The Neural Basis of Trajectory Computations in Rodent Posterior Parietal Cortex and Hippocampus

by

Mark Roger Bower

A Dissertation Submitted to the Faculty of the

GRADUATE INTERDISCIPLINARY PROGRAM IN NEUROSCIENCE

In Partial Fulfillment of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College of

THE UNIVERSITY OF ARIZONA

2003

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

THE UNIVERSITY OF ARIZONA ®
GRADUATE COLLEGE

As members of the Final Examination Committee, we certify that we have read the dissertation prepared by Mark Roger Bower entitled The neural basis of trajectory computations in rodent posterior parietal cortex and hippocampus

and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

Bruce L. McNaughton	16 Apr, 2003 Date
Carol A. Barnes	April 16, 2003 Date
Lynn Nadel	4/21/03 Date
Timothy W. Secomb	4/16/03 Date
Robert S. Sloviter	4/16/03 Date

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copy of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Dissertation Director Bruce L. McNaughton — 4/25/03 Date

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this thesis are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: 

ACKNOWLEDGEMENTS

The author would like to thank his parents, Jane and Roger Bower, for their unfailing support; Erik Ellsworth and Nathalie Gebara, two exceptional undergraduate assistants; Dr. Bruce McNaughton, my advisor, for creative suggestions when experiments did not produce results, and energetic criticism when they did; Dr. Arthur T. Winfree, for teaching me the importance of paradoxes; Dr. Partha Mitra and the staff of the Workshop on the Analysis of Neural Data (WAND), 1998 at the Marine Biological Laboratories in Woods Hole, MA for the original motivation to look at time-localized spectra of EEG and its relation to behavior; Dr. Jason Gerrard, for keeping me aware of life outside the lab; Dr. Ephron Rosenzweig, for lively discussions inside the lab; the congregation of Fountain of Life Lutheran church, for their kindness and love; Dr. David Euston, Nathan Insel, Kim Bohne, and Ryan Roop for assistance during data collection; and Regina Benton, for her unfailing love and encouragement to see this task through to completion.

DEDICATION

This dissertation is dedicated to the memories of two incredible individuals who changed my thoughts, my career and my life. Dr. Edmund A. Arbas believed in me when others only wondered, and provided a clear direction during a time of indecision. He also made me laugh. Dr. Arthur T. Winfree fired my imagination, and inspired a love of science that burns brighter within me each day. He also made me think. Both were taken from my life too soon. Neither will be forgotten.

"Then he said to Thomas, 'Put your finger here; see My hands. Reach out your hand and put it into My side. Stop doubting and believe.' " John 20:27

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

ABSTRACT

Space is a fundamental property of nature, so it is not surprising that the processing of spatial information involves many brain structures, including the posterior parietal cortex (PPC) and hippocampus. The PPC represents body position in multiple frames of reference, aligning the reference frames of different parts of the body to cooperate in the performance of a task. Hippocampal "place cells" increase their firing rates in specific locations, and place cell responses can be based on information encoded in multiple reference frames, including external, sensory-based reference frames or internal, memory-based frames. The computation of whole-body trajectories is common to all navigation tasks, but little is known about how this computation is performed. Rats were trained to run to distant targets that were presented either randomly or as segments of sequences. When rats were trained to make direct trajectories to distant targets that were presented in random order, no evidence was found that hippocampal place cells encoded distant goals. When rats learned a sequence of goals that contained a repeated segment, hippocampal place cell activity along the repeated segment remained the same. This showed that differential hippocampal codes were not required for rats to differentiate overlapping sequential contexts. Differential hippocampal codes could be formed, however, if rats learned sequences of goals under specific conditions, reflecting differential activity from neural structures outside the hippocampus. At the initiation of trajectories, the phase of hippocampal theta was reset at the peak acceleration of the rat. The activity of some parietal neurons was modulated by hippocampal theta, though the

magnitude of modulation was not as great as that for hippocampal units, and the preferred firing phase of parietal units differed from that of theta-modulated place cells. Two classes of units in PPC responded differentially during the early and late stages of trajectories: one class with an increased firing rate during the early stages of trajectories, and the other with an increased firing rate during the late stages. These results provide a foundation for future studies of parieto-hippocampal interactions during trajectory planning.

Chapter 1 INTRODUCTION TO THE DISSERTATION

This dissertation contains the results of experiments that were designed to explore what occurs in the parietal cortex and hippocampus when an animal makes a trajectory to a distant goal. Because this dissertation includes results from both parietal cortex and hippocampus, the review chapter highlights only those aspects of the literature that are fundamental to an understanding of the results presented in this dissertation. The task throughout these experiments remained constant; rats computed and then made uninterrupted trajectories to distant targets. The data obtained during these trajectories came from two structures (specifically, the posterior parietal cortex and hippocampus) and from multiple sources (e.g., multiple single units, EEG, and video tracking). To simplify subsequent chapters, the chapter on methods describes the terminology, tasks and data processing common to all experiments in this dissertation, while the methods section of each chapter adds chapter-specific experimental protocols and data analysis methods. Results from hippocampal studies are presented first, followed by results from parietal studies. The discussion of these results is combined in the final chapter.

Chapter 2 LITERATURE REVIEW

Neural Representations of Space

What must occur for the brain to compute a movement from a start to a goal? A goal must be selected and a neural correlate for that goal must be established somewhere in the brain, along with a neural correlate of the current location so that the two may be used in the computation of a trajectory from the start to the goal. To compute a trajectory between these points, these representations must be active either in one brain structure (either at the same time or in a sequence), or in multiple structures that are in communication (McNaughton et al., 1994a). Finally, the trajectory must be computed and represented in the neural activity of some structure in the brain, even if that representation is the implementation of the motor activity involved in the movement itself. In addition, representations of the current location, goal and trajectory must exist within a coordinate frame. A coordinate frame is described normally by orthogonal axes that intersect at a unique origin. The physical entities to which these axes are fixed (i.e., with which the axes move) describe the coordinate frame. Coordinate frames for neural activity can be grouped into two classes: "egocentric" coordinate frames are dependent on the location or orientation of the observer (e.g., head-centered, body-centered), while "allocentric" coordinate frames are external to the observer (e.g., room-centered). The description of a given space varies with the coordinate frame that is used, and the same physical space can be described in different ways, simultaneously. Different brain regions could encode the same variables in different reference frames; there is no single,

preferred definition of "space" for the entire brain. In this dissertation, the term "spatial" refers to the general class of variables placed on axes whose dimension is length, and does not refer to any specific type of coordinate frame ("egocentric" or "allocentric") to the exclusion of all others. "Spatial information" refers to the ability to localize a variable in a spatial reference frame. "Spatial computation" refers to mathematical operations within a spatial reference frame that manipulate spatial information. The most direct example is the computation of the distance and direction required to move from a start to a goal location, which are all referenced to the same coordinate frame. Finally, a "spatial task" includes any problem that can be solved by spatial computations using the spatial arrangement of distant cues (e.g., "Move to the edge of the arena by running towards a point equidistant between those two, distant cue cards.")

To identify brain regions that could be involved with spatial computations, the flow of spatial information through the brain must first be discussed, focusing on two structures, parietal cortex and the hippocampus. Then, various methods by which biologically plausible networks of neurons could encode spatial information and compute trajectories are reviewed, leading to a discussion on vector computation in neural networks. Finally, neuroanatomy is related to physiology of both parietal cortex and the hippocampus to consider how trajectory computations might occur.

Anatomy: how does spatial information flow?

Cortical pathways

The question of how brains encode space encompasses several different definitions of space, including visual space, personal space for the movement of limbs, and navigational space for the movement of the whole body. One seminal hypothesis by Ungerleider and Mishkin (1984) proposed that visual processing divides into two "streams" in primates: a ventral stream, which recognizes objects (the "what" stream), and a dorsal stream, which process the position of those objects in space (the "where" stream). The most abstract level of processing in these two streams is the inferior temporal cortex in the "what" stream, and the posterior parietal cortex (PPC) in the "where" stream (Figure 2-1), and both streams deliver information to the hippocampus.

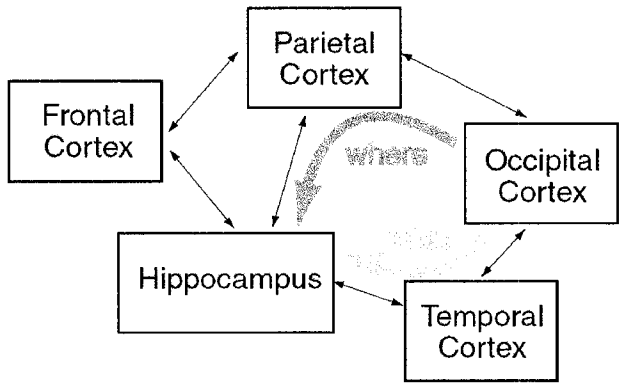

Figure 2-1. Schematic connectivity of mammalian cortex labeled with Ungerleider and Mishkin's visual streams. Spatial information describing the "where" of an object (dark gray) flows from occipital (visual) cortex to parietal cortex, while recognition information describing the "what" of an object (light gray) flows through temporal cortex.

A similar system might apply to spatial processing in rodents, suggesting that the two streams reunite in the hippocampus for the purpose of spatial navigation (McNaughton et al., 1989). Though there are differences between primate and rodent visual and spatial

processing systems (Goodale and Milner, 1982), recent work suggests that neuroanatomical similarities do exist (Insausti et al., 1997;Burwell and Amaral, 1998). Combining the theories of Ungerleider and Mishkin (1984) and McNaughton et al. (1989) suggests that PPC and hippocampus should contain trajectory-related activity in all movements, including directed, ballistic movements of the whole body to distant targets. It is reasonable to suppose, then, that trajectory computations either occur in one or both of these structures, or that the results of such computations occurring in other structures are reflected in them.

Human imaging

Support for such a theory can be found in imaging studies in humans performing spatial tasks. Spatial tasks are those whose solution requires the use of either the relative location of objects (e.g., knowing that an object is located between two landmarks), path integration (e.g., counting steps from a landmark in a specific direction), or some combination of the two. In a sense, a spatial task involves a representation of the ordering of cues or symbols, in addition to the representation of the cues themselves (Marr, 1970). Positron Emission Tomographic (PET) imaging shows that regional cerebral blood flow (rCBF) to the PPC increases when subjects adapt to distortion-inducing prisms, while visually tracking objects moving across a video screen (i.e., noting the relative change in position between a screen object and the screen edge) (Clower et al., 1996). Increased hippocampal rCBF in London taxi cab drivers was observed as they described how to get from one famous landmark to another, first by noting major landmarks along the way (i.e., relative locations), and then by noting

specific streets and turns required (Maguire et al., 1997). Increases were observed in rCBF in the hippocampus, the PPC and the basal ganglia of subjects asked to move through a virtual environment. In fact, hippocampal rCBF increased with increasing accuracy in the task (Maguire et al., 1998). Increased human rCBF during spatial tasks suggests involvement of PPC and hippocampus in spatial processing, but does not elucidate the mechanisms used by these structures in spatial computations.

Physiology: how is spatial information stored and read?

If either the hippocampus or the PPC is involved in trajectory computation, how might distant goals be represented, and how might trajectories be computed in neural structures? Georgopoulos et al. (1982) recorded from single units in non-human primate motor cortex during a reaching task, and found that within 60-80 msec of target onset, individual neurons responded selectively for movements in specific directions (Figure 2-2). Elevated firing rates were observed over a broad angle (showing increased firing rates over approximately 135°), and cells were found with optimal responses for movement in all directions. Georgopolous et al. suggested that, instead of each movement direction being associated with a specific subset of neurons, movement direction might be encoded as a function of a population of neurons with each neuron responding over a broad range of directions; i.e., with a broad "tuning curve". Broad tuning curves at the level of individual units could be summed to produce tightly focused tuning curves at the population level.

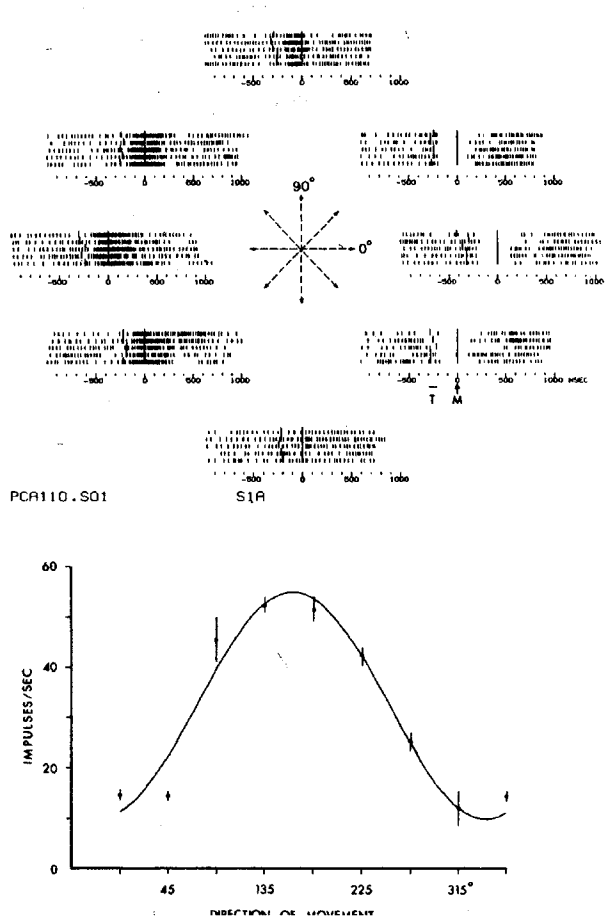

Figure 2-2. Single unit activity in prefrontal cortex correlates with the direction of arm movements over a broad range of directions. Movements were made in a two-dimensional plane from a fixed starting point to one of eight targets. (TOP) Raster plots for five trajectories to each target location. Each hash mark denotes a spike from a single neuron. Movement onset (judged as the time when arm velocity crossed a threshold) is marked by a vertical black bar. Note that firing rate increases just prior to the onset of movement, but only in the four directions from 90 to 225 degrees. Some inhibition of firing is present around the time of movement, especially for 0 and 315 degrees. (BOTTOM) Tuning curve for firing rate as a function of movement direction. Mean firing rates are plotted and fitted with a sinusoidal function. Again, the directions showing increased firing rate are seen to be quite broad, covering at least 135 degrees of movement. (Adapted from Georgopolous et al., 1982).

Encoding schemes

Understanding the relationship between neural activity and behavior is one of the central goals of neuroscience. Neural activity that is strongly tied to a behaviorally-relevant stimulus is called a neural correlate of that stimulus, and the search for neural correlates can be described as an attempt to "decode" patterns of neural activity. Often, this takes the form of asking, "What does a particular pattern of neural activity represent?" This question arises from a long history of studying neural activity in terms of input-output relationships, looking for the optimal input to generate the maximum output. Decoding, then, often takes the form of a response curve, which relates the firing rate of a unit (the "output") to some input variable (e.g., stimulus intensity, location in a maze, heading, etc.). Two general ways in which stimuli could be encoded in neural activity include variable encoding (Figure 2-3.A), in which a single variable is encoded as a monotonically increasing function of some input parameter, and value encoding (Figure 2-3.B), in which the response is maximized for a specific range of values (Ballard, 1986).

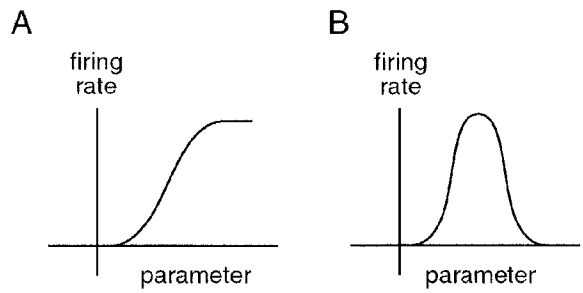

Figure 2-3. Basic neural response curves for encoding schemes. A) Variable encoding: the firing rate increases monotonically as a function of the stimulus input. B) Value encoding: the firing rate increases over a broad range of input values and peaks at one specific value; e.g., the directional response of units in motor cortex (Figure 2-2) (Modified from Ballard, 1986).

Reference frames

In the form of a mathematical function, response curves require an origin and coordinate axes. The biological equivalent is called a "frame of reference", that aspect of the body or environment to which neural responses are aligned, or centered. Reference frames can be viewed as a hierarchy, beginning with sensory organs at the lowest level (e.g., the eye, called a "retinotopic" reference frame, or the body surface, called a "somatotopic" reference frame), and then progress to successively higher frames of reference (e.g., from eye-centered to head-centered to body-centered). At the highest level of this hierarchy, a distinction can be made between responses that are body-centered, or "egocentric", which are centered on and moving with some of aspect of the self, and world-centered, or "allocentric", which are centered on the world and not dependent on location within an environment.

Linear associative mapping

Unless the current location, goal, and the trajectory between them are encoded in a single reference frame, some means must exist to transform one or more reference frames into another. One physiologically plausible mechanism for transforming one reference frame into another is through the use of associative maps (Kohonen, 1984). Associative maps amount to look-up tables that can be formed entirely by correlative mechanisms (i.e., on two or more events occurring at roughly the same time or in the same place) that are entirely local (i.e., not dependent upon an external "teacher" signal that computes error for all elements of a network). Linear associative maps are a variant that adds one constraint, namely that for a given input, I, the output, O, solves the linear

equation $I \times W = O$, where W is a matrix of potentially modifiable weights. Combining reference frames in this way makes new reference frames, performing a coordinate transformation in the process.

For example, given the current heading and location of the animal, P, and a stored memory of the goal, M, a trajectory, T, could be computed by subtracting the components of the goal vector from those of the present location, $T = M - P$. Neurons could perform this computation through the use of linear associative mapping (McNaughton et al., 1994a) (Figure 2-4).

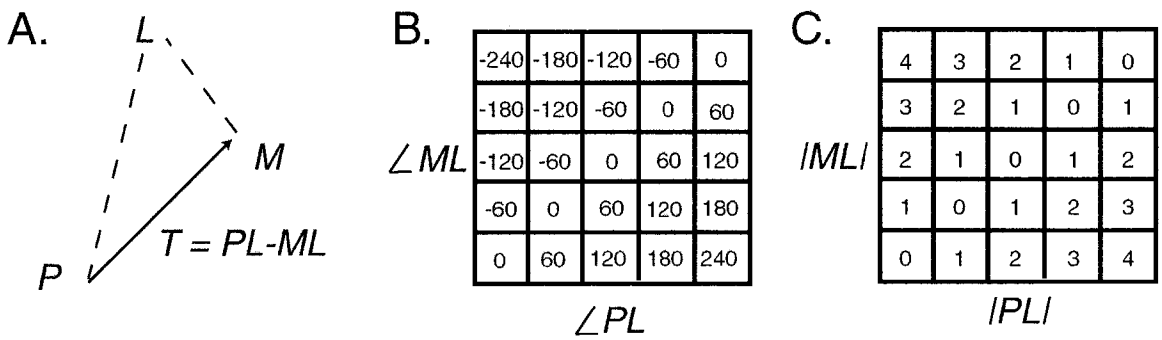

Figure 2-4. Trajectory computation through vector subtraction. (A) Consider a vector representation of present location, P, and the memory of a goal location, M, both encoded with respect to a landmark, L. The trajectory from P to M could be computed utilizing linear associative maps by separating direction and distance computations. (B) First, the representation of both present heading angle, ∠PL, and the heading recalled from memory, ∠ML, provide enough information to allow the computation of the required heading change. (C) Given proper heading, the distance is uniquely specified (McNaughton et al., 1994a).

Gain fields

How can the static synaptic weight matrices of a network of biological neurons compute the dynamic variables required for the broad array of movements made by animals? One possibility combines Ballard's response classes, multiplying the output of a value-encoding unit by the output of a variable-encoding unit to form a response called a

gain field. The first observation of gain fields occurred during studies of the response of light-sensitive units in PPC (Andersen and Mountcastle, 1983). Mountcastle had previously observed (Lynch et al., 1977) that parietal units encoded target locations as a function of the location of the image on the retina (i.e., to the left or right of the fovea). Re-phrased into Ballard's terms, the image generated a value-encoded response in retinotopic coordinates. Andersen and Mountcastle found that these units were also sensitive to the direction of gaze; i.e., to the angle of the eyes in the orbit, or a variable-encoded response in head-centered coordinates. Gaze direction and the tuned response of retinal location could then be combined by downstream units to shift parietal receptive fields in response to changes in gaze direction.

As an example, consider reaching for a coffee cup sitting directly in front of you, while you read the morning paper (Figure 2-5). As your eyes move left to right while reading, the image of the cup on your retina changes position, though the location of the cup is fixed relative to your body. While the brain knows that it is your eyes that are moving and not the cup, because of efference copy from the eye muscles, how does the brain maintain the correct angle to the cup relative to your body for reaching? Scaling the receptive fields of reach-direction units as a function of gaze direction shifts the peak of those fields, providing the invariant response required to maintain the reach direction in front of you, while your eyes change position relative to the cup.

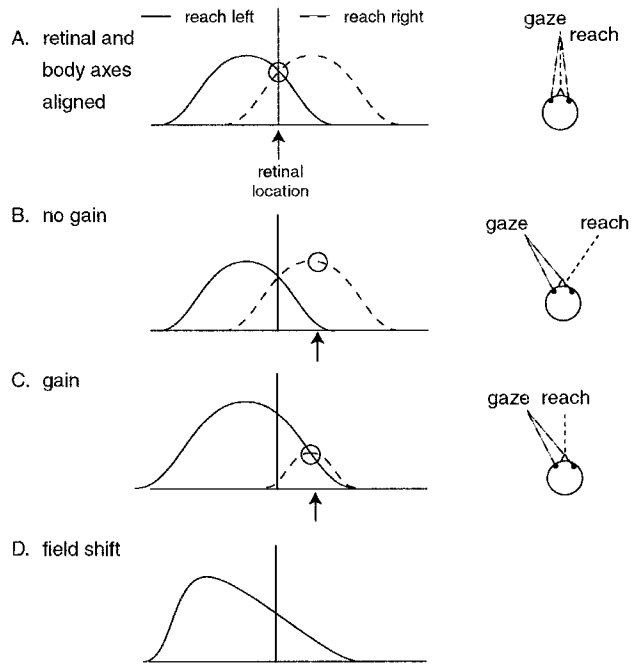

Figure 2-5. Gain fields transform receptive field reference frames. A) The response curves for two, hypothetical reach-related units, one for reaches to the left (solid line) and the other for reaches to the right (dotted line). When gaze and retinal frames are aligned, the image of an object directly in front of the observer falls on the center of the retina, causing equal drive between the two units, and a subsequent reach directly ahead. B) When the eyes gaze to the left, retinal and head reference frames become misaligned. An object in front of the observer now appears in the right field of vision, causing a greater response in "reach right" unit than in the "reach left" unit, and a subsequent error in the direction of reach. C) Gain fields compensate for gaze direction by scaling the two response curves. Though the image of the object still falls in the right field of view, the response of the two units remains balanced, causing a correct reach straight ahead. D) The peak of the response curve of a downstream unit (not necessarily involved in reaching) receiving gain modulated input from both reach-related units shifts in response to changes in gaze direction.

Artificial neural networks were shown to be capable of combining information regarding the retinal location of a stimulus and the position of the eyes to compute the spatial location of distal objects (Zipser and Andersen, 1988). The network contained three layers of units, a hidden layer connecting two input layers (containing retinal location and gaze direction) and an output layer (containing the direction of the target).

After training, units in the hidden layer responded to combinations of inputs in terms of gain fields, similar to what had been observed in monkey parietal cortex (Andersen et al., 1985). This showed that gain fields could be used to combine two reference frames into a third, transformed reference frame. That is, signals from the eye encoded in a retina-centered reference frame could be combined with signals from the eye muscles in a head-centered reference frame to transform the signals from the eye into a head-centered reference frame. Furthermore, this process could be repeated, transforming eye signals into body-centered and world-centered reference frames (Snyder et al., 1998).

Gain fields can align multiple reference frames without a "teacher"; i.e., by using statistical regularities within the input data (Salinas and Abbott, 1995). Unlike the Zipser and Andersen model that used an external input from which to compute errors during training, Salinas and Abbott showed that feedback of the consequences of neural activity through correlation-based learning rules (e.g., Hebbian learning, discussed below) were sufficient to generate arbitrary transformations between reference frames. Transformations obtained in this manner exhibited the additional, useful property of automatically aligning reference frames; i.e., a representation of an object in one reference frame, such as a retina-centered frame, could be used to compute a trajectory to that object in another reference frame, such as a hand-centered frame. Salinas and Abbott further showed that gain fields using attention as the variable-encoded gain generated spatially invariant representations of objects (Salinas and Abbott, 1997). They suggested that such a mechanism might explain the ability of neurons in Ungerleider and Mishkin's

"what" pathway to respond selectively to familiar objects regardless of their location, and neurons in the "where" pathway to respond selectively to places regardless of the object.

Vector subtraction

The ability of gain fields to shift receptive fields allows a mechanism whereby neural ensembles encoding vector quantities can be subtracted without the need for special neural machinery. If two populations of neurons project to the same structure and are modulated by independent gain mechanisms, then shifting one receptive field with respect to a different, fixed receptive field generates a difference between the two fields. This difference can be used to compute the trajectory from a reference point to the location of the shifted field (Figure 2-6).

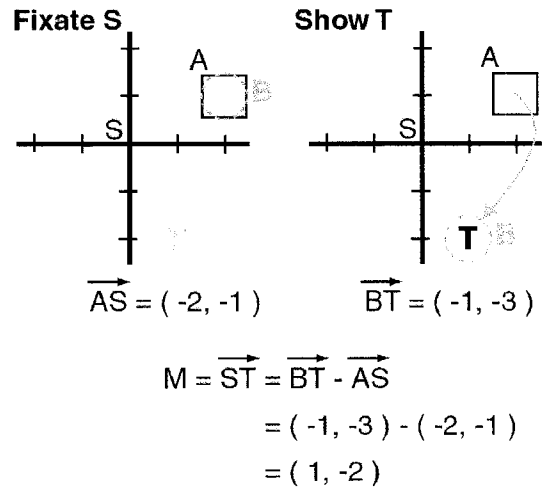

Figure 2-6. Vector subtraction by shifts in receptive fields. (LEFT) Visual receptive fields, A and B, are marginally activated during fixation at the origin, S. The eventual target (T in light gray) requires a movement down and to the right. (RIGHT) When the target appears, the receptive field of unit B is shifted to the location of the target, while field A does not shift. While the eyes remain fixated on S prior to movement onset, the vector components (here, assumed to be rectilinear) of the shift in field B, *BT*, and the displacement of A from the fixation point, *AS*, are sufficient to compute the required trajectory (BOTTOM) The computation described in vector notation.

In this manner, the same gain field machinery that allows information from multiple modalities to be combined into common reference frames (as in aligning gaze and reach frames) also allows the computation of trajectories to goal locations (as in units that sum gain field inputs). Looking at trajectory computation in this manner, the process of making a place the target of an impending movement shares many features with the representation of that place. That is, the representation of a goal and the trajectory required to reach it can now be seen as different aspects of the same computation. An example of this dual combination/computation is given by a model of the head direction system in rats (McNaughton et al., 1991;Skaggs et al., 1995). Based on a combination of inertial and visual cues, head direction cells respond selectively to the heading of rats. They have been observed in the anterior thalamus, hippocampus, post-subiculum and parietal cortex (for review, see (Taube, 1998a). The tuning curve of an individual head direction is broad, covering more than 90°. The combined activity of multiple head direction cells, however, can produce a precise encoding of head direction, in the same manner as a population of motor units produced accurate reaching directions, as shown by Georgopolous et al. (1982). Because it integrates head angular velocity, the head direction system is subject to drift, as shown by a continual increase in error with time as rats move in total darkness. McNaughton et al. (1991) and Skaggs et al. (1995) showed how visual input could correct for drift by causing the current estimate of head direction to jump to that imposed by sensory cues. The two representations that were combined in this case are current and corrected head direction. The computation involved both suppressing the current head direction and activating the intended head direction. Both of

these responses could be accomplished by gain fields: lowering the gain associated with the internal (current) head direction signal, and raising the gain associated with external (intended) head direction information. The changes in gain could result from estimates of the error between expected and observed sensory input.

Navigation

The methods of encoding and recalling information described in the previous section can be applied to the task of directed movement. Finding one's way around an environment is a routine occurrence in the life of every mobile creature, and multiple strategies exist by which neural systems may navigate their way to distant targets. Navigation computations can be grouped into three strategies with associated costs and benefits, arising through trade-offs between storage space and computational complexity: simpler ways of navigating require more things to be stored in memory, and vice-versa.

Sensory matching

The nervous system could provide navigational information by matching current activity patterns generated by sensory systems to stored patterns of past sensory events. While offering great simplicity, pattern matching does not store information efficiently. Recalling relationships between a given starting point and N stored locations requires a unique memory for each pair, producing $(N-1)^2$ patterns to be stored. This storage estimate does not include the storage redundancy that would be present in a system that utilizes multiple frames of reference. For example, a system that encodes space in both eye- and head-centered coordinates must either store sensory memories in both frames, or

provide some method for determining which frame to use selectively. This heavy storage load is offset, however, by the simplicity of spatial computations it allows. Trajectory computations could be performed by "look-up tables": "Given the current sensory information and the desired goal, turn until you see this specific pattern, then move this far."

Dead reckoning

Movement in space could be based on the relations between sensory cues and integrated movements, what navigators call "dead-reckoning". Movements are recorded as directions (either allocentric, compass-heading units such as "going north", or by egocentric, body-based units such as "turning right"), and distances (either by the integration of steps, or by measuring velocity and noting elapsed time). This reduces the storage requirement for information regarding pairs of landmarks, requiring only the distance and direction between each pair. In addition, dead reckoning can be extended to allow the computation of new trajectories from combinations of previous trajectories through the computation of "shortcuts". If you were directed to take three steps North, then three steps East, and then three steps South, you could quickly realize that the movements to the North and South were unnecessary, and just head East. You would reach the same destination, though facing East, instead of South.

Landmarks

The nervous system could use recognizable cues, or "landmarks", in the environment, both to identify the environment, and to tie patterns of activity to those

landmarks for future reference. Locations could be encoded with respect to a given landmark, greatly reducing the storage requirements of a sensory matching system. The cost of this, however, is a significant increase in the complexity of the system needed to compute trajectories. Current location and heading with respect to landmarks must be maintained, including some error-correction mechanism to prevent drift, and the system must be able to recognize landmarks from all directions under varying conditions. In addition, some mechanism must exist to convert current and intended locations into a desired trajectory.

Body-Centered Representations

Superior Colliculus

Hard-wired cross reference

Which strategies do actual neural systems appear to use? One way to address this question is to examine spatial coding in structures known to contain spatial representations. Single unit spatial representations are seen in the superior colliculus of mammals (the optic tectum in reptiles). In snakes, the tectum contains maps of visual, auditory and somatic space that are combined to generate orienting movements towards the source of signals (Hartline et al., 1978). The homologous structure in primates, the superior colliculus, controls orienting responses involving saccades (Mays and Sparks, 1980), as does the midbrain of the owl in orienting the head towards sound sources (Knudsen and Konishi, 1978). The colliculus/tectum provides an ideal example of a neural structure that merges input stored in multiple reference frames. Hartline et al.

mapped the receptive fields in the tectum of pit vipers that originated from visual input from the eyes and infrared input from the pit receptors, and found that the two maps (present in different layers of the tectum) were only roughly aligned. Lack of precise alignment suggested that precise trajectory computations required the combination of the response of multiple units to allow ballistic trajectories, i.e., a population code, much like Georgopolous et al. (1982) would later propose in monkey motor cortex. In the monkey superior colliculus, the primate homolog of the tectum, Mays and Sparks studied the encoding of eye saccades by delivering electrical stimulation during natural saccades, perturbing the natural movement. They found that, rather than encoding the movement in terms of distance and direction, saccades were made to a specific location in the orbit, much as what would be expected of a process controlled by an associative map.

Receptive fields in the colliculus are not fixed, but shift with changes in the relative position between receptor organs (Hartline et al., 1995), suggesting a receptor-centered reference frame for each map that is gain-modulated by receptor position. Shifts in receptive fields in collicular neurons show that trajectory computations can be achieved by reasonable neural mechanisms, and that these mechanisms need not be complicated.

Posterior Parietal Cortex

Anatomy

Neocortical circuits could perform associative mapping functions, and compute vector operations using gain-shifted receptive fields. Vector operations on spatial

variables would be predicted to occur in neocortical regions activated during spatial tasks, such as those in Ungerleider and Mishkin's "where" pathway, including the parietal cortex. In primates, parietal cortex resides caudal to the Central Sulcus, anterior to the Lunate Sulcus, and dorsal to the Temporal Sulcus and Temporal Cortex (Andersen, 1995). The most anterior portion of Parietal Cortex, Somatosensory Cortex (SSC), contains a representation of the body surface and projects caudally to multiple, associative parietal structures collectively called the Posterior Parietal Cortex (PPC). The parietal cortex mediates the flow of information between visual input from the occipital cortex and motor control in the frontal cortex (Caminiti et al., 1996). The PPC contains several subregions, which can be roughly divided into four areas. Moving caudally from the central sulcus : Area 5, the Intra-Parietal (IP) sulcus, Area 7, and the Medial Superior Temporal area (MST) (Figure 2-7). Area 7 is the primary source of parieto-hippocampal projections, both indirectly through perirhinal and parahippocampal cortices (Suzuki and Amaral, 1994;Lavenex et al., 2002), and possibly through direct cortico-hippocampal projections. (Rockland and Van Hoesen, 1999), though the strength of these latter projections is weak.

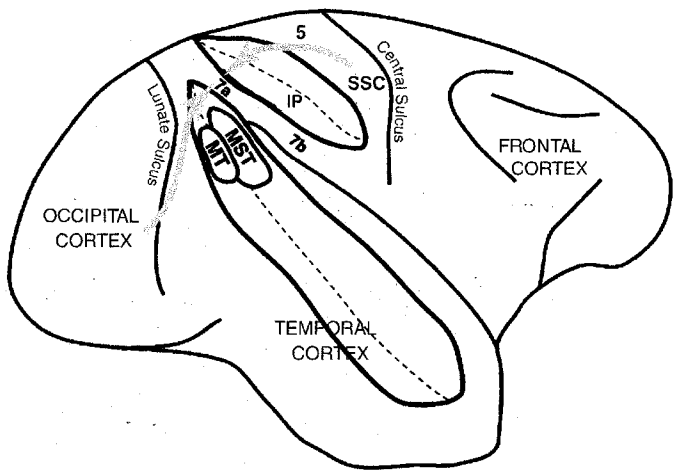

Figure 2-7. Schematic showing the convergence of information onto Area 7. Visual information from occipital cortex flows to the Medial Temporal (MT) and Medial Superior Temporal (MST) cortex, both of which are involved in forming representations that are independent of motion. These structures send motion-insensitive visual information to Area 7. Somatosensory information reflecting limb positions and movements flows from primary somatosensory cortex (SSC) to Area 5, and then into Intra-Parietal (IP) cortex, being transformed to successively more general coordinate frames (e.g., from hand to shoulder to body centered coordinates). IP cortex projects limb position information that has been transformed into body-centered coordinates to Area 7, which generates world-centered representations of the body (e.g., the position of a hand in relation to a nearby coffee mug). (Adapted from (Andersen, 1995).

Unlike the grooved, or "gyrencephalic", neocortex of primates, rodent neocortex cannot be partitioned by surface features because the cortex of the rat is smooth, or "lissencephalic". In the rat, cortical regions are based on similarities in cytoarchitecture (the presence of different cell types, and the physical features and connectivity of those cell types). Without obvious landmarks, the boundaries of cortical regions in the rat are not as clearly defined as in the primate (Zilles, 1985) (Figure 2-8).

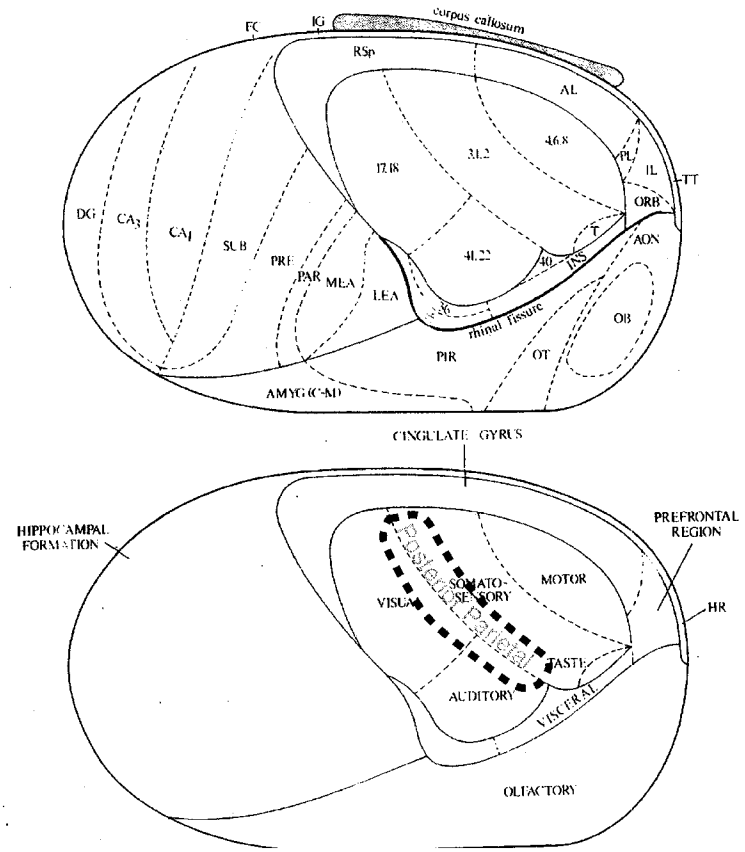

Figure 2-8. Schematic of the cortical sheet in the rat. The neocortex of the rat is smooth, so regions cannot be defined according to surface landmarks, as in primates. Regional designations are based on cellular architecture and connectivity. At the gross level, regions are defined by stereotaxic coordinates, which can vary depending on the author. Zilles (1985) placed the rodent PPC at the intersection of parietal and occipital cortex (areas 3, 17 and 18 in the figure). Burwell et al. (1995) restricted PPC to a region located entirely within somatosenory cortex.

The rodent PPC is not as clearly defined as is the primate PPC. Krieg (Krieg, 1946) identified the rodent PPC as a strip of cortex caudal to primary somatosensory cortex and rostral to primary visual cortex, which Krieg labeled Area 7 (Chen, 1989). Zilles (1985) considered the rodent PPC to lie at the intersection of several neocortical regions including medial and lateral extrastriate occipital cortices, Oc2MM and Oc2ML, and the primary somatosensory area. Burwell et al. (1995) placed the PPC in the caudal

portion of the primary somatosensory area, but excluded any occipital cortical areas. Cortical connectivity with the hippocampus in the rat, however, possesses many similarities with the connectivity of the primate hippocampus (Burwell and Amaral, 1998) (Figure 2-9).

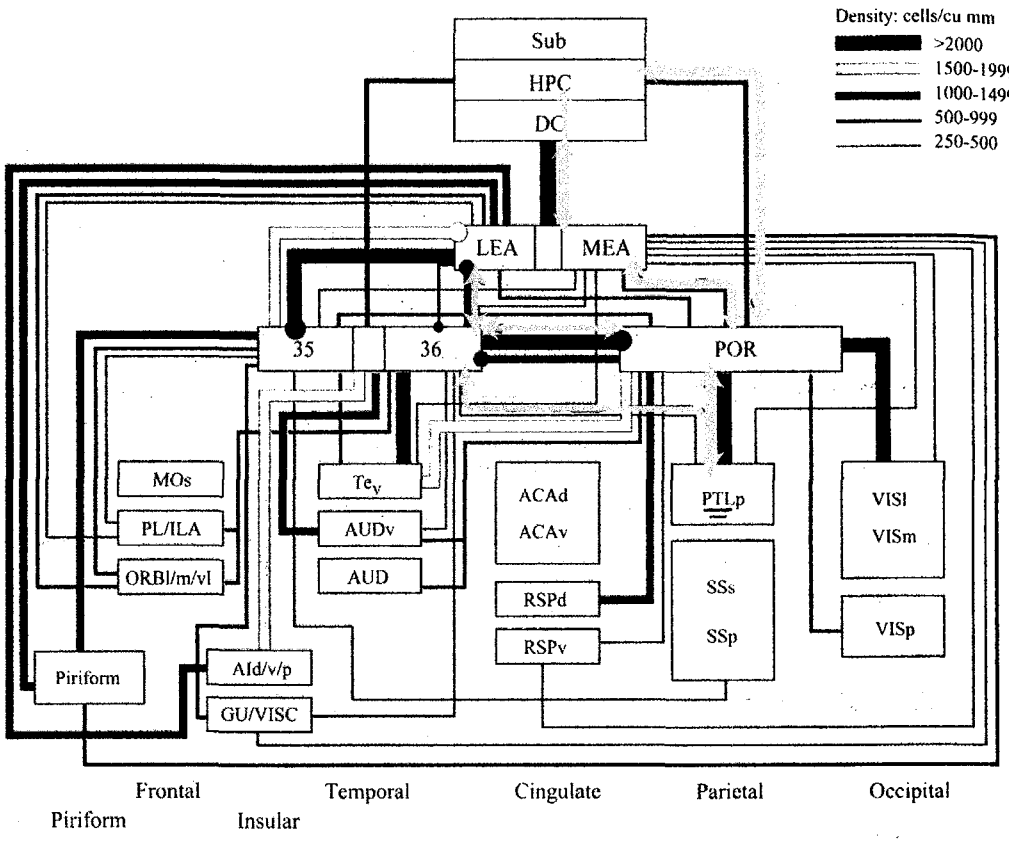

Figure 2-9. Schematic of cortical-hippocampal connectivity in the rat. The communication pathway between PPC (PTLp) and hippocampus (DG and HPC) are shown as gray arrows. PPC projects bi-directionally to both perirhinal cortex (areas 35 and 36) and postrhinal cortex (POR). Projections from postrhinal cortex join those from perirhinal cortex into the lateral entorhinal cortex (LEA), as well as the medial entorhinal cortex (MEA). Entorhinal cortex projects into the dentate gyrus (DG) and the hippocampus proper (HPC), including CA3 and CA1. CA1 and the Subiculum project to the deep layers of entorhinal cortex, and directly to postrhinal cortex. (Modified from Burwell and Amaral, 1998).

Rodent PPC does not project directly to the hippocampus; it projects to both peri- and post-rhinal cortices, which are homologous to the perirhinal and parahippocampal

cortices in primates (Burwell and Amaral, 1998). These cortices project to entorhinal cortex, which then projects to the dentate gyrus of the hippocampus. Projections from the hippocampus back to PPC go through the entorhinal and then perirhinal cortices. Projections to PPC from postrhinal cortex are not as strong as those going in the opposite direction (Burwell and Amaral, 1998). Both the PPC and hippocampus in rodents connect bi-directionally to several subcortical targets including the thalamus and basal ganglia (Burwell and Amaral, 1998). Unlike the entorhinal cortex in monkeys, which receive over 70% of its input from perirhinal and parahippocampal cortex (Suzuki and Amaral, 1994), the entorhinal cortex of rodents receives less than 15% of its input from perirhinal and postrhinal cortex (Burwell and Amaral, 1998). That is, the rodent hippocampus receives a greater proportion of its input from brain regions other than parietal or temporal cortex than do monkeys.

In humans, deficits following bilateral lesions of the PPC include errors in visually guided reaching (called "optic ataxia") and in the generation of saccades to foveate both fixed and moving targets (DeRenzi, 1982; (Hyvarinen, 1982), while unilateral parietal lesions (particularly of the right side) result in "hemineglect", i.e., inattention to one side of an object or scene (Bisiach et al., 1979). It should be noted, however, that hemineglect also follows unilateral frontal lesions (Guariglia et al., 1993), suggesting the parietal cortex cooperates with frontal structures in the representation of space (Burnod et al., 1999).

Physiology

In non-human primate PPC, single unit activity in somatosensory-related reference frames progress caudally from the central sulcus, and visually-related frames rostrally from the lunate sulcus, converging in Area 7 (Andersen, 1997). Area 5, the most rostral part of the PPC, receives projections from primary somatosensory cortex and represents arm movements during reaching in a body-centered reference frame (Lacquaniti et al., 1995) with axes of azimuth, elevation and distance of the hand from the body. This representation was independent of intention to move or not to move because activity persisted through delay periods in a "go/no go" task (Kalaska and Crammond, 1995). Caudal to area 5, three cortical subregions in the Intraparietal Sulcus (IP) progressively translate the reference frame from arm- to head- to eye-centered reference frames (Colby et al., 1993). It should be noted that these distinctions are not precisely limited to a given anatomical region, but rather are organized along a continuum from body- to eye-centered representations (Duhamel et al., 1997), and that the initiation of movement is not required for increased activity; i.e., activity in these structures is not strictly motor-related (Assad and Maunsell, 1995). Jumping to the caudal input to PPC, two areas of the temporal lobe (areas MST and MT) receive input from visual areas, respond to visual expansion and motion, and project to the most caudal part of the PPC, area 7 (Van Essen et al., 1981;Tanaka and Saito, 1989). Area 7 neurons are driven by multiple sensory modalities, including vision. Many neurons in Area 7 encode space in allocentric (or world-centered) reference frames, showing responses that are

insensitive to changes in head direction or gaze shifts (Snyder et al., 1998). Area 2v, which lies in the vicinity of the intersection between Areas 5 and 7, receives projections from vestibular nuclei (Fukushima, 1997), and produces short latency responses to stimulation of the eighth cranial nerve (Boisacq-Schepens and Hanus, 1972). The presence of physiological responses to multiple sensory modalities in many different reference frames in monkey parietal cortex has led to the hypothesis that at least one parietal function is to transform coordinates from one representation into another.

Much less is known about the role of the PPC cortex at the unit level in forming spatial representations in the rodent. Support for vector representations of space in the rodent PPC comes from the observation of units that produced direction-selective responses, both in allocentric coordinates (e.g., head direction cells) and egocentric coordinates (e.g., cells that responded to right and left turns) (Chen, 1989). Chen found few PPC neurons that responded to purely sensory or purely motor signals. Some PPC neurons expressed a directional bias that persisted in the dark. Directional cells in certain parietal regions rotated with distal cues, while those in others remained fixed, similar to "head direction cells" found in several limbic structures including the anterior thalamus and post-subiculum (Taube, 1995). The activity of other units in rodent PPC correlates with combinations of body movements and spatial location, producing "right turn at end of arm", "left turn at end of arm" and "running towards the center" cells (Chen, 1989). Units whose responses represent the conjunction of different types of responses are the hallmark of associative maps. McNaughton et al. (1996) proposed such activity could be used to compute trajectories using vector coding of landmarks; i.e., that positions could

be encoded as vectors (i.e., as distance and directions) to landmarks, and that trajectories could be computed by vector subtraction operations based on associative maps. What remains to be shown is the presence of units in rodent PPC that encode distance to visible landmarks.

One caveat to any search for a specific set of coordinates in a neural structure – such as representations in a goal-centered reference frame in rodent parietal cortex – is that reference frame transforms may be quite general and need not follow a specific sequence of transform steps. Salinas and Abbott (1997) showed that gain fields are capable of transforming sensory space directly into motor space, bypassing intermediate reference frames. An example of direct sensorimotor transformation was found in Area 5 in monkeys during a reaching task under visual control (Buneo et al., 2002). One means of solving this task through transformations of reference frames would be to encode the retinal signal in body-centered coordinates and subtract the body-centered coordinates of hand position. This would allow the hand and eye to move in the same coordinate space. Buneo et al., however, found that Area 5 contained representations that were modulated by both eye and hand position and not the body, implying that this region bypassed the need for an intermediate transformation step.

World-Centered Representations

Posterior Parietal Cortex

World-centered correlates in primates

The PPC as a coordinate transformer culminates in the transformation of body-centered frames into world-centered, or allocentric, frames. The degree of gain field modulation in non-human primates was quantified by computing the ratio of firing rate responses to visual stimuli falling on two, fixed retinal locations (Snyder et al., 1998). When the head was rotated, the ratio of responses at the two retinal locations differed in a significant number of units in area IP, but not in Area 7. This suggests that IP units encode activity in head-centered coordinates, because the degree of modulation varied with the position of the head with respect to the body. When both the head and body were rotated together, however, the degree of gain modulation remain fixed in IP, but differed for a significant number of units in Area 7, suggesting that Area 7 activity was modulated by the relative position of the entire body to the world. Gain modulation was dependent on the angular velocity of the head and body combined exceeding a threshold, suggesting that the world-centered frame of reference was dependent on vestibular input. It is not known whether parietal representations would be equally sensitive to translations within the environment without rotations.

Hippocampus

Anatomy

In primates, the hippocampus sits in the medial wall of the temporal lobe. The PPC (particularly Area 7) sends strong projections to the perirhinal and parahippocampal cortices, which in turn send strong projections to the hippocampus via the entorhinal cortex (Suzuki and Amaral, 1994). In addition, the hippocampus receives substantial inputs from temporal and frontal cortices, as well as many subcortical structures, including the amygdala, thalamus, hypothalamus and basal ganglia. Anatomically, the hippocampus receives highly processed input from a number of different association cortices and sensorimotor structures. The hippocampus can be thought of as the highest association area in the brain, or as "supramodal association cortex" (Mesulam et al., 1977). In a sense, this discussion now jumps from rising through successively higher association cortical areas to the highest area, skipping a number of intervening structures in the process, each of which likely makes significant contributions to spatial computations, but which will not be considered in this dissertation. The emphasis of this dissertation will not be to attempt to connect the different parietal and hippocampal responses during spatial computations, but rather to compare the two to understand the role of each.

Examining the neuroanatomy of the hippocampus proper, the basic structures and connectivity are preserved between non-human primates and rodents, though oriented differently in the rat brain (Figure 2-10).

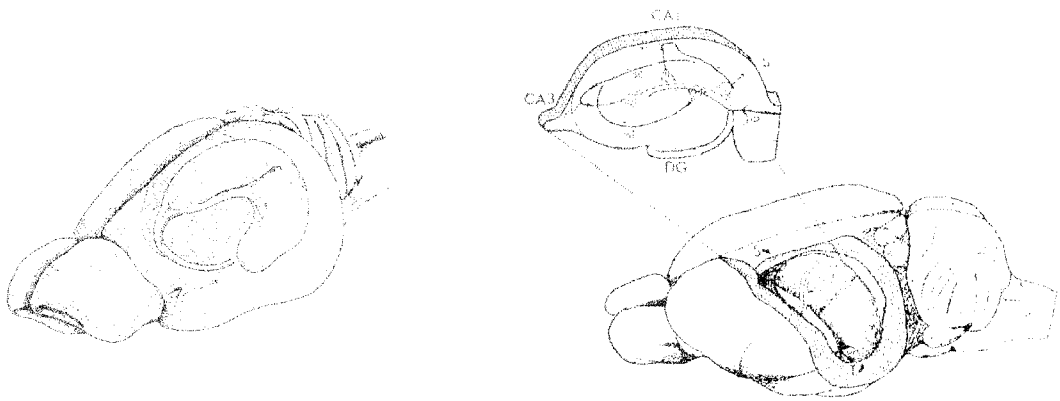

Figure 2-10. Two views of the hippocampus in the rodent brain. (LEFT) The hippocampus curves down the rostral-caudal axis bilaterally, connecting at the rostral pole. (RIGHT) In rodents, the dorsal portion of the hippocampus is situated beneath parietal cortex (removed in the drawing). This relation between the position of the two structures allows electrodes positioned at different depths from the dorsal surface to record from both structures. (From Burwell et al., 1995).

The hippocampus contains several substructures connected roughly in a loop projecting one into the next, beginning and ending with the Entorhinal Cortex (EC) (Figure 2-11). In some ways, the components of this loop-oriented structure resemble the input-output organization of the component structures of PPC, like a series of filters passing the results of one operation onto the next structure, by-passing structures in some cases. Unlike PPC, however, hippocampus includes two loops: one in CA3, which projects back onto itself, and the other including the processing of the entire hippocampus, projecting from and back onto the EC, though at different cortical levels.

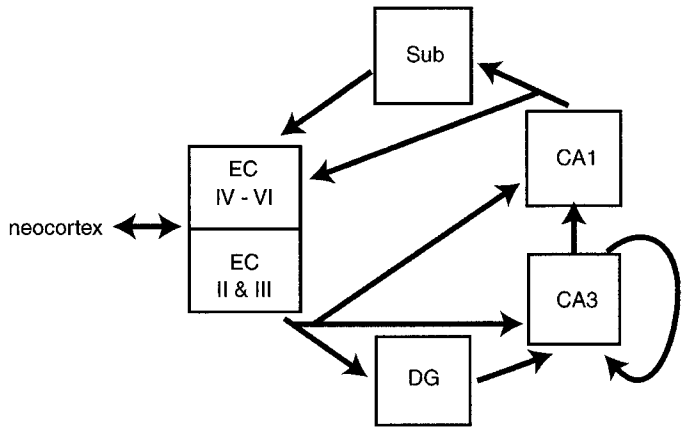

Figure 2-11. Schematic of connectivity within the rat hippocampus. Excitatory connections project roughly in a loop beginning with the superficial layer (those nearest the skull) of the Entorhinal Cortex (EC layers II and III), which project through the Perforant Path into the Dentate Gyrus (DG), CA3 and CA1. The DG projects to CA3 via the Mossy Fibers. CA3 projects both to CA1 via the Schaffer Collaterals, and back onto itself through longitudinal association fibers. CA1 (the region from which all hippocampal recordings discussed in this dissertation were made) projects to the Subiculum (Sub) and the deep layers of EC (layers IV-VI), completing the loop of primary projections. Note that information returned to the deep layers of EC can be both passed along to neocortical structures, and passed back into the hippocampus via communication with the superficial layers of EC.

Each sub-region of the hippocampus exhibits unique patterns of innervation and local connectivity that modify the canonical local circuit, which could provide insight into the computational processes performed by each (McNaughton and Morris, 1987;O'Reilly and McClelland, 1994;Treves and Rolls, 1991). Perforant Path projections from superficial layers of EC distribute information across a much larger population of DG granule cells. This expansion in the number of cells between EC and DG is thought to contribute to separating similar input patterns, allowing for the distinct encoding of similar events. DG projections to CA3 through the Mossy Fibers end in particularly efficacious synapses, which require fewer co-active inputs to drive post-synaptic neurons to fire. Such connections are well suited for imposing new patterns on top of ongoing activity, or for activating processes responsible for maintaining short-term impressions of

patterns (a sort of "photographic memory") (McNaughton and Morris, 1987). In addition, CA3 contains robust feedback connections onto itself, allowing ongoing CA3 activity patterns to be reinforced, or to allow the recall of complete patterns of activity from parts of the pattern. CA3 also contains some excitatory feedback projections to the granule cells of the DG (Penttonen et al., 1997). CA1 pyramidal cells project to both the Subiculum and deep layers of Entorhinal Cortex, and are the units from which all hippocampal recordings discussed in this dissertation were made. Several classes of interneurons provide inhibition within each structure in the hippocampus (for review see (Freund and Buzsaki, 1996), including two classes in particular that may play different computational roles: basket cells, which form synapses along the base of dendrites near the pyramidal cell body, and control synaptic input to the cell body through shunting inhibition, and chandelier cells, which form synapses on the axons of pyramidal cells, possibly controlling the overall level of activity in the hippocampus. In addition, inhibitory projections from CA1 provide feedback inhibition on both CA3 and DG (Sik et al., 1994). When compared to the laminar, columnar structure of cortical regions, hippocampal local circuits appear to be well suited for maintaining patterns of input activity in tightly controlled feedback loops (Marr, 1971).

Spatial deficits

Lesions

In humans, lesions of the hippocampus (often arising from stroke, tumor, anoxic event, gunshot wound or surgical intervention) consistently have produced spatial deficits

(Scoville and Milner, 1954;O'Keefe and Nadel, 1978;DeRenzi, 1982;McCarthy and Warrington, 1990;Squire, 1992), suggesting a fundamental role for hippocampus in spatial computations and memory. The landmark case of "H.M." (Scoville and Milner, 1954) identified a hippocampal role in the acquisition of new memories following the bilateral loss of medial temporal lobe regions including the hippocampus. "H.M." suffered from intractable seizures originating bilaterally in the medial temporal lobe. Following removal of the hippocampus (along with significant amounts of surrounding temporal lobe tissue), H.M. suffered profound anterograde amnesia (an inability to form new memories, including new spatial memories) though his IQ remained above average. In addition, H.M. suffered some retrograde amnesia (an inability to recall memories of events and places prior to the surgery), though the extent of this deficit remains the subject of debate (Corkin, 2002). H.M. was still capable, however, of drawing a detailed room layout of the house into which he and his parents moved several years after his surgery (Corkin, 2002). Corkin suggested that preservation of navigation skills in hippocampal amnesiacs depended on the familiarity of the environment. In a different case (patient "E.P."), however, a hippocampal patient who had moved to a new neighborhood following the lesion and had lived there for five years was unable to describe the spatial layout of the new neighborhood (Teng and Squire, 1999). This patient could recall remote spatial memories, but could not acquire new spatial memories. One explanation of the discrepancy in spatial recall of these two patients involves the extent of each lesion; the lesion in E.P. is extensive, and includes perirhinal and parahippocampal cortex. Another explanation involves the scale of the region recalled:

H.M. recalled the layout of a small home; E.P. was asked to recall the layout of an entire neighborhood. It is doubtful that E.P. had as much experience walking through his neighborhood as H.M. had walking through his home. Hippocampal lesions, however, did not prevent either subject from finding their way from a start to a goal within a new environment. This suggests that extra-hippocampal structures activate patterns of activity associated with navigation; i.e., the hippocampus is not the only structure involved in trajectory computations.

Spatial deficits have also been associated with parietal and hippocampal lesions in rats. Sutherland, Whishaw and Kolb (Kolb et al., 1983;Sutherland et al., 1983) tested spatial learning deficits in rats with PPC and hippocampal lesions on the Morris water maze, in which a platform is submerged just under the surface of a pool of milky water, which prevents the rat from seeing the platform. The submerged platform provides an escape for rats placed in the pool, but only if the platform can be located. Rats that reached the platform, either through random search during swimming or by being placed there by researchers, will subsequently swim to the platform when placed in the pool, apparently recalling the arrangement of distal visual cues to locate the hidden platform. Rats with parietal lesions took longer to learn the location of the platform than did controls, while rats with lesions to either the dentate gyrus or CA3 failed to learn the location of the maze at all. In a variant of this task, hippocampal and PPC lesioned rats were tested on variants of the Morris water maze (Save and Poucet, 2000). In the "distal" training condition, the location of the platform remained fixed with respect to the arrangement of distal cues, as in the standard Morris task. In the "proximal" training

condition, the location of the platform remained fixed with respect to three large cues placed inside the pool perimeter, but was rotated randomly (along with the proximal cues) with respect to distal cues between trials, reducing the utility of distal landmarks in aiding search. Hippocampal lesioned rats were slower to learn to the location of the platform in both conditions, while PPC lesioned rats were impaired in the proximal, but not the distal condition. The impairment of PPC lesioned rats in the proximal condition suggests that the PPC is involved in the integration of multi-modal information (e.g., visual and idiothetic). Their lack of impairment in the distal condition suggests that the PPC is not involved in recalling the overall context of the environment, or orienting within that environment.

Role in cellular learning

Hebb

Some brain mechanism or mechanisms must associate the disparate elements of experiences in a manner that allows subsequent recall of the whole experience from the parts. Hebb (1949) proposed a biologically plausible mechanism for associating neural responses based on the presence of simultaneous activity. This proposal has become known as Hebb's Postulate:

> "When an axon of cell *A* is near enough to excite cell *B* or repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that *A*'s efficiency, as one of the cells firing *B*, is increased." (Hebb, 1949)

A possible cellular basis for long-lasting synaptic change that would allow the formation of associations was discovered in the hippocampus (Bliss and Lomo, 1973). Bliss and Lomo observed long-lasting (> 4 hours) changes in synaptic efficacy in the dentate gyrus following high-frequency electrical stimulation of perforant path fibers, a phenomenon that became known as Long-Term Potentiation, or LTP. Subsequent work showed that this change required co-activation of inputs onto a common post-synaptic cell (McNaughton et al., 1978). McNaughton et al. stimulated perforant path fibers, while recording from granule cells in the dentate gyrus. At low stimulation intensity, synaptic responses in granule cells could be observed without the induction of LTP, showing that minimal stimulation of the post-synaptic cell was not sufficient to generate LTP. McNaughton et al. also showed that simultaneous stimulation of two components of the perforant path (the medial and lateral pathways) produced LTP at intensities that did not produce LTP when either pathway was stimulated, separately. Evidence for pre-synaptic specificity in LTP induction utilized the fact that the perforant path in each hemisphere projects to both the ipsilateral dentate gyrus and, to a lesser extent, to the contralateral dentate gyrus (Levy and Steward, 1979). A stimulation intensity that was capable of

inducing LTP in the ipsilateral dentate gyrus failed to induce LTP contralaterally. In addition, the induction of LTP in the ipsilateral pathway did not produce LTP for stimulation of the contralateral pathway. The biochemical cascade leading from synaptic activation to LTP induction includes several steps. LTP induction in the hippocampus can be blocked by antagonists of the N-Methyl-D-Aspartate (NMDA) glutamate receptor (Harris et al., 1984), and by inhibitors of protein synthesis (Stanton and Sarvey, 1984) (for review see (Malenka and Nicoll, 1999). To summarize, LTP induction requires co-activation of multiple inputs, is specific to those synapses common to both the pre- and post-synaptic induction pathways, and requires protein synthesis to produce a long-lasting increase in synaptic efficacy.

The relevance of synaptic enhancement through LTP to questions of spatial memory has been the subject of much debate (O'Keefe and Nadel, 1978;McNaughton et al., 1996;Eichenbaum et al., 1999). A connection between LTP and spatial memory has been shown by experiments that interfered with the expression of LTP. Saturating levels of LTP were associated with increased errors in spatial tasks (McNaughton et al., 1986), and maximal, electro-convulsive shock (which occludes subsequent LTP-related changes) produced spatial learning deficits in the Morris water task (Barnes et al., 1994). The infusion of a drug that selectively blocks AMPA/kainate receptors produced both deficits in spatial learning both during the acquisition of new spatial information, and during the recall of previously learned information in the Morris water maze (Riedel et al., 1999). Rats were implanted with minipumps containing either vehicle or the water-soluble AMPA receptor blocker, LY326325 ("LY"), which were disconnected following

the last of three days of training on the Morris water maze task. Rats from both groups were then challenged with an acute injection of LY prior to a task retention task two weeks after the completion of training. The performance of rats that received LY during either chronically during training or acutely prior to the retention task was impaired compared to vehicle controls. In a different group of rats, LY was administered either acutely or chronically over several days to rats after training on the task. Only rats in the group receiving chronic LY were impaired during the retention task over two weeks later. Combined, these results suggest that AMPA-mediated hippocampal activity is necessary for acquisition, storage and recall of allocentric spatial memories.

The role of LTP in the acquisition and recall of sequences has been studied both theoretically and experimentally. If LTP strengthens connections between neurons whose firing was closely correlated in time, then LTP should alter the connections between place cells with closely overlapping fields (Levy, 1996;Blum and Abbott, 1996). Levy realized that cognitive mapping was a sequence prediction problem, and so proposed sequence encoding as a function of the hippocampus. Blum and Abbott realized that rats running on circular tracks visit place fields in the same sequence on multiple occasions, and that this should cause a biased strengthening of connections in one direction, but not the other. As the synaptic strength between two cells grew stronger, the first cell should cause the second to fire sooner along the trajectory, causing the place field of the second to shift closer to that of the first, opposite the direction of running. This shift was observed subsequently in hippocampal place cell data by Mehta et al. (Mehta et al., 1997). It should be noted that this shift in the response of a receptive

field differs from that produced by gain fields in that it is not as dynamic, persisting beyond the termination signals that caused the shift.

Place cells: why they could represent place

The possibility of connecting the cellular basis of learning in the hippocampus to the spatial deficits observed in hippocampal lesions arose in the early seventies with the observation of single units in the hippocampus that responded preferentially when a rat visited a certain portion of an environment (O'Keefe and Dostrovsky, 1971). These units were given the name "place cells", and the localized region in which a given place cell became maximally active was called a place field. When multiple single units were recorded simultaneously from the rat hippocampus, place fields were distributed across the space that was accessible to the rat (Wilson and McNaughton, 1993). The spatial information contained in the firing rates of the set, or "ensemble", of units was used to reconstruct the location of the rat over 1 second epochs, and produced a mean error of 5 cm.

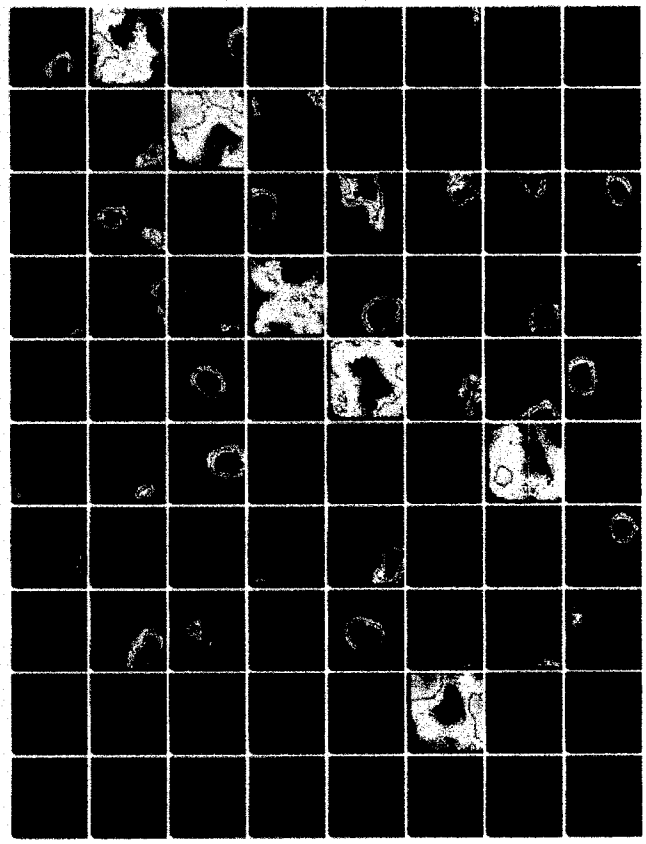

Figure 2-12. Spatial firing distributions of multiple, simultaneously recorded hippocampal pyramidal cells and interneurons in a freely moving rat. Each panel represents the spatial firing distribution of a single cell within a square, 62 x 62 cm box. The distribution was computed by normalizing the number of spikes that occurred in a given location by the amount of time the rat spent in that location. For cells with significant spatially related firing, red indicates the maximum firing rate. Blue indicates a lack of spatially significant activity. The panel in the leftmost column of the top row shows a typical spatial firing distribution of a hippocampal pyramdial cell (i.e., a place cell). The panel in the column second from left in the top row shows a typical spatial firing distribution of a hippocampal interneuron. Note that though the spatial firing of the interneuron is more distributed than that of the place cell, the activity does contain some degree of spatial information (Wilson and McNaughton, 1993).

The finding that ensembles of hippocampal place cells contained the spatial information necessary to represent each location in an environment uniquely supported an earlier theory regarding one possible function of the hippocampus. O'Keefe and Nadel (1978) combined synaptic, cellular and behavioral evidence to describe how the

hippocampus could contain a "cognitive map" of the current environment, for use in (among other things) learning spatial relationships of cues.

Early experiments designed to determine the factors that controlled place cell activity searched for external variables to which place cell activity could be tied, as had been the case for many sensory related structures including, for example, in PPC (Mountcastle et al., 1975). The simplest theory was that activity in a given place cell represented being in that place, in the way that activity in a PPC vision-related cell represented the presence of a specific stimuli in the visual environment. Such a direct mapping, however, was soon found to be inadequate for describing place cell properties. Individual units can respond to a different location in a new environment, or might not respond at all (O'Keefe and Conway, 1978). Though the removal of specific cues in a familiar environment could alter place cell activity (O'Keefe and Speakman, 1987), suggesting some degree of dependence on external cues, place field size and shape were invariant to changes in the size and shape of distant cues (Muller et al., 1987). In addition, Muller et al. found that changing the shape of the enclosure changed the location of place fields in an unpredictable manner, quite unlike the fixed response of neocortical units, which may shift but do not shift relative to one another. Place fields, which were initially active in a lighted environment and dependent on the location of distal cues, maintained their positions in subsequent darkness, suggesting that mechanisms other than visual input could control place fields (Quirk et al., 1990). Similarly, place fields active when the rat was placed into a darkened enclosure maintained their location when the arena was lighted, but differed from the arrangement

of fields observed in the same enclosure when it was illuminated during the entry of the rat. Clearly, place cell responses were neither independent of, nor directly tied to, environment stimuli.

Furthermore, the degree to which place cell properties depended on distant cues varied with exposure to a given environment. Place cells responded upon the initial entry into a place field (Hill, 1978), showing that place specificity was not strictly a learned response. Subsequent work, however, has shown that the location and boundaries of place fields can be modified significantly following the initial experience (Wilson and McNaughton, 1993), and can diverge significantly over time, even in the absence of differential reward (Lever et al., 2002). Experience can modify place cell responses by substituting a black cue card for a white card on the wall of an enclosed cylinder, and noting that place cell responses to the two cards diverged over subsequent days (Bostock et al., 1991). When rats were disoriented by being rotated prior to each entry into the same environment, place fields failed to follow subsequent cue card rotations (Knierim et al., 1995), suggesting that stable place-related activity was tied to the stability of both external cues and an animal's ability to recognize those cues in stable locations.

Other experiments have shown that place cell responses carry spatial information. Spatial information can be defined as the precision of the prediction of location, given knowledge of the rat's location when a spike occurred. Skaggs et al. (1996) used the following example: a spike that restricted the prediction of the rat's location to a specific half of the maze conveyed one bit of spatial information. A spike that restricted the prediction to one quarter of the maze conveyed two bits of spatial information.

Hippocampal responses cannot, however, be placed within a reference frame containing fixed distance metrics, as is the case for parietal responses. While preserving spatially restricted responses, hippocampal place fields can be stretched in response to changes in the environment. After allowing rats to become familiar with a fixed environment over several training sessions, the size of the environment was expanded by moving the walls further apart (O'Keefe and Burgess, 1996). Place fields stretched with the new locations of the walls, maintaining fixed positions relative to the intersections of walls, but occupying a larger space. Spikes from place cells with stretched place fields still contain significant information per spike; there should be no loss of ability to decode the location of the rat from an ensemble of stretched place fields, even though the coordinate system to which those place fields are attached has been warped.

Knowing one's current location alone is insufficient for the computation of a trajectory; information regarding direction is required, too. Two types of direction information are particularly relevant to trajectory computations: current heading and target heading. In rats, a fixed directional reference is provided by a system composed of "head direction cells" (Ranck, Jr., 1984). Each head direction cell responds across a broad directional tuning curve; i.e., using Ballard's definitions, head direction cells "value encode" the current heading of the rat (see Figure 2-2 for a similar example involving the directional responses of motor neurons). Across a population of head direction cells, the directional tuning curves are distributed around the azimuth of the rat, meaning that the current heading of the rat can be decoded from the ensemble of head direction cells (Taube, 1998b). Head direction cells have been identified in the hippocampus (Ranck,

Jr., 1984), anterior thalamus (Taube, 1998b), lateral dorsal nucleus of the thalamus (Mizumori and Williams, 1993), neocortex (Chen et al., 1994a) and striatum (Trullier et al., 1997). Like place cells, the tuning of head direction cells can rotate with rotations in local cues (Taube, 1998b), can be maintained in the absence of visual cues (McNaughton et al., 1994b), and yet can be disrupted by instability of visual cues (Knierim et al., 1995). When several head direction cells were recorded simultaneously, the heading of the rat could be decoded, similar to the manner in which an ensemble of place cells allows the decoding of location. Together, place fields and head direction cells provide a current location and heading with respect to distant cues, producing a complete, allocentric reference frame, which could be used for trajectory planning.

In conclusion, hippocampal place cells do not respond simply to specific environmental patterns, and their responses can be modified with experience. Place cell activity (particularly in conjunction with head direction cells) does contain sufficient information to serve as a "cognitive map" on which current and goal locations could be represented, and which could be used to compute trajectories between the two. These considerations have led to the hypothesis that place fields combine to form map-like representations of a given environment (McNaughton et al., 1996). Map-like representations could constitute a system capable of representing space in terms of location (rather than distance) and direction, and capable of stabilizing the activity of many individual units to maintain a representation of a maze or room much larger than the area of a single place field (O'Keefe and Nadel, 19778; Wilson and McNaughton, 1993).

Goals could activate individual neurons

Trajectories could be computed within the hippocampus itself if goal-related activity were present. Place fields could represent current location, head direction cells could maintain a fixed directional reference, and the trajectory could be computed using a feed-forward network similar to that used by Zipser and Anderesn (1986) to describe spatial transformations in parietal cortex. Activity similar to that expected of a "goal cell" has been observed in monkey hippocampal cells that respond to distant locations, and have been labeled "view cells" (Rolls et al., 1997). View cells were observed during single unit recordings in monkeys that were supported by a moveable carriage. The carriage supported and held the head, so that head direction and carriage direction were the same. Eye position was monitored by scleral search coils. Hippocampal view cells responded selectively when the monkey looked at a specific section of the room wall, irrespective of the current location of the monkey. Changes in view cell firing could not be explained by eye position, head direction or current location, but did correlate with the viewed wall location in allocentric coordinates (Georges-Francois et al., 1999). The response remained selective for sections of the wall even when the specific features of the wall were covered (Robertson et al., 1998), suggesting that the response was not due to salient features on the wall itself. Such units would posses the necessary information to allow trajectory computations within the hippocampus using vector subtraction mechanisms, such as shifting receptive fields. It should be noted, however, that some controversy exists regarding the presence place fields in the non-human primate hippocampus similar to those observed in rats, with some reports finding place cells (Ono

et al., 1993), and others not finding place cells (Georges-Francois et al., 1999). It is possible that "view cells" were not recorded in CA1, but rather were recorded from parahippocampal cortex. If goal locations, in addition to current location, activated unique patterns in the rodent hippocampus, trajectory computation could utilize even a brief period of goal-related activation. That is, the hippocampus might only represent goals transiently, at which times trajectories could be computed.

Goals could activate new reference frames

Another possibility is that the activation of goal-related activity leads to changes in place cell activity; that is, changes in the place fields of some units. Several experiments have shown that the relationship between place cells and fields can be altered by many different manipulations. When rats were trained to shuttle back and forth on a linear track, place fields observed during running in the two directions were uncorrelated (McNaughton et al., 1983a), suggesting that a change in reference frame had occurred at the ends of the track. Analysis of place field activity at the ends, however, did not show any discontinuity that would be expected if such a change in reference frame had occurred (Redish, 1999). Muller and Kubie (Muller et al., 1987) varied the size and shape of a cylindrical environment, along with a single cue card located on the wall. When the cue card was removed, the size and shape of place fields remained the same, but were rotated randomly across trials. When the walls of the maze were expanded, creating a larger diameter cylinder, the place fields of roughly 1/3 of the cells increased in size, but retained their position with respect to the cue card. When the floor plan of the maze was changed from circular to rectangular, no correlation could be found

between place fields from the previous, circular floor plan. Wilson and McNaughton (1993) recorded from rats in a square box with a single, moveable wall over several days, allowing the rat to become familiar with the environment. When the wall was removed, allowing the rat to explore a connected, novel space of roughly the same size, many place fields remained unchanged, but several in the familiar side of the maze were altered. Skaggs and McNaughton (1998) allowed rats to explore two square arenas, connected by a passageway, of identical size and color. Some cells maintained similar field locations in the two arenas, while the fields of others were dramatically altered. Changes in direction, task, distal cues, maze shape, and physical location can all lead to partially or completely different reference frames in the hippocampus, showing that hippocampal activity patterns can be controlled by many different variables.

Map-like representations

The term "reference frame" has been used in this review to describe activity in both parietal cortex and hippocampus. One could ask if the term means the same thing in regards to the neural responses of both structures. The response of units from both structures can be associated with a system of coordinates that is fixed with respect to some origin (e.g., the retina, hand, head, or in the case of "allocentric" coordinates, the environment), defining a response space. Within a given coordinate system, the mapping between physical variables and neural responses in both parietal cortex and hippocampus is continuous; i.e., units with overlapping responses produce similar responses to all inputs. In addition, in hippocampus, similarity of responses can be preserved in the presence of changes in environmental stimuli. For example, place fields can rotate as a

coherent group (i.e., place fields preserve neighbor relations) in response to the rotation of visual cues (Muller and Kubie, 1987) or follow rotations of a maze above the vestibular threshold (Sharp et al., 1995), though both of these responses depends to some extent on the familiarity of the environment. In addition, neural responses span the available space within that coordinate system; i.e., there do not appear to be any discontinuities that produce no neural response. In motor cortex, directional responses of motor units to arm movements are uniformly distributed throughout angle space (Georgopoulos et al., 1982), and hippocampal place fields appear to be distributed uniformly across physical space (O'Keefe and Nadel, 1978); but see (Hollup et al., 2001). Two challenges, however, can be raised to this assertion: first, it is difficult to disprove because one can always argue that any observed discontinuity could be represented by some neuron that was not recorded. Second, experimental observations in the hippocampus appear to contradict it. When recordings were made from rats as they ran between two separate, but visually similar rooms that were connected by a passageway, some place fields stopped at one wall and did not continue across the intervening space into the other room (cells 4, 12, 16 and 23 in Figure 2 and cells 1, 5 and 19 in Figure 3 of Skaggs and McNaughton, 1996). In addition, place fields near a barrier in an open arena do not extend through the barrier (unpublished observation, BL McNaughton and JJ Knierim). It could be argued that truncated place fields represent holes in the mapping of place field responses onto physical space. One difficulty with this interpretation, however, is that it requires that a special case be made for the presence of walls.

A simpler interpretation focuses on the continuity of the mappings between neural responses and the physical variable they represent. In parietal cortex, changes in responses due to gain fields alter the mapping between neural responses and the physical variables, but do not cause breaks in that mapping. If the response of two parietal units overlaps at one level of gain, that response remains overlapped at all levels of gain. In a mathematical sense, the mapping is continuous. The same cannot be said, however, for mappings in the hippocampus. Knowledge regarding the distance between the place fields of two cells in one reference frame provides no information regarding the distance between the place fields of those cells in another reference frame. In this sense, the definition of reference frame differs between parietal cortex and hippocampus. The use of the term "map-like representation" to describe ensemble properties of hippocampal place fields combines the properties of a continuous coordinate system within a given reference frame, though neighbor relationships are not preserved across different hippocampal reference frames. In this dissertation, the term "hippocampal map" means any distribution of place fields (including the absence of a place field) for a given set of place cells.

DNMS

In addition to place-related activity, a role for the hippocampus in non-spatial tasks has also been suggested. While it can be argued that all events occur at some location, and so must always contain some spatial component (O'Keefe, 1999), performance of tasks that emphasize object recognition regardless of the location of the object or its relation to surrounding spatial cues has been shown to be susceptible to

hippocampal damage. These studies are mentioned to avoid the misperception that the hippocampus is known to respond only during spatial tasks. One such task is the Delayed Non-Match to Sample (DNMS) task in which an animal is shown an object, and then after a delay is again shown that object in a new location, along with a new object. The animal is rewarded upon choosing the new, novel object. The choice is presented in a different location than the sample, to control for choice based on spatial information. Monkeys with lesions to perirhinal cortex cannot acquire this task (Meunier et al., 1993). Monkeys with lesions to the hippocampus show no deficit if they receive pre-training on the task prior to the lesion (Murray and Mishkin, 1998), while hippocampal monkeys that did not receive pre-training did show deficits (Zola and Squire, 2001). Combined, these results suggest that an intact hippocampus is required for acquisition of the DNMS task, but that the perirhinal cortex is capable of performing visual object discrimination if an intact hippocampus was present during training. In rats, Squire et al. (2001) found that hippocampal lesioned rats performed as well as controls when the delay between sample and choice was only four seconds, but were impaired significantly compared to controls at delays of one and two minutes. They proposed that the hippocampus was not necessary for learning the rule governing the task (the "non-match" rule), but was necessary for recall of the cued item following a significant delay. They suggested that performance was not impaired at four seconds due to ongoing activity patterns of the matched item in cortical structures. Data from both monkeys and rats suggest that acquisition of the DNMS tasks requires an intact hippocampus, but that long term memory for the task can be accomplished by extra-hippocampal structures.

Sequence-related activity

Hippocampal activity has been observed in relation to sequences of elements or places, particularly in the relation to sequences of locations visited during the course of a task. Place fields visited repeatedly in the same order while running on a linear track exhibited overlapping periods of spiking, and so exhibited temporally as well as spatially correlated firing. When the activity of these same units was recorded during sleep preceding and following the task of running on the maze, pairs of units with non-zero spatial correlations on the maze (units whose place fields overlap) showed increased temporal correlations during sleep following the task when compared to sleep prior to the task (Wilson and McNaughton, 1994). When the order of firing between pairs of correlated neurons was considered, the spatial order of firing during the task was preserved in the temporal order of firing during subsequent sleep during (Skaggs and McNaughton, 1996); (Lee and Wilson, 2002); i.e., if two place cells were visited in the order "place cell A, then place cell B" during the task, firing during subsequent sleep preserved this order.

Theta

Theta generation

The order in which place cells become active during running is encoded in another way, namely by the phase relationship between place cell spiking and a rhythmic (7-10 Hz), hippocampal oscillation called "theta" (O'Keefe and Recce, 1993), discussed below. Since the first description of theta (Saul and Davis, 1933), studies of the theta rhythm can be broken into two main categories: those concerned with the physiological generating mechanisms, and those concerned with behavioral correlates of theta (for an historical review, see (Bland, 1986). The mechanisms underlying the generation of theta are not completely understood, but include contributions from at least three structures: entorhinal cortex (EC), the medial septum, and CA3 (for review see (Buzsaki, 2002). As described previously in this review, the superficial layers of EC send glutamatergic projections to DG, CA3 and CA1 of the hippocampus. Urethane, which attenuates glutamatergic synaptic transmission to a small extent (Moroni et al., 1981), spares a 4-7 Hz hippocampal oscillation (Kramis et al., 1975). In contrast, when the cholinergic projections of the medial septum to the hippocampus are blocked by the application of the muscarinic antagonist, atropine, a 7-12 Hz oscillation remains. The 7-12 Hz oscillation arising from glutamatergic EC projections has been labeled the "atropine-resistant" component of theta, while the 4-7 Hz oscillation has been labeled the "atropine-sensitive" component. In addition to these extra-hippocampal theta generators, voltage-dependent oscillations in CA3 pyramidal cells recorded in hippocampal slices show that the hippocampus itself plays an active role in the generation of theta (Leung and Yim, 1986).

Theta during movement

A long history of literature has associated rhythmic hippocampal activity with voluntary movement (Vanderwolf and Heron, 1964;Yokota and Fujimori, 1964) and the coordination of sensory stimuli. In relation to different classifications of movement, Vanderwolf (1975) proposed two types of behaviorally modulated theta: Type I was associated with voluntary movements (e.g., running, jumping, but not grooming), was blocked by urethane anesthesia, but was resistant to atropine. Type II theta was associated with searching movements (e.g., sniffing and whisker flicking), and was blocked by atropine (Vanderwolf et al., 1975). Changes in the amplitude and frequency of theta have been observed during voluntary, goal-directed movements. When dogs were trained to avoid an aversive shock by pressing a lever for one conditioning stimulus, but to not press the lever for a second conditioning stimulus, theta was evoked only for the first condition; i.e., only when voluntary movement was required (Dalton and Black, 1968). In rats trained to avoid aversive shock by jumping out of a conditioning chamber, theta was present both before and during the jump, but increased in frequency, reaching a peak just prior to the jump (changing from 6 Hz to 10 Hz) (Vanderwolf, 1969). Rats trained to jump up to a ledge from an electrified grid on the floor of a conditioning chamber showed increases in both the amplitude and frequency of theta that peaked just prior to the initiation of the jump (Vanderwolf, 1969). When rats were trained to jump to different heights under different weight-loadings, the change in frequency correlated with the speed of the response and the distance jumped (Morris, 1976). In rats trained to lever press for rewarding electrical brain stimulation, theta was evoked at least one second prior to the lever press, and increased in amplitude and frequency just prior to the lever

press (Buno and Velluti, 1977). The increase in amplitude and frequency correlated with lever presses occurred only at a fixed theta phase, suggesting that the changes in amplitude and frequency were associated with a phase locking between movement and theta, though this result was not analyzed quantitatively. Quantitative analysis of theta phase and lever pressing for food reward showed significant phase locking of the lever press to the "negative wave" of theta (the half-cycle associated with maximal activation of hippocampal pyramidal cells) in seven of nine rats (Semba and Komisurak, 1978). Not all reports, however, have associated changes in theta with voluntary movements. When water-restricted rats were conditioned to run to a window for water reward, changes in theta frequency correlated more strongly with the presentation of the CS (a 500 Hz tone) than with motor activity towards the reward location (Buzsaki et al., 1979), suggesting that theta was associated more with the coordination of sensory information than with the coordination of motor activity. Rats trained on both a reference memory task (lever pressing associated with whether the cue was a light or a tone) and a working memory task (lever pressing based on whether two, sequential cues matched; i.e., "light-light") showed theta phase resetting only in the working memory task (Givens, 1996), even though similar, voluntary movements occurred in both tasks. It should be noted, however, that a subsequent study failed to replicate this result, finding that the phase of theta was reset by the onset of a light alone (i.e., the "reference memory" task) (Williams and Givens, 2003).

Phase precession

A relation between place cell activity and theta was established with the discovery of hippocampal phase precession. In rats shuttling back and forth on a linear track, spikes from place cell fired in progressively earlier phases of the theta cycle, "precessing" backwards in phase as the rat moved through the place field (O'Keefe and Recce, 1993). Phase precession occurred in most active CA1 pyramidal cells, was present in granule cells of the dentate gyrus, and did not change linearly, but rather accelerated in the rate of phase change as the field was traversed (Skaggs et al., 1996). This latter result suggested that the mechanism underlying phase precession could involve multiple components, and subsequent work has shown that the relation between theta phase and location within a place field exhibits a bimodal distribution (Yamaguchi et al., 2002).

Summary

Parietal cortex and hippocampus receive highly processed spatial information from multiple modalities. Lesions to either structure produce deficits in the performance of spatial tasks, and imaging studies show the involvement of both structures during the performance of spatial tasks. Single unit recordings in each structure show spatial correlates, though the properties of these correlates differ. Parietal cortex transforms reference frames, starting at or near the level of sensory receptors (such as the hand or the eye) and ending at the level of more abstract frames, including a world reference frame. Hippocampus maintains a map-like representation of an environment that associates ongoing activity in restricted, neuronal groups with the constellation of sensory cues and internal movement-related signals (possibly through an LTP-like mechanism) to distinguish different environmental contexts. Trajectory computations could utilize the

processing that occurs in each structure to represent starting location (place fields), goal locations (possibly by representing distant goals in a manner similar to hippocampal place fields, or by representing the distance to a goal in parietal cortex), and trajectories (possibly computed by vector subtraction through receptive field shifting in parietal cortex). Computations concerning the same trajectory could occur in both structures simultaneously, and be synchronized by oscillations in global activity, as signified by rhythmic field potentials (e.g., theta). Both parietal cortex and hippocampus have access to spatial information needed for such computations, but evidence regarding the role they play in such computations is lacking.

Implications for goal storage and recall

How might the known responses of parietal and hippocampal neurons relate to the encoding and recall of goal locations, and to the computation of trajectories to those goals? Trajectory computation encompasses several general questions in the field of neurophysiology: How are parts of the world represented in the brain? How are those representations stored? How are they recalled? How are the recollections used? Research on parietal cortex has focused on how representations of space are transformed from one reference frame to another (i.e., how representations are used), but many questions remain regarding how the results of such transformations are stored or recalled to aid future computations. Hippocampal research has focused on how activity patterns are formed and recalled, particularly in the case of place cells, but little is known regarding how these activity patterns are actually used by animals to perform tasks. In one of the

earliest models to combine the function of both neocortex and the hippocampus, David Marr suggested that neocortex performed unsupervised clustering to allow distinctions to be formed and made on different classes of stimuli (Marr, 1970). Marr then proposed that the hippocampus aids the storage and recall of members of those classes through the completion of incomplete versions of previously stored patterns (Marr, 1971). McClelland and McNaughton (1995) proposed that these functions were complimentary to each other in that the hippocampus completed partial patterns that had been previously stored in the neocortex, allowing learning during repetition to be cumulative. McClelland and McNaughton predicted that activity patterns in neocortex and hippocampus should be correlated during periods when the two structures were recalling and storing similar activity patterns, such as during sleep. Single units recorded simultaneously in both parietal cortex and hippocampus of rats showed that the pattern of correlations between parietal cortex and hippocampus in sleep following maze running matched the pattern observed during running more closely than the pattern observed in sleep prior to running (Qin et al., 1997). In non-human primates, a similar result was observed for correlations within and between neocortical regions. In motor, somatosensory, and parietal cortex (but not prefrontal cortex), the pattern of correlations during a rest period following the performance of a task matched the pattern of correlations observed during the task more closely than the pattern observed prior to the task (Hoffman and McNaughton, 2002). This suggests that the hippocampus may orchestrate the replay of activity patterns in neocortical structures, just as activity patterns are replayed within the hippocampus (Wilson and McNaughton, 1994).

Goals in the hippocampus

Codes for goals

The simplest hypothesis would be that activity related to a goal, either through sensory stimuli or recollection of a previous visit, activates specific hippocampal neurons. This activation could be similar to that of place cells responding to the current location of the rat. Goal-related activation could either activate place cells with place fields near the goal (McNaughton et al., 1994a), or cause a shift to a goal-centered frame of reference (Gothard et al., 1996a). The first hypothesis would be the simpler of the two, because it would not require the switching of map-like reference frames during the computation of a trajectory. This hypothesis would, however, require some other means of distinguishing goal-related activity from activity related to the current location. The latter hypothesis would require two different map-like reference frames during the computation of a trajectory: one for the current location of the rat, and the second for the intended location. This hypothesis would have the advantage of making a clear distinction between current position and goal locations. One possibility would be that the two reference frames are activated at different phases of theta.

Ensembles of place cell activity can be associated with context

Alternatively, goals could be expressed in the hippocampus through the control of place cell maps, activating goal-specific maps within the hippocampus. Several manipulations produce multiple maps for the same environment, , though the purpose of representing the same physical location with different distributions of place fields is unknown. In each case, the location of a place field of a given cell in one map provides

no information as to the location of the place field for that same cell in a different map. A given cell may even be active in one map, but not in another. Rats that were trained to search randomly in an open arena, and then to run first clockwise and then counter-clockwise towards the four compass points of the same arena (Markus et al., 1995). Though the environment remained fixed, the distribution of place fields was different for each of the three conditions.

Direct testing of whether place fields exist in the hippocampus of non-human primates, or computation using gain fields exists in rodent parietal cortex has been confounded by species-task differences. Parietal single-unit recordings regarding spatial computations are normally made while chaired, non-human primates perform saccade and reaching tasks, neither of which is well-suited to rodents; rodents do not make large saccades, or commonly use their forepaws to make extended reaches to distant objects under visual guidance. Hippocampal single-unit recordings involving open-field navigation are normally made in freely moving rodents, and are difficult to replicate in non-human primates due to a number of technical difficulties. These difficulties include the lack of sufficient telemetry bandwidth to remove the need for recording cables, which can be grasped easily by monkeys. The computation and execution of whole-body, ballistic trajectories in rodents offers an opportunity to simultaneously test parietal and hippocampal representations of space, because it combines the goal-directed tasks that are typical of parietal research with the spatial exploration tasks that are typical of hippocampal research.

Chapter 3 GENERAL METHODS

Pre-training

Four, male Brown-Norway/Fisher hybrid rats (6-9 months old, 250-350g) were food-deprived to 80% of their *ad libitum* weight. Rats were housed in plexiglass home cages, and maintained on a reversed, 24 hour light-dark cycle. Training and recording sessions occurred during the "dark" portion of this cycle. Rats were trained to find food pellets at one of eight, equally spaced reward zones on the edge of a 1.3 m diameter circular arena (one rat was trained to run to 16 equally spaced reward zones). Reward zones were marked by wooden clothespins, each containing a Light Emitting Diode (LED) and attached to the edge of the arena. Each LED was located approximately 5 cm above the table surface. A training session lasted approximately 30 minutes, and contained a number of separate trials. Each trial began when two, simultaneous cues were activated: a 4 kHz tone signaled the availability of reward somewhere in the arena, and the illumination of one LED marked the correct reward zone. To help rats locate the visual cue, the LED marking the correct reward zone flashed at 1 Hz. Rats were trained to run to the vicinity of the correct reward zone (within 10 cm), and were given food reward upon arrival. This training continued until each rat made direct trajectories to reward locations, and more than 50 rewards were received in a training session.

Surgery

NIH guidelines were followed for all surgical procedures. Each rat was injected with Bicillin (0.1 cc i.m. per hind leg, Wyeth Laboratories), anesthetized with Nembutal (sodium pentobarbital, 0.8 cc/kg, Abbott Laboratories), and placed in a stereotaxic holder. The skull was cleared of skin and fascia, craniotomies were opened over the stimulating and recording sites, and the stimulating electrodes and hyperdrive were implanted and cemented in place using dental acrylic. Two pairs of stimulating electrodes (two were placed into the Medial Forebrain Bundle (MFB, +0.25 mm AP, 1.9 mm ML, 8.5 mm DV, 19.5° projecting caudally). Each stimulating electrode consisted of two, teflon-coated, stainless steel wires (coated diameter .0045", 316SS3T, MedWire, Inc., Mt. Vernon, NY) twisted together with approximately 1mm of insulation removed at the tip. A "hyperdrive" (Gothard et al., 1996b) containing twelve tetrodes (McNaughton et al., 1983b) was placed dorsal to the hippocampus (∓ 2.0 mm ML, -4.0 mm AP relative to Bregma, over the right hemisphere in three rats (Rats 1, 3 and 4) and over the left hemisphere in one rat (Rat 2)). Each tetrode consisted of four, teflon-coated, nichrome wires (diameter 14μm, Kanthal Palm Coast, Inc., Palm Coast, FL) twisted together. This location is anterior to cortical areas Oc2MM and Oc2ML (Zilles, 1985), and is also known as area PTLp (Burwell et al., 1995). Rats were allowed to recover for one week following surgery. Food restrictions were removed following surgery, and rats were allowed to resume feeding *ad libitum* during all subsequent training and recording.

Stimulation parameter training

Because the motivation level achieved by MFB stimulation was difficult to evaluate subjectively, a testing protocol was developed to ensure equivalent, high motivation levels across rats. Following recovery from surgery, stimulation efficacy was evaluated by placing each rat in an operant conditioning chamber equipped to deliver MFB reward stimulation (Liebman and Cooper, 1989). Rats 1 and 2 were trained to lever press for reward, but each rat required one week of training to acquire this task. Rats 3 and 4 were trained to poke their nose into a small opening (interrupting an infrared beam) for reward, which was learned within two days by both rats. MFB stimulation consisted of a train of bipolar current pulses with a 100 μsec pulse width. Stimulation frequency (150 Hz) and duration (300 msec) were held constant, while current intensity was varied randomly from 50 to 120 μA across two-minute trials. The number of responses per second (lever presses or nose pokes) was plotted as a function of current amplitude and fitted with a third order polynomial using least squares approximation (Figure 3-1).

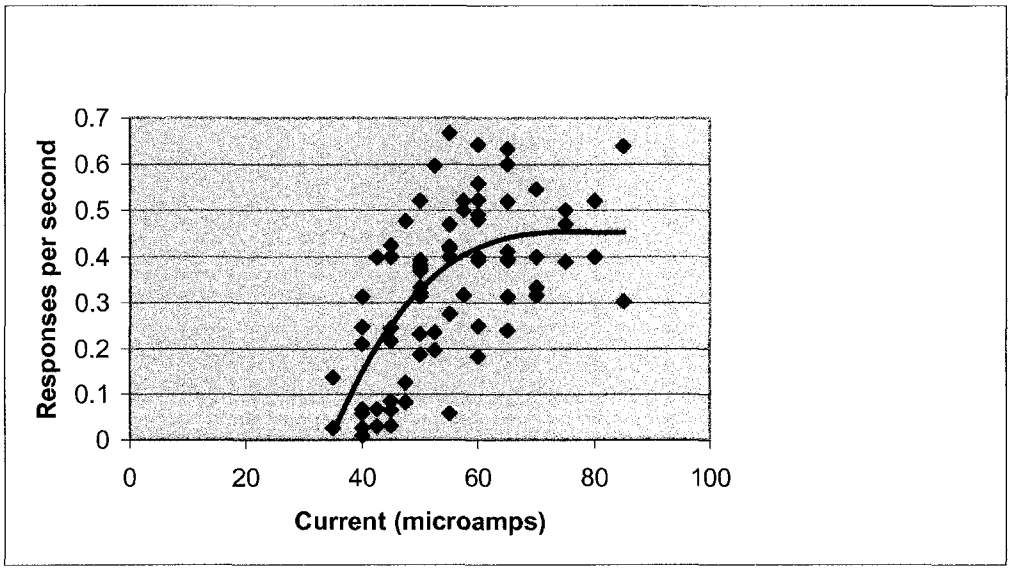

Figure 3-1. A quantitative method was used to determine Medial Forebrain Bundle (MFB) stimulation current intensity. The horizontal axis shows the stimulation current amplitude. This amplitude was varied randomly across two-minute sessions. The number of responses per second is shown on the vertical axis. Blue dots represent individual training sessions. The black line shows a least-square polynomial fit to the data. The smallest current that approximated the asymptotic response maximum was used during all subsequent sessions (in this case, at 70 µA).

The minimal current amplitude that approximated the asymptote of responses per minute was used during subsequent recording sessions. This amplitude ranged from 70 to 110 µamps. After determining the appropriate current setting, rats resumed training on the task learned during pre-training, but were rewarded with MFB stimulation, rather than with food.

Recording system

During recording sessions, the hyperdrive was connected to a recording headstage (Neuralynx, Inc.) that allowed low-noise transmission of neural data to the recording system, and connection from a stimulation isolation unit (SIU) (A365D-A, World

Precision Instruments, Inc., Sarasota, FL) to the MFB stimulating electrodes (Figure 3-2). The headstage also contained an array of LED's that could be detected by an overhead CCD camera, allowing the tracking of the position of the rat on the maze. All data were recorded using a Cheetah recording system (Neuralynx, Inc., Tucson, AZ) running in combination with a Pentium-based PC. Single unit data were recorded from each tetrode at 32 kHz and filtered between 0.6 and 6 kHz (Assembly Hunter amplifiers, Neuralynx, Inc., Tucson, AZ). Single units were recorded with respect to a reference electrode placed below cortex and above CA1, roughly 1 mm below the nominal brain surface. Video position was recorded using an overhead camera sampling at 60 frames/sec, which tracked the LED's located on the headstage. The same clock generated timestamps for both single unit activity and video data with a temporal resolution of 0.1 msec. Video spatial resolution was approximately 3 pixels/cm. Electroencephalogram (EEG, also known as Local Field Potential, or LFP) data were collected from each tetrode, along with data from an electrode located in the hippocampal fissure (roughly 320 μm below the CA1 cell body layer, *s. pyramidale*). EEG data were sampled at 1 kHz, filtered between 1 and 300 Hz, and were recorded with respect to the same reference electrode that served as the single unit reference electrode.

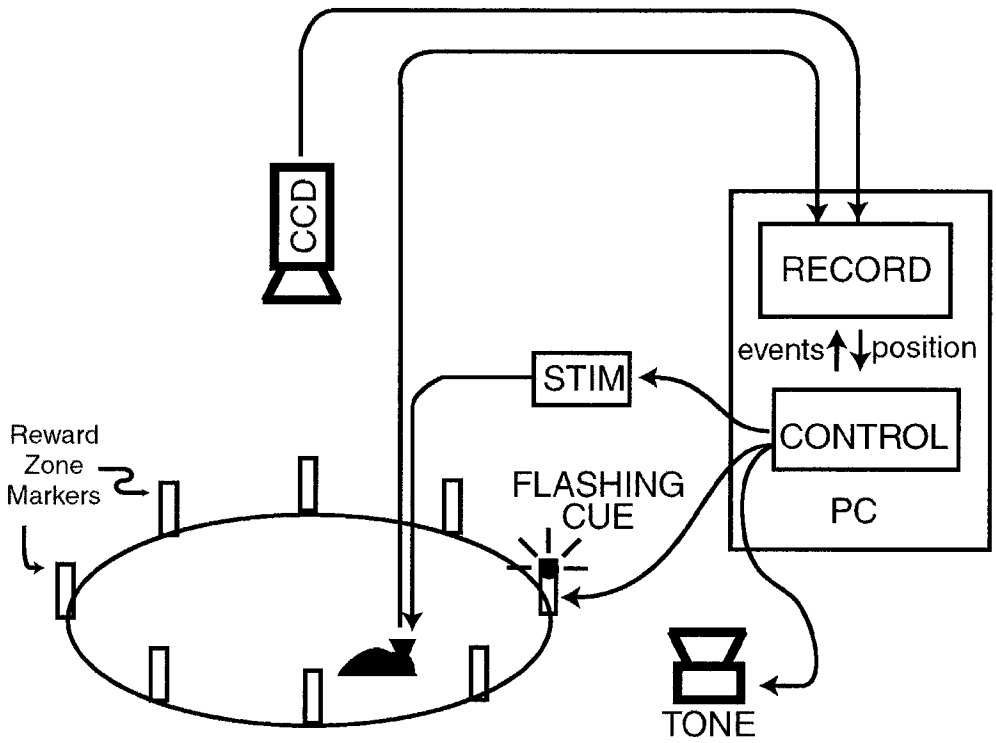

Figure 3-2. Schematic of experimental setup used to record data and control the experiment. Each rat moved freely in a 1.3 m diameter, open arena. During recording, the hyperdrive was connected to a recording headstage that allowed low-noise transmission of neural data, and also provided a connection from the stimulation isolation unit (SIU) to electrodes implanted in the Medial Forebrain Bundle (MFB) for electrical brain stimulation reward. The headstage also contained an array of Light-Emitting Diodes (LED's) that allowed the position of the rat to be tracked using an overhead CCD camera. Data from multiple single units and EEG were transmitted through a lightweight tether that was counter-balanced to reduce the weight load on the head of the rat. The experiment was controlled by a Visual Basic program ("CONTROL"). At the start of a single trial, the control program initiated a 500 msec, 4 kHz tone from a speaker located on the floor near the maze, and simultaneously illuminated an LED that was mounted in a wooden clothespin attached to the edge of the maze. This LED flashed at 1 Hz, and signaled the current reward location. At the same time, CONTROL sent a signal ("events") to RECORD, attaching a timestamp synchronized to both single unit and video data. (During the Sequence experiments, the onset of flashing was delayed by 5 sec, allowing the rat time to run to the correct reward zone without aid of visual cue.) Using video data from the CCD camera, RECORD tracked the location of the rat and sent this information to CONTROL. When the rat reached the vicinity of the current reward zone, CONTROL initiated a 300 msec reward stimulation ("STIM") of the MFB, simultaneously sending an event to RECORD, which attached a timestamp, marking the end of the trial. This timestamp was also used in post-processing to remove recording artifact arising from MFB stimulation.

Each tetrode was advanced through parietal cortex towards the hippocampus from 0 to 160 μm per day. While each tetrode was being advanced, the EEG signal from that tetrode was passed through an audio amplifier (Grass, Inc., West Warwick, RI), allowing

the progress of the electrode to be heard. On days when multiple single units were observed in parietal cortex, these units were recorded during that training session. As the cell body layer of CA1 was approached, 200-300 Hz oscillations ("ripples") could be heard over the audio monitor, and sharp waves observed on the oscilloscope (Model 310, Nicolet, Inc., Madison, WI). Tetrodes were then advanced into the cell body layer until multiple single units were observed.

Recording Sessions

Data were collected each day through a series of epochs collectively called a recording session. At the start of each recording session, an experiment control program (described in the next section) was started and initialized for the task to be run that day. The position of each possible reward zone around the edge of the maze was described by a polygon of points in video space centered on the image of the physical cue (i.e., the clothespin). During an experiment, the experiment control program monitored the spatial location of the rat, and noted when the position of the rat moved within the vertices of a reward zone polygon. The vertices of this polygon could be moved by the experimenter to make the reward zone as large or small as desired. In general, reward zones extended approximately 10 cm from each clothespin, and the same dimensions were fixed across sessions.

Following initialization of the control program, the rat was brought into the recording room and placed in a towel-lined bowl in the center of the recording arena. The headstage was attached to the hyperdrive, and electrical power was supplied to the

low-noise amplifiers and LED's on the headstage. The signal gain and threshold of each tetrode was checked to maximize the amplitude of units without exceeding the dynamic range of the amplifiers; i.e., without "saturating" the amplifiers. Data recorded was initiated and the rat was allowed to rest for 20-30 minutes in the towel-lined bowl. This epoch was called "Sleep 1," though there was no requirement that the rat had to sleep during this time. Following Sleep 1, the rat was moved from the bowl and allowed to run freely on the arena surface. The length of the lightweight tether was adjusted to allow free movement to all parts of the arena without an excess of slack when the rat was in the middle of the arena. The behavioral phase, called "Maze", began when the first cue (consisting of a light and tone) was given. Maze epoch lasted from 30 minutes to one hour. Following the Maze epoch, the rat was returned to the towel-lined bowl, and allowed to rest for another 20-30 minutes. This second rest session was called "Sleep 2." Following Sleep 2, the power to the headstage was turned off, the headstage was disconnected from the headcap of the rat, and the rat was returned to its home cage. A recorded checklist can be found in Appendix 1.

Experimental Control

To limit the time during which a trajectory could be computed, information regarding the next, correct reward zone was made available to the rat in a specific, controlled manner. A computer program controlled the delivery of all experimental cues to the maze and reward stimulation to the rat. For each experimental event (cue onset or stimulation), the program simultaneously sent a 5V, TTL pulse to the recording system.

The recording system assigned a timestamp with 0.1 msec resolution to each event and recorded the event and timestamp in an event file. Single unit recordings, the position of the rat obtained through video recordings, EEG data obtained from each tetrode, and all experimental events were given timestamps from the same clock, synchronizing all data streams.

Post-Processing

MFB stimulation produced an erroneous signal, or "artifact", in the record of each tetrode. The stimulation-induced artifact was removed by deleting all events occurring within one second following each stimulation. Because stimulation duration was 300 msec, this provided sufficient time for the fading of any transients due to the presence of high frequency stimulation signals. Putative single neurons were isolated using an interactive procedure on spike waveform and amplitude parameters, producing a collection of timestamps associated with each action potential, or "spike", from a given unit. Position during each video timeframe was extracted by finding the center of mass of the image of the headstage, and this position was associated with the timestamp of the video frame. Because both spike and position data shared timestamps generated by the same clock, the position of the rat on the maze at the moment that each spike occurred could be inferred by interpolating the position timestamps nearest to each spike timestamp.

Behavioral analysis

The behavior and motivation level of each rat varied during a recording session. Grooming sessions occurred during each session, producing individual trials that lasted for more than a minute. On some occasions, rats would simply rest for a few seconds. To ensure that data used for analysis came only from runs during which the rat was highly motivated to find the goal (referred to as being "on-task"), a set of criteria was established to remove trials containing other behaviors (referred to as being "off-task"). On-task runs satisfied three criteria: spatial accuracy, directed behavior, and temporal constraint. Spatial accuracy was judged by noting entry into zones other than the correct target zone. If the correct zone was not found initially, rats often searched surrounding edges of the maze, entering multiple wrong zones. In several cases, however, rats ran through a neighboring zone en route to the correct zone. A run passed the spatial accuracy criterion if the rat entered no more than one zone prior to the correct zone. The second constraint, referred to as "directed behavior", identified those runs where the rat stopped during the run, often in the middle of the maze. Such stops appeared to correlate with indecision regarding the correct zone, and led to multiple, shorter trajectories during a given run. Stops were identified by a drop in velocity below an average of 0.2 m/sec (max velocity during runs was approximately 0.6 m/sec), and runs with more than one such stop were considered to be invalid runs. Finally, on some runs, rats paused prior to the start of movement. To ensure that analyses were conducted only on runs where directed runs to the correct target were made, runs were considered valid only if the target zone was reached within five seconds of the tone onset. A characteristic valid run

(Figure 3-3) consisted of the following sequence of events: cue (tone/light) onset, movement onset, peak acceleration, peak velocity, and entry into the correct zone, followed immediately by MFB stimulation.

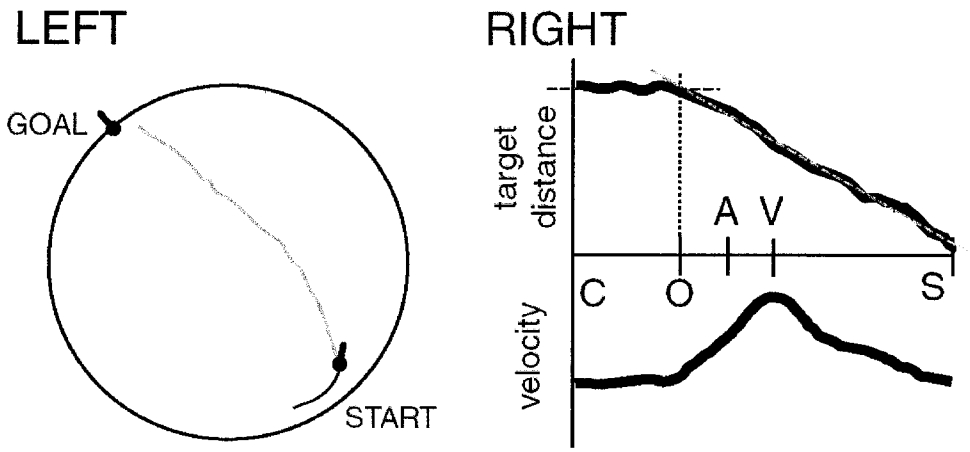

Figure 3-3. (LEFT) Schematic of a characteristic "valid" run. The red line shows the path taken by a rat from cue onset to the onset of uninterrupted movement towards the goal. The green line shows the path taken by the rat from the onset of uninterrupted movement to the arrival at the goal zone. (RIGHT) Computation of the key behavioral events during an on-task run. Each on-task run consisted of a series of events that could be identified either from computer-generated event time stamps or from the video recording. Each trial began with the simultaneous cue presentation of a 4 kHz tone, and illumination of the LED at the correct reward zone. (Upper Curve) The distance to the target computed for each video frame was computed by finding the Euclidean distance between the location of the rat in a video frame and the location of the rat at the end of the run. The onset of uninterrupted movement toward the goal was computed by finding the intersection of two lines fitted to the computed goal distance. This algorithm was based on the observation that while the rat was running to the goal, the distance to the goal decreased monotonically. The horizontal line (dotted red) approximates the distance to the goal prior to the onset of uninterrupted movement, during the time that the rat continued searching for the visual cue. The angled line (dotted green) approximates the distance to the goal during the uninterrupted trajectory. (Lower Curve) The velocity of the rat was computed for each video frame by finding the distance between the location of the rat in the preceding frame and the location in the following frame. Acceleration was computed by finding the difference in velocity for the preceding and following frames. The maximum acceleration following the onset of movement marked the moment of peak acceleration. The maximum velocity following the moment of peak acceleration marked the moment of peak velocity. Each trial ended with the delivery of Medial Forebrain Bundle (MFB) stimulation.

A valid trial consisted of a series of behavioral events. The components of a valid trial are defined in Figure 3-4. A trial consisted of a series of experiment control

timestamps and behavioral events. Each trial began with a cue onset and ended with reward stimulation. A series of behavioral events was collected in each "trajectory" beginning with trajectory onset and ending with reward stimulation.

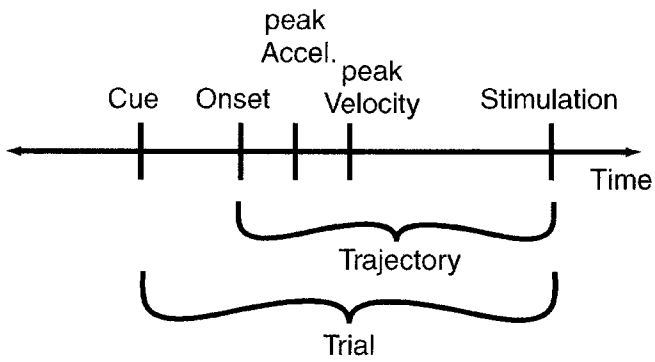

Figure 3-4. Trial-related definitions. Each trial during the Random task consisted of series of events in time, beginning with the "Cue". The rat located the target, and began moving toward that target at the trajectory "Onset". The rat accelerated toward the target, reaching a "peak acceleration" prior to reaching a "peak velocity". The rat arrived at the target and received reward, called "Stimulation".

Anatomical verification

Of the four rats whose data contributed to this dissertation, only one (rat 4) underwent histological verification that recording electrodes were located in CA1 of the hippocampus. The head cap of rats 1 and 2 became unstable due to bacterial infection and became detached, preventing accurate histological verification of the location of recording electrodes. Rat 3 was used in subsequent experiments in which the tetrodes were moved into the Dentate Gyrus, preventing histological verification that the recordings from rat 3 presented in this dissertation came from CA1. In addition, the physiological properties of putative hippocampal neurons were compared to the known

properties of hippocampal units (Ranck, Jr., 1973;McNaughton et al., 1983a). The physiological properties of putative parietal neurons were compared to those of hippocampal neurons, as described in Chapter 7 of this dissertation.

Chapter 4 RAT HIPPOCAMPAL PLAC CELLS DO NOT ANTICIPATE GOALS OF TRAJECTORIES

Abstract

Hippocampal units respond selectively to place and are an important component of the cognitive map theory, which proposes that the hippocampus provides a map-like representation of an environment. Such a cognitive map could be used to compute trajectories to distant goals within the hippocampus itself, provided distinct representations of current and intended locations could be maintained prior to the initiation of movement. A representation of distant objects must exist somewhere in the rodent brain because rats that have visually acquired a target can continue to that target in the dark, and can navigate to unseen goals from different starting locations. Multiple single units were recorded from hippocampus while rats ran for electrical brain stimulation reinforcement to distant targets, presented either in random order ("Random" task) or as segments of learned sequences (one of two "Sequence" tasks). During the Random task, no goal-related activity was observed during the early, planning stages of trajectories, and no correlation was observed between the number of spikes that occurred prior to entry into a place field and the intended, nearest approach to that place field. During two Sequence tasks, rats learned two, eight-element sequences of locations. Performance was disrupted by rotation of the sequences with respect to distal cues, suggesting that the rats used a spatial strategy to recall locations in each sequence. There was no difference, however, between the number of spikes occurring at different stages

of Sequence trajectories as compared to Random trajectories. In addition, during Sequence tasks, no difference in activity was observed between trajectories guided by visual cues and trajectories recalled from memory. These data suggest that place-related activity in the rodent hippocampus correlates specifically to the current location and context, and that the rodent hippocampus does not generate a representation of distant goals that could be used at the moment the trajectory to the goal is planned or initiated.

Introduction

How does the brain compute trajectories to a distant target? Physiologically plausible models of trajectory computation require that both current and intended locations be represented in neuronal activity (McNaughton et al., 1994a;Redish and Touretzky, 1994;Salinas and Abbott, 1995).

Although rats rely on vision less than do primates, the appearance and arrangement of distant, visual targets do impact rodent behavior during spatial tasks. Rats that have visually acquired a distant target and started a trajectory towards that target under lighted conditions continue to that target in subsequent darkness, and can locate targets based on the arrangement of local cues from different starting locations (Collett et al., 1986). Rats can also learn to navigate to unseen goals, such as the submerged platform in the Morris water maze (Morris, 1981); however, acquisition of a memory for the goal location requires an intact hippocampus (DiMattia and Kesner, 1988). In addition, distant objects can exert powerful control in stabilizing ensembles of hippocampal cells, rotating (Muller and Kubie, 1987;Bostock et al., 1991) or shifting

(Muller and Kubie, 1987) the location of place fields, and causing a change in coordinate frame of place cell activity (from start box-centered to goal-centered) midway through a trajectory (Gothard et al., 1996b). The discovery that specific neurons ("view cells") respond when a monkey looks toward a remote spatial location (Georges-Francois et al., 1999) raised the possibility that hippocampal activity may correlated to the location of distant goals, similar to the correlation observed between current location and place cell activity. View cells respond selectively to different, distant locations even when the monkey remains in one location; i.e., view cells respond "differentially" while the animal remains in a given location. Recently, two studies were reported in which differential hippocampal activity was found with respect to a fixed location, prior to subsequent turns in different directions (Wood et al., 2000;Frank et al., 2000), suggesting that different, distant goals might be represented in anticipation of the direction of the subsequent turn. When activity related to distant goals in the rodent hippocampus has been sought directly, however, (Bower et al., 2000;Fenton et al., 2002) no such activity has been observed. Fenton et al. trained rats to run to an unmarked location in the center of an arena, which caused the release of food pellets onto a random location in the arena, and then observed place cell activity as the rats returned to the food-release location. Place cell activity did not correlate with the location of the distant goal. This task was simplified, however, by having only a single goal to which the rat returned each time. Rodents are known to maintain return trajectories to a home location (Mittelstaedt and Mittelstaedt, 1980), possibly utilizing neural representations of "motor plans" outside the hippocampus. If hippocampal representations of distant goals arise during trajectory planning, they would

be more readily observable in tasks with an increased spatial processing load, either due to the presence of multiple possible targets, or by requiring that targets be recalled from memory.

Networks containing recurrent connections can perform vector subtraction, and can also recall memories of previously visited locations from incomplete input patterns (Marr, 1971). The presence of recurrent connections in the CA3 region of hippocampus suggests two different models for trajectory computations that could occur in the hippocampus (Figure 4-1).

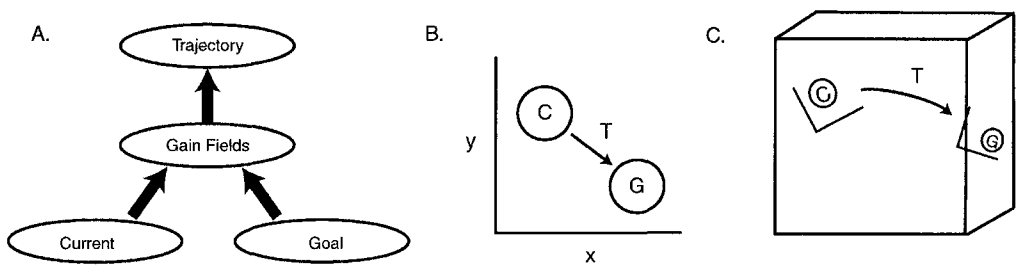

Figure 4-1. Possible network implementations for trajectory computation. A. Trajectories can be computed by using gain fields to perform coordinate transformations. In models of parietal cortex, information regarding the current location of, say, the hand that is stored in one coordinate system (e.g., "arm-centered") is combined with information regarding the goal that is stored in another coordinate system (e.g., "eye-centered"). Parietal cortex receives both inputs, and computes a trajectory to move the hand to the target in the "hidden layer" of the network, producing units with gain field responses that are similar to those of units recorded in parietal cortex (Zipser and Andersen, 1988). B. The hippocampus could compute trajectories in a similar manner by activating neural patterns for both the current and goal locations within the same coordinate frame. C. Alternatively, the hippocampus could contain multiple possible reference frames (the space inside the cube) and activate patterns related to the current and goal locations in different coordinate frames.

One model (Figure 4-1.B) computes a trajectory within the same hippocampal coordinate frame used by the encoding of current location. The other model (Figure 4-1.C) computes a trajectory across two different reference frames; i.e., using the terminology of Gothard et al. (1996b), current location could be encoded in "start box-

centered" coordinate frame, while the target could be encoded in a "goal-centered" coordinate frame.

Methods

Random task

Four rats were trained to run to one of eight zones (sixteen zones for rat 1) presented in random order (the "Random" task) for electrical stimulation of the Medial Forebrain Bundle (MFB). Reward availability was signaled by a 500 msec audio tone at 4 kHz, and by the onset of the LED (flashing at 1 Hz) at the reward location. This required the rat to find the flashing LED at the correct reward zone, and then make a trajectory to that zone. In addition, rats 1 and 2 were trained to run to one of two central reward locations when the audio cue was a 2 kHz tone, and the other central reward location when the audio cue was a 7 kHz tone. Neither rat learned to perform this task reliably, so trajectories to the central reward locations were dropped from analysis. Rats 3 and 4 were trained to run from and to cued zones at the edge of the maze.

Sequence tasks

After learning the Random task, three rats (all, but rat 1) were trained to run two different eight-segment sequences (the "Sequence" tasks). The definitions used to describe the components of sequence tasks in general are shown in Figure 4-2.

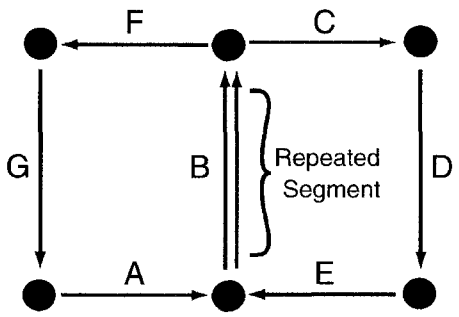

"Segment": The physical path between two zones.

"Sequence": An ordered set of segments, here ABCDEBFG.

"Repeated Segment": A segment occurring twice in a sequence.

"Choice Point": A start with two possible goals, e.g., the start of "C" and "F".

"Choice Context": A set of segments containing a repeated segment and a choice point.

"Trial": An individual run on a segment, as for the Random task.

"Lap": A set of trials containing one complete sequence, ABCDEBCFA.

"Block": A set of three laps.

Figure 4-2. Schematic and definitions of the components of a sequence task.

During sequence tasks, three complete laps through the sequence constituted a block of trials. Blocks alternated between "cued" and "non-cued" throughout the duration of the recording session. During cued blocks, reward availability was signaled by a tone and by the illumination of the LED at the correct reward site, as in the Random task. During non-cued blocks, the tone signaled reward availability, but the onset of the light was delayed by five seconds, giving the rat the opportunity to recall the correct reward site from memory. When the correct reward site was reached, the rat received MFB stimulation, followed by a one second delay, and then the presentation of the next reward site.

Evaluating sequence learning

Each rat learned two sequences (Figure 4-3), and was required to choose the correct target following the choice point with greater than 70% accuracy in the non-cued condition before beginning training on the next sequence. Accuracy was based on two criteria: zone entry and time. Zone entry applied only to trials beginning at the choice point, and required the rat to choose the correct zone for the context; e.g. during the Sequence 1 task, a zone entry error occurred if the rat entered the zone at the end of segment G while nominally running segment D. The time criterion used the combined time required to complete the sequential trials within a context, and was based on the time required for rats to complete a context during the cued condition. For Sequence 1, the time criterion for BCD and FCG was 10 seconds. For Sequence 2, the time criterion for ABCD and EBCF was 12 seconds.

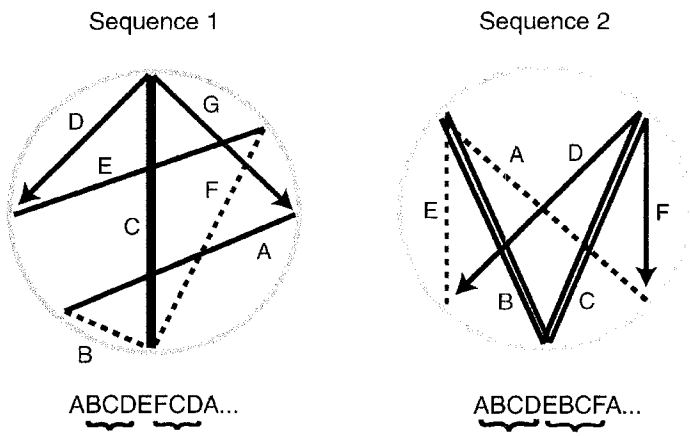

Figure 4-3. Schematics of two eight-segment sequences containing repeated segments as viewed from above the maze. Each segment is associated with a letter of the alphabet, and the associated string of letters is shown below each sequence. Each sequence can be broken into two "choice contexts," each associated with a turn in one of two directions at the end of the repeated segment (shown as red and blue lines on the maze, and brackets under the strings). The repeated segments are shown as parallel red and blue lines. (LEFT) Sequence 1 contains one repeated segment, C. (RIGHT) Sequence 2 contains two repeated segments in succession, BC.

Grouping trials

Trials were grouped according to two variables: computational stage, and physical location. When considered as a computational task, each trial consisted of a series of stages, based on behavioral events (Figure 4-4). The delay period between trials provided a "control" stage. In the Random task, the onset of the audio (and possibly a visual) cue provided sufficient information to allow the computation of a trajectory, marking the start of the "planning" stage. For Sequence tasks, the planning stage extended back to the end of the previous trial, because the rat was free to recall the next reward location without waiting for a cue. Each trajectory ended when the rat reached the goal and received reward stimulation, marking the end of the "finishing" stage. The end of the "planning" stage and start of the "finishing" stage could not be defined by any observation. Consequently, two behavioral events were selected as reasonable substitutes: The end of the "planning" stage was identified as the moment of peak acceleration during the trajectory. By this time, the goal of the trajectory was clear, suggesting that most, if not all, movement planning had been completed. The start of the "finishing" stage was identified as the moment of peak velocity, because the rat subsequently slowed in preparation for arrival at the goal.

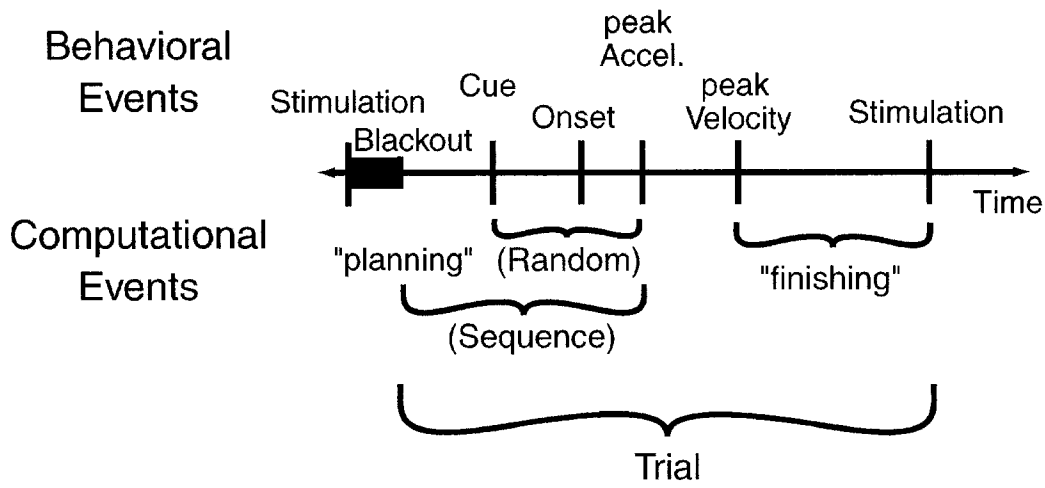

Figure 4-4. Schematic of the timeline during a trial. (Behavioral Events) Each trial consisted of a series of events that were marked by computer-generated timestamps. The first event was the end of the stimulation-artifact "Blackout" from the previous stimulation, followed by the "Cue" onset (the onset of the audio tone and the illumination of the LED at the target zone). Each trial ended with a computer-generated timestamp associated with the entry of the rat into the reward zone, and the computer-controlled delivery of MFB "Stimulation." Between these two timestamps, the occurrence of three events was computed using video tracking data, sampled at 60 frames/sec. The distance to the target was computed at each video sample. During those video samples when the rat was running to the target, the target distance decreased monotonically. Trajectory "Onset" was associated with the video sample at which this monotonic decrease in distance began. Velocity was computed as the Euclidean distance between video frames divided by the video frame rate, and acceleration was computed as the difference between velocities during consecutive video frames. During the video frames associated with an uninterrupted run to a target, acceleration reached a peak during a video frame ("peak Acceleration") preceding the video frame at which velocity reached a maximum ("peak Velocity"). (Computational Events) When considered as stages in a computation, the "planning" stage of each trajectory was associated with the time between the end of blackout (Sequence tasks) or cue onset (Random task) and peak acceleration. The "finishing" stage was associated with the time between "peak velocity" and the delivery of "stimulation" after the rat entered the correct goal zone.

For each hippocampal pyramidal cell recorded during a session, the total number of spikes that occurred between cue onset and peak acceleration was computed, constituting the spikes that occurred during the "planning" stage of the trajectory. Spikes occurring between the time of peak velocity and arrival at the goal were computed, constituting the spikes that occurred during the "finishing" stage of the trajectory.

Grouping trials: Random task

The second variable by which trajectories were grouped used the physical location of the rat at the beginning and end of the trial. This variable had different meanings for the Random and Sequence tasks. In the Random task, the rat could begin any given trial at any one of the reward zones, or could finish the trial at any one of the remaining zones. The same zone was never used as both the beginning and goal of a trial. Trials were grouped according to both the zone at which the trial began, called the "start", and the zone at which the trial ended, called the "goal". Combining the computational groupings ("planning" and "finishing") with the physical groupings ("start" and "goal") produced four groupings (Figure 4-5).

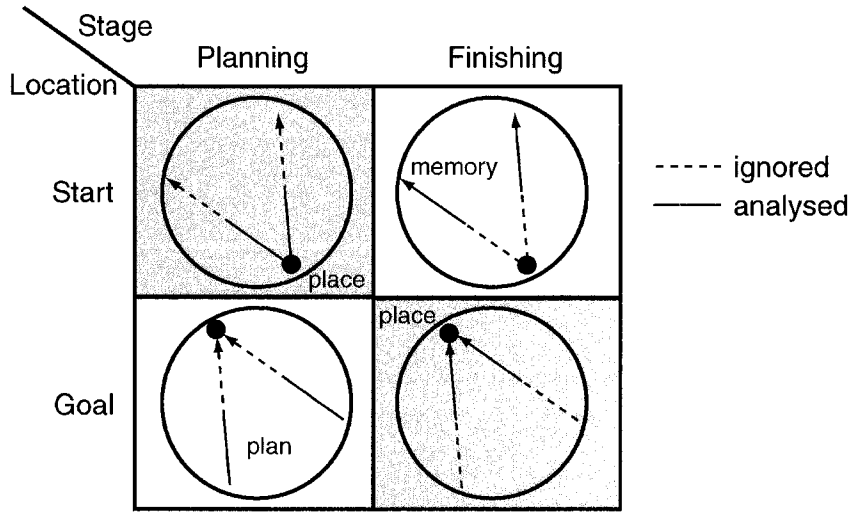

Figure 4-5. Different groupings according to physical location and computational stage. Each trajectory was divided into a "planning" or "finishing" stage. Spikes occurring during the stage of the trajectory shown as a solid line were analyzed in that condition, while spikes occurring during the stage shown as a dotted line were ignored in that condition. Spikes that were analyzed were counted and grouped according to either the start or goal of the trajectories. (Gray squares) Activity related to current location should be independent of either the forthcoming goal (upper left) or the origin of the segment (lower right). (White squares, lower left) Activity in anticipation of a common goal should produce similar activity during "planning" stages at different start locations. (White squares, upper right) Activity associated with a memory of a common starting location should produce similar activity during the "finishing" stages at different goal locations.

For each unit, spikes across all trajectories were grouped according to computational stage and physical location. Significant differences in spike counts relating to a particular zone were found using a Chi square test ($p=0.05$), where the degrees of freedom were two less than the total number of zones on the maze, because the start and goal zone were never the same. As an example, consider a unit with a place field near a specific zone. When the rat was near this zone, the activity of this unit would be expected to be significantly different than when the rat was near zones outside of the place field. This difference should be observed both during the "planning" stage when the rat starts the trial in the place field, and during the "finishing" stage when the rat runs into the place field (see the word "place" in the gray squares depicted in Figure 4-5). Alternatively, if a unit became activated during the planning of trajectories to a specific goal, then a significant difference in spike counts should be observed during the "planning" stage for trials grouped according to that goal (see the word "plan" depicted in Figure 4-5). Conversely, if a unit became activated after leaving a specific goal (i.e., perhaps as a memory of the starting location of the trajectory), then a significant difference in spike counts should be observed during the "finishing" stage for trials grouped according to "start" zone ("memory" in Figure 4-5).

Grouping trials: Sequence tasks

Trials during both sequential tasks were also grouped in different combinations. Unlike the random task, however, "Start" and "Goal" zones were fixed for many elements in the sequence. Because both sequence tasks involved repeated segments, however, trajectories from two specific starting locations ended at a common goal. In addition,

some trials during sequence tasks were cued by a flashing light at the reward location, while other trials were non-cued, requiring the rat to recall the next location in the sequence.

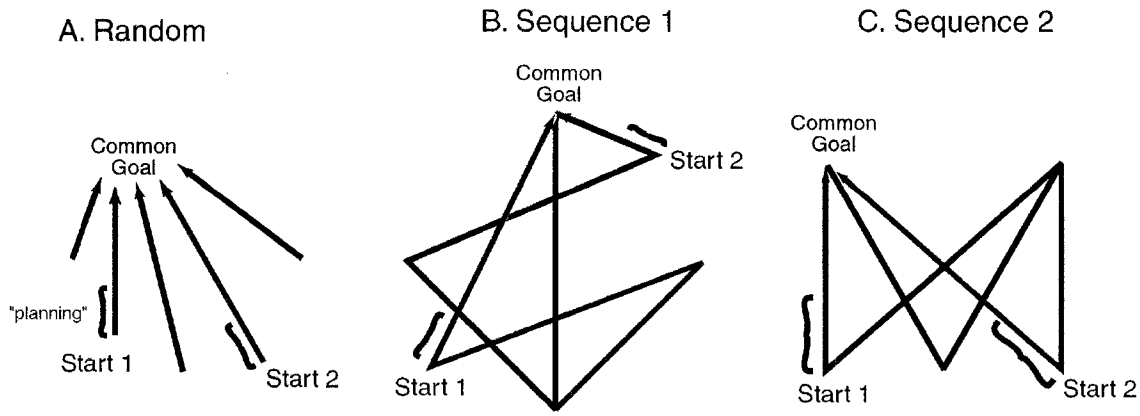

Figure 4-6. Schematic of "planning" stages located in different physical locations in each task. The trajectories from two starting points ("Start 1" and "Start 2") are noted in each case by red and blue lines. The nominal location of the rat during the "planning" stage in each trial is shown by the black brackets. (A) In the Random task, the goal location was chosen randomly for each unit. Starts 1 and 2 were chosen by finding the two starting zones having the largest number of trials ending at the common goal. (B and C) For the Sequence tasks, the two start zones were always on opposite sides of the maze.

Following the same procedure used to analyze runs in the Random task, spikes occurring during the "planning" stage of trials to a common goal were counted and grouped according to their starting location. For the Sequence tasks, a single zone was a common goal between two trajectories, so that starting locations were fixed (Figure 4-6). To allow a comparison of "planning"-stage activity during Random and Sequence tasks, goal zones were selected randomly in the Random task, and the two most frequent starting locations to that goal were identified. For comparison between Random and Sequence tasks, activity during the "planning" stage was considered significant for any single unit provided an average of more than one spike occurred for both units. This "one spike is enough" criterion was considered the minimum threshold for consistent

involvement of a unit in trajectory planning. ANOVAs were used to compare the "cued" and "non-cued" conditions during the Sequence tasks with the alpha level set at the .05 level.

Place fields

Place fields were identified using a variation of the criteria of Muller et al. (1987). The maze was divided into a 16x16 array of bins (approximately 10 x 10 cm per bin). Position data were binned and summed to generate a two-dimensional histogram of the time spent by the rat in each bin. For each unit, spikes were associated with a location on the maze by matching the timestamp of each spike event to the video timestamps. The spikes were then binned using the same 16x16 grid to generate a two-dimensional histogram of the total number of spikes that occurred in each bin. Ratemaps were generated for each unit by dividing the spike histogram by the position histogram, producing a mean firing rate for each spatial bin. A bin showed significant spatial selectivity if the mean firing rate in that bin exceeded the mean firing rate for the entire maze by more than one standard deviation. Place fields consisted of at least six such bins that were spatially connected along a side (not just a corner). These criteria were based on previously established criterion for place fields (Muller et al., 1987), and are here called the "Muller, Kubie and Ranck place field criteria". A single place field qualified a unit to be a place cell for that session, though some units exhibited more than one place field. On a given trial, a place cell was considered to have become active if it spiked at least three times within the place field associated with that cell.

Results

Random task

Behavior

Data from the Random task were obtained from four rats; data from two recording sessions were analyzed from each of two rats, and data from one recording session were analyzed from the other two rats. Data came from recording sessions in which at least 70% percent of all trials lasted less than 7 seconds (Figure 4-7). Across all rats, 1459 trials of 1680 satisfied this criterion. A total of 162 CA1 pyramidal cells were recorded.

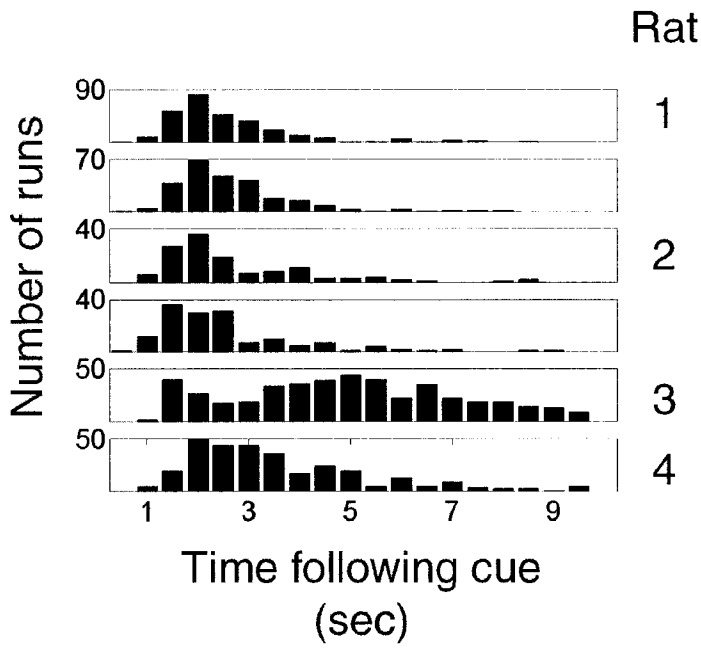

Figure 4-7. Distribution of times between cue onset and delivery of reward stimulation in the Random task. Each histogram represents the distribution of times during a single recording session. The horizontal axis shows the time (in seconds) between cue onset and delivery of reward stimulation when the rat arrived at the correct zone. The vertical axis shows the number of runs in a given bin. The longer, skewed distributions for rats 3 and 4 arose because these rats were required to run to and from zones located at the edge of the maze, while rats 1 and 2 started from zones located in the center of the maze and finished at zones located on the periphery of the maze; i.e., the trajectories made by rats 1 and 2 were shorter on average. Rats met the behavioral criterion for this task when they reached at least 70% of the targets within 7 seconds.

Once rats had acquired the task, their qualitative behavior became similar, or stereotyped, in several respects. Following each stimulation reward, they paused for approximately one second, turned toward the center of the maze, and ran inward, looking to the left and right. During this time, they engaged in stereotyped searching behaviors, including whisker flicking and sniffing. At trajectory onset, rats first oriented themselves to the target before accelerating towards it (Figure 4-8).

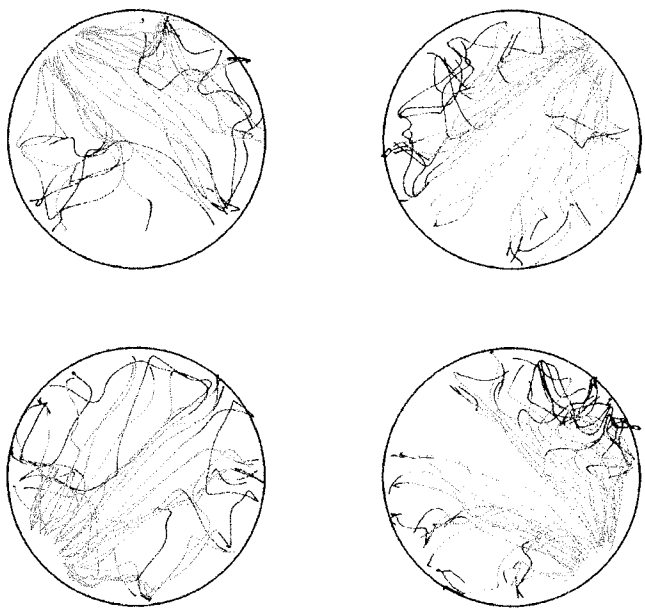

Figure 4-8. Examples of trajectories for a single session in the Random task grouped according to common target zones. Examples were taken from four target zones, and each figure shows all of the trajectories made to one of those zones. Red lines depict the location of the rat from cue onset to trajectory onset. Green lines depict the location of the rat from trajectory onset to the entry of the rat into the reward zone and subsequent delivery of reward stimulation.

The high level of motivation to locate the cued target was characterized by one behavior that occurred when the cued target was located directly behind the rat. Targets located behind the rat were outside the field of vision of the rat prior to the onset of

trajectory. In such cases, it was common for rats to look to the left and right, turn abruptly and run in the direction of the cued target. This suggested that rats had learned that a cued zone was present somewhere on the maze following a tone, and that if that zone was not visible then it must be behind them. Periods of intense searching were interrupted occasionally, however, by grooming bouts that lasted normally for less than a few seconds. Grooming bouts occurred two or three times during each recording session.

Place-related activity

During the six recording sessions, a total of 162 CA1 pyramidal cells were recorded: 14 and 5 from two sessions in rat 1, 40 and 34 from rat 2, 37 from rat 3 and 32 from rat 4. Fewer hippocampal cells were recorded from rat 1 because several tetrodes were used to record parietal units during those sessions. Because the information that was required to compute a trajectory was not present until the cue light was illuminated, any hippocampal activity associated with the planning of a trajectory to a specific target must have occurred after cue onset. If the motor plans are completed prior to or just after the onset of a trajectory, activity related to trajectory planning should be completed by the time the rat has reached its peak acceleration during that trajectory. These two time points provided a bounding window in which to search for activity related to trajectory planning. For each hippocampal pyramidal cell recorded during a session, the total number of spikes that occurred between cue onset and peak acceleration was computed, constituting the spikes that occurred during the "Early" stage of the trajectory. Spikes occurring between the time of peak velocity and arrival at the goal were computed, constituting the spikes that occurred during the "Late" stage of the trajectory. As

described in Figure 4-5, these stage-related spike counts were grouped according to four conditions. Two of these conditions provided test cases for this method of parsing trajectories. The first test condition grouped trajectories by start location and counted the number of spikes that occurred early in the trajectory. At the start of trajectories from a common location, cells with a place field near that location should fire robustly. The number of spikes per trial fired by that cell at that location should differ significantly from the number of spikes occurring at other starting locations. Of 162 hippocampal cells, 18 units fired significantly more spikes per trial early in trajectories coming from a particular zone (chi-squared test, $p<.05$), which is more than the number expected by chance (binomial test, $p=.05$). Likewise, at the end of trajectories towards a common goal, cells with place fields near that goal should fire more robustly than during trajectories to other goals. Of 162 hippocampal cells, 31 units fired significantly more spikes per trial late in trajectories to a particular zone (chi-squared test, $p<.05$). This number, too, is more than would be expected by chance (binomial test, $p=.05$). Across both conditions, 42 units exhibited at least one place field by Muller, Kubie and Ranck criteria.

Predictive activity

If distant goals activate hippocampal units during trajectory planning, then those units should be active at some time between cue onset and the moment of peak acceleration; i.e., "early" in the trajectory. Planning-related activity should depend solely on the goal, not on the current location. Mean firing rates during the "planning" stage were grouped according to the goal of each particular trajectory (see Figure 2-4, "plan").

Of 162 hippocampal pyramidal cells, only 2 cells (both recorded in rat 2 during a single session) showed significant goal-related activity during the "planning" stage (chi-square, $p<.05$). Of the 42 units that had place fields in the arena, two units showing predictive activity is less than the number of units expected by chance (binomial test, $p=.05$). Two other units (from a different session in rat 2) showed significant goal-related activity during the planning stage, but did not have place fields in the arena. In several instances, a given cell that exhibited two place fields within the arena often produced activity that could be interpreted as being predictive. When all trajectories through the place field are observed, however, this activity was explained more simply by the traversal of two place fields within a single trajectory (Figure 4-9).

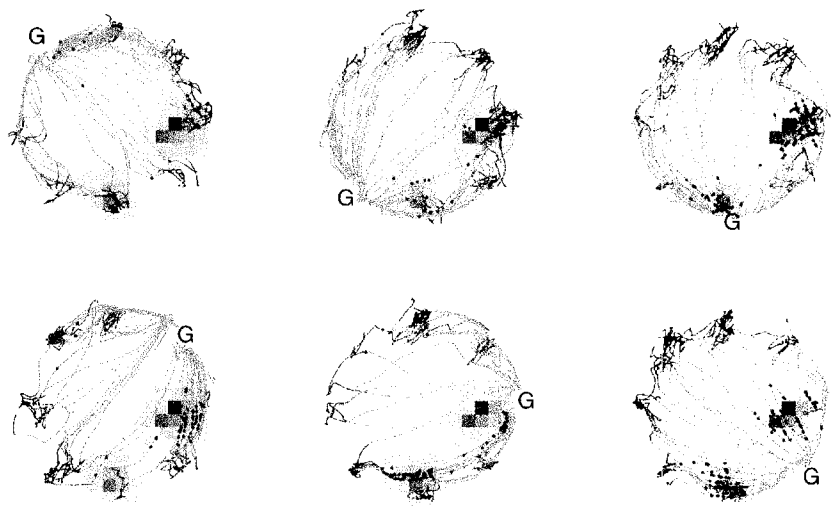

Figure 4-9. One unit with two place fields produce apparent anticipatory activity, but this can also be explained by place-related activity. When one field is located at the start of the trajectory, the unit responds at the start of the trajectory, but also responds whenever that field is crossed, regardless of the portion of the trajectory: start, middle or end.

Correlation of "planning" activity to place field distance

Future place field activity might be anticipated by place cells, even if only briefly, during the "planning" stage of trajectories. If a given place cell does anticipate future activation, then trajectories that will pass close to the center of that cell's place field (causing strong activation of the place cell) should be more likely to fire during the "planning" stage than trajectories far from the place field. For each place unit in a given recording session, spikes that occurred during the "planning" stage were counted for each trial, and the nearest approach to the center of the place field was computed. Of 162 hippocampal pyramidal cells recorded across all sessions, 42 exhibited at least one place field. Of 1077 trials during which a unit spiked three or more times while in the place field of that cell, 8432 spikes were associated with place field-related activity (an average of 7.83 spikes per trial). During these same trials, a total of 279 spikes occurred prior to entry into place fields (an average of 0.26 spikes per trial). There was no significant correlation between the nearest approach to a place field, and the number of spikes during trials prior to entry into the place field (r^2 value = .0023).

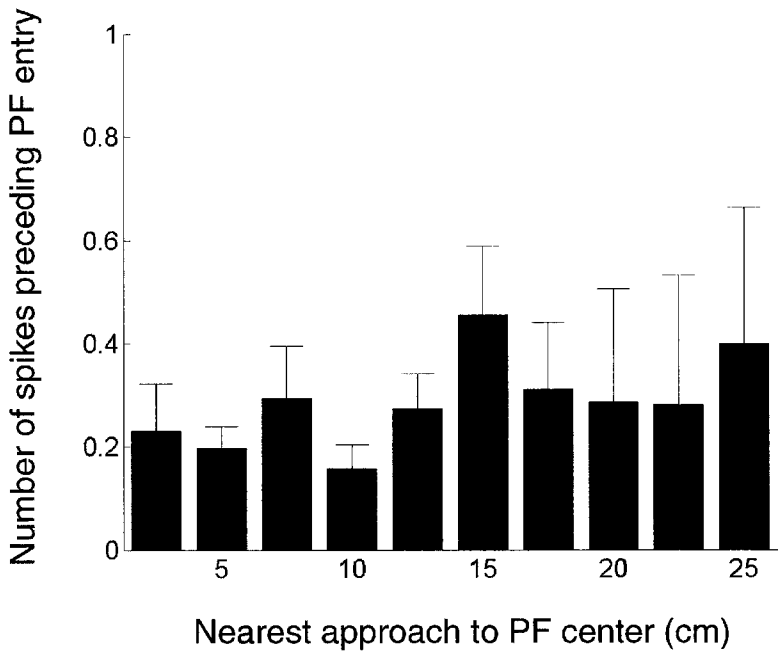

Figure 4-10. The proximity of trajectories to the center of place fields does not predict the number of spikes occurring prior to place field entry. The horizontal axis shows the distance of the nearest approach to the center of a place field. For a trial to be included in the analysis, at least three spikes must have occurred within the limits of the place field. The vertical axis shows the number of spikes that occurred prior to entry into the place field. Each data point was dithered by ± 0.5 spikes to allow visualization of overlapping points.

Sequential task

Learning

Two rats (rats 2 and 3) were trained to run Sequence 1, and three rats (rats 2, 3 and 4) were trained to run Sequence 2. Each rat had been trained previously to run to cued targets in the Random task. Trials in both Sequence tasks were grouped into blocks depending on whether the goal zone was marked by a light in conjunction with the audio tone marking the start of the trial (the "cued" condition) or whether the onset of the light was delayed by five seconds (the "non-cued" condition). Three cued blocks alternated with three non-cued blocks. A typical recording session contained 10-15 non-cued

blocks. Considering non-cued blocks of trials following the introduction of a new sequence (Figure 4-11, "New Sequence"), a significant difference in performance was observed in the first block following the change in sequence (t-test, $p>.05$), but no significant difference in performance was observed between the ten blocks preceding the change, and blocks 11-20 following the change (ANOVA, $p<.05$). That is, changing the sequence disrupted performance, but rats typically acquired new sequences within two recording sessions. To determine whether rats learned sequences using a motor strategy (i.e., as a series of left and right turns) or using a spatial strategy (i.e., running to targets based on their relationship to the constellation of extra-maze cues), rats were challenged with a rotated version of a familiar sequence (Figure 4-11, "Rotated Sequence"). The location of the reward zones in the sequence was rotated by 180° with respect to the extra-maze cues. The spatial relation between reward zones within the sequence was not changed. If rats learned sequences using a motor strategy, then the rotated version of the sequence could be performed immediately if the rats ignored distal, extra-maze cues. If rats did not ignore distal, extra-maze cues, however, the rotated sequence should appear to the rat as a new sequence. As in the case of a new sequence, performance was disrupted in the first block following the rotation (t-test, $p>.05$), but no significant difference in performance was observed between the ten blocks preceding the change, and blocks 11-20 following the rotation (ANOVA, $p<.05$). No significant difference in performance was observed for blocks 11-20 following a change to a new sequence and following a rotation of a known sequence (ANOVA, $p<.05$). That is, there was no difference between rats learning to run a new sequence and rats learning to run a rotated

version of a known sequence, suggesting that rats were using distal, extra-maze cues to learn the sequences.

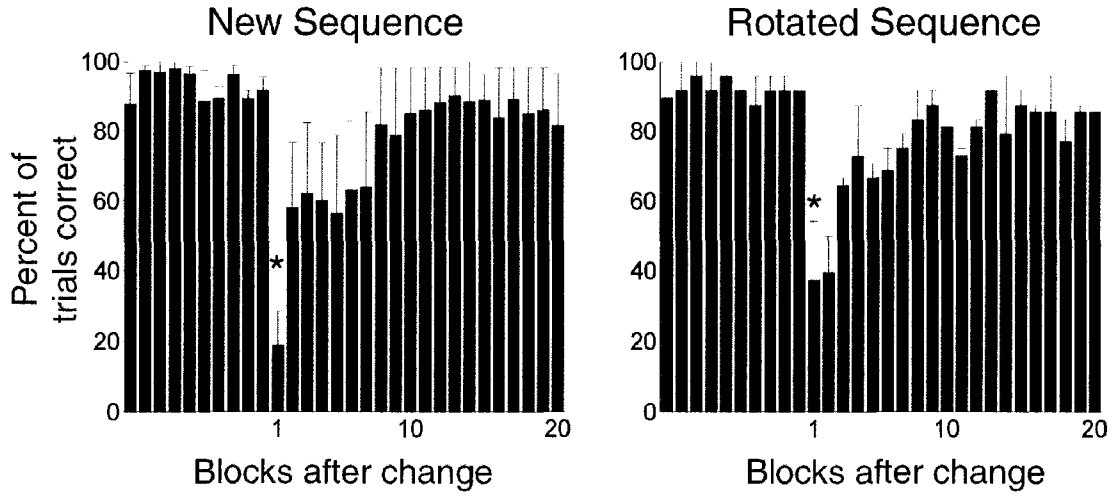

Figure 4-11. Changing to a new sequence or rotating the same sequence caused a disruption in performance, but the change was learned quickly. Two rats performing above a criterion (90% of all trials during a sequence) were exposed to two alternate sequences, each, an entirely new sequence and an 180° rotation of the original sequence. In both graphs, the horizontal axis shows the number of blocks (each containing 24 non-cued trials conducted over 2-3 days) following exposure to the alternate sequence. The vertical axis shows the percentage of trials that met the accuracy criteria. (New Sequence) Two rats were exposed to two new sequences, each (n=4). Each new sequence contained a new spatial arrangement of reward zones when compared to the original sequence. (Rotated Sequence) Two rats were exposed to one rotated sequence, each (n=2). The original sequence was rotated by 180° with respect to the recording room. All distal cues remained in their original locations. (*) Performance during the first block after the change in both conditions was significantly worse than the average performance on the preceding ten blocks (t-test, $p<.05$). There was no significant difference in performance levels between the conditions starting 10 blocks after the change in the sequences (ANOVA, $p>.05$).

Place-related activity

409 units hippocampal pyramidal units were recorded across 12 sessions in three rats performing one of two sequential tasks. Across all sessions, 3175 runs of 5167 (61.5%) total across all sessions were considered valid, and used for analysis. Unlike the first case of randomly presented goals, anticipatory spikes during a sequence should

accumulate at a fixed location. If this were the case, then one would expect to observe either longer place fields (if anticipatory spikes occurred in close proximity to the place field) or more place fields (if anticipatory activity were separated from the place field) per place cell in the case of sequential tasks. Place fields were computed for trajectories during the sequence task (Muller et al., 1987). For each unit exhibiting a place field on a given trial, the number of spikes occurring during that trial prior to entry into the place field was computed (Figure 4-12. When place field properties were compared across tasks, however, the only significant difference that was observed was that place fields during the sequence tasks had a smaller area (Figure 4-12.C). This effect can be attributed to restricted sampling of the place field, because the same trajectory was taken through the place field each time.

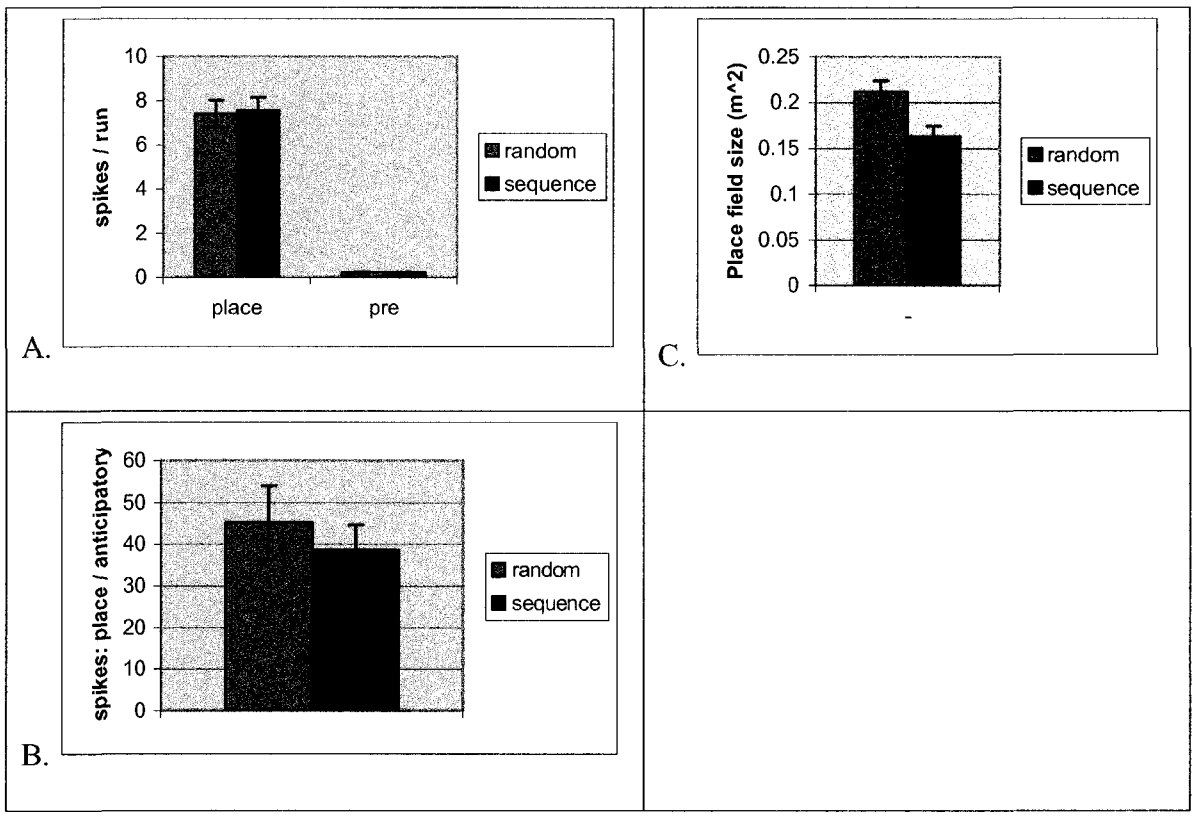

Figure 4-12. Comparison of place field properties during Random and Sequence tasks. The maze was divided into a grid of 16x16 bins, and both spike and position data were separated into these bins. The average firing rate within a bin equaled the number of spikes fired while in that bin divided by the amount of time spent in that bin. A place field consisted of at least six, spatially-connected bins with a firing rate at least one standard deviation above the average for all bins (Muller et al., 1987). A. (LEFT) Average number of spikes occurring within place fields. (RIGHT) Average number of spikes occurring prior to entry into a place field. B. Ratio of the number of spikes occurring with a place field to the number of spikes prior to entry into the place field. C. Average place field area, according to the criteria of Muller et al. (1987). The smaller area observed in the sequential tasks was due to the more restricted sampling during trajectories.

Comparison of Random vs. Sequence trials to a common target

The number of spikes occurring during the "planning" stage of trajectories was compared for the Random and Sequence 2 tasks. One session from both tasks was used from rats 2, 3 and 4. A threshold of one spike/trial during the "planning" stage of trajectories was set as the minimum requirement for reliable plan-related activity. For the

Random task, 32 of 102 units met this criterion (2 of 37 for rat 2, 12 of 32 for rat 3 and 18 of 23 for rat 4). Because no significant predictive activity was observed in the Random task, the criterion of one spike/trial is biased to produce too many "false positive" errors; i.e., the criterion is too easily met. For the Sequence task, 3 of 78 units met this criterion (1 of 42 for rat 2, 0 of 29 for rat 3 and 2 of 7 for rat 4). The number of units that reached criterion in Sequence 2 was significantly less than the number of units that reached criterion in the Random task (ANOVA, $p<.05$), suggesting that even a criterion that is biased to favor finding "planning" related activity failed to find such activity in rats running Sequence 2. Clear examples of units with place fields near goals on only one of two segments that converged on the same goal were observed (Figure 4-13).

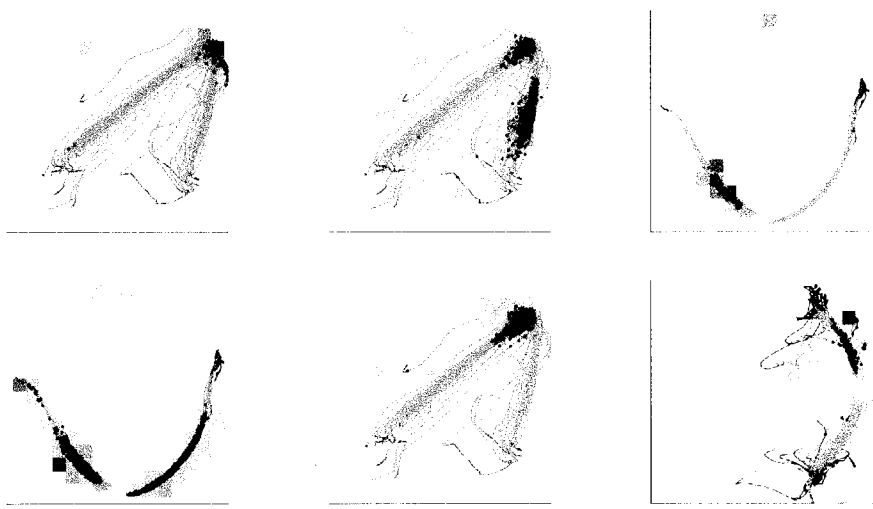

Figure 4-13. No predictive activity for place cells when goals are remembered. Greyscale boxes show the firing ratemap, blue dots are the location of the rat when a spike occurred, and green lines are trajectories. Here, trajectories and spikes are depicted from the stimulation at the previous reward to the current reward site, because rats learned this task very well and anticipated subsequent reward zones, arriving quickly after (and sometimes even before) the tone onset. For this reason, no distinction is made between red and green lined trajectories.

Discussion

Rats were trained to run to distant targets that were presented either in random order or as elements of a sequence. Place cell activity was observed in both conditions from CA1 of the hippocampus without any difference between conditions in place field size, or peak firing rate within the place field. No activity was observed that could be considered predictive of subsequent place cell firing when the rat entered a given place field. This was clearly evident in the case of randomly presented targets, because predictive activity could have been clearly distinguished from independent, place-related activity. For sequentially presented targets, no tendency for place cells in sequential tasks to exhibit more place fields than place cells in random tasks was observed, suggesting that no such masquerade occurred. Rather, hippocampal activity correlates predominantly with the current location of the rat, and is not predictive of remote places to which the rat is about to travel. This agrees with the failure to observe anticipatory activity in the hippocampus of rats running to and from a single, fixed location (Fenton et al., 2002).

An alternate possibility is that goal locations activate different patterns of hippocampal activity than that activated by the current location. If that were the case, then anticipatory activity would not correlate with activity at the current location in the random task, but should correlate with current activity in the sequential task, as the pattern of activity associated with the subsequent goal would be activated in roughly the

same physical location each time through the cycle. This possibility is not supported by the current data, however, due to the lack of either more place fields or more multi-field place cells in the sequential tasks when compared to the random task. Any theory postulating that hippocampal cells in the rat activate goal representations prior to arrival at the goal (i.e., at distances greater than a standard place field size) must be limited to a pattern of activity containing significantly fewer spikes than those activated by the current location. The upper bound on goal-related activity during the "planning" stage becomes even smaller when one considers that no difference was observed between the number of units that spiked an average of once per "planning" stage on the Random and Sequence tasks. The increased computational load involved in the recall of a series of locations did not activate more hippocampal than did running to cued locations.

These results suggest that the rodent hippocampus represents only the "here and now", and not goals or intentions of future movements. While it is known that intentions are reflected in the parietal cortex of primates (Assad and Maunsell, 1995), it is not known whether similar information is encoded in rodent neocortex. If such information is represented in neocortical structures in rats, this information is not reflected in hippocampal activity. While the presence of activity correlated with current location could be involved in trajectory computations in other structures, that computation does not appear to occur in the hippocampus.

Chapter 5 UNIQUE HIPPOCAMPAL CODES ARE NOT NECESSARY TO SEPARATE REPEATED ELEMENTS IN A SEQUENCE

Abstract

Learning sequences containing repeated elements requires differential neural activity to encode the sequential contexts of these elements. Where in the brain such distinctions are made is unknown. Spatial segments that repeat within a sequence of locations might be represented by different neural codes in the hippocampus, which contains spatially selective units ("place cells") and which is thought to separate, or "orthogonalize", similar inputs. In some alternation tasks, hippocampal place cell activity along repeated segments depends on whether the subsequent turn will be to the left or to the right (Wood et al., 2000), suggesting that hippocampal activity can provide sequential context information in some cases (Frank et al., 2000). In other instances of the same task, however, hippocampal activity on the repeated segment is not affected by sequential context (Lenck-Santini et al., 2001;Bower et al., 2000). To explore possible causes underlying these differences, rats were trained to remember and traverse sequences of locations containing repeated segments, under three different conditions: 1) reward was given at each target location; 2) moveable barriers were placed at the entry and exit of the repeated segment to direct the rat during training, and were removed once the sequence was learned; 3) no moveable barriers were used (as in condition 1), but reward was withheld at the entry or exit of the repeated segment. In the first condition, hippocampal

activity was identical along the repeated segment, indicating that complex alternation tasks can be learned in the absence of differential place cell codes in the hippocampus. Differential hippocampal activity was observed, however, in the latter two cases, suggesting that either long-term memory for distinct cues, the working memory of previously visited locations or the recollection of different goals can serve to separate hippocampal activity patterns at the same location.

Introduction

Experience constantly presents the brain with the problem of distinguishing similar events from novel ones. Combining multiple, similar experiences allows stable categories to be learned, while separating similar, novel experiences allows the appreciation of fine distinctions. Magnifying the differences between similar inputs, or "orthoganalizing" the inputs, increases the storage capacity of a network by making erroneous recall less likely (Marr, 1971;McNaughton and Morris, 1987;Treves and Rolls, 1991). Neural activity in the rodent hippocampus displays both characteristics: the capability of expressing a great deal of stability throughout a recording session and across many days (Thompson and Best, 1990), while also being capable of generating different neural codes for the same physical locations in different environmental and behavioral contexts (Bostock et al., 1991); (Skaggs and McNaughton, 1998;Knierim et al., 1995;Markus et al., 1995). What causes the hippocampus to complete or separate patterns, however, remains unclear. One example of this problem is the representation of repeated elements within a larger sequence (Figure 5-1).

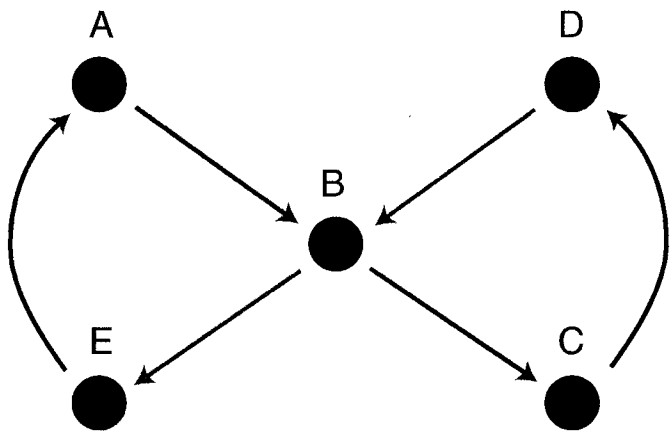

Figure 5-1. Repeated elements in a sequence. In the sequence ABCDBE..., when the current element is "B", correct prediction of the next element ("C" or "E") requires working memory for the previous element ("A" or "D") to modify the activity pattern at B. Correct sequential association requires that the neural activity patterns for B be different in the two contexts.

While little is known about how the hippocampus could separate repeated elements, several models of the hippocampus have proposed a role for the hippocampus in encoding sequences of places or events (McNaughton and Morris, 1987;Levy, 1996;Lisman, 1999;Wallenstein et al., 1998). Support for a hippocampal role in sequence encoding includes the observation that hippocampal lesions impair the ability of rats to recognize different sequences of odors (Agster et al., 2002;Fortin et al., 2002;Kesner et al., 2002). In the case of sequences containing repeated spatial segments, Wood et al. (2000) found differential activity in the hippocampus on the common arm of a T-maze (modified to allow return to the base of the T) prior to the choice point, even though all external variables were held constant once the task had been learned, suggesting that the distinction came from cues internal to the animal; i.e. working memory. A similar result was obtained by Frank et al. (Frank et al., 2000), who found differential activity at the choice point of a W-shaped maze, both prior to and following

the choice of turn direction. Two other reports, however, showed a lack of differential hippocampal activity during forced alternation (Bower et al., 2001;Lenck-Santini et al., 2002), suggesting that multiple factors influence the separation, or "orthogonalization", of hippocampal codes for a given location.

For the purposes of this chapter, we define the term "map" to mean a stable distribution of place field activity within a given environment (Gothard et al., 1996a). How are different maps separated for subsequent recall of one map and not another, and what necessitates the recall of a different map? Switching between maps is different than jumping to a new activity pattern within the same reference frame, as occurs in retina (Mays and Sparks, 1980), parietal cortex (Andersen et al., 1985), the head direction system (Skaggs et al., 1995), and within hippocampal maps (Gothard et al., 1996a). Working memory for the starting location of a current context has been shown to be capable of separating two similar contexts when local view information was identical (Skaggs and McNaughton, 1998). In rats trained to shuttle between two visually identical rooms that were connected by a hallway, the place fields of some cells differed in the two rooms. Given the spatially selectivity of hippocampal place cells, if a primary function of the hippocampus is to encode sequences (Lisman, 1999;Fortin et al., 2002), then sequences with repeated spatial segments, and hence overlapping contexts, might be expected to generate separate hippocampal patterns for each context.

Methods

Order of sequence training

Each rat had been trained previously to make trajectories to reward locations, presented in random order, that were cued visually by the illumination of a light-emitting diode (LED), and cued non-spatially by a 500 msec duration, 4 kHz tone. This task is described in this chapter as the "Random" task. "Sequence" tasks were presented to each rat by cueing reward zones in cyclic order, shown below the schematic of each sequence (Figure 5-2). After a rat completed a sequence three times (a "cued" block of sequences), a five second delay was inserted between the non-spatial, audio cue and the illumination of cue lights, giving the rat a chance to recall the next reward location and make a trajectory to it without a spatial, visual cue. After the rat completed the sequence three times with a delayed visual cue (a "non-cued" block of sequences), audio and visual cues were presented simultaneously, again, starting another block of "cued" sequences. Blocks of three, complete sequences alternated between being cued and non-cued throughout the duration of the recording session. Rats were considered to have learned a sequence when they reached a behavioral criterion of 70% correct turn direction at the choice point during non-cued trials.

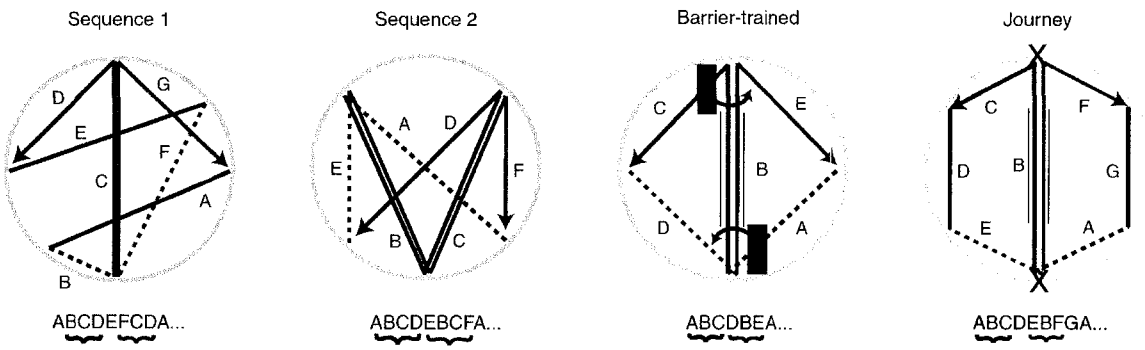

Figure 5-2. The top-down view of the arena showing the schematic for each "Sequence" task. Dotted lines represent the first segment of the two contexts. Counter-clockwise ("CCW") contexts contain a nominal CCW turn (blue) at the choice point, while clockwise ("CW") contexts contain a nominal CW turn (red). The repeated segment(s) are denoted by parallel red and blue lines. Arrows denote the trials that began at the "choice point", where the rat was required to make either a nominal CCW turn (blue arrow) or nominal CW turn (red arrow). Black lines show trials connecting two contexts that neither start nor end on part of a repeated segment. These trials can be considered to be part of either context. (Sequence 1) Segment C is the repeated segment. The rat received reward at the end of each segment. (Sequence 2) Segments B and C are the repeated segments, forming a "V" in the middle of the arena. The rat received reward at the end of each segment. (Barrier-trained) Segment B is the repeated segment. The rat received reward at the end of each segment. Thin, black lines along the repeated segment denote guide rails that restricted lateral movement. The black rectangles denote two, wooden barriers that barred entry into some segments during training. With the blocks located as shown in the figure, and the rat running on A, the barrier forced the rat to run along B. The next barrier then forced the rat to run along C (i.e., the clockwise context). At this time, both barriers were moved as indicated by the black arrows, forcing the rat to run thought D, B and E (the counter-clockwise context). Over the course of several training sessions, these blocks were moved further away from the entry and exit of B until they were removed from the arena, entirely. (Journey) Segment B is the repeated segment. Thin, black lines along the repeated segment denote guide rails. Unlike the other sequence tasks, electrical brain stimulation reward was given at the ends of B on a probabilistic basis. This probability was reduced across several sessions until the rat received no stimulation at either end of B.

Rats were trained to traverse one or more Sequence tasks from a set of four (Figure 5-2). Rats learned each sequence by running to illuminated LED's during the cued blocks of trials; the non-cued blocks of trials provided probe trials to evaluate the learning progress of the rat. "Sequence 1" consisted of eight segments, two of which followed the same physical path, but were followed by turns in opposite directions, nominally, either a clockwise (CW) or counter-clockwise (CCW) turn. For descriptive purposes, the series of segments leading into, through, and out of the repeated segment

are collectively referred to as "contexts". Contexts containing a nominal CW turn at the end of the repeated segment are referred to as "CW contexts", and those ending in a CCW turn are referred to as "CCW contexts". MFB stimulation was given at each reward location. "Sequence 2" also consisted of eight segments, but contained two, contiguous, repeated segments that followed the same physical path. MFB stimulation was given at each reward location. "Barrier-trained" rats learned to traverse a sequence of six segments containing one repeated segments, and were guided, during training, by two, moveable, wooden barriers. The barriers (wood blocks, 10x15x3 cm) were placed on the maze during the first 5-7 training sessions, and forced the rat to follow the correct path of one context. After the rat exited the repeated segment, the barriers were moved manually so as to force the rat to follow the correct path of the other context. The location of the barriers alternated with each context throughout the duration of a training session. Initially, the barriers were placed in close proximity to the path of the correct sequence, but were moved further from this path over subsequent training sessions, until the barriers were removed from the maze, entirely. A pair of guide rails (2 cm high, .7 m long, 10 cm separation) was added to the maze to limit the lateral movement of the rat along the repeated segment. The "Journey" sequence consisted of eight segments with one repeated segment. As in Sequence 1 and 2, no barriers were used during training for this task. Reward was given at both the beginning and end of the repeated segment on a probabilistic basis. The probability of reward was reduced across 5-7 sessions until no reward was given at either end of the repeated segment. A pair of guide rails (2 cm high,

.7 m long, 10 cm separation) was added to the maze to limit the lateral movement of the rat along the repeated segment.

Computation of binned mean firing rates

In analyses looking for differences in the firing rate of single units and populations of units on repeated segments, data from non-cued trials only were used. Sequential trials along the repeated segment(s) and following the choice point were grouped together into "contexts". Contexts were separated according to the nominal turn direction following the choice point. Data from contexts during which rats made a mistake, and chose the wrong reward zone, or did not reach the correct reward zone before the cue light was illuminated were dropped from further analysis.

Across all remaining non-cued trials, the average path taken by the rat through the repeated segment was divided into five spatial bins of equal length, similar to the procedure used by Wood et al. (2000). Position and spike data were assigned to one of these bins. For each trial, the binned firing rate of each unit was computed by dividing the number of spikes that occurred while the rat was in that bin by the cumulative amount of time the rat spent in that bin. Because the body orientation of the rat varied with turn direction at the beginning and end of the repeated segment, data in the first and last bins were dropped from further analysis. For each unit, the mean binned firing rate along the repeated segment was computed, and used to test for significant activity using an ANOVA ($p<.05$). In addition, binned firing rates were divided into two groups: one group for the CW context, and the other group for the CCW context. The mean binned firing rate was computed separately for the two groups, and tested using an ANOVA

($p<.05$). If the mean binned firing rate of a unit was significantly different in any one of the three bins, then the firing of that unit was considered to differ significantly (i.e., unit activity was "differential") along the repeated segment. The significance of the number of units with differential activity across the population was determined using a binomial test ($p=.05$).

Fisher Linear Discriminant

The degree of separation between the two groups was determined by the construction of a Fisher Linear Discriminant (FLD) (Duda et al., 2001) using a "leave one out" method. The binned firing rates (for the three middle bins) for all units were joined together to form a population vector for each trial; i.e., the population vector for each trial had a length of N units multiplied by three bins. Population vectors across trials were grouped according to CW and CCW half-sequences. A classifier was constructed using all of the population vectors from the two groups, with the exception of one vector from one group, which became the test vector for that particular classifier (Figure 5-3).

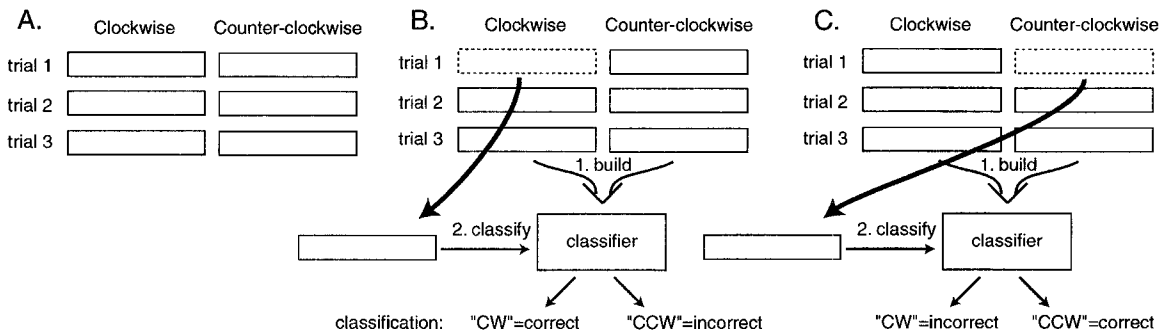

Figure 5-3. Classification using a "leave one out" method. A) The population vectors for each trial across all units and the middle bins are grouped according to context: "clockwise" and "counter-clockwise". B) One population vector (for the first, clockwise trial) is held aside as the test vector, while the remaining population vectors are used to build a classifier. The classifier finds the plane, or "discriminant", in N x 3 dimensional space that produces the best separation between the vectors of the two groups. After the discriminant is computed, the test vector that was originally held aside is classified as belonging to either the clockwise ("CW") or counter-clockwise ("CCW") group. Because the test vector originally belonged to the CW group, that classification is counted as a "correct" classification, while a CCW classification is counted as an "incorrect" classification. C) The previous test vector is returned to the original set of vectors, and the process is repeated for a vector from the counter-clockwise group. Because this test vector came from the counter-clockwise group, a CW classification is incorrect, and a CCW classification is correct. This process is repeated for each vector in both groups. The total percentage of correct classifications, or "separability score", ranges from 0 to 100%. If the two groups contain the same number of vectors, then the chance level is 50%.

The classifier was then used to decide which group gave the better match to the test vector; i.e., the test vector was "classified" as belonging to one group or the other. This process was repeated for each segment in both groups, where only the test vector was "left out" during the construction of the classifier. If the better match for the test vector was the group to which it belonged, this was considered a "correct" classification. Matching the other group was considered an "incorrect" classification. The percentage of correct classifications, or the "separability score", across both groups served to quantify the degree of separation between the two groups of population vectors. The best possible separability score for a single recording session was 100%. Separability scores were grouped by Sequence task, and a mean and standard error computed. The level of chance for each Sequence task was computed using the same classifier construction technique,

except that membership in the two groups was assigned randomly. Population vectors for a given Sequence task were considered to be separable if the original and chance separability scores were significantly different, as determined by an ANOVA (p=.05).

Results

Data were collected from three rats trained to run some or all of the sequences described in this chapter. All rats were trained first on the Random task, described in the previous chapter. Rat 1 was trained to run Sequence 1 and Sequence 2. Rat 2 was trained to run on all four sequences described in Methods, and the sequences were presented in the order listed there. Rat 3 was not trained on Sequence 1, but was trained on the other three sequences in reverse order; i.e., Journey, then Barrier-trained, followed by Sequence 2. Training on the sequences in reversed order was done to counter-balance for effects involving the order in which sequences were presented to the rats. Though data were acquired from sessions during training, the only data sets used for analyses in this chapter were those obtained on the last day of training on a particular sequence. Because the tetrodes were not moved after being positioned in the cell body layer of CA1, a great deal of unit overlap existed between data sets. One exception to this protocol concerns the analysis of data on the first day of training on the "Barrier-trained" and "Journey" tasks.

Sequence 1

Two rats were trained to run Sequence 1 (Figure 5-4). It became evident, however, that two problems existed with this task regarding the question of hippocampal

codes for repeated segments. The first problem was that neither rat learned to choose the correct direction in which to turn at the choice point on more than 80% of all non-cued trials (the behavioral criterion). The second problem was that the physical path taken by each rat during the counter-clockwise and clockwise half-sequences differed in the repeated segment. This problem led Wood et al. (2000) to narrow the alley of the repeated segment, forcing rats to take similar paths during both half-sequences. The open arena used in this experiment allowed rats to take variable trajectories. When units were observed to fire at significantly different rates during the repeated segment, it was not possible to determine whether that difference was due to differential coding or simply to the encoding of different places.

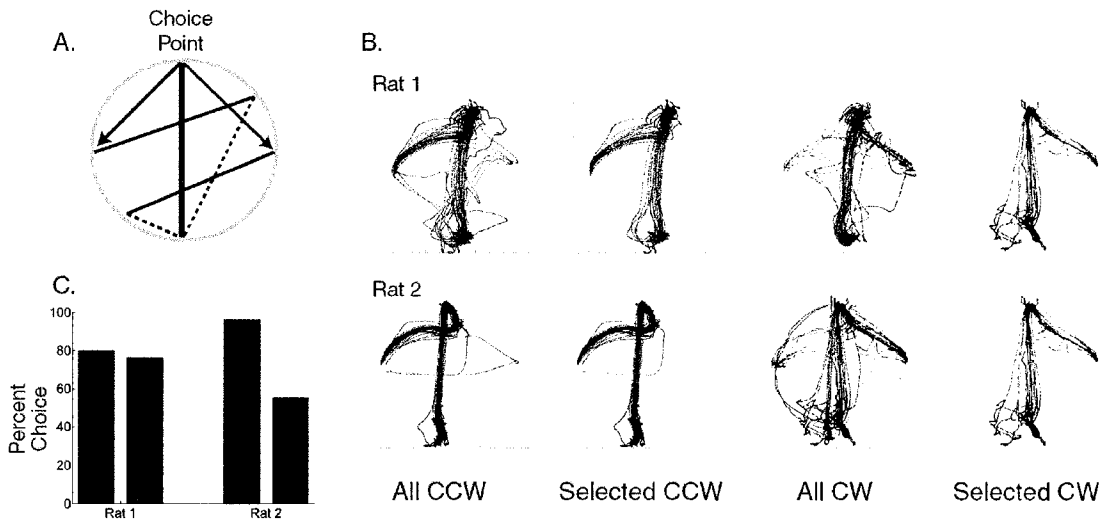

Figure 5-4. Behavioral analysis for the Sequence 1 task. (A) Top down view of the maze showing the schematic of Sequence 1. Arrows depict the different, possible paths following the choice point. Blue lines depict the counter-clockwise (CCW) half-sequence. Red lines depict the clockwise (CW) half-sequence. (B) Actual trajectories for the two rats that were trained on this task. The first column ("All CCW") shows all counter-clockwise trials leading up to and away from the choice point. The second column ("Selected CCW") shows those CCW trials that met the behavioral criterion. The third column ("All CW") shows all trials leading up to and away from the choice point. The fourth column ("Selected CW") shows those CW trials that met the behavioral criterion. (C) The percentage of trials for both rats that met the behavioral criterion.

Sequence 2

Three rats (rats 2, 3 and 4) were trained on Sequence 2. The behavioral criterion for deciding that the rat had learned the task required each rat to reach the correct target in the non-cued condition on 70% of all non-cued trials. To reach this criterion, rat 2 required six training sessions, rat 3 required eleven sessions, and rat 4 required ten sessions. 57, 29 and 13 units were recorded in each data set, respectively. Of these, 49, 29 and 10 units were identified as pyramidal cells, based on parameters as described in the chapter on General Methods. "Sequence 2" contained two repeated segments in succession, which were used to avoid the problem of context-dependent entries into repeated segments, observed with Sequence 1. In the schematic of Sequence 2 (Figure 5-5.A), blue and red paths show the paths following the choice point during counter-clockwise (CCW, blue) and clockwise (CW, red) turns, respectively. Actual paths taken by each rat (Figure 5-5.B) are shown first for all trials within each context, and then only for those trials that met the behavioral criterion and were used for subsequent analysis. The behavioral criterion had two components: the correct zone for the current context had to be entered first, and that zone had to be entered before the illumination of the visual cue (i.e., within 5 seconds after leaving the choice point). The first criterion (entering the correct one of two total zones) would make the level of "chance" performance 50%, but the additional constraint of reaching the correct zone within a time limit makes that level less than 50%. The bars graphs (Figure 5-5.C) show the behavioral performance for each rat. For the sessions used in subsequent analysis, each rat chose the correct zone following the choice point within the time limit on more than 70% of all non-cued trials.

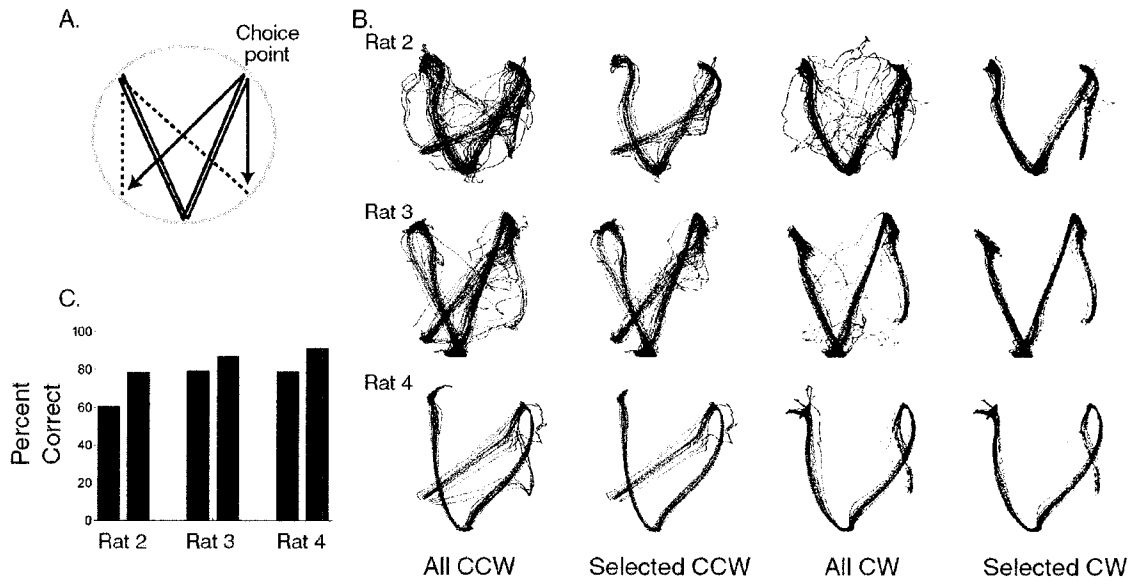

Figure 5-5. Behavioral analysis for the Sequence 2 task. (A) Top down view of the maze showing the schematic of Sequence 2. Arrows depict the different trajectories following the choice point. Blue lines depict the counter-clockwise (CCW) half-sequence. Red lines depict the clockwise (CW) half-sequence. (B) Actual trajectories for the three rats that were trained on this task. The first column ("All CCW") shows all counter-clockwise trials leading up to and away from the choice point. The second column ("Selected CCW") shows those CCW trials that met the behavioral criterion. The third column ("All CW") shows all trials leading up to and away from the choice point. The fourth column ("Selected CW") shows those CW trials that met the behavioral criterion. (C) The percentage of trials for both rats that met the behavioral criterion. Each rat was considered to have learned the task when two behavioral criteria were met on at least 70% of all non-cued trials: the correct zone following the choice point for the current context was chosen, and that zone was reached before the onset of the visual cue. That is, the rat had to recall the correct zone for the context, and reach it within 5 seconds. The presence of two criteria made the performance level of "chance" less than 50%, yet each rat performed at better than 70%.

Considering only the trials leading to the choice point (the second of the repeated segments), 27 of the 88 pyramidal cells across all three rats showed significant bin-

specific activity (ANOVA, p=.05) (9 of 49 for rat 2, 12 of 29 for rat 3, and 6 of 10 for rat 3). Of pyramidal cells with significant activity in the second repeated segment, 1 of 27 cells showed significantly different activity for CW and CCW contexts (ANOVA, p=.05), which was not significant according to a binomial test (p=.05) (Figure 5-6). Activity for the one unit with significantly different activity in the two contexts is shown in the first column for rat 4 in Figure 5-6.

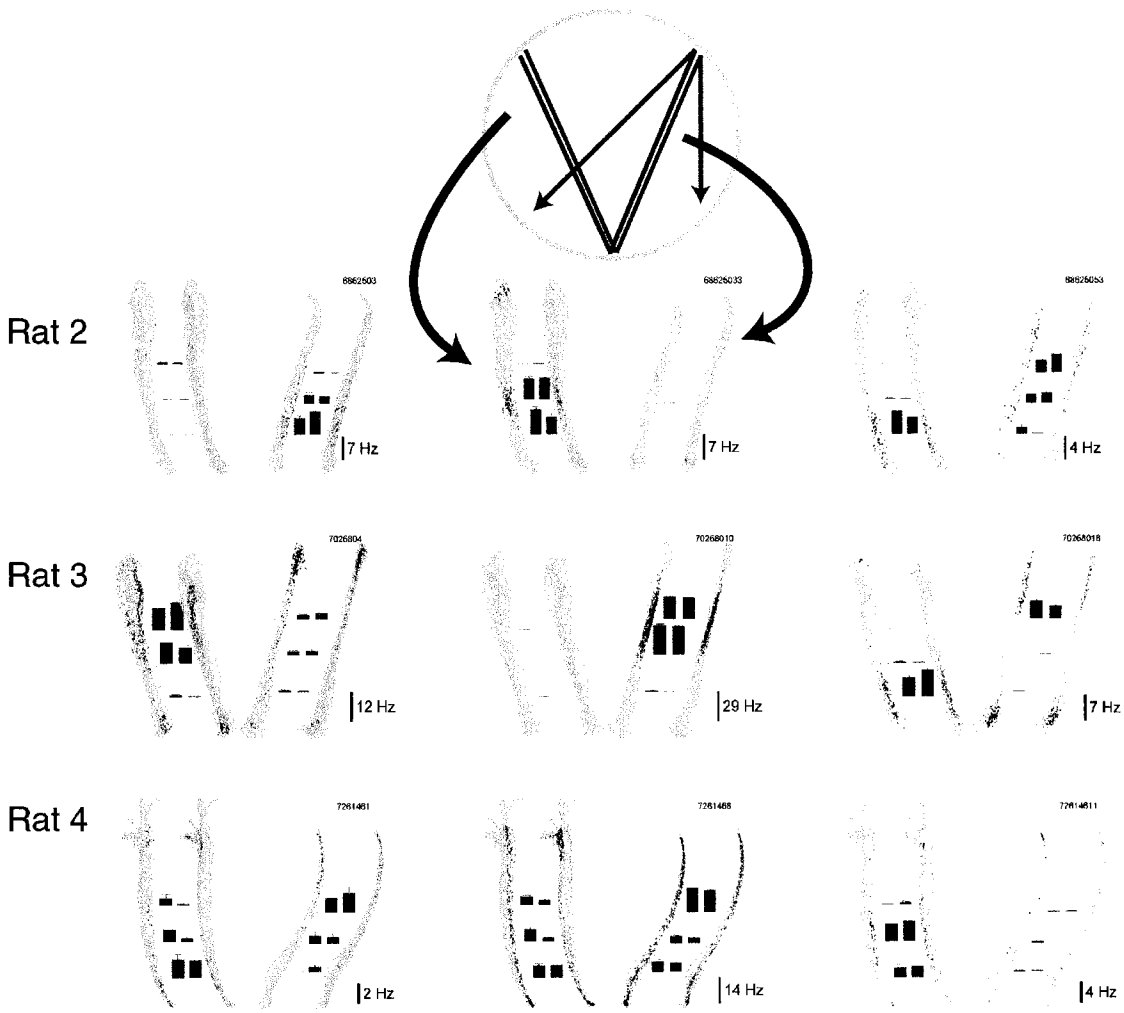

Figure 5-6. A complex alternation task is solved without differential encoding in the hippocampus. Each row depicts three different units during one recording session from the same rat. Each figure

shows the trajectories and binned firing rates from a single pyramidal cell. Within each figure, the actual path taken by the rat (shown in gray) during clockwise (light gray paths; blue dots and bars) and counter-clockwise (dark gray paths; red dots and bars) contexts are paired together and overlaid to show the similarity between the routes that were taken. In each figure, the pair of trajectories to the left shows the routes taken and spikes that occurred during the first of the repeated segments; the pair of trajectories to the right shows the routes taken and spikes that occurred during the second of the repeated segments (the segment leading to the choice point). Within each pair of trajectories, trials during counter-clockwise (CCW) contexts are depicted to the left, and clockwise (CW) context to the right. Spikes are overlaid on the trajectories according to the location of the rat when each spike occurred: (blue) spikes occurred during the counter-clockwise context, (red) spikes occurred during the clockwise context. The mean binned firing rates and standard errors are shown by bar graphs located between paired trajectories (CCW in blue, CW in red). The units displayed were chosen because they showed the largest difference between the two contexts. Scale bars show firing rate for each bar graph, respectively.

One possible explanation for the discrepancy between our result and that of Wood et al. (2000) involved the hemisphere from which hippocampal units were recorded. Several (but not all) units recorded in Wood et al. were recorded from the left hippocampus. For pyramidal units with significant activity in the second of the repeated segments, 0 of 12 units from the left hemisphere (rat 3) showed significantly different activity, while 1 of 15 units from the right hemisphere (rat 2 and 4) did so, which are both fewer than would be expected by chance (binomial test, p=.05).

Barrier-trained

Two rats (rats 3 and 4) were trained on a variant of the sequence tasks that involved the use of barriers during training. The lack of significantly different neural activity during repeated segments in Sequence 1 and Sequence 2 resulted from work ongoing (Bower et al., 2001) at the time of the publication of Wood et al. (2000). To address the discrepancy between our results and that of Wood et al. (2000) and Frank et al. (2000), we attempted to replicate the result of Wood et al. with our experimental setup. One possible explanation for the discrepancy in our results and those obtained by

Wood et al. centered on differences in how rats were trained to alternate turns. Wood et al. cued rats during training by blocking the paths that rats could take until they had learned the task. In addition, no reward was given on the repeated segment, but only on subsequent segments following a correct turn from the repeated segment. To determine whether either of these factors contributed to the differential activity observed by Wood et al., two different training methods were used: First, obstacles were placed at the entry and exit arms to the repeated segment, forcing the rat onto the correct arm, as well as providing cues that changed reliably with the direction of the impending turn. Second, without the use of obstacles, stimulation was withheld at the ends of the common arm, allowing the rat to make smooth, continuous entries into, trajectories through, and exits from the common arm. Results from the first variation, the "Barrier-trained" task, are described in this section.

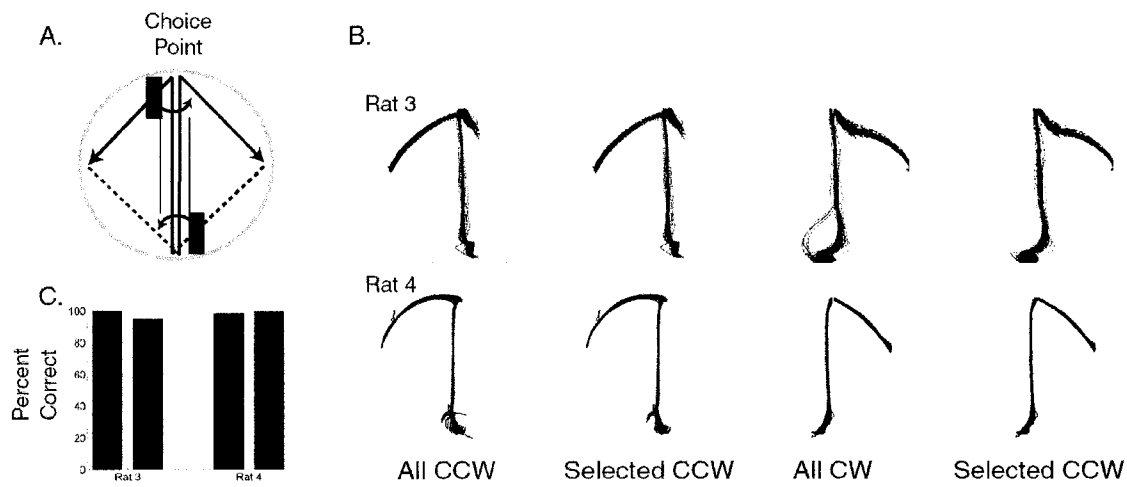

Figure 5-7. Behavioral analysis for the Barrier-trained task. (A) Top down view of the maze showing the schematic Barrier-training. Arrows depict the different trajectories following the choice point. Blue lines depict the counter-clockwise (CCW) half-sequence. Red lines depict the clockwise (CW) half-sequence. Black rectangles denote moveable wooden barriers that forced the rat to follow one of the two half-sequences. The barriers were removed after multiple training sessions. (B) Actual trajectories after the barriers had been removed for the two rats that were trained on this

task. The first column ("All CCW") shows all counter-clockwise trials leading up to and away from the choice point. The second column ("Selected CCW") shows those CCW trials that met the behavioral criteria (chose the correct zone first, and did so before the cue light was illuminated). The third column ("All CW") shows all trials leading up to and away from the choice point. The fourth column ("Selected CW") shows those CW trials that met the behavioral criterion. (C) The percentage of trials for both rats that met the behavioral criterion.

On the first day of training using barriers at the entry and exit of the repeated segment, 19 of 22 units recorded from rat 3 were identified as pyramidal cells, based on wave shape, firing rate, and ISI distribution. 16 of these pyramidal cells showed significant, bin-specific activity along the repeated segment (ANOVA, $p<.05$). Of these, 13 cells showed significantly different activity between the CW and CCW contexts (ANOVA, $p<.05$). This was a significant number of units (binomial test, $p=.05$).

From the first session during which barriers were present, rat 3 required 10 sessions and rat 4 required 9 sessions until the rats were running the task and the barriers were no longer present on the maze. Considering only the trials on the repeated segment, 40 of the 51 pyramidal cells across the two rats showed significant bin-specific activity (ANOVA, $p=.05$) (14 of 16 for rat 3, and 10 of 24 for rat 4). Of pyramidal cells with significant activity in the second repeated segment, 24 of 30 cells showed significantly different activity for CW and CCW contexts (10 of 14 for rat 3, and 12 of 19 for rat 4) (ANOVA, $p=.05$) (Figure 5-8). This was more units than would be expected by chance (binomial test, $p=.05$). As described in Wood et al., differential activity was observed both in the form of units with new place fields, and as significant changes in firing rate in similar locations (i.e., resulting from the same place field).

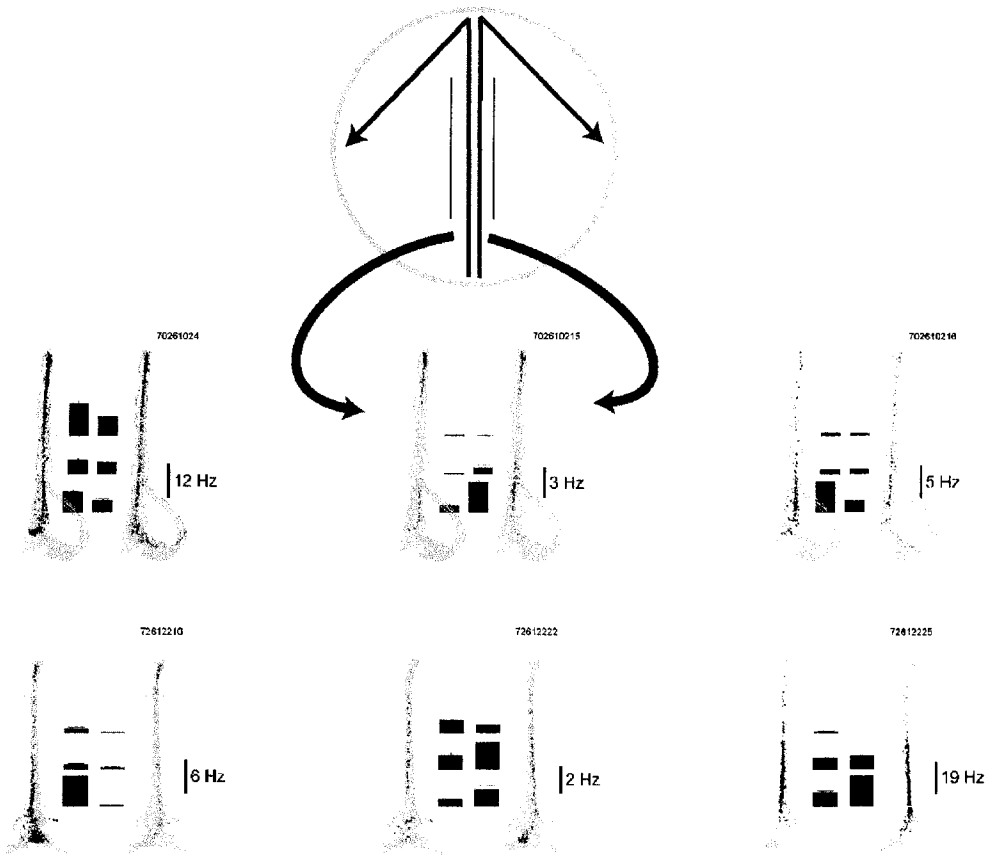

Figure 5-8. Differential activity following training with blocks. Scatter and rate plots for units with significantly different firing rates for trajectories (grey dots) ending in turns in different directions. As in Figure 5-6, blue dots show the location of the rat on the repeat segment when a unit spiked during trajectories prior to a left turn, and red dots show the location of the rat when a cell spiked during trajectories prior to a turn to the right. Bar graphs show binned mean firing rates. Scale bars show the firing rate for each unit.

Journey

When stimulation was withheld at the ends of the common segment, rats learned to traverse the common segment without stopping at the entry or exit (as was the case when those locations were rewarded). Rats were trained to perform this task over several days by decreasing the probability of reward at the entry and exit of the repeated segment

until the probability of reward reached zero. This required six and ten sessions, respectively, for the two rats trained on this task (Figure 5-9).

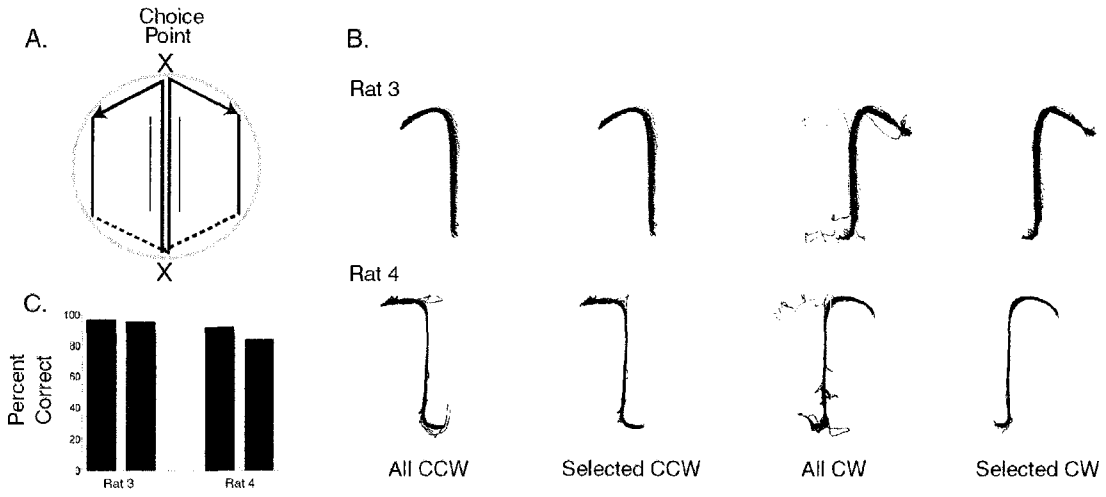

Figure 5-9. Behavioral analysis for the Journey task. (A) Top down view of the maze showing the schematic of the Journey task. Arrows depict the different trajectories following the choice point. Blue lines depict the counter-clockwise (CCW) half-sequence. Red lines depict the clockwise (CW) half-sequence. (B) Actual trajectories for the two rats that were trained on this task. The first column ("All CCW") shows all counter-clockwise trials leading up to and away from the choice point. The second column ("Selected CCW") shows those CCW trials that met the behavior criterion. The third column ("All CW") shows all trials leading up to and away from the choice point. The fourth column ("Selected CW") shows those CW trials that met the behavior criterion. (C) The percentage of trials for both rats that met the behavioral time criterion and were selected.

On the first day of training (when reward was given with a probability of 50% at the entry and exit of the repeated segment), 20 of the 25 units recorded from rat 3 were identified as pyramidal cells, based on wave shape, firing rate, and ISI distribution. Of the the 20 pyramidal cells, 9 cells showed significant, bin-related activity along the repeated segment (ANOVA,.<.05). Of these, 0 cells showed significantly different activity between the CW and CCW contexts (ANOVA, $p<.05$), which was less than the number of units that would be expected by chance (binomial test, $p=.05$).

Considering data sets when no reward stimulation was given at either the entrance and exit of the repeated segment, and considering only the trials on the repeated segment, 35 of the 50 pyramidal cells across the two rats showed significant bin-specific activity (ANOVA, p=.05) (16 of 17 for rat 3, and 19 of 33 for rat 4). Of pyramidal cells with significant activity in the second repeated segment, 24 of 30 cells showed significantly different activity for CW and CCW contexts (13 of 16 for rat 3, and 12 of 19 for rat 4) (ANOVA, p=.05) (Figure 5-8), which was significant according to a binomial test (p=.05). The number of pyramidal cells that exhibited significantly different behavior is more than would be expected by chance (binomial test, p=.05) (Figure 5-10). As described in Wood et al., differential activity was observed both in the form of units with new place fields, and as significant changes in firing rate in similar locations (i.e., resulting from the same place field).

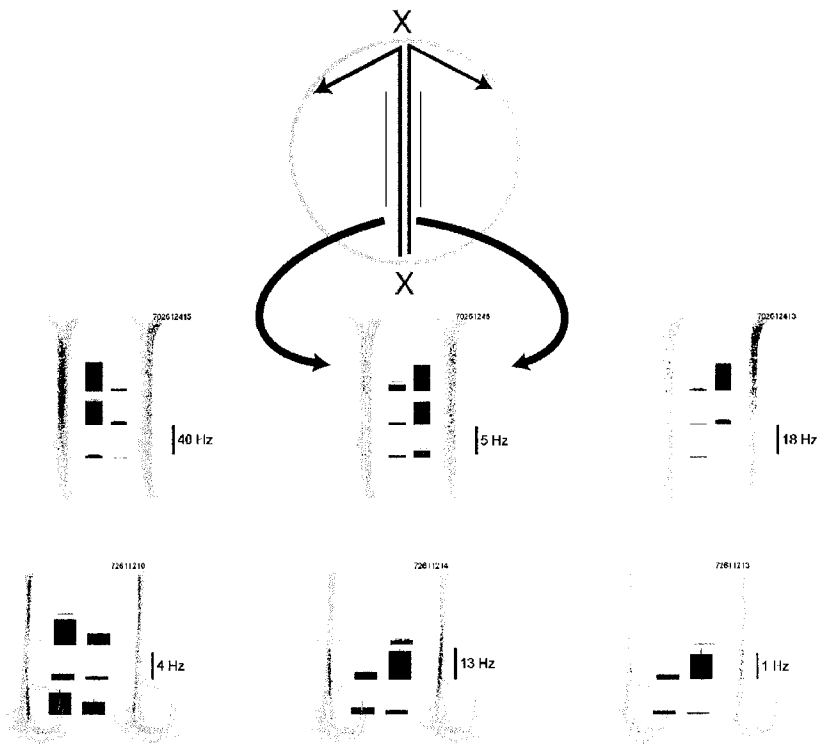

Figure 5-10. Differential activity for continuous trajectories through the repeat segment. No stimulation was given at the entry (bottom of each trajectory) or exit (top) of the repeat segment. As in Figure 5-6, blue dots show the location of the rat on the repeat segment when a unit spiked during trajectories prior to a left turn, and red dots show the location of the rat when a cell spiked during trajectories prior to a turn to the right. Bar graphs show binned mead firing rates. Scale bars show the firing rate for each unit.

Discriminant analysis

Different formulae have been proposed to quantify the difference in activity for ensembles of units during a repeated segment task (Wood et al., 2000;Lenck-Santini et al., 2001). Both measures approximate the optimal Fisher Linear Discriminant (FLD), which finds the best projection for separating the segments of two samples on the two sides of a dividing plane, or "discriminant". Whereas a correlation determines the degree of similarity between two samples, a discriminant determines the degree of "separability" between two samples (Duda et al., 2001). When applied to the population data from

Sequence 2, Barrier-trained, and Journey tasks, the latter two showed significantly larger separability scores than that for Sequence 2 (Figure 5-11). The separability score expected by chance was computed by randomizing class membership within the data for each session, computing a separability score for each session, and then comparing the separability of the original data set against that for the randomized data set (ANOVA, p=.05). Across all data sets, the chance level for separability was not different than 50% (t-test, p<.05).

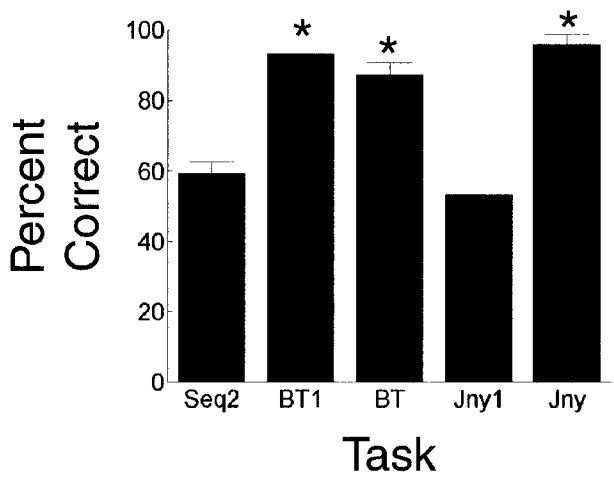

Figure 5-11. Ensemble activity differentiates the direction of subsequent turns for Barrier-Training and Journey tasks, but not for Sequence 2. During trials along the repeated segment, position was divided into five, equally spaced bins. For each unit, a mean firing rate within each bin was computed for each trajectory by dividing the number of spikes occurring in that bin by the total time spent in that bin. Binned firing rates across all units were joined into population vectors. The population vectors were grouped according to half-sequences. A Fisher Linear Discriminant (FLD) was constructed by the "leave one out" method, where the population vector not used in the construction of the discriminator became the test vector. The discriminant determined which of the two groups was the better match for the test vector. This process was repeated for each trial, and a total percentage of correct classifications for the session computed. The vertical axis shows the mean percentage of correct classifications (i.e., the "separability") across rats for each sequence task. "Seq2" describes the overall performance for Sequence2, "BT1" for the first day of Barrier-Training, "BT" for the last day of Barrier-Training, "Jny1" for the first day of Journey training, and "Jny" for the last day of Journey training. Asterisks denote performance that is significantly different from chance, as determined by randomizing group membership prior to construction of the FLD, and comparing using an ANOVA ("*" signifies p<0.05). The lack of an error bar for "Jny1" was due to having only one rat in this condition.

In addition to considering the separability of hippocampal codes following several days of training, the separability observed on the first training day was computed for the Barrier-trained and Journey task ("BT1" and "Jny1" in Figure 5-11). Separability on the first day of the Barrier-trained task was significantly greater than chance, while the separability observed for the first day of training on the Journey task was not. It should be noted that the behavioral scores of both rats 3 and 4 on the first day of training on the Journey task was below the behavioral criterion. Because these sessions were the first in which each rat had experienced probabilistic reward, both rats spent more time than usual searching near zones were reward had been expected and re-tracing paths when reward was not received.

Discussion

Rats learned to alternate left and right turns without cues during a complex spatial sequence task, though no differential neural activity was observed in the hippocampus. That is, alternation did not require a unique hippocampal code for each half-sequence. This agrees with the observation that hippocampal lesions do not impair the ability of rats to learn forced alternation tasks (Dalland, 1976). The task in the latter study, and others like it, however, was simple enough to have been solved in a non-spatial (i.e., response based) manner, which would explain why the rat's performance was not disrupted by hippocampal lesions. In the Sequence 2 task in the present study, a response-based strategy was not used by rats, because performance was disrupted by a rotation of the sequence with respect to distal cues. Although the dependence of Sequence 2 task

performance on an intact hippocampus remains to be verified directly, it is clear that memories for the spatial location of sequence elements were required, and that this spatial sequence memory does not require orthogonal codes in the hippocampus to separate repeated elements in the two contexts.

Differential activity in the repeated segment was observed, however, during the Barrier-trained and Journey tasks (in agreement with Wood et al., 2000 and Frank et al., 2000), even though the sequences in these tasks were computationally simpler. Both the Barrier-trained and Journey tasks possessed fewer repeated segments and had fewer intervening segments between repeated segments than Sequence 2, were symmetrical, and did not have any paths that crossed physically. The lack of orthogonalization during the initial training on the Journey task showed that the simplicity of the task configuration was not the critical variable.

These results agree with many others showing that hippocampal activity does not correlate directly with location. Rats running on linear tracks, in two similar environments connected by a hallway (Skaggs and McNaughton, 1998), performing different tasks (foraging and directed movement) in the same space (Markus et al., 1995), in presence of unstable distal cues (Muller and Kubie, 1987), and disorientation prior to maze entry (Knierim et al., 1995) all show different hippocampal firing patterns in the same physical place under different conditions. Unlike those studies, however, the present experiments show both orthogonalization and the lack thereof when the rat is moving in the same direction, through the same space, with the same distal cues, and while performing the same task.

A more likely explanation involves the interaction of the different types of information that were available to the hippocampus, while the rat was running along a repeated segment (Figure 5-12). In each task, local view information was present along the repeated segment, but provided no information to aid in separating the two contexts. Working memory for recently visited locations, or for future goals, however, was present. Working memory has been shown to be capable of orthogonalizing hippocampal codes for similar environments (Skaggs and McNaughton, 1998), and a similar ability could exist for activity related to recently visited locations or future goals. The presence of differential activity on the first day on the Barrier-trained tasks showed that the physical presence of the barriers was sufficient to orthogonalize hippocampal codes along the repeated segment, and suggests that recollections of the barriers from long-term memory should be capable of driving orthogonalization through the process of pattern completion.

	Local View	Working Memory	Long-Term Memory	Diff?
Sequence 2	✓	0	X	No
Journey (Day1)	✓	0	X	No
Journey	✓	✓	X	Yes
Barrier (Day 1)	✓	✓	X	Yes
Barrier	✓	0	✓	Yes

✓ Information present in CA1 during repeat
X Information not present in CA1
0 Hypothesis: not present in CA1

Figure 5-12. Comparison of information present along repeated segments and the presence of different hippocampal activity patterns. "Local View" stands for local view information, "Working Memory" for working memory for the start or goal of the current context, "LTM" for long-term memories of polarizing cues (i.e., the wooden barriers used during training), and "Diff?" for the presence of differential activity along the repeated segment. A large checkmark signifies the presence of information, or differential activity. "X" signifies the absence of information, or differential activity. "0" signifies that, by hypothesis, is not present in the hippocampus during the repeated segment.

This model suggests that orthogonalization failed to occur during Sequence 2 because the working memory for the start of the current context and activity related to the final goal were insufficient to drive orthogonal hippocampal codes, because the start and goal were separated from the repeated segment by reward stimulation, and were not part of the same, continuous path. The additional influence of long-term memory for the barriers during the Barrier-trained task, however, was enough to divide the code for the repeated segment. The lack of orthogonal codes on the first day of the Journey task arose because working memory and goal-related input were not strong enough to override the similarities in local view input. With further experience, however, these inputs separated hippocampal codes on the repeated segment.

One alternative explanation involves the direction of turns upon exiting the common arm. In both Sequence 1 and Sequence 2, three rats favored CW turns (rats 1, 2 and 3), and one favored CCW turns (rat 4), regardless of the nominal turn direction. This did not correlate with the hemisphere in which the hyperdrive was implanted (rat 3 was the only rat with a hyperdrive in the left hemisphere). The turn bias could have resulted from slight motor artifact induced by MFB stimulation, or from rats simply learning to turn in one direction early in training, and maintaining that direction across different tasks. This explanation, however, is not supported by the observation that a similar turn bias existed for each rat in all three sequences, including sequences in which differential activity was observed. Another alternative explanation is that the presence of rails to guide trajectories also aided the separation of hippocampal codes, but this explanation is not supported by the lack of orthogonalization during the first day on the Journey task

("Jny1" in Figure 5-11). Another possibility is that orthogonal hippocampal codes were required for rats to achieve the highest possible levels of performance. Rats chose the correct zone after the choice point most consistently on tasks that also showed orthogonal hippocampal codes; i.e., the Barrier-trained and Journey tasks. The best performing rat on the Sequence 2 task (rat 3), however, showed the least amount of separation between hippocampal maps along the repeated segment.

Hippocampal lesioned animals are impaired on tasks involving learning the sequence of odors, suggesting that the hippocampus plays a critical role in the representation of non-spatial sequences (Fortin et al., 2002; Agster et al., 2002; Kesner et al. 2002). These results can be reconciled with the results presented here by considering the training methods used in each experiment. In both Kesner et al. and Agster et al., the task could be solved by use of the strength of the memory trace for recent events, rather than by encoding or recalling the sequence itself. In Fortin et al. (2002), rats learned two sequences of odor pair discriminations containing repeated segments. Rats were trained to choose the odor that was rewarded prior to beginning comparisons common to both sequences. Rats lesioned using ibotenic acid were impaired at this task, while rats receiving radiofrequency lesions (the same method used by both Kesner et al. and Agster et al.) were less impaired and remained capable of learning the task following the hippocampal lesion. Fortin et al. suggested that the larger deficit following ibotenic acid could have resulted from damage of hippocampal target structures, particularly the entorhinal cortex. While this task approximates a non-spatial version of that described in our results, the results of Fortin et al. could also be explained by the method in which rats

learned each sequence of odors. The first sequence was presented individually, multiple times across several training sessions until a criterion was reached, and then the second sequence was presented. Errors following hippocampal lesions could then result from an inability to recall or activate the correct sequence stored in an extra-hippocampal structure, rather than a loss of the representation of the sequence. If this were the case, errors in lesioned animals should result primarily from perseverance.

Chapter 6 PHASE RESETTING OF HIPPOCAMPAL THETA TO ACCELERATION

Abstract

Computations involving distributed representations must synchronize, or coordinate, activity of those representations if they are to be used as combined inputs of a single computation. Oscillations in local field potentials provide one possible mechanism for such synchronization. In the rat hippocampus, local field potentials in the frequency range of 4-12 Hz, called the "theta rhythm", have been associated with different types of voluntary movements, sensory processing while the animal remains motionless, and REM sleep. In addition, the timing of place cell spikes within a cycle of the theta rhythm varies systematically with the rat's position in the place field of the cell. Firing shifts systematically in the theta cycle, i.e., "precesses" as a place field is traversed. Single unit activity in the subiculum, entorhinal cortex, and cingulate cortex is also modulated by theta, suggesting that theta may coordinate computations utilizing information stored within these structures. Many studies have shown that during brief movements, such as lever pressing or jumping, the amplitude and phase of theta both increase. Other studies have shown that the frequency decreases during brief periods of sensory coordination, such as discerning between audio and visual cues. Little is known, however, about theta in relation to movements over a longer duration, such as uninterrupted, whole-body

trajectories. In the present study, rats were trained to run to distant goals. The phase of theta was reset to the time of peak acceleration during trajectories.

Introduction

Theta is a periodic signal present in the hippocampus during specific behaviors, including voluntary movements (Vanderwolf and Heron, 1964); (Yokota and Fujimori, 1964) and the coordination of sensory stimulation. Goal-directed movements require the synchronization of limb movements with brain processes associated with navigation to goals, but little is known about any such synchronizing mechanism.

Theta is present throughout the duration of several "voluntary" movements, including running and jumping, but not "consumatory" behaviors, such as grooming, eating or copulation (Vanderwolf et al., 1975). The amplitude and frequency of theta increase prior to the initiation of discrete movements such as the pressing of a lever at the initiation of a jump (Vanderwolf, 1969). In rats trained to make vertical jumps, the magnitude of the change in amplitude and the rate of frequency change are related to the speed and distance of the jump (Morris, 1976). Changes in theta frequency have also been associated specifically with the coordination of sensory stimuli, rather than with movement. The phase of theta was observed to reset at light onset in rats trained to lever press following the onset of a cue light (Williams and Givens, 2003).

Theta reset during navigation tasks requiring whole-body movements could coordinate sensorimotor integration and brain processes related to navigation. If the hippocampus is involved in trajectory computation, even if only by providing a

representation of the current location of the rat, then the projection of this information to other structures might require some kind of coordination with the ongoing activity in those structures. The timing of place cell spike vary with respect to theta, firing earlier in the theta cycle as the place field is traversed (O'Keefe and Recce, 1993). This relationship is complex even for running at a constant speed (Skaggs et al., 1996;Yamaguchi et al., 2002).

Methods

Time-averaged theta

Theta was recorded at the hippocampal fissure across 4 rats (1221 runs from rat 1, 1173 runs from rat 2, 3912 runs from rat 3, and 2323 runs from rat 4) with respect to an electrode located in the Corpus Callosum. EEG data from each trial were centered on the time of trajectory onset and normalized so that the maximum amplitude within the EEG data for each trial was 1.0. The data were averaged, and the standard error of the mean computed.

Summing behavior-centered theta

Four events within each trajectory were used to center windows of EEG data: cue onset, the onset of the trajectory, the peak of acceleration, and the peak of velocity. For each event, windows of EEG data from each trial from 1 sec prior to 2 sec following the event were centered on the event. Each sample of windowed data was down-sampled to a uniform sampling frequency of 500 Hz, and digitally filtered between 2 and 25 Hz. The

data values within each window were normalized by the maximum amplitude of the signal within the window; i.e. the maximum amplitude within each window was 1, regardless of the event on which data in that window had been centered.

A time window with a duration of 1 second was moved along each data sample in 200 msec increments, starting one second prior to the centering event and continuing until two seconds following that event, generating 11 over-lapping windows. The spectrum of each window was then computed using discrete, prolate spheroid (DPS) kernels (the "pmtm" function in Matlab, Mathworks Inc., Boston, MA). One advantage of spectral estimation using DPS kernels is that the estimates are normally distributed, allowing direct computation of confidence limits, which were computed for each window in both the time and frequency domain.

Monte Carlo estimate of variance

A variance estimate for spectrograms was generated using a Monte Carlo simulation. A separate simulation was run for EEG windows centered on each behavioral event. After being centered on a particular event and down-sampled, the signal in each window was shifted by a random amount of time. The time shifts were uniformly distributed across the period of a theta cycle (100 msec), and a different time shift was used for each window. The average signal across all time-shifted windows was computed. This process was repeated 1,000 times. The group mean and standard error were computed. The mean and standard error were compared to the mean of the non-shifted signals to determine statistical significance.

Results

Time-averaged EEG

EEG data came from 8629 trials across 17 sessions in four rats (rat 1 had 1221 runs across four sessions; rat 2 had 1173 runs across four sessions; rat 3 had 3912 runs across five sessions; rat 4 had 2323 runs across four sessions). If activity were synchronized to the onset of trajectories, then the phase of the dominant frequency between 2 and 25 Hz (ostensibly theta) should be roughly equivalent across trials, and should produce significant peaks in the time-averaged signal due to constructive interference of waves. Significant peaks were observed in the time-averaged signal beginning approximately 400 msec after the onset of trajectories and continuing for approximately 700 msec (Figure 6-1).

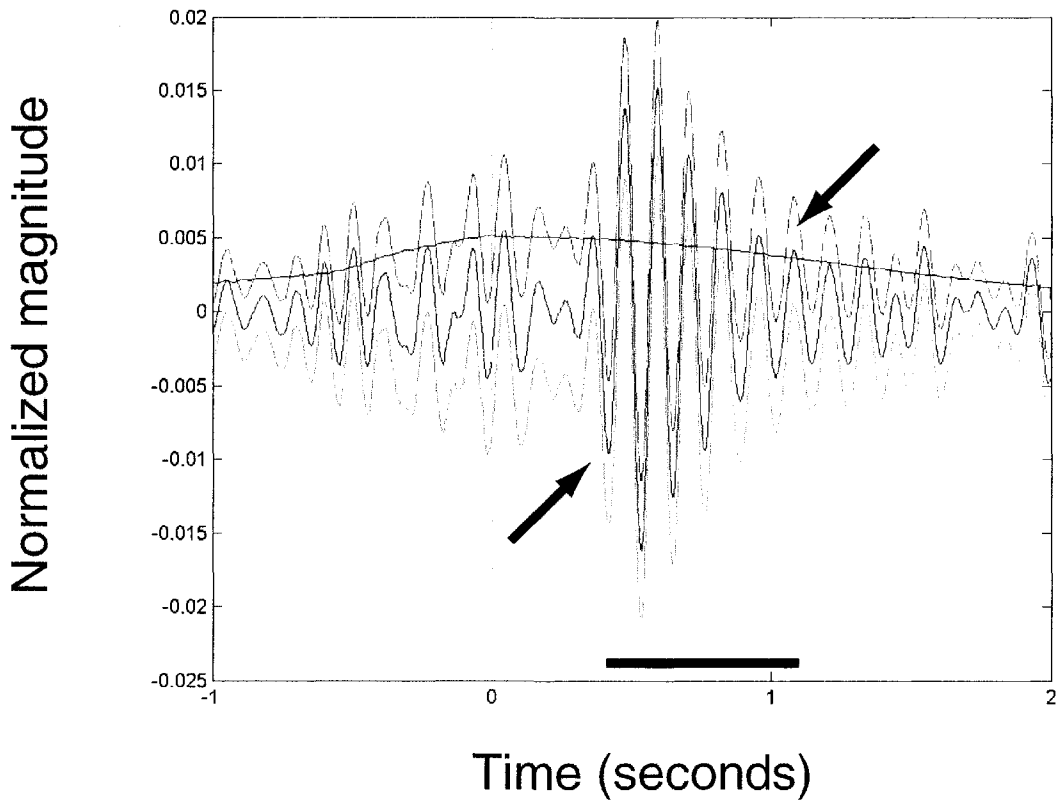

Figure 6-1. Hippocampal theta is reset around the time of the onset of trajectories. Hippocampal EEG is reset at the onset of trajectories, beginning roughly 400 msec (the horizontal, black bar) after the initiation of the trajectory (the vertical blue line centered on "0"). EEG samples were collected at roughly 1 kHz from the hippocampal fissure starting one second prior to and ending two seconds after the initiation of each trajectory as determined by video tracking of position (see General Methods) in 8629 trials across 17 sessions in four rats (rat 1 had 1221 runs across four sessions; rat 2 had 1173 runs across four sessions; rat 3 had 3912 runs across five sessions; rat 4 had 2323 runs across four sessions). The EEG data were filtered between 4 and 12 Hz, and then normalized to the peak filtered magnitude within each trial. The waveforms were then averaged across all trials (blue waveform), and the standard error of the mean (red and green waveforms) computed (the black waveform shows the magnitude of the confidence limits). Significant peaks due to synchronization across trials occur when neighboring valleys of the upper (red) confidence limit pass below the intervening peak of the lower (green) confidence limit, and vice-versa. The first and last such peaks (black arrows) bracket an extended period of time (black bar) running from roughly 400 to 1100 msec following the onset of trajectories ("0").

Because theta arises from multiple generating mechanisms that each project to different CA1 dendritic layers (Buzsaki, 2002), the relationship between theta phase and pyramidal cell spiking activity varies with depth. A more emphatic result might be

obtained by aligning theta phases across sessions, though the result obtained without aligning phases was substantial.

Comparison across different centering events

Because theta appeared to be synchronized to some event occurring several hundred msec after the onset of trajectories, EEG data centered at the moment of peak acceleration and peak velocity were compared to data centered at cue onset to see if either produced greater phase resetting. Data from 6 sessions across four rats (2 sessions from rat 1, 2 sessions from rat 2, 1 session from rat 3, and 1 session from rat 4) were centered on cue onset, time of peak acceleration and time of peak velocity. The data were time-averaged, and spectra computed for sliding windows with durations of one second. To allow comparison of spectral magnitudes across sessions, the peak of all spectra within a single recording session were normalized to the maximum spectral value within that session. Across the six sessions, the mean and standard error were computed (Figure 6-2). Windows of EEG data occurring at the time of peak acceleration and for approximately one second following that time showed significantly greater power in the theta frequency range of 4-12 Hz (three-way ANOVA, $p=0.05$).

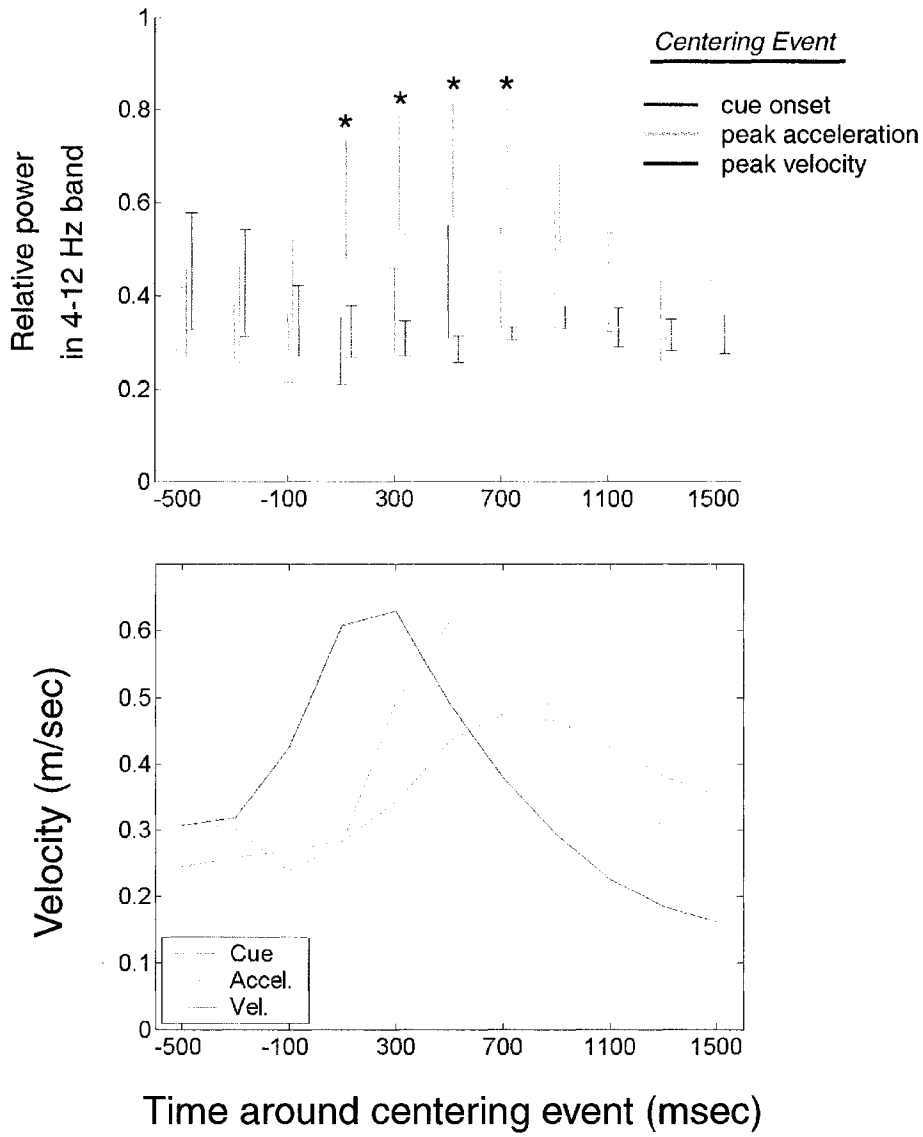

Figure 6-2. Time-averaged EEG centered on the peak of acceleration contains more relative power in the theta band. (TOP) EEG data for each run for six sessions across four rats (2 sessions from rat 1, 2 sessions from rat 2, 1 session from rat 3, 1 session from rat 4) were centered on three behavioral events and a time average computed for each. The resulting waveforms were divided into overlapping, one second-duration sub-windows, at 11 different time points, separated by 200 msec. The spectrum of each sub-window was computed, and the maximum value of the spectrum between 4 and 12 Hz was found, producing 11 different maxima for each of three different centering events (cue, acceleration and velocity). These 33 values were normalized to the largest value among them. The mean and standard error of these values across sessions were computed and plotted. Significant differences in means were computed using a three-way ANOVA, p=.05. (BOTTOM) The mean distance traveled within each sub-window. The mean velocity within a sub-window was greatest when the centering event was the moment of maximum velocity for each trajectory (blue line). The increase in the strength of theta of the summed EEG is not explained by an increase in velocity alone. That is, power in the theta band is greatest when the data were centered on peak acceleration, but the increase in theta power does not correlate with increased velocity.

Monte Carlo estimate of variance

To test whether the power in the theta frequency band arose from the phase coupling of the theta in the time domain, an estimate of the variance of spectra computed from time-averaged waveforms was computed using a Monte Carlo simulation. The data from each trial were shifted independently by ±50 msec (roughly one half of the period of theta), averaged in time, and then the spectrum computed. This process was repeated 1,000 times for data in each window, and a mean and standard error computed. The spectrum of the original, non-shifted data was compared to that from the randomly shifted data (Figure 6-3).

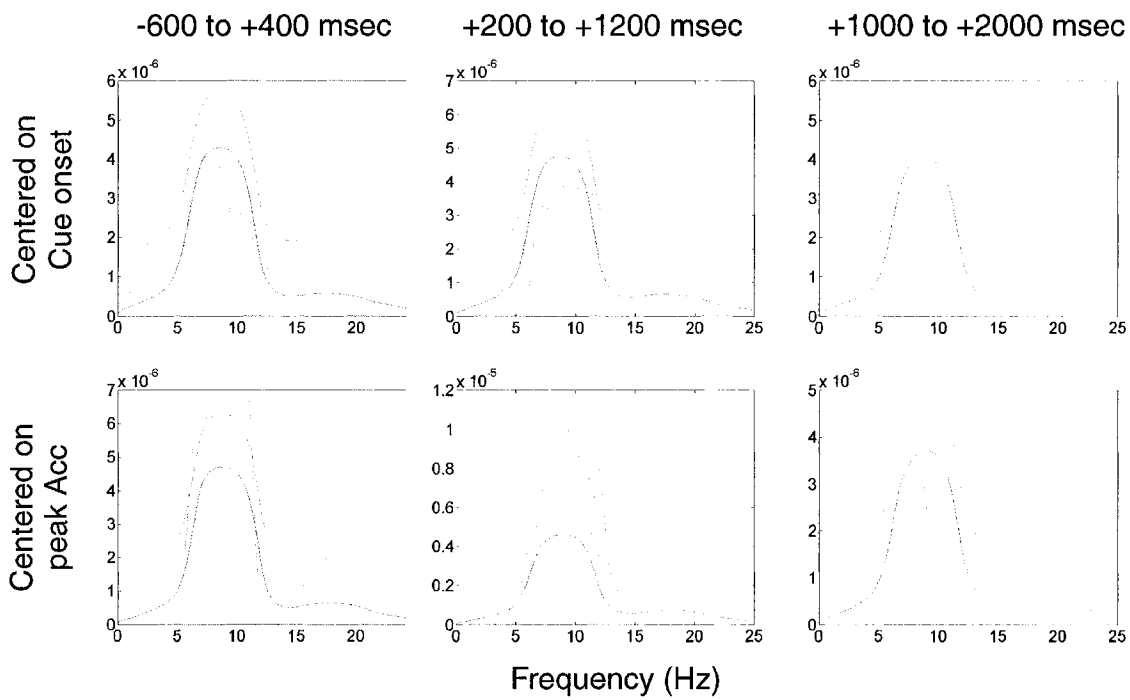

Figure 6-3. Significance of spectral peaks computed using a Monte Carlo simulation. EEG data from one session from rat 1 were centered on cue onset and peak acceleration, and divided into one-second duration windows, as described above. The top row shows the results from data centered on cue onset, and the bottom row shows the results from data centered on the peak of acceleration. Each column shows the results from data in different time windows. The data for each trial were shifted by ±50 msec (roughly half the period of a theta cycle), summed, and a spectrum computed. This was done 1,000 times for the data in each window. The mean (plotted in blue) and standard error of the mean (plotted in red) were computed. The spectrum of the original data was computed (plotted in green). A peak in the spectrum of the original data greater than one SEM above the mean was considered significant. A significant peak in the spectrum was observed only for the window containing data from 200 to 1200 msec following the time of peak acceleration.

Discussion

During ballistic trajectories to distant targets, the spectral power of hippocampal theta increases at the onset of movement and continues until the target is reached. Across trials, the phase of theta becomes consistently reset to the onset of the trajectory. The phase resetting is significant for data centered on the time of peak acceleration, but not for data centered on the time of the cue onset of on the time of peak velocity. This result agrees with the findings of other experiments regarding the reset of the phase of theta to discrete movements, such as lever pressing and jumping. These results expand those findings by showing that the resetting of theta during extended trajectories lasts for a substantial amount of time (at least one second)

Chapter 7 COMPARISON OF PARIETAL AND HIPPOCAMPAL PHYSIOLOGY

Introduction

Because parietal units were recorded from tetrodes that were subsequently moved into the hippocampus, histological verification of the depth of tetrode placement in posterior parietal cortex was not possible. This raised the possibility that tetrodes nominally located in parietal cortex might have been advanced into hippocampus, instead. The physiologic properties of hippocampal units observed during extra-cellular recordings are well-known (O'Keefe and Dostrovsky, 1971;McNaughton et al., 1983a;Muller et al., 1987). Two classes of hippocampal units can be identified: pyramidal cells and interneurons. Hippocampal pyramidal cells have a spike width of more than 300 μsec, a mean firing rate less than 5 Hz, and have an asymmetric wave shape; the early phase of a recorded spike is more rapid and peaked than the later phase. Hippocampal interneurons have a spike width less than 300 μsec, a mean firing rate greater than 5 Hz, and have a symmetric wave shape. In addition, hippocampal pyramidal cells that are active during running fire preferentially in restricted locations (their "place fields"), generating more spatial information per spike than interneurons, which spike across a larger area (Skaggs et al., 1996). Another property that differentiates hippocampal pyramidal cells and interneurons is the distribution of intervals between spikes. Hippocampal interneurons fire tonically (Schwartzkroin and Mathers, 1978). When the intervals between spikes are binned and the distribution

plotted as a histogram, the tonic firing of interneurons produces a unimodal distribution. Hippocampal pyramidal cells, however, spike in bursts, both during sleep and while running. Bursts during sleep are known as "complex spikes." (O'Keefe and Dostrovsky, 1971) Bursts during awake running are associated with place field activity. In both cases, precisely timed bursts of spikes produce a multi-peaked distribution of inter-spike intervals (ISI). Units nominally recorded from hippocampus during these experiments should possess properties similar to one of these two classes. The properties of parietal unit waveshape and ISI distribution during extra-cellular recordings in freely moving rats are not as well known.

Methods

Nominal tetrode depth

The depth of each tetrode was recorded as the tetrode was advanced through cortex towards the hippocampus. When implanted, each tetrode was embedded in a protective polymer layer (Silastic). Following implantation, each tetrode was advanced through the polymer, while the electrical resistance through the tetrode to the grounding screw was monitored. When the tetrode emerged from the polymer and made physical contact with the brain, the resistance of the tetrode dropped from more than 2 Mohms to less than 1 Mohm. The depth of the tetrode was recorded as nominal "zero" depth. Subsequent movements of each tetrode were recorded as a distance from the brain surface, i.e., from nominal zero. For any single recording session, then, each tetrode was associated with a depth from the brain surface. This nominal depth could be compared

with known anatomical depths of both parietal cortex and the hippocampus. In addition, the depth at which each tetrode reached the cell body layer of CA1 provided a second reference depth. Nominally, at the brain surface coordinates used in these experiments (4.0 mm posterior, 2.0 mm lateral to Bregma) parietal cortex extends 1.0-1.5 mm down from the brain surface. The Corpus Callosum extends another 0.5-1.0 mm below parietal cortex. The nominal depth of CA1 of the hippocampus is 2.0-2.5 mm below the cortical surface. Parietal units should be no more than 1.5 mm from the brain surface, and no less than 0.5 mm from CA1.

Spike width

Spike waveforms contained 32 voltage samples collected at 32 kHz; i.e., the total duration of the sampling window was 1 msec (Figure 7-1). The sampled data points were fit with a spline to approximate the waveform on a finer time scale. The width of a spike was computed by finding the time difference between the maximum of the waveform ("peak") and the minimum ("valley").

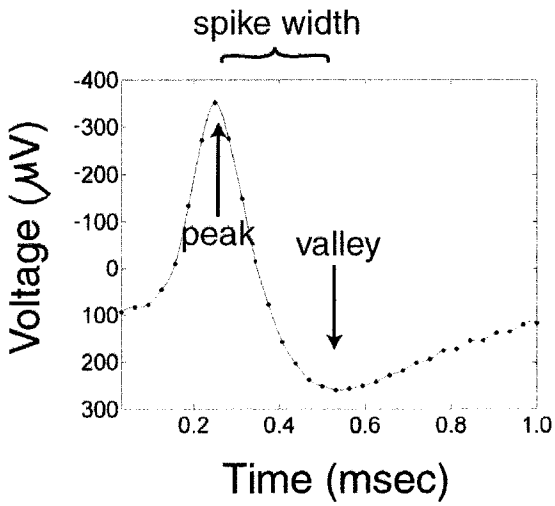

Figure 7-1. Characteristic waveform of a single unit during extra-cellular recording. The horizontal axis shows time. The vertical axis shows the voltage of the recorded spike. In the 1 msec of time shown in the figure, the voltage was sampled at 32 times (blue dots). These sampled points were fit with a spline (blue line). The time between the maximum of the fitted curve (peak) and the minimum (valley) was the spike width.

First inter-spike interval peak

A timestamp was associated with each spike from a given unit. The time difference between adjacent spikes (i.e., the "inter-spike interval", or "ISI") was computed across all spikes that occurred during a single recording session. A histogram was generated by binning the logarithm of the intervals. The histogram was smoothed by convolution using a Gaussian kernel. Local maxima ("peaks") were found. The peak in the smallest time bin was called the "first ISI peak" (Figure 7-2).

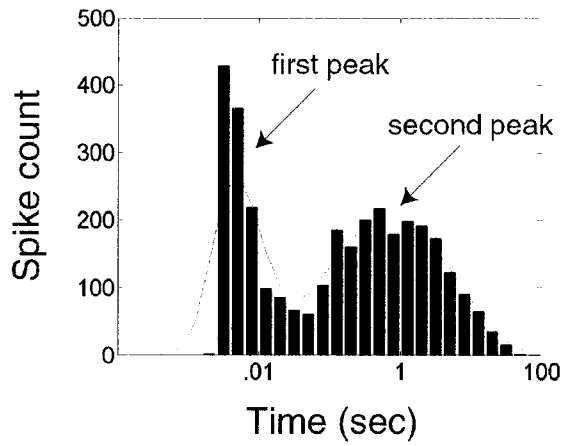

Figure 7-2. Computation of first Inter-Spike Interval (ISI) peak. The horizontal axis shows time on a logarithmic scale. The vertical axis shows the number of spikes in each bin. The spike counts were convolved with a Gaussian to smooth the data, and the local maxima ("peaks") were found. The first peak is the peak in the smallest time bin.

Firing Rate

The average firing rate of a unit was computed by dividing the number of spikes associated with a unit by the total recording time. Data from all three recording epochs (i.e., Sleep 1, Maze and Sleep 2) were included.

Spatial information

The spatial information contained in the spikes of a given unit was computed as described by Skaggs et al. (Skaggs et al., 1996). Briefly, spatial information describes how much information each spike contains regarding the spatial location of the rat. A unit that spiked only when the rat was in a particular location would have a high spatial information content. A unit that fired regardless of the location of the rat would have a low spatial information content. Spatial information was computed by binning both

occupancy and spikes in a two-dimensional histogram. The probability of being in each bin, p_i, was computed by dividing the amount of time spent in the bin by the total time. The average firing rate for each bin, λ_i, was computed by dividing the total number of spikes that occurred while the rat was in that bin by the time spent in that bin. The spatial information for a given unit was

$$S = \sum_i p_i \left(\frac{\lambda_i}{p_i} \right) \log \left(\frac{\lambda_i}{p_i} \right)$$

Results

Parietal unit data were recorded from two rats running to randomly presented goals across three recording sessions.

Nominal unit depth

Nominal tetrode depth at which the putative parietal units were recorded were computed with respect to the brain surface. The additional depth moved by each tetrode before it reached the CA1 cell body layer was found by taking the difference between the depth of the tetrode at which putative parietal units were recorded and the depth at which that tetrode encountered CA1. Parietal cortex in the rat extends 1.0 to 1.5 cm below the cortical surface. PPC is separated from Hippocampus by fiber tracts of the Corpus Callosum and Alveus by a distance of 0.5-1.0 mm (Paxinos and Watson, 1997). The nominal recording depth across all putative parietal units ranged from .64 to 1.04 cm below the cortical surface. Tetrodes that recorded putative parietal units moved at least 700 μm before being moved into the cell body layer of CA1. From Paxinos and Watson

(1997), units at putative cortical depths of 500-1,000 μm correspond to layers IV and V, while units deeper than 1,000 μm correspond to layer VI.

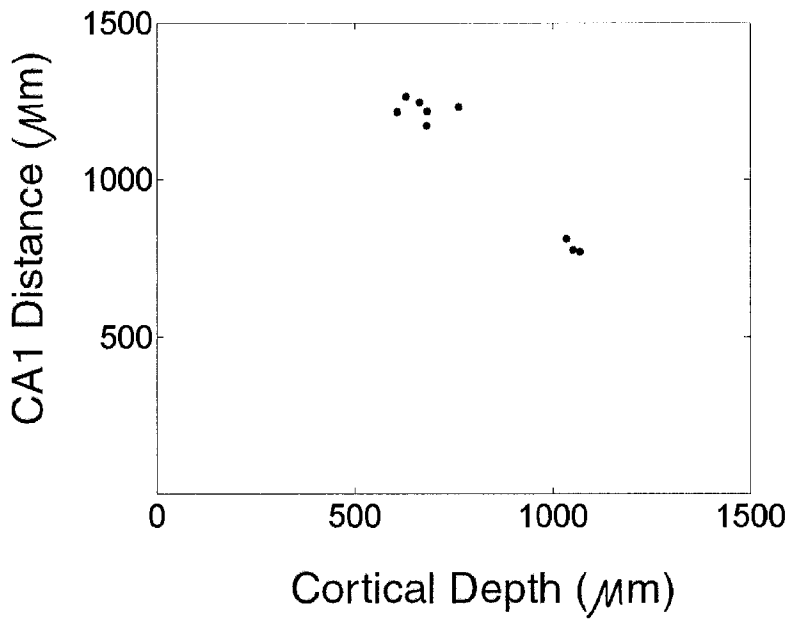

Figure 7-3. Nominal separation between the recording location of parietal units and the cell body layer of CA1. The horizontal axis shows the depth of the recording electrode from the brain surface. The vertical axis shows the additional depth moved by the recording electrode before being moved into the cell body layer of CA1. Coordinates were dithered by ± 50 μm to separate overlapping points.

Single unit physiology

A second indirect test examined the physiology of recorded units nominally located in parietal cortex and hippocampus. Physiological characteristics of hippocampal pyramidal cells and interneurons are well known (Ranck, Jr., 1973;McNaughton et al., 1983a). For hippocampal pyramidal cells, the spike width (the time from the peak to the valley of a spike waveform) is greater than 300 μsec, and the mean firing rate is less than 5 Hz. For hippocampal interneurons, the spike width is less than 300 μsec, and the mean

firing rate is greater than 5 Hz. In addition, the distribution of the time duration between spikes (the Inter-Spike Interval, ISI) for the two units is different. The ISI distribution of hippocampal pyramidal cells contains multiple peaks, due to presence of rapid, precisely timed groups of spikes known as "complex bursts". The ISI distribution of hippocampal interneurons is unimodal, reflecting tonic firing. Figure 7-4 shows the comparison for units that were recorded nominally from both parietal cortex and hippocampus.

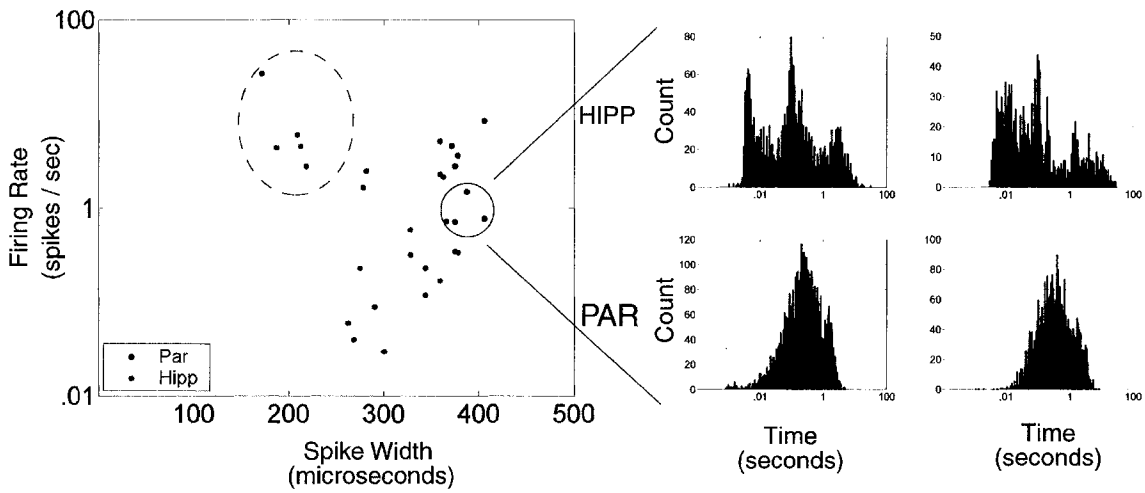

Figure 7-4. Comparison of spike width and firing rate. On the scatter plot, the horizontal axis shows the mean spike width (the time from the peak to the valley of the recorded spike) across all valid recording channels. The vertical axis shows the mean firing rate across the entire recording session. The blue dots represent the parietal units that showed significant theta modulation from two rats across three sessions: 8 units from one rat, and 1 unit from the second. The red dots represent a sample of units recorded from the hippocampus during the same sessions. (dotted circle) Units identified as interneurons by mean firing rate, spike width and waveform shape (McNaughton et al., 1983a) are distinguishable from the remaining hippocampal units, which were identified as pyramidal cells. (solid circle) Two units from parietal cortex and hippocampus are examined in more detail on the right side of the figure. The histograms were binned according to the logarithm of the Inter-Spike Interval (ISI) and plotted on a common time axis. (top row) The ISI distribution of hippocampal pyramidal cells is multi-peaked due to bursts of closely timed spikes. (bottom row) The ISI distribution for parietal units is unimodal due to a more constant temporal firing pattern. Peaks in the ISI distribution at times longer than 1 second are due to artifact introduced by blocking MFB stimulation artifact. Blocking MFB stimulation artifact created an artificial period of inactivity that generated the late peak.

A two-step process was used to separate parietal units from hippocampal pyramidal cells and interneurons. First, the first ISI peak latency and spatial information were used to identify hippocampal place cells. During extra-cellular recordings, the waveform of hippocampal pyramidal cells and putative parietal cells are similar, but the ISI distributions of the two cell types are different (Figure 7-4). The first ISI peak latency was used to along with spatial information to identify the characteristics of putative parietal units. Units in the hatched area had a first ISI peak greater than 20 msec and less than 2 bits/spike of spatial information.

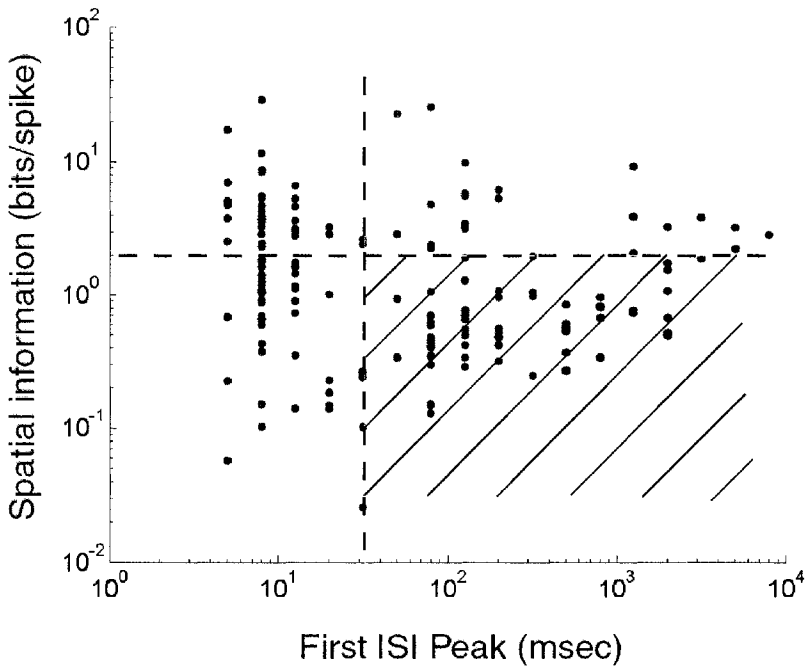

Figure 7-5. Comparison of spatial information and latency to first peak in the ISI distribution. The horizontal axis represents the time of the first peak in the ISI distribution. The vertical axis represents the spatial information for each unit for the duration of the task. (Blue) Hippocampal units with a latency less than 25 msec. (Red) Hippocampal units with a latency greater than 25 msec. (Blue) Parietal units. Parietal unit activity contains less than 2 bits/spike of spatial information and has a first ISI peak latency greater than 20 msec (hatched area).

The units selected in Figure 7-5 (those in the hatched area) were plotted as a function of spikewidth and mean firing rate (Figure 7-6). Putative parietal neurons showed a wide range of mean firing rates, but tended to have spike widths greater than 300 msec.

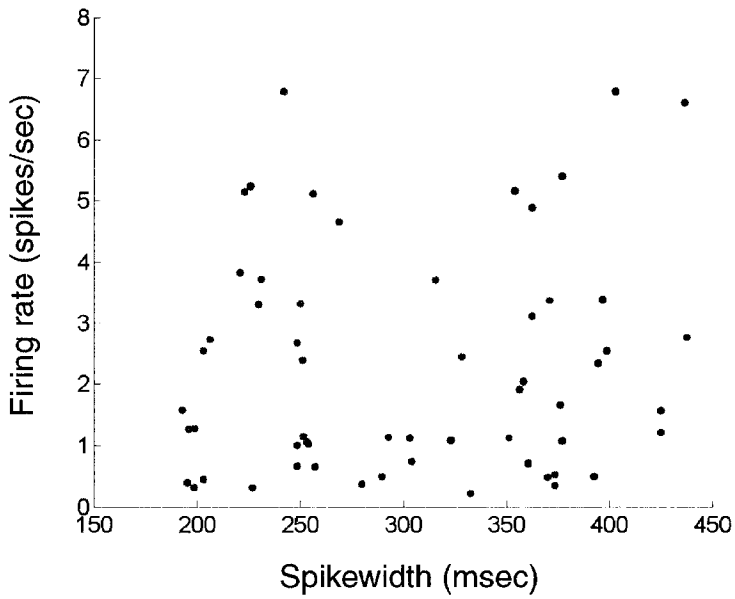

Figure 7-6. Spike width and firing rate for units with less than 2 bits/spike of spatial information and a first ISI peak latency greater than 20 msec. The horizontal axis shows the spike width (peak-to-valley time). The vertical axis shows the mean firing rate across the duration of the recording session. (red dots) hippocampal units. (blue dots) parietal units.

Discussion

Putative parietal units were identified on the basis of nominal tetrode depths and physiologic parameters. Parietal units were located within 1.5 mm of the brain surface, and not less than 0.7 mm from the cell body layer of CA1. The first ISI peak latency of

putative parietal units was greater than 20 msec, the spatial information was less than 2 bits/spike, and the spike width was greater than 300 msec.

Chapter 8 THETA MODULATION OF PARIETAL NEURONS

Abstract

Both the firing rate and the timing of individual neuronal spikes can convey information, and in some cases both do. An example of combined encoding is given by hippocampal place cells in rats, which display increased firing rates in specific locations, while also displaying phase-dependent spiking with respect to 6-10 Hz oscillations in field potential in different structures of the hippocampus ("theta"). Both firing rate and theta phase appear to convey allocentric spatial information: Firing rates of ensembles of neurons can be used to decode the spatial location of the rat, and the order of activation of multiple place cells can be decoded by noting phase differences with respect to theta between the spikes of different, activated place cells. In non-human primates, the firing rate of units in Posterior Parietal Cortex (PPC) is correlated with egocentric spatial processing, but it is unknown whether the timing of individual spikes also conveys information. To determine whether temporal information relevant to spatial orientation of navigation might be present in the rodent PPC, rats were trained to run to targets located at the edge of an open arena, while multiple single parietal units were recorded, along with hippocampal theta. Some parietal cells were found to respond preferentially to theta phase, in a manner similar to hippocampal place cells, but roughly 120° out-of-phase with hippocampal pyramidal cells. This suggests that theta may be involved in the synchronization of spatial information contained in PPC and hippocampus. Though theta synchronization of single unit activity has been observed in extra-hippocampal structures,

these results constitute the first evidence of phase locking to hippocampal theta by single units in a brain structure not associated with the limbic region.

Introduction

The nervous system controls our movements as we reach for objects, direct our gaze, or move towards distant targets. In non-human primates, a rich literature exists regarding the involvement of single unit activity in parietal cortex in the encoding, computation and learning of reaching and gaze trajectories to targets (for review see (Andersen, 1997;Colby and Goldberg, 1999). Less is known, however, about the role of parietal single unit activity during movement in rodents. Single units in rodent in parietal cortex respond to the conjunction of specific behaviors and locations (Chen, 1989), as well as to head direction (Chen et al., 1994b), and so could play a role in trajectory computation. Rodent hippocampal pyramidal cells are known to respond selectively to place (O'Keefe and Dostrovsky, 1971), and are phase locked to a 7-10 Hz oscillation in the local field potential called "theta" (O'Keefe and Recce, 1993). It has been proposed that hippocampus and neocortex provide "complementary learning systems" for the acquisition and long-term storage, respectively, of memory (Marr, 1970;Marr, 1971), particularly spatial memory (McClelland et al., 1995), but it remains unclear how such interactions support moment-by-moment navigation, such as the computation of a trajectory to a distant goal.

In rats, hippocampal theta is a regular, low frequency (4-12 Hz) oscillation present in the dentate gyrus, CA3 and CA1 during "exploratory" behaviors (e.g., running

and jumping), during "consumatory" behaviors (e.g., eating, grooming and copulation) (Vanderwolf et al., 1975), and during REM sleep (Winson, 1972). The function of theta remains largely unknown (Buzsaki, 2002), but the firing of hippocampal neurons is correlated with the phase of theta (Sinclair et al., 1982;Buzsaki et al., 1983). Within place fields, the correlation between spiking and theta phase changes as a function of location within the place field, firing earlier in the theta cycle as the place field is crossed (O'Keefe and Recce, 1993). Theta modulation of unit activity has been observed in several cortical structures outside the hippocampus including entorhinal cortex (Mitchell and Ranck, Jr., 1980) and in cingulate cortex (Leung and Bostock, 1987), all of which are associated with the so-called limbic system. Less is known about the timing of sensory or motor spikes with respect to large-scale oscillations in the brain. In non-human primates, single unit activity in visual cortex synchronizes to a local oscillation of 40-60 Hz ("gamma") in response to specific stimuli, and all units within a visual cortical column maintain this same phase relationship (Gray et al., 1989). In addition, cells in distant cortical columns that are activated by the same stimuli will show the same phase relationship (Gray et al., 1989), putatively "binding" features across specific sub-regions within a structure. Evidence that hippocampal and parietal activity interact has been found in unit recordings from both structures during sleep before and after a random foraging task. Prior work had shown that pairs of hippocampal units that were coactive during foraging showed increased correlations in subsequent sleep when compared to the correlations observed in prior sleep (Wilson and McNaughton, 1994). As in the case of co-active pairs of hippocampal units, parietal-hippocampal unit pairs that had correlated

activity during the foraging task showed increased correlations during sleep following the task when compared to correlations observed during sleep preceding the task (Qin et al., 1997). In addition, during sleep in the rat, cortical unit activity increases during hippocampal sharp waves, suggesting that periods of heightened activity within the hippocampus can lead to similar increases in activity levels across neocortex (Battaglia et al., 2001).

Methods

Physiology of units

Each unit discussed in this chapter met the criteria established for parietal units described in Chapter 6. Multiple single units were recorded using tetrodes contained in a "hyperdrive" (Gothard et al., 1996b) located –4.0 mm AP and ± 2.0 mm ML relative to Bregma. Each unit was recorded at a nominal depth of not more than 1.5 mm from the brain surface, and not less than 0.7 mm. The spike width of each unit was greater than 300 μsec, the spatial information was less than 2 bits/spike, and the first ISI peak was greater than 25 msec.

Computation of theta phase and degree of modulation

EEG data were recorded from the hippocampal fissure with respect to a ground electrode just below the cortex, and were sampled at 1 kHz. The EEG data were filtered between 2 and 12 using a digital filter with zero phase distortion. Theta phase was computed for each EEG sample using a Hilbert transform (Matlab, Mathworks Inc.,

Boston, MA). A phase was assigned to each spike by interpolation of theta phase at the EEG points preceding and following the spike time. Because theta is a periodic signal, the assignment of a "zero" phase was arbitrary. The phase associated with the maximum firing rate of hippocampal units was chosen as the "zero" phase. This could only be done directly for data sets in which both hippocampal and parietal units were recorded simultaneously, as was the case with Rat 1. In the case of Rat 3, all tetrodes were located in parietal cortex during the two recording sessions discussed in this section. Each of these tetrodes was moved into the hippocampus at a later date, and the zero phase for Rat 3 was approximated using hippocampal data from this later date. Spike phases were binned using 12 bins (30° per bin) across all valid runs. To determine the depth of theta modulation, a cosine function was fitted to the phase histogram of each unit, generating an amplitude and a phase that produced the best fit for each unit. The peak of the fitted curve was the mean value of the histogram height plus the amplitude of the fitted cosine, while the value of the valley was the mean value of the histogram heights minus the amplitude. Depth of modulation equaled $(peak - valley)/peak$ (Skaggs et al., 1996). The significance of modulation was determined by computing the circular mean and standard deviation, applying the Rayleigh test (p=.05) (Batschelet, 1981). Significant differences between circular means were tested using the Watson-Williams test (p=.05).

Results

Firing rates of some parietal units are modulated by hippocampal theta

Data were acquired from three rats running to randomly presented goals during two recording sessions each. 31 parietal units were recorded from one rat (rat 1 in this dissertation), 24 parietal units were recorded from the second rat (rat 3 in this dissertation), and 6 parietal units were recorded from the third rat (rat #6533, which will be called rat 5 in this dissertation). The firing rates of 19 units were significantly modulated by hippocampal theta, 12 of 31 units in rat 1, 5 of 24 units in rat 3, and 1 of 6 units in rat 5 (Figure 8-1). Directionality of phase response curves was computed using the circular mean and standard deviation (Batschelet, 1981), and significance established by the Rayleigh test ($p<.05$).

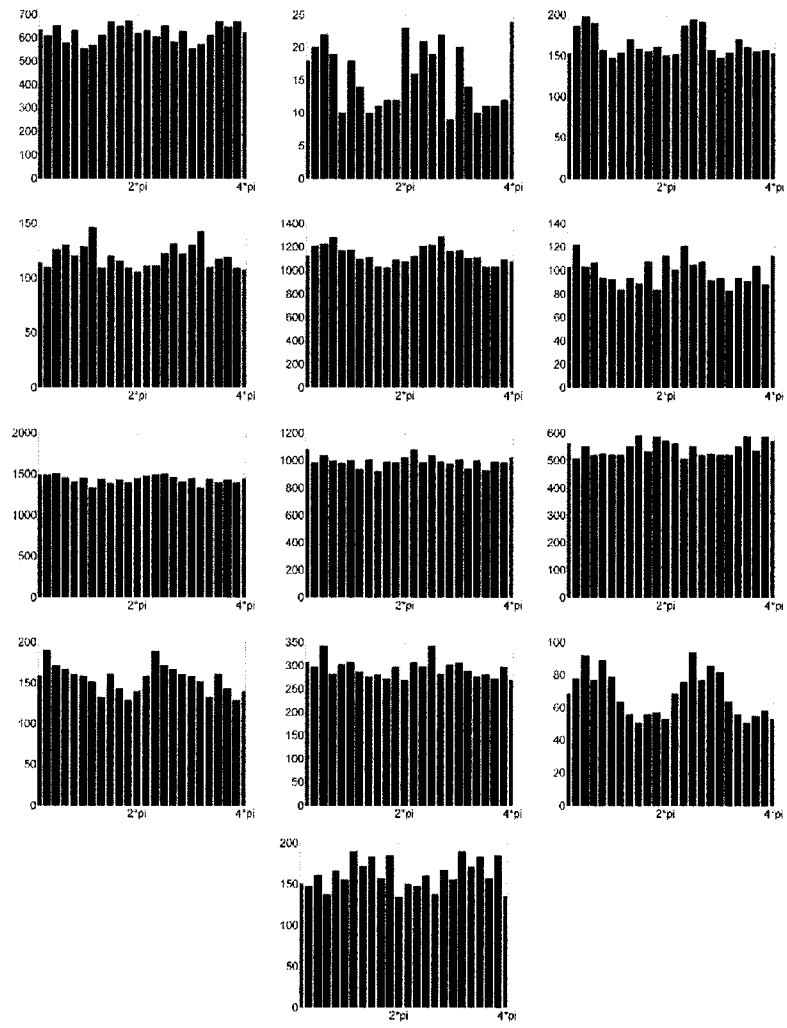

Figure 8-1. Theta-modulated firing rates of single units in Posterior Parietal Cortex during trajectories. The horizontal axis shows the phase of theta across two complete cycles. The vertical axis shows the total count of spikes for each bin. EEG was recorded from the hippocampal fissure, filtered between 2 and 12 Hz, and a phase associated with each sampled point using a Hilbert transform. A theta phase was assigned to each spike by interpolating the theta phase of the EEG samples before and after that spike. The phases were binned (bin width = 30°), and counted to generate a histogram. Significant modulation was found using the Rayleight test ($p<.05$).

Modulation depth

The histogram of each of the 18 theta-modulated units was fitted with a sinusoidal function. The magnitude of modulation was computed for each, and the results binned,

and compared to a sample of 15 hippocampal units obtained from rat 1, recorded simultaneously with the parietal units obtained from rat 1 (Figure 8-2).

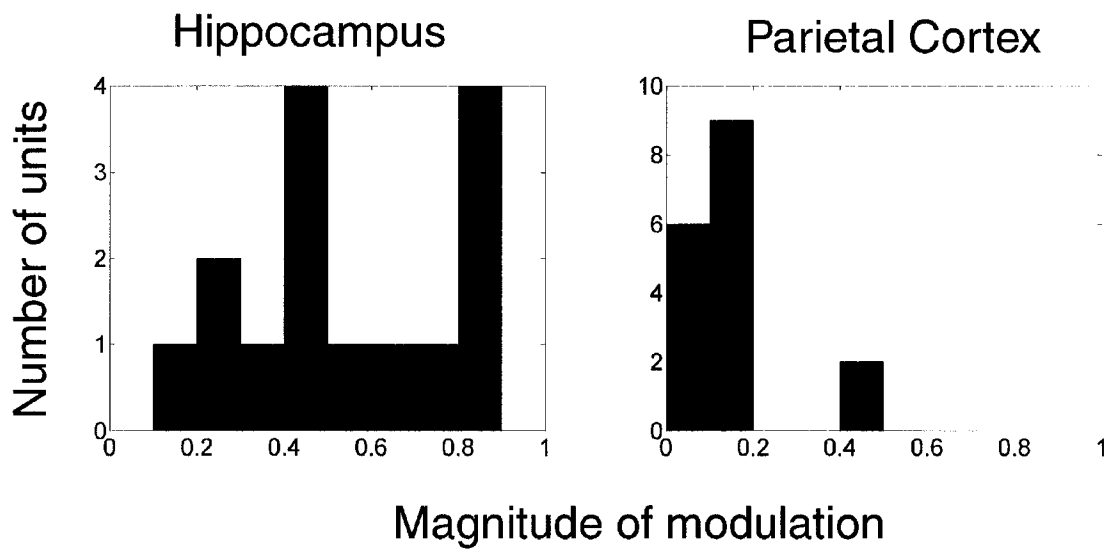

Figure 8-2. Comparison of magnitude of modulation between hippocampal and parietal units. The horizontal axis shows modulation as computed by (peak-valley)/peak. The vertical axis shows the number of units in each magnitude bin with significantly directional theta modulation (i.e., which were significantly modulated by theta (Rayleigh test, $p<.05$)).

Preferred phase

Simultaneous recordings of both hippocampal and parietal single units, in conjunction with hippocampal theta recorded at the fissure allowed a direct comparison of phase preference between theta-modulated units in both structures (Figure 8-3)

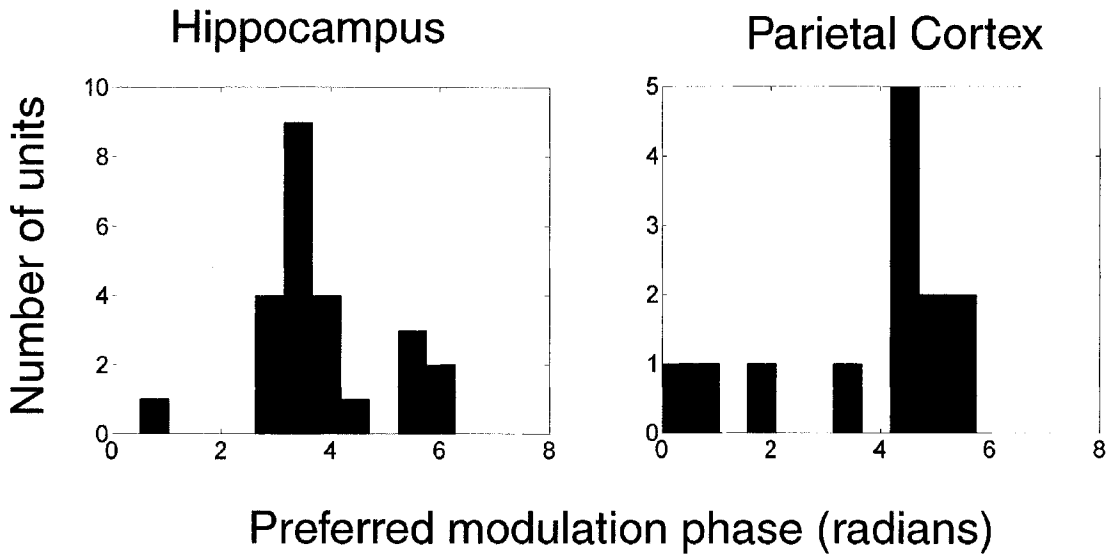

Figure 8-3. Phase preference for theta-modulated cells in both posterior parietal cortex and hippocampus. For each unit showing significant theta modulation, the preferred phase was found by computing a circular mean (Batschelet, 1981). To allow comparison across recording sessions, the nominal zero phase was established by finding the phase associated with maximum hippocampal pyramidal cell activity (Skaggs et al., 1996), and the phase of each unit adjusted accordingly. For rat 1, hippocampal units recorded simultaneously were used to adjust phase. For rat 2, phase was adjusted using hippocampal units recorded two weeks after parietal recordings were made.

The preferred phase theta-modulated parietal units showed a significant directionality (i.e., phase preference) for both rat 1 and rat 2 (Watson-Williams test, $p<.05$). The phase preference of the combined unit populations between rats 1 and 2, however, was not significant (see Discussion for a possible reason). For parietal and hippocampal recordings made simultaneously in rat 1, the preferred phase of parietal units lagged that for hippocampal units by approximately 120° (Figure 8-4).

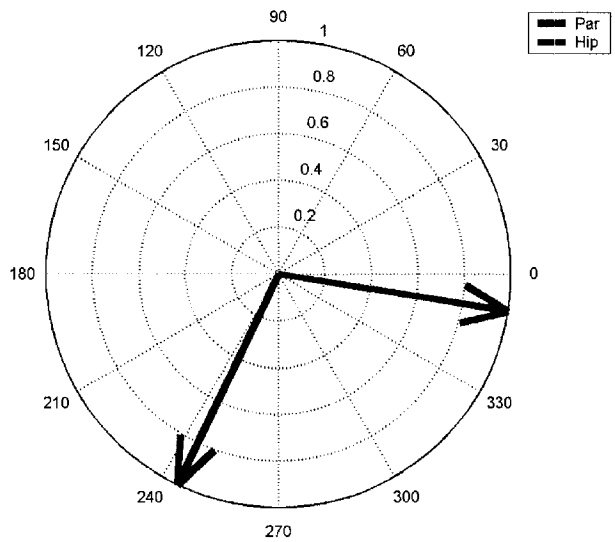

Figure 8-4. The preferred phase of theta-modulated parietal units lags that for hippocampal units by approximately 120°. 18 of 24 parietal units recorded across two sessions in rat 1 were significantly modulated by theta, and showed a significant phase preference (Rayleigh test, $p<.05$). When the preferred phase of theta-modulated parietal and hippocampal units were adjusted to a nominal zero corresponding to the phase of maximum hippocampal activity, the phase preferences were significantly different (Watson-Williams test, $p<.05$).

Discussion

During running to distant targets, the activity of some parietal cells was significantly modulated by theta. As in the case of hippocampal pyramidal cells, parietal cells show a phase preference for theta. The phase preference for parietal units differs, however, from that of simultaneously recorded hippocampal neurons. Though the phase preferences observed in the data from rat 1 and rat 2 were significant in both cases individually, the phase preference of the combined population was not. The lack of a significant phase preference for the combined population might have resulted from the

method used to provide a common "zero phase" across recording sessions. In the data from rats 1 and 5, hippocampal units were recorded simultaneously with parietal units. The theta phase associated with the maximum activation of hippocampal units defined a zero phase for parietal units (Skaggs et al., 1996). In the data from rat 2, no hippocampal data were recorded simultaneously with the parietal data. Hippocampal data used to provide a zero phase were recorded two weeks after the acquisition of the parietal data used in this study. It is possible that a shift in the electrode depth relative to the hippocampal fissure altered the phase relationship of theta to unit recordings, causing a change in the observed phase preference for parietal units recorded from rat 2.

Chapter 9 DISTANCE ENCODING BY PARIETAL NEURONS

Abstract

To interact with distant objects and places, the brain must represent extra-personal space and compute trajectories for the whole body in that space. Gerbils, once given information regarding the distance and direction to a target, can make direct trajectories to those targets, even if the trajectory is completed in darkness. This suggests that trajectories are computed and stored somewhere in the rodent brain prior to the onset of movement. Representations of space and trajectories in that space could be based on vector components (distance and direction). Neural correlates of direction, "head direction cells", have been observed in several structures in the rodent brain, including parietal cortex, but similar representations of distance have not been observed. To determine whether representations of distance exist in the rodent parietal cortex during whole-body trajectories, rats were trained to run to cued targets in an open arena for electrical brain stimulation, while multiple, single-unit activity was recorded. The responses of several units, however, were correlated with the progression of the trajectory; i.e., the firing rate of some units decreased during trajectories, while the firing rate of other units increased during trajectories. These data suggest that parietal cortex may encode target distance in several different reference frames, but too few units were recorded in these experiments to support a distance-encoding hypothesis.

Introduction

The nervous system controls our movements as we reach for objects, direct our gaze, or move towards distant targets. When running to distant targets, the brain must plan a trajectory through space prior to the onset of movement to correctly orient towards the target and to initiate an appropriate acceleration of movement. Such a plan must exist in rodents as they make direct, "ballistic" runs to distant targets based on information available only prior to the onset of movement, but little is known about how trajectories involving the whole body are encoded in the brain.

In non-human primates, a rich literature exists regarding single unit activity involved in the encoding, computation and learning of uninterrupted trajectories involving reaching and gazes to targets within arms reach (for review see (Andersen, 1997;Colby and Goldberg, 1999). During reaching tasks, parietal activity is correlated with target direction in various coordinate systems tied to different body structures (e.g., eye, hand, head). In addition, single unit activity in Area 7 (a part of the Posterior Parietal Cortex, PPC) is correlated with reference frames not attached to the body (i.e., in "allocentric" coordinates) (Snyder et al., 1998).

A literature search showed that evidence is lacking in the non-human primate PPC data regarding the representation of target distance. The most relevant studies include the encoding of semantic distance between numbers (Pinel et al., 2001), the representation of quantity in prefrontal cortex (Nieder et al., 2002), and the finding that right hemisphere lesions of parietal cortex do not disrupt distance estimation during walking with the eyes closed (i.e., path integration) (Philbeck et al., 2000). Interestingly, evidence was found in

human lesion data showing a dissociation between neglect of targets that are near (less than a few meters) and those that are more distant (Halligan and Marshall, 1991;Vuilleumier et al., 1998).

In rodents, lesions of parietal cortex impair whole-body navigation, specifically to goals marked by the spatial arrangement of cues (Kolb et al., 1983). Evidence that landmark distance was stored and recalled was observed when gerbils were trained to search for food reward based on local landmarks (Figure 9-1). During training, food reward was placed at a consistent distance and direction with respect to three landmarks. During probe trials after one or more landmarks had been removed, rats searched at a specific distance from the remaining landmarks (Collett et al., 1986).

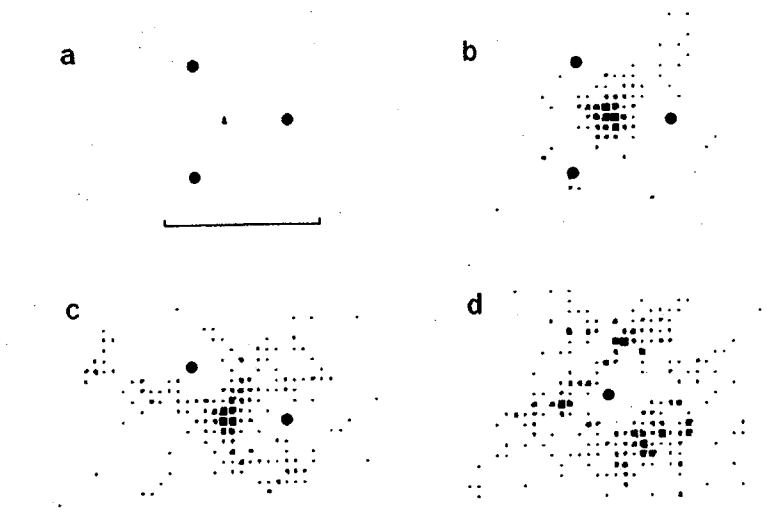

Figure 9-1. Gerbils search for food at a fixed distance from landmarks. (a) Food-deprived gerbils were trained to dig in sawdust for nuts (wedge-shaped object) located in the middle of three landmarks (black circles). (b) During probe trials, the nut was removed, and the gerbil were allowed to search for the food. The density of black dots shows the amount of time the gerbil spent searching in a specific location (a higher density of dots means more time was spent in that place). With all three landmarks present, the gerbil searched primarily in the middle of the three landmarks, where the food had been located in training trials. (c) During a probe trial with one landmark removed, the gerbil searched at a fixed distance from both remaining landmarks. (d) When a second landmark was removed, the gerbil searched at a fixed distance around the sole, remaining landmark. (Figure from (Collett et al., 1986)

Physiological evidence of distance encoding has been observed in the rat hippocampus using an enclosure with moveable walls (O'Keefe and Burgess, 1996). Rats were allowed to explore an arena equipped with moveable walls for several days, while multiple single hippocampal units were recorded. After several sessions, the walls of the enclosure were moved farther apart, increasing the search area available to the rat. Hippocampal place fields observed in the original environment were observed to stretch and, in some cases, split as the distance between the walls increased. The location of some place fields with respect to the distance to the walls remained fixed. While single unit activity relating to distance has not been observed in rodent parietal cortex, activity correlated with head direction has been observed (Chen et al., 1994b), as has activity modulated by a combination of place and behavior during running on radial eight arm maze (Chen, 1989). Little is known, however, about the response of single units in parietal cortex during uninterrupted trajectories of the whole body in an open arena.

Vector encoding of target locations offers an efficient and biologically plausible mechanism to explain how trajectories between start and goal locations could be computed (McNaughton et al., 1994a). If the goal is encoded as a vector recalled from memory, M, along with the current location percept, P, then the trajectory vector can be found by vector subtraction (Figure 9-2), $T = P - M$.

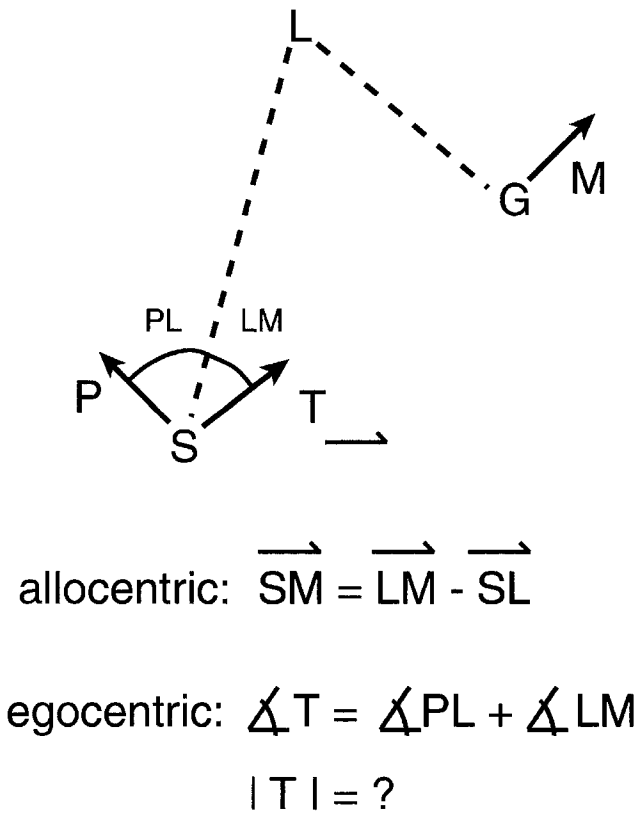

Figure 9-2. Trajectory computation in world-centered (allocentric) and body-centered (egocentric coordinates) using vector representations of current and goal locations. If the starting location, S, is encoded as a vector with respect to some distant landmark, L, and the goal, G, is also vector-encoded with respect to that landmark, then the length and angle of the trajectory can be computed by vector subtraction. This vector, however, must still be converted to an egocentric frame of reference to allow the body to move. The distance between the current and goal locations and the landmark could be learned over time. Alternatively, if no memory of the goal with respect to landmarks existed, the current sensory perception, P, could be used to compute the angle through which the body would need to turn to orient on the target. Prior to movement, however, it is unclear how the distance to the target might be estimated. Given the height of the head above a flat surface, distance could be estimated by the height of the target on the retina. Once movement towards the target begins, distance could be estimated by the motion of the target on the retina (the height of things on the retina will drop more quickly).

Vector subtraction could be accomplished biologically in several ways including through linear associative maps (McNaughton et al., 1994a), or through reciprocally-connected radial basis function units (Deneve et al., 2002). Linear associative maps of sensory space have been observed in the tectum of pit vipers (Hartline et al., 1978), in the

inferior colliculus of owls (Knudsen and Konishi, 1978), and in the superior colliculus of monkeys (Mays and Sparks, 1980). Models of spatial transformations in parietal cortex show how linear associative maps can combine information encoded in different reference frames, allowing vector subtraction (Zipser and Andersen, 1988;Salinas and Abbott, 1995), but do not address the representation of distance during uninterrupted trajectories. To address, multiple single unit recordings were made from the posterior parietal cortex of rats making uninterrupted trajectories to distant targets.

Methods

Behavior

The behavior and motivation level of each rat varied during a recording session. Grooming sessions occurred during each session, and produced individual trials that lasted for more than a minute. To ensure that data used for analysis came only from runs during which the rat was highly motivated to find the goal (i.e., "on-task"), a set of criteria was established to remove trials containing other behaviors (i.e., "off-task"). Off-task runs composed less than 20% of the total number of runs across rats (rat 1: 77 of 617 runs; rat 2: 42 of 127 runs; rat 4: 132 of 684 runs). On-task runs satisfied three criteria: spatial accuracy, directed behavior, and temporal constraint. Spatial accuracy was judged by noting entry into zones other than the correct target zone. If the correct zone was not found initially, rats often searched surrounding edges of the maze, entering multiple wrong zones. In several cases, however, rats ran through a neighboring zone en route to the correct zone. A run passed the spatial accuracy criterion if the rat entered at most one

zone prior to the correct zone. The second constraint, directed behavior, identified those runs where the rat stopped during the run, often in the middle of the maze. Such stops appeared to indicate indecision regarding the correct zone, and were accompanied by multiple, shorter trajectories during a given run. Stops were identified by a drop in velocity below an average of 0.2 m/sec (max velocity during runs was around 0.6 m/sec), and runs with more than one such stop were considered to be invalid runs. Finally, on some runs, rats paused prior to the start of movement as if temporarily unmotivated to run to the goal. To ensure that analyses were conducted only on data for directed runs to the correct target, runs were considered valid only if the target zone was reached within five seconds of the tone onset. A characteristic valid run (Figure 9-3) consisted of the following sequence of events: Cue (tone/light) onset, trajectory Onset, peak Acceleration, peak Velocity, and entry into the correct zone, followed immediately by MFB Stimulation.

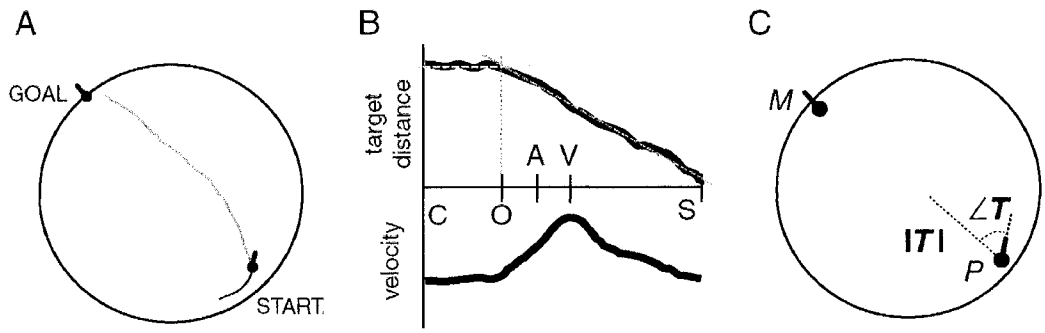

Figure 9-3. (LEFT) A characteristic "valid" run. (red) Path from cue onset to the start of ballistic movement towards the goal. (green) ballistic movement towards the goal. Head direction at the start and at the goal is shown by the direction of the flag. (MIDDLE) Each valid run consisted of a series of events which could be identified either from computer time stamps or from the video recording: (C)ue onset, trajectory (O)nset, peal (A)cceleration, peak (V)elocity, then (S)timulation upon entry of the correct zone. (RIGHT) Moving the body requires that trajectory coordinates (T, stored in terms of body-centered distance and direction) be computed from allocentric, world-centered coordinates (the present location in the environment, P, and the memory of the goal location, M, recalled from previous visits to that spot).

Raster plots

Raster plots were generated for both the actual and relative distance to the target from the onset of the trajectory. Actual distance to the target for each video frame was computed by finding the Euclidean distance between the rat's current position and the position of the rat when the target zone was entered. Relative distance to the target was computed by normalizing actual distances within each run; i.e., each run started with a relative distance of "1.0" and ended with a relative distance of "0.0". Histograms were generated by binning these data into 10 equally spaced bins. Raster plots as a function of time, rather than distance, were also generated for spikes occurring within a window extending from 500 msec prior to the onset of movement to 500 msec following the onset of movement.

Rate maps

Rate maps were constructed by dividing the maze into equally-spaced bins, and then dividing the number of spikes occurring in a given bin by the amount of time spent by the rat in that bin (Muller et al., 1987); (Wilson and McNaughton, 1993).

Task-stage related spike counts

Each trajectory was divided into three stages: cue onset to trajectory onset (CO), movement onset to peak velocity (OV), and peak velocity to stimulation (VS). The first stage, CO, included the time period when information regarding the correct reward zone was available, but the rat has not initiated movement towards that zone. The moment of peak velocity (V) divided the trajectory itself into early and late parts: trajectory onset to

peak velocity, OV, included the start of the trajectory; peak velocity to stimulation, VS, included the end of the trajectory. Mean firing rate for each stage was computed by dividing the number of spikes that occurred in that stage by the time duration of the stage. To allow comparison across units with different baseline firing rates, the firing rate of a given unit in all three stages was normalized by the firing rate in the first stage, CO.

Results

53 units were recorded over four sessions in three rats (30, 12 and 11, respectively). For each unit, spikes occurring during a trajectory (occurring after trajectory onset and before arrival at the goal) were selected for analysis. The selected spikes were then sorted according to whether they occurred while the rat was far from or near to the goal; i.e., according to the relative distance along the trajectory. Spikes occurring in the first ("early") and last ("late") third of the trajectory were counted and a ratio between the computed. Units firing 1.5 times as many spikes early in the trajectory (when the rat was far from the goal) were called "far" cells. Units firing 1.5 times as many spikes late in the trajectory (when the rat was near the goal) were called "near" cells. Of the 53 units analyzed, 4 units were identified as "near" (Figure 9-4), and 5 units were identified as "far" cells (Figure 9-5).

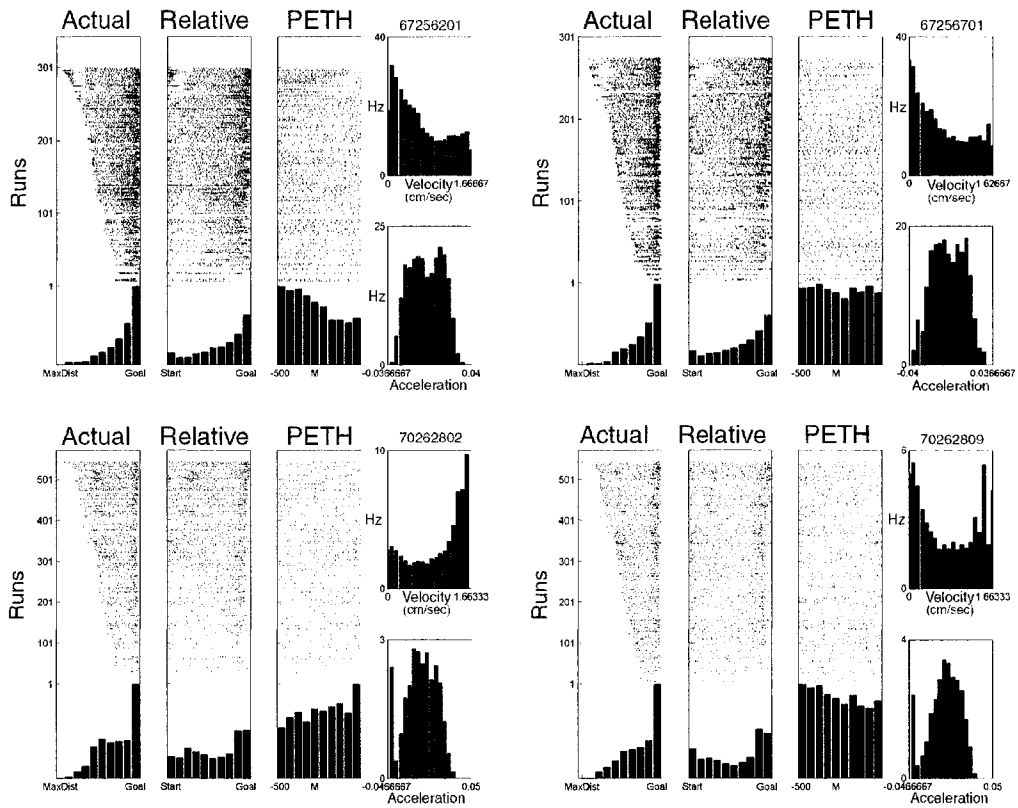

Figure 9-4. Raster plots and tuning curves for all "near" cells. Each panel shows the spikes from a single unit during all trajectories in a single recording session plotted as a function of five different variables. Raster plots are sorted according to the total distance of each trajectory, the longest runs being plotted at the top. The following descriptions apply to the data for each unit. (ACTUAL) Raster and histogram plotted as a function of the Euclidean distance to the goal. The maximum distance (1.3 m) is at the left, and the goal (0 m) is at the right. (RELATIVE) Raster and histogram plotted as a function of the normalized distance to the goal. The start of each trajectory is aligned to the left side, and the end of each trajectory is aligned to the right side. (PETH) Peri-event time histogram centered on the onset of movement, "M", towards the goal. The window extends from 500 msec before to 500 msec following the onset of the trajectory. (Velocity) The spikes occurring on all trajectories were binned as a function of velocity and counted (i.e., a velocity tuning curve). (Acceleration) The spikes occurring on all trajectories were binned as a function of acceleration and counted (i.e., an acceleration tuning curve).

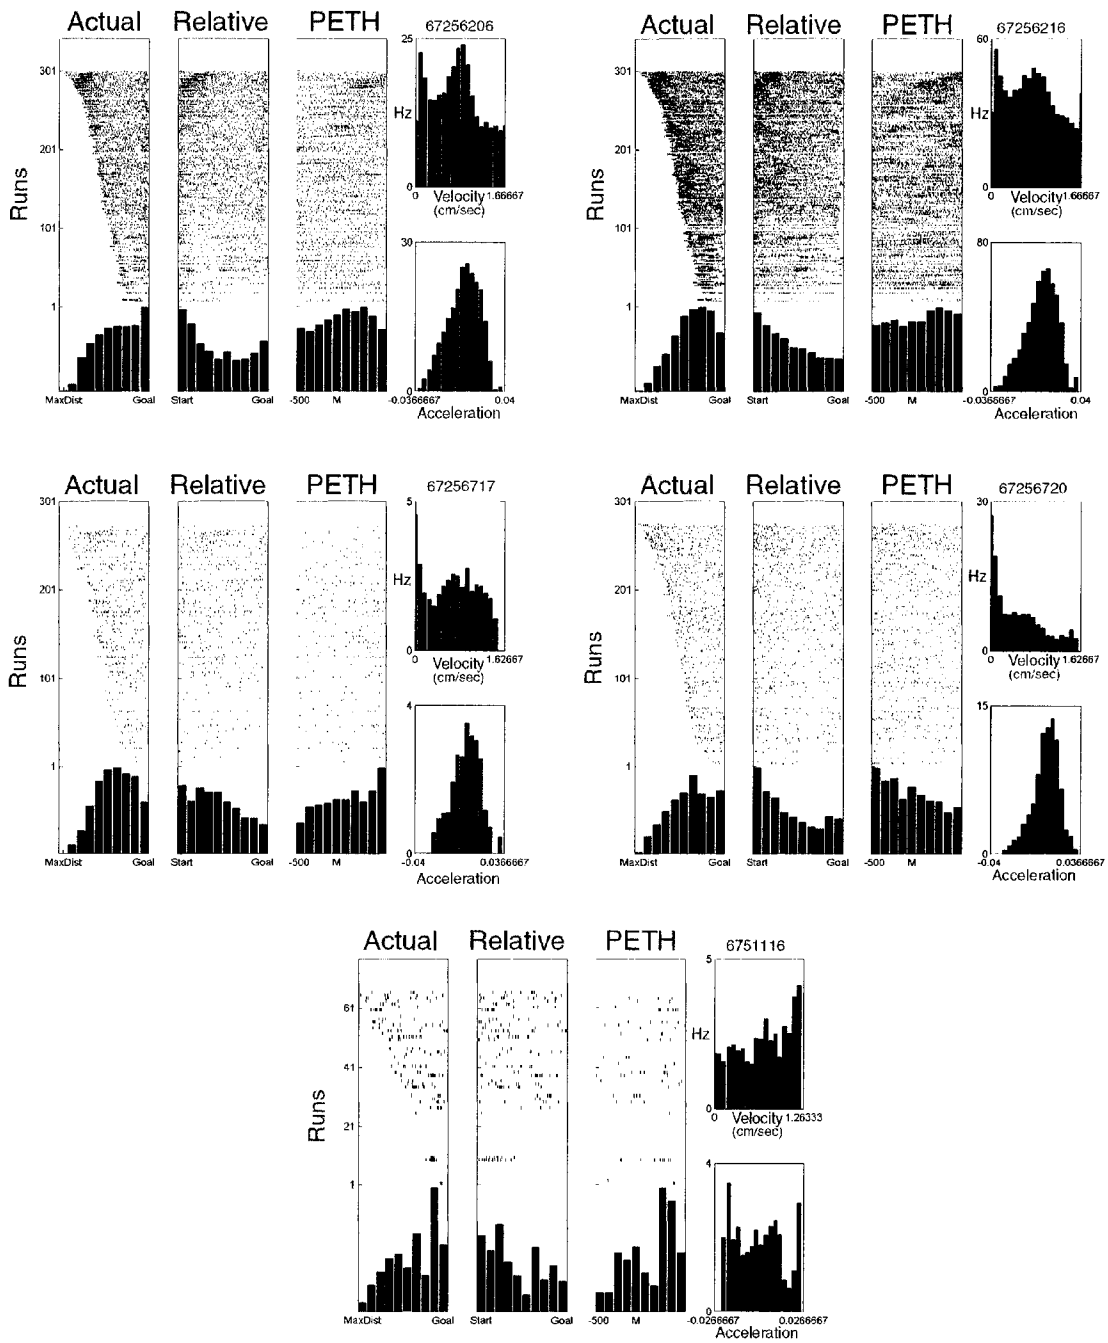

Figure 9-5. Raster plots and tuning curves for all "far" cells. As in Figure 9-4, each panel shows the spikes from a single unit during all trajectories in a single recording session plotted as a function of the actual and relative distance to the goal, and as a peri-event time histogram centered on the onset of the trajectory. (Velocity) The spikes occurring on all trajectories were binned as a function of velocity and counted (i.e., a velocity tuning curve). (Acceleration) The spikes occurring on all trajectories were binned as a function of acceleration and counted (i.e., an acceleration tuning curve).

Distribution of relative firing rates relative to task stage

If "near" and "far" cells encode the relative completion of a trajectory, then the firing rates of each class should peak at different stages of trajectories. That is, most of the spikes from "far" cells should occur early in the trajectory, between the time of movement onset and that of peak velocity (OV), while few of the spikes of "near" cells should occur during this period. Conversely, most "near cell" spikes should occur from the time of peak velocity to arrival at the goal and receipt of stimulation (VS). Normalizing within each trajectory by the total number of spikes (CS), and comparing to the relative amount of time spent in each stage shows how population activity shifts from "far" cells early in a trajectory to "near" cells later in the trajectory (Figure 9-6). Using a t-test, the relative firing rate of "far cells" during the OV stage was significantly different from the CO stage ($p<.001$, $t=3.7$, $df=513$). The relative firing rate of all "near cells" during the VS stage was significantly different from the CO stage ($p<.001$, $t=9.31$, $df=520$).

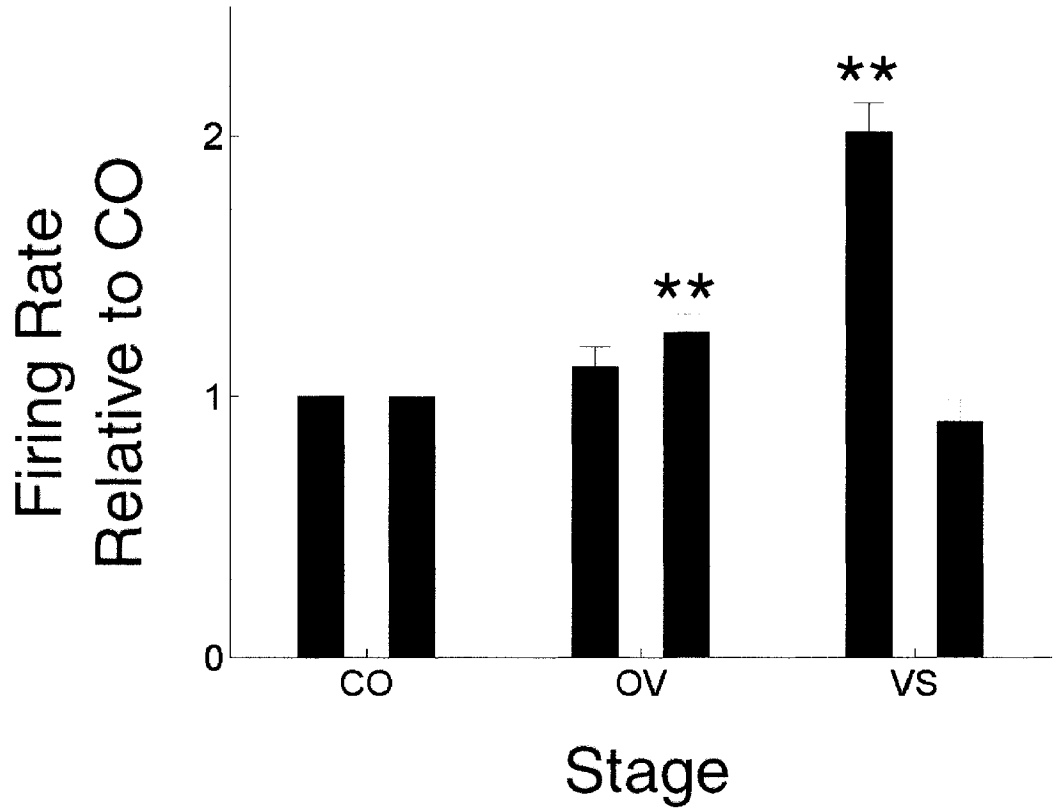

Figure 9-6. The relative firing rates for the two classes of units diverge during trajectories. Trajectories were divided into three bins: CO, from cue onset to trajectory onset; "OV", from trajectory onset to peak velocity; and "VS", from peak velocity to stimulation at the reward site. The mean firing rate for each unit was computed at each stage, and then normalized to the mean firing rate during CO. Cells were grouped as either "near cells" (black) or "far cells" (red). The mean firing rate and standard error of the mean were computed according to cell group.

Relation to cortical depth

Because the units analyzed in this chapter were recorded from tetrodes that were subsequently moved into the hippocampus, histological verification of tetrode placement in parietal cortex was not possible. Two indirect methods were used to eliminate the possibility that the units analyzed in this section were not actually located in the hippocampus. The first method used the nominal tetrode depth at which the putative parietal units were recorded, and the additional depth moved by the tetrode before

encountering hippocampal complex-bursting cells; i.e., before moving into the CA1 cell body layer.

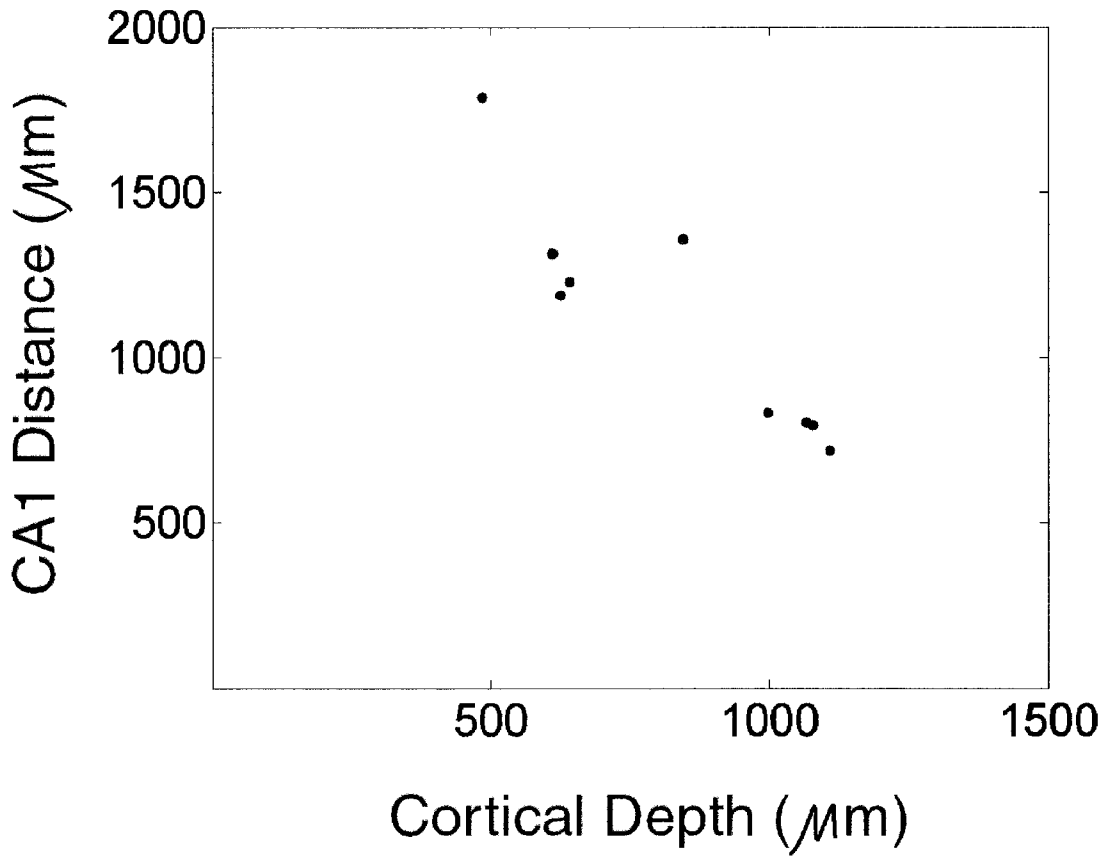

Figure 9-7. Cortical depth and additional depth to CA1 for both "near" and "far" cells. The horizontal axis shows the depth of the recording electrode from the brain surface. The vertical axis shows the additional depth moved by the recording electrode before being moved into the cell body layer of CA1. Blue dots represent "near" cells. Red dots represent "far" cells. Both coordinates for each unit were dithered by ± 50 μm to separate overlapping points.

Both "near" and "far" cell types were found roughly 500-1000 μm below the cortical surface, nominally cortical layers 4 through 6 (Paxinos and Watson, 1997). Units encountered at more shallow cortical depths were further from the CA1 cell body layer, and each unit was at 500 μm from the CA1 cell body layer, nominally the distance through the Corpus Callosum and Alveus.

Rate maps

The place specificity of both hippocampal pyramidal cells and interneurons can be seen using rate maps, as was shown by several examples in the earlier chapters describing hippocampal units in this dissertation. If the putative parietal units analyzed in this chapter were indeed hippocampal units, or the projecting axons of hippocampal units (as could be the case for units recorded in the Alveus), then their spatial rate maps should bear some resemblance to hippocampal rate maps. This was not the case (Figure 9-8).

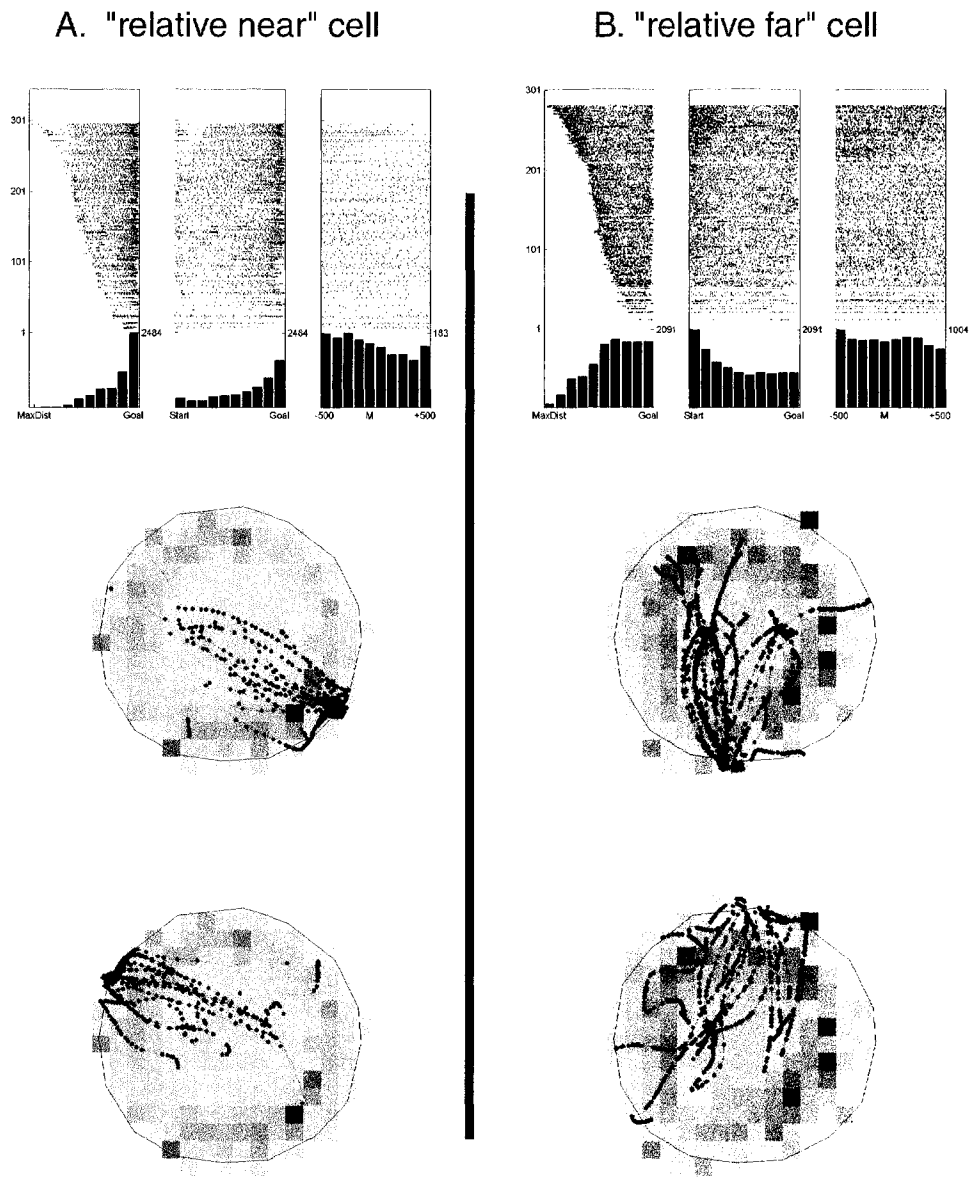

Figure 9-8. The rate maps of "near" and "far" cells show no place specificity, as would be expected of hippocampal units. Looking down onto the circular maze from above, rate maps (see Methods) for two units (A. "relative near" cell, B. "relative far" cell) are shown in gray scale. All trajectories to a single target are shown as red lines. In each case, trajectories to only two (of 16) targets on opposite sides of the maze are shown for clarity; no significant differences were observed during runs to other targets. The location of the rat at times when spikes occurred during a run are shown as blue dots. Note that the rate maps for both units appear as annuli, even though the distance-related responses are quite different.

The rate maps of both "near" and "far" cells peak for annuli around the center of the circular arena. This would be expected for "near" cells because they fired most near the edge when rat approached the edge-mounted target. Annular-peaked rate maps would also be expected for "far" cells because they fired most near the edge when the rat was on the opposite side of the maze from the target.

Discussion

In the rat, the distance to the target of a uninterrupted, whole-body trajectory is encoded in the activity of parietal units. Tuning curves for some units show that some units increase their firing rate as the rat nears the target, while other units decrease their firing rate, and these changes are independent of the physical length of the trajectory. In addition, other units generate a maximal response when the rat is a specific distance from the goal of the trajectory, independent of the stage of the task (i.e., whether at the start or the end of the trajectory). The firing rate of these different classes of units contains sufficient information to accurately determine the stage of the current trajectory (i.e., to differentiate the early stages of a trajectory from the later stages), and to estimate the true distance remaining to the target.

The presence of distance encoding units in PPC, along with the presence of head direction units (Chen et al., 1994a), suggests that vector representations of trajectories might be computed or stored there. Such a result would agree with similar findings from single-unit recordings in non-human primates during reaching and saccade tasks. The presence of different classes of distance-selective units (both "relative" and "absolute")

would also agree with findings in non-human primate recordings regarding the presence of multiple reference frames.

Alternatively, what we have here called distance-selective units could be visual-responsive cells that are responding to the targets themselves, either the physical structure or through response to the light itself. This does not appear to be the case, however, because the response of these units is not modulated by 1 Hz in synchrony to the light flashes. In addition, the possibility that these units are responding to the physical structure of the cues does not explain the distant-dependent response of these three classes of units.

These units could be encoding acceleration and velocity during the trajectories; the tuning curves for "far" cells peak at low valued, positive accelerations, while those of "near" cells peak at low values of velocity. This alternative hypothesis is hard to support for "fixed" cells due to the lack of firing at the start of short trajectories. The acceleration tuning curves of "far" cells look very similar to those of both "fixed" and "constant rate" neurons, suggesting that if "far" cells encode acceleration, then so do a large number of parietal neurons. In addition, the velocity and acceleration tuning curves are quite broad, suggesting that any tuning would be at best inexact. Finally, the tuning curves observed would require that these units responded maximally to a fixed velocity or acceleration, rather than being simply modulated by velocity and acceleration. A simpler interpretation is that "far" neurons, as well as "fixed" neurons are indeed distance-selective, and that the increased activity due to distance tends to occur when the rat is far from the target and accelerating. Likewise, "near" cells are indeed distance-selective, but the increased

response occurs at low velocities is due to the rat slowing down as it nears the target. According to this interpretation, tuned responses to velocity and acceleration are secondarily associated with the primary, distance-related correlates.

The experimental methods described here could show that vector coding of distance and direction could allow prediction of intended targets if more units could be recorded simultaneously, such as the use of higher density recording methods. Future recordings might also show that different, anatomical regions of rodent PPC represent distance information according to different reference frames, as observed in non-human primate recordings. This could explain the inter-rat variability observed in these experiments by suggesting that slight variations in the stereotaxic placement could lead to significant differences in the reference frame observed.

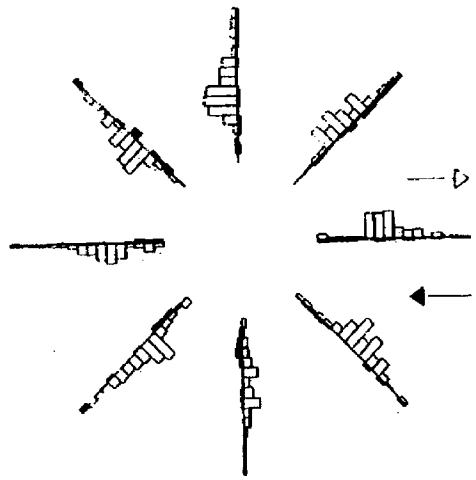

Figure 9-9. An example of parietal cell activity on the radial eight-arm maze. The spokes emanating from the center are the arms of the maze. Histograms on each arm show the binned mean firing rate of the cell on inbound (filled rectangles) and outbound (empty rectangles) journeys. One correlate of the activity is the distance to the end of the current trajectory, regardless of the arm, where this cell fires preferentially at the middle of the trajectory. (From (McNaughton et al., 1994b)).

Another possibility involves the observation of parietal cells that responded to the conjunction of behavior and location (Chen, 1989). Some parietal cells were observed to fire in the middle of each arm of a radial eight-arm maze, but only on the outward journey. The response of these units could be interpreted as being distance selective for the end of the arm. The response was independent of the heading of the rat, the current arm of the maze, or the previous arm that had been visited. The encoding of distance cannot explain all of the results regarding parietal unit responses observed by Chen. Cells that fired selectively during either right or left turns at the ends of each arm, for example, would not be explained by a distance hypothesis.

Chapter 10 DISCUSSION

Review of the main question

How does the brain compute trajectories to distant goals? The experiments described in this dissertation addressed several issues related to this central question, but not all; it is a broad question. In the experiments described in this dissertation, rats were trained to either locate distant beacons or recall the location of distant targets, and then run to those targets. While such a task might appear to be simple, it should be noted that rats required hundreds of trials to learn a sequence of locations, whereas rats can learn the Morris water maze task in five to ten trials (C.A. Barnes, personal communication). The results described in this dissertation resulted from the application of different analytical tools applied to data acquired from rats making uninterrupted trajectories. Any attempt to describe how the brain computes trajectories should recognize that each of the results described as occurring either in the hippocampus or parietal cortex, or in the activity of single units or in the combined activity of neuronal populations as described by EEG signals are occurring simultaneously and within a time scale of as little as fractions of a second.

Hippocampus during a trajectory

During trajectories to a cued goal, no evidence was found that neural activity in the hippocampus was related to the location of the intended goal. Hippocampal place cells near intended goals did not activate reliably around the time of trajectory onset. In

addition, no activity similar to that of non-human primate "view cells" was observed in the rodent hippocampus. In addition, no evidence was found to support a shift in reference frames between one centered on the current location of the rat (as expressed by place cells) and a frame centered on the goal. During trajectories to cued, distant landmarks, single unit activity in the rodent hippocampus appeared to be predominantly related to the current location of the rat.

Evidence was obtained, however, that the phase of hippocampal theta was reset at the time of peak acceleration. No evidence was observed to suggest that hippocampal activity was related directly to movement; i.e., the hippocampus is no more a "motor structure" than it is a "sensory structure". Just as hippocampal activity can be influenced by sensory stimuli, the phase resetting of theta to movement-related variables agrees with previous results suggesting that the hippocampus can also be influenced by motor activity (Vanderwolf and Heron, 1964;McNaughton et al., 1983a).

Hippocampus during sequences

One surprising result from these experiments was the observation that rats can learn to run to sequences of multiple locations, even when those sequences contained repeated elements. Rats learn these sequences quickly (within two half-hour recording sessions). The observation that rotations of these sequences disrupted performance as significantly as the presentation of an entirely new sequence suggested that at least some part of the behavioral strategy used to solve these sequences was spatial. This suggests that rats with hippocampal lesions would have difficulties learning to run to sequences of

locations. If rats were guided during their training, however, such as in Barrier-trained task, it seems reasonable that rats with hippocampal lesions could learn to run to sequences. The lack of correlation between the separability of hippocampal activity on repeated elements across all repeated-element tasks poses a direct challenge to any theory that ties the ability of rats to discriminate overlapping sequential contexts specifically to the hippocampus.

Parietal cortex during a trajectory

These experiments did not provide as many answers regarding the response of parietal neurons during trajectories as for hippocampal neurons. This resulted primarily from the use of recording technology that was not suited ideally for recording cortical units. Evidence was obtained, however, relating single unit activity in the parietal cortex to hippocampal theta, and suggests that some parietal units respond selectively to the early and late stages of trajectories. The former result is significant in that it is the first evidence for phase locking between hippocampal theta and units located in a brain structure that is not part of the limbic system. This suggests that in addition to the many roles that have been proposed for hippocampal theta, one should the role of synchronizing activity within structures that communicate directly with the hippocampus with more distal structures.

Future directions

The results described in this dissertation raise a number of questions for future work both in the parietal cortex and hippocampus. Of interest regarding the parietal

cortex is the relationship between parietal activity and hippocampal theta. While the degree of modulation of parietal activity is not as great as that for hippocampal activity, the influence of hippocampal theta on parietal activity during waking behaviors compliments the story of hippocampal-neocortical interaction during sleep. A significant question is whether theta modulation of parietal units serves to encode any information. The presence of distance selective units in parietal cortex suggests that computations similar to those that have been observed in monkey parietal cortex (namely, the presence of gain fields) might also occur in rat parietal cortex. The study of gain fields in rodents would require improvements in the measurement of head direction, because current, video-based techniques cannot capture the fine movements of the head in three dimensions with temporal precision.

Future work regarding the hippocampus includes several questions regarding how rats learn sequences of locations. In the experiments described in this dissertation, the longest sequence learned by rats involved only eight elements. How many elements could a rat learn? How complex might the repeated segments be? For instance, could a rat learn to run between two locations a certain number of times before moving on to another sequence (i.e., ABCBCDA)? In regards to the cause of differential hippocampal activity along the repeated segment during the Journey task, did this arise because of memories of different starting locations, or because of working memory in extra-hippocampal structures related to different goals? One way of addressing this question would be to repeat the experiment with probabilistic reward given at either the entrance or the exit of the repeated segment. If the memory of different starting locations is

significant in producing differential activity, then consistent reward at the entrance of the repeated segment should disrupt differential activity. If the presence of different goals is significant, then consistent reward at the exit of repeated segment should disrupt differential activity.

APPENDIX A

RECORDING CHECKLIST

I. Prepare the acquisition computer

Turn on purple box
Double-click on and start Cheetah
Click to "File", click "Open Config File"
Go to D:\Cheetah Data\7261
Double click on turnning.cfg file
Click on "Acquisition" menu at top,
then "Change Data Directory"
Click the "New Folder" icon
Make a new directory for today's run,
click on it, then hit "Save" as
"rat#_sess#", e.g. "7026_04"
Start acquisition by clicking "ACQ"

II. Prepare the stimulator

Turn on the Stimulus Isolation Unit (SIU)
Connect SIU to the gray box above it;
(rat #7261) red to S1, black to S2
Make sure the output control is set to "200"
Make sure the light is "on" for S2

III. Get the rodent

Plug in the rat; turn off the hot lights
Turn on the headstage: on, off, on

IV. Prepare Cheetah for recording

Set gains and thresholds for each TT
Go up one directory and save the
settings in E:\Lights_Tones\rat#
under "rat#_sess#.cfg"
Click on "video tracker" and make sure
the rat is being tracked w/o noise.
All except Pure Blue
Close the video tracker window
Minimize Cheetah
Double-click on Bower SeqControl" executable
Set "Zone Order" to "Random"
Set "Stim Dur" to "400"
Click on "Apply Settings"
Re-expand the Cheetah window

V. Begin recording

Turn on "REC"
In the "User Event Window" (UEW),
check "Single Keystroke"

Type "a" in the text window; hit "Del"
Write the clock time into the lab book.
Write down bad channels on all tetrodes
SLEEP 1 (20-30 minutes)
When done, write clock time in lab book
Type "b" in the UEW text window, the "Del"

VI. Running the rat

Set the SIU output to "ON"
Check that polarity remains on S2/black
Turn off the audio monitor
Put the rat on the maze
Speakers on floor should be unblocked
Type "a" in the UEW text window;
then clear w/ "Del"
Note the clock time in the lab notebook
Minimize Cheetah window
Click "Start" button in exp control;
you should hear a tone and see tracking
Help rat start running if he needs it
Move mouse every 100 rewards, or so.
MAZE
When rat has completed the days task,
hit "Stop" in "Position "window
Write the number of rewards in lab book
then exit the training program
Turn SIU output to "OFF"
Re-expand Cheetah window
Type "b" in UEW text window,
then clear w/ "Del"
Turn off the SIU
Put rat back into home dish

VII. Finishing the session

Type "a" in UEW window, and clear it
Write clock time in lab book
SLEEP 2
Type "b" in UEW window
Write clock time in lab book
Click "REC" to turn off recording
Click "ACQ" to stop acquisition
Go to "File", and click "Exit"
Turn off headstage power
Unplug rat and return him to his cage,
or restart Cheetah to turn tetrodes

VIII. Backup the data

Go into newly-created directory (DATA_DIR)
and zip VT1.dat by right-clicking on it then clicking on "Add to VT1.zip"
Copy data to network storage by dragging DATA_DIR to Seth\Lights

REFERENCES

1. Agster KL, Fortin NJ, Eichenbaum H (2002) The hippocampus and disambiguation of overlapping sequences. J Neurosci 22: 5760-5768.

2. Andersen RA (1995) Encoding of intention and spatial location in the posterior parietal cortex. Cereb Cortex 5: 457-469.

3. Andersen RA (1997) Multimodal integration for the representation of space in the posterior parietal cortex. Philos Trans R Soc Lond B Biol Sci 352: 1421-1428.

4. Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. Science 230: 456-458.

5. Andersen RA, Mountcastle VB (1983) The influence of the angle of gaze upon the excitability of the light- sensitive neurons of the posterior parietal cortex. J Neurosci 3: 532-548.

6. Assad JA, Maunsell JH (1995) Neuronal correlates of inferred motion in primate posterior parietal cortex. Nature 373: 518-521.

7. Ballard DH (1986) Cortical connections and parallel processing: Structure and function. The Behavioral and Brain Sciences 9: 67-120.

8. Barnes CA, Jung MW, McNaughton BL, Korol DL, Andreasson K, Worley PF (1994) LTP saturation and spatial learning disruption: effects of task variables and saturation levels. J Neurosci 14: 5793-5806.

9. Batschelet E (1981) Circular Statistics in Biology. London, England: Academic Press.

10. Battaglia, F. P., Sutherland, G. R., and McNaughton, B. L. Widespread modulation of neocortical cell activity during hippocampal sharp waves. Society for Neuroscience Abstracts . 2001.
Ref Type: Abstract

11. Bisiach E, Luzzatti C, Perani D (1979) Unilateral neglect, representational schema and consciousness. Brain 102: 609-618.

12. Bland BH (1986) The physiology and pharmacology of hippocampal formation theta rhythms. Prog Neurobiol 26: 1-54.

13. Bliss TV, Lomo T (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. J Physiol (Lond) 232: 331-356.

14. Blum KI, Abbott LF (1996) A model of spatial map formation in the hippocampus of the rat. Neural Comput 8: 85-93.

15. Boisacq-Schepens N, Hanus M (1972) Motor cortex vestibular responses in the chloralosed cats. Exp Brain Res 14: 539-549.

16. Bostock E, Muller RU, Kubie JL (1991) Experience-dependent modifications of hippocampal place cell firing. Hippocampus 1: 193-205.

17. Bower, M. R., Ellsworth, E. G., Insel, N., and McNaughton, B. L. Going there vs. being there: Do hippocampus or parietal cortex activate representations of remote goals? Society for Neuroscience Abstracts . 2000.
Ref Type: Abstract

18. Bower, M. R., Euston, D. R., Gebara, N., and McNaughton, B. L. The role of the hippocampus in disambiguating context in a sequence task. Society for Neuroscience Abstracts . 2001.
Ref Type: Abstract

19. Buneo CA, Jarvis MR, Batista AP, Andersen RA (2002) Direct visuomotor transformations for reaching. Nature 416: 632-636.

20. Buno W, Velluti JC (1977) Relationships of hippocampal theta cycles with bar pressing during self-stimulation. Physiology and Behavior 19: 615-621.

21. Burnod Y, Baraduc P, Battaglia-Mayer A, Guigon E, Koechlin E, Ferraina S, Lacquaniti F, Caminiti R (1999) Parieto-frontal coding of reaching: an integrated framework. Exp Brain Res 129: 325-346.

22. Burwell RD, Amaral DG (1998) Cortical afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. J Comp Neurol 398: 179-205.

23. Burwell RD, Witter MP, Amaral DG (1995) Perirhinal and postrhinal cortices of the rat: a review of the neuroanatomical literature and comparison with findings from the monkey brain [published erratum appears in Hippocampus 1996;6(3):340]. Hippocampus 5: 390-408.

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

25. Buzsaki G, Grastyan E, Tveritskaya IN, Czopf J (1979) Hippocampal evoked potentials and EEG changes during classical conditioning in the rat. Electroencephalography and Clinical Neurophysiology 47: 64-74.

26. Buzsaki G, Leung LS, Vanderwolf CH (1983) Cellular bases of hippocampal EEG in the behaving rat. Brain Res 6: 139-171.

27. Caminiti R, Ferraina S, Johnson PB (1996) The sources of visual information to the primate frontal lobe: a novel role for the superior parietal lobule. Cereb Cortex 6: 319-328.

28. Chen, L. L. Spatial representation in the rat Posterior Parietal Cortex. 1989.
Ref Type: Thesis/Dissertation

29. Chen LL, Lin LH, Barnes CA, McNaughton BL (1994b) Head-direction cells in the rat posterior cortex. II. Contributions of visual and ideothetic information to the directional firing. Exp Brain Res 101: 24-34.

30. Chen LL, Lin LH, Green EJ, Barnes CA, McNaughton BL (1994a) Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. Exp Brain Res 101: 8-23.

31. Clower DM, Hoffman JM, Votaw JR, Faber TL, Woods RP, Alexander GE (1996) Role of posterior parietal cortex in the recalibration of visually guided reaching. Nature 383: 618-621.

32. Colby CL, Duhamel JR, Goldberg ME (1993) Ventral intraparietal area of the macaque: anatomic location and visual response properties. J Neurophysiol 69: 902-914.

33. Colby CL, Goldberg ME (1999) Space and attention in parietal cortex. Annu Rev Neurosci 22: 319-349.

34. Collett TS, Cartwright BA, Smith BA (1986) Landmark learning and visuo-spatial memories in gerbils. J Comp Physiol [A] 158: 835-851.

35. Corkin S (2002) What's new with amnesic patient H.M.? Nature Reviews Neuroscience 3: 153-160.

36. Dalland T (1976) Response perseveration of rats with dorsal hippocampal lesions. Behav Biol 17: 473-484.

37. Dalton A, Black AH (1968) Hippocampal electrical activity during the operant conditioning of movement and refraining from movement. Communications in Behavioral Biology 2: 267-273.

38. Deneve S, Latham PE, Pouget A (2002) Reading population codes: a neural implementation of ideal observers. Nat Neurosci 2: 740-745.

39. DeRenzi E (1982) Disorders of Space Exploration and Cognition. New York: John Wiley & Sons.

40. DiMattia BD, Kesner RP (1988) Spatial cognitive maps: differential role of parietal cortex and hippocampal formation. Behav Neurosci 102: 471-480.

41. Duda RO, Hart PE, Stork, D.G. (2001) Pattern Classification. New York: John Wiley and Sons, Inc.

42. Duhamel JR, Bremmer F, BenHamed S, Graf W (1997) Spatial invariance of visual receptive fields in parietal cortex neurons. Nature 389: 845-848.

43. Eichenbaum H, Dudchenko P, Wood E, Shapiro M, Tanila H (1999) The hippocampus, memory, and place cells: is it spatial memory or a memory space? Neuron 23: 209-226.

44. Fenton AA, Bures J, Cimadevilla JM, Olypher AV, Wesierska M, Zinyuk L (2002) Hippocampal place cell activity during overtly purposeful behavior (in dissociated reference frames). In: The Neural Basis of Navigation (Sharp PE, ed), pp 59-80. Boston, MA.: Kluwer Academic Publishers.

45. Fortin NJ, Agster KL, Eichenbaum H (2002) Critical role of the hippocampus in memory for sequences of events. Nat Neurosci 5: 458-462.

46. Frank LM, Brown EN, Wilson M (2000) Trajectory encoding in the hippocampus and entorhinal cortex. Neuron 27: 169-178.

47. Freund TF, Buzsaki G (1996) Interneurons of the hippocampus. Hippocampus 6: 347-470.

48. Fukushima K (1997) Corticovestibular interactions: anatomy, electrophysiology, and functional considerations. Exp Brain Res 1-16.

49. Georges-Francois P, Rolls ET, Robertson RG (1999) Spatial view cells in the primate hippocampus: allocentric view not head direction or eye position or place. Cereb Cortex 9: 197-212.

50. Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci 2: 1527-1537.

51. Givens B (1996) Stimulus-evoked resetting of the dentate theta rhythm: relation to working memory. Neuroreport 8: 159-163.

52. Goodale MA, Milner AD (1982) Fractionating orientation behavior in rodents. In: Analysis of Visual Behavior (Ingle DJ, Goodale MA, Mansfield RJW, eds), pp 549-586. Boston: MIT Press.

53. Gothard KM, Skaggs WE, McNaughton BL (1996a) Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. J Neurosci 16: 8027-8040.

54. Gothard KM, Skaggs WE, Moore KM, McNaughton BL (1996b) Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. J Neurosci 16: 823-835.

55. Gray CM, Konig P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature 338: 334-337.

56. Guariglia C, Padovani A, Pantano P, Pizzamiglio L (1993) Unilateral neglect restricted to visual imagery [see comments]. Nature 364: 235-237.

57. Halligan PW, Marshall JC (1991) Spatial compression in visual neglect: a case study. Cortex 27: 623-629.

58. Harris EW, Ganong AH, Cotman CW (1984) Long-term potentiation in the hippocampus involves activation of N-methyl-D-aspartate receptors. Brain Res 323: 132-137.

59. Hartline PH, Kass L, Loop MS (1978) Merging of modalities in the optic tectum: infrared and visual integration in rattlesnakes. Science 199: 1225-1229.

60. Hartline PH, Vimal RL, King AJ, Kurylo DD, Northmore DP (1995) Effects of eye position on auditory localization and neural representation of space in superior colliculus of cats. Exp Brain Res 104: 402-408.

61. Hebb DO (1949) The Organization of Behavior. John Wiley and Sons.

62. Hill AJ (1978) First occurrence of hippocampal spatial firing in a new environment. Exp Neurol 62: 282-297.

63. Hoffman KL, McNaughton BL (2002) Coordinated reactivation of distributed memory traces in primate neocortex. Science 297: 2070-2073.

64. Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI (2001) Accumulation of hippocampal place fields at the goal location in an annular watermaze task. J Neurosci 21: 1635-1644.

65. Hyvarinen J (1982) The Parietal Cortex of Monkey and Man. Springer-Verlag.

66. Insausti R, Herrero MT, Witter MP (1997) Entorhinal cortex of the rat: cytoarchitectonic subdivisions and the origin and distribution of cortical efferents. Hippocampus 7: 146-183.

67. Kalaska JF, Crammond DJ (1995) Deciding not to GO: neuronal correlates of response selection in a GO/NOGO task in primate premotor and parietal cortex. Cereb Cortex 5: 410-428.

68. Kesner RP, Gilbert PE, Barua LA (2002) The role of the Hippocampus in memory for the temporal order of a sequence of odors. Behav Neurosci 116: 286-290.

69. Knierim JJ, Kudrimoti HS, McNaughton BL (1995) Place cells, head direction cells, and the learning of landmark stability. J Neurosci 15: 1648-1659.

70. Knudsen EI, Konishi M (1978) A neural map of auditory space in the owl. Science 200: 795-797.

71. Kohonen T (1984) Self-organization and associative memory. Berlin: Springer-Verlag.

72. Kolb B, Sutherland RJ, Whishaw IQ (1983) A comparison of the contributions of the frontal and parietal association cortex to spatial localization in rats. Behav Neurosci 97: 13-27.

73. Kramis R, Vanderwolf CH, Bland BH (1975) Two types of hippocampal rhytmical activity in both the rabbit and the rat: relations to behavior and effects of atropine, diethyl ether, urethane and pentobarbital. Exp Neurol 49: 58-85.

74. Krieg, W. J. S. Connections of the cerebral cortex. I. The albino rat. A topography of the cortical areas. J.Comp Neurol. 84, 221-275. 1946.
Ref Type: Abstract

75. Lacquaniti F, Guigon E, Bianchi L, Ferraina S, Caminiti R (1995) Representing spatial information for limb movement: role of area 5 in the monkey. Cereb Cortex 5: 391-409.

76. Lavenex P, Suzuki WA, Amaral DG (2002) Perirhinal and parahippocampal cortices of the macaque monkey: projections to the neocortex. J Comp Neurol 447: 394-420.

77. Lee AK, Wilson MA (2002) Memory of sequential experience in the hippocampus during slow wave sleep. Neuron 36: 1183-1194.

78. Lenck-Santini PP, Muller RU, Save E, Poucet B (2002) Relationships between place cell firing fields and navigational decisions by rats. J Neurosci 22: 9035-9047.

79. Lenck-Santini PP, Save E, Poucet B (2001) Evidence for a relationship between place-cell spatial firing and spatial memory performance. Hippocampus 11: 377-390.

80. Leung LS, Yim CY (1986) Intracellular records of theta rhythm in hippocampal CA1 cells of the rat. Brain Res 367: 323-327.

81. Leung LW, Bostock E (1987) Electrical activity of the cingulate cortex. I. Generating mechanisms and relations to behavior. Brain Res 407: 68-80.

82. Lever C, Wills T, Cacucci F, Burgess N, O'Keefe J (2002) Long-term plasticity in hippocampal place-cell representation of environmental geometry. Nature 416: 90-94.

83. Levy WB (1996) A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. Hippocampus 6: 579-590.

84. Levy WB, Steward O (1979) Synapses as associative memory elements in the hippocampal formation. Brain Res 175: 233-245.

85. Liebman JM, Cooper SJ (1989) The Neuropharmacological Basis of Reward. Oxford: Clarendon Press.

86. Lisman JE (1999) Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate-CA3 interactions. Neuron 22: 233-242.

87. Lynch JC, Mountcastle VB, Talbot WH, Yin TC (1977) Parietal lobe mechanisms for directed visual attention. J Neurophysiol 40: 362-389.

88. Maguire EA, Burgess N, Donnett JG, Frackowiak RS, Frith CD, O'Keefe J (1998) Knowing where and getting there: a human navigation network [see comments]. Science 280: 921-924.

89. Maguire EA, Frackowiak RSJ, Frith CD (1997) Recalling routes around london: activation of the right hippocampus in taxi drivers. J Neurosci 17: 7103-7110.

90. Malenka RC, Nicoll RA (1999) Long-term potentiation--a decade of progress? Science 285: 1870-1874.

91. Markus EJ, Qin YL, Leonard B, Skaggs WE, McNaughton BL, Barnes CA (1995) Interactions between location and task affect the spatial and directional firing of hippocampal neurons. J Neurosci 15: 7079-7094.

92. Marr D (1970) A theory for cerebral neocortex. Proc R Soc Lond B Biol Sci 176: 161-234.

93. Marr D (1971) Simple memory: a theory for archicortex. Philos Trans R Soc Lond B Biol Sci 262: 23-81.

94. Mays LE, Sparks DL (1980) Saccades are spatially, not retinocentrically, coded. Science 208: 1163-1165.

95. McCarthy RA, Warrington EK (1990) Cognitive Neuropsychology: A Clinical Introduction. Academic Press.

96. McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102: 419-457.

97. McNaughton BL, Barnes CA, Gerrard JL, Gothard K, Jung MW, Knierim JJ, Kudrimoti H, Qin Y, Skaggs WE, Suster M, Weaver KL (1996) Deciphering the hippocampal polyglot: the hippocampus as a path integration system. J Exp Biol 199 (Pt 1): 173-185.

98. McNaughton BL, Barnes CA, O'Keefe J (1983a) The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Exp Brain Res 52: 41-49.

99. McNaughton BL, Barnes CA, Rao G, Baldwin J, Rasmussen M (1986) Long-term enhancement of hippocampal synaptic transmission and the acquisition of spatial information. J Neurosci 6: 563-571.

100. McNaughton BL, Chen LL, Markus EJ (1991) "Dead Reckoning," landmark learning, and the sense of direction: A neurophysiological and computational hypothesis. Journal of Cognitive Neuroscience 3: 190-202.

101. McNaughton BL, Douglas RM, Goddard GV (1978) Synaptic enhancement in fascia dentata: cooperativity among coactive afferents. Brain Res 157: 277-293.

102. McNaughton BL, Knierim JJ, Wilson MA (1994a) Vector-encoding and the vestibular foundations of spatial cognition. In: The Cognitive Neurosciences (Gazzaniga M, ed), pp 585-595. Boston: MIT Press.

103. McNaughton BL, Leonard B, Chen LL (1989) Cortical-hippocampal interactions and cognitive mapping: A hypothesis based on reintegration of the parietal inferotemporal pathways for visual processing. Psychobiology 17: 230-235.

104. McNaughton BL, Mizumori SJ, Barnes CA, Leonard BJ, Marquis M, Green EJ (1994b) Cortical representation of motion during unrestrained spatial navigation in the rat. Cereb Cortex 4: 27-39.

105. McNaughton BL, Morris RGM (1987) Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends Neurosci 10: 408-415.

106. McNaughton BL, O'Keefe J, Barnes CA (1983b) The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. J Neurosci Methods 8: 391-397.

107. Mehta MR, Barnes CA, McNaughton BL (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. Proc Natl Acad Sci U S A 94: 8918-8921.

108. Mesulam MM, Van Hoesen GW, Pandya DN, Geschwind N (1977) Limbic and sensory connections of the inferior parietal lobule (area PG) in the rhesus monkey: a study with a new method for horseradish peroxidase histochemistry. Brain Res 136: 393-414.

109. Meunier M, Bachevalier J, Mishkin M, Murray EA (1993) Effects on visual recognition of combined and separate ablations of the entorhinal and perirhinal cortex in rhesus monkeys. J Neurosci 13: 5418-5432.

110. Mitchell SJ, Ranck JB, Jr. (1980) Generation of theta rhythym in medial entorhinal cortex of freely moving rats. Brain Res 189: 49-66.

111. Mittelstaedt ML, Mittelstaedt H (1980) Homing by Path Integration in a Mammal. Naturwissenschaften 67: S. 566.

112. Mizumori SJ, Williams JD (1993) Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. J Neurosci 13: 4015-4028.

113. Moroni F, Corradetti F, Moneti G, Perani D (1981) The release of endogenous GABA and glutamate from the cerebral cortex in the rat. Naunun-Schmiedeberg's Arch Pharmacol 316: 235-239.

114. Morris RGM (1976) Hippocampal EEG during a ballistic movement. Neurosci Lett 3: 102-102.

115. Morris RGM (1981) Spatial localization does not require the presence of local cues. Learning and Motivation 12: 239-261.

116. Mountcastle VB, Lynch JC, Georgopoulos A, Sakata H, Acuna C (1975) Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. J Neurophysiol 38: 871-908.

117. Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J Neurosci 7: 1951-1968.

118. Muller RU, Kubie JL, Ranck JB, Jr. (1987) Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. J Neurosci 7: 1935-1950.

119. Murray EA, Mishkin M (1998) Object recognition and location memory in monkeys with excitotoxic lesions of the amygdala and hippocampus. J Neurosci 18: 6568-6582.

120. Nieder A, Freedman DJ, Miller EK (2002) Representation of the quantity of visual items in the primate prefrontal cortex. Science 297: 1708-1711.

121. O'Keefe J (1999) Do hippocampal pyramidal cells signal non-spatial as well as spatial information? Hippocampus 9: 352-364.

122. O'Keefe J, Burgess N (1996) Geometric determinants of the place fields of hippocampal neurons. Nature 381: 425-428.

123. O'Keefe J, Conway DH (1978) Hippocampal place units in the freely moving rat: Why they fire where they fire. Exp Brain Res 31: 573-590.

124. O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res 34: 171-175.

125. O'Keefe J, Nadel L (1978) The Hippocampus as a Cognitive Map. Oxford: Clarendon Press.

126. O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3: 317-330.

127. O'Keefe J, Speakman A (1987) Single unit activity in the rat hippocampus during a spatial memory task. Exp Brain Res 68: 1-27.

128. O'Reilly RC, McClelland JL (1994) Hippocampal conjunctive encoding, storage, and recall: avoiding a trade- off. Hippocampus 4: 661-682.

129. Ono T, Tamura R, Nishijo H, Nakamura K (1993) Neural mechanisms of recognition and memory. In: Brain Mechanisms of Perception and Memory: from Neuron to Behavior. (Ono T, Squire LR, Raichle ME, Perret DI, Fukuda M, eds), pp 330-355. New York: Oxford University Press.

130. Paxinos G, Watson C (1997) The rat brain in stereotaxic coordinates.

131. Penttonen M, Kamondi A, Sik A, Acsady L, Buzsaki G (1997) Feed-forward and feed-back activation of the dentate gyrus in vivo during dentate spikes and sharp wave bursts. Hippocampus 7: 437-450.

132. Philbeck JW, Behrmann M, Black SE, Ebert P (2000) Intact spatial updating during locomotion after right posterior parietal lesions. Neuropsychologia 38: 950-963.

133. Pinel P, Dehaene S, Riviere D, LeBihan D (2001) Modulation of parietal activation by semantic distance in a number comparison task. Neuroimage 14: 1013-1026.

134. Qin YL, McNaughton BL, Skaggs WE, Barnes CA (1997) Memory reprocessing in corticocortical and hippocampocortical neuronal ensembles. Philos Trans R Soc Lond B Biol Sci 352: 1525-1533.

135. Quirk GJ, Muller RU, Kubie JL (1990) The firing of hippocampal place cells in the dark depends on the rat's recent experience. J Neurosci 10: 2008-2017.

136. Ranck JB, Jr. (1973) Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. Exp Neurol 41: 461-531.

137. Ranck JB, Jr. (1984) Head direction cells in the deep layer of dorsal presubiculum in freely moving rats. Society for Neuroscience Abstracts 10: 599.

138. Redish AD (1999) Beyond the Cognitive Map: From Place Cells to Episodic Memory. The MIT Press.

139. Redish AD, Touretzky DS (1994) The reaching task: evidence for vector arithmetic in the motor system? Biol Cybern 71: 307-317.

140. Riedel G, Micheau J, Lam AG, Roloff E, Martin SJ, Bridge H, Hoz L, Poeschel B, McCulloch J, Morris RG (1999) Reversible neural inactivation reveals hippocampal participation in several memory processes. Nat Neurosci 2: 898-905.

141. Robertson RG, Rolls ET, Georges-Francois P (1998) Spatial view cells in the primate hippocampus: effects of removal of view details. J Neurophysiol 79: 1145-1156.

142. Rockland KS, Van Hoesen GW (1999) Some temporal and parietal cortical connections converge in CA1 of the primate hippocampus. Cereb Cortex 9: 232-237.

143. Rolls ET, Robertson RG, Georges-Francois P (1997) Spatial view cells in the primate hippocampus. Eur J Neurosci 9: 1789-1794.

144. Salinas E, Abbott LF (1995) Transfer of coded information from sensory to motor networks. J Neurosci 15: 6461-6474.

145. Salinas E, Abbott LF (1997) Invariant visual responses from attentional gain fields. J Neurophysiol 77: 3267-3272.

146. Saul LJ, Davis H (1933) Action currents in the central nervous system. Arch Neurol (Chic) 29: 255-259.

147. Save E, Poucet B (2000) Involvement of the hippocampus and associative parietal cortex in the use of proximal and distal landmarks for navigation. Behav Brain Res 109: 195-206.

148. Schwartzkroin PA, Mathers LH (1978) Physiological and morphological identification of a nonpyramidal hippocampal cell type. Brain Res 157: 1-10.

149. Scoville WB, Milner B (1954) Loss of recent memory after bilateral hippocampal lesions. 1957 [classical article]. J Neuropsychiatry Clin Neurosci 12: 103-113.

150. Semba K, Komisurak BR (1978) Phase of the theta wave in relation to different limb movements in awake rats. Electroencephalography and Clinical Neurophysiology 44: 61-71.

151. Sharp PE, Blair HT, Etkin D, Tzanetos DB (1995) "Influences of vestibular and visual motion on the spatial firing patterns of hippocampal place cells. J Neurosci 15: 173-189.

152. Sik A, Ylinen A, Penttonen M, Buzsaki G (1994) Inhibitory CA1-CA3-hilar region feedback in the hippocampus. Science 265: 1722-1724.

153. Sinclair BR, Seto MG, Bland BH (1982) Theta-cells in CA1 and dentate layers of hippocampal formation: relations to slow-wave activity and motor behavior in the freely-moving rabbit. J Neurophysiol 48: 1214-1225.

154. Skaggs WE, Knierim JJ, Kudrimoti H, McNaughton BL (1995) A model of the neural basis of the rat's sense of direction. Adv Neural Inf Process Syst 7: 173-180.

155. Skaggs WE, McNaughton BL (1996) Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience [see comments]. Science 271: 1870-1873.

156. Skaggs WE, McNaughton BL (1998) Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. J Neurosci 18: 8455-8466.

157. Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus 6: 149-172.

158. Snyder LH, Grieve KL, Brotchie P, Andersen RA (1998) Separate body- and world-referenced representations of visual space in parietal cortex. Nature 394: 887-891.

159. Squire LR (1992) Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans [published erratum appears in Psychol Rev 1992 Jul;99(3):582]. Psychol Rev 99: 195-231.

160. Stanton PK, Sarvey JM (1984) Blockade of long-term potentiation in rat hippocampal CA1 region by inhibitors of protein synthesis. J Neurosci 4: 3080-3088.

161. Sutherland RJ, Whishaw IQ, Kolb B (1983) A behavioural analysis of spatial localization following electrolytic, kai. Behav Brain Res 7: 133-153.

162. Suzuki WA, Amaral DG (1994) Perirhinal and parahippocampal cortices of the macaque monkey: cortical afferents. J Comp Neurol 350: 497-533.

163. Tanaka K, Saito H (1989) Analysis of motion of the visual field by direction, expansion/contraction, and rotation cells clustered in the dorsal part of the medial superior temporal area of the macaque monkey. J Neurophysiol 62: 626-641.

164. Taube JS (1995) Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. J Neurosci 15: 70-86.

165. Taube JS (1998b) Head direction cells and the neurophysiological basis for a sense of direction. Prog Neurobiol 55: 225-256.

166. Taube JS (1998a) Head direction cells and the neurophysiological basis for a sense of direction. Prog Neurobiol 55: 225-256.

167. Teng E, Squire LR (1999) Memory for places learned long ago is intact after hippocampal damage. Nature 400: 675-677.

168. Thompson LT, Best PJ (1990) Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. Brain Res 509: 299-308.

169. Treves A, Rolls ET (1991) What determines the capacity of autoassociative memories in the brain? Network 2: 371-397.

170. Trullier O, Wiener SI, Berthoz A, Meyer JA (1997) Biologically based artificial navigation systems: review and prospects. Prog Neurobiol 51: 483-544.

171. Ungerleider LG, Mishkin M (1984) Two cortical visual systems. In: Analysis of Visual Behavior (Ingle DJ, Goodale MA, Mansfield RJW, eds), pp 549-586. Boston: MIT Press.

172. Van Essen DC, Maunsell JH, Bixby JL (1981) The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. J Comp Neurol 199: 293-326.

173. Vanderwolf CH (1969) Hippocampal electrical activity and voluntary movement in the rat. Electroencephalography and Clinical Neurophysiology 26: 407-418.

174. Vanderwolf CH, Heron W (1964) Electroencephalographic waves with voluntary movement: Study in the rat. Arch Neurol (Chic) 11: 379-384.

175. Vanderwolf CH, Kramis R, Gillespie LA, Bland BH (1975) Hippocampal rhythmical slow activity and neocortical low voltage fast activity: Relations to behavior. In: The hippocampus (Isaacson RL, Pribram KH, eds), pp 101-128. New York: Plenum Press.

176. Vuilleumier P, Valenza N, Mayer E, Reverdin A, Landis T (1998) Near and far visual space in unilateral neglect. Ann Neurol 43: 406-410.

177. Wallenstein GV, Eichenbaum H, Hasselmo ME (1998) The hippocampus as an associator of discontiguous events. Trends Neurosci 21: 317-323.

178. Williams JM, Givens B (2003) Stimulation-induced reset of hippocampal theta in the freely performing rat. Hippocampus 13: 109-116.

179. Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space [see comments] [published erratum appears in Science 1994 Apr 1;264(5155):16]. Science 261: 1055-1058.

180. Wilson MA, McNaughton BL (1994) Reactivation of hippocampal ensemble memories during sleep [see comments]. Science 265: 676-679.

181. Winson J (1972) Interspecies differences in the occurrence of theta. Behav Biol 7: 479-487.

182. Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H (2000) Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron 27: 623-633.

183. Yamaguchi Y, Aota Y, McNaughton BL, Lipa P (2002) Bimodality of theta phase precession in hippocampal place cells in freely moving rats. J Neurophysiol 87: 2629-2642.

184. Yokota T, Fujimori B (1964) Effects of brain-stem stimulation upon the hippocampal electrical activity, somatomotor reflexes and autonomic functions. Electroencephalography and Clinical Neurophysiology 16: 375-383.

185. Zilles K (1985) The Cortex of the Rat. Berlin: Springer.

186. Zipser D, Andersen RA (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. Nature 331: 679-684.

187. Zola SM, Squire LR (2001) Relationship between magnitude of damage to the hippocampus and impaired recognition memory in monkeys. Hippocampus 11: 92-98.