EC fired more regularly with the negative peaks of EC relative to cells in the superficial layers of EC. Therefore, using CA1 ripples as a reference point, this

evidence established the cause of ripple events in CA1 to be CA3 activity and the effect to be synchronous activation of parahippocampal targets.

The story, however, is not so simple. Csicsvari *et al.* (1999) showed that ripples could be categorized into two classes: slow ripples (100-130 Hz) and fast ripples (140-200Hz); and proposed that fast ripples were unlikely a result of CA3 input and more likely a result of intrinsic CA1 activity (Csicsvari et al., 1999). In 2000, Csicsvari *et al.* showed that small amplitude CA1 ripples were not coherent across electrodes while large amplitude CA1 ripples were. In the same paper, they also provided evidence for topography in ripple initiation where CA3b activity can induce ripples in CA1b, and CA3a activity can induce ripples in CA1a, and any other combination was unlikely. Therefore, if CA3 is driven in synchrony and in wide expanse, the sharp wave-ripple complex will more likely be global and easier to detect. In contrast, small amplitude ripples may be a result of local synchrony in the CA3 population, the CA1 population, or other input population, such as the EC via the temporoammonic pathway, for example.

The significance of the ripple event, whether large or small, is seen widely by the hippocampus community as a temporal window of opportunity for inducing changes in synaptic reorganization (Buzsaki, 1989). The synchronous nature of ensemble input into CA1 can minimize the difference in pre- and post-synaptic spike timing, and therefore maximize plasticity under the rules of spike-timing dependent plasticity (Bi & Poo, 1998). The physical location of the ripple event, on the medial-lateral and septal-temporal axes, may determine the quality of the information being replayed as dictated by the topographical projections from the entorhinal cortex (Witter et al., 1989; Hampson et al., 1999; Hafting et al., 2005; Sargolini et al., 2006; Witter, 2006). The occurrence of ripples during sleep and waking states supports the idea that the hippocampus is constantly engaged in the maintenance of memory when demands on the hippocampus are permissive (Siapas & Wilson, 1998; Sirota et al., 2003; O'Neill et al., 2006; Cheng & Frank, 2008). Together, these anatomical and physiological findings support the association of ripple activity with the unique and necessary function of the hippocampus in information encoding and memory consolidation. In order to better understand the role of ripple

associated activity as a process for memory maintenance, we have focused on two highly relevant and understated problems: (1) ripple characterization and (2) ripple detection.

The long standing definition of ripple events are amplitude based; anything above a certain power in the 100-250 Hz band is a ripple in practice. This definition, although practical, does not do justice to the richness of the ripple event and all the information that may be contained. Only two studies have made advancements in the analysis of ripples based on wavelet methods, but lack the resolution to identify fine timescale structure in individual ripple events (Sirota et al., 2003; Gillis et al., 2005). Using an adaptive filter for instantaneous frequency (IFQ) estimation (Chapter 5), we explore the possibility of using IFQ to enhance the definition of a ripple for more robust ripple detection and more accurate ripple characterization.

In this chapter, we show that changes in frequency within a ripple event can be on the order of thousands of Hz per second, the direction of frequency modulation (FM) can be either positive or negative, and the relative timing of FM can precede and follow the time of peak amplitude. These variations suggest that there exists more than one physiologically meaningful class of ripple events. We classified ripples based on average IFQ and average FM within a ripple (+/- 10ms about ripple peak) and performed statistical analyses. In terms of ripple detection, we combined ripple amplitude and frequency information and found that detection based on the combined feature set yields only a subset of the events compared to the use of amplitude information only. We will discuss the advantages and limitations of this novel method for ripple detection further. We considered the underlying structure of long (> 100 ms) and short ripple events (< 100ms) by detecting the number of "ripple units" per event. A linear regression analysis revealed the average duration of the ripple unit that was shorter than expected at 36ms; this result is discussed in the text.

METHODS AND MATERIALS

The data was recorded from the pyramidal cell layer of dorsal CA1 during a 60 minute sleep epoch, where a 4 month old, male, Long-Evans rat was resting on an elevated platform inside a sleep chamber that was only open to the ceiling. The LFP signal was recorded to disk, downsampled to 800 Hz, and then bandpass filtered using a 100 point FIR filter with cutoff frequencies at 100 Hz and 250 Hz. The band-pass filtered signal containing ripples is referred to as the ripple signal. The data is modeled using an autoregressive process of order 2 or an AR(2) process. The instantaneous frequency of the ripple signal was estimated twice, once with the normal KF and once with amplitude demodulated-KF (AD-KF) as in Chapter 5. Frequency estimates obtained from the KF are referred to as *censored* estimates since the tracking of frequency is dependent on ripple signal amplitude. Frequency estimates from the AD-KF are referred to as *uncensored estimates*. The censored frequency estimate can be interpreted as a signal which contains both amplitude and frequency information.

The Kalman filter parameters were initialized as follows: $n_0 = 3$, $C(n_0) = [y(2), y(1)]$, $\Sigma_w = \sigma_w^2 I_{2x2}$, $\Sigma_v = \sigma_v^2$, where $\sigma_w^2 = 0.005$, $\sigma_v^2 = 0.05$, and I_{2x2} is the 2x2 identity matrix for the normal KF. Two parameters were altered for the AD-KF, $\sigma_w^2 = 0.005$, $\sigma_v^2 = 0.1$. The initial state and state covariance matrix, $x(n_0 | n_0)$ and $\Sigma_x(n_0 | n_0)$, are initialized using the Yule-Walker equations and 10 seconds of data. We used the omega information criterion (OIC) to narrow down the range of σ_v^2 and σ_w^2 by using a double for-loop to explore a large parameter space (Chapter 5).

Ripple amplitude was measured by taking the largest peak of the ripple-band signal (no Hilbert transforms used here). Ripple frequency was taken as the average frequency in a +/- 10ms window centered about the time of the largest peak of the ripple oscillation. The ripple frequency modulation was taken as the average frequency modulation in a +/- 10ms window centered about the time of the largest peak of the ripple oscillation.

110

Ripple detection based on amplitude information was performed as follows. The absolute value of the Hilbert transformer was taken as the "ripple envelope". When the envelope exceeded 0.05mV, this flagged a ripple event. The start and end times of the ripple event were marked by the first crossing of 0.04mV to the left and right of the ripple peak. These thresholds correspond to 3 and 2.5 times the standard deviation of the signal. These thresholds were chosen by visual inspection as average ripple amplitude and SNR will vary as a function of electrode location.

RESULTS

Adaptive AR Modeling of Ripple Frequency

Sharp wave-ripple ripple complexes were observed in local field potentials (LFPs) recorded from the CA1 subfield of the dorsal hippocampus while the rat was in a state of immobility or slow wave sleep. An example of LFP from the data set can be found in Figure 4.1(A). The estimation of instantaneous frequency (IFQ) was performed twice on the original ripple signal, Figure 4.1, once using the KF and once using the AD-KF. The resulting estimates of IFQ clearly show that the tracking of frequency is *censored* when the KF is used and the ripple signal has low amplitude (Figure 4.1(D), black solid line). In contrast, the estimate obtained from the AD-KF is dynamic throughout (*uncensored*) and is consistent with the estimates from the KF when the amplitude of the ripple signal is large. In this work, the uncensored FM estimate is used to characterize activity only during ripple events. The censored FM signal is a biased estimate of frequency dynamics and practically serves as an indicator of FM during epochs of large amplitude.

Figure 4.1 : Example of ripples along with censored and uncensored frequency estimates. The uncensored frequency estimates are derived from the AD-KF, while censored frequency estimates are derived from the normal KF. (A) 500 ms trace of an LFP recording taken from stratum oriens of the CA1 region of the dorsal hippocampus. (B) Bandpass filtered signal between 100-250 Hz retains high frequency oscillations while removing sharp waves. (C) The ripple signal is normalized by its amplitude envelope to produce a sinusoid that is only modulated in frequency. (D) The censored-frequency estimate (solid black) given by the normal KF, and the uncensored-frequency estimate (dashed-gray) given by the AD-KF. (E) The change in frequency or FM derived from the respective frequency estimates in (D).
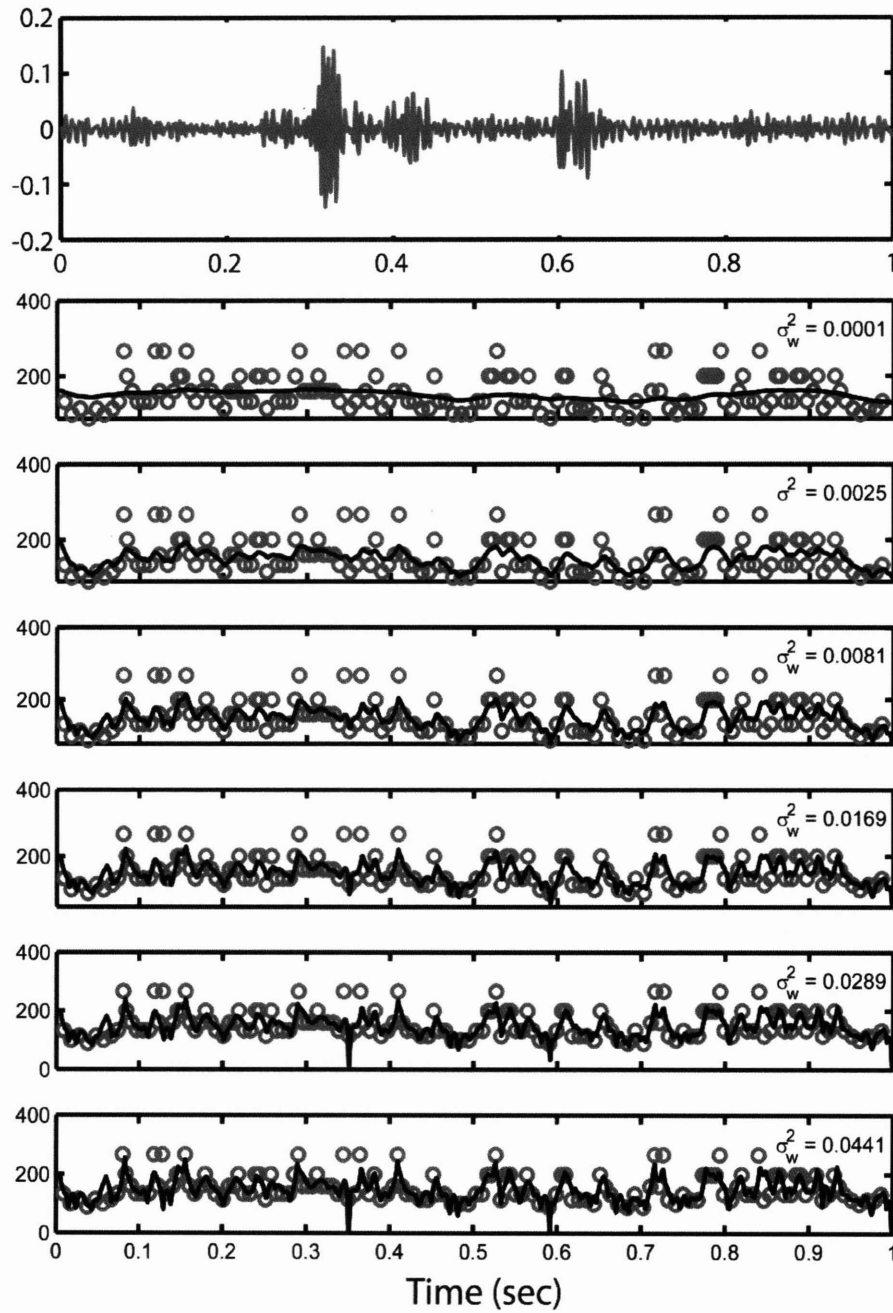
Figure 4.2 : Selection of KF parameters. The grey circles represent empirical estimates of frequency based on peaks in the ripple signal which has units of mV (top panel). The solid black line is the AD-KF (non-censored) estimate of frequency based on $\sigma_v^2 = 0.1$ and σ_w^2 stated in the respective plots.

113

The frequency modulation (FM) signal emphasizes the dynamics in IFQ (Figure 4.1(E)). One advantage of using FM to characterize ripple activity is that the absolute value of the IFQ is removed. Thus, FM within ripple events may reveal structure that might otherwise be obscured by variability in ripple frequency in general. In Figure 4.1(E), it is easy to see a pulsatile pattern in the FM signal (gray dashed) given by the AD-KF. Considering the power in the ripple band is very low relative to the total power of the raw LFP during times of non-ripple activity, these frequency fluctuations are most likely due to a "desynchronized state" or measurement noise. Therefore, the FM signal is only used when it can be referenced to an operating regime that has been identified by physiological measurements such as the LFP, MUA, and/or behavioral state.

The estimates of IFQ and FM are obtained by using a KF for AR(2) processes. The parameters of the KF and AD-KF are selected by matching the IFQ estimates with the empirical frequency estimates. The process is explained in detail in Chapter 5 and can be clearly seen in Figure 4.2 for the AD-KF. With σ_v^2 held fixed at 0.1, the value for σ_w^2 was altered systematically. The empirical estimates of time-varying frequency are extremely helpful in terms of guiding the parameter selection process. Large fluctuations in empirical frequency are mostly likely due to noise and can be disregarded. Sharp changes in the AD-KF estimates are indications of over-fitting. With these considerations, the most appropriate values for σ_w^2 lie between 0.002 and 0.008. We took the best value to be the midpoint at $\sigma_w^2 = 0.005$.

The frequency content of the ripple is modeled by an AR(2) process. In Figure 4.3(B), we can observe how the evolution of the AR process is driven by the ripple signal. Each line in panel (B) is the shape of the power spectral density (PSD). The pole magnitude in Figure 4.3(C), directly relates to sharpness of the peak in the PSD. Pole magnitudes near 1 represent oscillations while values near 0.5 represent noisy wideband activity.

Figure 4.3 : Dynamics of AR parameters over a 1 second interval. (A) A representative snapshot of activity in the ripple band (100-250Hz). (B) Dynamic PSD based on AD-KF estimates. There are 8 PSD's snapshots per every 10ms. (C) Pole magnitude informs us about the specificity of the oscillation. Lower magnitude marks larger sidelobes in the PSD, and, therefore, segments of noisy oscillations.

Characterization of Ripple Events

Figure 4.4 (A,B,C) are distributions of ripple amplitude, ripple frequency, and ripple frequency modulation. In panel (A), the amplitude distribution reflects a 0.05 mV cutoff for ripple detection; the edge on the left is not sharp because ripple detection is based on the Hilbert transform of the ripple signal while the amplitude is measured in the ripple signal. In panel (B), the distribution for ripple frequency has short tails due to the band limited nature of the ripple signal. In general, the range of frequency and FM do not differ greatly between noise and ripple events. What is markedly different is the correlation between frequency and FM Figure 4.4 (F,I). In panel (F), ripples clearly show a strong negative correlation between frequency and FM ($\rho = -0.1442$, p < 1e-10).
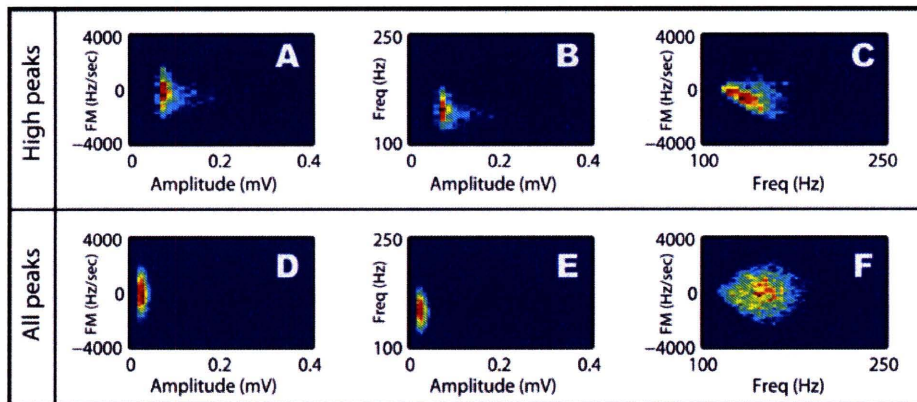
Figure 4.4 : Distribution of ripple characteristics. The threshold for ripple detection by amplitude was set to 0.05 mV (high peaks). Amplitudes were measured as the largest peak in the ripple oscillation event. The frequency of the ripple is defined to be the average frequency in a +/- 10ms window about the time of the maximum ripple peak. The frequency modulation is defined as the average FM in a +/- 10ms window about the time of the highest ripple peak. (D-F) 2-D histograms show the expected covariation in amplitude and frequency parameters during ripple events. (G-I) 2-D histograms show the expected covariation in amplitude and frequency parameters at random times. The correlation between frequency and FM during ripple events is $\rho = -0.1442$ (p < 1e-11), and during random times $\rho = 0.0005$ (p = 0.96)
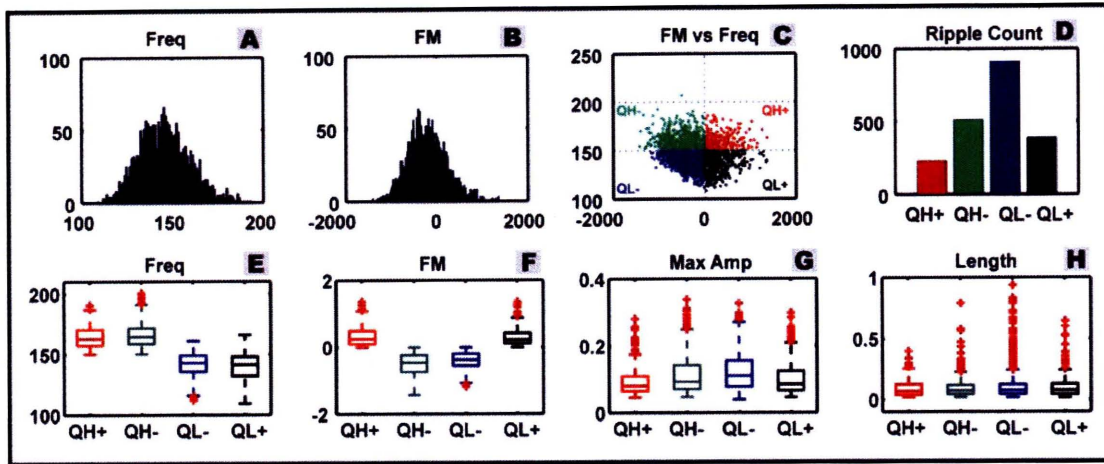
Figure 4.5 : Classification of ripple events based on frequency information. (A) Distribution of ripple frequency (B) Distribution of frequency modulation during ripple. (C) Classification of ripple events based on positive and negative FM, and low and high ripple frequency: QH+ represents high frequency and positive FM, QH- represents high frequency and negative FM, QL- represents low frequency and negative FM, QL- represents low frequency and positive FM. (D) Number of events for each class. In total, there are 2301 events. (E-H) Comparison of classes for each parameter distribution.

5% | 50% | 95%	QH+	QH-	QL-	QL+
Frequency (Hz)	150.7 | 158.3 | 179.9	150.9 | 158.9 | 179.5	123.5 | 138.6 | 148.8	119.1 | 136.6 | 148.7
FM (Hz/sec)	20.6 | 250.5 | 910.5	-1046 | -455.4 | -54	-816.3| -372.9 | -42.8	17.3 | 226.9 | 854.1
Amplitude (mV)	0.05 | 0.08 | 0.19	0.05 | 0.09 | 0.27	0.06 | 0.11 | 0.23	0.05 | 0.08 | 0.21
Length (ms)	23.7 | 67.4 | 229.3	27.5 | 74.9 | 231.0	26.2 | 77.4 | 318.3	25.0 | 78.6 | 289.8

Table 4.1 : Statistics of ripple features for each classes. The values in each cell represent 5%, 50%, and 95% of the population

117

In our inspection of ripple amplitude and frequency content, we found that ripples vary greatly in terms of envelope shape, frequency, and frequency modulation. In order to prevent fine time-scale structure from being averaged out, ripples were first grouped according to frequency and FM. Figure 4.5 summarizes the procedure for classification and the results of the clustering. Figure 4.5(A) is the distribution of ripple frequency; the point of partitioning was chosen to be the mode of the distribution (~150Hz). Although the distribution for FM is skewed to the left, the cutoff point was chosen to be 0 Hz/sec because it leads to straightforward interpretation of results, Figure 4.5(B). Out of the 2034 ripple events that were found, ~40% of the units fell into the low frequency, negative FM category (QL-); ~25% into the high frequency, negative FM category (QH-); ~20% into the low frequency, positive FM category; and ~15% in the high frequency, positive FM category. The frequency distribution for QH classes was significantly different from QL classes as expected. The distribution for FM was different for all class pairs except for QH+ and QL+ (bootstrap, p = 0.52). Amplitude distributions were similar for QL+ and QH+, QH- and QL-, QH- and QL+. The distributions for ripple length were not significantly different between any class pairs (Table 4.2).

	Frequency			FM			Amplitude			Length		
	QH-	QL-	QL+	QH-	QL-	QL+	QH-	QL-	QL+	QH-	QL-	QL+
QH+	0.6	<u>1E-11</u>	<u>1E-11</u>	<u>1E-11</u>	<u>1E-11</u>	0.52	<u>2.3E-2</u>	<u>6E-4</u>	0.3	0.6	0.25	0.25
QH-	X	<u>1E-11</u>	<u>1E-11</u>	X	<u>1E-7</u>	<u>1E-11</u>	X	5.6E-2	6E-2	X	0.18	0.33
QL-	X	X	4E-2	X	X	<u>1E-11</u>	X	X	<u>3E-4</u>	X	X	0.68

Table 4.2 : Results from pair-wise hypothesis tests (bootstrap of sample mean, two tailed, $\alpha = 0.05$). Tests which show a significant difference in the two distributions have the p-value underlined.
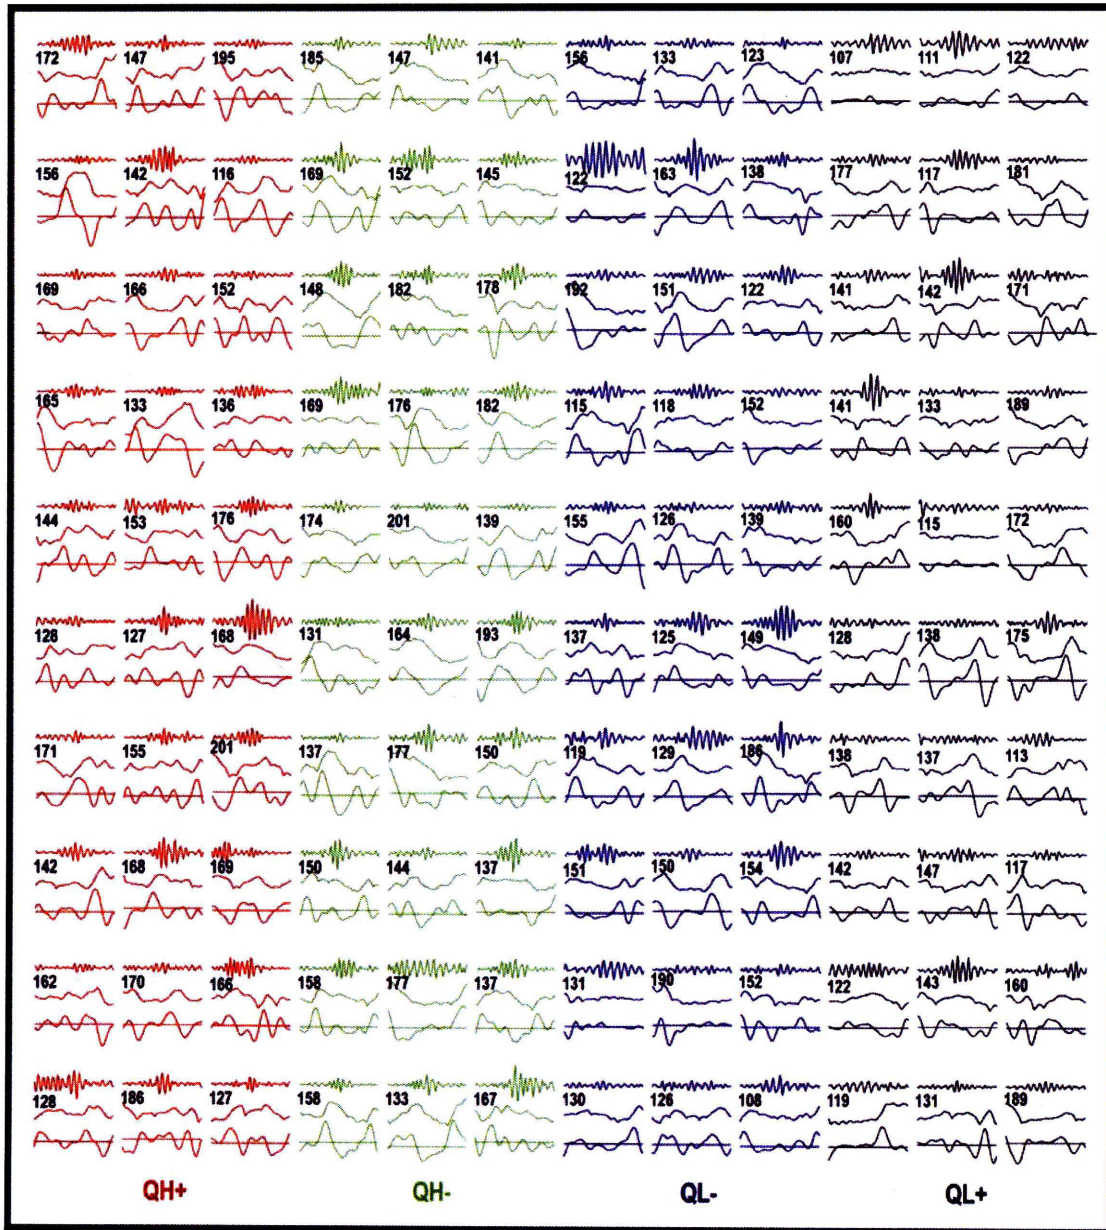
Figure 4.6 : Real data examples from each class. Each segment is 50ms long and is centered about the largest peak of the ripple. For each ripple event, the ripple, uncensored frequency estimate, uncensored FM estimate are shown. The number below each ripple trace is the instantaneous frequency at the beginning of the trace. The horizontal line through the FM estimate is the zero FM line.

Figure 4.6 contains raw data examples of ripples from different classes. Inspection of individual events in each class will reveal that some ripples may be better suited as part of other classes. In regard to this point, we have investigated different parameters for partitioning the events and found that no algorithm can use a low dimensional representation of ripple events and still provide perfect classification. The use of frequency and FM are best because the partition points are easily defined and the interpretation is direct.
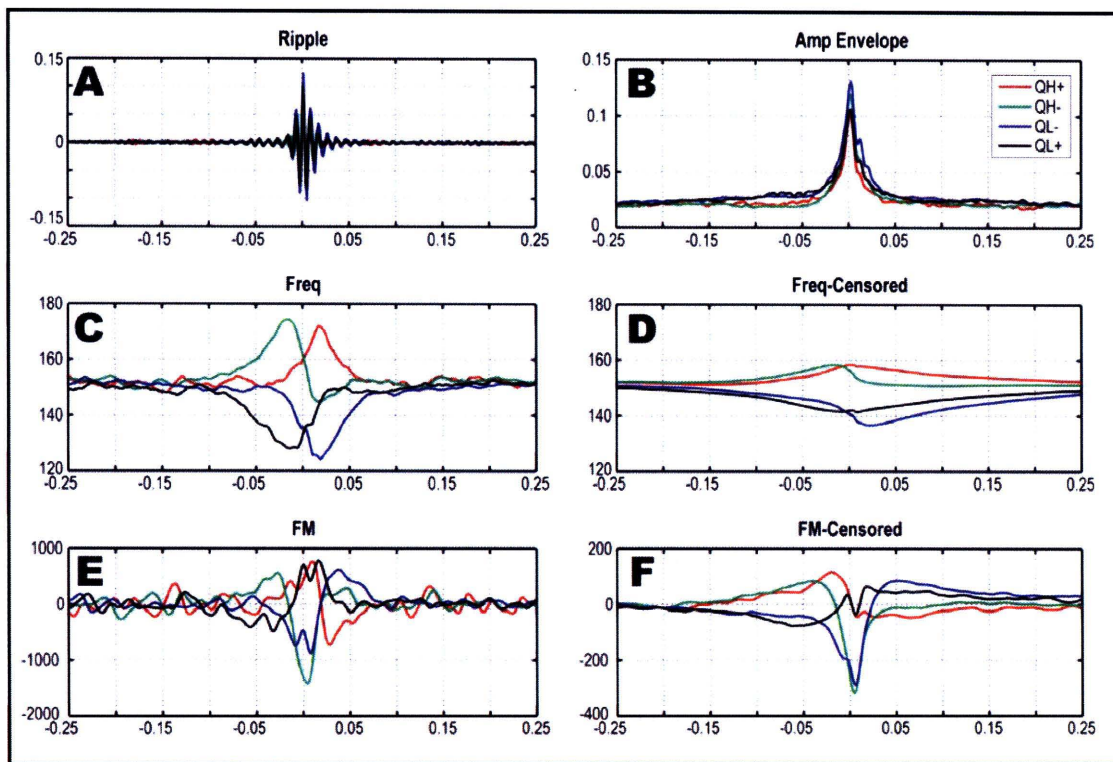
Figure 4.7 : Ripple triggered averages by class. (A) The average of the ripple signal (B) The average of the amplitude envelope of the ripple which is measured by taking absolute value of the Hilbert transformer output for each ripple event. (C) Average uncensored frequency. (D) Average censored frequency. (E) Average uncensored frequency modulation. (F) Average censored frequency modulation.

Ripple triggered averages aligned to the large peak of the ripple oscillations revealed distinct frequency profiles between classes. On average, FM occurred within a +/- 10ms window centered around the ripple peak, Figure 4.7(C). At the point of the ripple peak, the average values of frequency and FM are consistent with how the ripples were partitioned, Figure 4.7(C).

Censored estimates of ripple frequency are less dynamic, however, as seen in Figure 4.7(D), it is still easy to distinguish between high and low frequency classes. Average censored FM was clearly negative for the QL- and QH-, but was near zero for QL+ and QH+, Figure 4.7(F). The differences in average censored and non-censored FM suggest that the positive FM in QL+ and QH+ is likely to be caused by high frequency, low amplitude noise at the tail of the ripple. Whether or not the activity at the tail of the ripple is meaningful is addressed in the discussion.
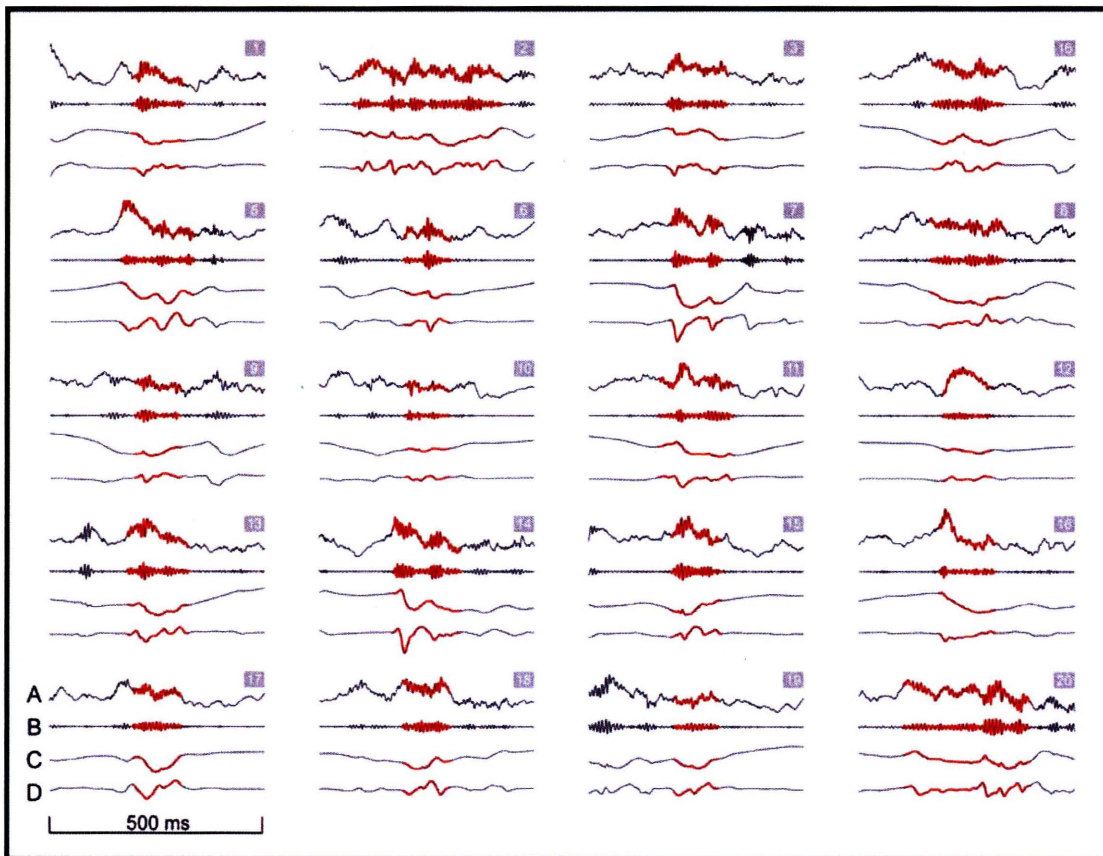
Figure 4.8 : Real data examples of long ripples. A long ripple is defined here to be a ripple event that is 100 ms or longer. For each event, the order of the traces from top to bottom are: A) raw LFP, B) ripple signal (100-250Hz), C) censored frequency estimate, and D) censored FM estimate. The red outline marks the ripple event and sets it apart from the larger data snapshot.

Characterization of Long Ripples

Ripples are thought to be the reflection of highly synchronous activity occurring in a window of 30ms-50ms (Chrobak & Buzsaki, 1996). Often, ripples of greater length can be found. The distribution of ripple length reveals that 50% of the ripples are less than 75 ms in duration, Figure 4.9(A). Of the other 50%, what do they really represent? We considered two possibilities: 1) the long ripple is a unitary event which constitutes a longer period of synchronous input from CA3 or 2) the long ripple is actually a chaining of multiple shorter ripples that overlap in time.

Figure 4.9 : Detection of ripple units in ripple bursts. (A) The distribution of ripple lengths as defined by a 0.05 mV upper detection threshold and 0.04 mV lower threshold. (B) Distribution of ripple unit count for all events (n=2034). (C) Correlation of ripple length and number of ripple units. On average, each unit is 36 ms in duration.

In order to investigate this, we examined by eye ripple events that were found to be greater than 100ms in duration. Some examples of these events are shown in Figure 4.8 along with their *censored* frequency and FM estimates. Censored estimates are used in this situation because they are less sensitive to noise than the uncensored estimates but still exhibit consistent structure within ripple classes. In general, longer ripples tend to co-occurred with multiple sharp waves associated activity in stratum oriens, however, a more definite statement would use sharp wave activity from stratum radiatum (trace A). Examples 12, 16, 19 of Figure 4.8 contain a single long SPW event or no SPW event at all.

Ripple triggered averages in Figure 4.7(D) reveal a unimodal frequency profile and unimodal FM profile that is consistent with a strong pulse like input (off-on-off) where during the ON-period, the timing of inter-spike intervals is consistent with spike adaptation (Dzhala & Staley, 2004).

Based on this observation, sharp changes in frequency during long ripples suggest the onset of a new ripple unit. Although all the events in Figure 4.8 fall into the long ripple category, not all of the events show clear multi-modal frequency dynamics. This could be due to the fact that the censored frequency estimate has slow frequency tracking when the amplitude of the ripple is low. Another possibility is that unitary ripple events can last for longer periods than what has been reported.

We used the structure of averaged ripple events to form a heuristic algorithm for estimating the number of ripple units in a long ripple. The algorithm found the minima (troughs) of the absolute value of the censored FM signal and required that the absolute value of the minima be less than 100 Hz/sec. If no trough meeting that criterion was found in the time segment defining the ripple, a single ripple unit was assumed. Otherwise, the estimated number of ripple units was equal to the number of troughs. The relationship between ripple length and number of ripple units is shown as a scatter plot in Figure 4.9(C) for all ripple events (n=2034). A linear regression analysis revealed that on average, each ripple unit is 36 ms in duration. This heuristic may not be robust enough to accurately state the number of units for each ripple event, but may serve well as a tool for characterizing the relationship between "ripple bursts" and other measures on a population level.

Ripple Detection using combined amplitude and frequency information

A prerequisite problem in the analysis of ripple activity is deciding when a ripple event begins and ends. The most common approach is to rectify the 100-250Hz band-limited LFP and set an upper and lower amplitude threshold. The upper threshold is used to locate the ripple itself while the lower threshold is used to determine the extent of the ripple in time. The first difficulty with this approach is there is no systematic method for determining the threshold. The second difficulty is a single threshold (upper or lower) can be optimal in a global sense but may not be optimal for shorter periods.

One way to address these issues is to use "more" information about ripple structure to detect ripples. In addition to improved detection of ripple boundaries, more information would allow for better separation of true ripple events from noise events. We investigated this possibility by using the *censored* FM signal as an indicator signal for ripple activity. The censored FM signal is greatest in absolute value when both the amplitude of the ripple is highest and the change in frequency during the ripple is highest. Mid-range values of censored FM would imply large FM in combination with low amplitude or high amplitude in combination with low FM. Figure 4.10 provides an example of how amplitude information (AI) alone compares with amplitude+frequency information (AFI).

Figure 4.10 : Comparison of ripple detection methods: Amplitude information (AI) vs. Amplitude + Frequency information (AFI). The top-most trace is the raw LFP and the peak of the sharp wave is marked as 0.76 mV. The ripple signal has been magnified 3X relative to the LFP trace. The ripple envelope can be used for ripple detection. In this schematic, any event that is above the dashed line would be considered to be a ripple event. In the bottom-most trace, the absolute value of the censored FM signal can be used to indicate when a ripple has occurred. What is not shown is the lower thresholds for determining the length of the ripple; for AI it is 0.4mV and AFI it is 200 Hz/sec.
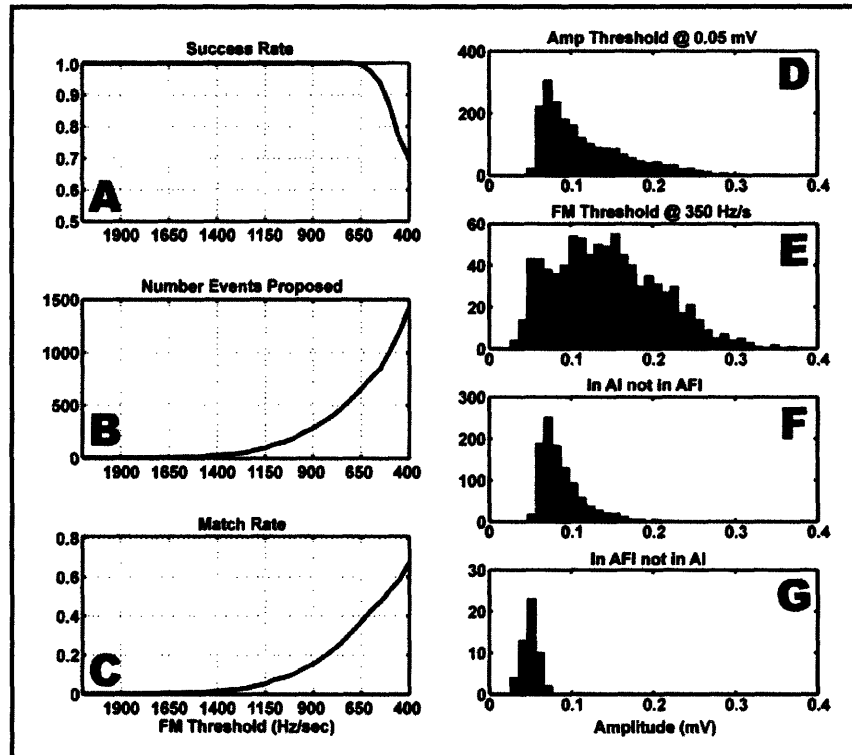
Figure 4.11 : Performance of frequency based ripple detection. Events found by amplitude threshold are considered to be the "true" set of events. The events proposed by FM threshold are tested to see if they overlap with true events. The lower threshold for determining the start and end of the ripple was 0.4mv for AI and 200 Hz/sec for AFI. (A) The success rate is the number of proposed events that match true events divided by the total number of events. (B) As the threshold is lowered, the number of proposed events increases. (C) The match rate is the number of true events accounted for by proposed events divided by the total number of true events. (D) Distribution of amplitude for ripple events found by setting an amplitude threshold at 0.05 mV. (E) Distribution of proposed event amplitudes found by setting a threshold at 350 Hz/second. (F) The amplitude distribution for true events that were not accounted for by proposed events. (G) The amplitude distribution of proposed events that did not match any true events.

Figure 4.12 : Examples of ripples from mutually exclusive categories. (Left) Ripple events corresponding to Panel (F) of Figure 4.11. (Right) Ripple events correspondi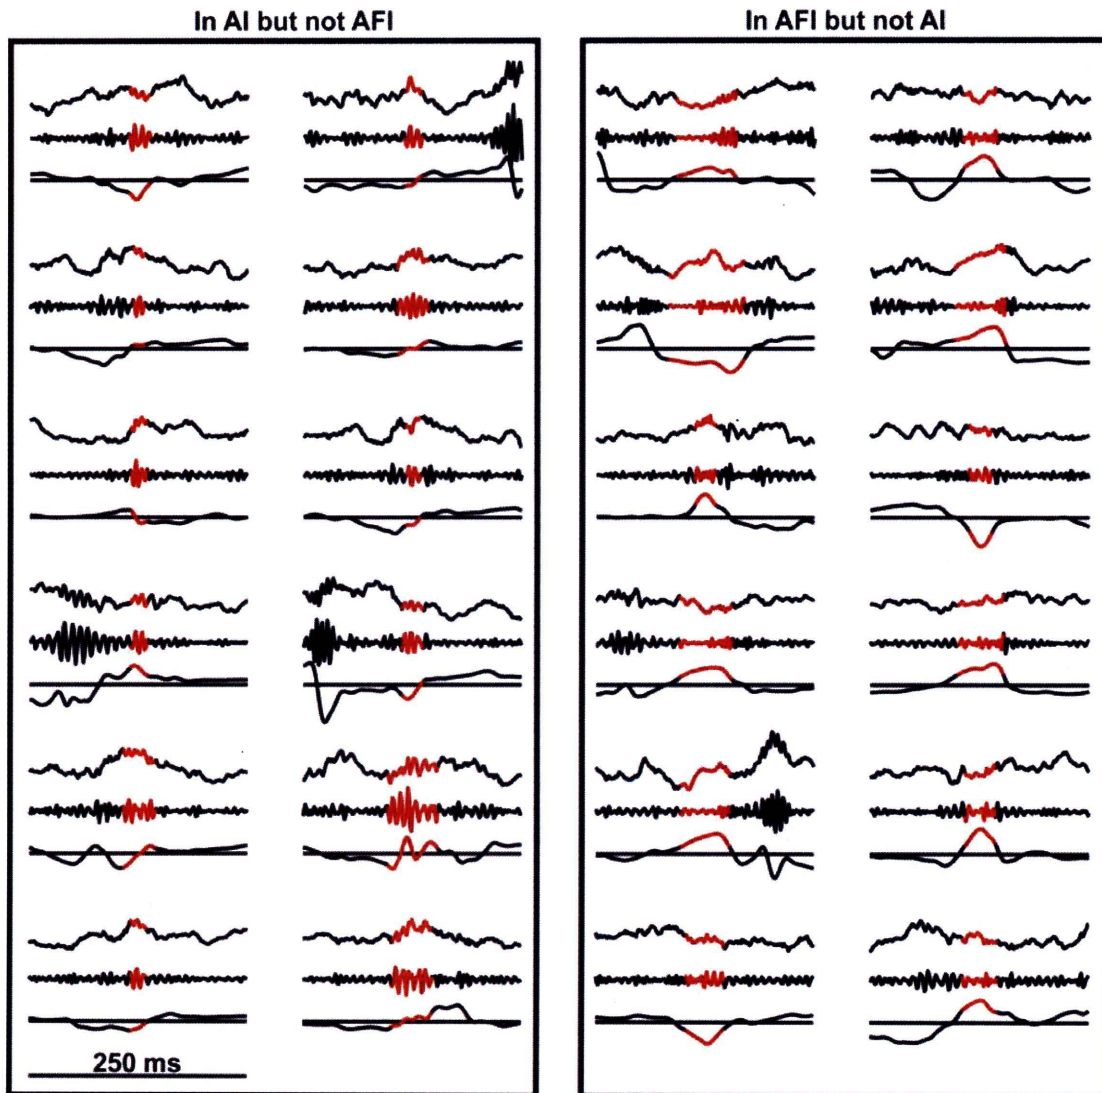ng to Panel (G) of Figure 4.11. Traces of ripple band activity are magnified 3X relative to raw LFP activity. The order of the traces for each example is raw LFP, ripple activity in the 100-250 Hz band, and censored FM. The horizontal line marks the point of zero FM.

The performance of the AFI over AI was measured by taking the events detected by AI to be "true" for the purposes of comparison. Events detected by systematically lowering the threshold for AFI (proposed) were compared against the AI events (true); if a proposed event overlapped with a true event by at least 50%, the proposed event was considered to be successful in identifying a true event. Two evaluation metrics were reported from this analysis: 1) success rate and 2) match rate. The success rate measured the number of proposed events matching true events divided by the total number of true events. The match rate was taken to be the number of true events accounted for by proposed events divided by the total number of true events. Figure 4.11 contains a summary of this analysis. At a threshold of approximately 650 Hz/sec the success rate begins to drop dramatically. At a 1500 Hz/sec threshold, the match rate begins to rise in an exponential fashion and at 650 Hz/sec, approximately 40% of true ripples were detected by using AFI. The ripple amplitude distribution of AFI detected events was more Gaussian distributed than AI detected events Figure 4.11(D,E).

Lets use the results when the AFI threshold was set to 350 Hz/sec as an example for discussion. Of the small number of AFI events that did not match any true events, what could they possibly represent? And of the 60% of true events that were not picked by AFI detection, why were they not detected by AFI? Figure 4.12 contains examples from the mutually exclusive populations. The ripples of Figure 4.12(G), in AFI and not AI, tended to be low amplitude events with large frequency modulation that curiously tended to occur at the onset of ripples but did not encompass the whole ripple itself.

The ripples of Figure 4.12(F), in AF and not AFI, represented medium to high amplitude ripples with below threshold FM. In Figure 4.13, the FM distribution of only the QH+ class was altered significantly. In contrast, the amplitude distributions for all the classes were altered except for the QH+ class. The distribution of ripple length is heavily influenced by the ripple detection signal and threshold. Therefore, the differences in ripple length are taken to be an artifact of the detection method. In Figure 4.7(F), the ripple-triggered average of censored-FM for QH+ and QL+ revealed that the censored-FM was very low in magnitude (Hz/sec) relative to QH- and QL-. Therefore, one would expect that

the events of QH+ and QL+ would not likely be picked up by AFI detection. Figure 4.13(A,C) reveals a 53.5% and 74.1% change in QH+ and QL+ population, while the change in QH- and QL- populations where 49.9% and 55.6%, respectively. This is also consistent with the fact that QH+ and QL+ have lower absolute values of FM relative to QH- and QL- classes (Table 4.1 and Table 4.2).

These analyses demonstrate that ripple detection based on amplitude and frequency information preferentially locates events with high amplitude and frequency modulation; and that a larger percentage of ripple events from the QH+ and QL+ classes are filtered out.
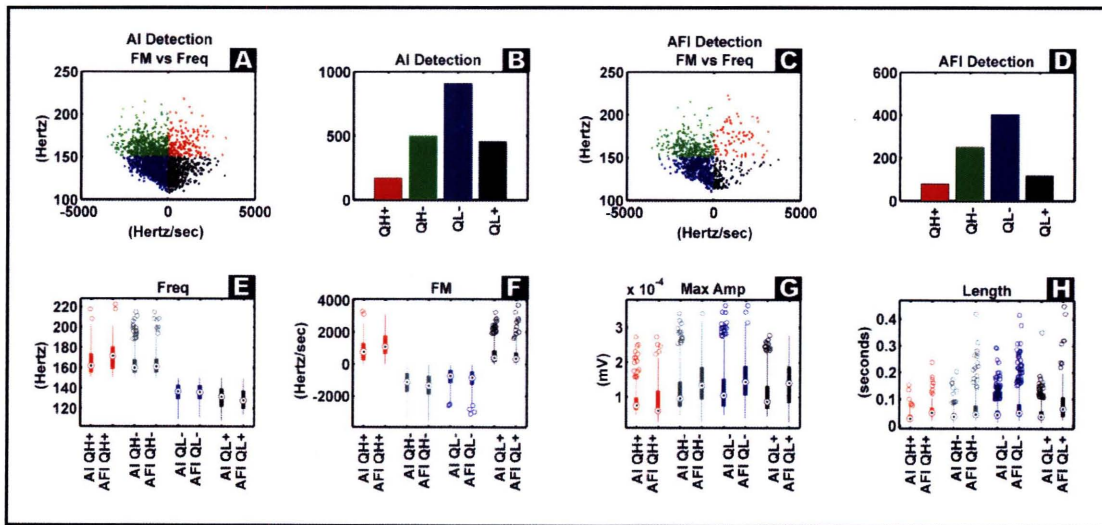
Figure 4.13 : Statistical analysis of ripple detection. (A) Classification of amplitude detected ripple events. (B) Distribution of events per class. (C) Classification of frequency detected ripple events. (D) Distribution of events per class. (E-H) Box-and-whisker plots for different parameters and both detection methods. Distribution of frequency was only different for QH+ class (p < 0.005). Distribution of FM was only different for QH+ class (p < 0.005). Distribution of amplitude was different for QH- (p < 0.001), QL- (p < 1e-7), and QL+ (p < 0.001). Distribution of event lengths were different for all classes (p < 1e-10).

Discussion

Forward and reverse replay during ripples are evidence that the temporal relationship of behaviorally relevant spike sequences are preserved in the hippocampus with high fidelity (Foster & Wilson, 2006). The mechanisms for replay are not well understood and are currently under investigation. This fine timescale analysis of ripple frequency dynamics provides a foundation for exploring questions that would otherwise be intractable. Before this technique, ripples were characterized in terms of amplitude and average frequency (Csicsvari et al., 1999). Based on the analysis in this chapter, it is now clear that frequency is a dynamic quantity within ripple events; both frequency and frequency modulation can vary greatly. Nevertheless, ripple frequency dynamics are stereotypic and the population of ripple events can be accurately described using a small number of ripple definitions or ripple classes.

As the analyses of ripples become more complex there will be a need for classification schemes that capture the significance of ripple-related features. Much like the classification of brain rhythms, having categories in which neural activity can be described allows for consistent comparisons to be made across different analyses and data sets. After inspecting thousands of ripples events, we concluded that a sufficiently robust method for defining ripple classes could use average frequency and average FM information for separation. Amplitude was not used as a feature for classification because LFP properties are generally sensitive to position in the hippocampus as well as electrode properties. Using the two features based on frequency, four ripple classes were defined (Figure 4.5). Visual inspection of ripple events in each class revealed that the classification scheme successfully grouped events that were qualitatively similar. These classes may represent distinct physiological processes and, at the least, are valuable as a framework for systematically exploring the function of ripples and discussing results.

Ripple classes are most meaningful when they reflect a distinct cellular and/or network mechanism of brain function. *In vitro* investigations of epileptiform activity in the hippocampal slice have revealed important contributions from receptor dynamics and

population synchrony in the generation of fast ripple oscillations in the hippocampus (Dzhala & Staley, 2004; Foffani et al., 2007; Spampanato & Mody, 2007). The activation of K^+ channels and ionotropic glutamate receptors were found to be critical for maintaining accurate spike timing and forming coherent field recorded fast ripples (Dzhala & Staley, 2004). Foffani et. al. (2007) concluded that fast ripple generation (200-600 Hz) is likely due to the desynchronization of single unit spike timing leading to a weak correlation in population activity at relatively shorter intervals; this suggests normal (non-epileptic) ripples are likely the result of summed activity from units that have reliable spike timing. Spampanato & Mody (2007) examined the activity of oriens lacunosum-moleculare interneurons (O-LM) and radiatum lacunosum-moleculare (R-LM) interneurons during fast ripples and found that O-LM interneurons were inhibited during ripples and R-LM interneurons were inhibited at the center of the ripples but not at the beginning and end.; their observation of O-LM cell inhibition during ripple events was consistent with *in vivo* observations of O-LM cell activity during ripple events (Klausberger et al., 2003). Klausberger and Somogyi obtained *in vivo* recordings from different interneuron types in the CA1 region of the anesthetized rat and report that during a ripple event, pyramidal and parvalbumin positive basket cells are most likely to fire at the middle of the ripple event; axo-axonic cells are most likely to fire during the first half of the ripple but not the second; bistratisfied cells are likely to fire throughout the ripple but have a preference for the beginning of the ripple event; and O-LM cells are most likely to fire at the beginning of the ripple, can fire at the end of the ripple, but have zero probability of firing near the peak of the ripple (Klausberger et al., 2003; Somogyi & Klausberger, 2005). The diversity of interneurons in terms of spatial domain and temporal firing patterns suggests that inhibition in the hippocampus provides structure and synchrony in a ripple-stage dependent manner. The dynamics of IFQ and FM during ripples may, therefore, be explained using the relative participation of specific interneuron types and pyramidal cells during the start, middle, and end of ripple.

The majority of ripple events contained negative FM over positive FM, therefore, most ripples ended at a lower frequency relative to the frequency at the ripple start. Data from

in vitro and *in vivo* recordings confirm that O-LM and axo-axonic cells are most active at the early stage of each ripple event while pyramidal cells are most active at the center of the ripple event (Klausberger et al., 2003; Maier et al., 2003; Behrens et al., 2005; Spampanato & Mody, 2007). Instantaneous frequency at the start of the ripple would therefore reflect the summed activity of O-LM and other interneurons. Frequency modulation during the ripple is most likely determined by synchronous spike frequency adaptation of individual units in response to both intrinsic receptor dynamics and network properties (Johnston & Wu, 1995). A characteristic of pyramidal cell bursting is with each successive spike the inter-spike interval is more variable than the first and towards the end of the burst the desynchronization of population activity emerges. It follows that the tail of the ripple may likely represent a decrease in synchrony between pyramidal cells and an increase in activity of axo-axonic and O-LM interneurons. The decrease in frequency during the middle of the ripple is most likely due to accommodation of single units as a population (Dzhala & Staley, 2004). Ripple events with positive FM have lower amplitude than ripple events with negative FM but have the same instantaneous frequency distribution (Figure 4.5). Examination of FM+ events, Figure 4.6, reveals a wide variety of possible subclasses. For example, some events resemble short gamma-like oscillations with steady amplitude and very little FM. Other events with positive FM resemble noise-like activity with no clear ripple structure. Based on the preceding argument, low or positive FM may result from input that is strong enough to elicit population synchrony but not strong enough to engage the same mechanisms for sharp-wave ripple complexes (Dzhala & Staley, 2004).

The classification scheme presented here is practical and directly interpretable. Other classification schemes might serve better depending on the question at hand. For example, the temporal relationship between frequency modulation and amplitude modulation could reflect the timing of inputs into the hippocampus or reflect when different circuits are engaged in the 3-stage (start, middle, end) ripple phenomena. In such a scenario, a feature that could be used to classify events is the difference between time of maximum FM and time of maximum amplitude. In addition, classification of ripples based on absolute FM

may serve an important purpose as well. Frequency modulation, whether positive or negative, may reflect an enhanced level of excitability that is a direct consequence of glutamatergic inputs to the hippocampus. Partitioning ripple events in this manner, may allow the investigator to examine only the events that reflect strong input and circuit dynamics within a particular range of timing.

We will now discuss the characterization of the "long ripple," a phenomena that is highly relevant to hippocampal circuit dynamics but has not been characterized in the literature. A long ripple is defined here as an event that is longer than the maximum expected length of a ripple (> 100ms). By extrapolating the expected frequency dynamics for short ripple events, we were able to show that long ripples are most likely caused by the chaining of individual "ripple units" in time. The *censored* frequency estimate showed unimodal frequency dynamics within a ripple event (one positive or negative peak over time). For long ripples, counting the number of peaks in the censored frequency signal yielded an estimate of the number of ripple units in the ripple. A linear regression population analysis of all detected ripple events revealed the expected length of a ripple unit to be approximately 36 ms which is shorter than the expected length of 50-100ms (Chrobak & Buzsaki, 1996). This difference may be caused by a combination of factors relating to ripple detection and ripple unit detection. If the lower threshold which determines the extent of the ripple in time is too high, the average length of the ripple will be smaller. In this analysis, a conservative lower threshold of 0.04mV could likely explain the result. In addition, the ripple unit detection procedure could be detecting false positives, which would increase the number estimated units per ripple event.

In order to investigate the possibility that IFQ and FM variation during ripple events may reflect different physiological processes, the process of ripple detection may be more useful if the process where more selective. By constraining the ripple detection problem using both amplitude and frequency content of ripples, the measure of success would therefore be weighted more heavily on characteristics of the event population. The detection algorithm presented in this chapter makes use of the censored estimate of instantaneous frequency which inherently contains both ripple amplitude and ripple

frequency information. The lowest value for threshold that was tested (200 Hz/sec) accounted for 70% of the events found by amplitude detection alone. The proposed method of detection was shown to preferentially reject ripples of low frequency and positive frequency modulation from all classes. Selectivity was non-uniform across classes. For example, events in the QH+ class were detected based on high FM (Figure 4.13(E-H)). Proposed events based on amplitude and frequency information that did not always match amplitude detected events tended to occur at the onset of ripple events in a region of low amplitude and high frequency modulation. At this time, it is unclear whether or not these events are physiologically meaningful. Therefore, when performing AFI based ripple detection the most optimal choice for choosing the detection threshold is to choose it based on a 100% match rate against AI detection (Figure 4.11). A comparison of amplitude and FM distributions between AI and AFI based ripple classes revealed that AFI based detection preferentially selected ripples of large amplitude and/or large frequency modulation (Figure 4.13).

In conclusion, the advancements made in this chapter directly address current challenges in the analysis of ripple activity. The adaptive frequency estimator(s) offers the opportunity to see structure in ripples that would otherwise be overlooked with current methodologies (Csicsvari et al., 1999; Sirota et al., 2003; Gillis et al., 2005). Ripples can now be classified using information that is independent of ripple amplitude, which can vary depending on electrode properties and position, to provide a vocabulary for comparing results between analyses, subjects, and even different studies. The variation in ripple frequency and frequency modulation suggest that the CA regions of the hippocampus are flexible in their response to synchronous activity, (Figure 4.4). Whether or not the variability in ripple properties is directly related to the nature of the CA1 input or a consequence of some intrinsic state variable is to be determined. The fact that frequency can shift greatly within a ripple event suggests that the ripple is a unit of activity that may be parsed even further. On one hand, long ripples may be decomposed into overlapping ripple units. On the other hand, ripple activity on a cycle-to-cycle basis could be driven on a cycle-to-cycle basis rather than being driven by a continuous burst. The quality of

forward and reverse replay demonstrates that the fidelity of information storage is high in the hippocampus; hence, each ripple could be as rich in information as any other window in time. These relationships between behavioral experience, spike history, and ripple generation are complex and not well understood, however, there is now a better analytical framework for tackling these questions.

Chapter 5 : CONCLUSION

The primary objectives of the thesis was two fold: 1) to better understand the relative contributions of HPC subregions (CA1, CA3, and dentate gyrus) in the process of encoding and consolidation of episodic and spatial memory and 2) to develop, implement, and apply adaptive filtering techniques which have the ability to operate on arbitrary timescales, in order to characterize the dynamic structure of neural data occurring on the true timescale of hippocampal memory processing.

We started by obtaining local field potential recordings and ensemble single-unit recordings from all three subregions of the HPC *in vivo* during spatial and episodic experience. We were able to collect ensemble single-units from CA1, CA3, and DG simultaneously from two rats; and in another two rats, we recorded from either CA1 only or dentate gyrus only. In addition, local field potential recordings were obtained from all the principal cell layers of the respective subregions, however, only recordings from CA1 were included in the analyses presented in the thesis. These data were recorded during food reward foraging on a linear track that expanded in length over the course of the experiment. By comparing activity across old and new track segments, we were able to make conclusions about hippocampal computations in relatively "new" versus relatively "old" spatial locations. Moreover, by examining hippocampal activity in a fixed location, but over time, we were able to make conclusions about hippocampal computations as a function of general experience.

The first step in performing these comparisons was to isolate spike trains representing single-unit activity. We then used an adaptive point process filter for tracking the shape change in place fields over time with sub-second resolution. In contrast to other studies that use mean rate over a large time window to characterize neuronal activity, we were able to extract the temporal evolution of maximum rate as a function of track location and use that to describe the neuron. This allowed for more precise characterization of rate modulation as a function of experience. After many attempts to utilize the maximum rate dynamics to find structure in the data, we found that the best way to represent neuron rate dynamics was to perform linear regressions on the maximum rate signal as a function of

experiment phase and track section. By pooling the regression slopes according to hippocampal recording location, location on the track, phase of the experiment, and running direction, we had essentially computed a compact feature set representing ensemble hippocampal rate modulation over the whole experiment. We then proceeded by performing standard statistical comparisons of the ensemble rate modulations to quantify hippocampal subregion specific differences.

In summary, our main finding in Chapter 2 was that on the initial exposure to a familiar environment, receptive fields in CA3 exhibited a greater magnitude in rate modulation over those of CA1 and DG in the home location on track, but over time these differences in rate modulation converged to similar non-zero values for all three subregions. The pronounced rate modulation in CA3 may be associated to the anatomy of the CA3 recurrent collaterals as well as the converging input from dentate gyrus and entorhinal cortex, and suggests an important role for CA3 in the rapid encoding of context. The convergence in subregion specific rate modulation suggests a constitutively active process of pattern separation that involves all three subregions of HPC. We hypothesize that these ongoing rate updates may reflect an automatic "time integration" process that may work in parallel with the path integration process to provide unique hippocampal outputs over time.

In Chapters 3 and 4, we focus on the analysis of local field potentials. At first, we were interested in using the local field potentials between subregions to make a statement about causality. In the signal processing literature, much focus is currently given to the development of analytical tools that can use local field potentials to infer the direction of information flow. We implemented and applied many of these methods and found that almost all of the causality measures we computed were too difficult to interpret. In the end, we concluded that these measures required significantly more time to develop than we were willing to invest, so we decided to take the work in another direction. Fortunately, we had focused on dynamic estimation of causality and therefore were left with analytic tools for tracking the frequency content in local field potentials on arbitrarily small timescales. Interestingly, when we searched the systems neuroscience literature we found

140

that there were virtually no reports on the quantification and interpretation of frequency modulation of local field potentials.

Chapter 3 of the thesis contains the derivation and some example applications of a Kalman filter for adaptively estimating the parameters of an AR model. As we developed the adaptive filter, we found that previously described methods based on autoregressive filters did not make accurate estimates of instantaneous frequency. Upon investigating the structure of the autoregressive model, we found two issues that led to erroneous measurements. First, large fluctuations in the oscillation amplitude can directly alter the adaptation speed of the Kalman filter, thereby, leading to biased estimates of instantaneous frequency and frequency modulation. Second, we found that having multiple poles in an autoregressive model can lead to interference when placing the poles. We addressed both of these issues by restricting the bandwidth of the signal, limiting the autoregressive model order to 2, and amplitude demodulating the signal prior to frequency estimation. Another significant advancement was the introduction of the frequency information criterion (OIC), which is a data driven criterion for choosing Kalman filter parameters. Together, we have developed a simple, robust, and semi-automatic approach to fine-timescale characterization of brain rhythms. We demonstrate the use and effectiveness of the algorithm by analyzing the theta rhythm during a linear track task. With our algorithm, we found consistent modulation of theta amplitude and frequency as a function of task phase consistent with previously published reports.

With the ability to measure instantaneous frequency and with the goal of understanding the replay phenomena (which tends to co-occur with ripple oscillations), we took the opportunity to develop a framework for capturing as much frequency information from each ripple event as possible. We found that ripple frequency and frequency modulation varied greatly and that ripple events could be classified using their average frequency and average frequency modulation. With the intent to compare ripple statistics over different experimental sessions, we developed a ripple classification scheme that would allow for consistent comparisons to be made across data sets. We also attempted to formalize the procedure for ripple detection since there is currently no standard procedure. Our attempt

combined both amplitude and frequency information through the Kalman estimate of frequency, and resulted in the ability to detect 40%-60% of the ripples found by amplitude thresholding alone. This marked decrease is not taken to be a negative result, but rather an indication that our detection mechanism is selective to events with high frequency modulation and/or high amplitude modulation. In addition, we analyzed "long ripples" with the goal of understanding there composition. Using the average structure for individual ripple events, we found that long ripple events were likely a result of multiple ripples occurring in close temporal proximity.

In summary, our goal to better understand the nature of episodic information encoding and consolidation in the hippocampus was attained in a number of ways. We successfully executed behavioral experiments where *in vivo* recordings of LFP and ensemble single units from CA1, CA3, and DG were recorded. With our adaptive point process filter, we were able to measure subregion specific ensemble rate modulation as a function of experience in time and in space on a continuous basis. We found that even in a very familiar environment, rate modulation was initially high in all subregions, but especially high in CA3. Towards the end of the experiment, the rate modulation decreased significantly from the initial value and the rate modulation between the subregions converged. These results suggest a unique role for CA3 in the initial stages of context familiarization and that the subregions of the HPC act coherently during constitutive episodic encoding in a familiarized environment. We developed a general, easy-to-operate, filter for estimating instantaneous frequency and frequency modulation of brain rhythms. We applied this to the problem of understanding mechanisms for memory consolidation via replay associated events. More specifically, we developed and tested a framework for quantifying fine timescale frequency structure in ripple events. This framework is novel and advances our ability to characterize/classify/detect ripple events, and make comparisons across experiments.

Chapter 6 : REFERENCES

Arnold, M., W. H. R. Miltner, et al. (1998). "Adaptive AR modeling of nonstationary time series by means of Kalman filtering." IEEE Trans Biomed Eng **45**(5): 553-562.

Axmacher, N., F. Mormann, et al. (2006). "Memory formation by neuronal synchronization." Brain Research Reviews **52**(1): 170-182.

Baccala, L. A. and K. Sameshima (2001). "Partial directed coherence: a new concept in neural structure determination." Biological Cybernetics **84**(6): 463-474.

Barbieri, R., L. M. Frank, et al. (2002). "Construction and analysis of non-Gaussian spatial models of neural spiking activity." Neurocomputing **44**: 309-314.

Barbieri, R., L. M. Frank, et al. (2000). "A time-dependent analysis of spatial information encoding in the rat hippocampus." Neurocomputing **32**: 629-635.

Barbieri, R., L. M. Frank, et al. (2001). "Diagnostic methods for statistical models of place cell spiking activity." Neurocomputing **38**: 1087-1093.

Barbieri, R., M. C. Quirk, et al. (2001). "Construction and analysis of non-Poisson stimulus-response models of neural spiking activity." Journal of Neuroscience Methods **105**(1): 25-37.

Bartels, R. H., J. C. Beatty, et al. (1987). An introduction to splines for use in computer graphics and geometric modeling. Los Altos, Calif., M. Kaufmann Publishers.

Bartos, M., I. Vida, et al. (2007). "Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks." Nature Reviews Neuroscience **8**(1): 45-56.

Behrens, C. J., L. P. van den Boom, et al. (2005). "Induction of sharp wave-ripple complexes in vitro and reorganization of hippocampal networks." Nature Neuroscience **8**(11): 1560-1567.

Bencheqroune, A., M. Benseddik, et al. (1999). "Tracking of time-varying frequency of sinusoidal signals." Signal Processing **78**(2): 191-199.

Bi, G. Q. and M. M. Poo (1998). "Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type." Journal of Neuroscience **18**(24): 10464-10472.

Bialek, W., F. Rieke, et al. (1991). "Reading a Neural Code." Science **252**(5014): 1854-1857.

Bland, B. H., S. D. Oddie, et al. (1999). "Mechanisms of neural synchrony in the septohippocampal pathways underlying hippocampal theta generation." Journal of Neuroscience 19(8): 3223-3237.

Bliss, T. V. P. and T. Lomo (1973). "Long-Lasting Potentiation of Synaptic Transmission in Dentate Area of Anesthetized Rabbit Following Stimulation of Perforant Path." Journal of Physiology-London 232(2): 331-356.

Bohlin, T. (1977). "Analysis of EEG signals with changing spectra using a short-word Kalman estimator." Mathematical Biosciences 35(3-4): 221-259.

Box, G. E. P., G. M. Jenkins, et al. (1994). Time series analysis : forecasting and control. Upper Saddle River, N.J., Prentice Hall.

Bragin, A., J. Engel, et al. (1999). "High-frequency oscillations in human brain." Hippocampus 9(2): 137-142.

Brown, E. N., R. Barbieri, et al. (2002). "The time-rescaling theorem and its application to neural spike train data analysis." Neural Computation 14(2): 325-346.

Brown, E. N., L. M. Frank, et al. (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, E. N., D. P. Nguyen, et al. (2001). "An analysis of neural receptive field plasticity by point process adaptive filtering." Proceedings of the National Academy of Sciences of the United States of America 98(21): 12261-12266.

Burgess, N. and J. O'Keefe (1996). "Neuronal computations underlying the firing of place cells and their role in navigation." Hippocampus 6(6): 749-762.

Buzsaki, G. (1989). "Two-stage model of memory trace formation: A role for "noisy" brain states." Neuroscience 31(3): 551-570.

Buzsaki, G. (2002). "Theta oscillations in the hippocampus." Neuron 33(3): 325-340.

Buzsáki, G. (2006). Rhythms of the brain. Oxford ; New York, Oxford University Press.

Buzsaki, G., L. W. S. Leung, et al. (1983). "Cellular Bases of Hippocampal Eeg in the Behaving Rat." Brain Research Reviews 6(2): 139-171.

Chen, Y. H., S. L. Bressler, et al. (2006). "Stochastic modeling of neurobiological time series: Power, coherence, Granger causality, and separation of evoked responses from ongoing activity." Chaos 16(2).

Cheng, S. and L. M. Frank (2008). "New experiences enhance coordinated neural activity in the hippocampus." Neuron 57(2): 303-13.

Chrobak, J. J. and G. Buzsaki (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, J. J. and G. Buzsaki (1996). High-Frequency Oscillations in the Output Networks of the Hippocampal-Entorhinal Axis of the Freely Behaving Rat. 16: 3056-3066.

Clemens, Z., M. Molle, et al. (2007). "Temporal coupling of parahippocampal ripples, sleep spindles and slow oscillations in humans." Brain 130(Pt 11): 2868-78.

Cravens, C. J., N. Vargas-Pinto, et al. (2006). "CA3 NMDA receptors are crucial for rapid and automatic representation of context memory." European Journal of Neuroscience 24(6): 1771-1780.

Csicsvari, J., H. Hirase, et al. (1999). "Fast network oscillations in the hippocampal CA1 region of the behaving rat." J Neurosci 19(16): RC20.

Csicsvari, J., H. Hirase, et al. (1999). "Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat." Journal of Neuroscience 19(1): 274-287.

Csicsvari, J., H. Hirase, et al. (2000). "Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events." Neuron 28(2): 585-594.

Czurko, A., H. Hirase, et al. (1999). "Sustained activation of hippocampal pyramidal cells by 'space clamping' in a running wheel." European Journal of Neuroscience 11(1): 344-352.

Czurko, A., H. Hirase, et al. (1999). Sustained activation of hippocampal pyramidal cells by 'space clamping' in a running wheel. 11: 344-352.

Dan, Y. and M. M. Poo (2004). "Spike timing-dependent plasticity of neural circuits." Neuron 44(1): 23-30.

Davis, S., S. P. Butcher, et al. (1992). "The Nmda Receptor Antagonist D-2-Amino-5-Phosphonopentanoate (D-Ap5) Impairs Spatial-Learning and Ltp Invivo at Intracerebral Concentrations Comparable to Those That Block Ltp Invitro." Journal of Neuroscience 12(1): 21-34.

de Waele, S. and P. M. T. Broersen (2003). "Order selection for vector autoregressive models." Ieee Transactions on Signal Processing 51(2): 427-433.

DeCoteau, W. E., C. Thorn, et al. (2007). "Learning-related coordination of striatal and hippocampal theta rhythms during acquisition of a procedural maze task."

Proceedings of the National Academy of Sciences of the United States of America 104(13): 5644-5649.

Diba, K. and G. Buzsaki (2007). "Forward and reverse hippocampal place-cell sequences during ripples." Nature Neuroscience 10: 1241-1242.

Djuric, P. M. and S. M. Kay (1992). "Order Selection of Autoregressive Models." Ieee Transactions on Signal Processing 40(11): 2829-2833.

Dzhala, V. I. and K. J. Staley (2004). "Mechanisms of fast ripples in the hippocampus." Journal of Neuroscience 24(40): 8896-8906.

Dzhala, V. I. and K. J. Staley (2004). "Mechanisms of fast ripples in the hippocampus." J Neurosci 24(40): 8896-906.

Eden, U. T., L. M. Frank, et al. (2004). "Dynamic analysis of neural encoding by point process adaptive filtering." Neural Computation 16(5): 971-998.

Ego-Stengel, V. and M. A. Wilson (2007). "Spatial selectivity and theta phase precession in CA1 interneurons." Hippocampus 17(2): 161-174.

Ergun, A., R. Barbieri, et al. (2007). "Construction of point process adaptive filter algorithms for neural systems using sequential Monte Carlo methods." Ieee Transactions on Biomedical Engineering 54(3): 419-428.

Foffani, G., A. M. Bianchi, et al. (2004). "Adaptive autoregressive identification with spectral power decomposition for studying movement-related activity in scalp EEG signals and basal ganglia local field potentials." J Neural Eng 1(3): 165-73.

Foffani, G., Y. G. Uzcategui, et al. (2007). "Reduced Spike-Timing Reliability Correlates with the Emergence of Fast Ripples in the Rat Epileptic Hippocampus." Neuron 55(6): 930-941.

Foster, D. J. and M. A. Wilson (2006). "Reverse replay of behavioural sequences in hippocampal place cells during the awake state." Nature 440(7084): 680-683.

Frank, L. M., E. N. Brown, et al. (2006). "Hippocampal and cortical place cell plasticity: implications for episodic memory." Hippocampus 16(9): 775-84.

Frank, L. M., E. N. Brown, et al. (2000). "Trajectory encoding in the hippocampus and entorhinal cortex." Neuron 27(1): 169-178.

Frank, L. M., U. T. Eden, et al. (2002). "Contrasting patterns of receptive field plasticity in the hippocampus and the entorhinal cortex: An adaptive filtering approach." Journal of Neuroscience 22(9): 3817-3830.

Freund, T. F. and G. Buzsaki (1996). "Interneurons of the hippocampus." Hippocampus 6(4): 347-470.

Gersch, W. (1970). "Spectral analysis of EEG's by autoregressive decomposition of time series." Mathematical Biosciences 7(1-2): 205-222.

Gilbert, P. E., R. P. Kesner, et al. (2001). "Dissociating hippocampal subregions: double dissociation between dentate gyrus and CA1." Hippocampus 11(6): 626-36.

Gillis, J. A., W. P. Luk, et al. (2005). "Characterizing in vitro hippocampal ripples using time-frequency analysis." Neurocomputing 65-66: 357-364.

Gold, A. E. and R. P. Kesner (2005). "The role of the CA3 subregion of the dorsal hippocampus in spatial pattern completion in the rat." Hippocampus 15(6): 808-14.

Gothard, K. M., K. L. Hoffman, et al. (2001). "Dentate gyrus and ca1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input." J Neurosci 21(18): 7284-92.

Granger, C. W. J. (1969). "Investigating Causal Relations by Econometric Models and Cross-spectral Methods." Econometrica 37(3): 424-438.

Gray, C. M., P. E. Maldonado, et al. (1995). "Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex." J Neurosci Methods 63(1-2): 43-54.

Guzowski, J. F., J. J. Knierim, et al. (2004). "Ensemble dynamics of hippocampal regions CA3 and CA1." Neuron 44(4): 581-4.

Hafting, T., M. Fyhn, et al. (2008). "Hippocampus-independent phase precession in entorhinal grid cells." Nature.

Hafting, T., M. Fyhn, et al. (2005). "Microstructure of a spatial map in the entorhinal cortex." Nature 436(7052): 801-6.

Hampson, R. E., J. D. Simeral, et al. (1999). "Distribution of spatial and nonspatial information in dorsal hippocampus." Nature 402(6762): 610-614.

Han, Z. S., E. H. Buhl, et al. (1993). "A high degree of spatial selectivity in the axonal and dendritic domains of physiologically identified local-circuit neurons in the dentate gyrus of the rat hippocampus." Eur J Neurosci 5(5): 395-410.

Hargreaves, E. L., D. Yoganarasimha, et al. (2007). "Cohesiveness of spatial and directional representations recorded from neural ensembles in the anterior thalamus, parasubiculum, medial entorhinal cortex, and hippocampus." Hippocampus 17(9): 826-841.

Harris, K. D., D. A. Henze, et al. (2002). "Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells." Nature 417(6890): 738-741.

Haykin, S. (1996). Adaptive Filter Theory. New Jersey, Prentice-Hall.

Henze, D. A., L. Wittner, et al. (2002). "Single granule cells reliably discharge targets in the hippocampal CA3 network in vivo." Nature Neuroscience 5(8): 790-795.

Hesse, W., E. Moller, et al. (2003). "The use of time-variant EEG Granger causality for inspecting directed interdependencies of neural assemblies." Journal of Neuroscience Methods 124(1): 27-44.

Hirase, H., A. Czurko, et al. (1999). "Firing rate and theta-phase coding by hippocampal pyramidal neurons during 'space clamping'." European Journal of Neuroscience 11(12): 4373-4380.

Holscher, C., R. Anwyl, et al. (1997). "Stimulation on the positive phase of hippocampal theta rhythm induces long-term potentiation that can be depotentiated by stimulation on the negative phase in area CA1 in vivo." Journal of Neuroscience 17(16): 6470-6477.

Huerta, P. T. and J. E. Lisman (1995). "Bidirectional Synaptic Plasticity Induced by a Single Burst During Cholinergic Theta-Oscillation in Ca1 in-Vitro." Neuron 15(5): 1053-1063.

Ivanov, V. and L. Kilian (2005). "A practitioner's guide to lag order selection for VAR impulse response analysis." Studies in Nonlinear Dynamics and Econometrics 9(1).

Jeffery, K. J. (2007). "Self-localization and the entorhinal-hippocampal system." Current Opinion in Neurobiology 17(6): 684-691.

Ji, D. Y. and M. A. Wilson (2007). "Coordinated memory replay in the visual cortex and hippocampus during sleep." Nature Neuroscience 10(1): 100-107.

Johnston, D. and S. M.-s. Wu (1995). Foundations of cellular neurophysiology. Cambridge, Mass., MIT Press.

Jonas, P., G. Major, et al. (1993). "Quantal Components of Unitary Epscs at the Mossy Fiber Synapse on Ca3 Pyramidal Cells of Rat Hippocampus." Journal of Physiology-London 472: 615-663.

Jones, M. W. and M. A. Wilson (2005). "Phase precession of medial prefrontal cortical activity relative to the hippocampal theta rhythm." Hippocampus 15(7): 867-73.

Jones, M. W. and M. A. Wilson (2005). "Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task." PLoS Biol 3(12): e402.

Kalman, R. (1960). "A new approach to linear filtering and prediction." Trans. ASM J. Basic Eng. 82.

Klausberger, T., P. J. Magill, et al. (2003). "Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo." Nature 421(6925): 844-848.

Ko, C. C. and C. P. Li (1994). "An Adaptive Iir Structure for the Separation, Enhancement, and Tracking of Multiple Sinusoids." Ieee Transactions on Signal Processing 42(10): 2832-2834.

Kus, R., J. S. Ginter, et al. (2006). "Propagation of EEG activity during finger movement and its imagination." Acta Neurobiologiae Experimentalis 66(3): 195-206.

Kus, R., M. Kaminski, et al. (2004). "Determination of EEG activity propagation: Pairwise versus multichannel estimate." Ieee Transactions on Biomedical Engineering 51(9): 1501-1510.

Lee, A. K. and M. A. Wilson (2002). "Memory of sequential experience in the hippocampus during slow wave sleep." Neuron 36(6): 1183-1194.

Lee, I., M. R. Hunsaker, et al. (2005). "The role of hippocampal subregions in detecting spatial novelty." Behavioral Neuroscience 119(1): 145-153.

Lee, I. and R. P. Kesner (2004). "Differential contributions of dorsal hippocampal subregions to memory acquisition and retrieval in contextual fear-conditioning." Hippocampus 14(3): 301-310.

Lee, I., G. Rao, et al. (2004). "A double dissociation between hippocampal subfields: differential time course of CA3 and CA1 place cells for processing changed environments." Neuron 42(5): 803-15.

Lee, I., D. Yoganarasimha, et al. (2004). "Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3." Nature 430(6998): 456-9.

Leutgeb, J. K., S. Leutgeb, et al. (2007). "Pattern separation in the dentate gyrus and CA3 of the hippocampus." Science 315(5814): 961-966.

Leutgeb, S. and J. K. Leutgeb (2007). "Pattern separation, pattern completion, and new neuronal codes within a continuous CA3 map." Learning & Memory 14(11): 745-757.

Leutgeb, S., J. K. Leutgeb, et al. (2006). "Fast rate coding in hippocampal CA3 cell ensembles." Hippocampus 16(9): 765-774.

Leutgeb, S., J. K. Leutgeb, et al. (2004). "Distinct ensemble codes in hippocampal areas CA3 and CA1." Science 305(5688): 1295-8.

Louie, K. and M. A. Wilson (2001). "Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep." Neuron **29**(1): 145-156.

Lutkepohl, H. (2005). New Introduction to Multiple Time Series Analysis. New York, Springer.

Maier, N., V. Nimmrich, et al. (2003). "Cellular and network mechanisms underlying spontaneous sharp wave-ripple complexes in mouse hippocampal slices." Journal of Physiology-London **550**(3): 873-887.

Mainardi, L. T., A. M. Bianchi, et al. (1995). "Pole-Tracking Algorithms for the Extraction of Time-Variant Heart-Rate-Variability Spectral Parameters." Ieee Transactions on Biomedical Engineering **42**(3): 250-259.

Manns, J. R. and H. Eichenbaum (2006). "Evolution of declarative memory." Hippocampus **16**(9): 795-808.

Markus, E. J., C. A. Barnes, et al. (1994). "Spatial Information-Content and Reliability of Hippocampal Ca1 Neurons - Effects of Visual Input." Hippocampus **4**(4): 410-421.

Marr, D. (1969). "A Theory of Cerebellar Cortex." Journal of Physiology (Cambridge) **202**(2): 437-470.

Marr, D. (1971). "Simple Memory - Theory for Archicortex." Philosophical Transactions of the Royal Society of London Series B-Biological Sciences **262**(841): 23-&.

McHugh, T. J., K. I. Blum, et al. (1996). "Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice." Cell **87**(7): 1339-1349.

McHugh, T. J., M. W. Jones, et al. (2007). "Dentate gyrus NMDA receptors mediate rapid pattern separation in the hippocampal network." Science **317**(5834): 94-99.

McHugh, T. J. and S. Tonegawa (2007). "Spatial exploration is required for the formation of contextual fear memory." Behavioral Neuroscience **121**(2): 335-339.

McNaughton, B. L. and R. G. M. Morris (1987). "Hippocampal Synaptic Enhancement and Information-Storage within a Distributed Memory System." Trends in Neurosciences **10**(10): 408-415.

Mehta, M. R., A. K. Lee, et al. (2002). "Role of experience and oscillations in transforming a rate code into a temporal code." Nature **417**(6890): 741-746.

Mehta, M. R. and M. A. Wilson (2000). "From hippocampus to V1: Effect of LTP on spatio-temporal dynamics of receptive fields." Neurocomputing **32**: 905-911.

Morris, R. G. M. (2006). "Elements of a neurobiological theory of hippocampal function: the role of synaptic plasticity, synaptic tagging and schemas." European Journal of Neuroscience 23(11): 2829-2846.

Morris, R. G. M., P. Garrud, et al. (1982). "Place Navigation Impaired in Rats with Hippocampal-Lesions." Nature 297(5868): 681-683.

Moser, E. I. and M. B. Moser (2003). "One-shot memory in hippocampal CA3 networks." Neuron 38(2): 147-8.

Muller, R. U. and J. L. Kubie (1987). "The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells." J Neurosci 7(7): 1951-68.

Muller, R. U., J. L. Kubie, et al. (1987). "Spatial firing patterns of hippocampal complex-spike cells in a fixed environment." J Neurosci 7(7): 1935-50.

Nakazawa, K., T. J. McHugh, et al. (2004). "NMDA receptors, place cells and hippocampal spatial memory." Nature Reviews Neuroscience 5(5): 361-372.

Nakazawa, K., M. C. Quirk, et al. (2002). "Requirement for hippocampal CA3 NMDA receptors in associative memory recall." Science 297(5579): 211-8.

Nakazawa, K., M. C. Quirk, et al. (2001). "Place cell analysis of CA3 NMDA receptor knockout mice." Society for Neuroscience Abstracts 27(2): 1975.

Nakazawa, K., L. D. Sun, et al. (2003). "Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience." Neuron 38(2): 305-15.

Neves, G., S. F. Cooke, et al. (2008). "Synaptic plasticity, memory and the hippocampus: a neural network approach to causality." Nature Reviews Neuroscience 9: 65-75.

Nicoll, R. A. and D. Schmitz (2005). "Synaptic plasticity at hippocampal mossy fibre synapses." Nature Reviews Neuroscience 6(11): 863-876.

O'Keefe, J. (1976). "Place Units in Hippocampus of Freely Moving Rat." Experimental Neurology 51(1): 78-109.

O'Keefe, J. and M. L. Recce (1993). "Phase relationship between hippocampal place units and the EEG theta rhythm." Hippocampus 3(3): 317-30.

O'Neill, J., T. Senior, et al. (2006). "Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior." Neuron 49(1): 143-155.

Okeefe, J. (1976). "Place Units in Hippocampus of Freely Moving Rat." Experimental Neurology 51(1): 78-109.

Okeefe, J. and M. L. Recce (1993). "Phase Relationship between Hippocampal Place Units and the Eeg Theta-Rhythm." Hippocampus 3(3): 317-330.

Oppenheim, A. V., R. W. Schafer, et al. (1999). Discrete-time signal processing. Upper Saddle River, N.J., Prentice Hall.

Palva, S. and J. M. Palva (2007). "New vistas for alpha-frequency band oscillations." Trends in Neurosciences 30(4): 150-158.

Percival, D. B. and A. T. Walden (1993). Spectral analysis for physical applications : multitaper and conventional univariate techniques. Cambridge ; New York, NY, USA, Cambridge University Press.

Quirk, G. J., R. U. Muller, et al. (1990). "The Firing of Hippocampal Place Cells in the Dark Depends on the Rats Recent Experience." Journal of Neuroscience 10(6): 2008-2017.

Remondes, M. and E. M. Schuman (2003). "Molecular mechanisms contributing to long-lasting synaptic plasticity at the temporoammonic-CA1 synapse." Learning & Memory 10(4): 247-252.

Remondes, M. and E. M. Schuman (2004). "Role for a cortical input to hippocampal area CA1 in the consolidation of a long-term memory." Nature 431(7009): 699-703.

Rolls, E. T. (1996). "A theory of hippocampal function in memory." Hippocampus 6(6): 601-620.

Rolls, E. T. and R. P. Kesner (2006). "A computational theory of hippocampal function, and empirical tests of the theory." Progress in Neurobiology 79(1): 1-48.

Rosenberg, J. R., D. M. Halliday, et al. (1998). "Identification of patterns of neuronal connectivity--partial spectra, partial coherence, and neuronal interactions." Journal of Neuroscience Methods 83(1): 57-72.

Rotenberg, A. and R. U. Muller (1997). "Variable place-cell coupling to a continuously viewed stimulus: evidence that the hippocampus acts as a perceptual system." Philosophical Transactions of the Royal Society of London Series B-Biological Sciences 352(1360): 1505-1513.

Salat, D. H., A. J. W. van der Kouwe, et al. (2006). "Neuroimaging HM: A 10-year follow-up examination." Hippocampus 16(11): 936-945.

Sameshima, K. and L. A. Baccala (1999). "Using partial directed coherence to describe neuronal ensemble interactions." Journal of Neuroscience Methods 94(1): 93-103.

Sargolini, F., M. Fyhn, et al. (2006). "Conjunctive representation of position, direction, and velocity in entorhinal cortex." Science 312(5774): 758-762.

Schelter, B., M. Winterhalder, et al. (2006). "Testing for directed influences among neural signals using partial directed coherence." Journal of Neuroscience Methods 152(1-2): 210-219.

Scoville, W. B. and B. Milner (1957). "Loss of Recent Memory after Bilateral Hippocampal Lesions." Journal of Neurology Neurosurgery and Psychiatry 20(1): 11-21.

Siapas, A. G. and M. A. Wilson (1998). "Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep." Neuron 21(5): 1123-1128.

Sinnamon, H. M. (2006). "Decline in hippocampal theta activity during cessation of locomotor approach sequences: Amplitude leads frequency and relates to instrumental behavior." Neuroscience 140(3): 779-790.

Sirota, A., J. Csicsvari, et al. (2003). "Communication between neocortex and hippocampus during sleep in rodents." Proceedings of the National Academy of Sciences of the United States of America 100(4): 2065-2069.

Skaggs, W. E., B. L. McNaughton, et al. (1996). "Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences." Hippocampus 6(2): 149-72.

Slawinska, U. and S. Kasicki (1998). "The frequency of rat's hippocampal theta rhythm is related to the speed of locomotion." Brain Research 796(1-2): 327-331.

Somogyi, P. and T. Klausberger (2005). "Defined types of cortical interneurone structure space and spike timing in the hippocampus." Journal of Physiology-London 562(1): 9-26.

Spampanato, J. and I. Mody (2007). "Spike timing of lacunosom-moleculare targeting interneurons and CA3 pyramidal cells during high-frequency network oscillations in vitro." Journal of Neurophysiology 98(1): 96-104.

Squire, L. R. and E. Lindenlaub (1990). The Biology of memory : Symposium Bernried, Germany, October 15th-19th, 1989. Stuttgart ; New York, F.K. Schattauer.

Squire, L. R., C. E. Stark, et al. (2004). "The medial temporal lobe." Annu Rev Neurosci 27: 279-306.

Squire, L. R. and S. Zola-Morgan (1991). "The Medial Temporal-Lobe Memory System." Science 253(5026): 1380-1386.

Tabuchi, E., A. B. Mulder, et al. (2003). "Reward value invariant place responses and reward site associated activity in hippocampal neurons of behaving rats." Hippocampus 13(1): 117-32.

Tarvainen, M. P., J. K. Hilturien, et al. (2004). "Estimation of nonstationary EEG with Kalman smoother approach: An application to event-related synchronization (ERS)." Ieee Transactions on Biomedical Engineering 51(3): 516-524.

Taube, J. S. (2007). "The head direction signal: Origins and sensory-motor integration." Annual Review of Neuroscience 30: 181-207.

Treves, A. and E. T. Rolls (1991). "What Determines the Capacity of Autoassociative Memories in the Brain." Network-Computation in Neural Systems 2(4): 371-397.

Tsien, J. Z., P. T. Huerta, et al. (1996). "The essential role of hippocampal CA1 NMDA receptor-dependent synaptic plasticity in spatial memory." Cell 87(7): 1327-1338.

Vanderwolf, C. H. (1969). "Hippocampal Electrical Activity and Voluntary Movement in Rat." Electroencephalography and Clinical Neurophysiology(26): 407-418.

Wilson, M. A. and B. L. McNaughton (1993). "Dynamics of the Hippocampal Ensemble Code for Space." Science 261(5124): 1055-1058.

Wilson, M. A. and B. L. McNaughton (1994). "Reactivation of Hippocampal Ensemble Memories During Sleep." Science 265(5172): 676-679.

Witter, M. P. (2006). "Connections of the subiculum of the rat: Topography in relation to columnar and laminar organization." Behavioural Brain Research 174(2): 251-264.

Witter, M. P., H. J. Groenewegen, et al. (1989). "Functional organization of the extrinsic and intrinsic circuitry of the parahippocampal region." Progress in Neurobiology 33(3): 161-253.

Witter, M. P. and F. G. Wouterlood (2002). The parahippocampal region : organization and role in cognitive function. Oxford ; New York, Oxford University Press.

Wyble, B. P., J. M. Hyman, et al. (2004). "Analysis of theta power in hippocampal EEG during bar pressing and running behavior in rats during distinct behavioral contexts." Hippocampus 14(5): 662-674.

Ylinen, A., A. Bragin, et al. (1995). "Sharp Wave-Associated High-Frequency Oscillation (200-Hz) in the Intact Hippocampus - Network and Intracellular Mechanisms." Journal of Neuroscience 15(1): 30-46.

Zalutsky, R. A. and R. A. Nicoll (1990). "Comparison of 2 Forms of Long-Term Potentiation in Single Hippocampal-Neurons." Science 248(4963): 1619-1624.

Zetterberg, L. H. (1969). "Estimation of parameters for a linear difference equation with application to EEG analysis." Mathematical Biosciences 5(3-4): 227-275.

Zetterberg, L. H. (1977). "Means and methods for processing of physiological signals with emphasis on EEG analysis." Adv Biol Med Phys 16: 41-91.

Zhang, K. C., I. Ginzburg, et al. (1998). "Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells." Journal of Neurophysiology 79(2): 1017-1044.

Zola-Morgan, S., L. R. Squire, et al. (1986). "Human Amnesia and the Medial Temporal Region - Enduring Memory Impairment Following a Bilateral Lesion Limited to Field Ca1 of the Hippocampus." Journal of Neuroscience 6(10): 2950-2967.