Spatial learning and localization in rodents: Behavioral simulations and results

by

Rushi Prafull Bhatt

A thesis submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
MASTER OF SCIENCE

Major: Computer Science
Major Professor: Vasant Honavar

Iowa State University
Ames, Iowa
2000

Graduate College
Iowa State University

This is to certify that the Master's thesis of

Rushi Prafull Bhatt

has met the thesis requirements of Iowa State University

Major Professor

For the Major Program

For the Graduate College

TABLE OF CONTENTS

CHAPTER 1. OVERVIEW

1.1 Introduction

The computational strategies used by animals to acquire and use spatial knowledge for navigation have long been the subject of study in Neuroscience, Cognitive Science, and related areas. A vast body of data from lesion studies and cellular recordings directly implicates the hippocampal formation in rodent spatial learning [ON78]. The present model is based on the anatomy and physiology of the rodent hippocampus [She74, CS92]. We draw inspiration from the locale hypothesis, which argues for the association of configurations of landmarks in the scene to the animal's own position estimates at different places in the environment [ON78]. The system that generates the animal's own position estimate using the vestibular as well as motor commands issued is referred to as the Path Integration (PI) system [ON78]. The PI is hypothesized to play an important role in spatial tasks in terms of disambiguation of conflicting spatial cues and generation of relational information between available landmarks. At the same time, PI system is an important component of route based navigation system as it can supply place labels to incoming sensory cues which can be effectively used to compute the sequence of actions that can be used to guide the animal to a goal.

From Various lesion studies it has been found that the hippocampus plays an important role in spatial tasks, that is, tasks that involve specific locations and geometric relationships between objects. Lesions to hippocampus in rodents produce a severe deficit in learning of new spatial tasks, while at the same time keeping the stimulus-response type of task learning largely intact [MGRO82, MW93a, DMW99].

Apart from its linkage with the spatial tasks, it has been found that the hippocampus plays an important role in incorporation and retention of new memories. It has been proposed by numerous researchers that the hippocampus works as a temporary storage for new memories [Mar71, Buz89, CE93]. These memories are then supposed to be transferred to a more permanent storage site. This idea is still under scrutiny by many researchers and is not completely accepted. Opponents of this transfer of memory idea hypothesize that the hippocampus works more like an index for retrieval of memories and is not simply a temporary storage site [NM97, ESYB96, Eic96].

Despite such great interest in spatial learning and the hippocampal formation, there are relatively few computational or conceptual models that explain this wealth of experimental data at a low level, and yet be general enough to shed new light on the general principles that guide learning and behavior. By low level we mean a level at which the performance of the hippocampal system is explained with respect to its performance on specific tasks and the degradation of performance on the same tasks upon lesion or inhibition of the hippocampus. In this thesis, we examine how the computational model developed by Balakrishnan and Honavar (1999) [Bal99] explains some of the important experiments on the rodent spatial learning and navigation system which critically involves the hippocampus and its surrounding brain regions.

1.2 Organization of Thesis

The next chapter presents an overview of the anatmoical and functional characteristics of the hippocampal formation. Chapter 5 discusses some existing models for explaining the function of the hippocampus. In Chapter 6 we discuss the computational model for rodent spatial learning and navigation developed by Balakrishnan and Bousquet, and Honavar[BBH98d]. This model has been utilized in the present thesis to perform some behavioral experiments that validate the model in the context of reproducibility of rodent behavior. These experiments will be discussed in Chapter 7. We then reach the main conclusions of this thesis and propose some related future work in Chapter 8.

CHAPTER 2. REVIEW OF FUNCTIONAL AND ANATOMICAL ORGANIZATION OF THE HIPPOCAMPAL FORMATION

2.1 Introduction

The hippocampal formation consists of the hippocampus proper (also called “Ammon’s Horn”, abbreviated as CA) and surrounding structures; dentate gyrus and subiculum. The hippocampus is a bi-lateral limbic structure, the details of which will be discussed in the following sections.

The hippocampal formation is believed to be one of the major sites responsible for processing and incorporation of new spatial information [ON78]. One of the reasons to believe so is based on lesion studies carried out on animals, where it was found that animals with lesions in the hippocampal region were unable to learn spatial tasks [ON78, MGRO82].

Similar findings are also available in human subjects. A famous example is that of patient H.M. who, after undergoing bilateral resection of temporal lobes was rendered completely unable to learn and remember information. At the same time H.M. had a normal short-term memory and was able to recall many events that happened before the surgery [Mil72, Mil73]. Further, it was also found that H.M. was able to learn new motor skills but he was unaware of his coming across the task before.

Another case that attracted much attention was that of patient R.B.. Patient R.B. was found to be unable to learn and remember new information. On the other hand, he showed normal conditioning, priming behavior. After the death of R.B. due to unrelated causes, it was found that the degeneration was mainly confined to the CA1 region of the hippocampus [ZMSA86].

The role played by the hippocampus in human navigation is still a matter of controversy.

Recent studies using positron emission tomography (PET) scan and regional cerebral blood flow (rCBF) has shown that specific subparts of the human hippocampus are active when subjects perform navigational tasks. The measurements were made while human subjects played a virtual reality game (Duke Nukem 3D) and it was found that the activity in the right hippocampus was proportionate to the accuracy of navigation. Activity level in the left hippocampus did not co-vary much with the accuracy of navigation, and it was interpreted that this region actively maintains the memory trace of the destination during navigation of recollection of paths taken during previous trials [MBD+98]. Although the foregoing study precisely states the brain regions active during spatial navigation and planning, the role of human hippocampus during spatial tasks involving physical locomotion are still unclear.

More recent studies in primates have shown existence of "Spatial view" cells in the primate hippocampus. Recordings from the monkey hippocampus were performed while the monkey walked actively in an obstacle-free laboratory room. The cells firings were found to have a high correlation with which area of the environment, in this case laboratory walls, the monkey was looking at. This characteristic is in contrast with the rodent hippocampus where the pyramidal cells fire when rodent is at a particular place in its environment regardless of its orientation [RTR+98]. Previous studies by the same group have shown that 9.3% of neurons in primate hippocampus respond to stimuli appearing at some but not other corners of a computer monitor. 2.4% of cells only fired when a novel stimulus at specific locations on the computer monitor were displayed for the monkey [RMC+89]. Furthermore, it has been shown that these Spatial view cells encode location of stimuli with reference to an allocentric frame and not the position of stimulus on the retina, head direction or the location of the monkey [GFRR99].

2.2 Inputs and Outputs in the Rodent Hippocampal Formation

The hippocampal formation receives highly processed sensory information from the cortical regions via the entorhinal cortex and from non-cortical regions via the fornix. It also projects back to the cortical regions via the entorhinal cortex and to the subcortical regions via fornix, as will be discussed shortly.

2.2.1 Inputs to the hippocampal formation

The inputs to the hippocampal formation can be further divided into those from the cortical regions and those from the non-cortical regions.

The inputs from the cortical regions come via projections onto the entorhinal cortex which in turn projects to the hippocampal formation. Tract tracing experiments have shown major pathways originating from the inferior temporal gyri (higher processing area for visual sensory information), parietal and temporal lobes (higher processing area for auditory sensory information) and also from the frontal lobes. The only exception here is the olfactory information. Direct projections, as opposed to projections from higher areas of sensory information processing in the cortex, from the olfactory bulb and pyriform cortex onto entorhinal cortex have been found. Most inputs from these cortical regions arrive into the superficial cortical layers of the entorhinal cortex. Entorhinal cortex is thus believed to be a place for multi-sensory information integration. It has been contended that the hippocampal formation receives inputs form virtually all higher cortical regions [CE93].

The entorhinal cortex also receives inputs from amygdala, medial septum, dorsal raphe nucleus, locus ceruleus and parts of the thalamus. Most of the pathways from the septum onto hippocampus are cholinergic as well as GABAergic [TM91]. A recent study has found select nuclear divisions of the amygdala project to the entorhinal cortex as well as hippocampus, subiculum and parasubiculum in a topographically ordered fashion [PRS$^+$99]. A coarse topographic specificity in terms of projections from other cortical as well as non cortical areas has also been observed. More rostral cortical areas project more heavily to more rostral portions of the parahippocampal cortex, and more caudal neocortical areas project more heavily to more caudal portions of the parahippocampal cortex [CE93].

The inputs from the non cortical structures come via the fornix. The fornix carries inputs from the thalamus, septum, hypothalamus and other brainstem nuclei. The subcortical inputs are believed to serve as modulatory signals that influence activity in the hippocampal formation [TM91].

Most of the inputs to the hippocampal formation are believed to be excitatory, except some inputs from the septum, as mentioned above.

Apart from that, the hippocampus receives connections from the other hippocampus through the commissural pathways that join the two hippocampi across the midline [CS92].It is reasonable to assume that cortical inputs carry time-dependent information arriving from the external events noted by the sensory systems [TM91].

2.2.2 Outputs from the hippocampal formation

The entorhinal cortex also projects back to the same cortical areas mentioned above. Most of the projections from entorhinal cortex to the cortical regions have been observed to be from cortical layers V/VI.

Entorhinal cortex also sends fibers back to the septum. Projections to the septum also originate from the pyramidal cells in the hippocampus as well as from non pyramidal cells from the same area. The pyramidal cells of CA3 and CA1 subdivisions of CA project to the lateral septum. These connections are excitatory. The non-pyramidal cells in the hippocampus project to the medial septum/diagonal band. The latter projections are believed to be inhibitory [TM91].

2.3 Anatomy of the Hippocampal Formation

The anatomical connections in the hippocampal formation are some of the most well-studied, partly due to the systematic structure in the connections themselves, and partly because of the numerous examples of neuronal plasticity that have been observed in this region.

The hippocampus is an elongated C shaped structure spanning from the septal nuclei to the temporal cortex. Figure 2.1 shows major synaptic connections found in slices taken perpendicular to to the long axis (septotemporal axis). These slices show a structure that resembles two interlocking ‘C’ shaped arrangements. One ‘C’ is known as the Dentate gyrus, within which the granule cell layer is the principle layer. The other ‘C’ is the hippocampus proper, and is also referred to Cornu Ammonis (abbreviated as CA) and is subdivided into regions CA1 through CA4. Such a structure is seen all along the septotemporal axis.

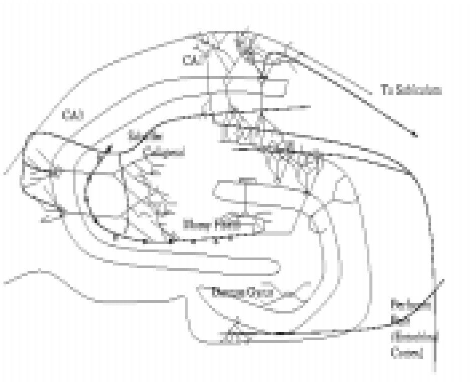

Figure 2.1 Schematic of connections in hippocampus

Using simultaneous stimulation and recording studies as well as neuronal tract tracing techniques, the major connection pathways in hippocampal formation are now well-understood [AW89].

In brief, the circuitry can be described as follows.

2.3.1 Entorhinal cortex and the perforant path

The entorhinal cortex is the origin of a strong projection (the perforant pathway) to the dentate gyrus and hippocampus [AW89].

Cortical layers II and III of the entorhinal cortex project via perforant path fibers into the dentate gyrus, the hippocampus and the molecular layer of the subiculum. Anterograde and retrograde tracing techniques have shown that projections from very constrained regions in the entorhinal cortex reach wide septotemporal areas of the dentate gyrus, as opposed to earlier suggestions by Andersen, Bliss and Skerde (1972) . It can also be concluded that a single layer of the dentate gyrus along the septotemporal axis is innervated by multiple focal points in the entorhinal cortex. One of the suggestions for the functional role is that these highly divergent connections result in a sparse, distributed encoding of the highly processed sensory information which is then delivered to the hippocampus via the dentate gyrus.

It should also be noted here that far reaching projections of the inhibitory 'basket cells' in the dentate gyrus have also been found [SDL78]. It can therefore be concluded that excitatory as well as inhibitory connections span over long distances over the septotemporal axis in the hippocampus.

2.3.2 Dentate gyrus and mossy fibers

The dentate gyrus receives projections from the entorhinal cortex through the perforant path. Dentate gyrus in turn provides input to the CA3 layer of the hippocampus proper.

"The dentate gyrus can be divided into three layers: The molecular layer, in which perforant path fibers terminate; the granule cell layer, which is populated by the principal cell type, the granule cell; and a deep or polymorphic layer which is populated by a variety of neuronal types. The granule cells give rise to the mossy fibers, which collateralize in the polymorphic layer and then enter the CA3 layer where they form en passant synapses with the proximal dendrites of the pyramidal cells" [AW89]. These mossy fiber synapses with CA3 are strong, which has led researchers to suggest that they provide the context [O'K89] or reference frame of the task to be performed [MBG^{+}96], by transformation of the sensory input activity arriving at entorhinal cortex into a non-overlapping activity pattern of granule cells, which in turn are conveyed to the pyramidal cells in the CA3 layer. It has been found that on an average, CA3 pyramidal cells get about 330,000 contacts from the mossy fibers collectively, or, a granule cell makes contact with around 14 CA3 pyramidal cells and each CA3 pyramidal cell is innervated by only about 46 granule cells [CS92].

2.3.3 CA3-CA1 connections

The hippocampus proper can be further divided into stratum moleculare, stratum radiatum, stratum of pyramidal cell bodies and stratum oriens as one progresses further from the septotemporal axis radially.

The CA3 region of the hippocampus primarily contains pyramidal cells that show a characteristic complex spiking behavior. These pyramidal cells are arranged in a single layer, in stratum of the pyramidal cell bodies. The dendrites of these cells descend into the deeper stratum oriens as well as into stratum radiatum and stratum lacunosum-moleculare.

The CA3 and CA1 pyramidal cells receive inputs from three different sources: (i) from cortical layer II of entorhinal cortex through the perforant path, synapses of which are formed in the uppermost parts of the apical dendrites of pyramidal cells in CA3. Some direct connections from cortical layer III of entorhinal cortex to CA1 apical dendrites have also been found. (ii) from dentate gyrus through mossy fibers which make connections to the proximal apical dendrites of CA3 pyramidal cells, and (iii) recurrent inputs from other CA3 pyramidal cells. See Figure 2.2 for a clear picture of the connections mentioned above. Unlike CA3 cells, CA1 cells do not project to other pyramidal cells of CA1 [CS92].

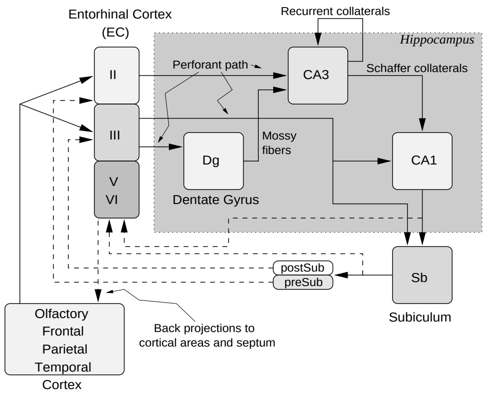

Figure 2.2 Schematic of major connection pathways in hippocampus

The number of recurrent collaterals on a CA3 pyramidal cell from other CA3 pyramidal cells is believed to be around 6000, or about 1.8% of the CA3 cell population. These recurrent connections are located on the dendritic tree spaced between the mossy and perforant inputs [CS92]. Some researchers have likened the structure of the recurrent collaterals to an auto-associative recurrent network suggesting that CA3 serves as a pattern completion device capable of recalling entire scenes from partially observed data [Mar71, Rol90]. However, others have suggested a hetero-association role, suggesting that these collaterals predict future activations of the neurons based on the current activations [McN89, MN89]. Some experimental evidence for this latter view is provided by [SM96].

In CA3, mossy fibers from the dentate gyrus project into a region just above the pyramidal cell layer. Axons from CA3 pyramidal cells then make highly collateralized connections that terminate within the CA3 layer and make strong projections into the CA1 layer via the Schaffer collaterals.

As discussed earlier, the CA1 pyramidal cells receive excitatory inputs from the entorhinal cortex via the perforant path and from the CA3 pyramidal cells through the Schaffer collaterals. Axons from the CA1 pyramidal neurons project via the alveus to the subiculum and also to the deep cortical layers of the entorhinal cortex. Subiculum also receives input from the entorhinal cortex and projects to the pre- and post-subiculum, the deep layers of the entorhinal cortex, and to the hypothalamus, septum, anterior thalamus and the cingulate cortex. All these connections are excitatory [CS92].

The afferents from brain-stem areas to the hippocampus proper have been found to synapse mostly in stratum lacunosum/moleculare of CA1 and CA3 and in a restricted part of the hilar zone under the granule cells in the dentate gyrus. These connections are believed to contribute towards most of the serotonin found in hippocampus [ON78].

It has also been found that for the CA3 pyramidal cells, majority of dendrites were located in stratum oriens, while almost the same amount of dendrites were present in stratum radiatum and stratum lacunosum-moleculare. For the CA1 region pyramidal cells, majority of dendrites were found in stratum radiatum. Presumably, more dendrites in a particular stratum would mean a greater number of en passant synapses as well as synapses with interneurons in the stratum in question [AW89].

CA3 and CA1 regions also contain interneurons that suppress the activity in pyramidal cells. It has been found that the pyramidal cells make excitatory connections to the basket cells present in the stratum of pyramidal cells, which in turn provide GABAergic inhibitory input back to the pyramidal cell. It has been found by experiments with CA1 cells that when Schaffer collaterals and commissural axons in stratum radiatum were stimulated, the range of frequencies under which LTP was produced increased in the presence of a GABA type A receptor agonist (muscimol), while LTP was induced only at very low frequencies in presence of GABA type A antagonist (picrotoxin) [SM99]. Thus, this inhibitory loop, presumably between pyramidal and basket cells is believed to be a controlling mechanism for LTP and LTD induction.

It has also been found that interneurons in the lacunosum-moleculare region of CA1 are not restricted to the CA1 region. The dendritic and axonal processes of some of these interneurons were seen ascending in stratum lacunosum-moleculare, crossing the hippocampal fissure, and coursing in stratum moleculare of the dentate gyrus. Stimulation of hippocampal afferents caused excitatory as well as inhibitory postsynaptic potentials in these interneurons. EPSPs were most effectively elicited by stimulation of fiber pathways in transverse slices, whereas IPSPs were predominantly evoked when major pathways were stimulated in longitudinal slices. Thus, these interneurons are different in characteristics from the interneurons (for example, basket cells) and the pyramidal cells [LS88a, LS88b].

It has recently been found that during rhythmic oscillations in area CA3, interneurons with similar dendritic and axonal arbors behave differently. One group of interneurons is powerfully excited by CA3 pyramidal cells, whereas two other interneuron groups were relatively unaffected by pyramidal cell firing. One of these groups of interneurons is inhibited by other local interneurons during the pyramidal cell bursts. Thus, morphologically similar interneurons are wired radically differently and hence produce very dissimilar firing characteristics [MWK98]. It has also been found that one group of these interneurons can undergo LTP while another group is incapable of undergoing LTP. mGluR system is supposed to govern this kind of behavior [Informal talk by Lacaille]

The major connections in the hippocampus can be summarized as shown in Figure 2.2 [BBH98a].

2.4 Physiological Properties of Hippocampal Cells

2.4.1 Spatiality and directionality of pyramidal cell firing in area CA

Apart from evidence from lesion studies, cellular recordings from pyramidal cells in the hippocampal formation of behaving rodents have show that many such cells fire in complex spike bursts only when the animal is in a constrained region of its environment. Such cells show a characteristic, complex spiking behavior which is distinct from other cells found in the vicinity of these cells. O'Keefe named them place cells and the corresponding regions where each is active, the place field [O'K76]. The specific regions in which these cells fire are well defined for a given environment and can be manipulated by changing the size and sensory cues available in the environment within which the experiments are carried out [ON78, OD71, TSE97].

Cells with such location-specific firing have been found in almost every major region of the hippocampal system, including the entorhinal cortex [QMKR92], the dentate gyrus [JM93], hippocampus proper [OD71, O'K76], the subiculum [BMM$^+$90, SG94], and the postsubiculum [Sha96].

In addition to place cells, head-direction cells have also been discovered. These cells respond to particular directions of the animal's head, irrespective of its location in the environment. Each such cell fires only when the animal's head faces one particular direction (over an approximately 90 degree range) in the horizontal plane. The firing of these cells can be altered by a complex interaction between visual and angular motion signals. Importantly, in every case reported to date, any manipulation that alters the reference direction of one of these cells results in a corresponding alteration in the reference direction for the whole system which is in contrast to hippocampal place cells where partial re-mapping of groups of cells encoding the same environment is possible. These cells were first discovered in the postsubicular area of the hippocampal formation [Ran84, TMR90a, TMR90b]. Since then, such directional cells have also been discovered in the retrosplenial cortex [CLBM94, CLG$^+$94], the anterior thalamus [Tau95, BS95a], and the laterodorsal thalamus [MW93b].

A number of experiments have been performed in order to determine the properties of the place and head-direction cells. It is now known that the spatial representation in the place cells is not grid-like, i.e., adjacent neurons are as likely to represent distant portions of the environment as close ones [O'K76, MKR87, OS87, O'K89, WM93]. Also, place cells are active in multiple places in the environment [OS87] and also in multiple environments [KR83, MKR87, MK87]. Further, places appear to be represented in the hippocampus using an ensemble code, i.e., a set of place cells appear to code for a place [WM93].

Experiments have also revealed that when the animal is introduced into a familiar environment, place fields are initialized based on visual cues and landmarks [MK87, MKR87, SKM90]. Once initialized, the place fields have been found to persist even if the visual cues are removed in the animal's presence [OS87], implying that place cell firing must also be maintained by a source other than visual stimulus. It has been found that place fields of CA1 cells are conserved in darkness, provided the animal is first allowed some exploration of the apparatus under illuminated conditions [MLC89, QMK90]. This has led to the hypothesis that place fields are maintained by ideothetic (self-motion) mechanisms, i.e., by the path integration system.

2.4.2 Spike characteristics of hippocampal place cells

It has been found that the place cells show a very characteristic firing pattern within the firing field. These cells fire in bursts that last for a maximum of 2 seconds and fire at a peak rate of around 20 action potentials. Further each burst contains around 10 to 20 action potentials. This firing pattern is further modulated by the ubiquitous theta wave modulation present whenever the animal is in a state of locomotion [ON78]. It has been found that even though the firing of these place cells is highly correlated with the location of the animal, the firing itself is not robust in the sense that across visits to the same place does not reliably produce the bursts. Even if the animal takes a path through the place field which is very similar to a previous path there is no guarantee that the same or a similar sequence of action potentials will be observed. In fact, this variability is found to be in excess of what one should expect if the generating process for these spikes is supposed to be a Poisson process with the probability of firing set to the mean firing rate of the place cell in question measured over the complete recording session [FM98]. Apart from this "excess variability" observed in place cells, it has been found that the hippocampal place cells are strongly modulated by the activity in the inhibitory interneurons. In a novel environment, it has been found that the amount of activity in these interneurons is low and hence the amount of synaptic inhibition on the place cells is significantly low which then grows as the familiarity of the animal to the environment increases [WM93].

2.5 Conclusion

The large body of experimental data found over the years begs a functional explanation. Some efforts in this direction have been made by a number of researchers [RT98, NM97, MBG$^+$96, CE93] which will be discussed in the following chapters. Most of these efforts have either been towards explaining a subset of experimental data or towards delivering a very high level theory which would be difficult to justify using neuroanatomical data without many assumptions which cannot be directly justified.

The precise mechanisms that give rise to this intriguing behavior of cells in and around the hippocampal formation as well as the reason for survival of a structure like the hippocampal formation in the evolutionary process are yet to be determined. The uniformity of the types of defects that arise across many species following damage to the hippocampal formation gives strong evidence in the favor of the idea that hippocampal formation is of prime importance to the overall process of memory incorporation and learning. Efforts towards a better understanding of the processes in hippocampal formation seem to have taken two separate paths. One school works on the lowest possible level to discover the mechanisms that govern long term potentiation (LTP) of cells in order to discover the molecular basis of LTP, while the other school works at a very high level systems approach where the firing characteristics of cells or the hippocampal formation is directly linked to behavior of animals, for example, see [Sha97, ON78, MGRO82]. The modeling efforts in this direction also have been limited to some high-level explanations of place field characteristics [BDJO97] which shed little light on the precise mechanisms required for the other types of cells in the hippocampus. At the other extreme, efforts are concentrated on some cell-level modeling techniques that only addresses induction of LTP etc [WLJS92].

Interesting avenues of research have also opened up in the direction of mutant and knockout models where the behavioral as well as in-vivo recordings shed new light on the underlying learning processes. Such experiments are extremely useful for pin-pointing the exact brain regions and their roles in memory incorporation [KHH$^+$98, MBT$^+$96, CGT$^+$98] to consider the set of experiments. Modeling efforts at this "middle level" that bridge the gap between cellular/molecular level theories of LTP induction and the large-scale behavioral level theories are worth pursuing in the light of these mutant and knock-out studies.

In conclusion, the overall functional role of the hippocampal formation is still largely unknown and researchers are only beginning to understand the various mechanisms and connections that give rise to learning and memory formation mediated by the hippocampal formation. Modeling efforts at all levels of explanation are in order at this point where there is an abundance of data which can not be explained with a single, coherent, set of theories of learning.

CHAPTER 3. SPATIAL LEARNING FROM A COMPUTATIONAL CONTEXT

3.1 Background

As we saw in Chapter 2, lesion and pharmacological studies performed by numerous experimenters suggested that hippocampus is critically involved in the formation of a "cognitive map". The cognitive map in this context denotes representation of objects in an animal's environment which is formed by storing relative spatial positions of different objects.

It is interesting to note that there is growing evidence that some storage of relative temporal positions or sequences of events are also stored in the hippocampus [NM97, SM96, Eic96]. Hence, hippocampus is no longer believed to be a static map of one's environment, but a region where new spatial as well as temporal information is learned and consolidated.

A number of diverse models for spatial learning and navigation exist in the literature. It is generally accepted that the hippocampus is the prime site linked with learning of a wide range of tasks, most of which have a distinct spatial component [ON78, Hea98]. Most models that deal with biologically inspired spatial learning, and therefore have the hippocampus as a major part, are conceptual models [MBG^{+}96, Mar71, RH96, NM97, Eic96], some are concrete computer simulations[Zip86, RT96, RT98, BA96] and some have also been implemented in robots[Bro85, BDJO97].

The above models mostly fall into two major categories. Either they are high level conceptual descriptions of underlying psychological or physiological phenomena [Mar71, Buz89, CE93, MW93a, NM97] or they are a low-level task oriented specification of hippocampal function [Zip86, RT98, BS95b]. There is a general lack of literature in the area of specifying exactly the computational requirements for learning a representation of ones spatial environment. These models will be discussed in more detail in Chapter 5.

In this chapter, the issue of representation of spatial information in order to perform localization and navigation in one's environment is discussed. What kind of information is required and how it should be represented by an animal (or a mobile robot, for that matter) in order to successfully navigate in its environment needs to be clearly defined. We discuss the aforementioned issues in the light of a biologically plausible computational model developed by BalakrishnanEtAl (1998) [Bal99, BBH98a] for rodent spatial navigation and localization. The model is based based on the locale system hypothesis suggested by O'Keefe and Nadel (1978) . The main idea behind locale system hypothesis is that landmarks, or more generally, objects are represented as relations between one another in terms of their relative positions or some other similar metric. Such a representation is supposed to arise due to a fusion of incoming sensory information of different modalities with the path integrator estimate. Here, from a spatial learning context, path integrator means a system that keeps track of the animal or mobile robots own position with respect to an allocentric frame of reference. In a more general case, path integrator can be considered as a series of operations performed on objects available to the animal to change their configuration, or, their relative position according to a fixed metric. In the spatial learning context, it is assumed that the path integrator uses movement commands issued to the motor system as well as the independent movement sensors (e.g. vestibular system in animals or acceleration/velocity sensors in mobile robots) as its input.

In the following sections, and in rest of this thesis, we discuss the spatial learning problem in the context of learning about one's physical environment for the purpose of localization and navigation. The arguments presented here can be extended for learning arbitrary relations between objects in one's environment, given a metric for measuring relative positions of objects.

3.2 Sensory Information Available to a Mobile Robot or a Navigating Animal

3.2.1 Linear distance based measurements

It is reasonable to assume that the primary information available to a navigating entity, an animal or a mobile robot, is the available landmarks in the environment. Any stable an sensorily distinct object in the environment can be considered to be a landmark for the purpose of learning a spatial environment. It has been substantiated by experiments [ON78] and simulations [RT96] that rodents use the perceived size of prominent objects and therefore presumably estimated distance of landmarks from their current position as landmarks for learning spatial environments. In an enclosed environment, distances of walls of the environment also serve for the learning purposes [BDJO97, OB96].

Distances of three unique landmarks in a two dimensional spatial environment are sufficient in order to find one's location in the environment, provided that the allocentric position, or in the case of animals with respect to position of goal or home, is available. It can be easily shown that in such a case an estimate of landmarks can be found by solving the following equations for x and y which are the unknown Cartesian coordinates of one's position.

$$d_1^2 = (x - x_1)^2 + (y - y_1)^2$$
$$d_2^2 = (x - x_2)^2 + (y - y_2)^2$$
$$d_3^2 = (x - x_3)^2 + (y - y_3)^2$$

In the above equations, d_1, d_2 an d_3 are the distances of landmarks from the subject's current position and (x_1, y_1), (x_2, y_2) and (x_3, y_3) are the coordinates of landmarks. After simple manipulations the solution of above equation reduces to solution of two linear equations. It can be seen above that the animal only needs to represent the allocentric location of three unique landmarks in order to successfully navigate. In case landmarks are not unique, it is required that there be at least three landmarks that are not arranged in a symmetric fashion.

3.2.2 Angular distance based measurements

In the case when the subject has compass information available to it, the task becomes significantly easier, as only the coordinates on one landmark need to be known in order to successfully navigate. It has been hypothesized that animals must use a strategy where distal cues are used to reset the compass direction estimate, while local cues are used for more precise position estimate [TSE97].

Indeed, distinct sets of cells have been found in the hippocampal formation of rodents whose firing has strong correlation with spatial nature of the task; the hippocampal pyramidal cell have strong correlation (among other parameters) with the position of the animal's head, regardless of its direction. These cells are therefore aptly named place cells. Also, the spatially constrained region in which these cells fire with high frequency are called place fields[ON78]. On the other hand Presubiculum[Ran84], Anterior Thalamic nucleus[Tau95], and Lateral Dorsal Nucleus [MW93b] of the Thalamus have cells that fire in preference of the animal's head regardless of its position in the environment. There is substantial evidence that anterior Thalamic head direction cells fire in anticipation of the head direction in the rat [BS95a]. This gives strong evidence in favor of a predict-observe-correct model for head direction system [RT96].

The above hypothesis, where there are are systems that keep track of compass information and allocentric positions of only a few prominent and stable landmarks in the environment, is attractive when the subject can judge its distance and angle with respect to a "north pole" from the landmarks. When information of other modalities like odor and sound are present. These modalities can only supply directional information about sound and odor sources. In such a case, it is still possible to obtain estimate of one's position if multiple, unique sensory modalities are available. Computations involved in such a case require access to an allocentric compass direction, as opposed to allocentric positions of landmarks. Also, the transformations required in order to find one's position do not remain linear as in the above case. Nevertheless, given a set of distal cues like far away and stable objects like mountains, sun etc. can be used to successfully reset compass direction, after which only angular distances need to be computed [Zip86].

3.3 Representation of Information

3.3.1 Storage

In any case, there nonlinear reverse mappings from the observations of landmarks, be it linear or angular, to places are required. Either significant amounts of computations are required at each point in order to determine place, or places need to be simply represented as snapshots of sensory information available at the different places. In the latter case, the problem reduces to that of storing these observation vectors along with the estimated positions of the subject whenever such sensory information is observed. Furthermore, it is more advantageous to store relationships between pairs, or more generally subsets, of available landmarks in terms of their linear or angular distances.

Whenever the subject needs to find its own position, it can make observations of either angular positions of the landmarks (in case compass information is available) or linear distance measurements of landmarks in case their allocentric position is known. After this, the closest match to this incoming measurement needs to be found from the previously stored observations. The place label associated with nearest the neighbor of observation from these previously stored vectors of observations can then be used to determine the subject's position.

This brings us to an important issue. How should one generate the position labels that are required for labeling these sensory information vectors? Navigating animals and mobile robots usually have some way of sensing their own displacement from the motor commands issued or from acceleration sensors, e.g. vestibular inputs in animals). This information can be used to keep track of subject's own motion which will henceforth be called path integration system. When the subject is first introduced in the environment the path integration system is reset to an arbitrary state and is associated to the first sensory observation vector available to the subject. After each motion step, the subject reaches a new place in its environment and therefore has a sensory information which is different from that available previously. At the same time subject has the updated path integrator estimate after the motion step is performed. If the sensory information at this new place is significantly different from the previous step, based on some discrimination metric, a new prototype vector for this new place can be stored along with the newly generated path integrator estimate. In such a fashion tuples containing $(sensory_information, path - integrator_estimate)$ can be generated and represented.

3.3.2 Retrieval

Upon re-entry in the environment the localization problem reduces to that of finding the closest match from the available sensory vectors and reseting the path integrator estimate to the associated label to these vectors. Computationally, this is a hard problem, mostly because there is no general way of fixing a distance measure between these stored sensory information vectors.

Even if the environment where subject is being introduced is identified straight away, search for sufficiently similar sensory information vector from those stored is time consuming. The problem becomes even more severe when multiple environments are stored by the subject and it also has to pick the right map before setting its path integrator estimate based on the stored sensory information for the map (the map selection problem).

Most of the techniques available in Vector Quantization literature and Associative Memory literature can be applied to tackle this sensory information vector storage problem [Koh89, Has95, Mar71].

3.3.3 Noise and inaccuracies in sensors and path integrator

The sensory measurements as well as the path integrator estimates are prone to noise. Effects of such noise on the path integrator can be seen in the drifting of place fields in darkness [MBG+96]. The model discussed in the next explicitly addresses the errors in sensory measurement and path integration estimates. The model effectively reduces these errors by applying a Kalman Filter [Kal60] like update mechanism.

3.4 Conclusion

In this chapter, we gave an overview of the type of information processing required of an animal or a mobile robot in order to successfully learn and represent spatial environments. Using the computational model proposed in this chapter, we have simulated some of the important spatial learning experiments performed on robots and have found the functioning of the model to be satisfactory [BBH98b, BBH98a].

In the following chapters, we discuss the model used for these behavioral simulations 6 and the results of these simulations 7.

We have completely ignored the issue of route-based or topological representation, where relations between places are stored instead of the positions of landmarks or expected sensory information available at different places in the environment. The route-based representation scheme has the advantage that the subject does not need to plan its route at every point of the environment. Once the subject knows the task at hand and is able to localize in its environment, it just needs to pick a program, or a sequence of actions, that can take it to a desired goal location. It is easy to see that such route based navigational systems can be generated in the framework discussed above. Upon reaching the reward sequence of recent activations of place cells along with the motor system commands that were issued can be stored as a program. Next time onwards, when faced with similar situation, the subject just needs to pick the right program and execute it in order to reach the goal. Thus, spatial learning can give rise to a stimulus-response type of behavior. Again, the issue of sorting these programs, and hence, learning a total order of these programs in terms of their relevance to the tasks at hand needs to be addressed.

It has been found that a hippocampal place cell is usually involved in representation of multiple environments and has place field characteristics that are radically different from environment to environment [ON78]. At the same time, manipulation of sensory cues can cause rapid and drastic re-mappings in the size, shape, and even existence of place fields [KR83, MCM91, SKM90]. It can be argued that many of the sensory cues are “characteristic” in the sense that they occur very frequently and therefore place cells must be capable of detecting this characteristic sensory information in various environments. Thus, by encoding only a few dimensions of the complete sensory vectors instead of complete sensory information, available storage can be utilized more efficiently. At the same time, many of the environments represented can be quickly eliminated simultaneously by non-firing of only a few place cells, thus reducing the complexity of the map selection problem problem mentioned above. Again, choosing the correct subset of features to be encoded is a problem that should be addressed if such encoding is to be made feasible.

Also, it is supposed that these place cells do not simply respond to the incoming sensory information but are also modulated by the path integrator estimate, so that the place cell fires only when the path integrator estimate is in accordance with the sensory information which should be present given the path integrator state [ON78]. Also, the place cells appear to be modulated by the context of the spatial task to be performed [FM98]. Whether such dependence on multiple, diverse controlling parameters is computationally more efficient needs to be seen.

CHAPTER 4. FORMALIZATION OF COGNITIVE MAPS FOR SPATIAL LEARNING AND NAVIGATION

In this chapter we elaborate on the intuitive idea of spatial learning and localization. We will define the minimum requirements for a system that is capable of learning, representing, and retrieving the spatial information in order to perform navigational tasks. We will develop these requirements based on the idea of locale hypothesis as developed by O'Keefe and Nadel (1978) .

In order to develop the formal requirements for spatial learning and localization, we accept the existence of two systems that feed into each other. The first system is the Local View (LV) which represents the current and (possibly) immediately preceding percepts of an animal or a mobile robot. The second system is the Path Integrator (PI) system that represents the current position of an animal or a mobile robot. These two systems together can be thought of as a "spatial map" for navigation and localization.

The inputs to these systems, or alternatively the spatial map, can be all or a subset of the input information discussed in Chapter 3.

Here, we formally discuss how this incoming information can be utilized to perform the navigation task. In order to do so, we divide the different operations performed on the map into two categories, namely, the intra-map operations and the inter-map operations.

In the discussion to follow, the variables denoted in bold typeface are vectors.

4.1 Preliminaries

The intra-map operations are closely related to the storage issues discussed in Chapter 3, section 3.3.1. These are the operations on the spatial map that enable the spatial learning system to incorporate information about different places in a given environment in order to successfully reset the PI upon re-entry. Let us first define the percepts available to the animal or mobile robot.

4.1.1 The local view

Consider L to be the set of landmarks in the environment, so that:

$$L = \{l_1, l_2, ..., l_n\}$$

We consider a perceptual observation $\mathbf{O_x}$ to be a set of vectors $\mathbf{o_{x,i}}$, each describing the observation of a landmark l_i place $\mathbf{x}$ in the environment. It is possible to have one (possibly unique) observation $\mathbf{O_x}$ for each place $\mathbf{x} \in \mathbf{M}$ where M is the current environment.

Hence, assuming that we do not have any missing observations of landmarks, $\mathbf{O_x}$ is an n-tuple $(\mathbf{o_{x,1}}, ..., \mathbf{o_{x,n}})$. Here, $\mathbf{o_{x,i}}$ is a vector that describes landmark i in the system. For the biological spatial navigation system, this is usually considered to be an activity pattern of neurons in the sensory system. For a robotic system, $\mathbf{o_{x,i}}$ can be denoted by a d_i dimensional feature vector such that $\mathbf{o_{x,i}} \in \Re^{\mathbf{n}}$.

To give a simple example, the observation $\mathbf{O_x}$ can be a sequence of landmarks scanned in an orderly fashion. As seen in Figure 4.1, at a given position, the observation could be a systematic scan of the environment. In such a case, $\mathbf{O_{x_1}} = \mathbf{dcab}$ while $\mathbf{O_{x_2}} = \mathbf{abcd}$.

Here, the observation *dcab* simply denotes the relative position of landmarks as observed from positions $\mathbf{x_1}, \mathbf{x_2}$. We diverge from the notation described earlier, for the sake of convenience in representation. The same information can be conveyed by representing each observation of landmark by the relative positions of other landmarks from a given landmark.

It can be shown that, for unique landmarks in general position, different regions of the environment can be identified uniquely using the landmark observations as described above.

For identification of regions outside of the convex hull formed by the landmarks in the environment, only the observations of landmarks on the hull are required, as it is evident from Figure 4.1. Each region outside of the convex hull made up of landmarks is labeled by the observation in the above convention for that region.

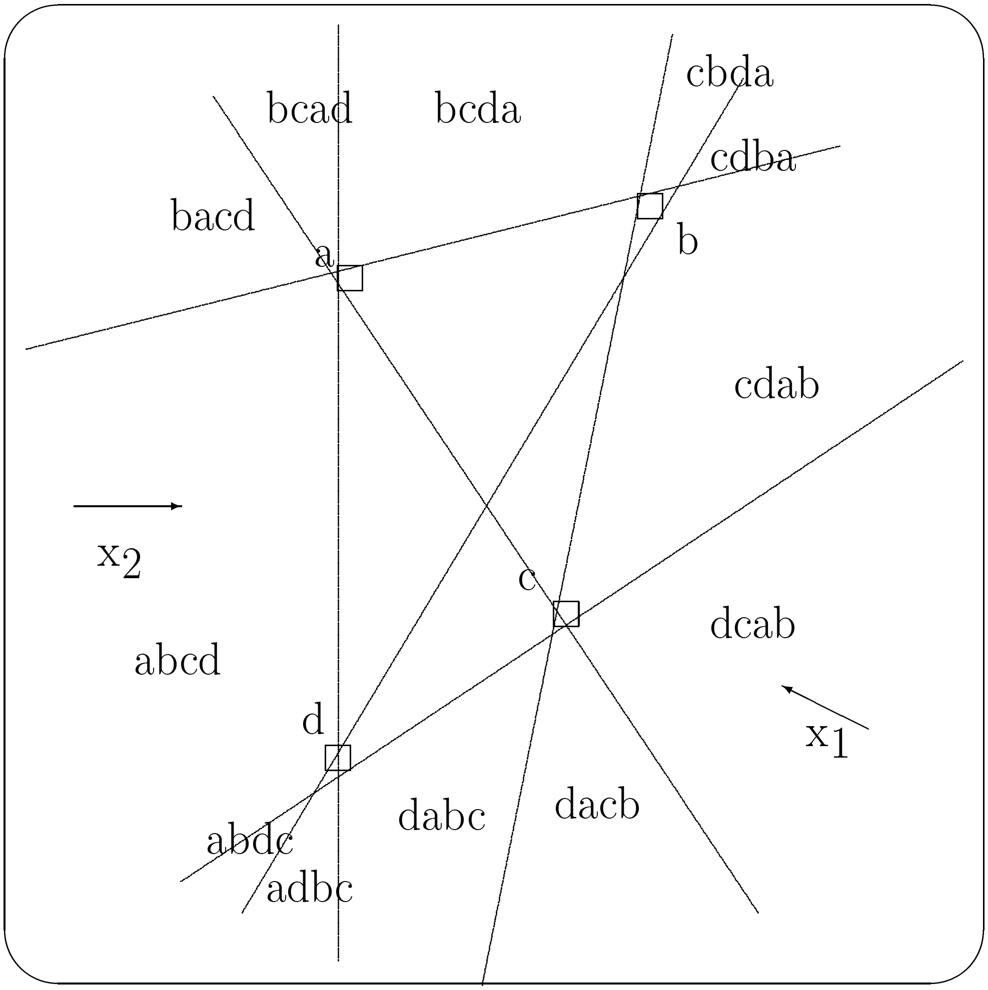

Figure 4.1 Observations of landmarks from different positions in an environment

For regions within the convex hull, angular measurements become necessary and therefore a richer observation is called for.

In the general case, the observation of the environment at time t does not only depend on the current sensory information, but also on a history of sensory inputs. Let us define the Finite-time Local View (LV_0^t) as follows.

Definition 4.1 (Finite-time Local View (LV_0^t)) *Finite-time Local View*

$$LV_0^t = \mathbf{O}_{\mathbf{x_0}}, \ldots, \mathbf{O}_{\mathbf{x_p}}$$

is a finite vector of observations from time $t = 0$ to $t = p$ where $t = 0$ is time at which either the agent was first introduced in the environment or the time at which the current context first became relevant. $t = p$ is the present time, or, the time of the latest observation available to the agent.

Similarly, we can also define the Infinite-time Local View ($LV^t_{-\infty}$) which contains the complete history of observations. We shall see later that the Infinite-time Local View has no direct value except that it is implicitly utilized by the animal while learning to perform actions for a given LV and PI state. Of course, assuming a finite age of the agent, LV cannot be infinite in length, but it is unbounded.

Definition 4.2 (Infinite-time Local View ($LV^t_{-\infty}$)**)** *Infinite-time Local View*

$$LV^0_{-\infty} = \mathbf{O_{x_{-e}}}, \ldots, \mathbf{O_{x_0}}, \ldots, \mathbf{O_{x_p}}$$

is an unbounded vector of observations from time $t = -e$ to $t = p$ where $t = -e$ the first known observation of any environment. Also, $t = p$ is the present time, or, the time of the latest observation available to the agent.

It is not possible for an agent to use $LV^t_{-\infty}$ as its local view of environment, because it grows with experiment and record of observations of unrelated environments is of little value in the present context of navigation. Instead, it only a finite history of observations that matters for decision making. Of course, $LV^t_{-\infty}$ matters in the sense that it is used while learning to select useful actions.

Our definition of LV^t_0 is still inadequate in the sense that not all elements of LV^t_0 will be useful for selection of actions. Furthermore, some observations are more important towards selection of action than others. This is true for observations $\mathbf{O_{x_i}}$ as well as the measurements of individual landmarks within each observation. Let us define appropriate weighting functions for the observations for taking this fact into account. The state of the PI system determines which observations of the environment are to be more useful for navigation than others. At the same time, past experiences of the animal dictate which landmarks are more stable or reliable for the task than others. For example, if the task is to find a route home, an antenna tower or a large building is more reliable for performing tasks than an truck or a certain color parked in the curbside. On the other hand, one might wish to use all the vehicles on the freeway as landmarks for the purpose of driving as opposed to the trees and buildings around the freeway, as they are much more important for completion of the navigational task. Let us consolidate all these experience and task driven facts into the function *Filter* as follows.

Definition 4.3 (Local View Filter) *Filter is a function that takes as its input the current PI state* $\mathbf{s}$ *and* LV_0^t *and returns a modulating vector* $\mathbf{W}$ *such the product* $\mathbf{W} \cdot \mathbf{LV_0^t}$ *gives an encoding of* LV_0^t *that is directly relevant for the present navigational task.*

$$\mathbf{W} = \mathbf{Filter}(\mathbf{s}, \mathbf{LV_0^t})$$

It is important to apply *Filter* to the local view before using it further in order to update the PI state or to learn action-selection, as it reduces the input feature space significantly by removing the irrelevant features in the environment. Of course, learning of the *Filter* function itself is very difficult in the general case.

From a connectionist point of view, *Filter* can be thought of as an inhibitory set of connections that suppresses the activity of sets of cells whose firing is redundant to the required sensory information for the task at hand. In other words, it is the regulatory network that does filtering of irrelevant information and contributes towards the stability of the network by decreasing the sum of excitation present in the network that could lead to unbounded oscillations because of excitatory lateral connection usually found in almost all brain areas.

4.1.2 The PI state

Let us now define what we mean by the PI state.

Definition 4.4 (PI State) *A PI state* $\mathbf{s}$ *is a tuple* $(\mathbf{s_p}, \mathbf{c})$ *where* $\mathbf{s_p}$ *is the state vector denoting the current position of the agent in the environment relative to an allocentric reference point, while* $\mathbf{c}$ *is the present context.*

The context is a combination of the task at hand and the current environment of the agent. Let us therefore split the context $\mathbf{c}$ into two parts, namely, the task specific context $\mathbf{c_t}$ and the environment specific context $\mathbf{c_e}$. The environment specific context is simply the encoding of the environment that the agent is currently in, out of the possible memories of environments it has access to.

Task specific context, on the other hand is more dynamic in the sense that it changes while the task is being performed. The agent needs to have access to information about the amount of task completion, the task at hand as well as the states of different drives like hunger, sleep etc. All these parameters are required to be constantly tracked in order to update task specific context.

In the following sections, we shall use the definitions described above to define some operations that need to be performed in order to store new information about places and at the same time keep a coherent trace of the pre-existing memory of an agent's environment.

At the same time, it is of prime importance to learn how to make the appropriate action in order to take the agent to it's desired goal. In the succeeding parts of this chapter we shall briefly discuss this issue.

In the case of rodents, it is by now accepted that the PI is distributed across various brain sections. As described in Chapter 2, there exist cells in the post subiculum, anterior thalamic nucleus, and retrosplenial cortex that show high firing frequency correlation with the allocentric direction as well as angular velocity of the animal's head. The latter suggests that the inputs to these cells must be from the vestibular system, as the vestibular system tracks the angular acceleration of the head apart form other parameters related to the animal's movements. Furthermore, some of these cells also show a preference for left turns as opposed to right turns or vice-a-versa. The other cells that are very widely studied are the pyramidal cells in the area CA of hippocampus that show strong dependence on place of the animal in its environment. These cells are commonly referred to as place-cells. Also, as seen in Chapter 2, these cells show an elevated firing rate only when the animal is at certain places in the environment. This is true even in environments that show symmetry; the place-cells fire only at a geometrically constrained places and mostly do not correspond to the symmetry of the environment. For example, an overwhelming majority of these place-cells fire only at one place instead of four even when the environment has four fold symmetry. Even across sessions, if the point of entry to environment is kept constant with a symmetrical environment, these place cell firing fields remain essentially constant. This means that the hippocampus either receives PI signal that breaks the symmetry of firing fields, or it is capable of generating its own PI signal using the head-direction and other vestibular signals. Figure 2.2 shows that hippocampus has outgoing excitatory connections to subiculum which in turn projects to post subiculum where head direction cells are in abundance. The other major incoming pathway to the hippocampus is through the fornix which conveys signals from the brain stem. All this evidence strongly implicates hippocampus to be a PI system or at least a major component thereof. For a stimulating discussion of the hippocampus as a PI system, refer to [MBG$^+$96].

Also, the place-cell firing fields have completely different shapes and sizes in different environments. Even within an environment, for linear or radial-arm mazes, the place-cell firing fields have directionality; they fire with higher frequency when the animal travels in one direction of the arm as opposed to the other. With environmental manipulations, it is also possible to change the firing fields of place-cells even during a single trial, regardless of the type of the environment [ON78, MK87, SKM90]. It is also found that in open environments, either all cells keep their firing fields or all cells change their fields simultaneously.

All these observations can be explained by accepting the existence of a distributed code for the context which is represented across the hippocampus. We can then maintain that the PI state and therefore the hippocampal attractor network state encodes the context as well. The context changes with environmental manipulations (dynamic re-mapping) or with the stages of task completion (directionality in radial arm mazes and re-mapping once reward is achieved in food hunting or systematic search tasks).

With the above in mind, we define the context as follows.

Definition 4.5 (Context) *The context* $\mathbf{c}$ *is a tuple* $(\mathbf{c_e}, \mathbf{c_t})$*, where* $\mathbf{c_e}$ *is the environment specific context, or the encoding that represents the current environment, while* $\mathbf{c_t}$ *is the internal* *state specific context.*

In the above definition, the internal state specific context encodes the current state of the agent's drives (e.g., hunger, thirst, sleepiness, ...) as well as the amount of completion of task so far. In that sense, internal state specific context is a time-dependent context.

4.1.3 The cognitive map

Using the above definitions of the PI and LV, it is now possible to define the cognitive map. The requirements for a cognitive map were discussed in Chapter 3. Briefly, the cognitive map should at the very least be useful in performing at least two operations. First, using the LV and possibly task specific context, it should be possible to find the complete PI state.

Second, it should be possible to fix the environment specific context. Simply put, the second requirement states that using the LV and the task specific context it should be possible to decide which of the possible environments the animal is currently in. For the purpose, we define the cognitive map as the mapping from LV to PI along with a pattern completion algorithm that completes the joint vector $LV \times PI$ given any partial information about LV and context.

Definition 4.6 (Cognitive Map) *Cognitive Map* $\mathbf{M}$ *is a system with two functional components:*

1. *Mapping: It is a mapping* $(Filter(LV_0^t), \mathbf{c_t}, \mathbf{c_e}) \rightarrow \mathbf{s}$
2. *Completion: It is a completion device that returns* $\mathbf{c_e}$ *or* $\mathbf{c_t}$ *using the LV and task specific context.* $(Filter(LV_0^t), \mathbf{c_t}) \rightarrow \mathbf{c_e}$

For the remainder of this chapter we shall use the definitions given so far to define operations on the cognitive map for acquisition and representation of new information about one's environment. We divide the learning and retrieval operations on the cognitive map in two parts; those that use a single representation of the environment (intra map operations), and those that require simultaneous access to representations of multiple environments (inter map operations).

In the discussion to follow, LV will denote the filtered local view, be it the finite or the infinite local view. The operations themselves will not vary with the type of LV used, but the learning algorithms and the underlying hardware will change widely based on the type of LV.

4.2 Intra Map Operations

4.2.1 Operations on LV

The major learning operation on LV is the one that associates the LV with the current PI state $\mathbf{s}$. We define the association as follows.

Definition 4.7 ($LV \times PI$) *The association (or mapping)* $LV \times PI$ *relates a PI state* $\mathbf{s}$ *to a local view LV and vice-a-versa, such that, given the LV measurement a set of possible PI states* $\mathbf{s}$ *can be determined. Similarly, given the PI state a set of possible LV measurements can be determined.*

An example implementation of such a $LV \times PI$ association is a bi-directional associative memory [Mar71, Has95]. It is always useful to incorporate any partial information available about the entity to be predicted using the $LV \times PI$ association. In order to utilize the association between PI and LV efficiently, we need an algorithm for effectively searching or determining the missing quantity or quantities using the available information.

We define the completion function that performs this operation as follows. Here, $\mathbf{s}'$ and LV' are the PI state and LV measurements available at a given place and time. Of course, some of the attributes within these quantities could be missing, or one of the quantities could be completely missing. For example, upon re-entry in an a-priori known environment, the $\mathbf{s_p}$ defined in Definition 4.4 could be missing or partially available. On the other hand, in order for the agent to keep track of the correctness of it's actions, the agent needs to correctly predict the LV and compare it with the incoming LV measurements. In such a case, the agent would decide to trust it's PI state in order to find what to expect from the environment. As we shall see in the following chapters, such predictions of PI state and/or LV observations can be used to refine the $LV \times PI$ mapping itself. Such refinement or update operations can help reduce measurement or inherent system noise.

With the above in mind, we define the pattern completion device as follows.

Definition 4.8 (Pattern Completion Function) *The pattern completion function, Complete, returns the tuple* $(\mathbf{s}, \mathbf{LV})$ *using the available, possibly partial, information about the PI state* $\mathbf{s}'$ *and LV measurement* LV'*:*

$$(\mathbf{s}, \mathbf{LV}) = \mathbf{Complete}(\mathbf{s}', \mathbf{LV}')$$

At the same time, whenever complete information about the PI state and LV is available, the agent needs to remember the $(\mathbf{s}, \mathbf{LV})$ pair for future calls to *Complete* function. Let us call the operation that remembers the $(\mathbf{s}, \mathbf{LV})$ pair the *Associate* function.

Definition 4.9 (Learning Association) *The Associate function receives the complete tuple* $(\mathbf{s}, \mathbf{LV})$ *as input and incorporates it into the current cognitive map* $\mathbf{M_t}$ *to return the updated cognitive map* $\mathbf{M_{t+1}}$*. After the associate function is called, subsequent calls to Complete in similar situations but with missing information can return more reliable completions of missing information:*

$$\mathbf{M_{t+1}} = \mathbf{Associate}_{\mathbf{M_t}}(\mathbf{s}, \mathbf{LV})$$

In the above definition, the subscript of *Associate* function denotes the representation or the state of cognitive map at the time of call to Associate function. The strategies used to associate PI state to LV can be drastically different based on the richness of representation required for the current context as well as the amount of information already stored in the cognitive map for the present task-specific context.

Apart form the *Associate* function that links the current PI state to LV, we also need to learn what kind of sensory information and how deep a history of sensory perceptions is really required for effectively completing the missing information regarding sensory information or PI state during subsequent calls to *Complete.* In other words, we need to learn the *Filter* function.

We define operation *LearnFilter* that accomplishes this task. Based on the current context as well as the current status of the cognitive map, *LearnFilter* figures out what weights for LV should be produced by the *Filter* function. We define the Filter learning operation as follows:

Definition 4.10 (Filter Learner) *The Filter learning function LearnFilter takes in the current task-specific context and the current state of cognitive map and updates the Filter function. This update is such that the updated Filter function would assign higher weights to LV attributes that matter more for the given* $\mathbf{c_t}$ *and* $\mathbf{c_e}$*.*

In the above definition of Learning Filter has to evaluate how prominent certain LV attributes are for a given environment and task as a whole and not for a given PI state. This requires a constant evaluation of the agent's progress in the environment and at the same time a constant observation of how good the LV attributes are as predictors of PI state. Of course, the learning of Filter is not is not restricted to the spatial learning system or even to the hippocampal formation. In order to evaluate the progress and success of the navigational task, we need to assess the rewards received at the end of the task as well as while the task is performed, and therefore other regions that specialize in prediction and evaluation of rewards are also be involved in the process.

4.2.2 Operations on PI

As the agent navigates in the environment it's position and the amount of completion of task changes. Therefore, we define another operation; which we will call *StateUpdate*. This operation returns the new PI state $\mathbf{s}'$ based on the history of PI states as well as the actions performed by the animal. If LV information is available to the agent while navigating, that information is also taken into account. Let us call the history, most probably finite, of PI states as $\mathbf{s_0^t}$. Here, time 0 is the time when PI was first reset, and time t is the time immediately prior to the adjustment operation. We define the *StateUpdate* operation as follows:

Definition 4.11 (PI State Update) *PI state update operation uses the available PI state history* $\mathbf{s_0^t}$*, the set of actions performed since the last PI state update (denoted by* $\mathbf{a}$*) and* *available LV information to return the next PI state* $\mathbf{s}'$*.*

$$\mathbf{s}' = \mathbf{StateUpdate}(\mathbf{s}_0^t, \mathbf{a}, \mathbf{LV})$$

To give a rough example of the *StateUpdate* operation, as the animal moves in its environment, the *StateUpdate* operation takes the stream of LV measurements coming in, as well as vestibular outputs which encode the action performed, to update its own position estimate in the environment as well as to an extent, the present task dependent context itself. It is crucial to appreciate the importance of the *StateUpdate* operation, as it potentially can help reduce the errors in observations and the previous PI state that also contains the context and information about the current environment. At the same time it functions as an important element in the mechanism that tracks the progress of task.

Simultaneously with the *StateUpdate* operation, it is necessary to keep track of the correctness of the present context and the task dependent context. As we shall see in the next section, this requires across-map information and to some extent also some domain knowledge about the task being performed.

4.3 Inter-Map Operations

In the foregoing section, we looked at different operations that need to be performed in order to learn new spatial information, represent it, and utilize it in order to self-localize in a spatial environment. Thus far, we have completely neglected the existence of multiple environments to choose from.

4.3.1 Choosing a context

Upon introduction to an environment, based on the task at hand as well as the memory of already learned environment, context needs to be fixed. Fixing of context can be divided into two parts.

In order to decide $\mathbf{c_t}$, we require the previous state of the PI, as it conveys information about the previous position as well as the previous context. At the same time, we also require the incoming sensory information in order to decide which of the tasks at hand can be performed first in case there are multiple equally important tasks to choose from.

Let us define the operation of choosing these contexts *FindContext* as follows:

Definition 4.12 (Context Finding Operation) *The context finding operation FindContext returns the context* $\mathbf{c}$ *which is a combination of the task specific context* $\mathbf{c_t}$ *and the environment specific context* $\mathbf{c_e}$ *based on the previous PI state(s), LV and the agent's internal state I.*

$$\mathbf{c} = \mathbf{c_t} + \mathbf{c_e} = \mathbf{FindContext_t}(\mathbf{s_0^t}, \mathbf{LV}, \mathbf{I}) + \mathbf{FindContext_e}(\mathbf{s_0^t}, \mathbf{LV}, \mathbf{E})$$

Here, $FindContext_t$ and $FindContext_e$ denote mechanisms that decide the task specific context and the environment context. In the above expression, $\mathbf{I}$ denotes the internal state of the agent, e.g. hunger, thirst, sleepiness etc. The internal state largely determines what task the agent should perform first in order to satisfy these drives.

In the above expression $\mathbf{E}$ denotes the set of environments to choose from. In order to choose an environment from multiple choices, we require a search through the representations of all available environments using the LV and the previous PI state $\mathbf{s}$. The previous PI state is required in order to discard some of the environments that the subject is highly unlikely to be in. In other words, $\mathbf{s}$ imposes some continuity conditions that helps reduce the search space. Once a set of candidate environment is selected, the problem now reduces to finding an environment that best matches the current observation. For the purpose, the operation *Complete* described in the foregoing section can be repeatedly called with successive substitutions of $\mathbf{c} = \mathbf{c_t} + \mathbf{c_e}$. Here, $\mathbf{c_e}$ is unknown. In this fashion, $\mathbf{c_e}$ that returns a coherent and valid PI state is found. Sometimes it may not be possible to fix $\mathbf{c_e}$ straight away, in which case multiple hypotheses about $\mathbf{c}$ need to be entertained simultaneously and be discarded as more observations arrive.

It has also been hypothesized that the hippocampus is involved in prediction of amount of completion of task. As the experience of the animal on specific tasks like navigating in linear or radial arm mazes increases, the place cells start to fire more and more in anticipation of future places. In other words, the place cells drift such that their firing is more and more correlated to the animal's future place [Eic96]. This suggests that the hippocampus is not only involved in encoding, but also in *prediction*. Similar phenomena have been found in the anterior thalamic nucleus (ATN) head-direction cells of the rat. These cells who strong correlation with the head direction of the animal around 200ms in the future [BS95a].

4.3.2 Learning the context

It is evident from the foregoing discussion that it is required to learn an encoding of the environment that facilitates fixing $\mathbf{c_e}$ as quickly and reliably as possible. In other words the *FindContext* operation defined in Definition 4.12 in the previous section is required to be learned. This operation is such that it should learn multiple *environments*, rather than possible PI states within an environment. Let us define a function named *LearnEnvironment* that maps observations to an environment. The *LearnEnvironment* operation is very similar to the *Associate* operation defined in Definition 4.9 except that it updates the computation performed in *FindContext* such that correct context can be fixed with higher and higher reliability as training progresses. It is very likely that multiple, similar environments exist, in which case the function *FindContext* can be many-to-many.

We need to bear in mind that every time *FindContext* fails to return a context $\mathbf{c}$ that is consistent with the current $\mathbf{c}$ and which satisfies the continuity with previous PI state, it can be assumed that the subject is in a new environment context. In such a case *LearnEnvironment* needs to be called at the end of, or even continuously during, the learning episode that results in learning the new environment. The operation *LearnContext* needs the new $LV \times PI$ mapping and labels the mapping with the present environment context $\mathbf{c_e}$ and $\mathbf{c_t}$. This is precisely the mapping *FindContext* requires to fix the context upon re-entry to the environment. Hence, we arrive at the definition of *LearnEnvironment* as follows:

Definition 4.13 (Learn Environment) *The function LearnEnvironment takes as its arguments the current* $LV \times PI$ *mapping and a finite sequence of previous PI states* $\mathbf{s_0^t}$ *along with the set of candidate contexts returned by FindContext. Using this information, LearnEnvironment labels the* $LV \times PI$ *mapping with the current context* $\mathbf{c}$.

4.3.3 Merging environment contexts

It is difficult to acquire all possible observations experience all task specific contexts for a given environment in a single learning episode. During one exploratory episode for a given task context $\mathbf{c_t}$, only a subset of all possible observations can be learned. In such a case, upon re-introduction into the environment, it is not possible to recognize the environment correctly using the observation that has never been seen before. Also, it is possible that two task contexts are also equivalent but the agent may not realize so until the end of the task, when it finds that the $LV \times PI$ mapping learned for the two different contexts were terribly similar.

With the above in mind, we need an operation that will recognize the equivalence of two contexts and modify $FindContext$ so that upon subsequent visits to the same context $\mathbf{c} = \mathbf{c_t} + \mathbf{c_e}$ the context can be correctly recognized. We call this the context merge operation, which modifies the $LV \times PI$ mapping discussed above in a way similar to $LearnEnvironment$. The $Merge$ operation is in the re-organization operation that consolidates multiple memories of the same environment or task context into one.

In the case of merging of task contexts, this can be thought of as acquisition of new procedural skills from declarative skills by classifying all these different tasks into equivalence classes.

We define the $Merge$ operation as follows.

Definition 4.14 (Map Merge) *The $Merge$ operation takes in a set of $LV \times PI$ operations that are contextually equivalent and returns a single $LV \times PI$ mapping along with a new encoding of context* $\mathbf{c_new}$ *that represents all the old, equivalent set of contexts.*

Once the $Merge$ operation is performed, subsequent calls to $FindContext$ return the new context whenever observations of LV, agent internal state and previous PI states for any of the old, equivalent contexts are encountered. For the $Merge$ operation to work, a similarity measure between multiple $LV \times PI$ mappings needs to be performed. This problem is one of the hardest to solve as it requires comparisons between newly learned environment for which encoding strategies and important features for encodings were hitherto unknown.

4.4 Selection of Actions

The ultimate aim of all of the foregoing learning is to produce behavior that enables the animal or mobile agent to correctly perform the task at hand. This requires learning of a mapping from the PI state to an action that leads to the goal. It should be remembered here that the goal depends on the task context. Therefore, the behavior to be learned here can be denoted as:

Definition 4.15 (Action Selection Mechanism)

$$\mathbf{a} = \mathbf{Action}(\mathbf{s})$$

Where **a** is the action or a sequence of actions that can lead the subject to a goal, for example, food at the end of the sequence of actions.

It is not clear whether the action should be selected continuously using the *Action* operation, or whether the *Action* operation simply returns a program P to be executed in order to reach a goal state. If a program is selected by *Action* and then executed separately, then there needs to be a program state **p**. The program execution can then be defined as a sequence of program states thorough which the behavioral component of the system needs to pass through to complete the behavioral task. In any case, we need another operation that keeps a check on the execution of the actions and ensure that the actions are performed correctly with respect to the task at hand. If the actions are performed incorrectly, these can be tracked by checking the trace of PI states with the sequence of PI states that the system needs to go through in order to complete the task. We define an *CorrectAction* operation that takes in the current observation, PI state, the action to be performed **a**, and the current state of the behavioral program **p** and returns correction $\delta\mathbf{a}$ to the action **a**.

At the same time, we also need a mechanism that learns the program itself. Most probably this is a completely separate system that is discussed extensively in the literature [RM86, KLM96, BS81].

CHAPTER 5. EXISTING COMPUTATIONAL MODELS

In this chapter we review some of the major computational modeling efforts for the function of the hippocampal formation. These models take some of the standard computational paradigms to propose the function of the hippocampal formation at various conceptual levels. Most of the models described here either explain some subsets of experimental evidence from experiments or rodents, or they explain the function of the hippocampus at a very abstract, general level. The abstract, general level modeling usually limits itself to storage and retrieval of arbitrary vectors and do not usually explain individual behavioral experiments. A more detailed description of hippocampal modeling efforts can be found in Trullier et. al. (1997) . As we discuss below, the modeling efforts for spatial learning and navigation can be classified into three major classes: Place-Response, Topological, and Metric based navigation models. We discuss some important models in each of these classes below.

5.1 Place-Response Based Navigation Models

In Place-Response based navigation models, each place is associated with an action that leads the animal closer to a desired goal location. One of the earliest models using this approach was developed by Zipser (1986) . In his model, Zipser proposed a model in which each of the simulated place cells was tuned to a specific landmark in the environment. The place cells were also directly linked with a goal cell which encoded a vector to goal location from the place where the given stimulus and consequently firing of the place cell would occur. In a population code version of the model a number of such place cells, each encoding an individual landmark or a subset of landmarks specify a place. A weighted average of goal cell encodings would then be used to navigate to an estimated goal location.

Brown and Sharp (1995) proposed a very similar model where the hippocampal place cells and the Subicular head direction cells converge to two clusters of cells in the nucleus accumbens (NA). Each cluster in NA in their model corresponds to a turn direction and the place cell and head direction cell connectivity and firing pattern directly maps to a turn command. Availability of a reward results in increase of synaptic strength in recently active connections, or in other words, activity dependent LTP is used for goal memory incorporation.

The major problem with the models described above is that multiple goal locations cannot be represented. Also, these models are highly susceptible to interference between representation of multiple environments in the same neural region. There is strong evidence of involvement of the same hippocampal pyramidal cell in representation of multiple environments [MBG$^+$96]. Apart from that, latent learning or pre-learning an environment in the absence of a goal location is not possible.

The model of Burgess and O'Keefe relaxes the severity of these problems significantly [BRO94, BO96]. Their model hypothesizes the function of the hippocampus, entorhinal cortex cells, head direction cells in the subiculum and has a set of goal cells which is distinct from the rest of the cells in their model. Their model uses winner-take-all type of learning in place cells and subicular cells, and a form of Hebbian or activity dependent learning between place cells, subicular cells, entorhinal cortex cells, and goal cell groups. They also use a spiking model of neuronal activity to model the EEG θ-rhythm dependent modulation of cell firing observed in the hippocampus. They hypothesize the existence of goal-cells downstream of the hippocampus, which are assumed to be tuned to different goals and updated based on reinforcements obtained at the goal locations. A goal is chosen based on the place cell and head direction cell activity as well as the entorhinal cortex cell activity which is meant to provide some sort of a context for accomplishing the navigational task.

5.2 Topological Navigation Models

The Topological Navigation scheme gives emphasis to storage of relations between encodings of places. In such a scheme, places can be represented in the form of directed graphs, with each link labeled by the action to be taken to reach a place from its immediate neighbor. In order to reach the goal, animal must perform a graph search from the current node to a specific node which can be labeled with a reward. The links can then be used to determine the actions to be performed in order to reach the reward. It is easy to see that this type of learning can incorporate latent learning, or learning in absence of explicitly defined goals and rewards, as well as can represent multiple goals. But it is unclear how well such models can predict the ability and computing time for planning to reach the goal from a behavioral and reaction time standpoint. Also, the models are not very scalable.

The model proposed by Muller et al. (1991) encodes the spatial map as a graph. In their model, place cells are allotted before the simulated animal (animat) is even introduced into the environment and are connected together by links that can be modified using synaptic plasticity. As the animat moves from place to place, these place cells fire in specific sequences and correlated activity dependent LTP links these place cells to each other. In such a fashion the animat can remember relationships between places and also subsequences of places visited over its training period.

The model of Schmajuk and Thieme (1992) also have proposed a very similar model. The difference between Schmajuk and Thieme model and Muller model is that the former model also has a Hebbian learning that associates these pre-configured place cells with another set of cells that they call “view nodes”. These view nodes are supposed to encode the sensory information at the particular place which is encoded by the place cells. Furthermore, their model also has a fast component that runs pre-learned routes in forward direction, thus determining the course of action to reach the goal. The slow component on the other hand simply predicts the next place based on the current place and the current intended motor command.

5.3 Metric Based Navigation Models

Although there is no clear-cut delineation between Topological representation and Metric based Navigation, modeling efforts in the literature have made a distinction between the two. This is mainly due to the manner in which these models are wired. It should be noted that the topological representation can easily be transformed to a metric based representation by utilizing the motor commands associated in order to reach one place from another. Similarly, metric based navigation can be transformed into a topological representation by generating motor commands and linking different places with such motor commands. Furthermore, there is no decisive way of finding which of the two kinds of encodings are really utilized in the brain. Perhaps it is a mixture of both as both have their own advantages and disadvantages. Topological representations are faster to learn and consolidate as routes to goals which can be quickly retrieved, executed and learned. On the other hand, Metric based representation has compactness of representation and flexibility in terms of computing detours and short cuts easily. Also, it is significantly easier to learn new reward locations in the Metric based representation as goals can simply be represented by their positions with respect to an allocentric frame of reference.

Prescott (1992) has developed a model of metric space representation in terms of local frames defined by subsets three landmarks each. By representing the position of a fourth landmark in the local frame, his model allows the creation of a database of relations between landmark locations. This model implicitly encodes the metric positions of the places, since any arbitrary goal location can be determined by a set of intervening local frame transformations that map the goal frame to the frame that the animat is in. Navigation to a goal requires the activation of frames that contain the goal location. Each of these active frames predict the relative position of a fourth landmark. These landmarks are then added to the list of visible (or available) landmarks and all frames activated by this new set of landmarks are retrieved from the relational database. This process is repeated until a frame containing the animat's current location is activated. At this point the set of frame transformations from the current frame to the goal frame can be used by the animat to navigate. Prescott also extended this approach to automatically determine the shortest trajectory to the goal.

The model used in this thesis to verify some behavioral experiments is closest to the one proposed by Redish and Touretzky [WTR94, RT96]. In their model, places are represented in terms of an ensemble of active place cells. Each place cell is tuned to the identities and egocentric positions of two randomly chosen landmarks visible from the current place. Response of the place cells also depends on the retinal angle, which can be taken as the visible size of the landmarks, and an internally generated path integrator estimate of the current location of the animat. Also, there is a certain degree of flexibility in what information is considered for determining the location of the animat. If the animat is reintroduced in the environment, it should not consider the state of the path integrator in its estimate of position. In this case, the place cells fire only in response to sensory inputs. A major difference between the Redish and Touretzky model and the model to be described later in this thesis is that the Redish and Touretzky model is capable of also resetting the head direction or angular orientation of path integrator while our model assumes that the head direction is available to the model by some external means. On the other hand, our model explicitly takes into account the sensor and movement errors and reduces these additive errors using a linear Kalman Filter [Kal60] scheme. It has been found that the place fields drift in the absence of visible landmarks. Such drift is supposed to be due to movement errors. Introduction of Kalman Filter like updates of position estimate, or the path integrator, and the estimates of place field centers in the current frame of reference effectively prevents such drifting in the presence of visible landmarks and at the same time allows the drifting in the absence of visible landmarks. The model used in performing the behavioral experiments is described in Chapter 6.

CHAPTER 6. COMPUTATIONAL MODEL FOR RODENT SPATIAL NAVIGATION AND LOCALIZATION

In this chapter we present the computational model for rodent spatial navigation and localization developed by Balakrishnan Et. Al. (1998) [Bal99, BBH97]. We describe the behavioral experiments performed using this model in the next chapter [BBH98b, BBH98a]. The model is based based on the locale system hypothesis suggested by O'Keefe and Nadel (1978) . The main idea behind locale system hypothesis is that landmarks, or more generally objects, are represented as relations to each other in terms of their relative positions. Such representation is supposed to arise due to a fusion of incoming sensory information of different modalities with the path integrator estimates of their positions with respect to an allocentric frame of reference. This allocentric frame of reference is supposed to be dependent on the type of environment, and independent of the current location of the animal in its environment.

6.1 The Basic Computational Model

The main idea behind the model discussed in this chapter is that the sensory measurements as well as the path integrator estimates are prone to noise. Effects of such noise on the path integrator can be seen in the drifting of place fields in darkness [MBG$^+$96]. The model discussed below explicitly addresses the errors in measurement and path integration estimates and effectively reduces these errors by applying a Kalman Filter [Kal60] like update mechanism.

Based on lesion and pharmacological studies performed by numerous experimenters, O'Keefe and Nadel (1978) suggested that hippocampus is critically involved in the formation of a "cognitive map". The cognitive map in their hypothesis is a representation of objects in an animal's environment which is formed by storing spatial positions of different objects relative to each other.

It is interesting to note that there is growing evidence that some storage of relative temporal positions or sequences of events are also stored in the hippocampus [NM97, SM96, Eic96].

Balakrishnan et. al. (1997) have developed a computational specification of this locale hypothesis that learns to encode landmarks in an environment and their spatial relationships with each other. This encoding can be utilized by an animat to self-localize, or find its position, in a familiar environment upon re introduction and then compute a trajectory to an estimated goal location [BBH97, BBH98d, BBH98c, BBH98b, BBH98a]. The system learns a set of distinct places in the environment and labels the center of each place with metric position estimates derived from the path integration system. This fusion of sensory and dead-reckoning information takes place in a functional model of the hippocampal formation.

The overall functioning of the model based on this specification is as follows. During exploration of its environment, the animat is supplied with sensory information about each distinct landmark in its environment. The animat estimates the distance and orientation of each landmark from its present position. For a new place visited, this information is used to allocate the first layer cells in the model. These cells correspond to the cortical layer II and III entorhinal cortex (EC) cells. Each unit in this first layer (henceforth EC layer) of the model responds to a specific type of landmark at a specific relative position to the animat. The activation of these cells is of a Gaussian shape centered at the first observed relative position of the landmark of that particular type. This is consistent with the observed properties of EC units in rodents [QMKR92]. If none of the current EC layer cells fires in response to a given landmark, a new EC layer cell tuned to that landmark is inducted in the model.

It is reasonable to assume that each place in the environment can be described by the way landmarks can be observed from that place, provided the landmarks are stable with respect to each other. Therefore, the concurrent activity of EC layer cells can be utilized as an encoding of relations between places. With a two layer network of cells it is easy to see that conjunctions of literals, truth of each of them denoting landmark presence at specific locations relative to animat, can be easily represented. Thus the second layer cells in the model, which correspond to the CA3 layer cells in the hippocampus, encode conjunctions of landmarks at specific relative positions to the current location of the animat. In other words, CA3 units encode the "snapshot" of the local view of the environment. EC layer cells simply encode features which in this case are the presence of absence of landmarks of specific type present at a specific relative position. Such an encoding has nothing to do with the semantic of the place itself. The firing of CA3 cells in response to sensory inputs at a given place thus constitutes an internal place-code for that place. During training if no CA3 cells fire above a predefined threshold in response to a given sensory input, a CA3 cell is inducted and tied to the currently active EC layer cells. The connection weight between the newly inducted CA3 cell and the active EC layer cells is proportionate to the activity level in the EC layer cells.

The path integrator is simply a two dimensional Gaussian random process with its mean updated at every motor command by an amount that is equal to the translation performed by the animat. This update is corrupted by a zero mean Gaussian noise to reflect the inaccuracies in the motor command. The mean of the path integrator is represented in Cartesian coordinates with an associated diagonal covariance matrix.

It has been proposed by researchers that CA3 layer cells work as pattern completion devices in situations where limited or corrupt sensory information is available [Rol96]. Other researchers believe that the pathway formed by direct connection from EC to CA3 and the parallel pathway from EC through dentate gyrus (Dg) onto CA3 combined perform this pattern completion-prediction role (Correspondence with Ali Minai).

The model also assumes that goals are learned and represented in terms of their metric positions. However, this goal representation is believed to reside outside the hippocampus. This is in accordance with the existence of place cells during pre-training when there is no reward present.

In addition to learning places by appropriately creating EC and CA3 units, each sensorily new place is also labeled by the current path integrator estimate. This is done by creating a new CA1 layer unit, linking it with the currently most active (or winner) CA3 layer cell, and labeling it with the metric position label from the path integrator. The metric label is simply the distance and angle encoded in Cartesian coordinates from the place where the path integrator was initialized, which in most cases is the point of introduction of the animat into the environment.

Each of the CA1 unit labels can also be considered a random variable with an associated covariance matrix, where covariance is computed between all other CA1 unit labels.

When the animal visits familiar places, incoming sensory inputs excite EC cells which in turn activate a place code in the CA3 layer. As multiple CA1 place codes may correspond to this CA3 code due to multiple places being sensorily the same, the path integrator estimate is used to determine the CA1 unit that represents the closest previously visited place. In the current implementation of the model a Mahalanobis distance test [DH73] is used to determine which CA1 layer unit is the closest match to the current path integrator estimate. The model then performs spatial localization by matching the predicted position of the animal which is the path integrator state with the observed position of the place field center, or the position given by the label of the closest currently active CA1 unit in the sense of Mahalanobis distance. A Kalman Filter like approach is used to update the current path integrator estimate as well the labels associated with all CA1 units in the current training episode. This is shown in Figure 6.1. The meaning of current training episode will become clear in Section 6.3

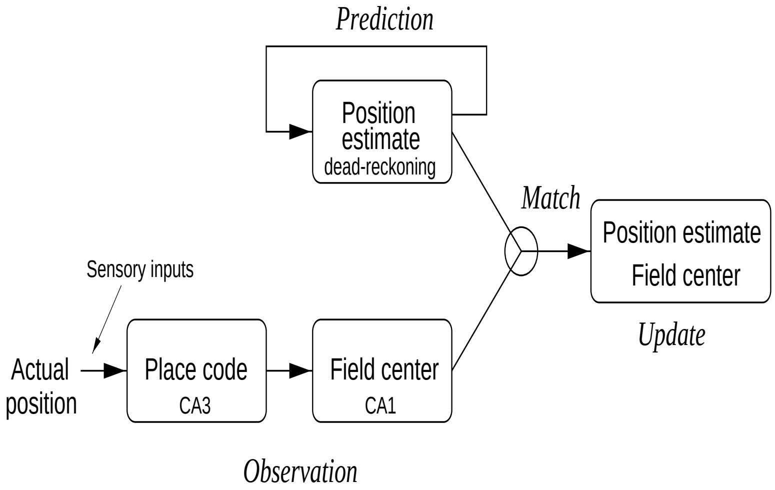

Figure 6.1 A schematic of hippocampal localization.

6.2 Learning and Updating Goal Locations

Since our computational model allows the animat to learn places in a metric framework, goals encountered by the animat can also be remembered in terms of their metric positions. This idea of computing and remembering goal locations from metric place estimates was first developed in [RT96]. In our work [BBH98a], the animats compute goal positions based on their current dead-reckoning estimates and the estimated distance to the goal. Thus, goals are represented in the same coordinate frame as the place field centers.

Since the goal locations are labeled in terms of path integrator state which is error prone, the position estimate of the goal location which is derived from path integrator is also prone to errors. In what follows we develop a mechanism of updating the estimated goal locations in a manner that minimizes the variance or the second order variability of the goal location.

Let G be a goal. When the animat visits the goal location for the first time, say in step j, the position estimate of the goal $\hat{x}_G$ is initialized to the current path integrator estimate of the animat $\hat{x}_{0,j}$. The variance of the goal estimate $\mathbf{C}_{GG}$ is set to the current variance of the animat's position estimate $\mathbf{C}_{00}$. Suppose the animat visits the same goal location again at time step k. The position estimate of the animat $\hat{x}_{0,k}$ may differ from $\hat{x}_G$ owing to dead-reckoning errors in the intervening animat motions. Further, one of these may be more accurate than the other, and the system must then appropriately update the goal position estimates taking their relative accuracy into account. Let us assume that the position $\hat{x}_G$ is updated using a linear combination shown in Equation 6.1.

$$\hat{x}_G^+ = \alpha.\hat{x}_G^- + (1-\alpha).\hat{x}_{0,k} \tag{6.1}$$

The value of α is determined by the one that minimizes the resulting variance of the goal estimate under the simplifying assumption that $\hat{x}_G^-$ and $\hat{x}_{0,k}$ are statistically independent.

$$\begin{aligned} Var(\hat{x}_G^+) &= \alpha^2 Var(\hat{x}_G^-) + Var((1-\alpha).\hat{x}_{0,k}) \qquad (6.2)\\ \mathbf{C}_{GG} &= \alpha^2.\mathbf{C}_{GG} + (1-\alpha)^2.\mathbf{C}_{00} \qquad (6.3)\end{aligned}$$

where $\mathbf{C}_{GG}$ is the variance of the goal position estimate while $\mathbf{C}_{00}$ is the variance of the animat's current dead-reckoning estimate. Minimizing the resulting variance,

$$\frac{\partial \mathbf{C}_{GG}}{\partial \alpha} = 0 = \quad 2.\alpha.\mathbf{C}_{GG} - 2.(1-\alpha).\mathbf{C}_{00} \tag{6.4}$$

$$\Leftrightarrow \quad \alpha = \quad \frac{\mathbf{C}_{00}}{\mathbf{C}_{00} + \mathbf{C}_{GG}} \tag{6.5}$$

Thus, the value of α given by Equation 6.5 minimizes the resulting variance (it can be easily verified that $\frac{\partial^2 \mathbf{C}_{GG}}{\partial \alpha^2} > 0$). Using this expression for α in Equation 6.1 leads us to the following update rules for goal position estimate and variance:

$$\hat{x}_G^+ = \frac{\mathbf{C}_{00}}{\mathbf{C}_{00} + \mathbf{C}_{GG}} \hat{x}_G^- + \frac{\mathbf{C}_{GG}}{\mathbf{C}_{00} + \mathbf{C}_{GG}} \hat{x}_{0,k} \tag{6.6}$$

$$\mathbf{C}_{GG} = \frac{\mathbf{C}_{GG}\mathbf{C}_{00}^T}{\mathbf{C}_{00} + \mathbf{C}_{GG}} \tag{6.7}$$

These update expressions are applied each time the animat reaches the goal location. Goal variance $\mathbf{C}_{GG}$ is initialized to ∞. When the animat reaches the goal for the first time (say, in step j), the above expressions automatically set $\hat{x}_G^+ = \hat{x_{0,j}}$ and $\mathbf{C}_{GG} = \mathbf{C}_{00}$. Thereupon, every visit to the goal location results in an update of the goal position estimate and its variance, based on the relative uncertainties in the goal and dead-reckoning estimates. This allows the animat to maintain reliable goal position estimates. It should also be pointed out that if goals are encountered in a particular frame f and if at a later point f is merged into another frame, the goal position estimate and its variance must be appropriately transformed into the other frame as will be described later.

6.2.1 Representing multiple goals

It is also possible to represent multiple goal locations in a manner similar to the one described above.

Assume that the animat reaches a goal location at at step k, at which point the path integrator estimate is $\hat{x}_{0,k}$. If the animat has a record of a goal visit and labeled its representation by an estimate $\hat{x}_G$ with variance given by $\mathbf{C}_{GG}$ a Mahalanobis test is performed as follows:

$$(\hat{x}_G - \hat{x}_{0,k})^T(\mathbf{C}_{GG} + \mathbf{C}_{00})(\hat{x}_G - \hat{x}_{0,k}) < \epsilon \tag{6.8}$$

where $\mathbf{C}_{GG} + \mathbf{C}_{00}$ is the covariance matrix of this test and ϵ is an appropriately chosen distance threshold based on the χ^2 test. Here $\hat{x}_G$ is taken to be the prediction of the goal position while $\hat{x}_{0,k}$ is the observed position of the goal.

If this test is satisfied, the goal encountered by the animat is said to be the same as the one visited previously (G). However, if this test fails, the animat is at a new goal location. The animat then learns this new goal location using the algorithm described earlier and adds this to its goal memory. Note that any number of goals can be learned and remembered by the animat.

In case of a re-visit to a goal, animat updates the closest available representation of the goal using the update rule developed in Section 6.2

6.2.2 Navigating to goals

Once the animat has learned the locations of different goals it can navigate to specific goals as required. Two processes are involved in the realization of such goal-directed behaviors. First, from its memory of different goals encountered and represented, the animat must choose one goal location to navigate to. Second, once an appropriate goal has been identified, the animat must move in such a manner as to approach the remembered location of the goal.

It is not clear which kinds of goal selection strategies are used by animals under different situations. We use a goal selection scheme that uses a combination of recency of goal encounter, its closeness to the current position of the animat, and a measure of the confidence associated with the goal which is updated upon re-visits to that goal location. We choose between the foregoing three strategies using a random drawing.

In the following chapter we will demonstrate that the implementation of this computational model provides a satisfactory computational simulation of some behavioral experiments performed with rodents.

6.3 Merging Different Learning Episodes (frame merging)

The model also has a capability to learn different parts of the environment over separate, independent trials and then merge these "frames" or episodes of learning whenever a place perceptually common to the current and any of the previously learned frames is found. More details about this process can be found in [Bal99]. Briefly, over multiple trials, animats allocate different sets of EC, CA3 and CA1 layer cells upon visits to new places. For frame or episode f_{old} which is a previously learned episode, CA1 layer cells are labeled by path integrator estimates for the integrator initialized based on the point of entry for f_{old}. During another training frame or episode f_{new} which starts at a previously unvisited place and with path integrator reset with the new point of entry as start point newly allocated CA1 cells are labeled by this new state sequence of path integrator and clearly these labels do not match with the labels on group of CA1 layer cells allocated during f_{old} which has some other place as path integrator origin. Whenever a CA3 cell linked to both a CA1 cell (or cells) in f_{old} as well as a CA1 cell (or cells) in f_{new}, clearly the animat is at a place which is perceptually the same as a previously visited place. In such a case a coordinate transformation is performed on the CA1 cells allocated during f_{old} so that their labels become consistent with the path integrator estimate during the current frame of reference. In other words, the origin of the old frame f_{old} is shifted to align with that of f_{new}. Also, the covariance matrix is also updated as described in Balakrishnan (1999) .

CHAPTER 7. BEHAVIORAL EXPERIMENTS AND RESULTS

7.1 Experiments of Collett et al. (1986)

We simulated the behavioral experiments of Collett et al. (1986) using the computational model of the hippocampal spatial learning described in Chapter 6. The experimental setup of Collett et al. (1986) consisted of a circular arena of diameter 3.5 meters placed inside a light-tight black painted room. The arena also contained cylinders of 70 cm height and 11 cm diameter painted black or white as landmarks. Gerbils were trained to locate a sunflower seed placed at fixed locations relative to different configurations of landmark placement. The floor of the arena was such that it prevented the gerbil from spotting the seed from a distance. The arena was illuminated by a single white light hung directly above the setup. It was found that the animals learned to self-localize and move towards the estimated position of reward over a number of training trials. Once the animals were trained, the reward was removed from the environment and landmark configurations were changed in order to provide the animal with conflicting or incomplete information for localization. Upon re-introduction in such a changed environment, animals searched for reward at its expected position as computed using this incomplete or altered sensory information. Video tracking system was used to record the trajectories that the animals took. This information was then used to compute histograms of time spent by animals at different places in the environment [CCS86]. We reproduce the same experiments below using the computational model described earlier.

In our simulations, we used a circular arena of radius 10 units. The walls of the environment did not contain any sensory stimuli. The landmarks, on the other hand, were assumed to be visible to the animat from all points in the arena. The sensory information, in this case the distance and orientation of landmarks from the current position of the animat, was assumed to be corrupted by a zero-mean Gaussian sensing error with standard deviation $\sigma_S = 0.01$. Each landmark at a specific relative position caused an EC unit to fire. A simultaneous activation of EC units caused firing of CA3 and CA1 layers. The animat motion was corrupted by zero-mean Gaussians with $\sigma_M = 0.5$ units. The animats also possessed means for dead-reckoning with errors modeled as zero-mean Gaussians with $\sigma_D = 0.05$ units.

For each trial, the animat was introduced into the arena at a randomly selected position and was allowed to perform 500 steps of sensing, processing and moving. If the animat happened to see the reward within this time period, it was made to approach and consume it. Each animat was subjected to five such training trials. In each trial the animat created representations of environment by associating sensory information at different locations in the environment with its dead reckoning position estimate. These representations were independent of each other across trials. In case the animat received the same sensory information, or in other words visited the same place assuming no perceptual aliasing, the independent representations over different trials were merged together to form a single unified representation. This consolidation process will henceforth be called frame-merging.

Note that such map consolidations are valid only in the case when there is no perceptual aliasing. There is no obvious way to circumvent this difficulty in the present model. One way to solve the frame-merging problem is to store partial paths, or particular sequences of sensory information stream, and merge frames only when the animat sees a long sequence of sensory information instead of a single snapshot of the environment. Such a method is viable in case of radial or linear maze situations, where frames can be merged only when the animat sees a particular sequence of sensory information only when it goes in a certain direction; say up the linear arm as opposed to down the arm. This explains gradual emergence of directionality observed in hippocampal place cells within arms of a radial arm maze while the lack of the same in open field or central area of radial arms observed by [Eic96].

Once training was complete, the animat was subjected to ten *testing trials,* in which the landmarks in the arena were manipulated in specific ways which resulted in sensory information that was partial or conflicting compared to the sensory information available to the animat during training trials. Furthermore, and the goal or reward was removed from the arena. Here, the animat was released at random positions in the arena with its dead-reckoning variance set to ∞. In other words, upon re-introduction the animat was completely uncertain about its dead reckoning position estimate. Animats were only allowed to self localize and navigate but were not allowed to induct any new units. A localized animat was allowed a maximum of 300 time steps to navigate to the estimated goal position. Since the goal was not found even after searching for 25 time steps at the goal location, animat chose another goal location. This method is consistent with claims by [RT98] where authors have hypothesized three distinct modes of hippocampal operation; acquisition when new memory traces are stored, recall where animal retrieves stored memories, and consolidation where memories are probably stored into more long-term storage probably located in the cerebral cortex.

For the training as well as testing trials, the trajectories followed by the animats were recorded. Also, the arena was decomposed into cells of size 0.33×0.33 and a count of the amount of time spent by the animats in each cell was kept. A normalized histogram for five animats was then plotted.

We simulated the one, two, and three landmark experiments of Collett et al. (1986), and the search distributions of our animats (Figures 7.1, 7.2, and 7.3) match rather closely with those of the gerbils. The large dark squares in the plots denote the landmarks.

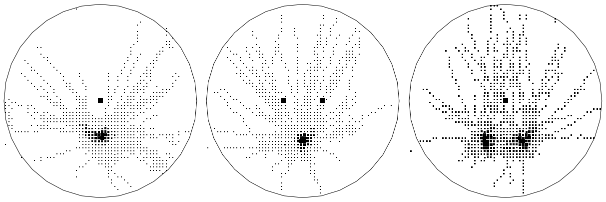

Figure 7.1 Left: One landmark experiment. Middle: Two landmarks experiment. Right: Two landmarks experiment with one landmark removed.

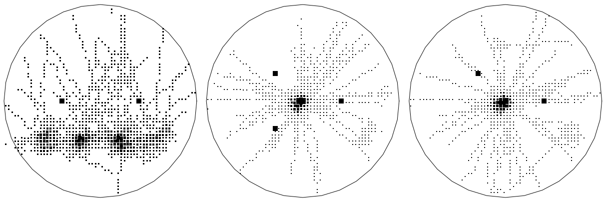

Figure 7.2 Left: Two landmarks experiment with landmark distance doubled. Middle: Three landmarks experiment. Right: Three landmarks with one removed.

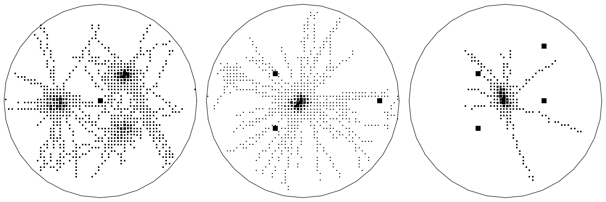

Figure 7.3 Left: Three landmarks with two removed. Middle: Three landmarks with one distance doubled. Right: Three landmarks with an extra landmark added.

7.2 Water-Maze Experiments of Morris (1981)

Morris (1981) experimented with male hooded rats of the Lister strain to demonstrate that rats are capable of rapidly learning to locate an object using distal cues.

A circular pool filled with opaque, milky water was used for the purpose. Objects present along the walls of the room served as distal cues. The pool was devoid of any objects except the escape platform. The escape platform was one of the following two kinds. First type was black colored, circular and protruding above the water, and therefore visible from a distance. Second type was white colored, circular and submerged in the water, thus virtually invisible.

The population of rats was divided into four groups of 8 individuals each. For *Cue + Place* group, the visible, black platform was used, which was always at the same location (NW, NE, SE, or SW) across all trials for a given rat. The second group was exactly like the first except that the white platform was used instead. This was designated the *Place* group. In *Cue-only* group, rats were trained using the black platform. However, in this case the platform was placed in one of the four positions, in an unpredictable sequence over trials. Finally, the

Place-Random group was similar to the Cue-only experiment except in the use of the white platform instead of the black one.

For each trial, the rats were released in the pool, and their trajectories were recorded along with the time taken to find the platform. Following 20 such trials over 3 days, the groups were further divided into subgroups of 4 individuals each. Each of these subgroups was subjected to 4 testing trials, of type A or B.

In Test A the platform was removed and search behavior was observed for 60 seconds. For Test B, rats of groups Cue + Place and Place were tested with the platform now placed in the quadrant diagonally opposite the one used in training. Rats of groups Cue-only and Place-Random were tested with the platform position held fixed. The escape behavior of the animals was then observed.

In our simulations, we used a circular arena of radius 3.75 units inside a square room measuring 20 by 20 units. Consistent with the ratio of pool and platform sizes in Morris' experiments, we chose the radius of our simulated platform to be 0.65 units. It was assumed that the animat could see the platform from a distance of 0.325 units. Four indistinguishable landmarks were used, one along each wall of the simulated room.

The sensing, motion, and dead-reckoning errors were same as in the foregoing experiments. We also assumed that rats swam slower than their normal walking speeds, and hence the size of motion step was set to 0.4 units.

As in the case of the original experiment, we allowed our animats four pre-training trials in which they randomly explored the environment for 100 steps without the platform present in the pool. During this stage, our spatial learning system allowed the animat to acquire a spatial map corresponding to the environment. In the training trials, the animats engaged in the goal seeking behavior. If the platform was not found at the particular goal location, the animat searched for 15 time steps before selecting another goal location and navigated towards it.

Groups of eight animats each were used in experiments corresponding to groups Cue-only, Cue + Place, Place, and Place-Random as in Morris (1981). The escape latencies for the first

20 training trials and the last four trials of Test B are shown in figure 7.4.

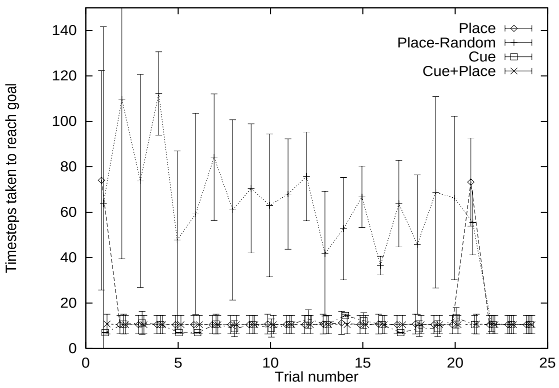

Figure 7.4 Escape Latencies while training and test B.

As seen in Figure 7.4, the Place group quickly learned the goal position. Furthermore, the Cue group achieved very small escape latencies. One reason for this is the fact that our simulation had a built-in mechanism to directly approach visible platforms from the start. Actual animals may not have such direct approach behaviors preprogrammed but may learn them with experience. Further, as with rats, our animats too perform poorly in the Place-Random experiment.

Figure 7.5 shows the paths taken during the first test trial by representative animats in different groups. Labels C+P, P, C and P/R denote group Cue + Place, Place, Cue-only and Place-Random respectively.

7.3 Discussion

The primary goal of the simulations was to test whether our computational model of hippocampal spatial learning and localization was capable of reproducing the behavior of gerbils. We simulated a number of experiments conducted by Collett et al. (1986) and by Morris (1981) .

It should be pointed out that our animats did not remember goals in terms of *independent vectors* to individual landmarks, as suggested by [CCS86]. Instead, places were remembered as independent vectors to landmarks, while the goal was simply remembered as a place.

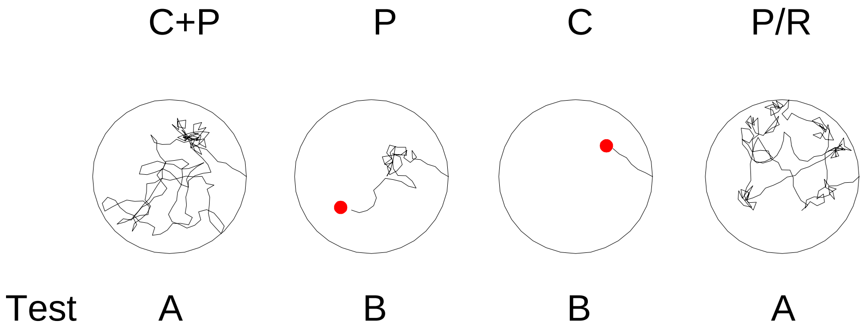

Figure 7.5 Trajectories followed by the animats (see text for detail).

In the process of simulating behavior, we identified an important issue, namely, how do animals choose one goal to approach from multiple ones that they might remember? Likelihood of storage of multiple goal locations in the hippocampus can be ruled out in the light of latent learning observed by researchers beginning from the early observations of Tolman [Tol48] and more comprehensively by Keith and McVety [KM88]. Storage of multiple maps based on the memory of reward location can more easily be ruled out in the light of fast acquisition of new reward places especially in Morris's Watermaze experiments [Mor81].

In order to simulate the Place-Random experiments of Morris, we had to incorporate a heuristic goal selection strategy. Our results using this mechanism closely parallel the behaviors observed by Morris. Indeed, our computational framework allows one to implement and test different hypotheses of goal selection. Such an approach can lead to a better understanding of goal selection processes in navigating rodents.

From Figure 7.6 it can be observed that Place and Cue + Place experiments indicate a strong spatial bias towards the training quadrant. While the former observation is consistent with the results of Morris, the latter is a surprise. However, this is a direct result of our spatial learning and navigation strategy, where we have assumed that the animat faithfully learns a

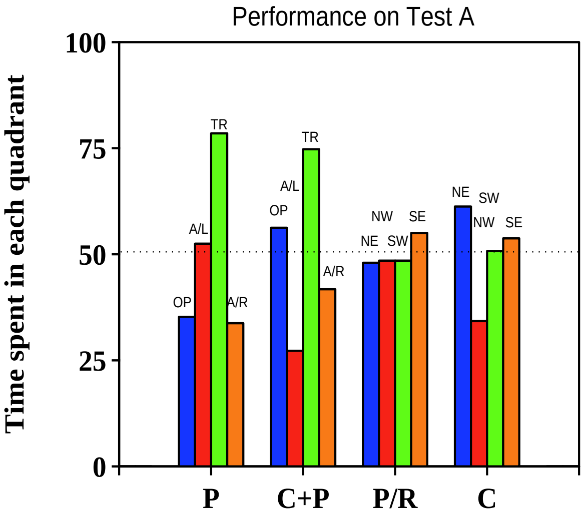

Figure 7.6 Performance on Test A. Histogram shows the duration of time spent by the animats in each quadrant. Here TR is the training quadrant, A/L and A/R are the adjacent quadrants to the left and right respectively, and OP is the opposite quadrant for groups Cue + Place and Place. The data from the other groups simply indicate the quadrant.

place map. There is a possibility that in the presence of reliable visual cues (e.g., platform), place learning may not be as reliable, since it is not even necessary. This hypothesis regarding differences in place learning in the presence or absence of reliable cues, remains to be studied.

7.4 Variable Tuning Widths of EC Layer Spatial Filters

It has also been found that the prominence given to landmarks in the environment depends on the animals recent experience and also, to some extent, the point of entry into the environment. Upon availability of extra cues that give multiple choices for fixing the Path Integration system, animals generally localize by giving more importance to the proximal as opposed to distal cues [SKM90]. This phenomenon can be explained by some simple extensions to the computational model described earlier as follows. The consequences and parallels to experimental evidence will be discussed later in this chapter.

EC layer cells in the present model act as spatial filters, responding to individual landmarks at specific positions relative to the animat. O'Keefe and Burgess [OB96] showed that the place cells in hippocampus can be modeled as a sum of Gaussians of varying variances where each Gaussian function encodes the distance to an edge of the environment along one of the two orthogonal axes. We have extended the spatial filters in EC in the light of the aforementioned work so that the tuning curves of such filters vary with the landmark distances along two orthogonal axes.

During exploration, if none of the EC units fire in response to an observed landmark position, an EC unit is recruited with the following activation function:

$$EC_i = \frac{1}{2\pi\ \sigma_1\ \sigma_2} \exp(-\frac{1}{2}\left(\frac{(x-\mu_x)^2}{\sigma_1^2} + \frac{(y-\mu_y)^2}{\sigma_1^2}\right))$$

$$\sigma_1 = \sigma_0\ \left(1 + 4\mu_1^2/R^2\right)$$

$$\sigma_2 = \sigma_0\ \left(1 + 4\mu_2^2/R^2\right)$$

Where μ_1 is the distance of the landmark in direction x_1 from the current position of animat, and similarly, μ_2 is the distance of the landmark in direction x_2 from the current position of animat.

For the purpose of simulation, we set σ_0 to 1.0 and R was set to 20, the diameter of the circular arena. As we shall see in what follows, this has an interesting effect on the localization behavior exhibited by animats.

It should also be noted that EC cell firing is also based on the landmark type, so that an EC unit firing signifies a landmark of a particular type at a particular relative position from the animat.

7.5 All-Or-None Connections Between EC and CA3 Layers

In the training phase, if none of the CA3 layer cells fire above a predetermined fraction of their peak firing level, a new CA3 layer cell is allocated. This newly created cell is then connected to the active EC layer cells. We have modified the connection weight assignment procedure of the existing model to reflect an all-or-none connection type. Rather than assigning weights proportionate to the activation of corresponding EC cells, we assign a weight of $1/nconn$ to each links, where $nconn$ is the total number of EC layer cells firing above their threshold levels. We set firing threshold of the newly allocated CA3 layer cell to 70% of its maximum possible weighted sum of incoming activations. During testing this threshold was reduced to 25% of the maximum possible activation level in order to allow animats to localize even in presence of partial sensory stimulus, or in other words, partial activity in the EC layer. Such a method has been found to be successful in modeling place-cell firing characteristics in simplified environments [OB96].

It has been observed that rodents give more importance to landmarks physically closer to their actual positions while localizing. Sharp and colleagues [SKM90] performed experiments on rodents in a cylindrical environment with a single cue card. After training, one more cue was added to the environment, producing a mirror symmetry in the environment. It was found that an overwhelming number of place-fields retained their shape and orientation with respect to only one of the two cues. Also, in most cases, place-fields were fixed relative to the cue that was nearest to the animal when it was first introduced in the environment.

7.6 Association of Rewards With Places

We have also extended the model to incorporate mechanisms that result in enhanced response of the EC layer neurons to landmark types that are closer to the reward locations. Whenever the animat receives a reward upon visiting a location, the maximum possible activations in EC layer cells are updated according to the following rule:

$$\delta w_j = \frac{1}{n-1}\,\Gamma\,\frac{\displaystyle\sum_{i=1, i\neq j}^{n} d_i}{\displaystyle\sum_{i=1}^{n} d_i}$$

where n is the number of types of landmarks present in the environment, $\Theta(i)$ is the total number of landmarks of type i present in the environment, and Γ is the amount subtracted from the landmark weights. Γ is computed as follows:

```
$\Gamma = 0$
for $i = 1$ to $n$
do
    if $w_i < \alpha\ \Theta(i)$
        $\Gamma = \Gamma + w_i$
        $w_i = 0$
    else
        $\Gamma = \Gamma + \alpha\ \Theta(i)$
        $w_i = w_i - \alpha\ \Theta(i)$
    endif
done
```

If multiple landmarks of same type are present, weights are altered by summing the distances of landmarks of similar types to the estimated goal location. The degenerate case of $n = 1$ is handled separately. It is clear that the weights remain unaltered if all landmarks are of the same kind, or, if all landmarks are equidistant from the goal. For the purpose of our simulations, α was set to 0.05 and the weights were initialized to 1.0.

The above rule gives more preference to landmark types that are near the goal location by removing a uniform amount α from weights assigned to each of the landmark types, and redistributing it so that a landmark type gains weight if such a landmark is near the goal. On the other hand, if there are multiple landmarks of the same type, or, if landmarks are far from the goal, such landmark type loses weight. It is also clear from the above equations that the sum of weights assigned to all landmark types remains unaltered.

It should be noted that the activation level is modulated uniformly across all EC cells that respond to a particular landmark type, and not just for EC cells that are active at the time of reward presentation.

We hypothesize that such a computation, which gives more weight to a particular *type* of sensory stimulus, takes place in the EC-Dg layers, as these layers get sensory information from the cortical areas as well as feedback connections from the Subiculum. The Subiculum is strongly believed to be part of the path-integration system [RT98]. Assuming a population code in the subiculum, it is conceivable that units encode path integration information which is reset using subsets of landmarks. This information can be used to supply a modulatory feedback signal to the units in EC and Dg.

7.7 Simulation Results

All simulation parameters and methods were identical to those in [BBH98a]. Briefly, the animats were introduced in an a-priori unknown environment that consisted of one or more landmarks. The landmarks could be identical or distinguishable from each other depending on the experiment being performed. Animats then explored their environments and allocated cells corresponding to different locations in the environment. The animats were also rewarded for visiting specific locations as they explored its environment. After a certain number of training trials, animats were removed from the environment, landmark positions were altered and the reward removed. When reintroduced in the environment, animats were able to re-localize, despite the change in configuration of the landmarks, using the available perceptual input, and moved toward the learned goal location.

7.8 Firing Characteristics of Units

In order to simplify analysis, for this part of experiments, animats were trained over a single training trial of 750 steps of random exploration.

As seen in Figure 7.8, animat consistently localized by giving more preference to the landmark physically closer to the point of entry into the environment. This phenomenon was not guaranteed with the scheme used in [BBH98a]. It should be noted, however, that the overall behavior displayed by animat stays unaltered with these enhancements, and we get search histograms similar to those in [BBH98a]. Figure 7.8 also shows the activation of CA1 layer place-cells once animats localized. It is important to note that the place-cell in question fired only in one of the two clusters over a single test trial. The place-cell in question fired

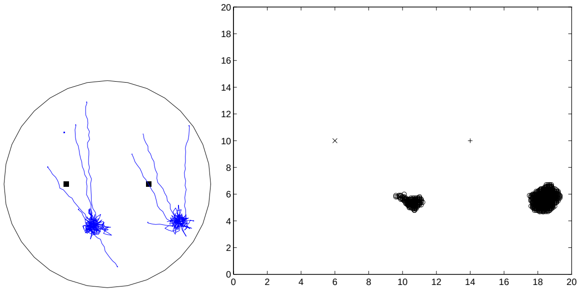

Figure 7.7 Left: Trajectories taken by an animat during test trials. Right: Superimposed place field firing regions during test trials.

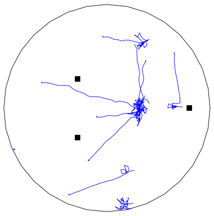

Figure 7.8 Left to right: Trajectories taken by animat when trained in an environment with three landmarks. The landmark on far right was distinguishable from rest. The landmark on far right was moved further while testing.

at the position cluster based around position $(12, 5)$ when the animat localized according to the landmark on the left, while the same place cell fired at places clustered around $(18, 6)$ when the animat localized using the landmark on the right. Interestingly enough, the right cluster is spread over a larger area, signifying that the place-cell in question fired over a larger area of the environment when animat localized using the right landmark. This effect can be explained by conflicting CA3 layer cell firing pattern due to the incorrect binding of CA1 unit activity with the path integration system for some of the activated CA1 units, resulting in a greater path integration estimate variance which in turn causes the CA1 unit in question pass the Mahalanobis distance test over a larger area of the environment. Figure 7.8 shows the trajectories during test trials, when one of the landmarks was distinct from the rest during training. In some of the trials animats were unable to localize because of lack of training in those regions of the environment. It can be seen in Figure 7.8 that in one of the test trials, animat localized solely based on the position of the right most landmark. The reason for such a behavior is discussed in the next section.

7.9 Landmark Prominence Based on Location and Uniqueness

The extension to the model that alter prominence assigned to landmark type is able to successfully replicate some of the behavioral results that were unaccounted for in [BBH98a], namely, the experiments where an array of three landmarks with different types of landmarks was transformed, in figure 9 c of [CCS86] as seen in Figure 7.9. The simulation parameters used here were identical to those in [BBH98a].

In addition, simulations demonstrate that the proposed extensions enable the animat to acquire associations between rewards and places and use them for goal-directed navigation.

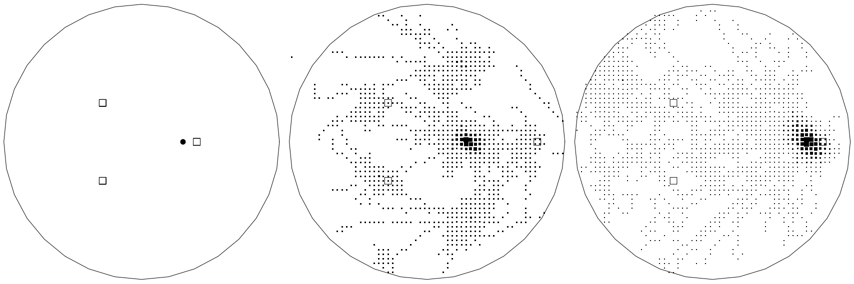

Figure 7.9 Top Left: Training Environment;Top Right: Normalized test histograms averaged over five animats with ten test trials each, when landmarks were indistinguishable from each other; Bottom: Right-most landmark distinguishable from the rest

As seen in Figure 7.9, during training the animats learned to give more weight to the type of landmark on the extreme right, due to its proximity to the goal as well as the uniqueness of its type. During testing, the right most landmark, which was distinguishable from the rest, was moved further towards right. Animats localized based on this unique landmark. Hence, a simple rule to associate the landmark type to a goal location was learned. Obviously, if all landmarks are identical, no such rule was learned, and the animats localized using a majority vote, as seen in Figure 7.9, top right. In Figure 7.3 , since only one visit to goal was allowed during training, the effect of prominence given by the animat to the unique landmark was not very pronounced.

CHAPTER 8. CONCLUSIONS

In this thesis we presented we developed a framework with associated operations required to learn and represent spatial information that helps one navigate in an a-priori unknown environment. We also presented some simulations of behavioral experiments using a model that roughly follows the framework for spatial learning. We showed that such a model performs very closely with the actual behavioral experiments performed on animals.

It turns out that the "Cognitive Map" is not just a single system, but is a result of many system working in synchrony. These systems can be broadly classified into the Path Integrator (PI) system and the Local View (LV) system, though at some points the boundary between PI and LV gets blurred.

We presented results of simulations based on experiments performed by Collett et. al. (1986). These experiments were mainly designed to measure the capabilities of rodents to learn and represent novel environments well enough to successfully navigate to a fixed reward location. It was found that our computational model performs equally well with model connections similar to the anatomical findings in the hippocampus and surrounding brain regions.

The next set of experiments, namely those of Morris (1981) were designed to find the types of strategies used by rodents to navigate when an escape platform was visible as opposed to invisible. Although it was not clearly proven in those experiments, it was evident that two distinct learning systems namely, spatial learning and stimulus-response based learning are utilized by animals. More conclusive evidence was later supplied by Morris and colleagues (1982) and even more conclusively by McDonald and White in 1993 . In this thesis we presented simulation results of Morris water-maze task with our computational model where we displayed very similar behavior of animats under circumstances similar to those in the original experiments. We also showed that two systems, one based on simple stimulus-response and other based on spatial learning are necessary and sufficient to reproduce the water-maze experiments. It should be noted here that more complex stimulus-response tasks, reinforcement learning, classical conditioning, and priming behavior cannot be reproduced by the present model. Entirely different brain regions and neural mechanisms are believed to be responsible for these behaviors[MW93a].

We also showed that multiple goals can be successfully represented and utilized for the navigational purposes by animats.

It remains to be seen how the learned environment can be transformed into a taxon or route-based navigation system [ON78]. This is a very interesting avenue of research because it can explain the emergence of an entirely different navigational strategy from this locale system based navigational system. Also, it is known that as experience of animal in an environment grows, the animal starts giving more emphasis to this taxon based system where a certain path to the goal is chosen as soon as the animal self-localizes in the environment, as opposed to a more exploratory approach where a route to goal is frequently computed. In a sense, it becomes a matter of choosing a right program to run in order to reach the goal as opposed to computing the path to the goal itself using a more general purpose map. This strategy is both computationally more efficient, as only a route search operation needs to be performed and at the same time it is more efficient storage wise in case only a few routes in an environment are frequently taken.

Another intriguing aspect is at the cellular level, where it has been found that although place cell firing itself is a robust estimate of animal's position in its environment [WM93], the action potential bursts of individual place cells are highly irregular. Similar paths taken in an environment produce very different results; sometimes the burst occurs while sometimes it doesn't [FM98]. Also this variance is excessive in the sense that the variance of the event that a burst occurs given a path in the environment is higher than a Poisson process with similar firing probability. This is an intriguing result as it is in direct contrast with a population code where an average (or more generally a simple function) of activity in a set of cells produces an estimate of the place. In such a case, all cells should on an average fire regularly for each event they encode in order to be robust estimators. This leads to interesting possibilities; perhaps the place cells also encode some yet to be determined signal. Another possibility is that the firing of place cells is simply an error signal that is supplied to another set of state estimators that represent the estimated position of the animals location in its environment. In other words, place cells act as comparators between what is being seen and what should be seen given the present path integrator state. In case only a subset of available perceptual cues are compared as measurements such sporadic firing of place cells can be explained. The latter hypothesis is not very far from what the model described here suggests. Further investigations in this direction are in order.

Exactly how this robust place-dependent or head-direction dependent firing of cells in the hippocampal system translates to navigational behavior is yet to be known. It can be safely assumed that the context, as described in Chapter 4 plays a major role in navigation. Experiments have found that there is rapid and drastic re-mapping of place field characteristics across environments or even within an environment if the animal performs a task that requires it to take specific paths in its environment repeatedly and in a fixed order. The exact role of the hippocampal place cells is yet to be known, whether they are more of a part of the PI system, keeping track of the animal's position relative to the completion of the task at hand, or whether it is a part of the LV system, responding to the current sensory information and conveying some encoding of it to the path integrator. Chances are that the hippocampal place cells are part of both systems, as their firing is very strongly modulated by both LV as well as the animal's position in the environment for the current task, and hence most probably, the PI state. This is what makes investigation of place cells so intriguing and interesting.

BIBLIOGRAPHY

[ABS72] D. Andersen, V. Bliss, and K. Skrede. Lamellar organization of hippocampal excitatory pathways. *Experimental Brain Research*, 13:222–238, 1972.

[AW89] D. Amaral and M. Witter. The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, 31(3):571–591, 1989.

[BA96] K. Blum and L. Abbott. A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, 8:85–93, 1996.

[Bal99] K. Balakrishnan. *Biologically Inspired Computational Structures and Processes for Autonomous Agents and Robots.* PhD dissertation, Iowa State University, Ames, IA, 1999.

[BBH97] K. Balakrishnan, O. Bousquet, and V. Honavar. Spatial learning and localization in animals: A computational model and its implications for mobile robots. Technical Report CS TR 97-20, Department of Computer Science, Iowa State University, Ames, IA 50011, 1997. (To appear in Adaptive Behavior).

[BBH98a] K. Balakrishnan, R. Bhatt, and V. Honavar. A computational model of rodent spatial learning and some behavioral experiments. In *Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society*, pages 102–107, Mahwah, New Jersey, 1998. Lawrence Erlbaum Associates.

[BBH98b] K. Balakrishnan, R. Bhatt, and V. Honavar. Spatial learning and localization in animals: A computational model and behavioral experiments. In *Proceedings of the*

Second European Conference on Cognitive Modelling, pages 112–119, Nottingham, UK, 1998. Nottingham University Press.

[BBH98c] K. Balakrishnan, R. Bhatt, and V. Honavar. Spatial learning in the rodent hippocampus: A computational model and some behavioral results, 1998. (in preparation).

[BBH98d] O. Bousquet, K. Balakrishnan, and V. Honavar. Is the hippocampal formation a kalman filter? In *Proceedings of the Pacific Symposium on Biocomputing '98*, 1998.

[BDJO97] N. Burgess, J. Donnett, K. Jeffery, and J. O'Keefe. Robotic and neuronal simulation of the hippocampus and rat navigation. *Philosophical Transactions of the Royal Society London B*, 352, 1997.

[BMM$^+$90] C. Barnes, B. McNaughton, S. Mizumori, B. Leonard, and L. Lei-Huey. Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in Brain Research*, 83:287–299, 1990.

[BO96] N. Burgess and J. O'Keefe. Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*, 6:749–762, 1996.

[Bro85] R. Brooks. Visual map making for a mobile robot. In *Proceedings of the IEEE Conference on Robotics and Automation*, 1985.

[BRO94] N. Burgess, M. Recce, and J. O'Keefe. A model of hippocampal function. *Neural Networks*, 7(6/7):1065–1081, 1994.

[BS81] A. Barto and R. Sutton. Associative search network: A reinforcement learning associative memory. *Biological Cybernetics*, 40:201–211, 1981.

[BS95a] H. Blair and P. Sharp. Anticipatory firing of anterior thalamic head direction cells: Evidence for a thalamocortical circuit that computes head direction in the rat. *Journal of Neuroscience*, 15:6260–6270, 1995.

[BS95b] M. Brown and P. Sharp. Simulation of spatial learning in the morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus*, 5:189–197, 1995.

[Buz89] G. Buzsaki. Two-state model of memory trace formation: A role for "noisy" brain states. *Neuroscience*, 31:551–570, 1989.

[CCS86] T. Collett, B. Cartwright, and B. Smith. Landmark learning and visuo-spatial memories in gerbils. *Journal of Neurophysiology A*, 158:835–851, 1986.

[CE93] N. Cohen and H. Eichenbaum. *Memory, Amnesia, and the Hippocampal System.* MIT Press/A Bradford Book, Cambridge, MA, 1993.

[CGT$^+$98] Y. Cho, K. Giese, H. Tanila, A. Silva, and Ho. Eichenbaum. Abnormal hippocampal spatial representaions in α-camkii^{T286A} and creb$^{\alpha\Delta-}$ mice. *Science*, 279:867–869, 1998.

[CLBM94] L. Chen, L-H. Lin, C. Barnes, and B. McNaughton. Head-direction cells in the rat posterior cortex. ii. contributions of visual and ideothetic information to the directional firing. *Experimental Brain Research*, 101:24–34, 1994.

[CLG$^+$94] L. Chen, L-H. Lin, E. Green, C. Barnes, and B. McNaughton. Head-direction cells in the rat posterior cortex. i. anatomical distribution and behavioral modulation. *Experimental Brain Research*, 101:8–23, 1994.

[CS92] P. Churchland and T. Sejnowski. *The Computational Brain.* MIT Press/A Bradford Book, Cambridge, MA, 1992.

[DH73] R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis.* John Wiley, New York, 1973.

[DMW99] B. Devan, R. McDonald, and N. White. Effects of medial and lateral caudate-putamen lesions on place and cue-guided behaviors in the water maze: relation ot thigmotaxis. *Behavioural Brain Research*, 100:5–14, 1999.

[Eic96] H. Eichenbaum. Is the rodent hippocampus just for 'place'? *Current Opinion in Neurobiology*, 6:187–195, 1996.

[ESYB96] H. Eichenbaum, G. Schoenbaum, B. Young, and M. Bunsey. Functional organization of the hippocampal memory system. *Proceedings of the National Academy of Sciences*, 93:13500–13507, 1996.

[FM98] A. Fenton and R. Muller. Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proceedings of the National Academy of Sciences*, 95:3182–3187, 1998.

[GFRR99] P. Georges-Franois, E. Rolls, and R. Robertson. Spatial view cells in the primate hippocampus: Allocentric view not head direction or eye position or place. *Cerebral Cortex*, 9:197–212, 1999.

[Has95] M. H. Hassoun. *Fundamentals of Artificial Neural Networks.* MIT Press, Cambridge, MA, 1995.

[Hea98] S. Healy. *Spatial Representation in Animals.* Oxford University Press, Cambridge, UK, 1998.

[JM93] M. Jung and B. McNaughton. Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*, 3:165–182, 1993.

[Kal60] R. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME*, 60, 1960.

[KHH$^+$98] C. Kentros, E. Hargreaves, R. Hawkins, E. Kandel, M. Shapiro, and R. Muller. Abolition of long-term stability of new hippocampal place cell maps by nmda receptor blockade. *Science*, 280:2121–2126, 1998.

[KLM96] L. Kaelbling, M. Littman, and A. Moore. Reinforcement learning: A survey. *Journal of AI Research*, 4:237–285, 1996.

[KM88] J. Keith and K. McVety. Latent place learning in a novel environment and the influences of prior training in rats. *Psychobiology*, 16(2):146–151, 1988.

[Koh89] T. Kohonen. *Self-Organization and Associative Memory.* Springer-Verlag, Berlin, 1989.

[KR83] J. Kubie and J. Ranck. Sensory-behavioral correlates in individual hippocampus neurons in three situations: Space and context. In W. Seifert, editor, *Neurobiology of the Hippocampus*, pages 433–447. Academic Press, London, 1983.

[LS88a] J.-C. Lacaille and P. Schwartzkroin. Stratum lacunosum-moleculare interneurons of hippocampal ca1 region: I. intracellular response characteristics, synaptic responses, and morphology. *Journal of Neuroscience*, 8:1400–1410, 1988.

[LS88b] J.-C. Lacaille and P. Schwartzkroin. Stratum lacunosum-moleculare interneurons of hippocampal ca1 region: Ii. intrasomatic and intradendritic recordings of local circuit synaptic interactions. *Journal of Neuroscience*, 8:1411–1424, 1988.

[Mar71] D. Marr. Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London*, 176:23–81, 1971.

[MBD$^+$98] E. Maguire, N. Burgess, J Donnett, R. Frackowiak, C. Frith, and J. O'Keefe. Knowing where and getting there: A human navigation network. *Science*, 280:921–924, 1998.

[MBG$^+$96] B. McNaughton, C. Barnes, J. Gerrard, K. Gothard, M. Jung, J. Knierim, H. Kudrimoti, Y. Qin, W. Skaggs, M. Suster, and K. Weaver. Deciphering the hippocampal polyglot: the hippocampus as a path-integration system. *The Journal of Experimental Biology*, 199(1):173–185, 1996.

[MBT$^+$96] PT. McHugh, K. Blum, J. Tsien, S. Tonegawa, and M. Wilson. Impaired hippocampal representation of space in ca1-specific nmdar1 knockout mice. *Cell*, 87:1339–1349, 1996.

[MCM91] B. McNaughton, L. Chen, and E. Markus. Dead reckoning, landmark learning, and the sense of direction: A neurophysiological and computational hypothesis. *Journal of Cognitive Neuroscience*, 3(2):190–202, 1991.

[McN89] B. McNaughton. Neuronal mechanisms for spatial computation and information storage. In L. Nadel, A. Cooper, P. Culicover, and R. Harnish, editors, *Neural Connections, Mental Computation*, pages 285–350. MIT Press, Cambridge, MA, 1989.

[MGRO82] R. Morris, P. Garrud, J. Rawlins, and J. O'Keefe. Place navigation impaired in rats with hippocampal lesions. *Nature*, 297:681–683, 1982.

[Mil72] B. Milner. Disorders of learning and memory after temporal lobe lesions in man. *Clinical Neurosurgery*, 19:421–446, 1972.

[Mil73] B. Milner. Hemispheric specialization: scope and limits. In F. Schmitt and F. Worden, editors, *The Neurosciences: Third Study Program*, pages 75–89. MIT Press, Cambridge, MA, 1973.

[MK87] R. Muller and J. Kubie. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7:1951–1968, 1987.

[MKR87] R. Muller, J. Kubie, and J. Ranck. Spatial firing patterns of hippocampus complex-spike cells in a fixed environment. *Journal of Neuroscience*, 7:1935–1950, 1987.

[MKS91] R. Muller, J. Kubie, and R. Saypoff. The hippocampus as a cognitive graph. *Hippocampus*, 1(3):243–246, 1991.

[MLC89] B. McNaughton, B. Leonard, and L. Chen. Cortical-hippocampal interactions and cognitive mapping: A hypothesis based on reintegration of the parietal and inferotemporal pathways for visual processing. *Psychobiology*, 17(3):230–235, 1989.

[MN89] B. McNaughton and L. Nadel. Hebb-marr networks and the neurobiological representation of action in space. In M. Gluck and D. Rumelhart, editors, *Neuroscience and Connectionist Theory*, pages 1–63. Lawrence Earlbaum Associates, Hillsdale, NJ, 1989.

[Mor81] R. Morris. Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12:239–261, 1981.

[MW93a] R. McDonald and N. White. A triple dissociation of memory systems: hippocampus, amygdala, and dorsal striatum. *Behavioral Neuroscience*, 103:3–22, 1993.

[MW93b] S. Mizumori and J. Williams. Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *Journal of Neuroscience*, 13:4015–4028, 1993.

[MWK98] L. McMahon, J. Williams, and J. Kauer. Functionally distinct groups of interneurons identified during rhythmic carbachol oscillations in hippocampus in vitro. *Journal of Neuroscience*, 18(15):5640–5641, 1998.

[NM97] L. Nadel and M. Moscovitch. Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology*, 7:217–227, 1997.

[OB96] J. O'Keefe and N. Burgess. Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381, 1996.

[OD71] J. O'Keefe and J. Dostrovsky. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 34:171–175, 1971.

[O'K76] J. O'Keefe. Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51:78–109, 1976.

[O'K89] J. O'Keefe. Computations the hippocampus might perform. In L. Nadel, L. Cooper, P. Culicover, and R. Harnish, editors, *Neural Connections, Mental Computation*, pages 225–284. MIT Press/Bradford Book, Cambridge, MA, 1989.

[ON78] J. O'Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Oxford:Clarendon Press, 1978.

[OS87] J. O'Keefe and A. Speakman. Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, 68:1–27, 1987.

[PM92] T. Prescott and J. Mayhew. Building long-range cognitive maps using local landmarks. In J. Meyer, H. Roitblat, and S. Wilson, editors, *Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 233–242, Cambridge, MA, 1992. MIT Press, Bradford Book.

[Pre95] T. Prescott. Spatial representations for navigation in animats. *Adaptive Behavior*, 4(2):85–123, 1995.

[PRS$^+$99] M. Pikkarainen, S. Ronkko, V. Savander, R. Insausti, and A. Pitkanen. Projections from the lateral, basal and accessory basal nuclei of the amygdala to the hippocampal formation in rat. *The Journal of comparative neurology*, 403:229–260, 1999.

[QMK90] G. Quirk, R. Muller, and J. Kubie. The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience*, 10:2008–2017, 1990.

[QMKR92] G. Quirk, R. Muller, J. Kubie, and J. Ranck. The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *Journal of Neuroscience*, 12(5):1945–1963, 1992.

[Ran84] J. Ranck. Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats. *Society for Neuroscience Abstracts*, 10:599, 1984.

[RH96] M. Recce and K. Harris. Memory for places: A navigational model in support of marr's theory of hippocampal function. *Hippocampus*, 6:735–748, 1996.

[RM86] D. E. Rumelhart and J. L. McClelland, editors. *Parallel Distributed Processing, Vol I-II.* MIT Press, Cambridge, MA, 1986.

[RMC+89] E. Rolls, Y. Miyashita, P. Cahusac, R. Kesner, H. Niki, J. Feigenbaum, and L Bach. Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *Journal of Neuroscience*, 9:1835–1845, 1989.

[Rol90] E. Rolls. Functions of the primate hippocampus in spatial processing and memory. In R. Kesner and D. Olton, editors, *Neurobiology of Comparative Cognition*, pages 339–362. Lawrence Earlbaum Associates, Hillsdale, NJ, 1990.

[Rol96] E. Rolls. The representation of space in the primate hippocampus and episodic memory. In T. Ono, B. McNaughton, S. Molotchnikoff, E. Rolls, and H. Nishijo, editors, *Perception, Memory, and Emotion: Frontiers in Neuroscience*, pages 375–400. Pergamon Press, 1996.

[RT96] A. Redish and D. Touretzky. Navigating with landmarks: Computing goal locations from place codes. In K. Ikeuchi and M. Veloso, editors, *Symbolic Visual Learning*. Oxford University Press, 1996.

[RT98] A. Redish and D. Touretzky. The role of the hippocampus in solving the morris water maze. *Neural Computation*, 10:73–111, 1998.

[RTR+98] E. Rolls, A. Treves, R. Robertson, P. Georges-Franois, and S. Panzeri. Information about spatial view in an ensemble of primate hippocampal cells. *Journal of Neurophysiology*, 79:1797–1813, 1998.

[SDL78] R. Struble, N. Desmond, and W. Levy. Anatomical evidence of interlamellar inhibition in the facia dentata of the rat. *Brain Research*, 152:580–585, 1978.

[SG94] P. Sharp and C. Green. Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *Journal of Neuroscience*, 14:2339–2356, 1994.

[Sha96] P. Sharp. Multiple spatial/behavioral correlates for cells in the rat postsubiculum: Multiple regression analysis and comparison to other hippocampal areas. *Cerebral Cortex*, 6:238–259, 1996.

[Sha97] P. Sharp. Subicular cells generate similar firing patterns in two geometrically and visually distinctive environments; comparison with hippocampal place cells. *Behavioral Brain Research*, 85:71–92, 1997.

[She74] G. Shepherd. *The synaptic organization of the brain: an introduction.* Oxford University Press, New York, 1974.

[SKM90] P. Sharp, J. Kubie, and R. Muller. Firing properties of hippocampal neurons in a visually symmetrical environment: Contributions of multiple sensory cues and mnemonic properties. *Journal of Neuroscience*, 10:2339–2356, 1990.

[SM96] W. Skaggs and B. McNaughton. Replay of neuronal firing sequence in rat hippocampus during sleep following spatial experience. *Science*, 271:1870–1873, 1996.

[SM99] P. Steele and M. Mauk. Inhibitory control of ltp and ltd: Stability of synapse strength. *Journal of Neurophysiology*, 81:1559–1566, 1999.

[ST92] N. Schmajuk and A. Thieme. Purposive behavior and cognitive mapping: A neural network model. *Biological Cybernetics*, 67:165–174, 1992.

[Tau95] J. Taube. Head-direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, 15:70–86, 1995.

[TM91] R. Traub and R. Miles. *Neuronal Networks of the Hippocampus.* Cambridge University Press, 1991.

[TMR90a] J. Taube, R. Muller, and J. Ranck. Head direction cells recorded from the postsubiculum in freely moving rats: I. description and quantitative analysis. *Journal of Neuroscience*, 10:420–435, 1990.

[TMR90b] J. Taube, R. Muller, and J. Ranck. Head direction cells recorded from the postsubiculum in freely moving rats: Ii. effects of environmental manipulations. *Journal of Neuroscience*, 10:436–447, 1990.

[Tol48] E. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55:189–208, 1948.

[TSE97] H. Tanila, M. Shapiro, and H. Eichenbaum. Discordance of spatial representation in ensembles of hippcoampal place eells. *Hippocampus*, pages 613–623, 1997.

[TWBM97] O. Trullier, S. Wiener, A. Berthoz, and J-A. Meyer. Biologically-based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51:483–544, 1997.

[WLJS92] J. Wathey, W. Lytton, J. Jester, and T. Sejnowski. Computer simulations of epsp-to-spike (e-s) potentiation in hippocampal ca1 pyramidal cells. *Journal of Neuroscience*, 12:607–618, 1992.

[WM93] M. Wilson and B. McNaughton. Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055–1058, 1993.

[WTR94] H. Wan, D. Touretzky, and A. Redish. Towards a computational theory of rat navigation. In M. Mozer, P. Smolensky, D. Touretzky, J. Elman, and A. Weigend, editors, *Proceedings of the 1993 Connectionist Models Summer School*, pages 11–19. Lawrence Earlbaum Associates, 1994.

[Zip86] D. Zipser. Biologically plausible models of place recognition and place location. In D. Rumelhart, J. McClelland, and the PDP Research Group, editors, *Parallel Distributed Processing: Explorations in the Micro-Structure of Cognition, Volume 2*, pages 432–470. MIT Press, Cambridge, MA, 1986.

[ZMSA86] S. Zola-Morgan, L. Squire, and D. Amaral. Human amnesia and the temporal lobe region: enduring memory impairment following a bilateral lesion limited to field ca1 of the hippocampus. *Journal of Neuroscience*, 6:2950–2967, 1986.