

Spatial Learning and Localization in Rodents: A Computational Model of the Hippocampus and its Implications for Mobile Robots

Karthik Balakrishnan
Allstate Research and Planning Center
321 Middlefield Road
Menlo Park, CA 94025, USA
kbala@allstate.com

Olivier Bousquet
Ecole Normale Supérieure des Télécommunications
Paris, France
Olivier.Bousquet@enst.fr

Vasant Honavar
Artificial Intelligence Research Laboratory
Department of Computer Science
Iowa State University, Ames, IA 50011, USA
honavar@cs.iastate.edu

Citation Information: Balakrishnan, K., Bousquet, O., and Honavar, V. (1998). Spatial Learning and Localization in Rodents: A Computational Model of the Hippocampus and its Implications for Mobile Robots. *Adaptive Behavior*. In press.

Abstract

The ability to acquire a representation of the spatial environment and the ability to localize within it are essential for successful navigation in *a-priori* unknown environments. The *hippocampal formation* is believed to play a key role in spatial learning and localization in animals in general and rodents in particular. This paper briefly reviews the relevant neurobiological and cognitive data, and their relation to computational models of spatial learning and localization used in contemporary mobile robots. It proposes a hippocampal model of spatial learning and localization, and characterizes it using a Kalman filter based tool for information fusion from multiple uncertain sources. The resulting model not only explains neurobiological and behavioral data from rodent experiments, but also allows a robot to learn a *place-based* metric representation of space and to localize itself in a stochastically optimal manner. The paper presents an algorithmic implementation of the model and results of several experiments that demonstrate its capabilities. These include the ability to disambiguate perceptually similar places, scale well with increasing errors, and the automatic acquisition of spatial information at multiple resolutions.

Keywords:

spatial learning, robot localization, hippocampal model, Kalman filter, probabilistic localization

Title for running head:

Spatial learning in rodents

1 Introduction

The ability to successfully navigate in a wide range of natural environments is essential to the survival of animals. Mobile robots need to be equipped with similar capabilities in order to perform the tasks expected of them in natural or man-made environments. This requires them to be able to acquire and use adequate representations of their spatial environments. Animals offer compelling existence proofs of such capabilities that have evolved in nature [Anderson, 1983; Schone, 1984] and challenge us to explore information processing mechanisms and computational architectures that can match their functionality, although they might be realized using different physical substrates with perhaps different design and performance constraints.

There is a large body of neuroanatomical, neurophysiological, and behavioral data on the possible role of different parts of the brain in general, and the *hippocampal formation* in particular, in spatial learning and navigation in animals [O’Keefe & Nadel, 1978; Schmajuk, 1997]. *Computational models*, although often simplified caricatures of their biological counterparts, offer an attractive approach to organizing, analyzing, abstracting, and exploring the implications of such data. In addition to suggesting new experiments designed to fill the gaps in our understanding of the systems being modeled, they are often very useful as sources of ideas for building artificial systems with comparable abilities. Against this background, a number of biologically inspired models of spatial learning, localization, and navigation have been proposed and implemented in robots and other automata [Kuipers & Byun, 1991; Mataric, 1992; Kortenkamp, 1993; Kuipers *et al.*, 1993; Wan *et al.*, 1994; Nehmzow, 1995; Blum & Abbott, 1996; Redish & Touretzky, 1996; Recce & Harris, 1996; Sharp *et al.*, 1996].

On the other hand, those involved in the design of autonomous robots are necessarily faced with multiple design and performance constraints imposed by the available technology and the task environments. Attempts to address the attendant engineering challenges in the design of such systems have led to the development of a broad range of mathematical and computational tools. These include information-theoretic characterizations of system complexities [Elfes, 1995; Koenig *et al.*, 1995], algorithms for the integration and use of noisy sensory data from multiple sensors [Ayache & Faugeras, 1987; Moutarlier & Chatila, 1989; Leonard & Durrant-Whyte, 1992; Hummel, 1995; Pagac *et al.*, 1995], and probabilistic localization approaches for mobile robots [Smith *et al.*, 1990; Crowley, 1995; Hebert *et al.*, 1995]. Use of such tools to analyze biologically

inspired models can often yield new insights into the capabilities and limitations of the underlying information processing structures and processes [Levy, 1989; Linsker, 1990; Treves & Rolls, 1991; Treves & Rolls, 1992].

The preceding discussion suggests that biologically inspired modeling efforts and the design of autonomous mobile robots can each benefit from the results and tools developed by the other. Against this background, this paper develops a biologically inspired model of spatial learning and localization, analyzes it using information fusion and probabilistic localization tools from robotics, and shows the contributions of the resulting model to both rodent spatial learning and robot navigation.

The rest of the paper is organized as follows: Section 2 identifies a set of processes that are required for goal-directed navigation in animals and robots. Different aspects of spatial learning and representation in robots and rodents are presented in Section 3, and the similarities between them are noted. This sets the stage for exploring useful synergies between approaches in these two domains. Section 4 provides a brief summary of the anatomical and physiological properties of the hippocampal formation and points to experimental results that implicate the hippocampus in rodent spatial learning. Our model of hippocampal spatial learning is then presented in Section 5. The need for probabilistic localization is clarified in Section 6 along with a summary of the Kalman filter approach to robot localization. Section 7 develops a Kalman filter framework for our computational model, while Section 8 presents algorithmic details of our implementation. The results of our simulation experiments are presented in Section 9. Other approaches related to the localization model developed in this paper are treated in Section 10. Finally, Section 11 provides a discussion including the contributions of this work to rodent and robot navigation. The Appendix presents some assumptions critical to the model developed in this paper and derives a number of results that aid the computational specification of our model.

2 Spatial Learning, Localization, and Navigation

Both animals and autonomous mobile robots need mechanisms to navigate purposefully in *a-priori* unknown or partially-known environments. However, for such behaviors to be possible, they must be endowed with mechanisms capable of answering the following questions (adapted from Levitt and Lawton (1990)):

1. Where am I? (*Localization*)
2. Where are other places relative to me? (*Spatial map*)
3. Where is the goal? (*Goal determination*)
4. How do I get to the goal from here? (*Path planning*)
5. How do I acquire places and goals? (*Spatial learning*)

As might be expected, the answers to these rather straightforward questions are intricately intertwined. The first question is concerned with the identification or recognition of the current place, a problem commonly referred to as *localization* in robotics. To aid localization, places must be represented and remembered in terms of sensory features that allow the animal or robot to quickly and *unambiguously* recognize different places. Quite often this is problematic due to *perceptual aliasing* in the environment. An artifact of the environment and/or sensor limitations, perceptual aliasing causes multiple places in the environment to appear sensorily identical, thus interfering with map learning and localization. Any mechanism that aids in the resolution of such ambiguities is thus of interest, as will be clarified later.

The second question deals with the representation of the spatial environment, which we refer to as a *spatial map*. This map encodes the relationship between places in the environment and could represent *topological*, *metric*, or *directional* information. It must be clarified that the map need not be *topographic* (e.g., maps in an atlas), so long as it *functionally* captures the relationships between places. Further, it is also possible, and often necessary, for the map to encode different kinds of relationships at differing levels of abstraction.

The third question concerns *goal determination* and includes acquisition, representation, and future identification of goals. Indeed, if the desired goal cannot be identified based on the current information state (sensory and memory representation) of the animal or robot, goal-directed navigation is not possible. Goals may be specified and represented in a number of ways, for instance, visual cues, proximity relationship to significant landmarks, metric position in a coordinate space, etc. [Anderson, 1983; Schone, 1984; Trullier *et al.*, 1997].

The computation of a navigation trajectory offers a means of answering the fourth question, which is referred to as *path planning* in the robotics literature [Latombe, 1991; Hwang & Ahuja,

1992]. The navigation trajectory that is computed depends on the nature of information encoded in the map (topological, metric, etc.), the specification of the goal, and importantly, on the kinds of computational mechanisms that the robot or animal has at its disposal. The trajectory may also be constrained by additional optimization concerns like shortest distance, least time, safest, etc., paths to the goal. Assuming the availability of appropriate computational resources, each representation scheme used in the spatial map might entail a different goal navigation strategy [Anderson, 1983; Schone, 1984; Trullier *et al.*, 1997]. For instance, if goals are represented by their positions in an underlying coordinate frame and the animal or robot knows its current position in this frame, the shortest trajectory to the goal is simply a vector-difference of the two positions.

If the animal or robot has reliable answers to these four questions, it can navigate in a purposeful manner to arbitrary goals. If the environments are *a-priori* known and static, places and maps can be pre-specified and goals can be pre-determined, making purposeful navigation an outcome of *genetically programmed behaviors* in animals and *pre-wired control* in robots. However, the real-worlds occupied by animals and most robots are dynamic and at best *partially-known*. In such cases, they must possess mechanisms to *learn* spatial maps of environments they encounter and *adapt* to dynamic changes within them. They must be capable of learning, representing, and updating places and goals, i.e., they must possess *spatial learning* ability.

The main focus of this paper is on spatial learning (including building of spatial maps) and localization aspects of goal-directed navigation. We briefly consider our current understanding of how these processes are realized by contemporary robots and animals. In the process, we will also explore similarities and relationships between them.

3 Aspects of Spatial Information Representation in Robots and Rodents

Since animals appear well-adapted to the spatial constraints of their ecological niches and demonstrate a wide variety of complex spatial behaviors, there is considerable interest in determining the structures and processes that make spatial learning possible. On the other hand, robots provide interesting examples of how spatial constraints are understood and addressed by current engineering approaches. This section presents our current understanding of spatial representation and learning processes in animals and contemporary robots. Not surprisingly, there appear to be close parallels

between the two.

3.1 Representation of Spatial Information in Robots

Contemporary robots represent spatial information in one of two broad ways: *metric* (or *location-based*) or *relation-based* [Moutarlier & Chatila, 1989], although some recent approaches have attempted to combine merits of the two [Levitt & Lawton, 1990; Kuipers & Byun, 1991; Kortenkamp, 1993; Thrun, 1996]. In the location-based scheme, places (or locations of objects) are all represented in the *same* coordinate frame with a *global origin*. This origin may be absolute in the space or chosen appropriately by the robot. In either case, the robot learns and represents the locations of places and objects with respect to this origin. With such a scheme the direct short-cut path between any two arbitrary places can be easily determined by simply comparing their metric locations in the common frame. However, extracting local relationships between adjacent places requires non-trivial computation on their metric representations. For instance, if we need to know the place that is north-west of place P_1 , we have to search through the entire map to identify a place that comes closest to being P_1 's north-west neighbor. Thus, location-based representations make it easy to determine the relationship between two given places, but make it hard to determine places that are in specific local relationships.

Relation-based approaches, on the other hand, represent *local relationships* between places, thereby capturing adjacency relations. Within such a framework it is quite trivial to determine a place that is in a specific local relationship to another (e.g., the place that is north-west of P_1). However, since there is no global coordinate frame representation, the relationship between two arbitrary places can only be determined in terms of the relationships among intermediate places. Depending on the size and form of the spatial map, this might require cumbersome computations.

An example of metric spatial representation in robots is the *occupancy grid* [Moravec & Elfes, 1985; Elfes, 1989] which represents the environment using a grid-like decomposition. Each grid cell corresponds to a portion of the environment and adjacent grid cells represent adjacent physical regions in the environment. The grid cells are associated with a figure of merit called the *probability of occupancy*, which is updated based on sensory inputs and appropriate sensor models [Elfes, 1989]. Since grid cells correspond to regions of physical space, the probability of occupancy is an indicator of the presence or absence of objects in the corresponding region of physical space. These

occupancy grids are useful for determining *free space* for uninhibited robot navigation and can often be easily learned from *sonar* data. *Inference grids* are recent generalizations of this concept that have been used to learn, represent, and manipulate different kinds of *features* (e.g., shape, color, texture, etc.) in addition to just occupancy [Elfes, 1992; Elfes, 1995]. Since the grids provide a complete map of objects and obstacles in the explored portion of the environment, an obstacle-free navigation trajectory can be computed to arbitrary goal locations on the map. However, since the grid represents the entire environment and not just the significant places, space is stored in an inefficient manner. Further, after each sensory measurement, the robot must update the occupancy probabilities of the entire grid (feature probabilities in the case of inference grids), making it computationally expensive to maintain these maps. Also, increasing the grid cell resolution (smaller grid cells) results in a quadratic increase in the number of cells required for the same 2D spatial environment. On the other hand, decreasing the grid resolution reduces the computational burden on the algorithm but results in a loss of information and a consequent increase in uncertainty [Elfes, 1992].

A popular example of relation-based spatial representation is the *topological map* which only represents *distinctive places* in the environment along with the local relationships between them [Kuipers, 1978]. This map can thus be thought of as a graph, where nodes represent distinct places and the edges denote the relationship between them [Brooks, 1985]. Distinctive places are usually defined using *sonar signatures* [Kuipers & Byun, 1991; Mataric, 1992; Kortenkamp, 1993], *image signatures* [Kortenkamp, 1993; Kortenkamp & Weymouth, 1994; Engelson, 1994], *panoramic views* [Tsuji & Li, 1993], *local occupancy grids* [Langley & Pfleger, 1995; Yamauchi & Langley, 1997], etc. Local relationships usually represented include the *directional* relationship between places [Nehmzow, 1995] or the directional relationship along with *distance* information [Kuipers & Byun, 1991]. Some topological approaches also encode metric information corresponding to the distinctive places (e.g., lengths of walls, corridors, etc.) [Mataric, 1992]. Such relation-based approaches produce a compact world representation since they only represent distinctive places and not the entire environment. Further, they are robust to global movement errors as they only represent local relationships between places [Brooks, 1985]. These topological graphs can be learned by exploration of the environment. However, navigation to a goal requires a search through the graph to determine a *route* from the current robot location to the goal, which can be computationally expensive for

large graphs.

The attendant advantages and limitations of metric and topological approaches to spatial representation has led some to consider *hybrid* approaches. This has led to work in *multi-level representations* of spatial information [Levitt & Lawton, 1990; Kuipers & Byun, 1991; Kortenkamp, 1993; Thrun, 1996], some of which are described in more detail in Section 10.2.

Although the spatial representations discussed above (and their many variants) are extensively used in contemporary robotics, a number of other issues remain to be addressed. Is a spatial scene sensed and represented simply as some form of a *signature* (e.g., a photograph-like snapshot) or is it represented as a *composite* of the individual objects (also called *landmarks*, *cues* etc.) making up the scene? Further, are *all* the landmarks in the scene treated equally or are some landmarks given priority over others? If they are indeed prioritized, what mechanism is used to assign these priorities? What role does landmark size, distinctiveness, function, stability/reliability, etc., play in spatial representation? Finally, is the representation of space homogeneous or is it heterogeneous, supporting spatial representations of multiple types at multiple levels of granularity?

Since these issues also apply to navigating animals, it is natural to ask how animals address them. The next section provides answers to some of these questions.

3.2 Spatial Learning and Representation in Rodents

Animal learning in general and spatial navigation in particular, have been the subject of intense study for several decades [Tolman, 1932; Hull, 1943; Tolman, 1948; Anderson, 1983; Mackintosh, 1983; Schone, 1984; Gallistel, 1990]. A large fraction of this effort has been devoted to the study of *rodent* navigation. This is not surprising in light of experimental evidence that space plays a dominant role in rodent behavior, making them natural subjects in studies of spatial learning [Hebb, 1949]. An important aspect of rodent spatial learning appears to be their tendency to learn *places* rather than individual *stimuli*. For instance, Hebb (1949) trained rats to run to a dish of food located at the edge of an open table. Once the rats were sufficiently trained, the table (and the dish) were rotated by 90° . It was found that the rats ran (at least once) to the prior location of the dish relative to the room even though the dish itself was visibly at a different position. In another experiment, rats were trained to jump from one platform to another to obtain a food reward. Once trained, the rats were found to jump into space when the second platform was moved to a different

location [Hebb, 1949]. Similar results confirming *place-learning* (as opposed to *cue-learning*) have also been observed by a number of other researchers [O’Keefe & Nadel, 1978].

Experiments have also revealed that rats are capable of *detour behaviors* and *latent learning*. In detour behavior, rats trained to follow a particular trajectory to a goal show themselves capable of adopting alternate trajectories (that are often *novel* and *optimal*) when the original trajectory is blocked or otherwise unavailable [Tolman & Honzik, 1930; Tolman *et al.*, 1946]. Latent learning refers to the ability of rodents to acquire a spatial representation of the environment in the *absence* of explicit goals [Tolman, 1932; Tolman *et al.*, 1946].

Rats have also been found to successfully swim to a submerged platform in a milky pool of water, thereby suggesting an ability to compute trajectories to *hidden goal locations* [Morris, 1981]. Other related experiments have established that rats appear to be able to deal with transformations of visual scenes due to locomotion and compute approach trajectories using *inverse transformations* [Keith & McVety, 1988]. Experiments with gerbils have led to the suggestion that these animals compute *vectors* to landmarks. Further, *independent* vectors appear to be computed for each of the landmarks [Collett *et al.*, 1986]. This has led to the suggestion of a vector-based representation of space, wherein a direct vector to a goal location can be computed by subtracting a vector from the goal location to a landmark from a vector to the same landmark from the current location [Collett *et al.*, 1986; McNaughton *et al.*, 1995].

Experiments have also provided many insights into the processing of sensory stimuli. For instance, it appears that rats give more importance to *remote sensory cues than to local ones*, possibly because remote cues are the *least-variable* objects in the environment [Hebb, 1949]. It has also been found that the *stability* of stimuli (landmarks) is critical for spatial learning, i.e., even if a goal is constantly and reliably associated with a landmark, rats fail to capture this relationship if the landmark (and the goal) are *unstable*, i.e., they are moved to different locations in each of the learning trials [Biegler & Morris, 1996]. Further, simply increasing the number of landmarks in the environment also does not help. This led the researchers to conclude that spatial learning is critically influenced by *stable* relationships rather than merely the number or salience landmarks [Biegler & Morris, 1996]. Similar results have also been reported by [Bennett, 1993a]. Other experimental manipulations have led to findings that *short* landmarks are often not remembered by the animals possibly because they tend to become *obscured* by intervening taller objects [Bennett,

1993b].

One of the earliest suggestions concerning the representation of spatial information was made by Tolman (1948). Based on extensive studies Tolman concluded that rats are capable of learning the locations of reward sites independently of the responses used to approach them. Further, by designing experiments that showed evidence of *metric computations* (e.g., length of maze arms, positions of arm joints, etc.) [Tolman & Honzik, 1930], he proposed that animals learn *cognitive maps* of their environments and that these maps allow the animals to extract *metric* position information [Tolman, 1948]. However, Hull explained the experimental observations of Tolman in terms of the *habit-family hierarchy*, arguing that spatial behavior arises strictly through the association of specific *responses* with particular *stimuli* [Hull, 1934a; Hull, 1934b]. O’Keefe and Nadel formalized the difference between these two theories of information representation, suggesting that rats possibly use two spatial representation and navigation systems: the *locale system* that corresponds to a cognitive map and represents spatial information in metric terms, and the *taxon system*, which is associated with stereotypic behaviors like route following [O’Keefe & Nadel, 1978]. Thus, while the former approach encodes metric information in the spatial map, the latter implicitly encodes local relationships.

3.3 Role of Path Integration in Rodent and Robot Navigation

Path integration, also known as *dead-reckoning*, refers to the process of updating an estimate of one’s own position based on the knowledge of direction, speed, and time of self-motion [Gallistel, 1990; Everett, 1995]. This usually involves integrating acceleration signals over time to obtain velocities, and the integration of velocity signals over time to obtain displacement vectors [Gallistel, 1990; Everett, 1995].

There is substantial evidence for path integration in rodents, primarily in the use of dead-reckoning for *homing* behaviors [Etienne, 1992; Gallistel, 1990]. For instance, gerbils have been found to return on a direct bearing to their nest after circuitous search trajectories, even in *complete darkness* [Mittelstaedt & Mittelstaedt, 1980; Mittelstaedt & Mittelstaedt, 1982]. Further, based on observations from experimental manipulations like rotation of the entire arena or a linear displacement of the nest, researchers have concluded that path integration is based on *inertial* directional information from the vestibular system and linear information involving proprioception

and/or efference copy from the animal's self-generated motion [Mittelstaedt & Mittelstaedt, 1980; Mittelstaedt & Mittelstaedt, 1982].

Similar experiments have also been performed by other researchers on golden hamsters [Etienne, 1985; Etienne, 1992; Etienne *et al.*, 1996], which demonstrate that without frequent visual corroborations, path integration systems rapidly accumulate errors. This has led to the suggestion that in the absence of visual information, path integration appears to be useful only for short exploratory excursions from a known site [Etienne *et al.*, 1996]. It has also been shown experimentally that in conflict situations between distant but familiar visual references and path integration, the animals appear to give priority to stable visual references [Etienne *et al.*, 1990].

Researchers have also found some evidence for the grounding of internal spatial representations in a *locomotion-based metric system* in non-homing behaviors [Collett *et al.*, 1986]. It is strongly suspected that these metric estimates are provided by the dead-reckoning system [McNaughton *et al.*, 1995; McNaughton *et al.*, 1996].

Dead-reckoning mechanisms are also used in robot navigation and a wide range of devices have been developed for this purpose [Everett, 1995]. As with animals, dead-reckoning is usually involved in *homing* behaviors, where the robot returns to its *home-base* after executing its spatial task. Dead-reckoning input has also been used in the building of metric spatial maps, with the dead-reckoning system providing a Cartesian coordinate representation of space. For instance, stochastic approaches for information fusion (e.g., using the Kalman filter) have been used for robot localization and world modeling [Ayache & Faugeras, 1987; Moutarlier & Chatila, 1989; Crowley, 1995]. In these approaches, dead-reckoning is used to provide estimates of robot position which are subsequently corrected based on observations. Recently, Yamauchi and Beer (1996) have defined a spatial representation scheme based on adaptive place networks, where Cartesian position estimates derived from dead-reckoning are used to represent regions of physical space.

Thus, there appear to be many similarities between spatial learning processes and representations used by rodents and contemporary robots. But what is the *locus* of these processes and representations in rodents? Do these processes in rodents have any implications for robot spatial learning? The following sections attempt to answer these questions.

4 Role of the Hippocampus in Rodent Spatial Learning

The hippocampal formation in animals has been strongly implicated in spatial learning based on two broad sources of evidence: *lesion studies* and *cellular recordings* from hippocampal cells. While lesion studies typically demonstrate the inability of hippocampal-lesioned animals to learn tasks of a spatial nature (e.g., mazes, object-place tasks, etc.), cellular recordings show correlated firings of hippocampal cells during the execution of such tasks. For instance, in the water-maze task (where rats have to swim to a submerged platform in a milky pool of water), rats with hippocampal damage are incapable of learning to navigate directly or efficiently to the hidden platform, although they appear perfectly capable of navigating accurately to a platform that is *visible* [Morris *et al.*, 1982; Sutherland *et al.*, 1982]. A number of other such lesion studies are detailed in [O’Keefe & Nadel, 1978; Jarrard, 1993; Fenton & Bures, 1994]. We will focus more on the anatomical and physiological properties of the hippocampal formation since they shed more light on the neural structures and processes used in rodent spatial learning and localization.

4.1 Anatomy of the Hippocampal Formation

The hippocampal formation is an association area of the brain that receives highly processed sensory information from the major associational areas of the cerebral cortex [Churchland & Sejnowski, 1992; Cohen & Eichenbaum, 1993]. As shown in Figure 1, these inputs arrive at a convergence area called the *entorhinal cortex* (EC), which itself is a part of a major convergence area called the *parahippocampal cortical* area [Squire *et al.*, 1989]. The hippocampal formation is composed of the *dentate gyrus* (Dg), and areas *CA3* and *CA1* of Ammon’s horn.

Insert Figure 1 near here

The Dg contains *granule cells* that receive input from the EC via the *perforant path*, and output to the CA3 via *mossy fibers*. The CA3 region of the hippocampus primarily contains *pyramidal* (or complex-spike) cells along with inhibitory interneurons like *basket cells*, *chandelier cells*, *mossy cells*, etc. [Traub & Miles, 1991]. These CA3 cells receive inputs from the EC through the perforant path, the Dg via mossy fibers, and other CA3 pyramidals via the recurrent collaterals. The CA1 region of the hippocampus also consists of pyramidal cells and interneurons. However, unlike CA3

cells, CA1 cells do not project to other cells in CA1. The CA1 pyramidal neurons receive inputs from the EC via the perforant path and from the CA3 pyramidal neurons through the *Schaffer collaterals*. Axons from the CA1 pyramidal neurons project via the *alveus* to the *subiculum* (Sb) and also to the EC. Sb also receives input from the EC and projects to the *pre-* and *post-subiculum*, the deep layers of the EC, and to the hypothalamus, septum, anterior thalamus and the cingulate cortex [Churchland & Sejnowski, 1992]. There is some evidence that the EC projects back to many of the cortical association areas from which it receives input and mediates information storage there [Rolls, 1996]. However, not much is known about these mechanisms. We will now present some physiological properties of hippocampal cells that are critical for understanding the design choices made in our computational model.

4.2 Physiological Properties of Hippocampal Cells

Cellular recordings from many regions of the brain, including the hippocampus, have revealed crucial properties of the underlying neuronal mechanisms. For instance, in their recordings from CA1 pyramidal cells of a rat hippocampus, O’Keefe and Dostrovsky found that the neurons were selectively active in particular regions of the environment of the moving rat [O’Keefe & Dostrovsky, 1971]. These cells thus appear to code for specific places and have been accordingly named *place cells*, with *place fields* denoting the corresponding regions over which they are active [O’Keefe, 1976]. Since their initial discovery, cells with such location-specific firing have been found in almost every major region of the hippocampal system, including the EC [Quirk *et al.*, 1992], the Dg [Jung & McNaughton, 1993], the hippocampus proper [O’Keefe & Dostrovsky, 1971; O’Keefe, 1976], the Sb [Barnes *et al.*, 1990; Sharp & Green, 1994], and the postsubiculum [Sharp, 1996].

In addition to place cells, *head-direction cells* have also been discovered in the hippocampal region. These cells appear to respond to particular directions of the animal’s head, irrespective of its location in the environment. Each cell fires only when the animal’s head faces one particular direction (over an approximately 90 degree range) in the horizontal plane, and their relative directional tuning appears to be independent of the pitch and roll of the head. Thus, these cells function like a *built-in compass*. These cells were first discovered in the postsubicular area of the hippocampal formation [Ranck, 1984; Taube *et al.*, 1990a; Taube *et al.*, 1990b]. More recently, such directional cells have also been discovered in the retrosplenial cortex [Chen *et al.*, 1994a; Chen *et al.*, 1994b],

the anterior thalamus [Taube, 1995; Blair & Sharp, 1995], and the laterodorsal thalamus [Mizumori & Williams, 1993].

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'Keefe, 1976; Muller *et al.*, 1987; O'Keefe & Speakman, 1987; O'Keefe, 1989; Wilson & McNaughton, 1993]. Also, place cells are active in multiple places in the environment [O'Keefe & Speakman, 1987] and also in multiple environments [Kubie & Ranck, 1983; Muller *et al.*, 1987; Muller & Kubie, 1987]. 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 [Wilson & McNaughton, 1993].

Experiments have also revealed that when the animal is introduced into a familiar environment, place fields are initialized based on visual cues and landmarks [Muller & Kubie, 1987; Muller *et al.*, 1987; Sharp *et al.*, 1990]. Once initialized, the place fields have been found to persist even if the visual cues are removed in the animal's presence [O'Keefe & Speakman, 1987], implying that place cell firing must also be maintained by a source other than visual stimulus. CA1 cell firings have also been found to be conserved in darkness, provided the animal is first allowed to explore the apparatus under illuminated conditions [McNaughton *et al.*, 1989; Quirk *et al.*, 1990]. This has led to the hypothesis that place fields are maintained by ideothetic (self-motion) mechanisms, in particular, by the dead-reckoning system.

Similarly, head-direction cells have also been found to be responsive to visual inputs and demonstrate a number of properties similar to the place cells described above.

Experiments have also shown that the motor system plays a critical role in the firing of place and head-direction cells. For instance, with restraints on active motion, both hippocampal place cell activity [Foster *et al.*, 1989] as well thalamic head-direction cell activity [Knierim *et al.*, 1995] have been observed to cease. Since the dead-reckoning system presumably receives input from the motor system (in the form of a motor efference copy, etc.) these experiments further implicate dead-reckoning in place and head-direction cell firing [McNaughton *et al.*, 1996]. Rats have also been found to develop stable, unique associations between visual stimuli and the dead-reckoning system, which allows them to automatically realign the dead-reckoning system when mismatches occur [Knierim *et al.*, 1995].

In summary, place cells and head-direction cells respond to sensory as well as dead-reckoning inputs. These cells are active in multiple environments and also active in multiple places in the same environment. The firing of these cells is conserved in darkness, provided the animal is first allowed to orient itself under lighted conditions. Finally, the firing of these cells is directly related to the motor system and any restraint on active motion ceases cell firing.

5 A Computational Model of Hippocampal Spatial Learning

Based on neurophysiological, anatomical, and behavioral data, O’Keefe and Nadel (1978) suggested that the hippocampal formation is the site of the cognitive map proposed by Tolman (1948). In what they termed the *locale system hypothesis*, they suggested that the hippocampal formation learns a map-like representation of space. Importantly, they hypothesized that the hippocampal cells integrate and associate sensory inputs with dead-reckoning information generated by the animal, thereby encoding the map using a metric representation scheme [O’Keefe & Nadel, 1978].

We have developed a computational specification of this locale hypothesis that learns, represents, and updates a metric spatial map [Bousquet *et al.*, 1998]. The system learns a map of *distinct* places in the environment and labels the *center* of each place with metric position estimates derived from dead-reckoning. This fusion of sensory and dead-reckoning information takes place in the hippocampal cells, as shown in the sketch of the model in Figure 2.

Insert Figure 2 near here

The overall functioning of our model is as follows. During exploration of its environment, the animal detects and recognizes landmarks, and estimates their relative positions. This *processed* landmark information is used to *organize* (or create) units in the EC layer in such a way that units respond to specific landmarks appearing at particular positions *relative* to the animal. Thus, EC cells encode and respond to *vectors* to specific landmarks which is consistent with the observed properties of EC cells in the hippocampus [Quirk *et al.*, 1992]. Since places are characterized by their relationship to a set of landmarks visible from that place and EC units encode vectors to specific landmarks, concurrent activity of a set of EC units captures the landmark relationships at that place. CA3 place cells are then organized (or created) to capture this EC activity, thereby

representing places. The firing of CA3 cells in response to sensory inputs at a given place thus constitutes an *internal place-code* for that place. The CA1 cells then associate this internal place-code with metric position labels from the dead-reckoning system. Thus, while EC and CA3 cells enable the learning of places using sensory information, the CA1 layer learns *centers* of these places in terms of dead-reckoning position estimates.

Another possible role for the CA1 layer is suggested by *perceptual aliasing* problems encountered by robots, wherein multiple places in the environment appear sensorily identical. Given the limitations on the sensory abilities of a number of animals and the commonalities in the features of many habitats (e.g., vast deserts, thick forests, etc.), it is reasonable to assume that animals frequently encounter such aliasing problems. The fact that they successfully deal with their environments suggests that they possess mechanisms for resolving such ambiguities. We hypothesize that perceptually identical places (and hence place codes in the CA3 layer) are *disambiguated* in the CA1 layer by means of dead-reckoning information. Thus, perceptually aliased places activate the *same* population of cells in the CA3 layer but different populations in the CA1 layer.

It should be noted that this model ascribes different roles to the CA3 and CA1 layers, which leads to differences in the firing of CA3 and CA1 units. However, neurobiological experiments to date have found little difference in the firing properties of CA3 and CA1 cells. In Section 11 we offer some explanations to resolve this discrepancy.

The model also assumes that goals are learned and represented in terms of their metric positions in the animal's dead-reckoning framework. However, this goal representation is believed to reside outside the hippocampus [O'Keefe, 1989; Redish & Touretzky, 1996].

In addition to learning places by appropriately creating EC, CA3, and CA1 units, the model also performs localization. When the animal visits familiar places, incoming sensory inputs activate a place code in the CA3 layer. As multiple CA1 place codes may respond to this CA3 code (due to perceptual aliasing), dead-reckoning information is used to determine the CA1 code with place field center closest to the dead-reckoning based position estimate of the animal. The hippocampal system then performs spatial localization by *matching* the *predicted* position of the animal (from the dead-reckoning system) with the *observed* position of the place field center (dead-reckoning estimate stored with the activated CA1 place code). Based on this match, the dead-reckoning

estimate as well as the place field center can be updated, as shown in Figure 3.

Insert Figure 3 near here

In our model, place fields are quickly formed through exploration of the environment. The CA1 place cells are driven by both sensory as well as dead-reckoning inputs. Thus, they can be manipulated by changes in the landmark configuration and yet continue to fire in darkness. In a familiar environment, the animal places high confidence in its dead-reckoning based position estimate. Under these circumstances, even if some landmarks are removed, the high confidence in path integration overrides the changed sensory activation, allowing the animal to still correctly identify the place. We believe that changes in the EC-CA3 (as well as EC-Dg) synapses soon align the place code with the new sensory input. Upon reintroduction into a familiar environment, the animal initially distrusts its dead-reckoning estimate. As described earlier, the animal uses sensory information to localize, thereby initializing its dead-reckoning system. In darkness, without the appropriate updates of Figure 3, the position estimates accumulate errors and place cell firing in CA1 drifts.

6 Need for Probabilistic Localization

Following the locale hypothesis, our spatial learning model of the hippocampus integrates information from two streams: the sensory inputs and the dead-reckoning system. However, it should be noted that information provided by both these streams is uncertain. Sensory systems of animals accommodate considerable errors (e.g., in the estimation of distance and direction to visible objects, recognition of objects, etc.). Dead-reckoning is also prone to estimation errors and drifts, and the very fact that place cell (and head-direction cell) firings drift in darkness is suggestive of errors in dead-reckoning. Thus, in order for the hippocampus to perform robust spatial localization using these uncertain information sources, it must necessarily be capable of appropriately handling these uncertainties. To borrow a term from the robotics literature, the hippocampus must be capable of *probabilistic localization*. Although several hippocampal models of spatial learning have been proposed, some of them closely related to our own [Redish & Touretzky, 1996; Redish & Touretzky, 1998], none of them explicitly address the question of handling uncertainty in data. In order

to satisfactorily characterize hippocampal spatial learning, we need a probabilistic framework for addressing uncertain information fusion.

Since mobile robots also have to deal with uncertainties in sensing and action, a number of probabilistic localization approaches have been developed in that context. One such localization tool is the *Kalman filter* [Kalman, 1960; Gelb, 1974; Crowley, 1995] (or some extension or generalization of it), which allows the robot to build and maintain a *stochastic spatial map* of its environment [Smith *et al.*, 1990; Moutarlier & Chatila, 1989; Leonard & Durrant-Whyte, 1992], propagate sensory and motion uncertainties, and localize in *stochastically-optimal* ways [Ayache & Faugeras, 1987; Hebert *et al.*, 1995]. As we will demonstrate shortly, our model of hippocampal spatial learning and Kalman filter based robot localization are very similar in principle. This encourages us to explore extensions of the Kalman filter framework to characterize hippocampal spatial learning. But first, we digress briefly to look at the use of Kalman filter based approaches in robot localization.

6.1 Robot Localization Using a Kalman Filter

The Kalman filter technique for robot localization typically maintains a stochastic spatial map of the robot’s environment at each discrete time-step k . This stochastic spatial map, denoted by a state vector \mathbf{x}_k , includes an estimate of the robot’s current position and possibly the estimated positions of other landmarks in the robot’s environment. It is assumed that the *system model* is specified which denotes the change in the state based on robot motion:

$$\mathbf{x}_k = \Phi_{k-1}\mathbf{x}_{k-1} + \mathbf{u}_{k-1} + \mathbf{v}_{k-1} \quad (1)$$

Here, Φ_{k-1} is the transformation of state based on robot motion (e.g., translation, rotation, etc.), \mathbf{u}_{k-1} is the linear change in state based on the intended robot motion, and \mathbf{v}_{k-1} is a *zero-mean* motion error with covariance matrix \mathbf{Q}_{k-1} . Here, \mathbf{v}_{k-1} represents the fact that in real-world environments, the intended motion of the robot (\mathbf{u}_{k-1}) is offset by errors like wheel-slippage, friction, improper tire inflation, etc. This makes the *actual* robot motion different from its *intended* one.

Kalman filter also requires a *measurement model* to be specified, which captures the relationship between measurements (or observations) and the current state of the robot \mathbf{x}_k :

$$\mathbf{z}_k = \mathbf{H}_k\mathbf{x}_k + \mathbf{w}_k \quad (2)$$

Here, \mathbf{H}_k represents the *sensor model* that maps states into observations. Since real-world measurements are not perfect, measurement error, \mathbf{w}_k , is modeled as a *zero-mean* noise with covariance matrix \mathbf{R}_k .

Given these two models, the Kalman filter stores an *estimate* of the current state $\hat{\mathbf{x}}_k$ and its associated covariance matrix $\mathbf{P}_k = E[(\mathbf{x}_k - \hat{\mathbf{x}}_k)(\mathbf{x}_k - \hat{\mathbf{x}}_k)^T]$, and updates the estimates by making *state predictions* and matching them with *observations*. Suppose the current state estimate is $\hat{\mathbf{x}}_{k-1}^+$ at time step $k - 1$, with the covariance matrix \mathbf{P}_{k-1}^+ (the ‘+’ superscript indicates *updated* or corrected values of these variables at step $k - 1$). Based on the intended robot motion \mathbf{u}_{k-1} , the Kalman filter predicts the new state of the system ($\hat{\mathbf{x}}_k^-$), using Equation 1 (the superscript ‘-’ stands for the predicted and as yet *uncorrected* value at step k). Since the motion error \mathbf{v}_{k-1} has a zero mean, this prediction is given by $\hat{\mathbf{x}}_k^- = \Phi_{k-1}\hat{\mathbf{x}}_{k-1}^+ + \mathbf{u}_{k-1}$. The corresponding covariance matrix is updated to $\mathbf{P}_k^- = \Phi_{k-1}\mathbf{P}_{k-1}^+\Phi_{k-1}^T + \mathbf{Q}_{k-1}$. Based on this state prediction and using the sensor model \mathbf{H} from Equation 2, the system predicts the measurement $\hat{\mathbf{z}}_k = \mathbf{H}_k\hat{\mathbf{x}}_k^-$ using an estimate of zero for the measurement error \mathbf{w}_k . This is the sensory input (or observation) the robot *expects to receive* at its new position. Based on actual measurement (or observation) \mathbf{z}_k , the Kalman filter then updates the state estimate and covariance matrix as follows (refer to [Gelb, 1974] for details of the derivation):

$$\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- + \mathbf{K}_k(\mathbf{z}_k - \hat{\mathbf{z}}_k) \quad (3)$$

$$\mathbf{P}_k^+ = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_k^- \quad (4)$$

where $\mathbf{K}_k = \mathbf{P}_k^-\mathbf{H}_k^T(\mathbf{H}_k\mathbf{P}_k^-\mathbf{H}_k^T + \mathbf{R}_k)^{-1}$ is referred to as the *Kalman gain* and $\mathbf{z}_k - \hat{\mathbf{z}}_k$ is called the *innovation*. This Kalman filter formulation is recursive and the process simply repeats as the robot moves. Kalman filter can thus be described by the schematic shown in Figure 4.

Insert Figure 4 near here

It can be shown that the Kalman filter update expressions shown above correspond to an estimation process that is *optimal* (minimum variance, maximum likelihood, etc.), provided the system (Φ) and measurement (\mathbf{H}) models are linear and the noises \mathbf{v} and \mathbf{w} are *Gaussians* with the following properties [Jazwinski, 1970; Gelb, 1974; Maybeck, 1990].

1. *Zero mean*: $E(\mathbf{v}_i) = 0$ and $E(\mathbf{w}_j) = 0$.

2. *White*: $E(\mathbf{v}_i \mathbf{v}_j^T) = \delta_{ij} \mathbf{Q}_i$ and $E(\mathbf{w}_i \mathbf{w}_j^T) = \delta_{ij} \mathbf{R}_i$, where δ_{ij} is the *Kronecker delta function* ($\delta_{ij} = 1$ iff $i=j$, and 0 otherwise) and \mathbf{Q}_i is the covariance matrix of the random noise \mathbf{v}_i and \mathbf{R}_i is the covariance matrix of the random noise \mathbf{w}_i .
3. *Uncorrelated*: $E(\mathbf{v}_i \mathbf{w}_j) = 0$ for all i, j ;

7 Hippocampal Kalman Filtering

It can be observed from Figures 3 and 4 that our model of hippocampal function and the Kalman filter both share the same *predict-observe-match-update* principle. Further, Kalman filter provides a framework for performing stochastically optimal updates even in the presence of prediction and observation errors. Since our goal is to develop a framework for uncertain information fusion in our hippocampal spatial learning model, it is interesting to explore whether hippocampal function could be described in terms of the Kalman filter. If so, one can use the Kalman filter to realize computational models of hippocampal function. In what follows, we will develop hippocampal localization model based on the Kalman filter.

7.1 State Vector Representation

In Kalman filter based approaches to robot localization the spatial map is usually represented in the form of a state vector that contains the position of the robot and the positions of landmarks in the robot’s environment. Since spatial localization in the hippocampus appears to be based on *place recognition*, we have to use a spatial map representation that somehow characterizes these places. As we have mentioned earlier, the CA3 layer in our computational model *learns* internal place-codes while the CA1 layer *learns* the *centers* of these places in terms of the animal’s dead-reckoning estimates. Thus, a natural candidate for spatial map representation in our model is the use of a state vector composed of the estimated *centers* of places encountered by the animal and represented in the CA1 layer. To keep the discussion simple, we will henceforth assume that place codes in CA3 and CA1 are represented by single units, although in reality ensembles of units are known to code for place [Wilson & McNaughton, 1993]. Thus, the state vector at time step k is given by:

$$\mathbf{x}_k = [x_{0,k}, x_1, \dots, x_n]^T \tag{5}$$

where $x_{0,k}$ denotes the position of the animal, x_i denotes the center of place field i , and n is the number of places that have been visited by the animal. In this paper we assume that these positions (animal position and place field centers) are specified in 2D Cartesian coordinates, i.e., $x_i = (x_{i_x}, x_{i_y})$. We do not consider the *orientation* of the animal (or a place) for two reasons. First, this simplifies the model. Second, our computational model currently does not have a mechanism for learning and updating the orientation or head-direction estimates in rodents.

The covariance matrix associated with this state vector, denoted by \mathbf{P}_k , is given by:

$$\mathbf{P}_k = \begin{pmatrix} \mathbf{C}_{00} & \mathbf{C}_{01} & \dots & \mathbf{C}_{0n} \\ \mathbf{C}_{10} & \mathbf{C}_{11} & \dots & \mathbf{C}_{1n} \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \mathbf{C}_{n0} & \mathbf{C}_{n1} & \dots & \mathbf{C}_{nn} \end{pmatrix}$$

where

$$\mathbf{C}_{ij} = \begin{pmatrix} C_{i_x j_x} & C_{i_y j_x} \\ C_{i_x j_y} & C_{i_y j_y} \end{pmatrix}$$

denotes the covariance between the 2D Cartesian representations of the state elements $x_i = (x_{i_x}, x_{i_y})$ and $x_j = (x_{j_x}, x_{j_y})$.

When a new place is visited, the state vector is augmented by the center of this new place and the state estimate and its covariance matrix are modified accordingly.

7.2 The System Model

As described earlier, the system model in the Kalman filter is used to predict the new state vector based on robot motion. Importantly, the state prediction is based on the *intended* motion of the robot, while the *actual* motion is influenced by a host of environmental factors (e.g., friction, wheel-slippage, etc.). Thus, errors appear in the actual motion of the robot, not the prediction. In contrast, the prediction of the animal's position in our model is assumed to be provided by *active* dead-reckoning, i.e., the position is estimated based on the actual motion of the animal (by integrating acceleration and velocity signals). Any errors in actual animal motion are then assumed to be appropriately *tracked* by the dead-reckoning system. However, dead-reckoning is itself error prone, which leads to the following modified system models:

$$\begin{aligned} \mathbf{x}_k &= \Phi_{k-1} \mathbf{x}_{k-1} + \mathbf{u}_{k-1} && \text{Actual animal motion} \\ \hat{\mathbf{x}}_k^- &= \Phi_{k-1} \hat{\mathbf{x}}_{k-1}^+ + \mathbf{u}_{k-1} + \mathbf{v}_{k-1} && \text{Dead-reckoning estimate} \end{aligned}$$

Here, \mathbf{u}_{k-1} is the actual motion of the animal (including its intended motion and motion errors). The dead-reckoning system produces a position estimate based on the animal’s actual motion \mathbf{u}_{k-1} with dead-reckoning errors being characterized by a zero-mean white noise \mathbf{v}_{k-1} with covariance matrix \mathbf{Q}_{k-1} . As explained earlier, this formulation is different from the one used in Kalman filter based robot localization (Section 6.1) because the error is now in the prediction rather than in the motion.

In our model, place field centers are learned in terms of their position relative to the origin of the animal’s dead-reckoning coordinate system. We also assume that when the animal moves, only its dead-reckoning position estimate changes. The place field centers are not directly affected by animal motion. These assumptions lead to the following simplification of the system model: $\mathbf{u}_k = [u_k, 0, \dots, 0]^T$ and $\mathbf{v}_k = [v_k, 0, \dots, 0]^T$. Further, if we assume that the animal only translates in the direction it is facing or turns in the same place, we can take $\Phi_k = \mathbf{I}$, where \mathbf{I} is the identity matrix, leading to the following *linear* system models for animal motion and dead-reckoning:

$$\begin{aligned} x_{0,k} &= x_{0,k-1} + u_{k-1} && \text{Actual animal motion} \\ \hat{x}_{0,k}^- &= \hat{x}_{0,k-1}^+ + u_{k-1} + v_{k-1} && \text{Dead-reckoning estimate} \end{aligned}$$

For simplicity, we also assume that the variance of the dead-reckoning error is the same along both the coordinate dimensions of the 2-D spatial representation. Further, errors along the two dimensions are considered independent (their covariance is zero).

7.3 The Measurement Model

The measurement model in the Kalman filter specifies the relationship between the state vector and the measurements the robot would make when in that state. In Kalman filter approaches for robot localization, \mathbf{H} is referred to as a *sensor model* since it allows one to determine the sensory inputs the robot would receive at any given place. The robot then predicts and observes sensory inputs, and updates the system state based on the resulting mismatch. However, such sensor models are hard to specify as they implicitly assume that the robot knows the locations of different sensory features (objects) in the environment. We will have to explore other alternatives for hippocampal Kalman filtering.

Since our model stores metric positions of place field centers in the CA1 cells, one possibility is to somehow predict and observe metric positions, thereby performing matches in *position space* rather than *sensory space*. For instance, we could use sensory inputs to determine the place i_k

where the animal is currently located, and treat \hat{x}_{i_k} (the *predicted* center of this recognized place) as the observation. The dead-reckoning generated position estimate, $\hat{x}_{0,k}$, could then be taken as the prediction, and these two quantities could be matched. However, this is not feasible because the observation (\hat{x}_{i_k}) and prediction ($\hat{x}_{0,k}$) are correlated (since \hat{x}_{i_k} is some previous value of $\hat{x}_{0,k}$) and the Kalman filter requires system and measurement errors to be uncorrelated.

Another possibility is to take the measurement model to be: $z_k = x_{0,k} - x_{i_k} + w_k$ which represents a vector from the center of place i_k , where the animal finds itself at time step k , to its current position ($x_{0,k}$). Although we can predict this measurement based on the estimated values of $\hat{x}_{0,k}$ and \hat{x}_{i_k} , we cannot observe z_k since we do not know the true center of the place field (x_{i_k}) or the exact position of the animal ($x_{0,k}$) in the environment. Indeed, that is the very reason why we require a spatial learning mechanism.

However, this problem can be circumvented by specifying a measurement function that always observes $z_k = x_{0,k} - x_{i_k} + w_k = 0$, which is equivalent to saying that the measurement model always observes the animal at the center of the corresponding place field. This measurement function constrains the form of the random error to $w_k = x_{i_k} - x_{0,k}$, as shown in Proposition 3 in Appendix B. Further, Proposition 4 shows that w_k has zero-mean provided the animal navigates randomly between place fields or moves from one place field center to another. Thus, the measurement error has zero mean, as required by the Kalman filter. However, it turns out that this error is *autocorrelated* and hence is not a white sequence (Proposition 5). This autocorrelation is difficult to measure, and though we could use extended state vectors to estimate this correlation, leading to an augmented Kalman filter [Jazwinski, 1970], we choose to ignore it. The justification for this is provided by the navigation behavior of the animal. If the animal moves to random positions in the place fields this autocorrelation will be zero, since the current displacement of the animal from the center of the current place field will not be related to its displacement in the previous step (due to the intervening random motion). On the other hand, if the animal moves purposefully from one place field center to the other, the correlation will be very small since the magnitude of the errors (w_k) can be expected to be small. Since it is reasonable to assume that the animal moves randomly during exploration and subsequently from one place field center to the next, we can ignore this autocorrelation term in our computations. With this caveat, it can be verified that the measurement error is an uncorrelated, white Gaussian noise with zero mean. Thus, this

measurement function can be used in the Kalman filter framework.

Thus, if the sensory input at a given place corresponds to a place i_k , we use $\mathbf{H}_k \mathbf{x}_k = x_{0,k} - x_{i_k}$ and $\mathbf{w}_k = w_k$ to obtain the measurement model and predicted measurements as:

$$\begin{aligned} z_k &= x_{0,k} - x_{i_k} + w_k = 0 && \text{Observed measurement} \\ \hat{z}_k &= \hat{x}_{0,k}^- - \hat{x}_{i_k}^- && \text{Predicted measurement} \end{aligned}$$

The covariance matrix \mathbf{R}_k associated with this measurement error is given by:

$$\mathbf{R}_k = \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}$$

where $\sigma^2 = R^2 \sigma_S^2 - 2\sigma_L^2 \ln(\Gamma_{EC})$, as shown in Proposition 6 in Appendix B.

7.4 Update Expressions

Given the state vector representation and the system and measurement models described above, we can easily verify that the requirements of a Kalman filter (i.e., linear models and zero-mean, white, uncorrelated Gaussian errors) are satisfied. Thus, the Kalman filter update expressions of Section 6.1 can be directly applied to our computational model. Hippocampal Kalman filtering then proceeds as follows:

1. Based on the actual animal motion, the dead-reckoning system generates a position estimate at time step k .

$$\hat{x}_{0,k}^- = \hat{x}_{0,k-1}^+ + u_{k-1} + v_{k-1}$$

The elements of the covariance matrix associated with the state are also updated as follows:

$$\mathbf{P}_k^- = \mathbf{P}_{k-1}^+ + \mathbf{Q}_{k-1}$$

where \mathbf{Q}_{k-1} is the covariance matrix of the dead-reckoning error and is given by:

$$\mathbf{Q}_{k-1} = \begin{pmatrix} \mathbf{q}_{k-1} & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & \dots & 0 \end{pmatrix}$$

The covariance update then reduces to the following:

$$\begin{aligned} \mathbf{C}_{00}^- &= \mathbf{C}_{00}^+ + \mathbf{q}_{k-1} \\ \mathbf{C}_{ij}^- &= \mathbf{C}_{ij}^+ && \forall i, j \text{ not both } 0 \end{aligned}$$

2. Based on its sensory inputs, suppose the animal identifies its current place as i_k . Using our model of hippocampal Kalman filtering, $z_k = 0$ and $\hat{z}_k = \hat{x}_{0,k}^- - \hat{x}_{i_k}^-$.
3. By matching z_k and \hat{z}_k , we update the dead-reckoning position estimate and its variance as follows.

$$\begin{aligned}\hat{x}_{0,k}^+ &= \hat{x}_{0,k}^- - (\mathbf{C}_{00}^- - \mathbf{C}_{0i_k}^-)(\mathbf{C}_{00}^- - 2\mathbf{C}_{0i_k}^- + \mathbf{C}_{i_k i_k}^- + \mathbf{R}_k)^{-1}(\hat{x}_{0,k}^- - \hat{x}_{i_k}^-) \\ \mathbf{C}_{00}^+ &= \mathbf{C}_{00}^- - (\mathbf{C}_{00}^- - \mathbf{C}_{0i_k}^-)(\mathbf{C}_{00}^- - 2\mathbf{C}_{0i_k}^- + \mathbf{C}_{i_k i_k}^- + \mathbf{R}_k)^{-1}(\mathbf{C}_{00}^- - \mathbf{C}_{0i_k}^-)^T\end{aligned}$$

where \mathbf{R}_k is the covariance matrix of the measurement error.

4. Also, each place field center m , and its associated covariances are also updated as follows:

$$\begin{aligned}\hat{x}_m^+ &= \hat{x}_m^- - (\mathbf{C}_{m0}^- - \mathbf{C}_{mi_k}^-)(\mathbf{C}_{00}^- - 2\mathbf{C}_{0i_k}^- + \mathbf{C}_{i_k i_k}^- + \mathbf{R}_k)^{-1}(\hat{x}_{0,k}^- - \hat{x}_{i_k}^-) \\ \mathbf{C}_{mn}^+ &= \mathbf{C}_{mn}^- - (\mathbf{C}_{m0}^- - \mathbf{C}_{mi_k}^-)(\mathbf{C}_{00}^- - 2\mathbf{C}_{0i_k}^- + \mathbf{C}_{i_k i_k}^- + \mathbf{R}_k)^{-1}(\mathbf{C}_{n0}^- - \mathbf{C}_{ni_k}^-)^T\end{aligned}$$

Thus, based on the match between the predicted and observed measurements, all the place field centers and their covariances are suitably updated.

7.5 Distinguishing Perceptually Similar Places

Often, different locations in the environment produce the same sensory input (perceptual aliasing) and we need mechanisms to handle such cases. In our model, we suggested the possibility that these aliasing problems in CA3 are resolved by CA1 using dead-reckoning information. It appears that we can use another tool from robotics to elegantly make such distinctions, namely, the *Mahalanobis distance* [Ayache & Faugeras, 1987]. Mahalanobis distance is a metric that computes the difference between predicted and observed values and normalizes them by their covariance. This distance measure has a χ^2 distribution with q degrees of freedom where q is the rank of the covariance matrix [Ayache & Faugeras, 1987; Crowley, 1995]. Since the errors in our system are assumed to be Gaussian and the system equations are linear, we can use the Mahalanobis test to perform matches [Ayache & Faugeras, 1987].

Suppose the animal is currently at a position $x_{0,k}$ and the sensory inputs activate a CA3 place code. Let us assume that this place code is associated with a CA1 code i_k with estimated place field center \hat{x}_{i_k} . Given the current estimate of the animal's position $\hat{x}_{0,k}$, we perform the following

test:

$$(\hat{x}_{0,k}^- - \hat{x}_{i_k}^-)^T (\mathbf{C}_{i_k i_k}^- + \mathbf{C}_{00}^- - 2\mathbf{C}_{0i_k}^- + \mathbf{R}_k)^{-1} (\hat{x}_{0,k}^- - \hat{x}_{i_k}^-) < \epsilon \quad (6)$$

where $(\mathbf{C}_{i_k i_k}^- + \mathbf{C}_{00}^- - 2\mathbf{C}_{0i_k}^- + \mathbf{R}_k)$ is the covariance between prediction $(\hat{x}_{0,k}^- - \hat{x}_{i_k}^-)$ and observation (0) (Proposition 7) and ϵ is a threshold that is chosen appropriately (Proposition 8).

If Equation 6 holds, we assert that the current place has indeed been visited before and that the CA1 unit i_k represents the place. However, if this test fails, it implies that the animal is now at a *new* place that perceptually resembles place i_k visited earlier. In this case, we recruit a new place cell in the CA1 layer and include its parameters in the state vector. Thus, our system recruits multiple units in the CA1 layer that respond to the same sensory input in the CA3 layer but are tuned to different centers.

Thus, the hippocampal Kalman filtering model described above, learns and updates a stochastic spatial map of the environment. In addition to appropriately handling the uncertainties in the sensory and dead-reckoning information streams, it also uses the Mahalanobis distance measure to disambiguate perceptually similar places. We will now describe an implementation of this model and present results of various simulation experiments.

8 Implementation

The localization model described above was implemented on a simulated robot. In this section we discuss the implementation of the proposed computational model, including the algorithms used for place learning and localization.

8.1 Sensory Inputs

We assume that our robot is endowed with a set of sensors that provide information about landmarks in the environment along with their positions relative to the robot. These sensors are *virtual* and can be implemented on real robots through a variety of physical devices. In our simulations the landmarks are assumed to be recognized accurately. However, the range information is assumed to be imprecise and the position of the landmark relative to the robot, (x_l, y_l) , is considered to be corrupted by a zero-mean, white Gaussian noise with standard deviation σ_S per unit distance (along each of the two coordinate dimensions).

8.2 Vector Representations in EC Cells

In our model, EC place cells function as *spatial filters*, responding to particular kinds of landmarks at specific positions relative to the animal. We use an EC representation where each cell stores some internal description of the landmark to which it responds, as well as the Cartesian position of the landmark with respect to the animal. This landmark position is considered *allocentric*, i.e., it is not dependent on the direction in which the robot is facing. This stored information is matched against incoming sensory information. While the landmark identity is matched directly, the currently observed landmark position is matched with the stored position using a *two-dimensional Gaussian function* with variance σ_L^2 . Thus, given that the landmark descriptions match, each EC cell fires over a roughly circular region around the stored position (called the EC field center) with the firing being Gaussian in nature (stronger near the center and gradually decreasing towards the periphery). The algorithm for determining EC layer firing is given below:

1. Set the EC output to zero, i.e., \forall EC cells i , $ECOutput[i]=0$.
2. For each landmark L that is currently sensed at a position (x_L, y_L) relative to the robot do:
 - (a) Determine the activation, $ECAct[i]$, of each EC cell i by the matching procedure described above.
 - (b) If none of the cells produce an activation greater than a threshold (Γ_{EC}), recruit a new EC cell j . Set its landmark description to the identity of landmark L and its field center to (x_L, y_L) . Set $ECOutput[j]=1.0$. Exit.
 - (c) Else, determine each cell i that fires above the threshold Γ_{EC} .
Set $ECOutput[i] = \max \{ECOutput[i], ECAct[i]\}$.

Thus, EC units are created when landmarks are observed in positions not seen before. Also, each EC cell responds to landmarks in an ellipsoidal region about the EC center. This region is roughly circular in our simulations with radius $r \leq \sqrt{-2\sigma_L^2 \ln(\Gamma_{EC})}$, as shown in Proposition 1.

8.3 Place Codes in CA3

For simplicity, our implementation assumes that the place codes in CA3 are given by single units rather than ensembles of them. Each CA3 cell is connected to EC cells that are active in a given

place. The algorithm used to set up CA3 cells is given below:

1. Determine the output of the EC layer using the algorithm described in Section 8.2.
2. Compute the output of each CA3 cell i as follows:

$$\text{CA3Output}[i] = \sum_j \text{ECOutput}[j] \times w_{ij},$$
 where w_{ij} is the weight between EC unit j and CA3 unit i .
3. If none of the CA3 units produce an activation greater than Γ_{CA3} , recruit a new CA3 unit i and connect it to each EC unit j with a weight given by:

$$w_{ij} = \frac{\text{ECOutput}[j]}{\sum_k \text{ECOutput}[k]}. \text{ Exit.}$$
4. Else, determine the CA3 unit k that has the highest activation of all and declare it the *winner*.

Thus, a place is said to be recognized if the winning CA3 place unit produces an activation greater than the threshold Γ_{CA3} . If all the CA3 units have an activation below the threshold, the robot is said to be in a new place and a new CA3 unit is created.

8.4 Unique Place Recognition in CA1

Each CA1 unit in our model is connected to a CA3 unit and to dead-reckoning information derived from the EC. Note that for the sake of simplicity, we again use single units to denote the place cells of CA1. The algorithm for determining CA1 activation is as follows:

1. If a new CA3 unit was created, recruit a new CA1 unit and connect it to the newly created CA3 unit with a weight of 1. Also, store the x and y coordinates of the current position estimate from dead-reckoning with the new CA1 unit. Exit.
2. Else, for each CA1 unit i connected to the winning CA3 unit, compute the Mahalanobis distance between the current position estimate from dead-reckoning and the position estimate stored at unit i .
3. Determine the winner, i.e., the CA1 unit j with the least Mahalanobis distance.
4. If this distance is less than Γ_{CA1} update the position estimate as well as the stored place field centers using the Kalman filter update rules. Exit.

5. Else, create a new CA1 unit and connect it to the winning unit in the CA3 layer with a weight of 1. Store the current dead-reckoning position estimates with the new CA1 unit.

Thus, CA1 units are created when the animal visits a new place or if the Mahalanobis distance of the best matching place is greater than the threshold Γ_{CA1} . This scheme leads to multiple CA1 units connected to the same CA3 unit, one for each of the perceptually aliased regions represented by the CA3 unit.

9 Experiments

This section provides details of the simulation environment used and the experiments that were performed to evaluate the spatial learning model developed in this paper.

9.1 Experimental Setup

The simulation environment for our experiments consists of a largely open room of size 20×20 units, with impenetrable, delimiting walls. The room also contains six identical landmarks, as shown in Figure 5. In our experiments the robot represents landmarks in its spatial map but not the walls. Consistent with the inputs to the animal hippocampus, we assume that detection and recognition of landmarks are performed elsewhere and only recognized landmark information is provided to the spatial learning system. Further, in most experiments, the sensor ranges are assumed to be limited, which makes some of the landmarks undetectable from certain positions in the room. For instance, with maximum sensing range limited to 10 units, the robot is capable of only detecting landmark 1 or 2 in the shaded regions shown in Figure 5. Since all the landmarks are identical, these two regions provide one example of perceptual aliasing in this environment. However, it should be noted that increasing the maximum sensing range (e.g., to 30 units) removes this aliasing.

Insert Figure 5 near here

The robot is assumed to have a compass (equivalently, the animal is assumed to have a faithful head-direction estimate) that provides the *allocentric* bearings (with respect to true north or the origin of the head-direction system) of the sensed landmarks. In addition to recognizing the

landmarks and their bearings, the robot sensors also provide estimated distances to the sensed landmarks. However, this range sensing is assumed to be erroneous and the error is modeled as a zero-mean Gaussian with standard deviation $\sigma_S = 0.01$ per unit distance. This range error of 1% per unit distance is compatible with contemporary sonar-based distance ranging [Everett, 1995].

As described earlier, the sensory information indicating the range and allocentric bearings of detected landmarks are provided as input to the EC layer. The activations of the EC, CA3, and CA1 layers are then computed using the algorithms described in Section 8. Our simulations used a value of $\sigma_L^2 = 1.0$ as the variance of the distance-matching Gaussians of the EC units. The threshold of the EC units was taken to be $\Gamma_{EC} = 0.6$. As shown in Proposition 1, these choices allow the EC cells to respond to landmarks within a circular region of radius 1.01 units from the EC field center. The activation threshold of CA3 units was also taken to be $\Gamma_{CA3} = 0.6$. Given the activation algorithm of units in the CA3 layer (Section 8.3), this choice produces CA3 place fields of equal or smaller size than the EC fields. Finally, the threshold for the Mahalanobis distance test in the CA1 layer was taken to be 4.61. The rationale for this is provided in Proposition 8.

The robot is introduced into this environment at some position (usually randomly chosen) with a random orientation. The dead-reckoning estimates of the robot are set to $(0, 0)$ signifying that the origin of dead-reckoning corresponds to this point of first introduction of the robot in the arena. The robot is then allowed to perform an *a-priori* fixed number of sensing, learning, and action steps. At each such step, the robot obtains sensory inputs as described earlier and executes a movement action. Elsewhere we have shown results of a robot following a circular trajectory [Bousquet *et al.*, 1998]. In the experiments reported in this paper, the robot executes a random-walk, i.e., it turns in a random-direction and moves forward in that direction by `motionStep` = 1 unit. However, these attempted motions are considered erroneous with the motion errors being modeled in the form of zero-mean Gaussians with $\sigma_M = 0.5$. The robot is assumed to have an *active dead-reckoning* mechanism to track its *actual* motion rather than its *intended* one. The dead-reckoning errors are modeled in the form of zero-mean Gaussians with $\sigma_D = 0.1$.

Once the robot has learned a place map it can be removed from the environment (kidnapped) and reintroduced at some other arbitrary place. If the robot has learned a reasonably accurate place map, i.e., it has explored the environment well, it must be capable of relocalizing accurately. However, if the environment contains perceptually aliased places the robot will have problems in

relocalizing. It must then move in a fashion so as to resolve the ambiguity as quickly as possible. Although some techniques have been developed for this problem (e.g., [Dudek *et al.*, 1995; Schuierer, 1997] for polygonal map representations), we have chosen to let the robot simply continue its random-walk until it finds a perceptually unique place at which to localize.

The robot relocalizes as follows. When it is kidnapped, the robot sets the variance of its position estimate to ∞ , signaling a complete distrust of its dead-reckoning position estimate. It also zeros the covariances between its position estimate and the place field centers that it has learned, thereby accepting its learned map as an accurate source of information. Upon reintroduction in the environment, it performs a random-walk. At each step it obtains sensory inputs and applies them to the EC layer of its hippocampal model. Using the algorithms described in Section 8, the activations of EC and CA3 layers are then computed. If a single CA3 unit has an activation above the threshold Γ_{CA3} and the rest have activations below the threshold, then the robot is *potentially* at a perceptually *unique* place. Further, if this CA3 winner is associated with only one CA1 unit, then the robot is at a perceptually unique place. The robot then localizes using the Kalman filter mechanism, thereby effectively updating its dead-reckoning position estimate to the center of the place where it is at that time. However, if either of these conditions are not true, i.e., multiple CA3 units produce above-threshold activation or the CA3 winner is associated with multiple CA1 units, then the robot is experiencing perceptual aliasing and it does not localize. It simply continues its random motion.

We now describe the experiments performed with this setup.

9.2 Spatial Learning in the Presence of Perceptual Aliasing

We will now present results that demonstrate the ability of our model to distinguish between perceptually aliased places and learn reliable spatial maps via exploration. In this experiment the maximum range of the robot sensors was limited to 10 units. As explained earlier, this creates perceptual aliasing in two regions of the simulation environment as shown in Figure 5.

In the exploration phase, the robot was introduced at perceptually aliased region B as shown in Figure 6. The picture on the left in Figure 6 shows the place field centers in the CA1 layer and the circles around the position estimates represent the 2.5σ boundary, where σ^2 is the variance of the corresponding place field estimate. The rationale for displaying the 2.5σ boundary is discussed in

Proposition 9. The actual position of the robot is shown by “▲” with the orientation of the triangle denoting the robot orientation. Further, the dead-reckoned estimate of the robot’s position is shown by “◇” with the dashed-circle denoting the 2.5σ boundary (where σ is the variance of this position estimate). The picture on the right in Figure 6 shows the *place map*, i.e., the firing fields of all the place units created in the CA1 layer. Since this figure shows the state of the system after one motion step, only one CA1 unit has been created thus far and the place map contains only one CA1 field. This CA1 unit (and the corresponding CA3 unit) are activated by just one EC unit which responds to the sole landmark visible from the robot’s start position, namely landmark 2. Thus, the size and shape of this CA1 field is the same as that of the corresponding CA3 field, which is same as the field of the sole EC unit.

Insert Figure 6 near here

Figure 7 shows the result of spatial learning after 10 random moves by the robot. Due to the presence of dead-reckoning errors the difference between actual and estimated positions of the robot starts increasing, as can be observed in the picture on the left. Further, the variance of the robot’s position estimate (denoted by the dashed circle around “◇”) and the variances of successively created CA1 units (gray circles) can be seen to increase. This is because the Kalman filter reduces variances of the estimates only when the robot revisits familiar places, and the robot has not done that in its current trajectory. The place map in Figure 7 also shows that the firing fields of units can vastly differ in shape and size. This is because CA3 (and hence CA1) fields are conjunctions of the firing fields of EC units. Thus, if multiple EC units are active at any given place, the corresponding CA3 (and hence CA1) fields will be an *intersection* of these EC firing fields and as a result, may be much smaller than individual EC fields.

Insert Figure 7 near here

Figure 8 shows the state of the system after 100 steps of exploration. As can be readily observed, there is a significant drift between the true position of the robot and its dead-reckoned position estimate. Further, the variance of the position estimate is also quite large. This is largely because the random-walking robot has not revisited places it witnessed early in its trajectory and Kalman

filter update mechanisms are unable to kick in. Also worth noting is the fact that the robot is now in region A, which is perceptually identical to the region where the robot originally started its motions (region B). However, the Mahalanobis test ensures that this perceptually familiar place is classified as a new (and different) location than the robot's starting position. Thus, the learning system successfully distinguishes between such perceptually aliased places based on its current *confidence* (measured in terms of the reciprocal of its variance) in its position estimate.

Insert Figure 8 near here

Specifically, CA3 unit 0 was created by the system to respond to the sensory inputs available at its starting position in region B. CA1 unit 0 was created to represent this starting position in terms of dead-reckoning coordinates. As can be observed from Figure 9, CA3 unit 0 also responds to the sensory inputs at the robot's current position in the perceptually aliased region A. However, the Mahalanobis test fails during the comparison of the place field center of CA1 unit 0 and the robot's current dead-reckoning estimate in region A, leading to the creation of a new CA1 unit numbered 61. Thus, perceptual ambiguities in the sensory space, reflected in the firing of CA3 unit 0, are resolved through the use of dead-reckoning information in the CA1 layer. (Lighter shades in Figure 9 correspond to regions of stronger firing.)

Insert Figure 9 near here

Figure 10 shows the state of the system after 200 exploration steps. It can be observed that the drift between estimated and true robot position has increased even further. What is more important is that the system also realizes this and reflects it by the corresponding increase in the variance of this estimate (large dashed-circle). Thus, the robot acknowledges the uncertainty in the estimate of its own position. However, the random-walking robot seems to be headed towards the region where it first began its random exploration. Since it has comparatively more accurate place field estimates in this region, revisiting places here will allow the robot to perform significant

corrections of its position estimate.

Insert Figure 10 near here

Figure 11 shows the state of the system at the end of its training (or exploration) phase i.e., the completion of 250 motion steps. It can be seen that the robot has revisited places encountered in the early part of its exploration. At such locations, Kalman filter updates lead to significant corrections of the position estimate of the robot and the place field centers (as well as their associated variances and covariances). Thus, not only do the variance circles shown in the picture on the left in Figure 11 move (with changes in place field centers and position estimate) but they also shrink (decrease in variance). Notice the significant improvement in the accuracy of the dead-reckoning estimate of the robot. With frequent visits to very familiar places the robot can update its place map to arbitrary levels of accuracy.

Insert Figure 11 near here

The trajectory followed by the random-walking robot during this exploration phase is shown in Figure 12. Here, “•” denotes the starting position of the robot and “o” denotes the final robot position.

Insert Figure 12 near here

Thus, our model allows the robot to learn places and their metric positions, distinguish between perceptually similar places, and learn fairly accurate place maps. Further, the model allows the robot to effectively correct its estimates when it revisits familiar places. Importantly, the model allows the robot to *autonomously* learn and update these spatial maps with little input from the user.

9.3 Relocalization in the Presence of Perceptual Aliasing

We demonstrated the ability of our spatial learning and localization model in learning accurate metric place maps even in the face of perceptual aliasing. We will now show the ability of our model to *localize* when introduced into familiar environments.

In this experiment the robot was first allowed 250 steps of exploration and map-building. In this phase it learned and updated places, as described in the previous section. Once the exploration phase was complete, the robot was kidnapped and introduced at the perceptually aliased region A. The variance of the position estimate of the robot was set to ∞ to signal a complete distrust of its dead-reckoning position estimates. As explained earlier, the robot moves randomly until a perceptually *unique* place is found. A place is deemed to be unique if only one CA3 unit fires above the threshold and is associated with only one CA1 unit. The robot then uses the Kalman filter to relocalize at that place.

Figure 13 shows the state of the robot after one and five steps of this relocalizing procedure. During this period the robot does not localize because it detects perceptual aliasing via above-threshold firing of multiple CA3 units or by the presence of multiple CA1 units connected to the CA3 winner. It thus keeps moving randomly.

Insert Figure 13 near here

After 10 steps the robot finally reaches a place south-west of landmark 1. This place is perceptually unique and the robot relocalizes, as shown in Figure 14. As can be observed, this localization leads to a correction of the current position estimate of the robot as well as a reduction of the position estimate variance. Further, once the robot has relocalized, it can move around and further improve its position estimate and variance, as reflected in the system state after 50 relocalization steps (Figure 14).

Insert Figure 14 near here

Thus, our model permits the robot to relocalize correctly even in the presence of perceptual aliasing. It must be borne in mind that this approach will not always work. For instance, if the robot encounters only *one* of the two perceptually aliased regions during its exploration, but is introduced at the *other* one after being kidnapped, it will localize incorrectly. This problem may be alleviated to some extent by extending the relocalization strategy to localize the robot only after validating a series of sensorily unique places. However, this approach will lead to other problems, e.g., the inability of the robot to localize in dynamic environments (where some sensory features

have changed), etc.

9.4 Empirical Evaluation of the Spatial Learning Model

We have also performed a number of experiments to empirically evaluate and validate our model of spatial learning. The primary attributes under study were the usefulness or need for spatial learning, the scalability of the system (particularly from the point of view of handling increased sensing and dead-reckoning errors), the effect of these parameters on the numbers of EC, CA3, and CA1 units required, and the representation of space learned by this system. These experiments were performed by choosing 100 random exploration trajectories for the robot. The robot was then allowed to execute 250 exploration steps along each of these trajectories, under the different parameter values constituting the experiment, and the results were recorded. By choosing 100 fixed trajectories for the experiments, we basically guaranteed that the path taken by the robot and the sensory inputs it received during its motion were identical for each of the experiments. Any observed differences in the results are then necessarily due to the particular parameter values used. We present some of these results in the following sections.

9.4.1 Usefulness of Spatial Learning

Our model performs spatial learning and localization. It not only learns places in terms of dead-reckoning position estimates but also updates and corrects the robot's position estimate using the Kalman filter. Is this necessary? Can the robot be expected to perform well without such a learning and localization mechanism?

Figure 15 shows the performance of the robot with and without the spatial learning mechanism. In this case, the error in the final position of the robot, at the completion of its 250 steps of exploration, was the measure of performance. This error was measured as the Euclidean distance between the final *true* and *estimated* position of the robot. As can be observed, without spatial learning and localization, the distance error increases almost linearly with increase in the standard deviation of the Gaussian dead-reckoning error. For instance, with dead-reckoning error of variance 1, the error without learning is approximately 20 units on an average, while the system with learning has a net error of approximately 5 units. This error of 20 units is immense considering the fact that the size of the simulation environment was 20×20 units. Further, it can be readily observed

that the two curves are increasingly divergent.

Insert Figure 15 near here

Similar results are obtained if the sensory range of the robot is set to 30 units, i.e., even if the robot is capable of observing all the landmarks from all positions in the room. However, with increased sensing range, the errors in the final robot position are slightly lesser than earlier, as shown in Figure 16. This is directly attributable to the sensor range, since with a range of 30 units there is no perceptual aliasing in the environment, which allows dead-reckoning drifts to be corrected properly. However, when the sensor range is limited to 10 units, there is perceptual aliasing in the simulation environment, as shown in Figure 5. Under such circumstances large dead-reckoning variances cause the Mahalanobis test to be satisfied for places that are significantly distant but sensorily similar, leading to wrong updates of the robot’s position estimate. This contributes to the increase in the final position error observed in Figure 16 for sensor range 10.

Insert Figure 16 near here

Given the fact dead-reckoning using contemporary, affordable devices is quite erroneous [Everett, 1995], it is clear that the use of our spatial learning and localization mechanism is of immense value. Although our spatial learning model does not eliminate the position errors completely, it significantly reduces the errors by updating the robot’s position estimate using the Kalman filter.

9.4.2 Importance of Updating Place Field Centers

A key feature of the proposed hippocampal Kalman filtering model is its ability to update not only the dead-reckoned position estimate of the robot but also the place field centers of the learned map. The advantages offered by this can be observed in Figure 17, which shows the distance error in the final robot position as a function of the dead-reckoning error, with and without the update of place field centers. Updates can be seen to lead to lower errors. For instance, with $\sigma_D = 0.6$, updates lead to an average distance error of 3 units, while without updates the average error is approximately 4 units. This difference is significant considering the size of the room which is only

20× 20 units.

Insert Figure 17 near here

Thus, the ability to update not only the robot’s position estimate but also the estimated positions of the different places, is an important advantage offered by our computational model. Although other hippocampal models in the literature support metric spatial learning [Redish & Touretzky, 1996; Redish & Touretzky, 1998], they do not have mechanisms to update position estimates or place field centers.

9.4.3 Effect of Range Sensing Errors on the Numbers of Units

We have also explored the effect of increased sensing error on the numbers of EC, CA3, and CA1 units required in our model. As can be observed in Figure 18, the *mean* number of EC units required (averaged over the 100 trajectories), grow almost linearly with the standard deviation of the Gaussian range sensing error. For instance, with a standard deviation of 0.2 the system only requires 150 EC units, while with a deviation of 0.8 it requires 250 EC units. Since the robots followed the same set of 100 trajectories in each case, the increase in the number of EC units required is a direct consequence of the increase in the range sensing error. This is explained by the fact that large sensor range errors produce radically different vectors to the same landmark from the same position. Since the EC fields have a limited firing field, new EC units are recruited to represent these different vectors, leading to their increased numbers.

Insert Figure 18 near here

It can also be observed that the numbers of CA3 and CA1 units increase very slowly beyond a certain level of error. This is because the sensory inputs in the EC layer are still able to activate the same place code in the CA3 layer, and consequently in the CA1 layer. Thus, sensory uncertainty is *absorbed* by the EC layer, and does not affect the scalability of the CA3 and CA1 layers. It can also be seen that there is a difference in the numbers of CA3 and CA1 units. This is due to the sensory range of 10 units, which causes perceptual aliasing in the environment. At such aliased locations, our model recruits multiple CA1 units that respond to the same CA3 unit.

Figure 18 also shows another interesting phenomenon. For range errors less than $\sigma_S = 0.4$,

the number of EC units in the model is less than (or marginally greater than) the number of CA3 and CA1 units, while for larger errors EC units are significantly more in number. This is because with such small errors the fields of EC units created at a given place (each responding to one observed landmark) coincide almost perfectly. Consequently, the CA3 unit place field, which is an intersection of the EC fields, is as large as the individual EC field. Thus, fewer CA3 (and hence CA1) units are required to cover the same region of the environment. In fact, with very small range error values (e.g., $\sigma_S \leq 0.05$) the number of CA3/CA1 units are lesser than the number of EC units largely due to this efficiency of space coverage. It should also be clarified that with sensor range restricted to 10 units, the number of CA1 units required by our model is always greater than the number of CA3 units since many of the 100 trajectories pass through the perceptually aliased regions. Thus, the apparent closeness in the numbers of CA3 and CA1 units for low values of σ_S in Figure 18, is simply an artifact of the scale chosen for representing that axis.

Similar effects are observed when the robot sensors have a range of 30 units, as shown in Figure 19. It should be noticed that the lack of perceptual aliasing in this environment leads to equal numbers of CA3 and CA1 units in the model. As before, the range error is effectively absorbed by the EC layer, with little increase in the numbers of CA3 and CA1 units beyond a certain point.

Insert Figure 19 near here

Thus, the size of the spatial map learned by our model (denoted by CA1 fields), scales quite well with increase in range sensing error.

9.4.4 Efficiency of Spatial Coding

Our spatial learning and localization system automatically develops a *multi-granular* representation of space with frequently visited places (or thoroughly explored regions) being represented at a *finer granularity* than regions where the robot spent a few passing moments. Since spatial maps represented at a finer granularity allow more accurate localization, while coarser spatial representations save on storage space, the ability to automatically develop a multi-granular spatial representation is a valuable property of our model. This multi-granular representation arises as a result of range-sensing errors and their consequent role in the creation of new EC units, the algorithm for determining the activations of CA3 units, the role of the Mahalanobis distance in the creation of

CA1 units, and importantly, the variance updates performed by the Kalman filter. For instance, Figure 20 shows the firing fields of four CA1 units. As can be observed, the fields differ in size with the larger firing fields corresponding to CA1 units created in rarely visited areas and smaller fields in regions quite frequently visited by the robot during its exploration run.

Insert Figure 20 near here

Kalman filter updates provide one way in which this multi-granular representation of space can arise. Suppose the robot is exploring a small region of the environment. It first creates a certain number of place cells for this region. Now, as it keeps revisiting the places in this region, the Kalman filter updates the estimates of the field centers and reduces their variances. As the variances decrease, the Mahalanobis distance measure increases until it fails at some point. Even though the robot is in a familiar place the system now creates another CA1 unit, thereby distributing the firing field of the original place between two place cells. And this process continues.

It should be noted that this change in granularity is an *emergent property* of our system and arises automatically through the experience of the robot in the environment.

10 Related Work

In this section we briefly explore the relationship between the proposed model and approaches to hippocampal spatial learning and robot localization.

10.1 Computational Models of Hippocampus-Based Navigation

Based on neurobiological and behavioral data, a number of computational models of hippocampal spatial learning and navigation have been proposed. These models can be distinguished based on the nature of spatial representation used and the form of navigation behaviors supported by them. Trullier et al. (1997), have distinguished three such classes of approaches: *place-recognition triggered response*, *topological navigation*, and *metric navigation*. A detailed discussion of these approaches is beyond the scope of this paper, but they may be summarized as follows.

In place-recognition triggered response, each place is associated with a navigation response that the animal performs when in that place. Usually the response is associated with a goal and the animal automatically executes the response thereby approaching the associated goal. Since

goals are implicit in these systems, so is the *motivation* [O’Keefe & Nadel, 1978]. A number of hippocampal models of spatial navigation fall into this category [Zipser, 1986; Burgess *et al.*, 1994; Brown & Sharp, 1995; Blum & Abbott, 1996; Burgess & O’Keefe, 1996].

In topological navigation, the animal learns places and associates adjacent places with the motion required to get from one place to the other. In such cases, goal-directed navigation reduces to determining a path from the current place to the place that houses the goal, using a procedure akin to a *graph search*. A number of models of spatial learning have this topological navigation flavor [Muller *et al.*, 1991; Kuipers & Byun, 1991; Mataric, 1992; Schmajuk & Thieme, 1992; Schmajuk & Blair, 1993; Bachevalier & Waxman, 1994; Scholkopf & Mallot, 1995].

Metric navigation approaches allow the animal to learn places in terms of their position in some metric framework computed by the animal. Thereupon, goal-directed navigation requires determining the current metric position of the animal, the position of the goal, and a computation to determine the approach vector from the current position to the goal (vector subtraction). Different approaches to metric navigation are discussed in [Prescott & Mayhew, 1992; Worden, 1992; Prescott, 1995; Wan *et al.*, 1994; Redish & Touretzky, 1996].

For a further comparison of these approaches, the reader is referred to [Trullier *et al.*, 1997; Balakrishnan, 1998]. It should also be noted that there is a close correspondence between these categories and those originally identified by O’Keefe and Nadel (1978), namely, *orientation-taxon hypothesis* (for place-recognition triggered response), and the *locale hypothesis* (for topological and metric navigation).

As can be observed, different aspects of neurophysiological, neuroanatomical, and behavioral data have led to different kinds of locale-based and taxon-based computational models of hippocampal spatial learning. But is hippocampal spatial learning locale-based or taxon-based? Although this is still a much-debated issue, a number of researchers have suggested that the hippocampus represents (or mediates the representation of) *both* locale as well as taxon information [O’Keefe & Nadel, 1978; Redish & Touretzky, 1998]. We too believe that the navigating animal has access to spatial information represented in multiple forms and that it chooses an appropriate navigation strategy (or some combination of strategies), in ways suited to its current navigational goals.

10.1.1 Models of Redish and Touretzky

Among the computational models of hippocampus-based spatial learning and navigation that have been proposed, the model developed in this paper bears the closest resemblance to the *self-localization* model of [Wan *et al.*, 1994; Redish & Touretzky, 1996]. In their computational framework, an exploring animat creates place cells that are tuned to the identities and egocentric positions of *two* randomly chosen landmarks visible from the current place. The place cells also respond to the *retinal* angle between two visible landmarks (possibly different from the two chosen above) and to the dead-reckoning generated position estimate of the animat. Each of these inputs is a Gaussian function and the overall firing of a place cell is simply a product of these Gaussians. Further, the unit firing is *fuzzy*, in the sense that unknown quantities simply drop out of the product. For instance, if the animat is removed from the environment and then reintroduced back in it, it does not have accurate dead-reckoning estimates. In this case, the place cells fire only in response to sensory inputs and the Gaussian function corresponding to the position estimate match simply drops out of the calculations. By storing egocentric positions of landmarks and an angle between two landmarks, their model can additionally perform head-direction resets [Redish & Touretzky, 1996]. It has been used to provide simulation accounts of some key behavioral experiments with gerbils reported in [Collett *et al.*, 1986].

Recently, Redish and Touretzky (1998) have extended this locale-based self-localization model to provide a more comprehensive account of rodent spatial learning and memory. They suggest that the hippocampus participates in both locale as well as *route-based* learning of space. While the locale-based learning component in this model is quite similar to the one in their earlier model [Redish & Touretzky, 1996], the route-based learning component is new. This extension uses the CA3 recurrent collaterals to store specific paths (routes) traversed by the animat, and plays an important role in transferring these routes to *long-term storage* via *slow cortical learning*. This enables them to not only explain the performance of rats in the Morris water maze [Morris, 1981], but also provide accounts of *anterograde* and *retrograde* amnesia caused by hippocampal lesions.

It should be noted that the locale-based self-localization model of Redish and Touretzky (1996; 1998) is closely related to the model developed in this paper since both are based on the locale hypothesis of O’Keefe and Nadel (1978) and on its implicated substrate in the hippocampus. Thus, both these models develop a functional characterization of the hippocampus wherein dead-reckoning

information is combined with sensory inputs to create a metric representation of space. Further, both models assume that sensory inputs are provided to the place learning system in terms of recognized landmark information along with their *vector* position relative to the animal. Finally, both assume that *goals* are represented in terms of their metric positions and goal-directed navigation is simply a matter of localizing in the metric framework and then proceeding to the remembered position of the goal.

However, the two models also differ in some important regards. The model developed in this paper incorporates robust mechanisms for information fusion in the presence of uncertainty in the sensory and dead-reckoning input streams. Further, by formulating the place learning and localization problem within the Kalman filter framework, the model additionally derives expressions that update the position estimate as well as the metric spatial map (place field centers) in a *stochastically optimal* fashion. The proposed model easily lends itself to extensions that support incremental acquisition and assimilation of multiple *local* maps [Balakrishnan *et al.*, 1998b].

10.2 Related Approaches for Robot Localization

Since localization plays a critical role in robot navigation, it is hardly surprising that numerous techniques have been developed for robot spatial learning and localization. A number of these approaches appear closely related to the localization system described in this paper. In the following sections we will briefly discuss and compare some approaches.

10.2.1 Approaches Related to Kalman Filtering

Since robots perceive aspects of their environments through their sensors, any attempt at world-modeling must be preceded by sensory information fusion (also called sensor fusion). Smith, Self, and Cheeseman (1990) argued strongly for the use of Bayesian estimation theory in robot spatial representation (called the *stochastic map*), and showed that an *optimal* information fusion approach was equivalent to a Kalman filter [Kalman, 1960] under certain conditions. However, the linearity assumptions of Kalman filter equations (see Section 6.1) were rather limiting in real-world applications. This led to the development of the *Extended Kalman filter* framework for building 3D maps of the environment [Ayache & Faugeras, 1987]. They use a Taylor-series expansion of the non-linear function and truncate it at the first derivative, thereby obtaining a linear *approximation* of the original function. Moutarlier and Chatila (1989) developed an alternate stochastic information

fusion framework for map-building and showed that Kalman filter was simply a special case. A good discussion of Kalman filter approaches to robot localization can be found in [Crowley, 1995], while the necessity of, and the difficulties associated with, maintaining correlations of the state variables of a stochastic map are detailed in [Hebert *et al.*, 1995].

Kalman filter approaches for robot localization require a *sensor model of the environment* in the measurement function. This sensor model provides the sensory inputs (or measurement) that the robot would receive when in any given position. Conventional approaches use a measurement function that specifies the *relative positions* of landmarks with respect to the robot. In such cases, the landmarks are represented in a *robot-centered* frame and the landmark positions are updated based on robot motion. These updated positions serve as predictions while the observed positions of the landmarks serve as observations in the Kalman filter framework [Smith *et al.*, 1990; Ayache & Faugeras, 1987; Leonard & Durrant-Whyte, 1992; Moutarlier & Chatila, 1989; Crowley, 1995].

However, if the environment contains multiple identical landmarks and sensor ranges are limited, this measurement function leads to matching problems (i.e., which state vector element should be matched against the sensed landmark?) and to consequent localization problems in the presence of perceptual aliasing. These problems also arise when a kidnapped robot has to relocalize. In contrast, the model described and implemented in this paper performs *place-based* localization. Since places are more distinct and distinguishable than relative landmark positions, it is easier to relocalize with a place-based system. Some researchers have recently begun exploring place-based extensions of Kalman filter approaches [Crowley, 1995; Hebert *et al.*, 1995].

Another key distinction of our model is that it does not require an *explicit* sensor model and instead relies on a simple neural system for recognizing places. By labeling places with dead-reckoning inputs, the system develops a kind of *inverse sensor model* that produces a place code (and a position estimate) in response to sensory input. Thus, instead of matching in sensory space our localization system performs matches in position space.

10.2.2 Approaches Based on Cognitive Mapping Theories

A number of cognitive mapping theories of spatial representation have been proposed that espouse organization of spatial information in multiple forms and levels [Kaplan, 1973a; Kaplan, 1973b; Kuipers, 1978]. For instance, Kuipers proposed a three level representation of space called the

spatial semantic hierarchy (SSH) where the lowest level consists of control rules that define *distinctive places* as some property of sensory inputs. The next higher level creates a topological representation by linking distinctive places by *travel edges* and the highest level in the hierarchy contains a *geometric representation* of the robot’s sensory environment [Kuipers & Byun, 1991]. This cognitive spatial learning architecture was implemented on a simulated robot NX [Kuipers & Byun, 1991] and later extended to physical robots [Kuipers *et al.*, 1993].

Kortenkamp’s RPLAN is an implementation of the PLAN (Prototypes, Location and Associative Networks) model of human cognitive mapping [Chown *et al.*, 1995; Kaplan, 1973a; Kaplan, 1973b]. In his implementation, sonar-based detection of *gateways* and vision-based detection of *scenes* are combined using a Bayesian network to provide robust definitions of place. The system then uses these place definitions to construct routes, networks of routes, and global spatial representations involving many such networks [Kortenkamp, 1993].

Such multi-level spatial representations are also found in Qualnav [Levitt & Lawton, 1990]. This simulated robot represents regions of space called *viewframes* at the lowest level. Viewframes are composed of relative angles and estimated distances to landmarks visible from the current place. At the next higher level, pairs of landmarks are used to define a virtual boundary called the *landmark pair boundary* (LPB). These LPBs lead to a topological division of space called *orientation regions*. Together, these two approaches allow for specific localization of the robot using viewframes within more general orientation regions.

The models described here provide implementations of cognitive mapping theories while our model attempts a computational characterization of a particular brain region (the hippocampus). The models differ in this regard. However, it is interesting to observe that the spatial representation hierarchies proposed in these cognitive theories closely resemble the different forms of information purported to exist in the hippocampus (Section 10.1). For instance, our computational model learns a metric representation of space. As we will outline in Section 11, we strongly believe that the CA3 recurrent collaterals encode a *topological* representation of space, while the subiculum and the post-subicular regions learn a *directional* map of the environment. Thus, our view of hippocampal function has the same flavor as the approaches based on cognitive mapping theories.

10.2.3 Neurobiological Models of Robot Navigation

Some neurobiological models of spatial learning have also been implemented on mobile robots. Mataric (1992) implemented a topological *place graph* loosely based on the place unit model of McNaughton and Nadel (1990). In this work, the robot explored the environment and built a topological place graph. The nodes in the graph represented *landmarks* and the links encoded the robot motions required to move between places. Places in the model were associated with individual landmarks like walls, corridors, etc. Further, place units in the graph were associated with a set of instructions that were executed whenever the corresponding unit was activated (i.e., the robot reached the corresponding place). Once the map was built, goal-directed navigation was performed by *spreading activation* backwards from the goal node and following the path to the *strongest activated* neighbor of the current node.

Bachelder and Waxman (1994) also implemented a spatial learning system on a mobile robot called MAVIN (mobile adaptive visual navigator). This robot wandered through its environment recognizing places based on the configuration of landmarks visible from each location. Landmarks were recognized using a Seibert-Waxman neural 3D object recognition system [Seibert & Waxman, 1992] which was trained to recognize the landmarks *prior* to the exploration by the robot. As the robot moved, an associative neural network used a Hebbian rule to learn the movements that led from one place to the other.

Recce and Harris (1996) recently implemented a robot navigation system based on interactions of neocortical and hippocampal theories. In their model, the hippocampus functioned as an *auto-associative* memory for matching *egocentric* maps constructed in the neocortex. The hippocampus stored snapshots of this egocentric map and the activity of head-direction cells was used to determine the best map rotation to match the snapshots stored in the hippocampus.

The approaches described above are implementations of computational accounts of hippocampal spatial learning. However, all of them deal with a topological representation of space unlike the metric ones created in our model. Further, even though the models were implemented on real robots, none of the models had any explicit mechanisms for handling sensory and motion errors. In contrast, our model explicitly fuses uncertain information in a stochastically optimal manner.

11 Discussion

In this paper we have described a computational model of the locale hypothesis of hippocampal spatial learning [O’Keefe & Nadel, 1978]. The hippocampal system learns places and associates them with dead-reckoning estimates, thereby learning a metric map of the environment. This metric representation allows the animal to navigate to arbitrary goals and locations on a direct trajectory. However, in order for such a representation to exist, the dead-reckoning estimates of the animal must be faithful and reliable, even in the face of perceptual aliasing (when multiple places appear sensorily identical). Further, since the sensory and dead-reckoning input streams can have uncertainty associated with them, it is also imperative that the model be capable of handling such uncertainties.

Based on the parallel between the requirements of our computational model and the Kalman filter approach for probabilistic robot localization, we developed a Kalman filter framework for our localization model. Not only does this approach handle uncertain data, but also provides stochastically optimal rules for information updates in the model. The model learns places based on sensory data, and through the Kalman filter process, also learns the centers of different places in terms of dead-reckoning information. Using a related tool, the Mahalanobis distance, the system is also capable of distinguishing between perceptually similar places and learns and localizes in their presence. The place field centers as well as the current dead-reckoning position are represented by estimates (with associated covariances) and are updated appropriately when the animal visits or revisits places. With frequent revisits to a known place, the uncertainty associated with that place field center can be seen to decrease and the estimated place field center converges to the true center. The system also handles relocalization (kidnapped robot problem) very easily by setting the variance of the robot’s position estimate to ∞ and using the same Kalman filter approach.

This computational model makes some interesting contributions to robot localization. Primarily, it provides a place-based extension of the Kalman filter by allowing the robot to learn places as conjunctions of landmark positions and estimating the centers of these places in terms of dead-reckoning coordinates. The spatial map learned using the model is multi-granular in nature, with frequently visited regions being represented at a finer granularity level. Thus, the model offers advantages of both metric as well as relational representations of space discussed in Section 3.1. In addition, the model uses an inverse-sensor model to perform matches in position space rather

than sensory space. Thus, it does not require a sensor-model which is often the prime difficulty in the conventional Kalman filter. Finally, unlike pure sensor-based systems, the model allows the robot to disambiguate between perceptually similar places, and unlike pure dead-reckoning based systems [Yamauchi & Beer, 1996], it can *relocalize* after being kidnapped.

Although we have not shown any results to this effect in this paper, goal-directed navigation is easily realized in our system. Goals encountered by the robot are represented in terms of their locations computed from the dead-reckoned position estimate of the robot and the relative position of the goal. Once such locations are learned, the robot can directly navigate to any goal location by computing a vector-difference of the goal position and its current position. We have developed mechanisms for learning, updating, selecting, and navigating to *multiple* goals within an environment. We have also extended our computational framework to support *incremental* learning and *integration* of local spatial maps. These extensions have allowed us to reproduce a number of behavioral experiments using our computational model, which are detailed in [Balakrishnan *et al.*, 1998b; Balakrishnan *et al.*, 1998a; Balakrishnan *et al.*, 1998c; Balakrishnan, 1998].

Although we have not fully developed it yet, we strongly believe that the hippocampal formation represents spatial information at many different levels. We believe that the CA3 recurrent collaterals use motion information to associate places, thereby capturing a topological description of the environment. We also suspect that the place cells recently discovered in the subiculum, encode egocentric information, making it possible to initialize and reset head-direction estimates in the post-subicular area. These regions can thus be thought of as encoding a *directional-map* of the environment. Some researchers have begun exploring similar hypotheses [Redish *et al.*, 1996; Redish & Touretzky, 1998]

Our model makes several neurobiological and behavioral predictions. For instance, our model implies that perceptually identical places in the same environment will excite the same population of cells in the CA3 layer but different populations in the CA1 layer. However, no differences have been observed in the firings of CA3 and CA1 units. How can we explain this inconsistency? It is our opinion that the behavioral experiments performed to date did not allow perceptual aliasing to occur. Even in rodent experiments in mazes of different kinds, the maze arms are perceptually different if we assume that the animals were not disoriented, i.e., they had accurate head-direction estimates. If environments supporting perceptual aliasing can be designed, our hypothesis can be

tested through appropriate recordings from CA1 and CA3 cells.

Also, in our current model of hippocampal spatial function, CA1 firings are maintained by dead-reckoning in darkness. However, CA3 cells have also been found to fire in darkness but they are not influenced by dead-reckoning in our model. How can the firing of these cells be explained by our model? A number of hypotheses can be forwarded. First, it is possible that the recurrent collaterals encode a topological relationship between places and thus provide a different (second?) dead-reckoning system. Based on animal motion, these recurrent collaterals excite cells that encode the neighboring place. Another explanation could be that CA3 cells associated with CA1 cells at a given place are activated by *top-down* activations and hence fire in darkness. Finally, it is also possible that position information is encoded in cell firings further upstream (e.g., in the EC cells). This is consistent with the observations of Feigenbaum and Rolls (1991), who have found evidence of egocentric as well as allocentric spatial encodings in the primate hippocampus. These explanations need to be explored in more detail.

Our model also predicts that animals can reduce their computational complexity of localization by navigating either randomly or from one place field center to another. As we elaborated in Section 7.3, navigating in this fashion allows the hippocampal Kalman filter to ignore correlations between measurement errors. A related behavioral prediction is that an exploring animal will search in gradually expanding trajectories from its point of origin (or nest), since this allows the propagation of reliable dead-reckoning estimates further and further away from its starting position. These predictions can be verified through appropriately designed experiments.

The suggestion that hippocampus performs computations that can be modeled by a Kalman filter is consistent with the hypothesis that most forms of animal learning is driven by discrepancies between predictions and observations (Cormier, 1989). If indeed the Kalman filter is a reasonable model of this aspect of hippocampal function, it raises the question as to what are the neural substrates of such computation. Based on structure of the hippocampus, O'Keefe (1989) suggested that the CA1 layer is capable of performing matrix inversions through iterative computations. This provides a possible realization of the the matrix inversions required by our model in the the Kalman gain and the Mahalanobis distance calculations within the CA1 layer. Buzsaki (1989) proposed a theory of hippocampal learning according to which *sharp waves* observed in resting rats are indicative of *consolidation* of memory traces via successive firings of CA3 populations affected

by the CA3 recurrent collaterals. This might provide the vehicle for the propagation and update of the state estimates and covariances required by our model.

12 Acknowledgements

We would like to thank the anonymous reviewers for their comments and suggestions that helped improve the content and organization of the paper. We are also thankful for the valuable feedback provided by David Touretzky on an earlier draft of this paper and by Rushi Bhatt on many intermediate ones. This research was partially supported by grants from the National Science Foundation (NSF IRI-9409580) and the John Deere Foundation, to Vasant Honavar and an IBM fellowship to Karthik Balakrishnan. This work was completed during Karthik Balakrishnan's doctoral studies in Computer Science at the Iowa State University Artificial Intelligence Laboratory under the direction of Vasant Honavar. Olivier Bousquet contributed to this work during his visit to Iowa State University during the summer of 1997, which was made possible by Ecole Polytechnique, France.

References

- Anderson, E. (1983). *Animals as Navigators*. Van Nostrand Reinhold Company.
- Ayache, N., & Faugeras, O. (1987). Maintaining Representation of the Environment of a Mobile Robot. *In: Proceedings of the International Symposium on Robotics Research*.
- Bachelder, I., & Waxman, A. (1994). Mobile Robot Visual Mapping and Localization: A View-Based Neurocomputational Architecture that Emulates Hippocampal Place Learning. *Neural Networks*, **7**, 1083–1099.
- Balakrishna, R., & Ghosal, A. (1995). Modeling of Slip for Wheeled Mobile Robots. *IEEE Transactions on Robotics and Automation*, **11**(1), 126–132.
- Balakrishnan, K. (1998). *Biologically Inspired Computational Structures and Processes for Autonomous Agents and Robots*. Ph.D. thesis, Iowa State University, Ames, IA. (forthcoming).
- Balakrishnan, K., Bhatt, R., & Honavar, V. (1998a). A Computational Model of Rodent Spatial Learning and Some Behavioral Experiments. *In: Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society*.
- Balakrishnan, K., Bhatt, R., & Honavar, V. (1998b). Spatial Learning and Localization in Animals: A Computational Model and Behavioral Experiments. *Pages 112–119 of: Proceedings of the Second European Conference on Cognitive Modelling*. Nottingham, UK: Nottingham University Press.
- Balakrishnan, K., Bhatt, R., & Honavar, V. (1998c). *Spatial Learning in the Rodent Hippocampus: A Computational Model and Some Behavioral Results*. (in preparation).
- Barnes, C., McNaughton, B., Mizumori, S., Leonard, B., & Lei-Huey, L. (1990). Comparison of Spatial and Temporal Characteristics of Neuronal Activity in Sequential Stages of Hippocampal Processing. *Progress in Brain Research*, **83**, 287–299.
- Bennett, A. (1993a). Remembering Landmarks. *Nature*, **364**, 293–294.
- Bennett, A. (1993b). Spatial Memory in a Food Storing Corvid. I. Near Tall Landmarks are Primarily Used. *Journal of Comparative Physiology A*, **173**, 193–207.
- Biegler, R., & Morris, R. (1996). Landmark Stability: Studies Exploring Whether the Perceived Stability of the Environment Influences Spatial Representation. *The Journal of Experimental Biology*, **199**(1), 187–193.
- Blair, H., & Sharp, P. (1995). 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.
- Blum, K., & Abbott, L. (1996). A Model of Spatial Map Formation in the Hippocampus of the Rat. *Neural Computation*, **8**, 85–93.
- Bousquet, O., Balakrishnan, K., & Honavar, V. (1998). Is the Hippocampal Formation a Kalman Filter? *In: Proceedings of the Pacific Symposium on Biocomputing '98*.

- Brooks, R. (1985). Visual Map Making for a Mobile Robot. *In: Proceedings of the IEEE Conference on Robotics and Automation.*
- Brown, M., & Sharp, P. (1995). 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.
- Burgess, N., & O’Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*, **6**, 749–762.
- Burgess, N., Recce, M., & O’Keefe, J. (1994). A Model of Hippocampal Function. *Neural Networks*, **7**(6/7), 1065–1081.
- Buzsaki, G. (1989). Two-State Model of Memory Trace Formation: A Role for "Noisy" Brain States. *Neuroscience*, **31**, 551–570.
- Chen, L., Lin, L-H., Green, E., Barnes, C., & McNaughton, B. (1994b). Head-Direction Cells in the Rat Posterior Cortex. I. Anatomical Distribution and Behavioral Modulation. *Experimental Brain Research*, **101**, 8–23.
- Chen, L., Lin, L-H., Barnes, C., & McNaughton, B. (1994a). 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.
- Chown, E., Kaplan, S., & Kortenkamp, D. (1995). Prototypes, Location and Associative Networks (PLAN): Towards a Unified Theory of Cognitive Mapping. *The Journal of Cognitive Science*, **19**(1).
- Churchland, P., & Sejnowski, T. (1992). *The Computational Brain*. Cambridge, MA: MIT Press/A Bradford Book.
- Cohen, N., & Eichenbaum, H. (1993). *Memory, Amnesia, and the Hippocampal System*. Cambridge, MA: MIT Press/A Bradford Book.
- Collett, T., Cartwright, B., & Smith, B. (1986). Landmark Learning and Visuo-Spatial Memories in Gerbils. *Journal of Neurophysiology A*, **158**, 835–851.
- Cormier, S.M. (1989). *Basic Processes of Learning, Cognition, and Motivation*. Mahwah, NJ: Lawrence Erlbaum.
- Crowley, J. (1995). Mathematical Foundations of Navigation and Perception for an Autonomous Mobile Robot. *Pages 9–51 of: Proceedings of the International Workshop on Reasoning with Uncertainty in Robotics*. Springer-Verlag.
- Dudek, G., Romanik, K., & Whitesides, S. (1995). Localizing a Robot with Minimum Travel. *Pages 437–446 of: Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms*.
- Elfes, A. (1989). *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. Ph.D. thesis, Electrical and Computer Engineering Department, Carnegie-Mellon University.

- Elfes, A. (1992). Dynamic Control of Robot Perception Using Multi-Property Inference Grids. *In: Proceedings of the 1992 IEEE International Conference on Robotics and Automation*. IEEE Press.
- Elfes, A. (1995). Robot Navigation: Integrating Perception, Environmental Constraints and Task Execution Within a Probabilistic Framework. *Pages 93–130 of: Proceedings of the International Workshop on Reasoning with Uncertainty in Robotics*. Springer-Verlag.
- Engelson, S. (1994). *Passive Map Learning and Visual Place Recognition*. Ph.D. thesis, Yale University, Department of Computer Science.
- Etienne, A. (1985). The Control of Short-Distance Homing in the Golden Hamster. *Pages 233–251 of: Ellen, P., & Thinus-Blanc, C. (eds), Cognitive Processes and Spatial Orientation in Animals and Man*. Boston, MA: Martinus Nijhoff.
- Etienne, A. (1992). Navigation of a Small Mammal by Dead Reckoning and Local Cues. *Current Directions in Psychological Science*, **1**(2), 48–52.
- Etienne, A., Teroni, V., Hurni, C., & Portenier, V. (1990). The Effect of A Single Light Cue on Homing Behaviour of the Golden Hamster. *Animal Behavior*, **39**, 17–41.
- Etienne, A., Maurer, R., & Seguinot, V. (1996). Path Integration in Mammals and its Interaction with Visual Landmarks. *The Journal of Experimental Biology*, **199**(1), 201–209.
- Everett, H. (1995). *Sensors for Mobile Robots: Theory and Application*. Wellesley, MA: A. K. Peters Ltd.
- Feigenbaum, J., & Rolls, E. (1991). Allocentric and Egocentric Spatial Information Processing in the Hippocampal Formation of the Behaving Primate. *Psychobiology*, **19**, 21–40.
- Fenton, A., & Bures, J. (1994). Interhippocampal Transfer of Place Navigation Monocularly Acquired by Rats During Unilateral Functional Ablation of the Dorsal Hippocampus and Visual Cortex with Lidocaine. *Neuroscience*, **58**, 481–491.
- Foster, T., Castro, C., & McNaughton, B. (1989). Spatial Selectivity of Rat Hippocampal Neurons: Dependence on Preparedness for movement. *Science*, **244**, 1580–1582.
- Gallistel, C. (1990). *The Organization of Learning*. Cambridge, MA: Bradford-MIT Press.
- Gelb, A. (1974). *Applied Optimal Estimation*. MIT Press.
- Hebb, D. (1949). *The Organization of Behavior*. New York: John Wiley & Sons.
- Hebert, P., Betge-Brezetz, S., & Chatila, R. (1995). Probabilistic Map Learning: Necessity and Difficulties. *Pages 307–321 of: Proceedings of the International Workshop on Reasoning with Uncertainty in Robotics*. Springer-Verlag.
- Hull, C. (1934a). The Concept of Habit-Family Hierarchy and Maze Learning. I. *Psychological Review*, **41**, 33–54.

- Hull, C. (1934b). The Concept of Habit-Family Hierarchy and Maze Learning. II. *Psychological Review*, **41**, 134–152.
- Hull, C. (1943). *Principles of Behavior*. New York, NY: Appleton-Century-Crofts.
- Hummel, R. (1995). Uncertainty Reasoning in Object Recognition by Image Processing. *Pages 131–145 of: Proceedings of the International Workshop on Reasoning with Uncertainty in Robotics*. Springer-Verlag.
- Hwang, Y., & Ahuja, N. (1992). Gross Motion Planning – A Survey. *ACM Computing Surveys*, **24**(3), 219–291.
- Jarrard, L. (1993). On the Role of the Hippocampus in Learning and Memory in the Rat. *Behavioral Neural Biology*, **3**, 279–287.
- Jazwinski, A. (1970). *Stochastic Processes and Filtering Theory*. New York, NY: Academic Press.
- Johnson, R., & Bhattacharyya, G. (1996). *Statistics: Principles and Methods*. John Wiley and Sons.
- Jung, M., & McNaughton, B. (1993). Spatial Selectivity of Unit Activity in the Hippocampal Granular Layer. *Hippocampus*, **3**, 165–182.
- Kalman, R. (1960). A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME*, **60**.
- Kaplan, S. (1973a). Cognitive Maps, Human Needs, and the Designed Environment. *In: Preiser, W. (ed), Environmental Design Research, Vol 1*. Dowden, Hutchinson, and Ross.
- Kaplan, S. (1973b). Cognitive Maps in Perception and Thought. *In: Downs, M., & Stea, D. (eds), Image and Environment*. Aldine Publishing Company.
- Keith, J., & McVety, K. (1988). Latent Place Learning in a Novel Environment and the Influences of Prior Training in Rats. *Psychobiology*, **16**(2), 146–151.
- Knierim, J., Kudrimoti, H., & McNaughton, B. (1995). Hippocampal Place Fields, the Internal Compass, and the Learning of Landmark Stability. *Journal of Neuroscience*, **15**, 1648–1659.
- Koenig, S., Goodwin, R., & Simmons, R. (1995). Robot Navigation with Markov Models: A Framework for Path Planning and Learning with Limited Computational Resources. *Pages 322–337 of: Proceedings of the International Workshop on Reasoning with Uncertainty in Robotics*. Springer-Verlag.
- Kortenkamp, D. (1993). *Cognitive Maps for Mobile Robots: A Representation for Mapping and Navigation*. Ph.D. thesis, University of Michigan, Electrical Engineering and Computer Science Department.
- Kortenkamp, D., & Weymouth, T. (1994). Topological Mapping for Mobile Robots Using a Combination of Sonar and Vision Sensing. *In: AAAI-94: Proceedings of the American Association of Artificial Intelligence*.

- Kubie, J., & Ranck, J. (1983). Sensory-Behavioral Correlates in Individual Hippocampus Neurons in Three Situations: Space and Context. *Pages 433-447 of: Seifert, W. (ed), Neurobiology of the Hippocampus*. London: Academic Press.
- Kuipers, B. (1978). Modeling Spatial Knowledge. *Cognitive Science*, **2**, 129-153.
- Kuipers, B., & Byun, Y-T. (1991). A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations. *Robotics and Autonomous Systems*, **8**.
- Kuipers, B., Froom, R., Lee, W., & Pierce, D. (1993). The Semantic Hierarchy in Robot Learning. *Chap. 6, pages 141-170 of: Connell, J., & Mahadevan, S. (eds), Robot Learning*. MA: Kluwer Academic Publishers.
- Langley, P., & Pfleger, K. (1995). Acquisition of Place Knowledge Through Case-Based Learning. *In: Proceedings of the International Conference on Machine Learning*.
- Latombe, J-C. (1991). *Robot Motion Planning*. Kluwer Academic.
- Leonard, J., & Durrant-Whyte, H. (1992). Dynamic Map Building for an Autonomous Mobile Robot. *International Journal of Robotics Research*, **11**(4).
- Levitt, T., & Lawton, D. (1990). Qualitative Navigation for Mobile Robots. *Artificial Intelligence*, **44**(3), 305-360.
- Levy, W. (1989). A Computational Approach to Hippocampal Function. *The Psychology of Learning and Motivation*, **23**, 243-305.
- Linsker, R. (1990). Perceptual Neural Organization: Some Approaches Based on Network Models and Information Theory. *Annual Review of Neurosciences*, **13**, 257-281.
- Mackintosh, N. (1983). *Conditioning and Associative Learning*. Clarendon.
- Mataric, M. (1992). Integration of Representation into Goal-Driven Behavior-Based Robots. *IEEE Transactions on Robotics and Automation*, **8**(3).
- Maybeck, P. (1990). The Kalman Filter: An Introduction to Concepts. *In: Cox, I.J., & Wilfong, G. T. (eds), Autonomous Robot Vehicle*. Springer Verlag.
- McNaughton, B., & Nadel, L. (1990). Hebb-Marr Networks and the Neurobiological Representation of Action in Space. *Pages 1-63 of: Gluck, M., & Rumelhart, D. (eds), Neuroscience and Connectionist Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- McNaughton, B., Leonard, B., & Chen, L. (1989). 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.
- McNaughton, B., Knierim, J., & Wilson, M. (1995). Vector Encoding and the Vestibular Foundations of Spatial Cognition: A Neurophysiological and Computational Hypothesis. *Chap. 37 of: Gazzaniga, M. (ed), The Cognitive Neurosciences*. Boston:MIT Press.

- McNaughton, B., Barnes, C., Gerrard, J., Gothard, K., Jung, M., Knierim, J., Kudrimoti, H., Qin, Y., Skaggs, W., Suster, M., & Weaver, K. (1996). Deciphering the Hippocampal Polyglot: the Hippocampus as a Path-Integration System. *The Journal of Experimental Biology*, **199**(1), 173–185.
- Mittelstaedt, H., & Mittelstaedt, M. (1982). Homing by Path Integration in a Mammal. *Pages 290–297 of: Papi, F., & Wallraff, H. (eds), Avian Navigation*, vol. 67. Berlin: Springer-Verlag.
- Mittelstaedt, M., & Mittelstaedt, H. (1980). Homing by Path Integration in a Mammal. *Naturwissenschaften*, **67**, 566–567.
- Mizumori, S., & Williams, J. (1993). Directionally Selective Mnemonic Properties of Neurons in the Lateral Dorsal Nucleus of the Thalamus of Rats. *Journal of Neuroscience*, **13**, 4015–4028.
- Moravec, H., & Elfes, A. (1985). High Resolution Maps from Wide Angle Sonar. *Pages 116–121 of: Proceedings of the 1985 IEEE International Conference on Robotics and Automation*.
- Morris, R. (1981). Spatial Localization Does Not Require the Presence of Local Cues. *Learning and Motivation*, **12**, 239–261.
- Morris, R., Garrud, P., Rawlins, J., & O’Keefe, J. (1982). Place Navigation Impaired in Rats with Hippocampal Lesions. *Nature*, **297**, 681–683.
- Moutarlier, P., & Chatila, R. (1989). Stochastic Multisensory Data Fusion for Mobile Robot Location and Environment Modeling. *In: The Fifth International Symposium on Robotics Research*.
- Muller, R., & Kubie, J. (1987). The Effects of Changes in the Environment on the Spatial Firing of Hippocampal Complex-Spike Cells. *Journal of Neuroscience*, **7**, 1951–1968.
- Muller, R., Kubie, J., & Ranck, J. (1987). Spatial Firing Patterns of Hippocampus Complex-Spike Cells in a Fixed Environment. *Journal of Neuroscience*, **7**, 1935–1950.
- Muller, R., Kubie, J., & Saypoff, R. (1991). The Hippocampus as a Cognitive Graph. *Hippocampus*, **1**(3), 243–246.
- Nehmzow, U. (1995). Animal and Robot Navigation. *Robotics and Autonomous Systems*, **15**, 71–81.
- O’Keefe, J. (1976). Place Units in the Hippocampus of the Freely Moving Rat. *Experimental Neurology*, **51**, 78–109.
- O’Keefe, J. (1989). Computations the Hippocampus Might Perform. *Pages 225–284 of: Nadel, L., Cooper, L., Culicover, P., & Harnish, R. (eds), Neural Connections, Mental Computation*. Cambridge, MA: MIT Press/Bradford Book.
- O’Keefe, J., & Dostrovsky, J. (1971). The Hippocampus as a Spatial Map: Preliminary Evidence from Unit Activity in the Freely Moving Rat. *Brain Research*, **34**, 171–175.
- O’Keefe, J., & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford:Clarendon Press.

- O'Keefe, J., & Speakman, A. (1987). Single Unit Activity in the Rat Hippocampus During a Spatial Memory Task. *Experimental Brain Research*, **68**, 1–27.
- Pagac, D., Nebot, E., & Durrant-Whyte, H. (1995). An Evidential Approach to Probabilistic Map-Building. *In: Reasoning with Uncertainty in Robotics, Proceedings of the International Workshop RUR'95, Amsterdam, The Netherlands*. Lecture Notes in AI, v.1093, Springer Verlag.
- Prescott, T. (1995). Spatial Representations for Navigation in Animals. *Adaptive Behavior*, **4**(2), 85–123.
- Prescott, T., & Mayhew, J. (1992). Building Long-Range Cognitive Maps Using Local Landmarks. *Pages 233–242 of: Meyer, J., Roitblat, H., & Wilson, S. (eds), Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press, Bradford Book.
- Quirk, G., Muller, R., & Kubie, J. (1990). The Firing of Hippocampal Place Cells in the Dark Depends on the Rat's Recent Experience. *Journal of Neuroscience*, **10**, 2008–2017.
- Quirk, G., Muller, R., Kubie, J., & Ranck, J. (1992). The Positional Firing Properties of Medial Entorhinal Neurons: Description and Comparison with Hippocampal Place Cells. *Journal of Neuroscience*, **12**(5), 1945–1963.
- Ranck, J. (1984). Head Direction Cells in the Deep Cell Layer of Dorsal Presubiculum in Freely Moving Rats. *Society for Neuroscience Abstracts*, **10**, 599.
- Recce, M., & Harris, K. (1996). Memory for Places: A Navigational Model in Support of Marr's Theory of Hippocampal Function. *Hippocampus*, **6**, 735–748.
- Redish, A., & Touretzky, D. (1996). Navigating with Landmarks: Computing Goal Locations from Place Codes. *In: Ikeuchi, K., & Veloso, M. (eds), Symbolic Visual Learning*. Oxford University Press.
- Redish, A., & Touretzky, D. (1998). The Role of the Hippocampus in Solving the Morris Water Maze. *Neural Computation*, **10**, 73–111.
- Redish, A., Elga, A., & Touretzky, D. (1996). A Coupled Attractor Model of the Rodent Head Direction System. *Network*, **7**(4), 671–685.
- Rolls, E. (1996). The Representation of Space in the Primate Hippocampus and Episodic Memory. *Pages 375–400 of: Ono, T., McNaughton, B., Molotchnikoff, S., Rolls, E., & Nishijo, H. (eds), Perception, Memory, and Emotion: Frontiers in Neuroscience*. Pergamon Press.
- Schmajuk, N. (1997). *Animal Learning and Cognition: A Neural Network Approach*. Cambridge, UK: Cambridge University Press.
- Schmajuk, N., & Blair, H. (1993). Place Learning and the Dynamics of Spatial Navigation: A Neural Network Approach. *Adaptive Behavior*, **1**, 355–387.
- Schmajuk, N., & Thieme, A. (1992). Purposive Behavior and Cognitive Mapping: A Neural Network Model. *Biological Cybernetics*, **67**, 165–174.

- Scholkopf, B., & Mallot, H. (1995). View-Based Cognitive Mapping and Path Planning. *Adaptive Behavior*, **3**(3), 311–348.
- Schone, H. (1984). *Spatial Orientation: The Spatial Control of Behavior in Animals and Man*. Princeton, NJ: Princeton University Press.
- Schuijver, S. (1997). Efficient Robot Self-Localization in Simple Polygons. *Pages 20–22 of: Proceedings of the 13th European Workshop on Computational Geometry*.
- Seibert, M., & Waxman, A. (1992). Adaptive 3D Object Recognition from Multiple Views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **14**, 107–124.
- Sharp, P. (1996). Multiple Spatial/Behavioral Correlates for Cells in the Rat Postsubiculum: Multiple Regression Analysis and Comparison to Other Hippocampal Areas. *Cerebral Cortex*, **6**, 238–259.
- Sharp, P., & Green, C. (1994). Spatial Correlates of Firing Patterns of Single Cells in the Subiculum of the Freely Moving Rat. *Journal of Neuroscience*, **14**, 2339–2356.
- Sharp, P., Kubie, J., & Muller, R. (1990). Firing Properties of Hippocampal Neurons in a Visually Symmetrical Environment: Contributions of Multiple Sensory Cues and Mnemonic Properties. *Journal of Neuroscience*, **10**, 2339–2356.
- Sharp, P., Blair, H., & Brown, M. (1996). Neural Network Modeling of the Hippocampal Formation Spatial Signals and Their Possible Role in Navigation: A Modular Approach. *Hippocampus*, **6**, 720–734.
- Smith, R., Self, M., & Cheeseman, P. (1990). Estimating Uncertain Spatial Relationships in Robotics. In: Cox, I., & Wilfong, G. (eds), *Autonomous Robot Vehicles*. Springer-Verlag.
- Squire, L., Shimamura, A., & Amaral, D. (1989). Memory and the Hippocampus. *Pages 208–239 of: Byrne, J., & Berry, W. (eds), Neural Models of Plasticity: Theoretical and Empirical Approaches*. New York, NY: Academic Press.
- Sutherland, R., Wishaw, I., & Kolb, B. (1982). A Behavioural Analysis of Spatial Localization Following Electrolytic, Kainate- or Colchicine-Induced Damage to the Hippocampal Formation in the Rat. *Behavioural Brain Research*, **7**, 133–153.
- Taube, J. (1995). Head-Direction Cells Recorded in the Anterior Thalamic Nuclei of Freely Moving Rats. *Journal of Neuroscience*, **15**, 70–86.
- Taube, J., Muller, R., & Ranck, J. (1990a). Head Direction Cells Recorded from the Postsubiculum in Freely Moving Rats: I. Description and Quantitative Analysis. *Journal of Neuroscience*, **10**, 420–435.
- Taube, J., Muller, R., & Ranck, J. (1990b). Head Direction Cells Recorded from the Postsubiculum in Freely Moving Rats: II. Effects of Environmental Manipulations. *Journal of Neuroscience*, **10**, 436–447.
- Thrun, S. (1996). *A Bayesian Approach to Landmark Discovery and Active Perception in Mobile Robot Navigation*. Tech. rept. CMU-CS-96-122. Carnegie Mellon University, Pittsburgh, USA.

- Tolman, E. (1932). *Purposive Behavior in Animals and Men*. New York, NY: Irvington Publishers.
- Tolman, E. (1948). Cognitive Maps in Rats and Men. *Psychological Review*, **55**, 189–208.
- Tolman, E., & Honzik, C. (1930). "Insight" in Rats. *University of California Publications in Psychology*, **4**(14), 215–232.
- Tolman, E., Ritchie, B., & Kalish, D. (1946). Studies in Spatial Learning. I. Orientation and the Short-Cut. *Journal of Experimental Psychology*, **36**, 13–24.
- Traub, R., & Miles, R. (1991). *Neuronal Networks of the Hippocampus*. Cambridge University Press.
- Treves, A., & Rolls, E. (1991). What Determines the Capacity of Autoassociative Memories in the Brain? *Network*, **2**, 371–397.
- Treves, A., & Rolls, E. (1992). Computational Constraints Suggest the Need for Two Distinct Input Systems to the Hippocampal CA3 Network. *Hippocampus*, **2**, 189–199.
- Trullier, O., Wiener, S., Berthoz, A., & Meyer, J-A. (1997). Biologically-based Artificial Navigation Systems: Review and Prospects. *Progress in Neurobiology*, **51**, 483–544.
- Tsuji, S., & Li, S. (1993). Memorizing and Representing Route Scenes. *In: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. MIT Press.
- Wan, H., Touretzky, D., & Redish, A. (1994). Towards a Computational Theory of Rat Navigation. *Pages 11–19 of: Mozer, M., Smolensky, P., Touretzky, D., Elman, J., & Weigend, A. (eds), Proceedings of the 1993 Connectionist Models Summer School*. Lawrence Erlbaum Associates.
- Wilson, M., & McNaughton, B. (1993). Dynamics of the Hippocampal Ensemble Code for Space. *Science*, **261**, 1055–1058.
- Worden, R. (1992). Navigation by Fragment Fitting: A Theory of Hippocampal Function. *Hippocampus*, **2**(2), 165–188.
- Yamauchi, B., & Beer, R. (1996). Spatial Learning for Navigation in Dynamic Environments. *IEEE Transactions on Systems, Man, and Cybernetics - Part B, Special Issue on Robot Learning*, **26**.
- Yamauchi, B., & Langley, P. (1997). Place recognition in Dynamic Environments. *Journal of Robotic Systems, Special Issue on Mobile Robots*, **14**(2), 107–120.
- Zipser, D. (1986). Biologically Plausible Models of Place Recognition and Place Location. *Pages 432–470 of: Rumelhart, D., McClelland, J., & the PDP Research Group (eds), Parallel Distributed Processing: Explorations in the Micro-Structure of Cognition, Volume 2*. Cambridge, MA: MIT Press.

Appendix A

Some Assumptions in Our Model

Some key assumptions have to be made for our computational model to be realizable. These assumptions are clarified in this section.

1. **Recognized landmark information**

Our model assumes that the EC inputs consist of recognized landmarks along with information about their allocentric position relative to the robot. Our model required this because of its analogy to the hippocampal formation, which is known to derive highly processed inputs from other associational areas of the brain. Since object and scene recognition are still big problems in contemporary robotics, is this approach only a simulation toy that is unimplementable on real robots? We suggest not. It might be observed that the functioning of our localization system depends on the notion of a place but does not depend on the exact mechanisms for defining a place representation. Thus, any prudent choice of a place representation would work. For example, a *local occupancy-grid* representation of a place [Langley & Pfleger, 1995; Recce & Harris, 1996; Yamauchi & Langley, 1997] can be used in our model, making it implementable on a mobile robot with sonar sensors.

2. **Reliance on dead-reckoning**

The model discussed in this paper assumes that the robot has a reliable dead-reckoning system with known error models. Dead-reckoning in contemporary robots is usually performed through odometric techniques that use optical and magnetic wheel encoders. Since the odometric approaches derive their navigational parameters from wheel rotation, they are subject to problems arising from slippage, tread wear, and/or improper tire inflation [Balakrishna & Ghosal, 1995; Everett, 1995]. However, Doppler and inertial navigation systems can often be used to considerably reduce these sources of errors [Everett, 1995]. *Doppler navigation* systems operate on the principle of the Doppler shift in frequency, observed when radiated energy reflects off a surface that is moving with respect to the emitter. *Inertial navigation* systems rely on a set of mechanisms that continuously sense minute accelerations in each of the three directional axes. These are then integrated over time to yield velocity and position. Such systems, although capable of producing reasonably precise dead-reckoning, currently have limited applications mainly due to their prohibitive cost [Everett, 1995]. However, with advances in technology, it is likely that accurate and affordable dead-reckoning devices will soon become available, making our approach quite practical.

3. **Maintaining correlations**

With probabilistic localization approaches like the Kalman filter, it is imperative that the correlations between the state variables be maintained correctly. Estimating and maintaining these correlations is often difficult owing to the approximation errors stemming from linearizing the (usually) non-linear system and measurement functions, biases on the robot position, and the computational complexity associated with updating state vectors containing a large number of elements [Hebert *et al.*, 1995]. However, these correlations and variances cannot be neglected since doing so will lead to inconsistencies in the uncertainties [Hebert *et al.*, 1995]. In our model, we ignore the autocorrelation of the measurement noise. However, as we

argued before, if the animal navigates randomly or purposefully between place field centers, this correlation term will be negligible.

Appendix B

Some Useful Results for Hippocampal Kalman Filtering

Proposition 1 (Size of the EC fields) *In two dimensional space the EC field, or the region over which it responds, is roughly circular with a radius of $r \leq \sqrt{-2\sigma_L^2 \ln(\Gamma_{EC})}$.*

If the variance of the EC cell distance matching Gaussians is given by σ_L^2 and the EC cell is said to recognize a landmark if the cell activation is greater than the threshold of Γ_{EC} , then the range over which the EC cell with center (c_x, c_y) responds is given by:

$$\exp^{-\frac{(x-c_x)^2+(y-c_y)^2}{2\sigma_L^2}} \geq \Gamma_{EC} \iff (x-c_x)^2 + (y-c_y)^2 \leq -2\sigma_L^2 \ln(\Gamma_{EC}) \iff r^2 \leq -2\sigma_L^2 \ln(\Gamma_{EC})$$

Thus, with $\sigma_L^2 = 1$ and $\Gamma_{EC} = 0.6$, we obtain $r \leq 1.01$, i.e., the EC unit responds to a landmark within a circular region of radius 1.01 from its center.

Proposition 2 (Variance of an EC cell) *The variance associated with an EC unit is given by $\sigma^2 = R^2\sigma_S^2 - 2\sigma_L^2 \ln(\Gamma_{EC})$*

The variance associated with an EC unit firing has two components. The first component is a measure of the uncertainty associated with the EC unit center, while the second component is the imprecision tolerated by the EC firing field. Suppose the sensory range error is characterized by a zero-mean, white Gaussian with standard deviation σ_S (per unit distance), and the maximum sensory range is R . When a new EC unit is recruited and its center is set up, the imprecision (or uncertainty) associated with the center is characterized by the sensing uncertainty and has a variance of $R^2\sigma_S^2$. From Proposition 1, the position uncertainty within the EC field is given by $r^2 \leq -2\sigma_L^2 \ln(\Gamma_{EC})$. Since these two sources of uncertainty are independent, the net variance associated with an EC cell is given by: $\sigma^2 = R^2\sigma_S^2 - 2\sigma_L^2 \ln(\Gamma_{EC})$. Note that this is the variance along each coordinate axis of the Cartesian system.

Proposition 3 (Form of the measurement error) *The measurement error is given by $w_k = x_{i_k} - x_{0,k}$.*

In our model the measurement function is assumed to represent the distance between the current position of the animal ($x_{0,k}$) and the center of the place field that it is currently in (x_{i_k}), corrupted by some random measurement noise. Thus

$$z_k = x_{0,k} - x_{i_k} + w_k$$

Since the true centers of the place fields are not known and the true position of the animal is also unknown (and is the very subject of this localization process), z_k cannot be measured. In our model we assume that the animal always measures 0, i.e., if the animal is in place i_k the system assumes that it must be at the center of that place. This constrains the form of the random noise to:

$$z_k = x_{0,k} - x_{i_k} + w_k = 0 \implies w_k = x_{i_k} - x_{0,k}$$

Thus, the Kalman filter in our model uses observed measurement of 0 and predicted measurement of $\hat{x}_{0,k} - \hat{x}_{i_k}$ to perform the required updates.

Proposition 4 (Measurement error has zero mean) $E(w_k) = 0$

In order to use the Kalman filter, we have to show that our measurement noise has zero mean. As shown in Proposition 3, the measurement noise is given by $w_k = x_{i_k} - x_{0,k}$. Thus:

$$E(w_k) = \int \sum_{\text{places } p} (x_p - x_{0,k}) \Pr(x_{0,k}/p) \Pr(p/k) dx_{0,k}$$

$$\iff E(w_k) = \sum_{\text{places } p} \underbrace{\int (x_p - x_{0,k}) \Pr(x_{0,k}/p) dx_{0,k}}_0 \Pr(p/k) \iff E(w_k) = 0$$

Notice that the noise is zero-mean only if there is no prior bias on the position of the animal in the place field, i.e., each position in the place field is either equi-probable or the animal navigates from one place field center to another.

Proposition 5 (Measurement error is autocorrelated) *The measurement error w_k is autocorrelated, i.e., $E(w_k w_{k-1}^T) \neq 0$.*

At time step $k - 1$, suppose the robot was at the position $x_{0,k-1}$. Suppose the sensor inputs determined that the robot was in place i_{k-1} with place field center $x_{i_{k-1}}$. Then:

$$w_{k-1} = x_{i_{k-1}} - x_{0,k-1} \quad (7)$$

Since the robot moves by u_{k-1} (and assuming linear robot transformations), its position at step k is given by:

$$x_{0,k} = x_{0,k-1} + u_{k-1} \quad (8)$$

Suppose this position corresponds to place i_k with center x_{i_k} , then:

$$w_k = x_{i_k} - x_{0,k} \quad (9)$$

Using equations 8 and 7 in 9, we get:

$$w_k = x_{i_k} - x_{i_{k-1}} - u_{k-1} + w_{k-1}$$

As can be observed, w_k depends on w_{k-1} , and hence is an autocorrelated sequence.

Proposition 6 (Variance of the measurement error) *The variance of the measurement error is given by*

$$\mathbf{R}_k = \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}$$

where $\sigma^2 = R^2 \sigma_S^2 - 2\sigma_L^2 \ln(\Gamma_{EC})$.

Since the measurement error is given by $w_k = x_{i_k} - x_{0,k}$, i.e., the deviation of the current animal position from the CA3 field center, the variance of this error is the variance associated with CA3 fields. Now, CA3 fields are weighted intersections of EC fields. Thus, in the worst case (i.e., CA3 and EC fields have the same size) the variance of CA3 fields equal the variance of EC fields. Since

the variance of EC cells is given by $\sigma^2 = R^2\sigma_S^2 - 2\sigma_L^2 \ln(\Gamma_{EC})$ along each coordinate axis of the Cartesian system (Proposition 2), and the covariance between the axes are assumed to be zeros, the variance of the measurement error becomes:

$$\mathbf{R}_k = \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}$$

Proposition 7 (Covariance in the Mahalanobis test) *The covariance matrix of the Mahalanobis distance test used in Section 7.5 is given by $(\mathbf{C}_{i_k i_k} + \mathbf{C}_{00} - 2\mathbf{C}_{0i_k} + \mathbf{R}_k)$*

Suppose we are considering CA1 unit i_k with true place field center x_{i_k} and estimated center \hat{x}_{i_k} . Since predictions in our model are given by $\hat{z}_k = \hat{x}_{0,k} - \hat{x}_{i_k}$ while the observations are given by $z_k = x_{0,k} - x_{i_k} + w_k$, we can determine the covariance between the predicted and observed measurements of the Mahalanobis test for CA1 unit i_k .

$$\begin{aligned} & E((\hat{z}_k - z_k)(\hat{z}_k - z_k)^T) \\ \iff & E((\hat{x}_{0,k} - \hat{x}_{i_k} - x_{0,k} + x_{i_k} - w_k)(\hat{x}_{0,k} - \hat{x}_{i_k} - x_{0,k} + x_{i_k} - w_k)^T) \\ \iff & E((\epsilon_{i_k} - \epsilon_{0,k} - w_k)(\epsilon_{i_k} - \epsilon_{0,k} - w_k)^T) \\ \iff & \mathbf{C}_{i_k i_k} + \mathbf{C}_{00} - 2\mathbf{C}_{0i_k} + \mathbf{R}_k \end{aligned}$$

where $\epsilon_{i_k} = x_{i_k} - \hat{x}_{i_k}$ and $\epsilon_{0,k} = x_{0,k} - \hat{x}_{0,k}$. We also assume that the dead-reckoning and measurement noises are independent, i.e. $E(\epsilon_{i_k} w_k^T) = E(\epsilon_{0,k} w_k^T) = E(w_k \epsilon_{i_k}^T) = E(w_k \epsilon_{0,k}^T) = 0$.

Proposition 8 (Chi-square (χ^2) distribution and the Mahalanobis test) *The distance threshold for the Mahalanobis test is 4.61*

The probability density function of a χ^2 distribution is an asymmetric curve with a long right-hand tail. Given the number of degrees of freedom, we can determine the value of the distribution (χ_α^2) such that the area under the curve to the right of it is α [Johnson & Bhattacharyya, 1996]. Since the Mahalanobis distance has a chi-square distribution, we can choose a value of the distance such that the area to the right of it is, say, 10%. Since the covariance matrix of the Mahalanobis test shown in proposition 7 has a rank of 2, we determine the value of the chi-square distribution with 2 degrees of freedom such that the area to the right of it is 10%. This value is 4.61 [Johnson & Bhattacharyya, 1996].

Thus, in our model, if we compute the Mahalanobis distance and declare that the match is correct if this distance happens to be less than 4.61, then we can be 90% confident of *not* rejecting actually correct matches.

Proposition 9 (The 2.5 σ boundary) *The 2.5 σ boundary around an estimated parameter is expected to include the true value with probability greater than 84%*

According to Chebyshev's inequality, if μ and σ are the mean and standard deviation of a random variable X , then for any positive constant k :

$$\Pr(|X - \mu| < k\sigma) \geq 1 - \frac{1}{k^2}$$

Thus, if x_i is the estimated parameter with current estimate \hat{x}_i and variance σ , using $k = 2.5$ in Chebyshev's theorem yields:

$$\Pr(|x_i - \hat{x}_i| < 2.5\sigma) \geq 0.84$$

Thus, the 2.5σ boundary is expected to include the true value with probability greater than 84%.

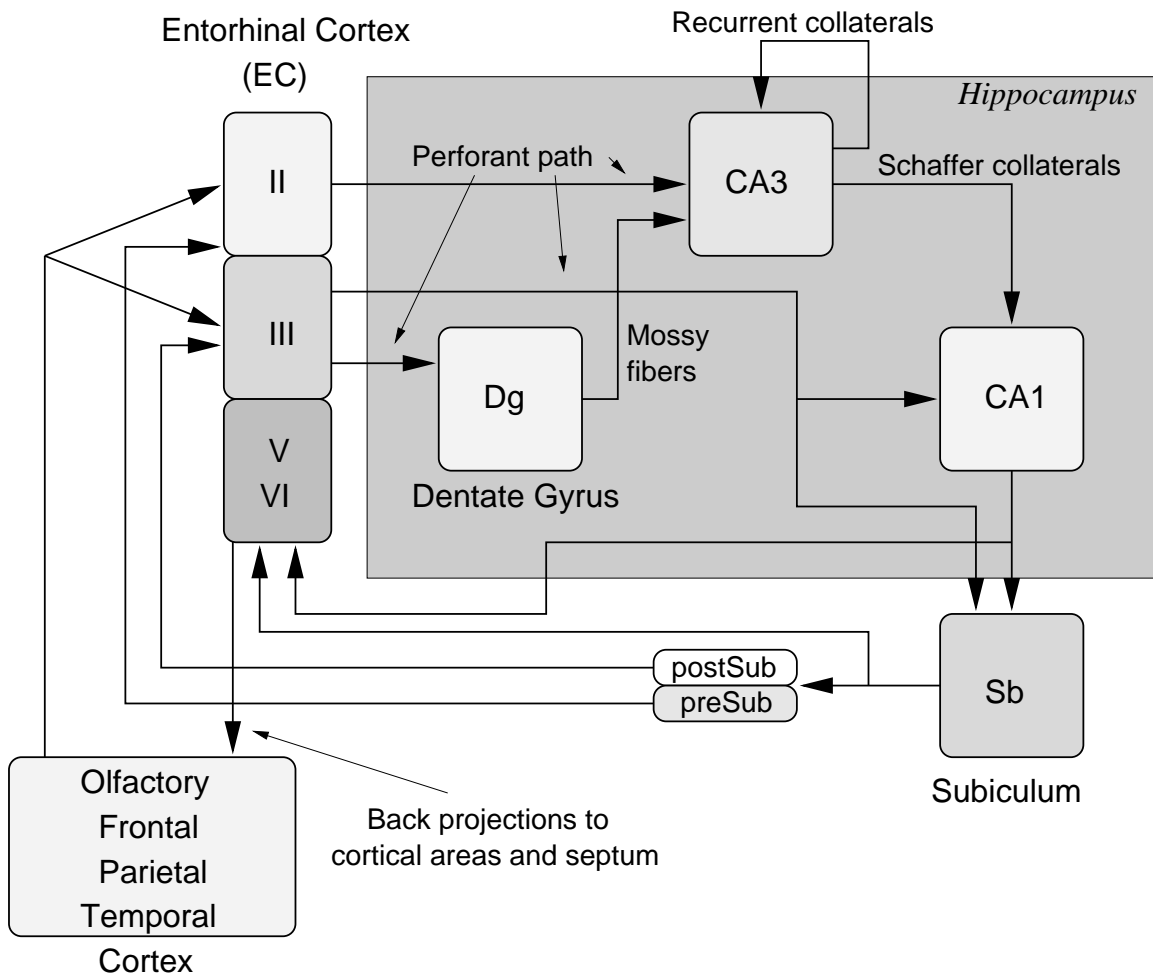


Figure 1: Anatomy of the hippocampal formation.

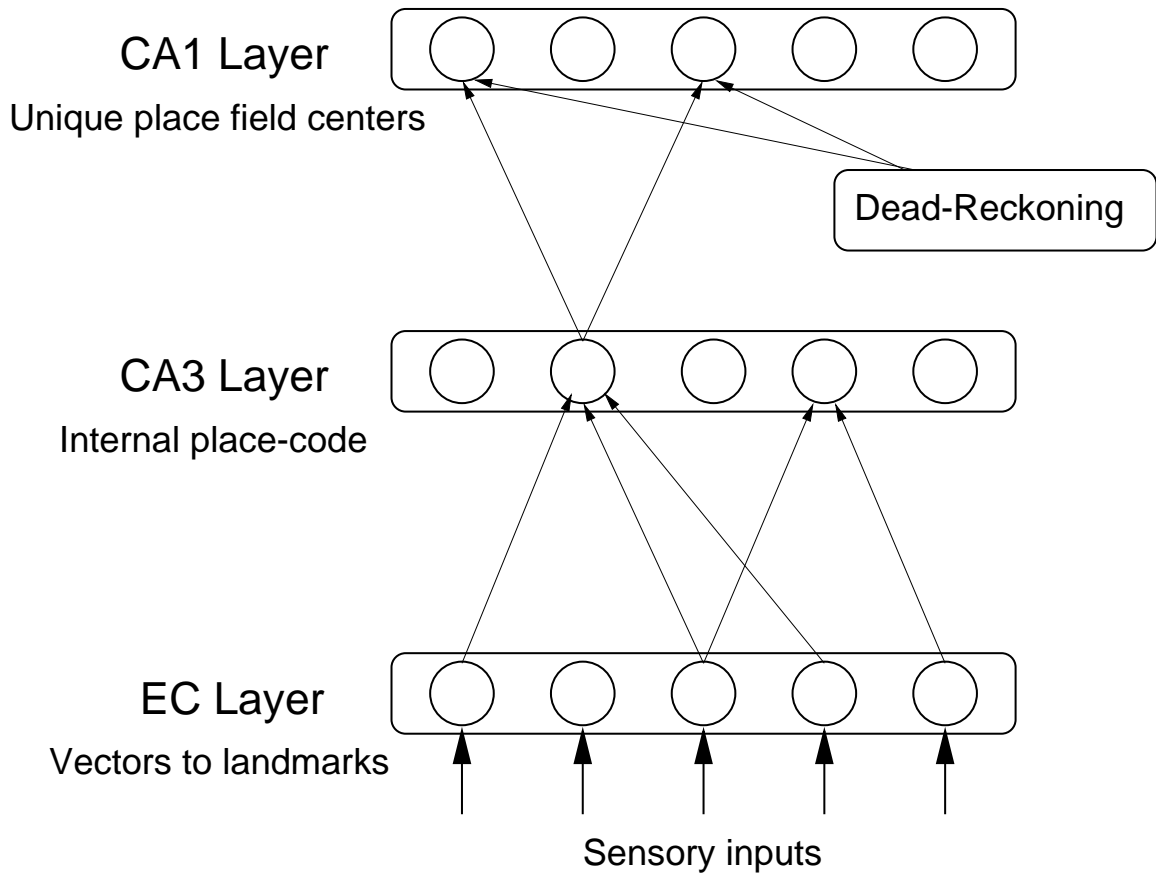


Figure 2: Computational model of the hippocampus.

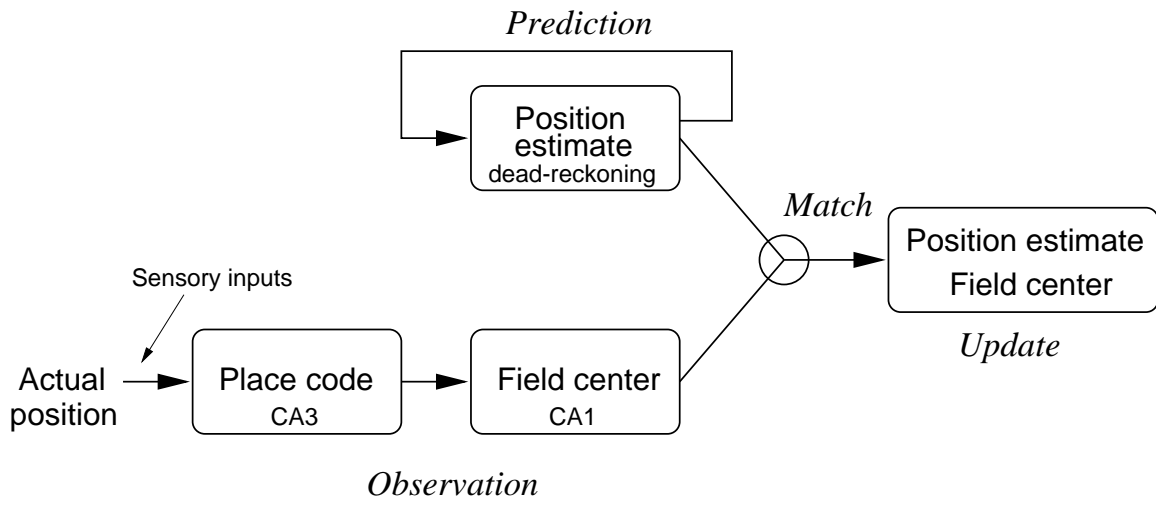


Figure 3: A schematic of hippocampal localization.

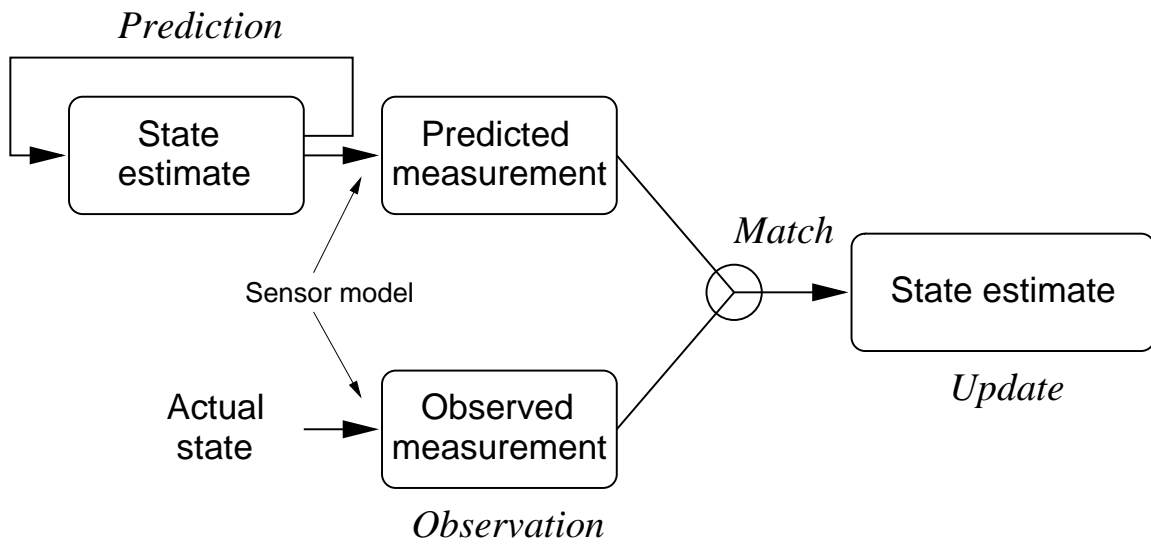


Figure 4: A schematic of Kalman filter

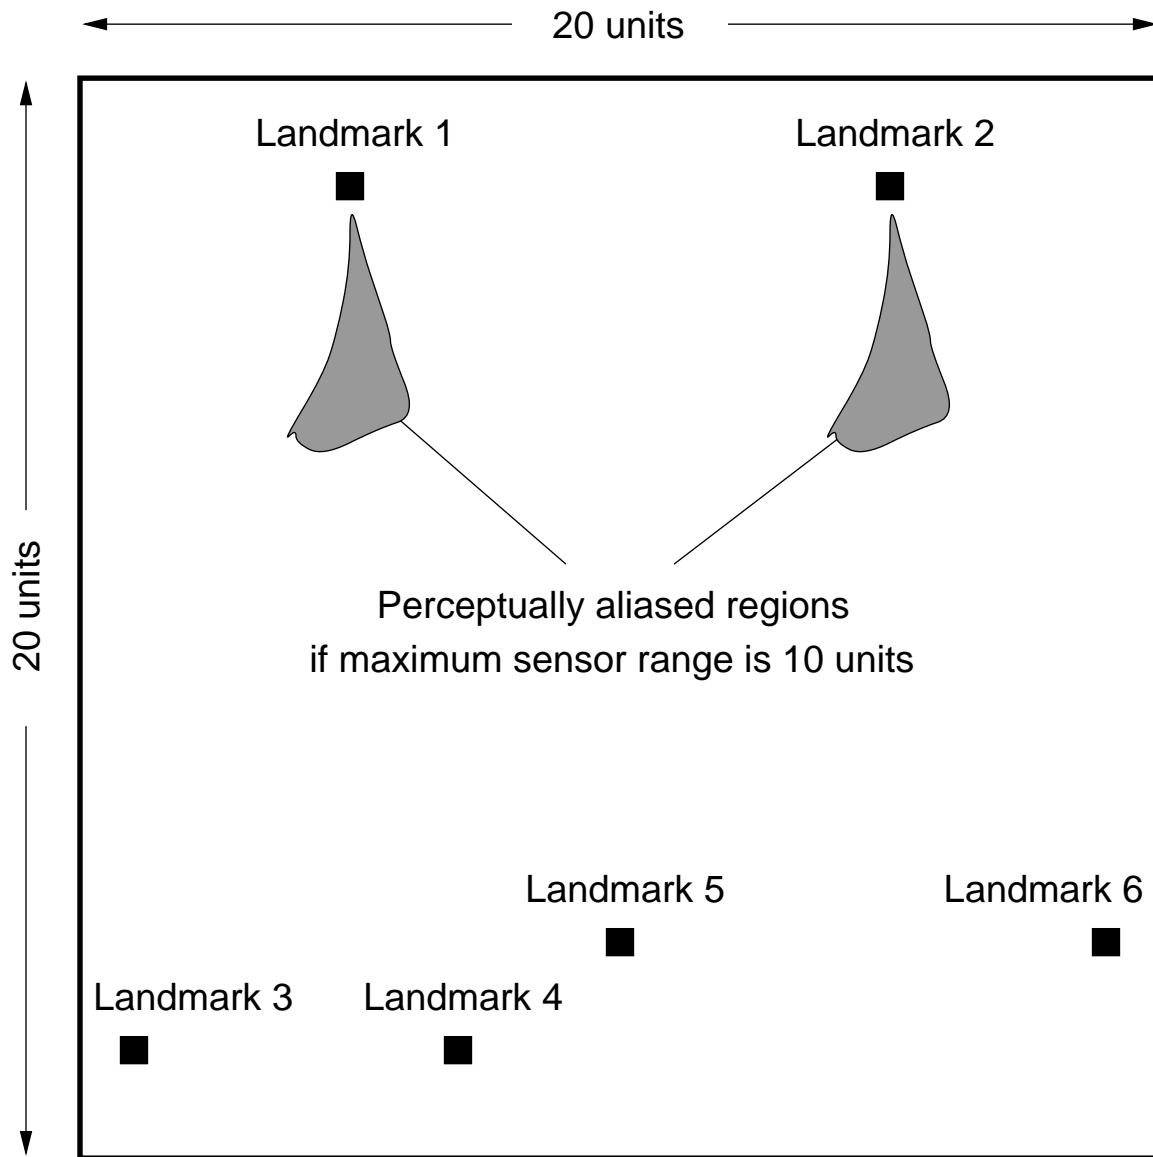
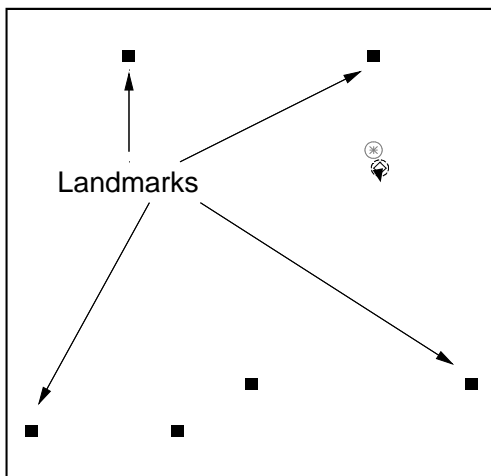


Figure 5: The simulation environment with six identical landmarks. The walls are impenetrable. If the maximum sensor range is limited to 10 units, the shaded regions are perceptually aliased, i.e., the robot receives identical sensory inputs in the two regions.

Place field estimates after one motion step



Place map after one motion step


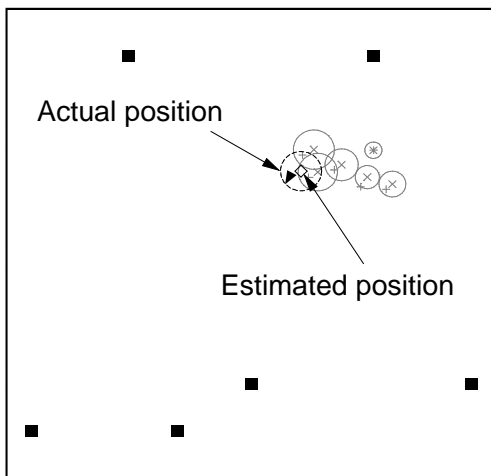


Figure 6: State of the system after one step. Place field centers along with their variances (displayed as 2.5σ circles) are shown on the left, while the place map (or CA1 firing field) is shown on the right.

Place field estimates after
10 motion steps



Place map after
10 motion steps

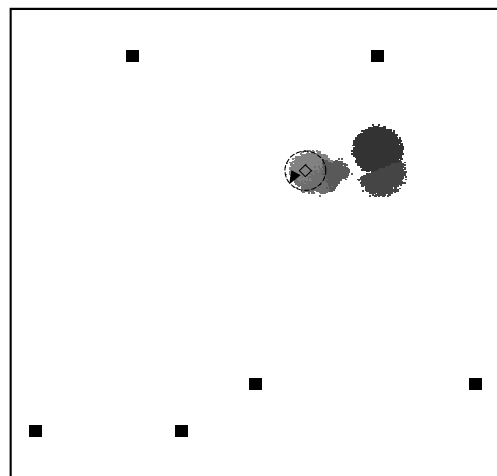



Figure 7: State of the system after 10 steps of exploration. The variance of the place field estimates steadily increase without revisits to places visited earlier. Place map shows firing fields of different shapes and sizes.

Place field estimates after
100 motion steps



Place map after
100 motion steps




Figure 8: State of the system after 100 exploration steps. The robot is now at a place that is perceptually similar to its starting place. However, the Mahalanobis test correctly classifies this as a new place. Some of the CA1 cells have extremely small firing fields.





Figure 9: Firing fields of CA3 and CA1 units. CA3 unit 0 responds to perceptually aliased places. This ambiguity is resolved by the Mahalanobis distance in the CA1 layer with units 0 and 61 responding to the two aliased places.

Place field estimates after
200 motion steps



Place map after
200 motion steps





Figure 10: After 200 exploration steps the error between the actual and estimated positions of the robot is quite alarming. However, the robot cannot correct this error without revisiting places for which it has accurate estimates.

Place field estimates after
250 motion steps



Place map after
250 motion steps




Figure 11: System state at the end of exploration. The robot finally revisits familiar places and significantly corrects its position estimate error.




Figure 12: Trajectory followed by the robot during the exploration phase.

Kidnapped robot relocalization

Upon reintroduction

5 motion steps


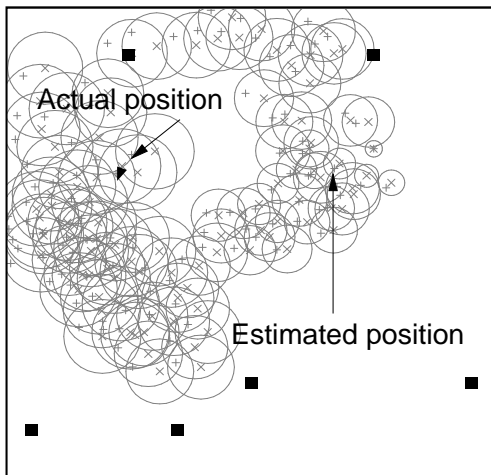
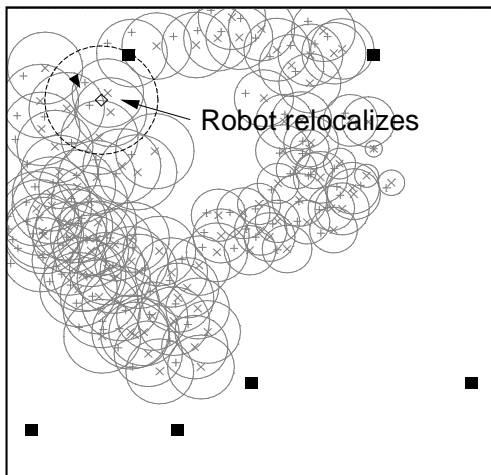


Figure 13: The kidnapped robot does not relocalize because it detects perceptual aliasing. It simply moves randomly looking for a sensorily unique place at which to localize.

Kidnapped robot relocalization

10 motion steps



50 motion steps




Figure 14: The Kalman filter relocalizes the robot at a perceptually unique place. The localized robot can further improve its position estimate, as shown in the picture on the right.

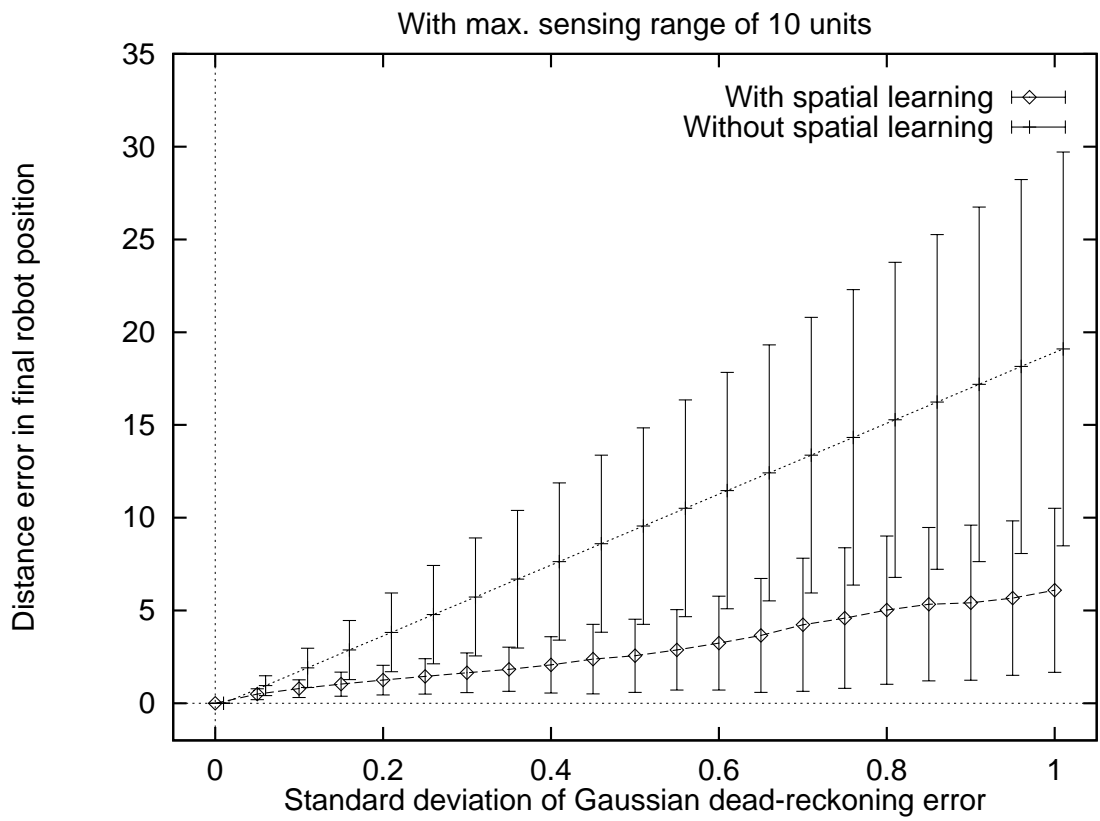


Figure 15: Comparison of final robot position error with and without spatial learning ability. The plot shows the mean and standard deviation of the final position error over 100 trajectories.

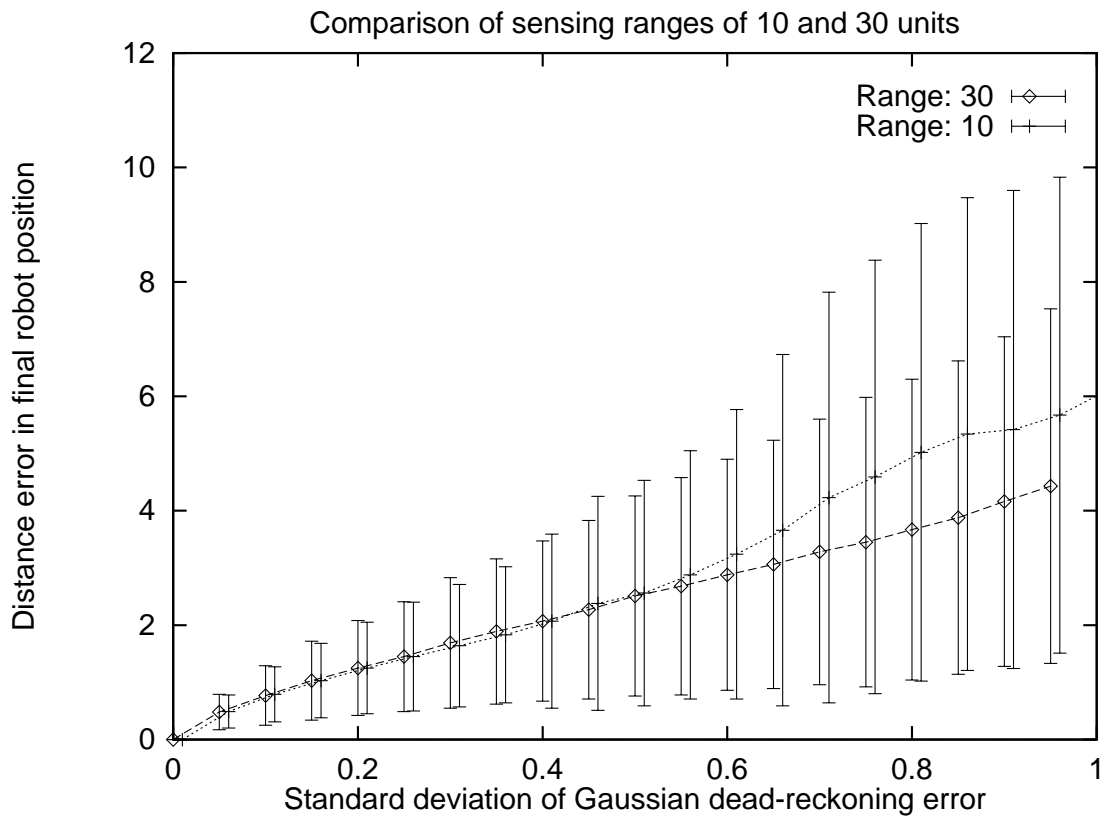


Figure 16: Error in the final robot position with sensor ranges restricted to 10 and 30 units. With a range of 30 units, there is no perceptual aliasing in the environment, which leads to better localization. With a range of 10 units perceptual aliasing interferes with faithful map learning, thereby conceding larger position errors.

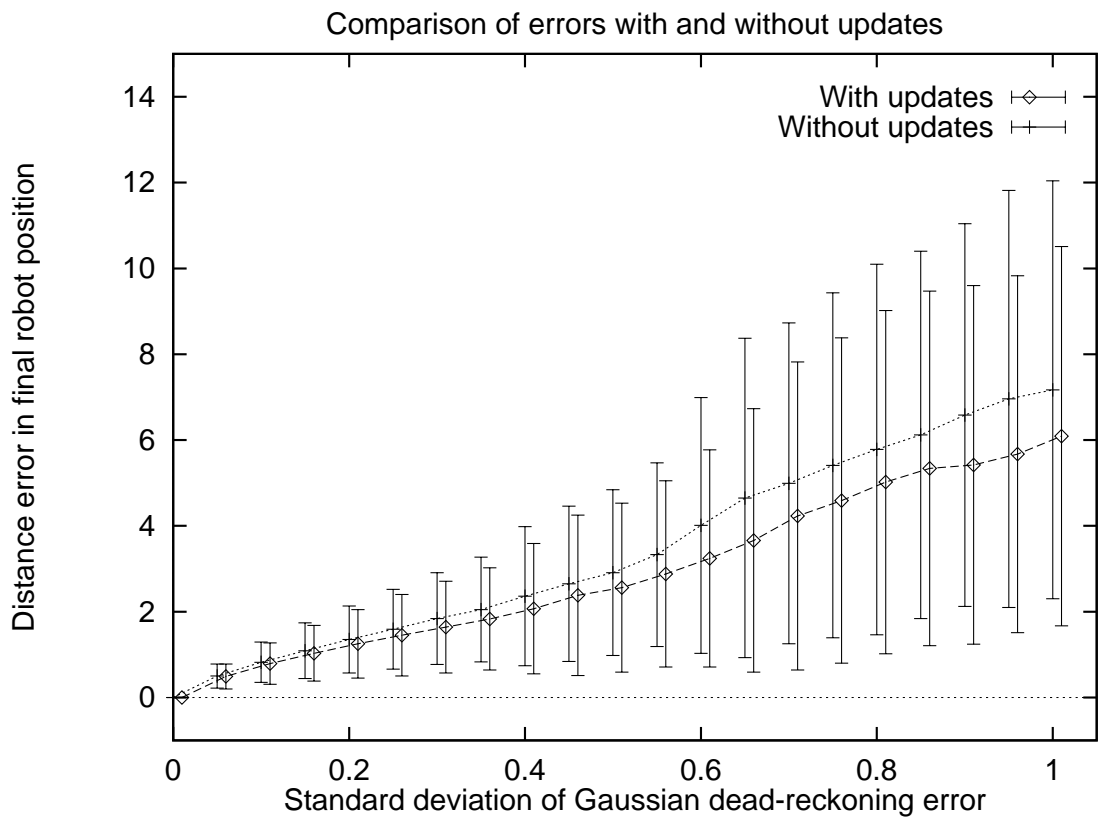


Figure 17: Hippocampal Kalman filtering updates place field centers in addition to the robot's position estimate, leading to lower localization error.




Figure 18: Mean numbers of EC, CA3, and CA1 units created by the spatial learning system. The numbers of CA3 and CA1 units scale up well with increase in range error.

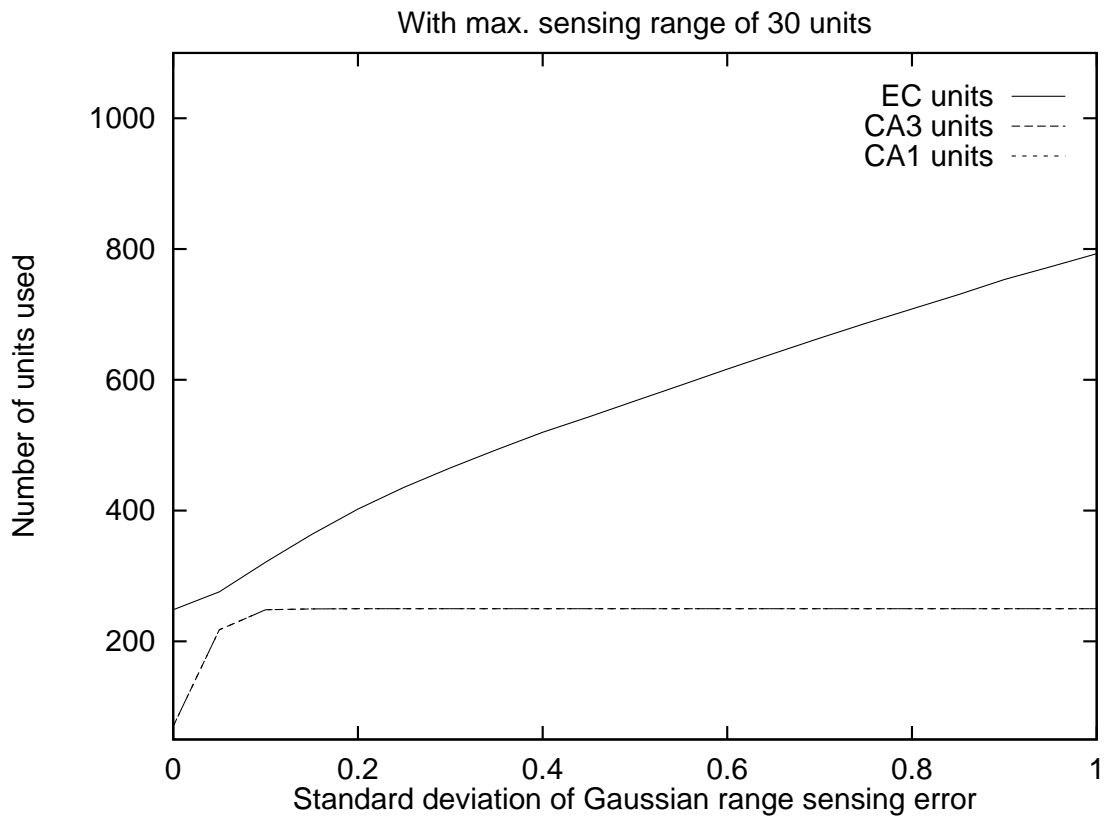


Figure 19: Increasing range error affects the numbers of EC units required but not the number of places represented (indicated by the number of CA3 and CA1 units). Since there is no perceptual aliasing with sensor range of 30 units, the number of CA3 units equal the number of CA1 units.




Figure 20: Multi-granular representation of space. Frequently visited regions of the environment are represented by more CA1 units with smaller firing fields, while cursorily explored regions are represented by fewer CA1 units with larger firing fields.