ARTICLE Communicated by Michael Hasselmo

Latent Attractors: A Model for Context-Dependent Place Representations in the Hippocampus

Simona Doboli
Ali A. Minai
Complex Adaptive Systems Laboratory, ECECS Department, University of Cincinnati, Cincinnati, OH 45221-0030, U.S.A.

Phillip J. Best
Laboratory of Cognitive Psychobiology, Department of Psychology, Miami University, Oxford, OH 45056, U.S.A.

Cells throughout the rodent hippocampal system show place-specific patterns of firing called *place fields*, creating a coarse-coded representation of location. The dependencies of this place code—or cognitive map—on sensory cues have been investigated extensively, and several computational models have been developed to explain them. However, place representations also exhibit strong dependence on spatial and behavioral context, and identical sensory environments can produce very different place codes in different situations. Several recent studies have proposed models for the computational basis of this phenomenon, but it is still not completely understood. In this article, we present a very simple connectionist model for producing context-dependent place representations in the hippocampus. We propose that context dependence arises in the dentate gyrus-hilus (DGH) system, which functions as a dynamic selector, disposing a small group of granule and pyramidal cells to fire in response to afferent stimulus while depressing the rest. It is hypothesized that the DGH system dynamics has "latent attractors," which are unmasked by the afferent input and channel system activity into subpopulations of cells in the DG, CA3, and other hippocampal regions as observed experimentally. The proposed model shows that a minimally structured hippocampus-like system can robustly produce context-dependent place codes with realistic attributes.

1 Introduction

The hippocampus has been a region of great interest for neuroscientists and psychobiologists because of its apparently central role in memory and cognition. In primates (including humans), the hippocampus is crucial to the formation and consolidation of episodic memories (Scoville & Milner, 1957; Squire, 1992; Squire & Zola-Morgan, 1988), and this function proba-

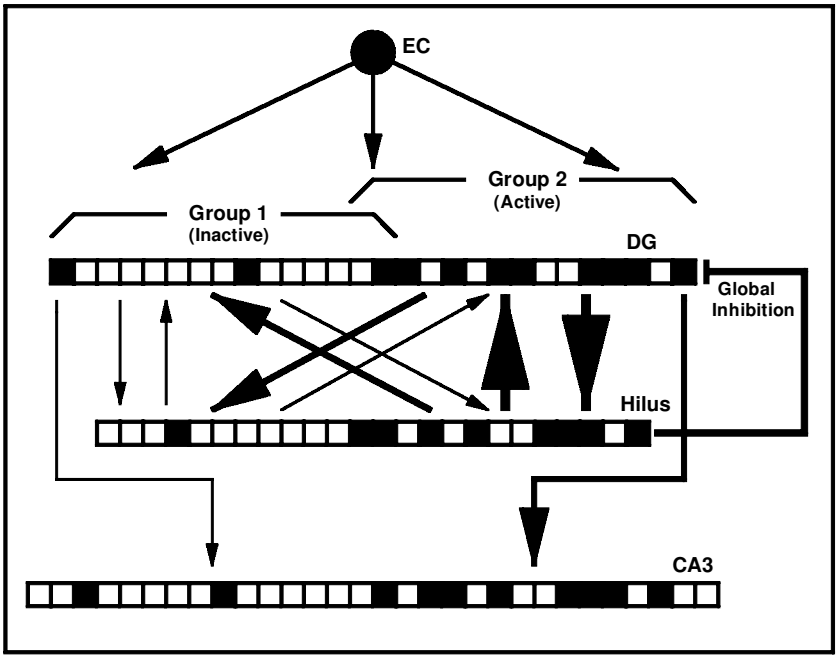

Figure 1: Schematic of the network model. The DG and H layers are configured into randomly chosen cell groups with possible overlap. The EC projects broadly to both DG and CA3, and the DG projects to CA3. There are no explicit groups in CA3, but DG-to-CA3 projections are very sparse. The figure shows only two DG/H groups: Group 1 is nominally inactive and Group 2 active. Filled squares symbolize cell activity, and arrows indicate excitatory signals between layers, with line width representing signal strength. Note that cells in each group are placed contiguously only for clarity: The model does not assume topographic organization of cell groups.

bly extends to other mammals as well (Eichenbaum, Kuperstein, Fagan, & Nagode, 1987). In rodents, the hippocampus has been implicated specifically in spatial cognition (Morris, Garrud, Rawlins, & O'Keefe, 1982) and is hypothesized to form a spatial cognitive map (O'Keefe & Nadel, 1978). Primary cells in the dentate gyrus (DG), CA3, and CA1 regions of the hippocampus (see Figure 1) have activity tuned to an animal's spatial location (O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1978; Muller, Kubie, & Ranck, 1987; Jung & McNaughton, 1993). These cells are called *place cells*, and the spatial regions where they are active are termed *place fields*. Together, the activity pattern of primary cells in CA3 or CA1 forms a coarse-coded, distributed representation of the animal's position. We call this a *place representation* or *place code*.

The sensory, behavioral, and motivational correlates of place fields throughout the hippocampus have been elucidated experimentally (Miller & Best, 1980; Hill & Best, 1981; Muller & Kubie, 1987; O'Keefe & Speakman, 1987; Thompson & Best, 1989, 1990; Sharp, Kubie, & Muller, 1990; Quirk, Muller, & Kubie, 1990; Quirk, Muller, Kubie, & Ranck, 1992; Bostock, Muller, & Kubie, 1991; Wilson & McNaughton, 1993; Markus et al., 1995; Gothard, Skaggs, Moore, & McNaughton, 1996; Gothard, Skaggs, & McNaughton, 1996; O'Keefe & Burgess, 1996; Barnes et al., 1997; Rotenberg and Muller, 1997). An important aspect of place representations that is not yet completely understood is their dependence on spatial or behavioral context. Several studies have shown that CA3/CA1 place representations, while strongly controlled by sensory cues (O'Keefe & Nadel, 1978; Hill & Best, 1981; Muller & Kubie, 1987), are not causally dependent on these cues and are affected significantly by an animal's perception of its situation (Quirk et al., 1990; Bostock et al., 1991; Markus et al., 1995; Markus, Qin, McNaughton, & Barnes, 1994; Gothard, Skaggs, Moore, & McNaughton, 1996; Gothard, Skaggs, & McNaughton, 1996). This phenomenon has been described as the dependence of place codes on frames of reference (Gothard, Skaggs, Moore, & McNaughton, 1996; Touretzky & Redish, 1996; Samsonovich & McNaughton, 1997). Our notion of context dependence refers in particular to the observation that identical sensory environments can produce different place representations without affecting the intraenvironment stability or spatial coherence of these representations (Quirk et al., 1990), suggesting the existence of a multistable mechanism in the hippocampus (Barnes et al., 1997). In this article, we address the issue of context-dependent place representations in the hippocampal system in an abstract connectionist framework.

2 Theoretical Development and Justification

2.1 Motivating Experimental Results. This research is motivated by several observations about hippocampal place fields:

1. All regions of the hippocampus show place fields (O'Keefe & Nadel, 1978; Muller et al., 1987; Jung & McNaughton, 1993). Entorhinal cortex (EC) neurons also have broadly tuned and noisy place fields (Barnes et al., 1990; Quirk, Muller, Kubie, & Ranck, 1992).
2. Only a subset (20–30%) of cells in regions CA1 and CA3 show place activity in any specific environment, and this active set appears to be randomly selected for each situation (Thompson & Best, 1989, 1990; Kubie & Muller, 1991; Muller, Kubie, & Saypoff, 1991; Wilson & McNaughton, 1993). Silent cells have alsio been reported in the DG (Jung & McNaughton, 1993), although to our knowledge, there is no systematic study of this effect. EC neurons, in contrast, appear to fire in almost all environments, with their place specificity more closely tied

to sensory stimulus than is the case for CA3/CA1 neurons (Quirk et al., 1992).

3. Within an environment, place fields are controlled primarily by distal visual cues (O'Keefe & Nadel, 1978; Miller & Best, 1980; Hill & Best, 1981; Muller & Kubie, 1987; Sharp et al., 1990; Bostock et al., 1991). However, removal of these cues in the animal's presence does not disrupt the fields (O'Keefe & Speakman, 1987; Quirk et al., 1990).

4. The place representation is established very early in an episode. Even in novel environments, animals have full-fledged place fields soon after entry (Hill, 1978) although they might take some time to consolidate (Wilson & McNaughton, 1993).

5. Different place representations are active in different environments (Thompson & Best, 1989; Kubie & Muller, 1991) and behavioral situations (Markus et al., 1994, 1995). However, the choice of representation seems to depend on the animal's perception of its environment rather than purely on sensory input. Identical environments—and even the same environment—can support distinct place representations in different contexts (Quirk et al., 1990).

6. Once a place representation has been activated in the hippocampus, it is very stable against large and cognitively dissonant changes in visual cues (Quirk et al., 1990). When it does change, the switch is usually to a significantly different global representation with a different set of active cells rather than a gradual shift (Bostock et al., 1991; Gothard, Skaggs, Moore, & McNaughton, 1996; Rotenberg & Muller, 1997). However, there are also data suggesting that the remapping leaves some place fields unchanged (Wilson & McNaughton, 1993; Shapiro, Tanila, & Eichenbaum, 1997; Tanila, Shapiro, & Eichenbaum, 1997; Knierim, Kudrimoti, & McNaughton, 1998; Skaggs & McNaughton, 1998), while others may switch their dependencies from one set of cues to another (Shapiro et al., 1997; Tanila et al., 1997; Skaggs & McNaughton, 1998). Also, local alterations to place fields can occur if the topology of the environment is altered locally (e.g., by the insertion of barriers) (Muller & Kubie, 1987).

7. Old rats show a propensity to activate an incorrect place representation upon return to a familiar environment after a hiatus, but when the correct representation is activated, it appears to show no degradation (Barnes et al., 1997). This suggests a multistable system of internally stable cognitive maps with a recognition-based selection of the correct one. Barnes et al. (1997) hypothesize that activation of incorrect maps in old rats may reflect a failure of recognition due to degradation in associative learning.

2.2 Primary Hypothesis. Based on the above observations, we hypothesize that DG and CA3 place fields form via the combination of two effects:

1. *Informational:* Specific sensitivity of individual cells to conjunctions and disjunctions of features from sensory, vestibular, directional, motivational, and motor information.
2. *Contextual:* Global selection of a cell group for activity in the particular environment and task at hand.

Similar ideas are inherent in some previously proposed models of the hippocampus, most notably in the work of Redish and Touretzky (Touretzky & Redish, 1996; Redish & Touretzky, 1997; Redish, 1997; Redish & Touretzky, 1999) and Samsonovich & McNaughton (1997). The term *context* throughout this article refers to a long-lasting bias that affects place representations for the duration of an episode (Quirk et al., 1990; Markus et al., 1994, 1995), unless disrupted by a significant mismatch between expectation and reality. This bias implicitly encodes the animal's perception of its situation, signifying a particular set of relationships, expectations, and conditions—and possibly, goals, actions, and consequences—that characterize the entire episode in a stable way and remain invariant over its duration. This should be distinguished from the transient influence of a system's recent dynamics on its current state (e.g., the influence of an animal's recent positions and actions on its current position). The latter type of context dependence has been proposed as a means for disambiguating stimulus sequences of finite length (Minai, Barrows, & Levy, 1994; Wu, Baxter, & Levy, 1995; Levy & Wu, 1996; Levy, 1996), and is quite likely to exist in the hippocampus. However, global context dependence requires a mechanism of much greater stability.

The issues involved in coding global context are illustrated well by the requirements of effective place representations. There are two conflicting imperatives:

> *Intraenvironment consistency.* Within a specific environment, the place representation must vary smoothly over space to encode consistent spatial information and allow generalization. If the place code is to serve as a robust representation of the animal's location, proximate locations should be coded similarly, with the similarity declining monotonically with distance. This is precisely what place fields achieve (as demonstrated by measures of spatial coherence in place fields; Muller et al., 1991; Sharp, 1997). The stability of place fields, in turn, requires dependence of pyramidal cell activity on information about the animal's location: sensory input (O'Keefe & Nadel, 1978; Muller & Kubie, 1987), vestibular cues (Hill & Best, 1981), and almost certainly, the estimate from a path integration (PI) system (Etienne, Maurer, & Sèguinot, 1996), which updates the animal's estimate of its location internally based on self-motion (Touretzky & Redish, 1996; Redish & Touretzky, 1997; Redish, 1997; Samsonovich & McNaughton, 1997).

Interenvironment discrimination. In two similar—even identical—environments with different significance, the place representations must be very different in order to discriminate between the situations. This means that the place representation must not depend exclusively on sensory and vestibular cues; it must also be shaped by the animal's perception of the global context it is in. This is borne out by the experiments of Quirk et al. (1990), who showed that an animal can have two distinct place representations of the same sensory environment in different contexts. The fact that place representations switch in a discontinuous manner (Quirk et al., 1990; Bostock et al., 1990; Gothard, Skaggs, Moore, & McNaughton, 1996; Gothard, Skaggs, & McNaughton, 1996; Barnes et al., 1997) also supports a global context-dependence hypothesis.

How does a system resolve these conflicting requirements? We propose that the hippocampus does so by using a two-part strategy. First, when a distinct context is recognized, it triggers the stable priming of a subgroup of dentate gyrus granule cells for firing and depresses the rest. Resetting this stable priming then requires a major disruptive event (e.g., removal from the environment). Once the selected group has been primed, the actual firing of cells within this group depends on sensory, vestibular, and path integration cues, thus producing place fields and a smooth representation of the environment in the DG (Minai & Best, 1998; Doboli, Minai, & Best, 1999). Since the projection from DG to CA3 is very tightly focused, the granule cells in the selected group excite only a subset of CA3 pyramidal cells, leading to smooth place fields in this subset and silence in the rest of the CA3. By selecting different groups in different contexts, the animal can thus create very different—but internally consistent—representations of sensorily identical environments.

This conjecture—which we call the *latent attractor hypothesis* for reasons that will be clear—is consistent with the statistical characteristics of place representations seen experimentally (Thompson & Best, 1989, 1990; Wilson & McNaughton, 1993). Indeed, a variant of this hypothesis has been suggested previously (though not formally developed) by Kubie and Muller (1991). It is also implicit in the chart theory advanced by Samsonovich and McNaughton (1997). However, our theory differs from these previous proposals in three major respects: (1) the computational mechanism we propose is very simple, with minimal preconfiguration of place-related network structure in the DG or CA3; (2) the stable priming is seen as originating in the dentate gyrus-hilus (DGH) system rather than the CA3; and (3) the system underlying the priming process is a two-layer (disynaptic) rather than a one-layer (monosynaptic) recurrent network (we use the terms *two-layer* and *one-layer* in their connectionist sense and not in the anatomical sense of cortical layers). The justification and implications of these choices are discussed below. More recently, Samsonovich (1999) has independently

proposed a simpler context-selection mechanism similar to ours in principle. However, it is still embodied in a one-layer CA3 model network. In contrast, we suggest that context dependence in CA3 place representations is not produced internally, but simply reflects the context dependence already present in the afferent input from the DG. We show that, based on current (mostly indirect) evidence, placing context selection in the DGH is logically and biologically reasonable.

2.3 Why the DGH System?. Previous suggestions about the encoding of context via multistable attractors have proposed CA3 as the likeliest locus (Kubie & Muller, 1991; McNaughton et al., 1996; Shen & McNaughton, 1996; Redish, 1997; Samsonovich & McNaughton, 1997; Samsonovich, 1999). However, we prefer the DGH system for several reasons:

- Granule cells show highly place-specific behavior like CA3 cells (Jung & McNaughton, 1993). There is also evidence that some granule cells are silent in a given environment (Jung & McNaughton, 1993).
- DG granule cells show the same context dependencies as CA3/CA1 place cells (Markus et al., 1995), suggesting that context is established by this point in the trisynaptic pathway.
- There is potential evidence of attractor dynamics in the DGH system (Shen, McNaughton, & Barnes, 1997).
- The DGH system has recurrent excitatory connectivity (Buckmaster & Schwartzkroin, 1994; Scharfman, 1995; Jackson & Scharfman, 1996; Golarai and Sutula, 1996) and a very structured inhibitory system (Han, Buhl, Lörinczi, & Somogyi, 1993).
- The recurrent connectivity in DGH is disynaptic (unlike CA3) and highly divergent and convergent (Scharfman, 1991), making it more suitable than the CA3 for implementing flexible multiattractor dynamics. In particular, the disynaptic recurrence allows the biasing function (needed for the latent attractor) to be located in the hilus and the coding function (needed for the projection to CA3) in the DG. This is important because biasing requires high source activity and diffuse efferent projection, while coding for CA3 representations requires sparse activity and a very focused efferent projection. A two-layer system implements this very well, and the required characteristics correspond well to the known anatomy and physiology of the DGH.
- The associational input from mossy cells to the granule cells shows NMDA-dependent LTP (Hetherington, Austin, & Shapiro, 1994).
- DG has a very large population of cells (Amaral, Ishizuka, & Claiborne, 1990) and low intrinsic activity, which would allow for many more attractors than CA3.

- DG granule cells and their mossy fiber projections are organized in a lamellar fashion (Andersen, Bliss, & Skrede, 1971; Amaral & Witter, 1989), which has been compared to a piano keyboard (Sloviter & Brisman, 1995). There is also strong evidence for mutual inhibition between lamellae (Struble, Desmond, & Levy, 1978; Sloviter, 1994). Finally, the mossy cell projection to granule cells is highly divergent. This architecture appears well suited to the formation of mutually inhibitory granule cell groups via the associative modification of mossy-to-granule cell synapses linking different lamellae.
- Using the CA3 recurrent connections for implementing context attractors for multiple environments might interfere with their use in pattern completion and associative memory recall (Treves & Rolls, 1992; Hasselmo, Schnell, & Barkai, 1995; Hasselmo, Wyble, & Wallenstein, 1996). Redish and Touretzky (1998) have shown that a CA3-like network can simultaneously accommodate pattern completion, memory recall, and contextual attractors, but the effect of this multiple functionality on system capacity is not yet known.

Essentially, our model proposes a division of labor between EC, DGH, and CA3. Many functions, including associative recall (Marr, 1971; McNaughton & Morris, 1987; O'Reilly & McClelland, 1994; Rolls, 1996; Treves & Rolls, 1992; Hasselmo et al., 1995, 1996), sequence completion (Minai & Levy, 1993; Minai et al., 1994; Wu, Baxter, & Levy, 1995; Levy & Wu, 1996; Levy, 1996; Sohal & Hasselmo, 1998; Lisman, 1999), path selection (Muller et al., 1991; Muller & Stead, 1996; Blum & Abbott, 1996), path integration (Samsonovich & McNaughton, 1997), and spontaneous replay of learned memories (Minai & Levy, 1993; Wilson & McNaughton, 1994; Shen & McNaughton, 1996; Skaggs & McNaughton, 1996; Redish & Touretzky, 1998) have been imputed to CA3 by various researchers. Meanwhile, most models have automatically assumed that the DG performs orthogonalization of sensory inputs (Treves & Rolls, 1992; O'Reilly & McClelland, 1994), and the hilus has been neglected completely. A notable exception to this is the recent model proposed by Lisman (1999), which suggests that the recurrent DGH functions as a pattern completion module in a system for sequence recall. Our hypothesis elegantly addresses three issues raised by these approaches. First, it provides CA3 place representations with the requisite context dependence (due to DG input) and informational contingencies (due to EC input) while leaving the CA3 free to perform other functions such as associating places along a path, disambiguating stimulus sequences, and recalling previously learned representations. Second, it leaves the orthogonalization function for the DG in place, with the subtle distinction that similar stimuli are orthogonalized only if they occur in different contexts or environments. This obviates the need to block DG activity during associative recall as hypothesized in earlier models (O'Reilly & McClelland, 1994; Redish, 1997; Minai, 1997). Finally, the hypothesis provides an

elegant explanation for the very rich recurrent architecture of the DGH system.

3 Computational Model

3.1 Functional Description. We model the DGH system as a recurrent system of granule-mossy-granule cell connections with various inhibitory subsystems. Both the DG and hilus are seen as organized into distributed non-disjoint excitatory cell groups (Andersen et al., 1971; Soltesz, Bourassa, & Deschenes, 1993; Sloviter & Brisman, 1995). Cells in the kth DG or hilus group tend to activate cells in the corresponding group in the other region, while each region uniformly inhibits all cells in the other region in proportion to its own activity. The net effect is to make the groups mutually inhibitory across the system. Due to competitive inhibition, the total allowed number of simultaneously active cells within a region is constrained to be fewer than the number of cells in a group. Since the groups are mutually inhibitory, a disproportionately high level of activity in one group tends to inhibit other groups strongly, quashing any incipient activity there and, as a consequence, perpetuating its own high level of activity. Thus, once a group "wins" the intergroup competition, activity will tend to stay within this group until it is strongly perturbed. In this sense, each group is an "attractor" of the system, but the whole attractor never becomes active. We hypothesize that the firing thresholds of granule cells are high enough that the system is capable of intrinsic (nonafferent driven) activity only under conditions of explicit disinhibition (Scharfman, 1991; Buckmaster & Schwartzkroin, 1995). When afferent activity from the EC is present, the effect of the recurrent system is to focus firing within a metastable subset of granule cells. We call this the *latent attractor effect*, since the effect of the EC activity is channeled by an attractor that cannot sustain itself autonomously, but whose effect is unmasked by the afferent input.

Formally, we define a *latent attractor* for a recurrent network of binary neurons as a nontrivial pattern of activity that is not realized (or only partially realized) at any particular time, but would be self-sustaining and attracting within a subdomain of the network's phase space under conditions of sufficient disinhibition or external stimulus. It is termed "latent" because it is implicit in the connectivity structure of the network but requires external stimulus to become visible, much as the peaks in a dark landscape may be picked out by a searchlight. This is different from the situation in other recurrent networks—including disynaptic ones (Carpenter & Grossberg, 1987a, 1987b; Kosko, 1988)—where the entire attractor is activated.

The activity of the EC cells in our hypothetical model encodes the various "informational" contingencies, including (but not limited to) conjunctions and disjunctions between sensory cues, vestibular information, and information from the path integration system. There is strong evidence for location-specific activity in EC neurons (Barnes et al., 1990; Quirk et al.,

1992), and it has been hypothesized (Touretzky & Redish, 1996; Redish & Touretzky, 1997) that the EC forms a stage in the rat's path integration system. Although we do not model these here, the firing contingencies of EC neurons are seen as being encoded through LTP-based learning and competitive mechanisms as suggested by Shapiro et al. (1997). Since the EC provides input to both DG and CA3, the same contingencies appear in neurons of those regions, and all would respond to cue manipulations in a heterogeneous manner consistent with the multiplicity of informational inputs to the system (Shapiro et al., 1997; Tanila et al., 1997; Knierim et al., 1998).

At the beginning of an episode, a particular pattern of input from the EC (a "recognition stimulus") excites DG granule cells disproportionately within a specific group, focusing future activity within this groups. For the duration of the episode, then, granule cells are driven primarily by the afferent input, but cells that are not part of the selected group are effectively suppressed. As the animal moves through its environment, the granule cell activity is projected onto the CA3 pyramidal cells where, in combination with the EC input, it appears as context-specific place fields. Place is thus coded with a two-level hierarchical representation. At the first, coarser level, each context is represented by the selection of active cell groups. Within this, individual cells are fired by sensory and ideothetic afferent information.

The disynaptic recurrence in our model is important because subcortical inputs (e.g., theta from the septum) can precisely control the recurrent loop via modulation of mossy and inhibitory cell excitability. Direct evidence for exactly such control exists (Deadwyler, West, & Robinson, 1981; Moser, 1996). However, we do not address this issue here.

3.2 Neuron Equations. The environment for our simulations is an $M \times M$ grid, (u, v), on which the simulated animal moves randomly. The position of the animal at time t is denoted by $L(t) = (L_u(t), L_v(t))$. The network has four layers: EC, DG, H (hilus), and CA3.

3.2.1 EC Layer. The EC layer has N_{EC} neurons, and is modeled at a purely phenomenological level, based on the known characteristics of EC place fields (Barnes et al., 1990; Quirk et al., 1992). Each cell, i, is given a broad, noisy gaussian place field defined by:

$$\begin{aligned} f_i(t) = exp[&-a_i(L_u(t) - c_i^u + q_u)^2 - b_i(L_v(t) - c_i^v + q_v)^2 \\ &+ d_i\sqrt{a_i}(L_u(t) - c_i^u)\sqrt{b_i}(L_v(t) - c_i^v)\ ], \end{aligned} \tag{3.1}$$

where $\bar{c}_i = (c_i^u, c_i^v)$ is center for i's field, a_i, b_i are field size and shape parameters, $d_i \in [-1,\ 1]$ is the field orientation parameter, and q_u, q_v are 0-mean positional noise with variance σ_q^2. All of these parameters are chosen randomly within reasonable ranges for each cell. As the animal moves on the

grid, neuron i is fired noisily, as

$$z_i(t) = f_i(t) + \eta(t) + \kappa, \tag{3.2}$$

where $\eta(t)$ is zero-mean uniform noise with variance s^2 and κ is a baseline activity level (set to a small value). The value of $z_i(t)$ is clipped to 0 if noise would make it negative. While rather simple, this method provides a sufficiently accurate model of the broad, noisy EC place fields seen experimentally (Quirk et al., 1992).

3.2.2 DG Layer. The DG layer has N_{DG} simulated granule cells that get excitatory input from the EC and H layers, and a broad inhibitory input from the H layer. In addition, the DG layer is assumed to have competitive recurrent inhibition, which maintains activity within a narrow range. This recurrent inhibition is not modeled explicitly, but is reflected in the stochastic K-of-N firing rule. The activation of granule cell i at time t is given by:

$$\begin{aligned} y_i(t) = {} & g_{EC-DG} \sum_{j \in EC} w_{ij} z_j(t) + g_{H-DG} \sum_{j \in H} w_{ij} z_j(t-1) \\ & - G_{H-DG} \sum_{j \in H} z_j(t-1), \end{aligned} \tag{3.3}$$

where w_{ij} is the connection strength from neuron j to neuron i, and g_{EC-DG}, g_{H-DG}, and G_{H-DG} are gain parameters that are set randomly within narrow ranges. The firing of cell i is based on the following rule:

$$z_i(t) = \begin{cases} 1 & \text{with probability } \rho_i(t) \\ 0 & \text{otherwise} \end{cases} \tag{3.4}$$

where $\rho_i(t)$ is defined as follows. Let $\Gamma_1(t)$ be the set of K_{DG} most excited DG cells at time t, $\Gamma_2(t)$ the set of the next K_{DG} most excited ones, and $\Gamma_3(t)$ the remaining neurons. Then,

$$\rho_i(t) = \begin{cases} r_1^{DG} & \text{if } i \in \Gamma_1(t) \\ r_2^{DG} & \text{if } i \in \Gamma_2(t) \\ r_3^{DG} & \text{if } i \in \Gamma_3(t) \end{cases} \tag{3.5}$$

where r_1^{DG} is close to 1, r_2^{DG} is a small value close to 0, and r_3^{DG} is an even smaller value, which indicates the sporadic activity of noncoding neurons. The quantity $\beta_{DG} = r_1^{DG} / \sum_k r_k^{DG}$ quantifies the reliability of the K_{DG}-of-N_{DG} competitive firing mechanism. Qualitatively, the rule means that almost all of the K_{DG} most excited neurons fire, a few of the next K_{DG} most excited ones fire, and the rest fire only sporadically, giving an approximate net

firing of K_{DG} neurons per time step. Neurons with $y_i(t) \leq 0$ are explicitly barred from firing, so that under very weak stimulus, the total firing may be significantly below K_{DG}. However, we do not consider this case here. Consistent with experimentally observed data, K_{DG} is taken to be a small fraction (< 5%).

3.2.3 H Layer. The H layer has N_H mossy cells, which receive excitatory input from the DG layer. The activation of cell i at time t is given by:

$$y_i(t) = g_{DG-H} \sum_{j \in DG} w_{ij} z_j(t), \tag{3.6}$$

where g_{DG-H} is a gain parameter. Firing is K_H-of-N_H implemented as for the DG layer with parameters r_k^H.

3.2.4 CA3 Layer. The CA3 layer has N_{CA3} simulated pyramidal cells, which get very sparse excitatory input from the DG layer and a diffuse excitatory input from the EC. The DG-CA3 synapses are taken to be much stronger than EC-CA3 ones, so the net effects of the two projections are approximately balanced. The CA3 also has extensive recurrent connectivity. However, these connections are not included in our current model in order to avoid obscuring our main hypothesis with additional effects. The recurrent projection might, of course, interact significantly with the context, but we see it primarily as a substrate for associative recall of stored, context-dependent representations (Marr, 1971; Rolls, 1996; Hasselmo et al., 1995, 1996) and perhaps for path learning (Muller & Stead, 1996; Blum & Abbott, 1996; Mehta, Barnes, & McNaughton, 1997).

The activation of CA3 cell i at time t is given by:

$$\begin{aligned} y_i(t) = {} & g_{EC-CA3} \sum_{j \in EC} w_{ij} z_j(t) + g_{DG-CA3} \sum_{j \in DG} w_{ij} z_j(t) \\ & - G_{DG-CA3} \sum_{j \in DG} z_j(t) \end{aligned} \tag{3.7}$$

The firing of neurons is implemented by a competitive K_{CA3}-of-N_{CA3} rule similar to the one for DG with parameters r_k^{CA3}.

3.3 Network Architecture. The EC-DG and EC-CA3 connections are selected randomly with each DG cell receiving input from a fraction C_{EC-DG} of EC cells and each CA3 cell from a fraction C_{EC-CA3}. The connections are made such that the fan-outs from individual EC neurons are approximately equal. The actual weight values are chosen randomly in a fixed range. The heart of our model is the architecture of the DG-hilus subsystem, which is as follows.

We choose m cell groups, $Q_1, \ldots, Q_m$, of excitatory cells in both the DG and H layers. All groups in DG have size n_{DG}, and those in H have size n_H. The cells in each group are chosen randomly and independently of other groups. A uniform random connection pattern is established between the DG and H layers in both directions, with the connectivity fractions given by C_{DG-H} and C_{H-DG} for DG-to-H and H-to-DG connections, respectively. As in the case of EC-DG and EC-CA3 connections, the fan-outs of presynaptic neurons are approximately balanced. The synaptic weights for these connections are then set as follows:

DG → H Weights

$$w_{ij} = \begin{cases} h_{DG-H} & \text{if } i \in Q_k \text{ and } j \in Q_k \text{ for some } k = 1, \ldots, m \\ l_{DG-H} & \text{else} \end{cases} \tag{3.8}$$

H → DG Weights:

$$w_{ij} = \begin{cases} h_{H-DG} & \text{if } i \in Q_k \text{ and } j \in Q_k \text{ for some } k = 1, \ldots, m, \\ l_{H-DG} & \text{else} \end{cases} \tag{3.9}$$

where $h_{DG-H} \gg l_{DG-H}$ and $h_{H-DG} \gg l_{H-DG}$. This very simple connectivity, which is essentially equivalent to the embedding of each group as an attractor via a two-level Hebbian rule (Willshaw, Buneman & Longuet-Higgins, 1969), is sufficient to set up the required system behavior. The competitive firing rule then ensures that only one attractor is active at a time, and that only in part.

The connections from DG to CA3 are very sparse, with connection probability C_{DG-CA3}, but the gains, g_{EC-CA3} and g_{DG-CA3} are set such that the effect of EC and DG inputs on the CA3 is roughly equal.

4 Model Evaluation Methods

We evaluate our model system by studying whether it satisfies the two conflicting requirements needed for place representation: intraenvironment consistency and interenvironment discrimination. The simulated animal is run in two sensorily identical environments, A and B (both have identical EC place fields), for a large number of steps (typically 5000–10000) along random paths. The initial input to the DG layer is significantly different in the two cases, signaling the difference of context at the time of entry. To simulate repeated trials in the same environment, the context input is similar each time, simulating the initial sensory input whose recognition causes the animal to decide on a specific context. A randomly chosen subset, S, of CA3 cells is monitored as the animal moves about, and the empirical place fields for these cells are reconstructed as follows. First, the conditional probability

of cell i firing at a location is estimated as:

$$P^1_{i|u,v} \equiv P(z_i(t) = 1 \mid L_u(t) = u, L_v(t) = v) \approx \frac{\sum_t z_i(t)\delta_{uv}(t)}{\sum_t \delta_{uv}(t)}, \tag{4.1}$$

where $\delta_{uv}(t) = 1$ if $L_u(t) = u,\ L_v(t) = v$, and 0 else. If $\sum_t \delta_{uv} = 0$ (location never visited), $P^1_{i|u,v}$ is also set to 0. The place field for cell i is then calculated as:

$$f_i^e(u,v) = \begin{cases} \alpha & : \quad P^1_{i|u,v} < \alpha \\ P^1_{i|u,v} & : \quad \alpha \le P^1_{i|u,v} \le \omega \\ \omega & : \quad P^1_{i|u,v} > \omega, \end{cases} \tag{4.2}$$

where $e \in \{A, B\}$ denotes the environment. Parameters $\alpha = r_3^{CA3}$ and $\omega = r_1^{CA3}$ are used to model the fact that no CA3 cell, under our firing rule, has probability less than r_3^{CA3} or greater than r_1^{CA3} of firing at any location, but this may not be picked up in the finite sample used to estimate $P^1_{i|u,v}$.

Intraenvironment consistency is measured by the extent to which the reconstructed place fields allow accurate localization of the simulated animal (Wilson & McNaughton, 1993; Zhang, Ginzburg, McNaughton, & Sejnowski, 1998; Brown, Frank, Tang, Quirk, & Wilson, 1998). This is done using a Bayesian maximum likelihood formulation as follows. The activity of cells in set S_e is monitored as the simulated animal moves on a random path in environment e to get the activity vectors $Z(t) = \{z_k(t)\};\ k \in S_e$. The set $S_e \subset S$ includes only those cells that have significant place fields in environment e. This is determined as follows. The activity of a cell in every 3×3 neighborhood of the environment is considered, and the number of neighborhoods in which at least 7 of the 9 cells have above average firing is counted. If this number exceeds 10 (out of a possible 324 neighborhoods for a 20×20 environment), the cell is deemed to have a place field and is included in S_e. We also impose a continuity constraint (Zhang et al., 1998) based on the previous estimated position in order to avoid unrealistic jumps from one step to another. Thus, the posterior probability function at each location depends not only on the current firing pattern but also on the previous reconstructed position $(\hat{L}_u(t-1), \hat{L}_v(t-1))$:

$$P(u, v \mid Z(t), \hat{L}_u(t-1), \hat{L}_v(t-1)) = CP(u, v \mid Z(t)) \\ P(\hat{L}_u(t-1), \hat{L}_v(t-1) \mid u, v) \tag{4.3}$$

where C is a normalizing factor, which does not depend on (u, v), and

$$P(u, v \mid Z(t)) = \prod_k [P(u, v \mid z_k = 1) z_k(t) + P(u, v \mid z_k = 0)(1 - z_k(t))] \tag{4.4}$$

with

$$\begin{aligned} P(u, v \mid z_k = 1) &= P(z_k = 1 \mid u, v) P(u, v) / P(z_k = 1) \\ &= f_k^e(u, v) P(u, v) / \langle z_k \rangle \qquad (4.5) \\ P(u, v \mid z_k = 0) &= P(z_k = 0 \mid u, v) P(u, v) / P(z_k = 0) \\ &= [1 - f_k^e(u, v)] P(u, v) / (1 - \langle z_k \rangle) \qquad (4.6) \end{aligned}$$

and

$$\begin{aligned} &P(\hat{L}_u(t-1), \hat{L}_v(t-1) \mid u, v) \\ &\quad \sim exp\left[-\|(\hat{L}_u(t-1), \hat{L}_v(t-1)) - (u, v)\|^2 / 2\sigma^2 \right]. \qquad (4.7) \end{aligned}$$

Angular brackets indicate time averages and $\|.\|$ the Euclidian distance. The value of σ represents the width of the gaussian function centered at the previous estimated position.

Assuming $P(u, v) = M^{-2}$ (all locations equiprobable), we calculate the estimate of the current location:

$$(u^*(t), v^*(t)) = \arg \max_{(u,v)} P(u, v, \mid Z(t), \hat{L}_u(t-1), \hat{L}_v(t-1)). \qquad (4.8)$$

As the animal moves along its path, the error (Euclidean distance) between the estimated and true positions is calculated, and its mean value is used as a measure of localization performance:

$$\lambda = \frac{1}{l} \sum_{t=1}^{l} \|(L_u(t), L_v(t)) - (\hat{L}_u(t), \hat{L}_v(t))\| \qquad (4.9)$$

where l is the length of the path.

Interenvironment discrimination is measured using the mean correlation coefficient of the reconstructed place fields for the two environments:

$$\begin{aligned} &\xi(A, B) \\ &\quad = \frac{1}{|S'|} \sum_{i=1}^{|S'|} \frac{\begin{array}{c} M^2 \sum_{u,v} f_i^A(u, v) f_i^B(u, v) \\ -(\sum_{u,v} f_i^A(u, v))(\sum_{u,v} f_i^B(u, v)) \end{array}}{\begin{array}{c} \sqrt{M^2 \sum_{u,v} f_i^A(u, v)^2 - (\sum_{u,v} f_i^A(u, v))^2} \\ \times \sqrt{M^2 \sum_{u,v} f_i^B(u, v)^2 - (\sum_{u,v} f_i^B(u, v))^2} \end{array}}, \qquad (4.10) \end{aligned}$$

where S' is the set of monitored cells that have place fields in at least one of the environments (i.e., $S' = S_A \cup S_B$), and $|S'|$ denotes the cardinality of S'. A low value for $\xi(A, B)$ indicates that the place representations for environments A and B are significantly different. However, because of the pervasive

noise in the system, the value of $\xi(A, B)$ itself is not a sufficient indicator of discrimination and must be compared with the correlation between place representations within the same environment during two different episodes. To this end, we also calculate a quantity $\xi(A, A)$ analogous to $\xi(A, B)$ but calculated over two sessions in environment A with the same initial condition in both sessions. A high positive $\xi(A, A)$ indicates that the place representation in environment A is stable over different sessions. The interenvironment discrimination achieved by the system is then measured as

$$D(A, B) = \langle \xi(A, A) \rangle - \langle \xi(A, B) \rangle, \tag{4.11}$$

where angular brackets indicate averaging over many independent A and B environments.

A significantly positive $D(A, B)$ indicates that the system actually discriminates between A and B, while a value close to zero means that the difference in place codes for A and B arises only due to the noise within the system.

In order to provide a baseline for our evaluations and determine the utility of the latent attractor architecture, we compare our system with one that is identical to ours in every way except that the DGH system does not have groups. The actual connectivity profiles of all postsynaptic neurons are matched exactly for the two systems. The goal is to determine how much the group architecture helps with context encoding beyond what one might get from a system without groups. The system with groups is henceforth termed the *latent attractor* (LA) model, while the network without groups is called the *no latent attractors* (NLA) model. The independent parameter of interest is defined as $R = g_{EC-DG}/g_{H-DG}$, which quantifies the relative strength of the afferent (informational) and recurrent (contextual) inputs to the DG layer.

5 Results

We evaluate our model by simulating LA and NLA networks with layers of the following sizes: $N_{EC} = 200$, $N_{DG} = 1000$, $N_H = 500$, and $N_{CA3} = 300$. The H-to-DG gain, g_{H-DG}, is fixed at 0.5, and g_{EC-DG} is varied systematically to change R. The LA network has 10 cell groups, with 100 cells per group in DG and 50 cells per group in H. This implies that the mean overlap between DG groups is 10, about 60 cells in each group belong to at least one other group, and about 350 cells are not part of any group. The values of other parameters are as given in the appendix.

Figure 2 shows the empirically calculated localization and discrimination values for the LA and NLA systems as a function of R. The two systems are quite similar in terms of localization. When R is small, the DG gets little sensory information, leading to poorly formed CA3 place fields (see Figure 3a) and poor localization. As R increases, the place fields improve

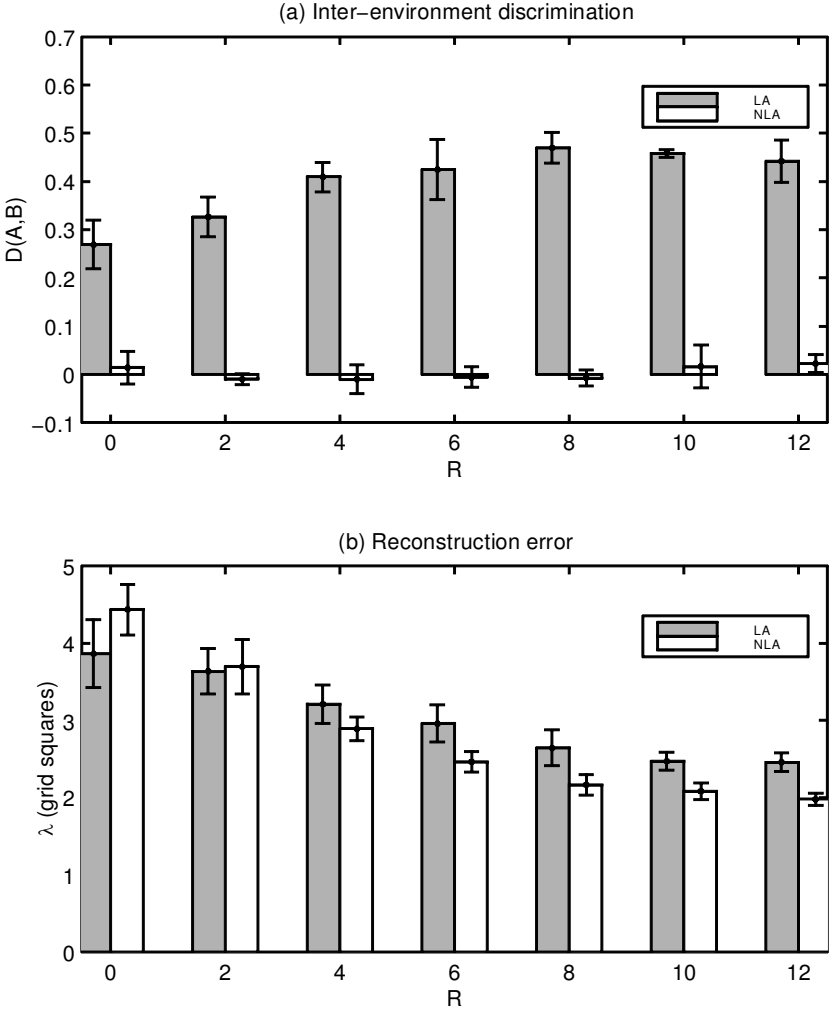

Figure 2: (a) The interenvironment discrimination, $D(A, B)$, for the LA and NLA systems. Environments A and B (size 20×20) have identical EC place fields and differ only in the initial afferent stimulus given to the DG. The data are based on output from 200 CA3 cells whose place fields were reconstructed over 5000 steps in each environment. (b) The localization accuracy over 100 steps in Environment A for both systems. Only place fields meeting a quality criterion were used in localization. Each bar was averaged over five runs. Error bars indicate one standard deviation above and below the mean.

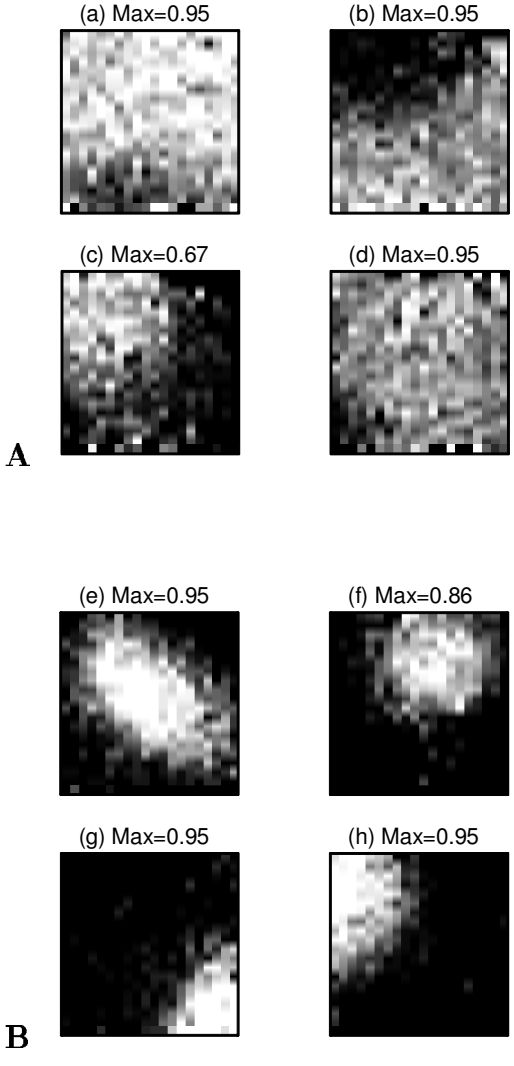

Figure 3: (A) Place-dependent activity for weak EC input to the DG layer ($R = 1.0$). Graphs a and b are for two typical CA3 cells the LA system, and graphs c and d for two NLA system cells. Note the lack of coherent place activity for both systems, which explains poor localization. (B) Place-dependent activity for strong EC input to the DG layer ($R = 12.0$). Graphs e and f are for two CA3 cells in the LA system and graphs g and h for two NLA system cells. Note that the fields are well formed, and localization is good. Each figure uses 24 gray-scale levels between 0 (black) and Max (white), where Max (indicated above each graph) is the peak probability of firing for the cell at any location in the environment.

(see Figures 3b and 4), and so does localization. The NLA system localizes slightly better than the LA system because its place code is distributed over more active neurons. However, this is inconsistent with the observation that only a small percentage of CA3 place cells are active in any one environment. A point worth noting is that the CA3 place representation in the LA system shows interenvironment discrimination even though the CA3 is structurally homogeneous and does receive a strong direct input from the EC. Thus, the partial influence from the DG is enough to produce a high degree of context dependence in CA3 place fields. This is reflected in the place fields for environments A and B, as shown in Figure 4.

The NLA system shows no interenvironment discrimination at any value of R. Although $\xi(A, B)$ is fairly small for low R, it is matched almost exactly by $\xi(A, A)$. Thus, the lack of correlation between A and B place fields for low R does not indicate any memory of the initial difference in context, but simply reflects the noise inherent in the system. As R increases and more sensory information (which is identical for A and B) gets through to the DG, the correlation increases identically for both the A-A and A-B cases. These results show that the recurrent loop in the DGH subsystem is not able to provide an adequate memory of the difference in initial conditions over the course of a session.

The LA system, in contrast, shows strong interenvironment discrimination at all R levels, though the value drops slightly, as expected, at high R. With increasing R values (not shown), we expect discrimination to be lost eventually, but the LA system clearly provides high discrimination and good localization over a broad range of R values.

One point of note is that even at best, $D(A, B) < 1$, whereas one might expect a value close to 1. The reason is that $\xi(A, A) < 1$ due to factors such as noise and path dependence. We now consider this issue in more detail. Ideally, a noise-free system run along identical paths in identical environments with the same initial conditions will produce perfectly correlated place fields ($\xi(A, A) = 1$). Thus, the factors that reduce the correlation below 1 are:

- *Path dependence.* Difference in place fields reconstructed over the different paths taken in the two sessions. This occurs because the recurrence in the DGH system makes DG (and, thus, CA3) activity dependent on previous states of the system.
- *Context dependence.* Difference in place fields due to the difference in initial conditions (indicating context) for the two sessions.
- *Noise.* Difference in place fields for the two sessions due to independent noise sources: spatial noise in EC fields; additive noise in EC cell activity; and stochasticity in competitive firing of DG, H, and CA3 neurons.
- *Differences in afferent input.* Place code differences due to real differences in the afferent sensory input.

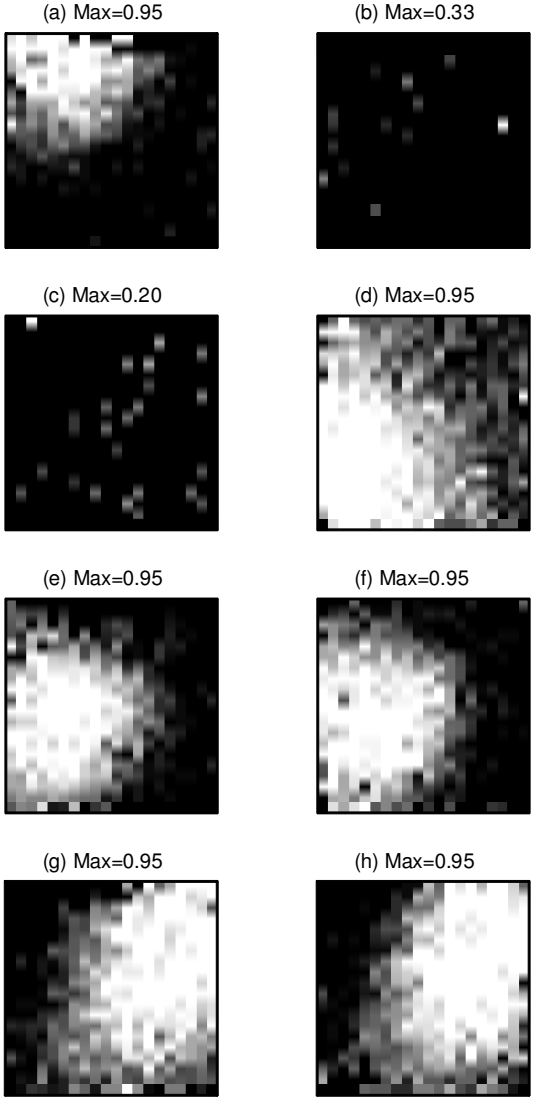

Figure 4: Activity patterns of the same cells in Environments A (left column) and B (right column) for $R = 6.0$. Note that the cells in the LA system (graphs a–d) have distinct place fields in the two environments, while cells in the NLA system (graphs e–h) do not, showing that the NLA system does not discriminate well on the basis of context. Many CA3 cells in the LA system did not have place fields in one or both environments, while almost all NLA cells had similar fields in both, which is contrary to experimental observation. Gray-scale coding is as in Figure 3.

Simply looking at $D(A, B)$ (or even $\xi(A, A)$ and $\xi(A, B)$) does not indicate the degree to which these factors affect the system's discrimination quality. The data shown in Figure 5 attempt to do this. The leftmost bars in graphs a and b show $\xi(A, A)$ in noise-free LA and NLA systems, respectively. All three sources of noise are eliminated, and the EC place fields and the initial conditions are identical for the two sessions. Thus, the residual lack of correlation, $1-\xi(A, A)$, can be ascribed completely to path dependence. Clearly, the NLA network is much more susceptible to path dependence than the LA network, though both improve as R increases and the similar sensory input from EC dominates the effects of the DGH recurrence.

The second set of bars in both graphs show $\xi(A, A)$ with all sources of noise present, but the EC place fields and initial conditions are again identical in the two sessions. These are the $\xi(A, A)$ values used in determining $D(A, B)$ for Figure 2. A comparison with the leftmost bars shows how much additional loss of correlation is caused by noise. Both systems are affected significantly (as expected).

Finally, the rightmost set of bars in the two graphs show $\xi(A, B)$, that is, place code correlation when EC inputs are the same but the initial conditions (contexts) are different and all noise sources are present. The NLA system correlations show no change from the situation with identical context (middle bars), while the LA system correlations drop dramatically, indicating context-based discrimination of place codes.

These results verify that dissimilarity in NLA representations for environments A and B is accounted for wholly by effects of noise and path dependence. For completeness, we also simulated the cases where the sensory input was different (different EC place fields) but the initial condition was the same in the two sessions, and the case when both EC fields and the initial conditions were different. Both LA and NLA systems produced substantially distinct representations in these cases, showing that differences in sensory input remain salient (we argue later that path integration can partially override these).

Finally, we also briefly considered the issue of network capacity for latent attractors. A full investigation of capacity (such as that carried out by Battaglia & Treves, 1998, for the charts system) is beyond the scope of this article, but we used simulations to study how the number of stable latent attractors in DG changed as the size of attractors, n_{DG}, increased relative to the system size, N_{DG}. This was parameterized by the ratio $\zeta = n_{DG}/N_{DG}$. The gain parameters were adjusted to keep the expected excitation to DG neurons in the active group constant. Group stability was evaluated by running the LA system in environment A for 110 steps and checking the DG activity pattern for steps 101 through 110. The degree of confinement to the initially activated (i.e., correct) group, Q^*, was calculated as

$$\psi(Q^*) = \frac{1}{10}\sum_{t=101}^{110}\left\{\frac{n_{DG}}{K(t)}\left[\frac{\sum_{i\in Q^*} z_i(t)}{n_{DG}} - \frac{\sum_{i\notin Q^*} z_i(t)}{N_{DG}-n_{DG}}\right]\right\}, \tag{5.1}$$

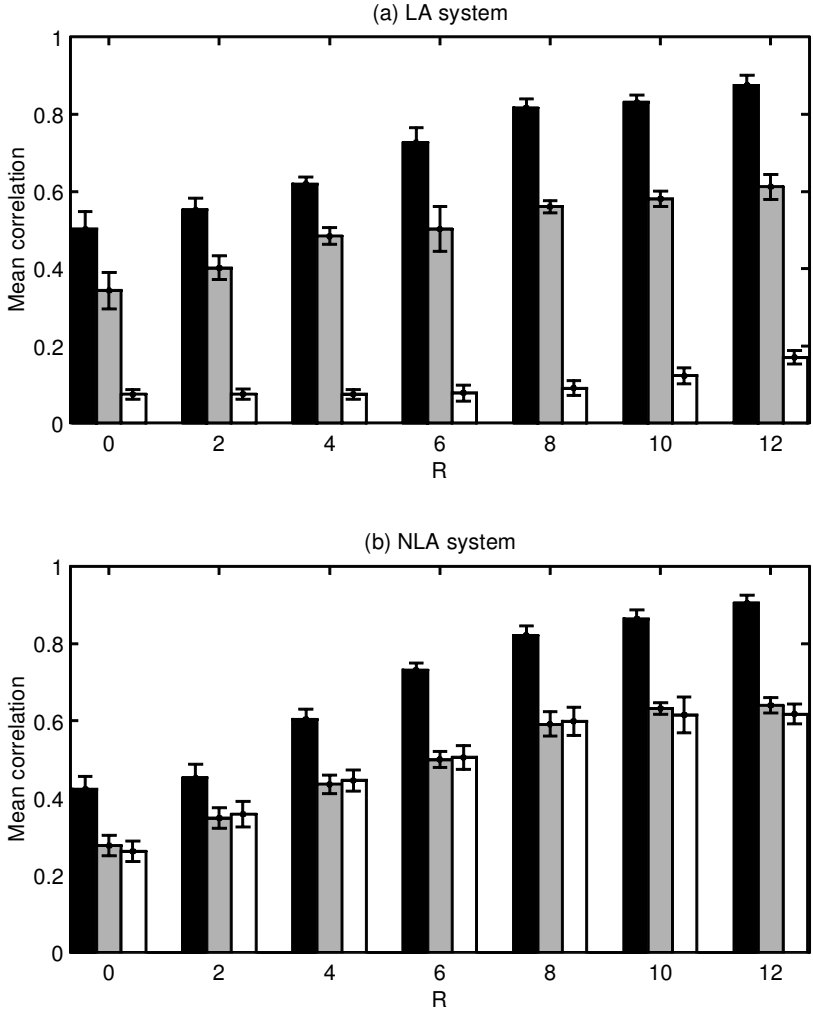

Figure 5: Mean correlations between reconstructed place fields for two sessions. The stimulus and network conditions in the two sessions are: (1) Black bars: Different paths, identical initial states (contexts), identical EC fields, and no noise; (2) Gray bars: Different paths, identical initial states (contexts), identical EC fields, and noise; (3) White bars: Different paths, different initial states (contexts), identical EC fields, and noise. See text for discussion.

where $K(t)$ is the number of DG neurons firing at step t. A value of $\psi(Q^*) = 1$ indicates total confinement of activity to group Q^*, while $\psi(Q^*) = 0$ means that confinement is no better than for activity distributed randomly over

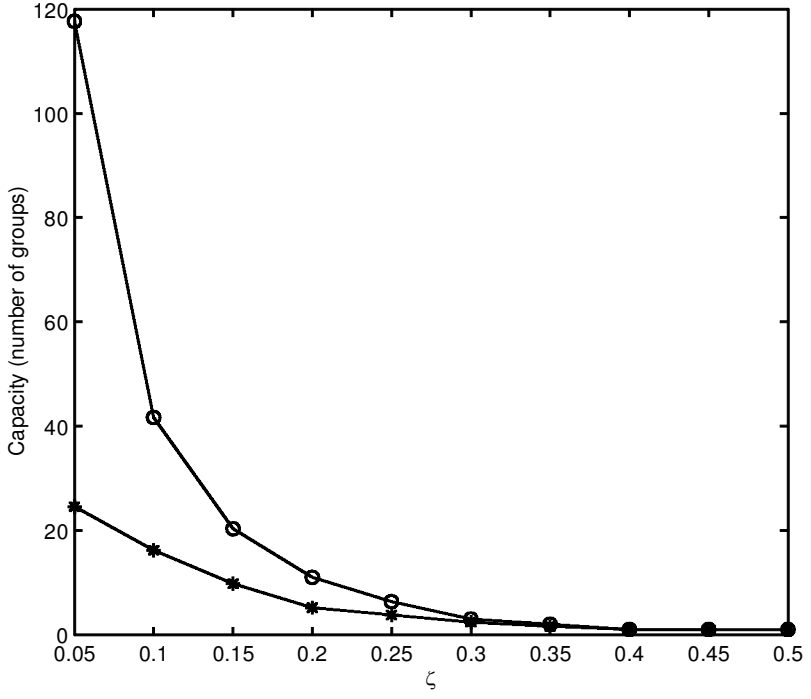

Figure 6: Capacity of the DGH system as a function of ζ, the size of the embedded latent attractors relative to network size. Network sizes are: $N_{EC} = 400$, $N_{DG} = 2000$, $N_H = 1000$ ($\circ$); $N_{EC} = 200$, $N_{DG} = 1000$, $N_H = 500$ ($*$) (the size used in all other simulations). All connectivities are as shown in the appendix. See the text for details.

the whole DG. For each ζ value, we determined the maximum number of groups that could be stabilized, where stabilization was defined by the following criterion: the mean ψ over all groups was at least 0.85, and no group had a ψ value less than 0.7. This was taken as the network's capacity. Figure 6 shows this capacity as a function of ζ for networks of two sizes. In both cases, we find that capacity decreases exponentially with ζ, but when groups are small compared to the neuron populations, a large number of latent attractors can be stored.

6 Discussion

6.1 Overall Significance of the Model. The primary result of our modeling is that the inclusion of a simple structured recurrent module (the DGH subsystem) in the hippocampal complex is sufficient to produce long-term,

stable, and reliable context dependence in CA3 place fields. However, an unstructured recurrent module is insufficient for this purpose.

The LA model we present is a very simple—even simplistic—one and does not account for many important aspects of hippocampal function. Its purpose, however, is to illustrate the latent attractors idea in the simplest possible terms within the context of hippocampal place representations. There has been at least one previous model that proposes mechanisms similar to ours: the charts model of Samsonovich and McNaughton (1997). While elegant and insightful, it involves considerably more detailed preconfiguration of the system. A primary aim of our model is to demonstrate that a much simpler system is capable of essentially the same type of behavior as that hypothesized in the charts model. Unlike the charts model, we do not assume any intrinsic, noise-free place specificity in our DG or CA3 cells, but allow it to emerge through convergence of EC afferents and competitive firing. In our model, place-specific firing is explicitly produced only in the EC layer and is very noisy and unreliable. Also, unlike the charts model, we do not configure the recurrent connectivity in our system to reflect the place fields of individual cells in different contexts. Our DGH system is connected randomly, the groups comprising the attractors are randomly chosen, and connection weights have only two possible values (rather than a continuum). The results indicate that although the detailed chart structure does produce better location estimates and more compact place fields, only minimal network structure is required to obtain the degree of localization, place specificity, and context dependence seen in actual recordings (Zhang et al., 1998). The same conclusion is reached by Samsonovich (1999) for a model similar to ours in principle but located in CA3 with connections specified via an off-line Hebbian prescription.

In spite of its simplicity, our model captures many essential properties of hippocampal place representations, including the following:

- The model accounts for different, randomly determined populations of active cells and silent cells in different environments (Thompson & Best, 1989, 1990).

- The model provides a simple explanation of the remapping phenomenon (Bostock et al., 1991; Shapiro et al., 1997; Tanila et al., 1997; Knierim et al., 1998), whereby place representations switch in a discontinuous manner. In our model, this can be seen as a switch in the selected groups. Partial remappings can be explained by assuming that the overlaps between some groups are larger than usual, or by more complex connectivity patterns that allow multiple groups to be active together and switch separately. We do not explore this issue in the current simulations.

- The model provides a simple mechanism for producing different place representations in identical sensory environments based on initial ex-

pectation (context) (Quirk et al., 1990), without interfering unduly with the reliable reproducibility of each representation within its context.

- The model shows how broad, noisy, and unreliable place fields in the EC can lead to much more reliable, context-specific place representations in the CA3 (Quirk et al., 1992). The increased accuracy, of course, is a direct result of averaging over multiple EC afferents performed by each DG and CA3 cell. The context dependence is a consequence of group selection in the DG.

6.2 Predictions. As a means of verifying (or rejecting) our hypotheses, we offer a set of experimental predictions that could be of interest to experimentalists:

DG place fields will show the same phenomena of "silent cells" and coding of similar environments by uncorrelated cell groups as the CA3/CA1. This prediction, if true, would mean that the phenomena observed in CA3/CA1 place codes (Thompson & Best, 1989) are already present at the DG stage in the trisynaptic pathway. There is already some evidence for silent cells (Jung & McNaughton, 1993) and context-dependence (Markus et al., 1995) in DG place representations, but, to our knowledge, there is no systematic study of the correlatedness of DG place codes for similar environments (i.e., whether these codes are more EC-like—Quirk et al., 1992—or more CA3-like—Thompson & Best, 1989; Kubie & Muller, 1991). A potential problem is that of locating silent granule cells, since granule cell firing is rather sparse and easily swamped by interneuron activity. However, disinhibition by bicuculline has been used successfully to elicit activity—albeit abnormal—from inactive granule cells (see, e.g., Sloviter and Brisman, 1995), and in vivo granule cell identification procedures have also been used successfully by Jung and McNaughton (1993).

CA3/CA1 place fields following a lesion of the DG using colchicene (McNaughton et al., 1989) would show significantly reduced context-dependence. It is known (McNaughton, Barnes, & O'Keefe, 1983; Knierim & McNaughton, 1995) that, at least in strongly directional environments such as arm mazes, CA3/CA1 place fields persist after DG lesion—presumably because of the direct input from EC to CA3. There is also evidence that EC place representations are not context dependent in the sense that similar environments produce similar EC representations (Quirk et al., 1992) (though Mizumori, Ward, & Lavoie, 1992 report that they do show directional firing on arm mazes). This contrasts with the context dependence seen in the CA3/CA1 fields (O'Keefe & Nadel, 1978; Muller & Kubie, 1987; Thompson & Best, 1989; Kubie & Muller, 1991; Quirk et al., 1990). If the DGH is responsible for this context dependence in CA3/CA1, its lesion should eliminate it, or at least reduce it significantly.

CA3/CA1 place codes in rats with mossy cell loss will show significantly reduced context dependence. Rats with hilar cell loss—usually due to kainic acid

application—are used as a model for temporal-lobe epilepsy (Nadler, Perry, & Cotman, 1978, 1980). It is widely believed that the cells lost include mossy cells, whose loss leads to a disruption in the pattern of inhibition to granule cells (Sloviter, 1987, 1994; Sloviter & Brisman, 1995; Gibbs, Shumate, & Coulter, 1997). Our hypothesis predicts that rats with hilar cell loss will also show less context dependence in their CA3/CA1 place codes.

Environments with similar entry points will be coded by cell groups with greater than usual overlap. This is a somewhat speculative prediction, but is included because it would be very easy to test and, if correct, would say a great deal. The idea is that if cell groups for representing novel environments are chosen at the time of entry (as hypothesized), environments with identical or similar entry areas would tend to get assigned to the same DGH cell group. Subsequently encountered differences in sensory input within each environment would produce different place representations in the two cases, but using the same group of cells, implying the prediction. There are several caveats. First, there may be other mechanisms that ensure random rather than stimulus-based group assignment. Second, group assignment during the first encounter with an environment may happen too slowly to show the effect. Third, the phenomenon might not happen reliably and may be observable only over a large enough sample of environment pairs.

6.3 Extensions and Open Issues. The work reported in this article addresses only some of the issues that our model is directed toward exploring. These include the consideration of more subtle latent attractor schemes, the explicit use of pattern recognition in the EC-DG projection to trigger appropriate groups, and the inclusion of path integration in the model. The last issue is of special significance because some of the experimental results that motivated this research require path integration for adequate modeling. For example, the model in its current form cannot directly duplicate the light and dark results of Quirk et al. (1990), though it certainly suggests a detailed explanation. As indicated earlier, we agree with the hypothesis (Touretzky & Redish, 1996; Redish & Touretzky, 1997) that path integration information enters the hippocampus through the EC. The simplest mechanism would be if the positional information reconstructed by the PI system directly elicited approximately correct firing in the EC (Redish & Touretzky, 1998). The rest of the system could then function as under normal circumstances. This framework can also be used to address the issues of hierarchical conflict resolution between distal sensory, local sensory, and ideothetic information that apparently occurs in the hippocampus (Hill & Best, 1981; Shapiro et al., 1997; Knierim et al., 1998). If EC cells are sensitive to subsets of cues, individual cues (local and distal), vestibular cues, PI estimates, and so on, the same contingencies will be reflected in DG and CA3 place fields, which might explain the heterogeneous and complex changes seen in place representations in pathologically inconsistent environments (Bures et al., 1997; Rotenberg & Muller, 1997; Knierim et al., 1998).

An important aspect of the hippocampal system that is not addressed in this work is the function of the CA3 recurrent connections. We essentially agree with previous proposals (Marr, 1971; Treves & Rolls, 1992; O'Reilly & McClelland, 1994; Hasselmo et al., 1994, 1996; Blum & Abbott, 1996; Rolls, 1996; Redish & Touretzky, 1998) that the CA3 functions as an autoassociator performing the dual roles of pattern completion (to produce correct place codes from degraded afferent information) and sequence learning (Minai & Levy, 1993; Levy, 1996; Blum & Abbott, 1996; Skaggs & McNaughton, 1996; Redish & Touretzky, 1998; Sohal & Hasselmo, 1998). It has been suggested (Hasselmo & Schnell, 1994; Redish & Touretzky, 1998) that the recurrent CA3 connections are used only during associative recall and are blocked by cholinergic (and perhaps even GABAergic) modulation during learning (Hasselmo & Schnell, 1994; Hasselmo et al., 1995, 1996). Our current model focuses on only the response of the CA3 place cells to afferent sensory information. Once a PI system is incorporated into the model, the inclusion of CA3 recurrent connections will become more relevant since pattern completion would be needed to resolve conflicts between sensory and PI information (Redish & Touretzky, 1998).

Another issue closely related with the role of CA3 is the 5–12 Hz theta rhythm of the hippocampal EEG (Vanderwolf, 1969). The discrete time steps in our model can be seen as corresponding to single theta cycles. We do not include any explicit provision for theta in our system, but our assumption of updating the afferent stimulus once every theta cycle is consistent with evidence that rats sample sensory information at the theta frequency (Greenstein, Pavlides, & Winson 1988). Long-term potentiation (LTP), which is probably the basis for neural learning, is also optimized when stimuli are spaced one theta cycle apart (Larson, Wong, & Lynch, 1986). It has been suggested (Lisman & Idiart, 1995) that the phase of the theta cycle may index multiple patterns simultaneously active in a hippocampal associative memory. An alternative possibility is that a single stored pattern may be retrieved recurrently via multiple high-frequency (gamma) oscillations nested in each theta cycle. Such a process could be very useful in eliciting previously stored CA3 place codes in response to partially correct input based on sensory stimuli and path integration (Redish and Touretzky, 1998). We do not include this in our current model but envisage doing so in the future.

Appendix

The following parameter values were used in the simulations reported.

M	Size of environment	20
N_{EC}	Number of neurons in layer EC	200
N_{DG}	Number of neurons in layer DG	1000
N_H	Number of neurons in layer H	500

N_{CA3}	Number of neurons in layer CA3	300
m (LA system only)	Number of cell groups	10
n_{DG} (LA system only)	Number of cells in each DG group	100
n_H (LA system only)	Number of cells in each H group	50
C_{EC-DG}	EC → DG connectivity	0.05
C_{EC-CA3}	EC → CA3 connectivity	0.07
C_{DG-CA3}	DG → CA3 connectivity	0.003
C_{DG-H}	DG → H connectivity	0.6
C_{H-DG}	H → DG connectivity	0.6
	EC → DG weights	U[0, 1]
	EC → CA3 weights	U[0.01, 0.1]
	DG → CA3 weights	U[0.4, 0.6]
h_{DG-H}	Within-group DG → H weight	1.0
l_{DG-H}	Cross-group DG → H weight	0.01
h_{H-DG}	Within-group H → DG weight	1.0
l_{H-DG}	Cross-group H → DG weight	0.01
a_i	EC place field width parameter	U[0.004, 0.006]
b_i	EC place field width parameter	U[0.004, 0.006]
c_i^u, c_i^v	EC place field center locations	U[-9, 29]
d_i	EC place fields orientation parameter	U[-1.0, 1.0]
σ_q^2	EC place field positional noise variance	1.0
s^2	EC firing rate noise variance	0.01
κ	EC baseline firing rate	0.0
g_{EC-DG}	EC → DG excitatory gain	$[0 \cdots 7]$
g_{H-DG}	H → DG excitatory gain	0.5
G_{H-DG}	H → DG inhibitory gain	0.2
K_{DG}	Nominal number of active DG neurons	40
r_1^{DG}	Probability of firing in the K_{DG} most active DG neurons	0.95
r_2^{DG}	Probability of firing in the next K_{DG} most active DG neurons	0.05
r_3^{DG}	Probability of firing in the remaining $N_{DG} - 2K_{DG}$ DG neurons	0.003
g_{DG-H}	DG → H excitatory gain	1.0
K_H	Nominal number of active H neurons	20
r_1^H	Probability of firing in the K_H most active H neurons	0.95
r_2^H	Probability of firing in the next K_H most active H neurons	0.05
r_3^H	Probability of firing in the remaining $N_H - 2K_H$ H neurons	0.003
g_{EC-CA3}	EC → CA3 excitatory gain	1.0

g_{DG-CA3}	DG → CA3 excitatory gain	1.0
G_{DG-CA3}	DG → CA3 inhibitory gain	0.01
K_{CA3}	Nominal number of active CA3 neurons	15
r_1^{CA3}	Probability of firing in the K_{CA3} most active CA3 neurons	0.95
r_2^{CA3}	Probability of firing in the next K_{CA3} most active CA3 neurons	0.05
r_3^{CA3}	Probability of firing in the remaining $N_{CA3} - 2K_{CA3}$ CA3 neurons	0.003
α	Minimum place field firing rate	r_3^{CA3}
ω	Maximum place field firing rate	r_1^{CA3}

Acknowledgments

This research was supported by grants no. IBN-9634424 (A.A.M. and P.J.B.), IBN-9808664 (A.A.M.), and IBN-9816612 (P.J.B.) from the National Science Foundation. We thank John Lisman, David Redish, Alexei Samsonovich, and Dave Touretzky for providing preprints. Discussions with many people have contributed to this work, but A.A.M. especially acknowledges André Fenton, Alex Guazzelli, Mike Hasselmo, Jaideep Kapur, Chip Levy, John Lisman, Mayank Mehta, David Redish, Alexei Samsonovich, Lokendra Shastri, Dave Touretzky, Aaron White, and Brad Wyble. Detailed feedback from three anonymous reviewers significantly enhanced the quality of this report.

References

Amaral, D. G., Ishizuka, N., & Claiborne, B. (1990). Neurons, numbers and the hippocampal network. *Progress in Brain Research, 83*, 1–11.

Amaral, D. G., & Witter, M. P. (1989). The three dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience, 31*, 571–591.

Andersen, P., Bliss, T. V. P., & Skrede, K. K. (1971). Lamellar organization of hippocampal excitatory pathways. *Exp. Brain Res., 13*, 222–238.

Barnes, C. A., Suster, M. S., Shen, J., & McNaughton, B. L. (1997). Multistability of cognitive maps in the hippocampus of old rats. *Nature, 388*, 272–275.

Barnes, C. A., McNaughton, B. L., Mizumori, S. J. Y., Leonard, B. W., & Lin, L-H. (1990). Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in Brain Research, 87*, 287–300.

Battaglia, F. P., & Treves, A. (1998). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Phys. Rev. E, 58*, 7738–7753.

Blum, K. I., & Abbott, L. F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation, 8*, 85–93.

Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus, 1*, 193–206.

Brown, E. N., Frank, L. M., Tang, T., Quirk, M. C., & Wilson, M. A. (1998). A statistical paradigm for neural spike-train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J. Neurosci, 18*, 7411–7425.

Buckmaster, P. S., & Schwartzkroin, P. A. (1994). Hippocampal mossy cell function: A speculative view. *Hippocampus, 4*, 393–402.

Buckmaster, P. S., & Schwartzkroin, P. A. (1995). Interneurons and inhibition in the dentate gyrus of the rat *in vivo*. *J. Neurosci., 15*, 774–789.

Bures, J., Fenton, A. A., Kaminsky, Y., Rossier, J., Sacchetti, A., & Zinyuk, L. (1997). Dissociation of exteroceptive and ideothetic orientation cues: Effect on hippocampal place cells and place navigation. *Phil. Trans. R. Soc. Lond. B, 352*, 1515–1524.

Carpenter, G. A., & Grossberg, S. (1987a). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Comp. Vision, Graphics, and Im. Proc., 37*, 54–115.

Carpenter, G. A., & Grossberg, S. (1987b). ART 2: Self-organization of stable category recognition codes for analog input patterns. *Applied Optics, 26*, 4919–4930.

Deadwyler, S. A., West, M. O., & Robinson, J. H. (1981). Entorhinal and septal inputs differentially control sensory-evoked responses in the rat dentate gyrus. *Science, 211*, 1181–1183.

Doboli, S., Minai, A. A., & Best, P. J. (1999). A latent attractors model of context-selection in the dentate gyrus-hilus system. *Neurocomputing*, 26–27, 671–676.

Eichenbaum, H., Kuperstein, M., Fagan, A., & Nagode, J. (1987) Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *J. Neurosci., 7*, 716–732

Etienne, A. S., Maurer, R., & Sèguinot, V. (1996). Path integration in mammals and its interaction with visual landmarks. *J. Exper. Biol., 199*, 201–209.

Gibbs III, J. W., Shumate, M. D., & Coulter, D. A. (1997). Differential epilepsy-associated alterations in postsynaptic $GABA_A$ receptor function in dentate granule and CA1 neurons. *J. Neurophysiol., 77*, 1924–1938.

Golarai, G., & Sutula, T. P. (1996). Bilateral organization of parallel and serial pathways in the dentate gyrus demonstrated by current-source density analysis in the rat. *J. Neurophysiol., 75*, 329–342.

Gothard, K. M., Skaggs, W. E., Moore, K. E., & McNaughton, B. L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci., 16*, 823–835.

Gothard, K. L., Skaggs, W. E., & McNaughton, B. L. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: Interaction between path integration and environmental cues. *J. Neurosci., 16*, 8027–8040.

Greenstein, Y. J., Pavlides, C., & Winson, J. (1988). Long-term potentiation in the dentate gyrus is preferentially induced at theta rhythm periodicity. *Brain Res., 438*, 331–334.

Han, Z.-S., Buhl, E. H., Lörinczi, Z., & Somogyi, P. (1993). A high degree of spatial selectivity in the axonal and dendritic domains of physiologically identified

local-circuit neurons in the dentate gyrus of the rat hippocampus. *Europ. J. Neurosci., 5*, 395–410.

Hasselmo, M. E., & Schnell, E. (1994). Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: Computational modeling and brain slice physiology. *J. Neurosci., 14*, 3898–3914.

Hasselmo, M. E., Schnell, E., & Barkai, E. (1995). Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in hippocampal region CA3. *J. Neurosci., 15*, 5249–5262.

Hasselmo, M. E., Wyble, B. P., & Wallenstein, G. V. (1996). Encoding and retrieval of episodic memories: Role of cholinergic and GABAergic modulation in the hippocampus. *Hippocampus, 6*, 693–708.

Hetherington, P. A., Austin, K. B., & Shapiro, M. L. (1994). Ipsilateral associational pathway in the dentate gyrus: An excitatory feedback system that supports N-methyl-D-aspartate-dependent long-term potentiation. *Hippocampus, 4*, 422–438.

Hill, A. J. (1978). First occurrence of hippocampal spatial firing in a new environment. *Exp. Neurol., 62*, 282–297.

Hill, A. J., & Best, P. J. (1981). The effect of deafness and blindness on the spatial correlates of hippocampal unit activity in the rat. *Exper. Neurol., 74*, 204–207.

Jackson, M. B., & Scharfman, H. E. (1996). Positive feedback from hilar mossy cells to granule cells in the dentate gyrus revealed by voltage-sensitive dye and microelectrode recording. *J. Neurophysiol., 76*, 601–616.

Jung, M. W., & McNaughton, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus, 3*, 165–182.

Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1998). Interactions between ideothetic cues and external landmarks in the control of place cells and head-direction cells. *J. Neurophysiol., 80*, 425–446.

Knierim, J. J., & McNaughton, B. L. (1995). Differential effects of dentate gyrus lesions on pyramidal cell firing in 1- and 2-dimensional spatial tasks. *Soc. Neurosci. Abstr., 21*, 940.

Kosko, B. (1988). Bidirectional associative memories. *IEEE Trans. Syst. Man. Cyber., 18*, 49–60.

Kubie, J. L., & Muller, R. U. (1991). Multiple representations in the hippocampus. *Hippocampus, 1*, 240–242.

Larson, J., Wong, D., & Lynch, G. (1986). Patterned stimulation at the theta frequency is optimal for the induction of long-term potentiation. *Brain Res., 368*, 347–350.

Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus, 6*, 579–591.

Levy, W. B., & Wu, X. (1996). The relationship of local context cues to sequence length memory capacity. *Network: Computation in Neural Systems, 7*, 371–384.

Lisman, J. E. (1999). Relating hippocampal circuitry to function: The role of reciprocal dentate-CA3 interaction in the recall of sequences. *Neuron, 22*, 233–242.

Lisman, J. E., & Idiart, M. A. P. (1995). Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science, 267*, 1512–1515.

Markus, E. J., Qin, Y.-L., McNaughton, B. L., & Barnes, C. A. (1994). Place fields

are affected by behavioral and spatial constraints. *Cog. Neurosci. Soc. Abstr., 1*, 91.

Markus, E. J., Qin, Y.-L., Leonard, B., Skaggs, W. E., McNaughton, B. L., & Barnes, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci., 15*, 7079–7094.

Marr, D. (1971). Simple memory: A theory for archicortex. *Phil. Trans. R. Soc. Lond. B, 262*, 23–81.

McNaughton, B. L., Barnes, C. A., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely moving rats. *Experimental Brain Research, 52*, 41–49.

McNaughton, B. L., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J., Kudrimoti, H., Qin, Y., Skaggs, W. E., Suster, M., & Weaver, K. L. (1996). Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *J. Exp. Biol., 199*, 173–185.

McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosci., 10*, 408–415.

Mehta, M., Barnes, C. A., & McNaughton, B. L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. USA, 94*, 8918–8921

Miller, V. M., & Best, P. J. (1980). Spatial correlates of hippocampal unit activity are altered by lesions of the fornix and entorhinal cortex. *Brain Res., 194*, 311–323.

Minai, A. A. (1997). Control of CA3 place fields by the dentate gyrus: A neural network model. In: J. M. Bower (Ed.), *Computational neuroscience: Trends in Research 1997: Proceedings of the Fifth Computational Neuroscience Meeting* pp. 411–416.

Minai, A. A., Barrows, G. L., & Levy, W. B. (1994). Disambiguation of pattern sequences with recurrent networks. In *Proc. World Congress on Neural Networks, San Diego* (Vol. 4, pp 176–180).

Minai, A. A., & Best, P. J. (1998). Encoding spatial context: A hypothesis on the function of the dentate gyrus-hilus system. In *Proc. Int. Joint Conf. on Neural Networks, Anchorage* (pp 587–592).

Minai, A. A., & Levy, W. B. (1993). Sequence learning in a single trial. In *Proc. World Congress on Neural Networks, Portland* Vol. 2 pp 505–508).

Mizumori, S. J. Y., Ward, K. E., & Lavoie, A. M. (1992). Medial septal modulation of entorhinal single unit activity in anesthetized and freely moving rats. *Brain Res., 570*, 188–197.

Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature, 297*, 681–683.

Moser, E. I. (1996). Altered inhibition of dentate granule cells during spatial learning in an exploration task. *J. Neurosci., 16*, 1247–1259.

Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing properties of hippocampal place cells. *J. Neurosci., 7*, 1951–1968.

Muller, R. U., Kubie, J. L., & Ranck, J. B., Jr. (1987). Spatial firing patterns of

hippocampal complex-spike cells in a fixed environment. *J. Neurosci., 7*, 1935–1950.
Muller, R. U., Kubie, J. L., & Saypoff, R. (1991). The hippocampus as a cognitive graph (abridged version). *Hippocampus, 1*, 243–246.
Muller, R. U., & Stead, M. (1996). Hippocampal place cells connected by Hebbian synapses can solve spatial problems. *Hippocampus, 6*, 709–719.
Nadler, J. V., Perry, B. W., & Cotman, C. W. (1978). Intraventricular kainic acid preferentially destroys hippocampal pyramidal cells. *Nature, 271*, 676–677.
Nadler, J. V., Perry, B. W., and Cotman, C. W. (1980). Selective reinnervation of hippocampal area CA1 and the fascia dentata after destruction of CA3-CA4 afferents with kainic acid. *Brain Res., 182*, 1–9.
O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature, 381*, 425–428.
O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Res., 34*, 171–175.
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. *Exp. Br. Res., 68*, 1–27.
O'Reilly, R. C., & McClelland, J. L (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus, 4*, 661–682.
Quirk, G. J., Muller, R. U., & Kubie, J. L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J. Neurosci., 10*, 2008–2017.
Quirk, G. J., Muller, R. U., Kubie, J. L., & Ranck, J. B., Jr. (1992). The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *J. Neurosci., 12*, 1945–1963.
Redish, A. D. (1997). *Beyond the cognitive map: Contributions to a computational neuroscience theory of rodent navigation*. Unpublished doctoral dissertation, Carnegie Mellon University, Pittsburgh, PA.
Redish, A. D., & Touretzky, D. S. (1997). Cognitive maps beyond the hippocampus. *Hippocampus, 7*, 15–35.
Redish, A. D., & Touretzky, D. S. (1998). The role of the hippocampus in solving the Morris water maze. *Neural Computation, 10*, 73–111.
Redish, A. D., & Touretzky, D. S. (1999). Separating hippocampal maps. In N. Burgess, K. Jeffery, & J. O'Keefe (Eds.), *Spatial functions of the hippocampal formation and the parietal cortex*. New York: Oxford University Press.
Rolls, E. T. (1996). A theory of hippocampal function in memory. *Hippocampus, 6*, 601–620.
Rotenberg, A., & Muller, R. U. (1997). Variable place-cell coupling to a continuously viewed stimulus: Evidence that the hippocampus acts as a perceptual system. *Phil. Trans. R. Soc. Lond. B, 352*, 1505–1513.
Samsonovich, A. (1999). *Theory of "partial remapping" in the hippocampal representation of space*. Unpublished manuscript, University of Arizona.
Samsonovich, A., & McNaughton, B. L. (1997). Path integration and cognitive

mapping in a continuous attractor neural network model. *J. Neurosci., 17*, 5900–5920.

Scharfman, H. E. (1991). Dentate hilar cells with dendrites in the molecular layer have lower thresholds for synaptic activation by perforant path than granule cells. *J. Neurosci., 11*, 1660–1673.

Scharfman, H. E. (1995). Electrophysiological evidence that dentate hilar mossy cells are excitatory and innervate both granule cells and interneurons. *J. Neurophysiol., 74*, 179–194.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Psychiatr., 20*, 11–21.

Shapiro, M. L., Tanila, H., & Eichenbaum, H. (1997). Cues that hippocampal place cells encode: Dynamic and hierarchical representation of local and distal stimuli. *Hippocampus, 7*, 624–642

Sharp, P. E. (1997). Subicular cells generate similar spatial firing patterns in two geometrically and visually distinctive environments: Comparison with hippocampal place cells. *Behav. Brain Res., 85*, 71–92.

Sharp, P. E., Kubie, J. L., & Muller, R. U. (1990). Firing properties of hippocampal neurons in a visually symmetrical environment: Contributions of multiple sensory cues and mnemonic processes. *J. Neurosci., 10*, 3093–3105.

Shen, B., & McNaughton, B. L. (1996). Modeling the spontaneous reactivation of experience-specific hippocampal cell assemblies during sleep. *Hippocampus, 6*, 685–693.

Shen, J., McNaughton, B. L., & Barnes, C. A. (1997). Reactivation of neuronal ensembles in hippocampal dentate gyrus during sleep following spatial experience. *Soc. Neurosci. Abstr., 23*, 506.

Skaggs, W. E., & McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science, 271*, 1870–1873.

Skaggs, W. E., & McNaughton, B. L. (1998). Spatial firing properties of hipocampal CA1 populations in an environment containing two visually identical regions. *J. Neurosci., 18*, 8455–8466.

Sloviter, R. S. (1987). Decreased hippocampal inhibition and a selective loss of interneurons in experimental epilepsy. *Science, 235*, 73–76.

Sloviter, R. S. (1994). The functional organization of the hippocampal dentate gyrus and its relevance to the pathogenesis of temporal lobe epilepsy. *Ann. Neurol., 35*, 640–654.

Sloviter, R. S., & Brisman, J. L. (1995). Lateral inhibition and granule cell synchrony in the rat hippocampal dentate gyrus. *J. Neurosci., 15*, 811–820.

Sohal, V. S., & Hasselmo, M. E. (1998). $GABA_B$ modulation improves sequence disambiguation in computational models of hippocampal region CA3. *Hippocampus, 8*, 171–193.

Soltesz, I., Bourassa, J., & Deschenes, M. (1993). The behavior of mossy cells of the rat dentate gyrus during theta oscillations in vivo. *Neuroscience, 57*, 555–564.

Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psych. Rev., 99*, 195–231.

Squire, L. R., & Zola-Morgan, S. (1988). Memory: Brain systems and behavior.

Trends in Neurosci., 11, 170–175.
Struble, R. G., Desmond, N. L., & Levy, W. B. (1978). Anatomical evidence for interlamellar inhibition in the fascia dentata. *Brain Res., 152*, 580–585.
Tanila, H., Shapiro, M. L., & Eichenbaum, H. (1997). Discordance of spatial representation in ensembles of hippocampal place cells. *Hippocampus, 7*, 613–623.
Thompson, L. T., & Best, P. J. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. *J. Neurosci., 9*, 2382–2390.
Thompson, L. T., & Best, P. J. (1990). Long term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Res., 509*, 299–318.
Touretzky, D. S., & Redish, A. D. (1996). Theory of rodent navigation based on interacting representations of space. *Hippocampus, 6*, 247–270.
Treves, A., & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus, 2*, 189–200.
Vanderwolf, C. H. (1969). Hippocampal electrical activity and voluntary movement in the rat. *EEG Clin. Neurophysiol., 26*, 407–418.
Willshaw, D. J., Buneman, O. P., & Longuet-Higgins, H. C. (1969). Non-holographic associative memory. *Nature, 222*, 960–962.
Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science, 261*, 1055–1058.
Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science, 265*, 676–679.
Wu, X., Baxter, R. A., & Levy, W. B. (1995). Context codes and the effect of noisy learning on a simplified hippocampal CA3 model. *Biol. Cybern., 74*, 159–165.
Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *J. Neurophysiol., 79*, 1017–1044.

Received September 11, 1998; accepted May 5, 1999.