Communicated by Mikhail Tsodyks

The Role of the Hippocampus in Solving the Morris Water Maze

A. David Redish
David S. Touretzky
Computer Science Department and Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213-3891, U.S.A.

We suggest that the hippocampus plays two roles that allow rodents to solve the hidden-platform water maze: self-localization and route replay.

When an animal explores an environment such as the water maze, the combination of place fields and correlational (Hebbian) long-term potentiation produces a weight matrix in the CA3 recurrent collaterals such that cells with overlapping place fields are more strongly interconnected than cells with nonoverlapping fields. When combined with global inhibition, this forms an attractor with coherent representations of position as stable states. When biased by local view information, this allows the animal to determine its position relative to the goal when it returns to the environment. We call this *self-localization.*

When an animal traces specific routes within an environment, the weights in the CA3 recurrent collaterals become asymmetric. We show that this stores these routes in the recurrent collaterals. When primed with noise in the absence of sensory input, a coherent representation of position still forms in the CA3 population, but then that representation drifts, retracing a route.

We show that these two mechanisms can coexist and form a basis for memory consolidation, explaining the anterograde and limited retrograde amnesia seen following hippocampal lesions.

1 Amnesia Following Hippocampal Lesions

Hippocampal lesions in humans produce devastating impairments in declarative memory (memories of specific items, events, or episodes) (Scoville & Milner, 1957; Squire & Zola-Morgan, 1988; Cohen & Eichenbaum, 1993; Zola-Morgan & Squire, 1993). Although these patients perform immediate recall tasks normally, they are strongly impaired at times greater than a few minutes. In addition to the anterograde amnesia, these amnesias extend backward in time to recently before the occurrence of the lesion, but they leave early memories intact (Squire & Zola-Morgan, 1988). Similar results have been seen in nonhuman primates with hippocampal lesions (Squire & Zola-Morgan, 1988; Zola-Morgan & Squire, 1993).

The theory proposed to explain these data is that the hippocampus serves as a temporary store for memory (Marr, 1969; Buzsáki, 1989; Cohen & Eichenbaum, 1993; McClelland, McNaughton, & O'Reilly, 1995). However, no models of hippocampal function in specific memory tasks exist; all published models of declarative memory demonstrate storage and retrieval of arbitrary binary vectors (Marr, 1971; Rolls, 1989; Gluck & Myers, 1993; Alvarez & Squire, 1994; Hasselmo & Schnell, 1994; O'Reilly & McClelland, 1994; Levy, 1996; Rolls, 1996). Although these theories can address general principles involved in memory, they cannot address the role of the hippocampus in specific tasks. This makes it difficult to compare their results to real experiments or to generate testable predictions from them.

Anterograde and limited retrograde amnesias after hippocampal lesion are also seen in rats tested in the Morris water maze (Sutherland & Hoesing, 1993). The Morris water maze consists of a submerged platform placed somewhere within a pool of water made opaque with milk or chalk (Morris, Garrud, Rawlins, & O'Keefe, 1982). When placed in this pool, rats try to find a way out; they initially swim randomly until they find the platform and climb out. Normal rats quickly learn the location of the platform: if the platform is removed, the rats search at the place where the platform had been (Morris et al., 1982). Rats with hippocampal lesions cannot learn this task (Morris et al., 1982; Morris, Schenk, Tweedie, & Jarrard, 1990; McDonald & White, 1994). If the rats are trained on the task first and then given a hippocampal lesion 1 week later, they show profound deficits; however, the same lesion 12 weeks after training produces much smaller deficits (Sutherland & Hoesing, 1993). Here, then, is a specific amnesia result that can be modeled in detail.

2 Modeling Memory in the Morris Water Maze

We suggest that rats trained on the Morris water maze use two different mechanisms to find the hidden platform, one locale based and one route based. The two mechanisms can be subdivided into five steps, the first two of which are locale based and the last three route based. These five steps occur in order (we address possible ways to sidestep this in section 4.2):

1. *Exploration*. The animal familiarizes itself with the environment.
2. *Self-localization*. Upon reentry into the environment, the animal must determine its location relative to the platform. From this information, it can determine the direction it needs to swim in order to reach the platform.
3. *Route learning*. When the animal travels along a specific path, routes are stored in the recurrent connections of hippocampal area CA3.
4. *Replay of routes during sleep*. During sleep, the recent routes are replayed because, when primed with noise, the hippocampal formation settles

to a stable representation of a location, which then drifts along the routes stored in the CA3 recurrent connections. McNaughton, Skaggs, and Wilson (Wilson & McNaughton, 1994; Skaggs & McNaughton, 1996) have reported data supporting this hypothesis: simultaneous extracellular recordings from hippocampal pyramidal cells have shown that cells tend to fire in the same sequence during slow-wave sleep as they did during recent experience in an environment. We discuss this and its implications for the theory in section 2.4.

5. *Consolidation*. The "dreamed" routes are transferred to long-term storage by a slowly learning cortical network.

Anterograde amnesia occurs because long-term memory requires a hippocampus in order to learn the routes. Retrograde amnesia occurs when long-term memory has not been completely trained at the time the hippocampus is lesioned. Once the routes have been stored in long-term memory, the animal can solve the task without a hippocampus.

Our previous work laid out a theory of the role of the hippocampus in navigation (Touretzky & Redish, 1996; Redish & Touretzky, 1997) (see Figure 1). The key components of the expanded theory presented here are as follows:

- Path integration occurs via a loop including the superficial layers of the entorhinal cortex (ECs), the subiculum (Sub), and the parasubiculum (PaS).[1]
- Sensory cues (called the *local view*, but not solely visual) enter the hippocampal formation from high-level sensory association areas, referred to here as HLS.
- The path integration and local view representations are first combined in ECs, but any conflicts are resolved by the recurrent connections in CA3.
- On reentry into a familiar environment, competitive dynamics in the hippocampus allows the system to settle to a coherent place code even with ambiguous sensory cues.[2] This coherent code resets the path integrator so that multiple experiences of the same environment are compatible with each other.
- During sleep, recurrent connections within the hippocampus force a coherent code to form from noise, but due to asymmetric connection strengths produced during training, the represented location precesses

[1] Path integration is the ability to return to a starting point, even after a long, circuitous path, using only idiothetic cues (Mittelstaedt & Mittelstaedt, 1980; Gallistel, 1990; Etienne, 1992).

[2] The place code is coherent if all neural activities are consistent with a representation of the same location in space.

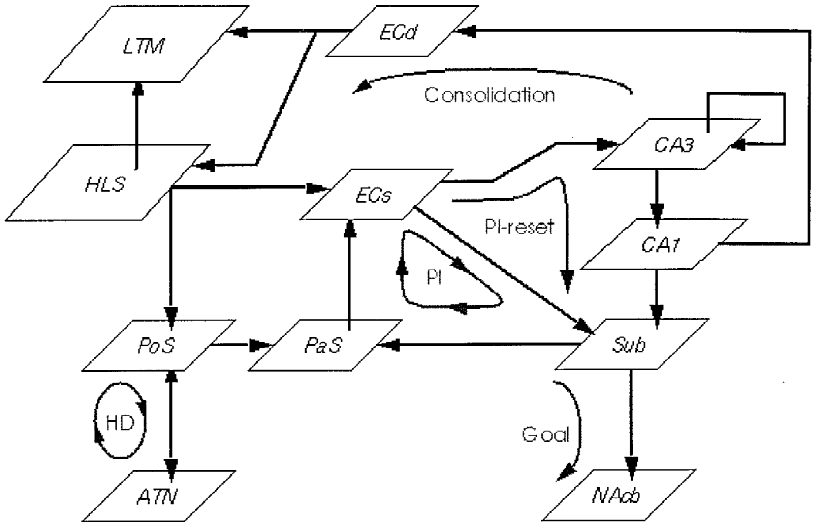

Figure 1: Sensory information (local view) enters the system by the high-level sensory areas, (HLS). Path integration (PI) occurs via a loop including the superficial layers of the entorhinal cortex (ECs), the subiculum (Sub), and the parasubiculum (PaS), updated by head direction information in the postsubiculum (PoS). Head direction (HD) is updated via a loop including the anterior thalamic nuclei (ATN). The hippocampus proper (CA3, CA1) serves to reset the path integrator (PI-reset) from local view information on reentry into a familiar environment. During sleep, routes are read out from hippocampus through the deep layers of entorhinal cortex (ECd) into long-term memory (LTM). The goal trajectory (Goal) is calculated in nucleus accumbens (NAcb) and requires information from the subiculum, transmitted via the fornix. Not all anatomical connections are shown, and some structures are omitted.

along the learned routes, effectively replaying routes traversed during the task.

- Slowly learning cortical networks (long term memory, or LTM) can be trained by these "dreams" of routes traveled so that eventually the animal can perform the task without a hippocampus.

The theory requires the hippocampus to show two major modes of operation: a storage mode and a recall mode. The hippocampus does show two modes of activity (Vanderwolf & Leung, 1983; Buzsáki, 1989; Stewart & Fox, 1990; Vanderwolf, 1990), seemingly gated by the presence of acetylcholine (ACh) (Vanderwolf, 1990; Huerta & Lisman, 1993; Hasselmo, 1995) arriving from septal input via the fimbria (Stewart & Fox, 1990) and referred to by the very different electroencephalogram (EEG) traces recorded during the

two modes. During motion, in the presence of ACh, the hippocampal EEG shows a 7–12 Hz rhythm called theta; during rest and slow-wave sleep, in the absence of ACh, the hippocampal EEG shows irregular activity, called LIA (large-amplitude irregular activity), characterized by short-duration sharp waves (Vanderwolf, 1971).

We now proceed to detail how our theory accounts for each of the five steps discussed above, reviewing the data supporting the theory and simulation results demonstrating each point. Additional simulation details are given in appendix A.

2.1 Exploration: Learning the Cognitive Graph. We begin by showing that the combination of random exploration and correlational (Hebbian) learning produces a weight matrix in the CA3 recurrent connections that is appropriate for the competitive dynamics needed for self-localization. Following Muller, Kubie, and Saypoff (1991), we call this connection function the cognitive graph: the synaptic efficacies between place cells are inversely related to the distance between the centers of their place fields.

As the animal wanders around the environment, cells with nearby place fields are more likely to be coactive than cells with well-separated fields. Combined with correlational long-term potentiation (LTP), in which the synaptic efficacy is increased when both cells are simultaneously active, after a session of wandering around an environment, the intrahippocampal connections will be inversely proportional to the distance between place field centers (Muller et al., 1991). LTP has been shown extensively in the recurrent connections in CA3 and in the Schaffer collaterals connecting CA3 to CA1 (see Landfield & Deadwyler, 1988).

Also supporting this theory are data showing the effect of ACh: while suppressing neuronal transmission in intrahippocampal synapses, ACh enhances LTP in them (Hasselmo & Schnell, 1994). We make the simplifying assumption that ACh shuts off the CA3 recurrent connections completely. Experiments in hippocampal slices show that it diminishes synaptic efficacy across the Schaffer collaterals by approximately 90 percent, while diminishing the efficacy of the perforant path (inputs from ECs) by only approximately 45 percent (Hasselmo & Schnell, 1994). ACh presumably is present during theta mode, while the animal is moving about the environment. Disruption of ACh has been found to shift the hippocampus out of theta, while cholinergic agonists shift the hippocampus into theta mode (Huerta & Lisman, 1993). LTP produced by hippocampal stimulation during theta or at intervals corresponding to the theta frequency is much stronger than similar stimulation during nontheta (LIA) (Larson, Wong, & Lynch, 1986; Larson & Lynch, 1989; Greenstein, Pavlides, & Winson, 1988; Pavlides, Greenstein, Grudman, & Winson, 1988; Huerta & Lisman, 1993).

Simulations. The network used to demonstrate the generation of the cognitive graph consisted of a limited version of the total model presented

in Figure 1. It included HLS, PI, HD, ECs, and CA3/1 components. Specific neuronal model details are given in appendix A.

The HD component consisted of a 1D (circular) neural array. At every time step, the currently represented value was calculated by a directional mean, the represented value was updated by the angular velocity, and then a new (idealized) HD representation was generated. This allowed us to simulate the neural interactions between the head direction and other representations without the computational time required to simulate the actual head direction update mechanism. We have previously shown that a postsubiculum-anterior thalamic nuclei head direction model can track real angular velocity sequences accurately (Redish, Elga, & Touretzky, 1996). The PI simulation was similar but used a two-dimensional (2D) (toroidal) neural sheet.

We also simulated the HLS component as a bump on a 2D neural sheet (as in Samsonovich & McNaughton, in press) and assumed that at every point in the environment, the position of the animal was correctly represented by the population. We did this because it is not clear what aspects of the environment are included in the local view; any set of spatial information about landmarks sufficient to localize the animal to a point will do. For some experiments, such as those described in section 2.2, there was more than one peak in the local view representation. This allowed us to ask questions about the ability of the system to handle ambiguous inputs without having to build a complicated, speculative model of the visual system of the rodent.

The EC's population was also simulated as a 2D neural sheet receiving input from the HLS and PI components.

Because we do not separately simulate CA3 and CA1, we refer to the combined simulated population as CA3/1 but refer to CA3 and CA1 in our discussions of the theory proper. The CA3/1 population consisted of a 2D neural sheet connected to the PI representation by one-to-one connections. A sparse random connection pattern works just as well, but by using a one-to-one pattern, we know the center of the place field for each CA3 neuron.

According to our theory, every time an animal enters an environment, it self-localizes as best it can by a sharp wave (LIA mode). We do not measure EEG in our simulations, but the self-localization sequence begins with a high proportion of the CA3 cells active at low levels, and settles to a small number of highly active cells within approximately 100 ms (see section 2.2). We believe this corresponds to a single sharp wave (Buzsáki, 1989).

We thus begin exploration by first placing the simulated animal at a random location and triggering a 100-ms sharp wave. Since this is a novel environment, there are no stored associations in the intrahippocampal connections, and the sharp wave serves only to reset the path integrator to a random point. This random location becomes the reference point for the environment, and the origin for the path integrator coordinate system. The animal then explores the environment by wandering randomly, during which LTP occurs in the intrahippocampal connections.

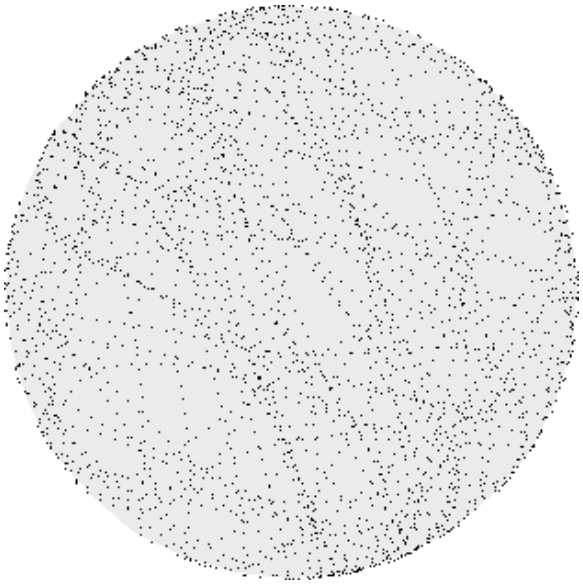

Figure 2: Route traveled by simulated rodent while exploring a 100-cm diameter circular environment for 5 minutes. Dots indicate position sampled every 10 seconds. Gray area denotes arena.

Two effects must occur for the animal to have familiarized itself with the environment. First, a mapping from local views (in HLS) to path integrator coordinates (in ECs) must be learned. Second, because local views may be ambiguous, the cognitive graph must be learned in the recurrent connections of CA3.

We show that the appropriate connection function appears within the first minutes of exploration. Figure 2 shows the track of the simulated animal wandering over the environment during the 5 minutes of exploration. The animal has clearly covered the entire environment. Figure 3 shows a scatter plot of the learned synaptic weights as a function of the distance between each pair of units in the CA3/1 population. The synaptic efficacy between two cells is, on average, inversely related to the distance between the centers of their place fields. A similar plot of HLS-to-ECs synapse strengths would show that local view representations (in HLS) have been associated with path integrator coordinates (in ECs).

Although the connection function appears quickly in our simulations, we used very large learning rates in order to minimize computation time. We do not know whether more realistic learning rates would allow the function to be acquired so quickly. If they did, the random trajectories shown by

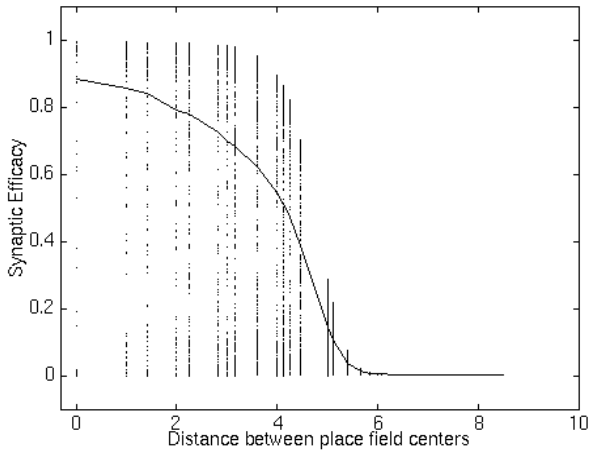

Figure 3: Scatter plot of learned synaptic weights as a function of distance between pairs of units. Distance is in cm. Line indicates the mean.

animals first placed in the water maze (with no knowledge of the goal location) would be sufficient to "explore" the environment.

2.2 Self-Localization and Place Field Stability. In order to navigate within a familiar environment, an animal must use a consistent representation of position from session to session. Although visual cues can serve to inform the animal of its initial position, if they are ambiguous, there must be a mechanism to settle on a consistent representation of location. We believe intrinsic competitive dynamics force the hippocampus to settle on a coherent code. These dynamics can serve as a disambiguation mechanism and can reproduce the search pattern that gerbils make when faced with ambiguous local cues (Touretzky & Redish, 1996).

We suggest that the competitive dynamics realized in the rodent proceeds thusly: subiculum, parasubiculum, hipppocampus, and entorhinal cortex are initially noisy; sensory cues in HLS passed through ECs into the hippocampus proper bias the random firing rates with candidate locations. The recurrent connections in CA3 allow one of these candidate locations to win out, forming a coherent place code in hippocampus. The connections between CA1 and subiculum reset the path integrator to the correct representation of the animal's location in path integrator coordinates. This happens in the course of a single sharp wave during LIA. In our simulations, the place code in CA3/1 is coherent within 50 to 100 ms. Figure 4 shows the first 70 ms of a simulated sharp wave.

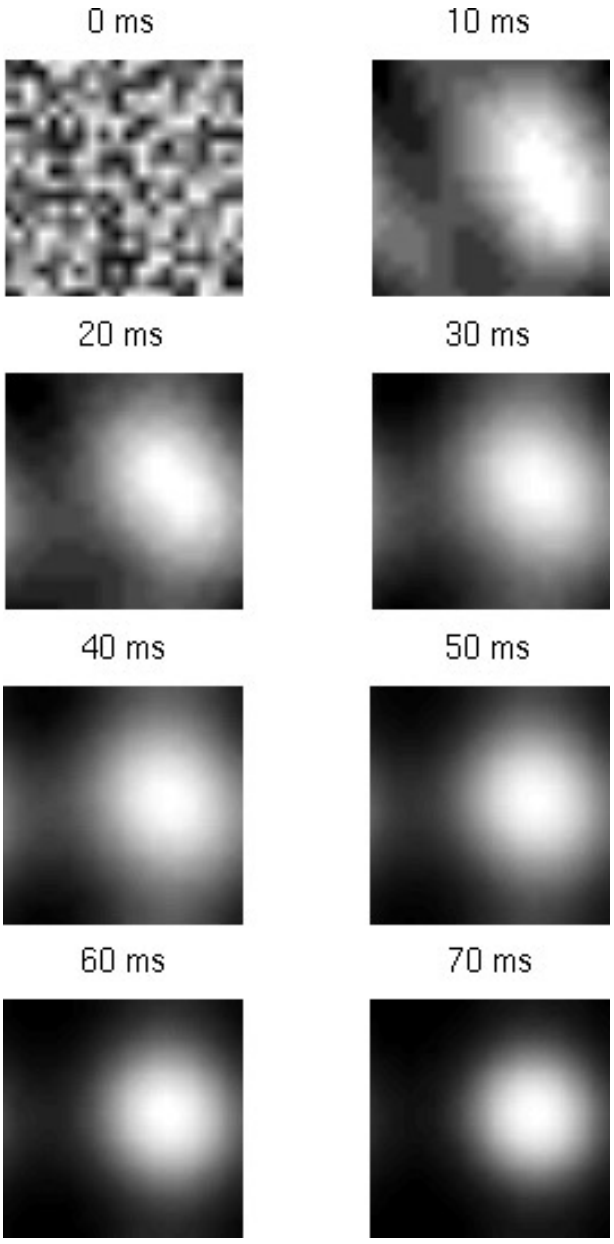

Figure 4: Starting from random noise, a coherent place code forms in less than 50 ms. Plot shows firing rates of CA3/1 place cells. Cells are laid out in a 2D sheet with their locations in the sheet corresponding to the centers of their place fields in the environment. Intensity values have been interpolated for clarity. White indicates high firing rate, black low.

During a sharpwave, place cells do not show normal place fields; many cells are simultaneously active (many more than would normally be active during theta) (Buzsáki, 1989). Because ACh is not present, synaptic efficacy between CA3 cells is presumably at full strength, allowing the system to settle from an initially noisy state to a coherent representation of the animal's location. Once this representation is coherent, the path integrator (in subiculum, receiving strong connections from CA1) is driven by the now coherent representation of location in CA1 and is effectively reset. The animal can now navigate around the environment.

Simulations. The network used to demonstrate self-localization used a similar architecture to that set out in section 2.1. The 2D neural sheets were enlarged to 20×20, and the hippocampal simulation was more detailed. We simulated the CA3/1 population as two pools, one excitatory and one inhibitory (labeled CAE and CAI, respectively, in appendix A) . The excitatory neurons were interconnected within and between pools by an idealization of the connection function learned in section 2.1 (a gaussian with a standard deviation of 20 cm). We had to use an idealization because our networks are small relative to those in the actual rodent brain. Inhibitory CA3/1 neurons were broadly connected to both the excitatory and inhibitory pools. Essentially, this connection structure corresponds to local excitation and global inhibition.

We measured the ability of this self-localization process to handle ambiguities in the local view by locking three bumps into the HLS representation. This simulates three "candidate positions" in the local view. This ambiguous local view representation is resolved in the CA3/1 representation into a coherent representation of position similar to that shown in Figure 4.

2.3 Route Learning. Given a representation of the animal's current location in the environment and a representation of the current goal, the animal should be able to calculate the direction to take to reach the goal. The nucleus accumbens receives information about current needs and desires from the amygdala (Davis, Rainnie, & Cassell, 1994) and information about current location via the fornix (Witter, Ostendorf, & Groenwegen, 1990) and is optimally situated to perform this function. This function of the nucleus accumbens was first suggested by Mogenson (1984), and a model showing its feasibility has been presented by Brown and Sharp (1995). (See Redish and Touretzky (1997) for a review.)

There are three neurophysiological effects that allow the hippocampus to store routes as the animal travels. First, the postsynaptic potential (PSP) has a nonzero time constant. As an animal travels from the place field of one neuron (say, *a*) to another (say, that of neuron *b*), neuron *a* continues to have an effect on the firing rate of neuron *b,* but when the animal was in place field *a*, neuron *b* did not have an effect on neuron *a.*

Second, imagine the animal at an instant along the route taken, passing

through a place field centered slightly off the route. This cell will have a firing rate somewhere between its maximum and minimum firing rates. Cells with place fields closer to the animal's position will have higher firing rates, and cells with place fields farther will have lower rates. This means that the output connection function from the neuron in question will be biased asymmetrically toward the path traveled.

Finally, as the animal moves through the place field, the timing of the spikes fired by that cell precesses with respect to the theta rhythm: cells behind the animal fire early in the cycle, while cells ahead of the animal fire late in the cycle (O'Keefe & Recce, 1993; Skaggs, 1995; Skaggs & McNaughton, 1996; Skaggs, McNaughton, Wilson, & Barnes, 1996). Thus the represented position sweeps across the actual position from back to front with each theta cycle. When combined with the biophysical time course of LTP, this phase precession will also favor connections pointing along routes to a goal (Blum & Abbott, 1996).

Simulations. The route-learning simulation consisted of the same network as used in section 2.2, with the addition of a new hippocampal mode. The simulation parameters as described in section 2.2 correspond to LIA mode, while the simulation parameters used for the route-learning simulation correspond to theta mode (see appendix A).

We do not explicitly model the nucleus accumbens. Instead we compare the subicular representation at the goal and the current subicular represenation, and then simulate travel in a straight line until the animal reaches either the goal or a wall. Figure 5 shows the paths traveled to reach the goal from the four cardinal points. These are the four routes that will be stored in the CA3/1 population.

We model the asymmetric nature of LTP by making the learning rule dependent on the synaptic drive of the presynaptic neuron and the firing rate of the postsynaptic neuron (see equation A.4). The synaptic drive S_i of neuron i is the effect of that neuron on all the neurons on which it synapses divided by the synaptic weight over each synapse (Pinto, Brumberg, Simons, & Ermentrout, 1996; see appendix A). It can be understood as a model of the postsynaptic potential or as a decaying memory of recent firing rates shown by neuron i, with a decay time constant of τ_i.

We do not model phase precession as an emergent result of a complex process; instead we assume that phase precession exists and show that, when combined with the asymmetric temporal nature of LTP, routes are stored in the recurrent connections of the hippocampus. In order to produce phase precession, we derive the preferred phase of each CA3 neuron using the approximation in Figure 6. We then define the firing rate of each neuron at time t as

$$F_i(t) = e^{-(\tilde{\theta}_i(t)-\theta(t))^2/\rho^2} \cdot \hat{F}_i(t), \tag{2.1}$$

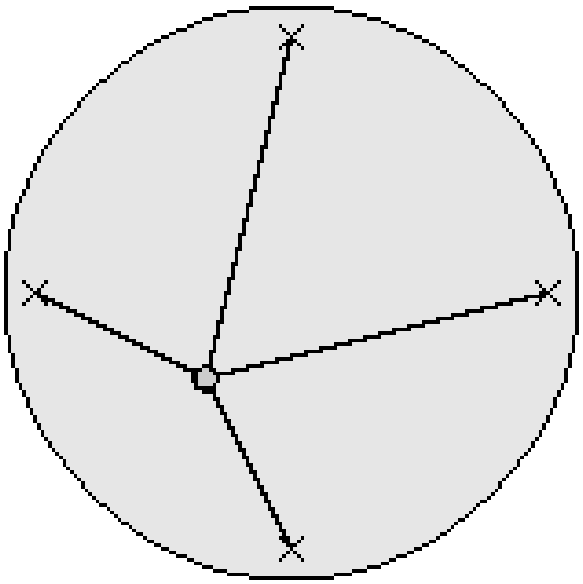

Figure 5: Four routes to the goal location. In order to demonstrate the accuracy of the simulation, the direction to the goal was determined by comparing the representation in subiculum with the prior subicular representations of the goal location. Lines indicate trajectories taken by the simulated animal to reach the goal (indicated by small circle). An x has been drawn at the initial location of the animal in each position. These routes are stored in the CA3/1 structure via LTP. Gray area denotes arena.

where $\tilde{\theta}_i(t)$ is the preferred phase of neuron i, $\theta(t)$ is the current phase of the theta rhythm, ρ is a constant, and $\hat{F}_i(t)$ is the peak firing rate at $\theta(t) = \tilde{\theta}(t)$. We assume a theta rhythm with a frequency of 7 Hz, so $\theta(t) = (\frac{7 \cdot 360^\circ}{sec} \cdot t) \bmod 360^\circ$. $\hat{F}_i(t)$ is determined by equation A.2 (see appendix A).

This makes the representation of position sweep from behind the animal to in front of it with each theta cycle as it does in the real animal (O'Keefe & Recce, 1993; Skaggs et al., 1996). We do not claim this as a model of how phase precession is actually generated in the rodent hippocampus, only that it produces a phase precession effect so that routes can be stored in the CA3 recurrent connections.

These effects combine to store routes in the recurrent connections of CA3. They produce a vector field pointing toward the path and then leading to the goal. Figure 7 shows the routes stored by an animal traversing the four paths in Figure 5.

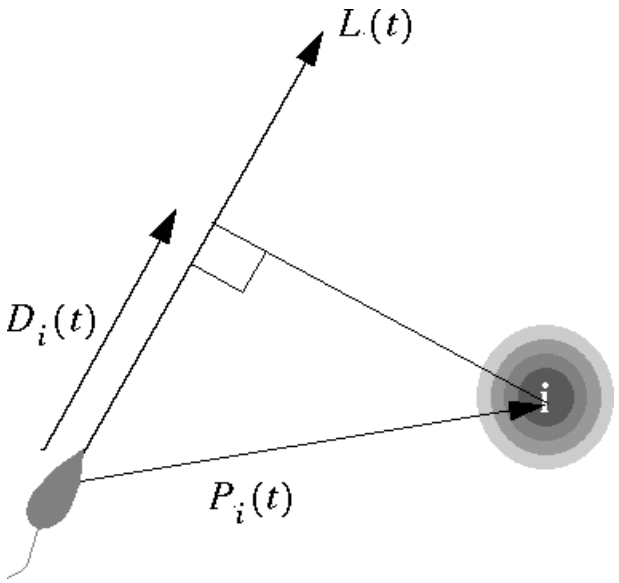

Figure 6: How we simulated phase precession. Let $L(t)$ be a ray originating at the simulated rodent's current position (as represented by the pyramidal cells in CA3), pointing in the direction of its current heading (as represented by the cells in postsubiculum). Let $P_i(t)$ be a vector from the represented position of the rodent to the center of the place field of place cell i, and $D_i(t)$ be the projection of $P_i(t)$ onto $L(t)$. Then the preferred phase of neuron i, $\tilde{\theta}_i(t)$ is proportional to $D_i(t)$: $\tilde{\theta}_i(t) = K \cdot D_i(t)$, where K is a scale factor chosen to be small enough that the phase precession will not wrap around ($K = 1.2$ deg/cm in our simulations). Thus, cells with place fields behind the represented position (in CA3) fire earlier in the theta cycle, and cells ahead of the represented position fire later. We do not claim this as a model of how phase precession is actually generated in the rodent brain, only that it produces a phase precession effect so that routes can be stored in the CA3 recurrent connections.

2.4 Replay of Routes During Sleep. When there is sensory input into the hippocampus and the hippocampus is in LIA mode (i.e., the animal is awake and looking around but not moving), sensory cues enter the system via HLS and ECs, and those CA3 cells that are consistent with the current local view will be more active than those that are not. This biases CA3 to settle to a place code consistent with the local view and thus can serve as a self-localization procedure.

On the other hand, when there is no sensory input, this bias will be absent, but due to the recurrent connections in CA3, the hippocampus will still settle into a coherent activity pattern. Due to the asymmetric connections that were stored when the animal traversed the routes to the goal, the place

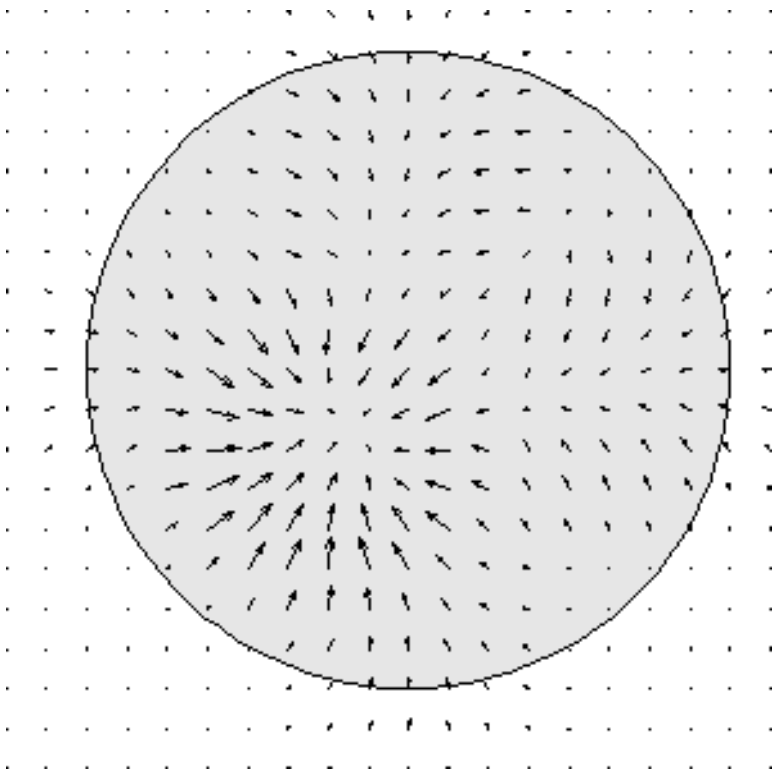

Figure 7: Vector field of routes to a goal stored in the recurrent connections of the model CA3. For each cell *j*, we calculated the center of mass of the output connection weights, and plotted an arrow from the place field center toward the center of mass. Length of arrow corresponds to linearly scaled distance to center of mass of the output connection weights.

code will precess along one of these remembered routes. The bias provided by the sensory input should be enough to keep the system from precessing when awake, but in the absence of sensory input (during sleep), there is nothing to stop the precession.

During sleep, when sharp waves occur without sensory input, we expect to see replay of routes. This is shown in Figure 8. Given an initial noisy state, a coherent code forms within half a second, and then over the next few seconds, it drifts along a remembered route.

Data supporting a replay of recent experience in hippocampus during sleep were first seen by Pavlides and Winson (1989). They showed that cells with recently visited place fields were more active during REM sleep than other cells whose place fields had not been recently visited. Wilson and

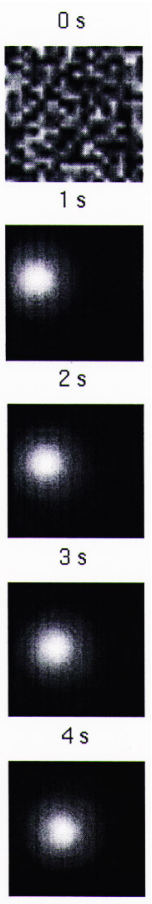

Figure 8: Replay of routes during LIA without sensory input. A coherent code forms quickly and then slowly drifts to the goal over the subsequent few seconds.

McNaughton (1994) showed that during slow-wave sleep (SWS), cells that showed correlated firing during a session in an environment (because their place fields overlapped) also showed a stronger correlation during sleep immediately after the session.

Skaggs and McNaughton (1996) explicitly examined the temporal nature of replay during sharp waves in slow-wave sleep. They defined the temporal bias B_{ij} between two cells i and j to be the difference between the integrated cross-correlation for the 200 ms after each spike of cell j and the integrated cross-correlation for the 200 ms before each spike of cell j. Thus, if cell i generally fires after cell j rather than before, B_{ij} will be greater than 0.

They report that the temporal bias during sleep after running on a linear track is strongly correlated with the temporal bias seen while the animal was running on the track. Although this does not conclusively demonstrate replay of routes during sleep, it provides strong circumstantial evidence for the hypothesis.

The original indication that REM sleep affected learning came from REM sleep deprivation studies (Hennevin, Hars, Maho, & Bloch, 1995). Because REM sleep is so distinctive, it is easy to wake a subject up whenever it goes into REM sleep. There is a window during which animals must be able to get REM sleep in order to learn a task such as the water maze (Smith, 1995). However, this REM deprivation window is usually measured in terms of a few hours, not a few weeks (Smith, 1995).

Although we have presented the model as if the replaying of routes occurs during a single sharp wave, the time course of the replay presented in Figure 8 is more compatible with a settling process that occurs during REM sleep: a single sharp wave lasts for 100 ms, while a REM sleep episode can last for a few minutes (Vanderwolf, 1990).

During sleep, rats first enter slow-wave sleep. LIA occurs throughout the hippocampus, and irregular slow activity occurs throughout the cortex. With the onset of REM sleep, all of the animal's muscles go limp, and theta appears in hippcampus, while low-voltage fast activity (LVFA) occurs in the cortex. This bout of REM sleep lasts for about 2 minutes (Vanderwolf, 1990)—plenty of time to retrace a route.

One possibility is that during the final sharp wave of the LIA block, the system settles on a coherent location and then throughout the following bout of REM sleep, the system retraces the route.

Another possibility is suggested by data from Smith (1995) and Kudrimoti et al. (1996). Smith (1995) reviews data that there is a post-REM NMDA-window in which NMDA blockers such as MK-801 will also disrupt learning. For example, animals with REM deprivation windows of 5 to 8 hours posttraining had NMDA windows of 9 to 12 hours posttraining, and animals with REM deprivation windows of 1 to 4 hours posttraining had NMDA windows of 5 to 16 hours posttraining.

This is particularly intriguing given recent data from Kudrimoti et al. (1996). They measured the increase in correlation of cells whose place fields overlapped during a recent exploration of an environment during three blocks of LIA in SWS. Blocks 1 and 2 occurred before the first bout of REM sleep, and block 3 occurred soon after it. They found that the correlation in block 1 was higher than the correlation in block 2, but that the correlation during block 3 was larger than block 2 and comparable to block 1. An intriguing possibility is that replay occurs during both SWS and REM. During SWS, information is written out to cortex and partially erased from the hippocampus, while during REM, the hippocampal information is strengthened so that it can be read out again during SWS. More work is clearly needed looking at the representation of location in hippocampus during sleep to explore

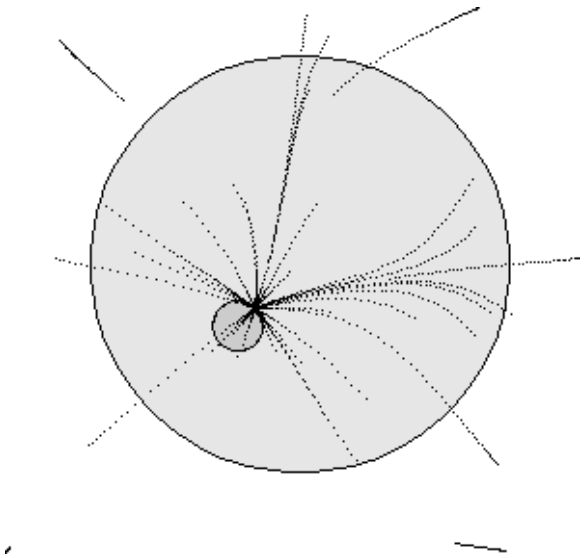

Figure 9: Believed position of the animal during 20 5-second dreams, sampled at 100-ms intervals. Each point indicates the position represented in the model CA3/1 at each sampled time.

whether the replay occurs during slow-wave or REM sleep and what the corresponding roles of slow-wave and REM sleep are.

Simulations. The route replay simulations used for Figures 8 and 9 were identical to the self-localization simulation of section 2.2, but the weight matrix used was the asymmetric one produced by LTP in section 2.3. In order to simulate "sleep," the HLS representation was set to zero.

2.5 Consolidation. If we hypothesize a cortical memory capable of storing routes that receives input from HLS and hippocampus and learns during slow-wave sleep, we can explain the anterograde and limited retrograde amnesia that follows hippocampal lesions applied before, soon after, or long after training in the Morris water maze (Sutherland & Hoesing, 1993).

Anterograde amnesia occurs because long-term memory requires a hippocampus in order to learn the routes. Retrograde amnesia occurs when long-term memory has not been completely trained at the time the hippocampus is lesioned. The retrograde amnesia reaches back only a limited time because once the routes have been stored in long-term memory, the animal can solve the task without a hippocampus.

Our theory suggests that the hippocampal cognitive map, necessary for learning the task with a hidden platform, is not written out to long-term

memory (LTM), but rather that LTM only stores the needed routes. Once the routes are stored in LTM, the representation maintained in the hippocampus is no longer required to perform the task successfully. In effect, the system learns the direct mapping from local views (HLS) to directions (LTM).

This hypothesis requires an output from hippocampus to cortex. There is such a pathway, leading from CA1 into the deep layers of entorhinal cortex (ECd), and then back into a variety of cortical areas (Amaral, 1993; Witter, 1993). Data from ischemic subjects show that damage restricted to CA1 is sufficient to produce the anterograde and limited retrograde amnesia discussed above (Squire & Zola-Morgan, 1988; Zola-Morgan & Squire, 1993; Wood, Mumby, Pinel, & Phillips, 1996).

This hypothesis also requires that ECs and ECd be active at different times: As the intrahippocampal connections are learning (during theta), activity should not be transmitted through ECd. But as routes are replayed during sleep (during LIA), ECs should be silent. Consistent with this hypothesis are data showing that while ECs cells are phase locked to the theta rhythm, they are uncorrelated to LIA, and conversely, while ECd cells are uncorrelated to theta, they are correlated to the LIA EEG signal (Chrobak & Buzsáki, 1994). During the self-localization procedure, this theory requires the hippocampus to show LIA also, but ECs cells should fire at a constant rate and would still be uncorrelated to LIA.

Simulations. In order to demonstrate the viability of the consolidation proposal, we simulated LTM as a three-layer feedforward network (HLS, LTM_H, and LTM_O) and trained the weights (HLS $\longrightarrow$ LTM_H and $\mathrm{LTM}_H \longrightarrow$ LTM_O) by the standard backpropagation algorithm (Rumelhart, Hinton, & Williams, 1986; Hertz, Krogh, & Palmer, 1991).

The training regimen for the LTM network consisted of 40 dreams, each of which were 5-second sequences of CA3/1 representations as the network was allowed to settle from noise and progress along the stored routes as described in the previous section. Each sequence was sampled at 100 ms intervals (see Figure 9). We used these representations to regenerate HLS representations and then generated directions of motion $\phi(t)$ by comparing each sample with its successor. We regenerated HLS(t) by taking the represented position in CA3 at time t and setting the HLS representation to the representation of that position. In the animal, we expect this regeneration process to occur from feedback connections between ECd and parietal cortex. Because CA3/1 is a 2D neural sheet, we can generate $\phi(t)$ by subtracting the center of mass of the CA3/1 representation at time t from the center of mass at time $t+\Delta t$. These input-output pairs $\langle \mathrm{HLS}(t), \phi(t) \rangle$ formed the training set for the LTM network.

The three-layer LTM network had 400 inputs, corresponding to each of the HLS neurons, 40 hidden units (LTM_H), and 2 outputs (LTM_O) representing $\cos(\phi)$ and $\sin(\phi)$. We do not claim that cortical representations look anything like this or that cortical learning occurs by backpropogation;

we only want to demonstrate that enough information is contained in the dreams to train a cortical network on the desired input-output function. We trained the network on repeated presentations of 40 dreams (1000 epochs, 40,000 total presentations). Each dream was sampled at 1 Hz, so that five input-output pairs were selected from each dream to train the network. This was done because subsequent pairs from a single dream are very similar, and this helped minimize computation time. In actual animals, we expect the entire dream to be used, and we expect each dream to be unique.

To test the long-term memory, we placed the simulated animal at evenly spaced positions in the environment (5-cm spacing) and measured the direction that would be taken at each location. Figure 10 shows the output vector in LTM_O at each location. The vector field shows interpolations between the routes because the HLS representations are distributed over approximately 50 to 100 neurons with similar place field centers. Weights from cells with place field centers situated between two routes end up trained to the vector average of the two routes.

To demonstrate that the network successfully learned the task, we simulated placing the animal at each of the four cardinal locations and moving it in the direction represented in LTM_O. As the simulated animal moved through the environment, it was allowed to continue following the direction represented in LTM_O (which changed as the representation in HLS changed). If the animal hit the wall, it turned clockwise if the direction represented in LTM_O was to the right of the normal vector to the wall and left otherwise. It turned in 45-degree increments until it could proceed. This has the effect of moving the animal along the wall. If the length of the LTM_O vector was zero, the animal proceeded along the shortest path to the wall until the LTM_O vector was nonzero or the animal hit the wall. The simulation ended after 60 seconds, or when the animal was within 10 cm of the goal. Figure 10 shows the paths taken from the four starting locations.

The explanation for the amnesias seen after hippocampal lesions can be clearly seen in Figure 10. If the hippocampus is lesioned before LTM has been trained, the simulated animal cannot solve the task at all (anterograde amnesia). If the hippocampus is lesioned after LTM has been partially trained, the animal at least approaches the goal but shows severe impairments (retrograde amnesia). Finally, if LTM has been fully trained by hippocampal dreams, the animal can solve the task accurately even with hippocampal lesions (retrograde amnesia is limited).

3 Predictions

This theory makes a number of predictions. We detail four here:

- *Coherency of the place code during LIA*. We hypothesize that the self-localization that must occur on reentry into a familiar environment is

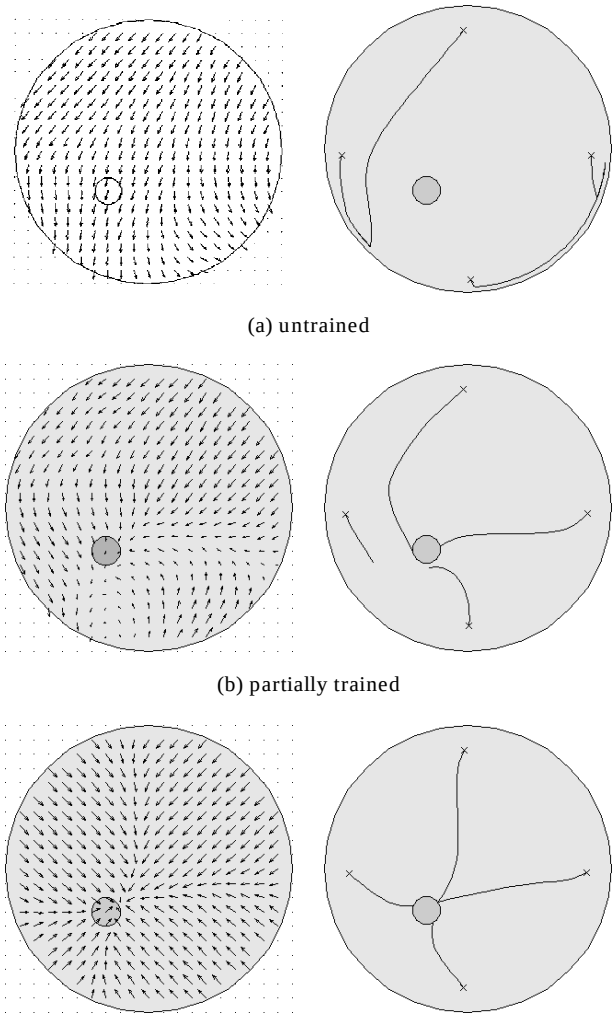

(a) untrained

(b) partially trained

(c) completely trained

Figure 10: Memory consolidation in a network trained on dreams. (Left) The vector field represented at each location by the long-term memory network. (Right) Paths taken by a simulated animal with hippocampal lesions, navigating using just its cortical route memory. Light gray area denotes arena; dark gray area denotes platform.

realized during LIA. This predicts that rodents reentering a familiar environment will show LIA before movement and that the place code will begin in an incoherent state and become coherent through the course of the LIA. Coherency can be rigorously defined as the inverse

width of the confidence interval of the location represented by the place code.

- *Automaticity of overlearned tasks.* We hypothesize that an animal does not store its entire cognitive map in long-term memory; it stores only routes. This predicts a difference in the ability to handle obstacles in normal and hippocampal animals, even after the span of the limited retrograde amnesia (12 weeks in the task above). Normal animals have little trouble planning trajectories to reach goals even with barriers (Tolman, 1948; Poucet, Thinus-blanc, & Chapuis, 1983; Leonard & McNaughton, 1990; Chapuis & Scardigli, 1993), but this requires a representation of the animal's position in the environment (O'Keefe & Nadel, 1978). Thus this requires the cognitive map and the hippocampus. Our simplified model of the mechanism to plan the path to the goal cannot plan paths around barriers. However, a more realistic model of nucleus accumbens should be able to plan such trajectories. Nevertheless, even a more complex goal memory model capable of planning paths around barriers would require a representation of current position. An animal dependent on a route memory (an animal with a trained LTM, but no hippocampus) would be severely impaired when a barrier is placed across the route.

- *Navigation without a hippocampus.* A number of alternative models of rodent navigation (e.g., Blum & Abbott, 1996; Samsonovich & McNaughton, in press) have suggested that the hippocampus is used in online navigation. Our model suggests that the hippocampus is necessary only for self-localization upon entry into a familiar environment. This means that rodents with hippocampal lesions should be able to perform tasks that do not require switching environments. For example, rodents should demonstrate path integration, such as the Mittelstaedt and Mittelstaedt (1980) pup retrieval task, even with hippocampal lesions. Presumably an animal in such a situation does not have to relocalize at any point in the task.

  Normally, rodents in the water maze task are carried back to a home cage to rest between trials or are disoriented by carrying them around the room before placing them at the starting point. Our model suggests that hippocampal animals may be able to perform the Morris water maze task if it is simplified in the following way: place the animal at the goal location (the hidden platform) and then smoothly carry the animal directly to the release point. A number of researchers have explored the Morris water maze with hippocampal and subicular lesions (Morris et al., 1982, 1990; McDonald & White, 1994), but we know of no experiments examining this simplified task.

- *Phase precession precedes asymmetric connection strengths.* Tsodyks, Skaggs, Sejnowski, & McNaughton (1996) present a model of how

phase precession can be generated by asymmetric connections. Although we do not present a model of how phase precession is generated, our theory requires that phase precession precede the asymmetry of connection strengths. This means that one should see phase precession from the first trajectory through an environment. This does not imply that once the connection strengths are asymmetric, they cannot enhance the phase precession effect. Once the connection strengths are asymmetric, a positive feedback situation exists between the asymmetry of the connection strengths and phase precession, which can enhance both.

4 Discussion

4.1 The Dual Role Played by the Hippocampus. We have suggested here that the hippocampus plays two roles:

1. *Self-localization.* On reentry into a familiar environment, self-localization is accomplished in the course of a sharp wave in the hippocampal formation.

2. *Route replay.* As the animal traverses routes in a familiar environment, these routes are stored in the recurrent connections of CA3. These routes are replayed during sleep.

These two roles can coexist because when there are sensory signals providing candidate biases to CA3, they are enough to counter the drift provided by the asymmetric connections. In the absence of sensory input, there is nothing to stop the drift, and the representation replays a recent route.

Blum and Abbott (1996) predicted that if asymmetric connections are present (such as is necessary to produce the route replay effect), then the place fields of cells should shift backward along much-repeated paths. Mehta and McNaughton (in press) have confirmed this prediction: the place fields of cells do shift backward along a much-repeated path as the animal runs along that path. They trained a rat to run in a loop on a rectangular elevated maze and recorded place cells in CA1. They found that the place fields shifted backward over the course of a single session.

This would seem to cause a problem for the self-localization role because the place fields would not represent the same location at the end of the run. However, Mehta, and McNaughton also showed that while the distance along the track covered by the place field shifted by almost 50 percent (approximately 5–7 cm), the center of mass of the field shifted very little (20 percent; approximately 2–3 cm). This duality is what we need for the dual-role hippocampus (self-localization / route storage). Because the number of cells overlapped by this expanded place field increases due to the increase in area covered, longer routes can be stored, but because the center of mass

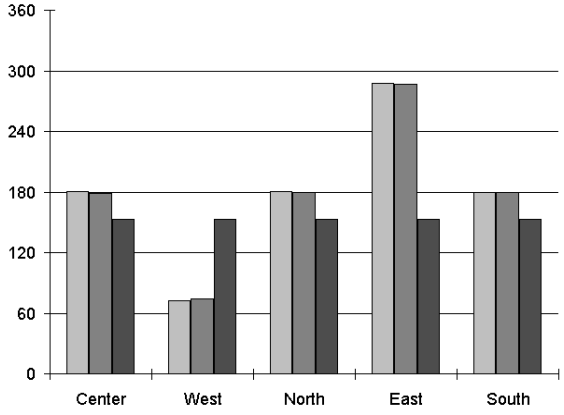

Figure 11: x-coordinate (on a toroidal map) of the final representation of the parallel relaxation process with (light bars) local view input and symmetric weights only, (medium bars) local view input and both symmetric and asymmetric weights, and (dark bars) both symmetric and asymmetric weights without local view input. y-coordinates (not shown) are similar.

of the field does not shift much, the self-localization process will not be affected very much.

Figure 11 shows the x-coordinate of the CA3/1 representation after a simulated sharp wave occurred under three conditions:

- *Light bars.* When the simulated animal was at five different locations in an environment but the recurrent connection matrix between CA3 excitatory cells consisted of symmetric weights only. (This is the simulation used in section 2.2.)
- *Medium bars.* When the simulated animal was at the same five locations but the CA3 recurrent connection matrix consisted of both symmetric and asymmetric weights. (This is the simulation used in section 2.4, but with a nonzero HLS representation.)
- *Dark bars.* In the absense of sensory input but when the representation began near each of the four locations above. (This is the simulation used in section 2.4.)

There is no difference at all between the results shown in the light and medium bars, indicating that the local view input is sufficient to hold the representation in place. On the other hand, when there is no sensory input to hold the representation in place, it drifts (see Figure 8). Even though the representation started near one of the five initial locations, in the absence of sensory input, by the end of the process, it had drifted to the goal location

(see also Figure 9). This shows that the two modes are in fact compatible (at least in simulation).

4.2 Why Can't the Animal Learn the Task Without a Hippocampus? Even with hippocampal lesions, animals can still learn simplified versions of the Morris water maze. If, for example, the platform is visible or the animal only begins from a single starting point or there is a landmark indicating the direction to swim (Eichenbaum, Steward, & Morris, 1990; McDonald & White, 1994), the animal can perform at near-normal levels. These simplified versions of the task allow routes to be learned based on single cues (the visible platform, a single motor program, a directional cue).

Brown and Sharp (1995) built a model of nucleus accumbens that uses reinforcement learning (Sutton, 1992) to back the direction-to-goal function from near the goal to locations distant from it. Their model learns to associate directions with representations of position. They use place cells as inputs, but there is no reason that the same process would not work with local view representations.

Because animals cannot learn to solve the Morris water maze without a hippocampus, we hypothesize that such a mechanism (as envisioned by Brown and Sharp) does not exist separate from the hippocampus. A number of authors have suggested that there are multiple memory and navigation systems (see, for example, O'Keefe & Nadel, 1978; Squire, 1992; Cohen & Eichenbaum, 1993). The Morris water maze data can be explained by assumming that there are three systems that allow animals to navigate within an environment:

1. *Locale system.* Animals can use constellations of distal cues to localize themselves on a "cognitive map"; combined with a memory of the goal location on the same map, the animals can navigate to the goal. This system requires a hippocampus. We detailed a model of this system in sections 2.1 and 2.2.

2. *Taxon system.* Animals can use visual cues to drive navigation directly toward or away from the cue. Evidence exists that this system requires the caudate nucleus (McDonald & White, 1994; Packard & McGaugh, 1992; Packard, 1994). We have not included a model of the taxon system in this article. Because there is no visible cue in the hidden-platform water maze task, it would not help the animal find the platform.

3. *Route system.* Routes stored in the hippocampus can be written out to the cortex, so that directions necessary to reach a goal are associated with local views. This is the system detailed in section 2.5 (see also section 4.3). This system requires training for each step the animal must take; it cannot learn to associate local views with directions to distant goals without hippocampal help (through route replay).

If there were a way to show the animal the route to the goal, it might be possible to train the route system even without a hippocampus. Whishaw, Cassell, and Jarrard (1995) and Schallert, Day, Weisend, and Sutherland (1996) both showed ways to train the route system directly and found that animals could learn to solve the water maze even with hippocampal lesions. Whishaw et al. (1995) trained animals with fimbria / fornix lesions to find a visible platform and then removed the visible platform. These animals concentrated their search where the platform had been. Schallert et al. (1996) used animals with kainate-colchicine hippocampal lesions. The animals were first trained with a large platform that filled almost the entire maze. Once the animals could reach that platform reliably, it was shrunk trial by trial until it was the same size as a typical platform in a water maze task. Again, the animals could learn to solve the water maze without a hippocampus.

4.3 Where Is the Route System? Although the data are not yet conclusive, we suggest that the most likely candidate for anatomical instantiation of the route system is from posterior parietal to posterior cingulate cortex.

There is a lot of evidence that posterior parietal cortex (PPC) supports the local view subsystem (we have labeled this *HLS* in Figure 1). Rodents with PPC lesions have extreme difficulties with spatial tasks, including the radial maze task, the hidden-platform water maze, and the visual-platform water maze (Kolb & Walkey, 1987). Rodents with PPC lesions are very poor at trajectories; when performing the water maze, they never improve their initial trajectory, even with visual cues informing the location of the platform (Kolb, 1990). DiMattia and Kesner (1988) report that PPC lesions forced the animals to revert to a random search strategy.

McNaughton et al. (1994) report that the firing rates of cells recorded from rodent posterior parietal cortex while an animal traversed a radial arm maze were correlated with whether the animal was turning (some cells were tuned to left turns, some to right turns, some to straight-ahead movements), as well as whether the animal was progressing toward or away from the center of the maze. Chen, Lin, Green, Barnes, and McNaughton (1994) found that approximately 5 to 10 percent of the cells were sensitive to direction on the radial arm maze, and of those, half required a cue in order to show a directional tuning. These cells fired only if the cue was present. But if the cue was present and then removed, some of the cells continued to show a tuning to direction, as if the cell remembered the direction the cue had been. They tended to show a broader tuning to direction than normal head direction cells (Chen et al., 1994).

If the parietal cortex forms the input to the route system, where is the output? We can identify requirements for the output components of the route system:

- It must receive input from the local view (parietal cortex).

- It must receive input from the hippocampus.
- It must receive input from head direction representations (postsubiculum and the anterior thalamic nuclei).
- It must send efferent projections to motor structures.
- It must be able to represent intended actions and directions of motion.
- Combined lesions of the area with hippocampus should cause impairments in navigation tasks. Although lesions of the area alone might not cause severe impairments, combined lesions should produce devastating results.

Posterior cingulate cortex (also known as retrosplenial cortex or area 29) is the best candidate for the route storage in long-term memory (labeled *LTM* in Figure 1). It is bidirectionally connected with parietal cortex (Vogt, 1985) and with postsubiculum (Wyss & van Groen, 1992; Finch, 1993), receives input from the anterior thalamic nuclei (Sripanidkulchai & Wyss, 1986; van Groen, Vogt, & Wyss, 1993) and subiculum (Vogt, 1985; Wyss & van Groen, 1992), and sends projections to motor cortex (Finch, 1993). Single cell recordings from posterior cingulate cortex show (rare) correlations to head direction and (more common) correlations to behavior (Chen et al., 1994).

Also supporting this hypothesis are data from Sutherland and Hoesing (1993), in which they showed that both cingulate and parietal lesions continue to have devastating effects, even if they occur 12 weeks after training. They suggest that posterior parietal cortex represents stimulus information, and posterior cingulate cortex then transforms this representation back into motor commands.

4.4 Retrograde Amnesia. Anterograde amnesia can occur when a system necessary for memory storage breaks down, and retrograde amnesia can occur when the actual storage site is damaged, but temporally graded retrograde amnesia implies a consolidation mechansism.

The role of the hippocampus in consolidation has been hotly debated recently (Nadel, 1991; Squire, 1992; Cohen & Eichenbaum, 1993; Weisend, Astur, & Sutherland, 1996; Koerner, Thomas, Weisend, & Sutherland, 1996; Bohbot, Liu, Thurm, Nadel, & Bures, 1996a). There are two time courses for consolidation that need to be handled separately. Short-term consolidation (STC) occurs over the course of hours (0.5 to 72) and long-term consolidation (LTC) occurs over the course of days to years.

STC is most likely to occur as a consequence of the biophysical time course of LTP. If LTP is disrupted—for example, chemically (Smith, 1995), by electroconvulsive shock (Bohbot, Otahal, Liu, Nadel, & Bures, 1996b), or by lesions (Bohbot et al., 1996a)—then memories are lost. But the window in which these kinds of memories can be disrupted is short. (See McGaugh (1966) for a review.)

On the other hand, LTC occurs because memories are stored in a temporary buffer and then written out from that buffer into long-term storage. The hypothesis that hippocampus serves as the temporary buffer was first proposed by Marr (1971). More recently, the theory has been extended by a number of researchers (Squire, 1992; Cohen & Eichenbaum, 1993; McClelland et al., 1995). In this article, we have tried to show how the hippocampus can store and replay memories that would allow an animal to solve a specific task (i.e., the Morris water maze).

Although the existence of STC seems to be well accepted, the existence of LTC has been called into question recently. We therefore review the relevant data on LTC in rodents and refer readers to other reviews of the nonrodent literature (Squire, 1992; Cohen & Eichenbaum, 1993; Rempel-Clower, Zola, Squire, & Amaral, 1996).

The Sutherland and Hoesing (1993) result cited in section 1 would seem to lay the question to rest conclusively, but the lesion was made with colchicine (Sutherland, personal communication), which selectively targets dentate gyrus granule cells and tends to spare the CA3 and CA1 fields.

Weisend et al. (1996) report that rats with hippocampal lesions show complete retrograde amnesia even out to 36 weeks; however, the lesions included large portions of dorsal subiculum. Dorsal subiculum is likely to play a role in path integration (Redish & Touretzky, 1997), which could affect navigation even in well-consolidated tasks. Bohbot et al. (1996a) found that in their experiments, the "temporal gradient of retrograde amnesia" is better explained as an effect of time since surgery: the lateral ventricles expand as time progresses postsurgery, compressing nearby structures. Koerner et al. (1996) explicitly tested the limited nature of retrograde amnesia in the hidden-platform water maze by comparing (1) animals that learned the task and waited 12 weeks, (2) animals that learned the task and received repetitions each week for 12 weeks, and (3) animals that learned 13 different hidden-platform water mazes over those 13 weeks. After hippocampal lesions (which again encroached on dorsal subiculum) only group 2 showed any performance above chance.

Cho, Berracochea, and Jaffard (1993) report definite retrograde amnesia in mice: animals with bilateral ibotenate entorhinal cortex lesions showed significant impairments of two-choice discriminations (on an eight-arm radial maze) for discriminations learned 0.5 to 2 weeks prior to surgery, but no impairment at all for those learned 4 or 8 weeks prior to surgery. In fact, the animals were better at the 4- or 8-week old discriminations than at the 2-week discriminations. Normals are better at recently learned discriminations. This suggests a consolidation time measured in weeks, but with entorhinal, not hippocampal, lesions.

Winocur (1990) examined social food preference in rats with bilateral electrolytic hippocampal lesions. Normal rats, when paired with a demonstrator rat that has recently eaten a food, prefer that food over other foods. Rats that first acquired the preference (from the demonstrator rat) and then

experienced surgery (a dorsal hippocampal lesion) showed a significant impairment when fewer than 2 days intervened between acquisition and surgery but were normal when 5 or 10 days intervened.

On the other hand, Bolhuis, Stewart, and Forrest (1994) examined animals trained on the water maze with two retention intervals between training and surgery: 2 to 4 days and 14 weeks. Animals were given either ibotenate hippocampal or subicular lesions. They found no temporal gradient to retrograde amnesia. Squire and Alvarez (1995) point out that animals were at chance for both time periods, and these floor effects may have precluded detection of a temporal gradient to the retrograde amnesia.

Although the extent of the temporal gradient of retrograde amnesia after hippocampal damage is still an open question, we present this model as an explanation of long-term consolidation, such as that described by Sutherland and Hoesing (1993), Winocur (1990), and Cho et al. (1993).

4.5 Related Work. The suggestion that correlational LTP in the CA3 recurrent connections combined with random exploration of an environment would produce a connection function where the learned synaptic weight is inversely related to distance between place field centers was first made by Muller et al. (1991). McNaughton, Skaggs, and Wilson (Wilson & McNaughton, 1994; Skaggs & McNaughton, 1996) have shown that after exploring an arena, cells with overlapping place fields in that arena are more likely to be correlated during subsequent sleep than those with nonoverlapping fields. However, none of these authors suggested the competitive dynamics that produces the self-localization mechanism in our explanation.

Shen and McNaughton (1997) have examined this competitive dynamics, with a model of attractor dynamics similar to that presented here in which cells with gaussian place fields exhibit correlational (Hebbian) learning. When these two effects are combined with random exploration, a local-excitation weight matrix is formed (see section 2.1). They demonstrate that when presented with random input and allowed to settle to a stable state, cells with recently visited place fields are more active than other cells, corresponding to data from Pavlides and Winson (1989) and Wilson and McNaughton (1994). They did not, however, look at sequences, as we have in section 2.4.

Samsonovich and McNaughton (in press) and Zhang (1996) have suggested that path integration can occur by the motion of a hill of activity on a two-dimensional neural sheet, and both demonstrate that dynamics such as we assume will produce a coherent hill from noise. However, both of these models of path integration differ significantly from ours. Zhang makes no claims about the anatomical implementation of the mechanism, and Samsonovich and McNaughton require the hippocampus for online navigation. Our model makes a different anatomical claim about the location of the path integration mechanism (see Redish & Touretzky, 1997, for more detail) and requires the hippocampus for self-localization upon reentry into an envi-

ronment only. In addition, Samsonovich and McNaughton explicitly reject the proposition that the hippocampus learns the CA3 recurrent connection strengths.

Tsodyks and Sejnowski (1995) examined a one-dimensional version of these dynamics, but looked at only a small simulation with a nonuniform distribution of place fields. This produces a situtation in which a few of the coherent representations are stronger attractors than the others. The dynamics of the inhomogeneous system are very different from that examined here and would predict a very nonhomogeneous distribution of both place field locations in the environment and place field sizes.

Blum and Abbott (1996) have shown that LTP combined with phase precession of place cells along the theta rhythm can also store routes to a goal. However, their theory requires the hippocampus for online navigation and is incompatible with the limited nature of the retrograde amnesia that we set out to explain.

Like the Burgess, Recce, and O'Keefe (1994) model, we hypothesize that phase precession is a consequence of an effect that exists prior to the asymmetric connection strengths. This puts our model in direct contrast to that of Tsodyks et al. (1996). Whereas we hypothesize that the asymmetric connections are a partial consequence of the phase precession, Tsodyks et al. (1996) have shown that the phase precession may in fact be generated by the asymmetric connections. These two properties may form a sort of feedback loop, which serves to strengthen the asymmetric connections.

Levy and his colleagues (Levy, 1996; Levy & Wu, 1996; Wu, Baxter, & Levy, 1996) have explored sequence learning in recurrent hippocampus-like networks and have demonstrated that cells with temporally extended effects can store and replay sequences of arbitrary binary vectors. They have not examined route learning in realistic tasks.

Finally, the hypothesis that hippocampus serves as a temporary store for memories that are replayed during sleep has been proposed by a number of authors (Marr, 1971; Cohen & Eichenbaum, 1993; McClelland et al., 1995), and Buzsáki (1989) specifically suggests that memories are stored during theta mode and recalled during LIA. But none of these authors addresses navigation, and none is able to model a specific task.

By concentrating on a single well-defined task, we have been able to look at all components of a memory-related task in detail, including how the hippocampus is used in solving the task (exploration and self-localization), storing the information needed to transfer the task to long-term memory (route learning), and training long-term memory to perform the task (replay of routes during sleep and consolidation). We believe this is a first step toward understanding the role of the hippocampus in long-term consolidation.

Appendix A: Simulation Details

Our model is shown in Figure 1. For simplicity we assume that each population consists of a 2D array of neurons with evenly spaced place field centers. We also assume a single environment. We leave relaxing these assumptions to future work, but refer readers to the discussion of *reference frames* in our previous work (Redish & Touretzky, 1997; Touretzky & Redish, 1996).

A.1 Neuronal Model. This neuronal model is more realistic than standard integrate-and-fire models but more abstract than compartmental models. It consists of three equations. For an extensive discussion of the derivation, see Wilson and Cowan (1972) and Pinto et al. (1996).

$$V_i(t) = \gamma_i + \sum_j w_{ij} S_j(t) \tag{A.1}$$

$$F_i(t) = \frac{1 + \tanh(V_i(t))}{2} \tag{A.2}$$

$$\tau_i \frac{dS_i(t)}{dt} = -S_i(t) + F_i(t) \tag{A.3}$$

$V_i(t)$ is proportional to the voltage in neuron i at time t; γ_i is a tonic inhibitory input to neuron i necessary for the dynamics of the system; w_{ij} is the synaptic weight from neuron j to neuron i; $F_i(t)$ is the firing rate at time t; and $S_i(t)$ is the synaptic drive of neuron i. τ_i is a time constant, proportional to the decay of the postsynaptic potential produced by neuron i in neurons on which it synapses. Synaptic drive is a nonmeasurable property of a neuron and can be understood as the effect of the neuron on all the neurons on which it synapses divided by the synaptic weight across each synapse.

As has been shown by Pinto et al. (1996), equations A.1 through A.3 form a consistent neuronal model that can be understood as describing either a continuous approximation to a single neuron with $F_i(t)$ being the probability of firing at time t, or a population of neurons with $V_i(t)$ being proportional to the average voltage and $F_i(t)$ the fraction of neurons in the population firing a spike at time t.

During movement, certain connections within the model are allowed to learn based on an LTP-type Hebbian learning rule in which a connection weight is increased if the product of the synaptic drive of the presynaptic neuron and the firing rate of the postsynaptic neuron is greater than the current weight:

$$\eta \frac{dw_{ij}}{dt} = [-w_{ij} + (S_j \cdot F_i)^2]^+. \tag{A.4}$$

Table A.1: Exporation Parameters I: Brain Structures

Brain Structure	Neurons	Parameters
High-level sensory (HLS)	204	12×12 (place) + 60 (direction)
Postsubiculum (PoS)	60	Tracks head direction
Subiculum / parasubiculum (Sub / PaS)	144	12×12; tracks position by path integration
Superficial entorhinal cortex (ECs)	144	$\tau = 1$ ms, $\gamma = -2$
CA3/1 pyramidal pool (CAE)	144	12×12, $\tau = 10$ ms, $\gamma = -1.5$

Table A.2: Exploration Parameters II: Connections

Connections, Exploration simulation	
Input $\longrightarrow$ Output	Parameters
HLS $\longrightarrow$ ECs	$\eta = 0.5, \sigma = 20$ cm
CAE $\longrightarrow$ CAE	$\eta = 0.5, \sigma = 20$ cm
ECs $\longrightarrow$ CAE	$w = 5.0$, one-to-one
ECs $\longrightarrow$ Sub / PaS	$w = 1.0, \sigma = 2$ cm
Sub / PaS $\longrightarrow$ ECs	$w = 2.0, \sigma = 2$ cm

η is a learning rate, and $[\ ]^+$ signifies rectification at 0, so weights can only increase. This is a simplification, and we plan to include LTD in future work.

Connections involving the LTM network were used in modeling consolidation and trained separately by the standard backpropagation of error algorithm (Rumelhart et al., 1986; Hertz et al., 1991).

A.2 Parameters and Specifics. We used a time step of 1 ms in all simulations.

A.2.1 Exploration. We do not simulate the mechanisms by which these structures track head direction or position but refer readers to already published models that demonstrate the viability of the assumed mechanisms (Redish et al., 1996; Redish & Touretzky, 1997; Samsonovich & McNaughton, in press; Zhang, 1996).

In order to sidestep the complexities involved in interpreting visual input, we assume that HLS represents the candidate positions given by the local view as a hill of activity on a two-dimensional toroidal grid (12×12 neurons) and the candidate head directions given as a hill of activity on a one-dimensional ring (60 neurons). HLS is assumed to be reset correctly to show the entering location and head direction and to track the position and direction accurately with movement through the environment. This simulates the local view.

The parameters for the exploration simulation are listed in Tables A.1 and A.2.

Table A.3: Self-Localization Parameters I: Brain Structures

Brain Structure	Neurons	Parameters
High-level sensory (HLS)	460	$(20 \times 20 + 60)$
Postsubiculum (PoS)	60	Tracks head direction
Subiculum / parasubiculum (Sub / PaS)	400	20×20; tracks position by path integration
Superficial entorhinal cortex (ECs)	400	$\tau = 1$ ms, $\gamma = -2$
CA3 / 1 pyramidal pool (CAE)	144	20×20, $\tau = 10$ ms, $\gamma = -1.5$, $\rho = 25^\circ$
CA3 / 1 inhibitory pool (CAI)	400	$\tau = 2$ ms, $\gamma = -7.5$

Table A.4: Self-Localization Parameters II: Connections

Connections	
Input $\longrightarrow$ Output	Parameters
HLS $\longrightarrow$ ECs	$w = 5.0, \sigma = 20$ cm
HLS $\longrightarrow$ PoS	$w = 2.0, \sigma = 20$ cm
ECs $\longrightarrow$ CAE	$w = 5.0$, one-to-one
ECs $\longrightarrow$ Sub / PaS	$w = 1.0, \sigma = 2$ cm
Sub / PaS $\longrightarrow$ ECs	$w = 2.0, \sigma = 2$ cm
CAE $\longrightarrow$ CAE	$w = 5.0, \sigma = 20$ cm
CAE $\longrightarrow$ CAI	$w = 16.0, \sigma = 20$ cm
CAE $\longrightarrow$ Sub	$w = 5.0$, one-to-one
CAI $\longrightarrow$ CAE	$w = -8.0, \sigma = 200$ cm
CAI $\longrightarrow$ CAI	$w = -12.0, \sigma = 200$ cm

A.2.2 Self-localization. Connections indicated as one-to-one connect input cells to output cells with corresponding locations on an evenly spaced grid. This was done for computational speed; a random, sparse connection function produces similar results. Other connections are full connections in which the connection weight from input cell j to output cell i, w_{ij}, falls off as a gaussian function of distance between the place field centers of the two cells i and j.

The parameters for the self-localization simulation are listed in Tables A.3 and A.4.

A.2.3 Route learning. The model used in the route-learning simulations was identical to the self-localization model with the change that the intrahippocampal connections had zero weight and the CAE $\longrightarrow$ CAE connection had a learning rate of $\eta = 0.01$.

A.2.4 Dreaming. The simulations used for dreaming were identical to those used to demonstrate self-localization, but the weight matrix used was that trained in section 2.3. In order to simulate "sleep," the HLS represen-

Table A.5: Consolidation Parameters

Backpropagation network simulating cortex		
Name	Role	Neurons
HLS	Input layer	400
LTM_H	Hidden layer	40
LTM_O	Output layer	2

tation was set to zero.

A.2.5 Consolidation. The HLS $\longrightarrow$ LTM_H and $LTM_H \longrightarrow LTM_O$ connections were fully connected and trained by the standard backpropagation of error algorithm. See the text for details.

The parameters for the backpropagation network used for the consolidation simulation are listed in Table A.5.

Acknowledgments

We thank two anonymous reviewers as well as Tad Blair, Bard Ermentrout, Bruce McNaughton, David Pinto, Alexei Samsonovich, Bill Skaggs, Matt Wilson, and Kechen Zhang for helpful discussions. We also acknowledge our fellow attendees at the workshop, From Place Cells to Episodic Memory: Models of the Hippocampus, at the Computational Neuroscience 1996 conference, particularly David August, Howard Eichenbaum, Michael Hasselmo, Chip Levy, and Ali Minai. This work has previously appeared in abstract form (Redish & Touretzky, 1996). This work was supported by National Science Foundation grant IBN-9631336.

References

Alvarez, P., & Squire, L. R. (1994). Memory consolidation and the medial temporal lobe: A simple network model. *Proceedings of the National Academy of Sciences, U.S.A., 91*, 7041–7045.

Amaral, D. G. (1993). Emerging principles of intrinsic hippocampal organization. *Current Opinion in Neurobiology, 3*, 225–229.

Blum, K. I., & Abbott, L. F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation, 8*(1), 85–93.

Bohbot, V., Liu, Z., Thurm, S. L., Nadel, L., & Bures, J. (1996a). Spatial memory? Never without a hippocampus. *Society for Neuroscience Abstracts, 22*, 1873.

Bohbot, V., Otahal, P., Liu, Z., Nadel, L., & Bures, J. (1996b). Electroconvulsive shock and lidocaine reveal rapid consolidation of spatial working memory in the water maze. *Proceedings of the National Academy of Sciences, U.S.A., 93*, 4016–4019.

Bolhuis, J. J., Stewart, C. A., & Forrest, E. M. (1994). Retrograde amnesia and memory reactivation in rats with ibotenate lesions to the hippocampus or subiculum. *Quarterly Journal of Experimental Psychology, 47B*(2), 129–150.

Brown, M. A., & Sharp, P. E. (1995). Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus, 5*(3), 171–188.

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

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

Chapuis, N., & Scardigli, P. (1993). Shortcut ability in hamsters (*mesocricetus auratus*): The role of environmental and kinesthetic information. *Animal Learning and Behavior, 21*(3), 255–265.

Chen, L. L., Lin, L. H., Green, E. J., Barnes, C. A., & McNaughton, B. L. (1994). Head-direction cells in the rat posterior cortex: I. Anatomical distribution and behavioral modulation. *Experimental Brain Research, 101*, 8–23.

Cho, Y. H., Berracochea, D., & Jaffard, R. (1993). Extended temporal gradient for the retrograde and anterograde amnesia produced by ibotenate entorhinal cortex lesions in mice. *Journal of Neuroscience, 13*(4), 1759–1766.

Chrobak, J. J., & Buzsáki, G. (1994). Selective activation of deep layer (V-VI) retrohippocampal neurons during hippocampal sharp waves in the behaving rat. *Journal of Neuroscience, 14*(10), 6160–6170.

Cohen, N. J., & Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system*. Cambridge, MA: MIT Press.

Davis, M., Rainnie, D., & Cassell, M. (1994). Neurotransmission in the rat amygdala related to fear and anxiety. *Trends in Neurosciences, 17*(5), 208–214.

DiMattia, B. V. D., & Kesner, R. P. (1988). Spatial cognitive maps: Differential role of parietal cortex and hippocampal formation. *Behavioral Neuroscience, 102*(4), 471–480.

Eichenbaum, H., Stewart, C., and Morris, R. G. M. (1990). Hippocampal representation in place learning. *Journal of Neuroscience, 10*(11), 3531–3542.

Etienne, A. S. (1992). Navigation of a small mammal by dead reckoning and local cues. *Current Directions in Psychological Science, 1*(2), 48–52.

Finch, D. M. (1993). Hippocampal, subicular, and entorhinal afferents and synaptic integration in rodent cingulate cortex. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 224–248). Boston: Birkhauser.

Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.

Gluck, M. A., & Myers, C. E. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus, 3*(4), 491–516.

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

Hasselmo, M. E. (1995). Neuromodulation and cortical function: Modeling the physiological basis of behavior. *Behavioral and Brain Research, 67*(1), 1–27.

Hasselmo, M. E., & Schnell, E. (1994). Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: Com-

putational modeling and brain slice physiology. *Journal of Neuroscience, 14*(6), 3898–3914.

Hennevin, E., Hars, B., Maho, C., & Bloch, V. (1995). Processing of learned information in paradoxical sleep: Relevance for memory. *Behavioural and Brain Research, 69*, 125–135.

Hertz, J., Krogh, A., & Palmer, R. G. (1991). *Introduction to the theory of neural computation*. Reading, MA: Addison-Wesley.

Huerta, P. T., & Lisman, J. E. (1993). Heightened synaptic plasticity of hippocampal CA1 neurons during a cholinergically induced rhythmic state. *Nature, 364*, 723–725.

Koerner, A., Thomas, M. J., Weisend, M. P., & Sutherland, R. J. (1996). Hippocampal-dependent memory consolidation: An evaluation of three hypotheses. *Society for Neuroscience Abstracts, 22*, 1118.

Kolb, B. (1990). Posterior parietal and temporal association cortex. In B. Kolb and R. C. Tees (Eds.), *The cerebral cortex of the rat* (pp. 459–471). Cambridge, MA: MIT Press.

Kolb, B., & Walkey, J. (1987). Behavioural and anatomical studies of the posterior parietal cortex in the rat. *Behavioral and Brain Research, 23*, 127–145.

Kudrimoti, H. S., Skaggs, W. E., Barnes, C. A., McNaughton, B. L., Gerrard, J. L., Suster, M. S., & Weaver, K. L. (1996). REM sleep and the reactivation of recent correlation patterns in hippocampal neuronal ensembles. *Society for Neuroscience Abstracts, 22*, 1871.

Landfield, P. W., & Deadwyler, S. A. (Eds.) (1988). *Long-term potentiation: From biophysics to behavior*. New York: Alan R. Liss.

Larson, J., & Lynch, G. (1989). Theta pattern stimulation and the induction of LTP: The sequence in which synapses are stimulated determines the degree to which they potentiate. *Brain Research, 489*, 49–58.

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

Leonard, B., & McNaughton, B. L. (1990). Spatial representation in the rat: Conceptual, behavioral, and neurophysiological perspectives. In R. P. Kesner & D. S. Olton (Eds.), *Neurobiology of comparative cognition* (pp. 363–422). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London, Series B, 262*, 23–81.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102*(3), 419–457.

McDonald, R. J., & White, N. M. (1994). Parallel information processing in the water maze: Evidence for independent memory systems involving dorsal striatum and hippocampus. *Behavioral and Neural Biology, 61*, 260–270.

McGaugh, J. M. (1966). Time-dependent processes in memory storage. *Science, 153*, 1351–1358.

McNaughton, B. L., Mizumori, S. J. Y., Barnes, C. A., Leonard, B. J., Marquis, M., & Green, E. J. (1994). Cortical representation of motion during unrestrained spatial navigation in the rat. *Cerebral Cortex, 4*(1), 27–39.

Mehta, M. R., Barnes, C. A., & McNaughton, B. L. (in press). An evidence for Hebbian synaptic potentiation in hippocampus during behavior. In M. Mozer, M. Jordan, and T. Petsche (Eds.), *Advances in Neural Information Processing, 9*. Cambridge, MA: MIT Press.

Mittelstaedt, M. L., & Mittelstaedt, H. (1980). Homing by path integration in a mammal. *Naturwissenschaften, 67*, 566–567.

Mogenson, G. J. (1984). Limbic-motor integration—with emphasis on initiation of exploratory and goal-directed locamotion. In R. Bandler (Ed.), *Modulation of sensorimotor activity during alterations in behavioral states* (pp. 121–138). New York: Alan R. Liss.

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.

Morris, R. G. M., Schenk, F., Tweedie, F., and Jarrard, L. E. (1990). Ibotenate lesions of hippocampus and / or subiculum: Dissociating components of allocentric spatial learning. *European Journal of Neuroscience, 2*, 1016–1028.

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

Nadel, L. (Ed.) (1991). The hippocampus and space revisited [Special issue]. *Hippocampus, 1*(3).

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.

O'Keefe, J., & Recce, M. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus, 3*, 317–330.

O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus, 4*(6), 661–682.

Packard, M. G. (1994). "Place" versus "response" learning debate revisited in the brain. *Society for Neuroscience Abstracts, 20*, 1016.

Packard, M. G., & McGaugh, J. L. (1992). Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience, 106*(3), 439–446.

Pavlides, C., Greenstein, J., Grudman, M., & Winson, J. (1988). Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of θ-rhythm. *Brain Research, 439*, 383–387.

Pavlides, C., & Winson, J. (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *Journal of Neuroscience, 9*(8), 2907–2918.

Pinto, D. J., Brumberg, J. C., Simons, D. J., & Ermentrout, G. B. (1996). A quantitative population model of whisker barrels: Re-examining the Wilson-Cowan equations. *Journal of Computational Neuroscience, 3*(3), 247–264.

Poucet, B., Thinus-blanc, C., & Chapuis, N. (1983). Route planning in cats, in relation to the visibility of the goal. *Animal Behavior, 31*, 594–599.

Redish, A. D., Elga, A. N. & Touretzky, D. S. (1996). A coupled attractor model of the rodent head direction system. *Network, 7*(4), 671–685.

Redish, A. D., & Touretzky, D. S. (1996). Details of a comprehensive theory of rodent navigation. *Society for Neuroscience Abstracts, 22*, 678.

Redish, A. D., & Touretzky, D. S. (1997). Cognitive maps beyond the hippocampus. *Hippocampus, 7*(1), 15–35.

Rempel-Clower, N. L., Zola, S. M., Squire, L. R., & Amaral, D. G. (1996). Three cases of enduring memory impairment after bilateral damage limited to the hippocampal formation. *Journal of Neuroscience, 16*, 5233–5255.

Rolls, E. T. (1989). The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In R. Durbin, C. Miall, & G. Mitchison (Eds.), *The computing neuron* (pp. 125–159). Reading, MA: Addison-Wesley.

Rolls, E. T. (1996). A theory of hippocampal function in memory. *Hippocampus, 6*, 601–620.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propogation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (volume 1, pp. 318–362). Cambridge, MA: MIT Press.

Samsonovich, A., & McNaughton, B. L. (in press). Path integration and cognitive mapping in a continuous attractor model of the hippocampus. *Journal of Neuroscience*.

Schallert, T., Day, L. B., Weisend, M., & Sutherland, R. J. (1996). Spatial learning by hippocampal rats in the Morris water task. *Society for Neuroscience Abstracts, 22*, 678.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry, 20*, 11–21.

Shen, B., & McNaughton, B. L. (1997). Modeling the spontaneous reactivation of experience-specific hippocampal cell assembles during sleep. *Hippocampus, 6*(6), 685–693.

Skaggs, W. E. (1995). Relations between the theta rhythm and activity patterns of hippocampal neurons. Unpublished doctoral dissertation, University of Arizona.

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., Wilson, M. A., & Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus, 6*(2), 149–173.

Smith, C. (1995). Sleep states and memory processes. *Behavioural and Brain Research, 69*, 137–145.

Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychology Review, 99*(2), 195–231.

Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: A neurobiological perspective. *Current Opinion in Neurobiology, 5*, 169–177.

Squire, L. R., & Zola-Morgan, S. (1988). Memory: Brain systems and behavior. *Trends in Neurosciences, 11*(4), 170–175.

Sripanidkulchai, K., & Wyss, J. M. (1986). Thalamic projections to retrosplenial cortex in the rat. *Journal of Comparative Neurology, 254*, 143–165.

Stewart, M., & Fox, S. E. (1990). Do septal neruons pace the hippocampal theta rhythm? *Trends in Neurosciences, 13*(5), 163–168.

Sutherland, R. J., & Hoesing, J. M. (1993). Posterior cingulate cortex and spatial memory: A microlimnology analysis. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 461–477). Boston: Birkhauser.

Sutton, R. S. (Ed.) (1992). Reinforcement learning [Special issue]. *Machine Learning, 8*(3 / 4).

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review, 55*, 189–208.

Touretzky, D. S., & Redish, A. D. (1996). A theory of rodent navigation based on interacting representations of space. *Hippocampus, 6*(3), 247–270.

Tsodyks, M., & Sejnowski, T. J. (1995). Associative memory and hippocampal place cells. *International Journal of Neural Systems, 6* (Suppl.), 81–86.

Tsodyks, M. V., Skaggs, W. E., Sejnowski, T. J., & McNaughton, B. L. (1996). Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocampus, 6*(3), 271–280.

van Groen, T., Vogt, B. A., & Wyss, J. M. (1993). Interconnections between the thalamus and retrosplenial cortex in the rodent brain. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 123–150). Boston: Birkhauser.

Vanderwolf, C. H. (1971). Limbic-diencephalic mechanisms of voluntary movement. *Psychological Review, 78*(2), 83–113.

Vanderwolf, C. H. (1990). An introduction to the electrical activity of the cerebral cortex: Relations to behavior and control by subcortical inputs. In B. Kolb & R. C. Tees (Eds.), *The cerebral cortex of the rat* (pp. 151–189). Cambridge, MA: MIT Press.

Vanderwolf, C. H., & Leung, (1983). Hippocampal rhythmical slow activity: A brief history and the effects of entorhinal lesions and phencyclidine. In W. Seifert (Ed.), *Neurobiology of the hippocampus* (pp. 275–302). New York: Academic Press.

Vogt, B. A. (1985). Cingulate cortex. In A. Peters & E. Jones (Eds.), *Association and auditory cortices* (pp. 89–150). New York: Plenum Press.

Weisend, M. P., Astur, R. S., & Sutherland, R. J. (1996). The specificity and temporal characteristics of retrograde amnesia after hippocampal lesions. *Society for Neuroscience Abstracts, 22*, 1118.

Whishaw, I. Q., Cassel, J-C., & Jarrard, L. E. (1995). Rats with fimbria-fornix lesions display a place response in a swimming pool: A dissociation between getting there and knowing where. *Journal of Neuroscience, 15*(8), 5779–5788.

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal, 12*(1), 1–24.

Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science, 265*, 676–679.

Winocur, G. (1990). Anterograde and retrograde amnesia in rats with dorsal hippocampal or dorsomedial thalamic lesions. *Behavioural and Brain Research, 38*, 145–154.

Witter, M. P. (1993). Organization of the entorhinal-hippocampal system: A review of current anatomical data. *Hippocampus, 3*, 33–44.

Witter, M. P., Ostendorf, R. H., & Groenwegen, H. J. (1990). Heterogeneity in the dorsal subiculum of the rat. Distinct neuronal zones project to different cortical and subcortical targets. *European Journal of Neuroscience, 2*, 718–725.

Wood, E. R., Mumby, D. G., Pinel, J. P. J., & Phillips, A. G. (1996). Impaired object recognition memory in rats following ischemia-induced damage to the hippocampus. *Behavioral Neuroscience, 107*(1), 51–62.

Wu, X., Baxter, R. A., & Levy, W. B. (1996). Context codes and the effect of noisy learning on a simplified hippocampal CA3 model. *Biological Cybernetics, 74*, 159–165.

Wyss, J. M., & van Groen, T. (1992). Connections between the retrosplenial cortex and the hippocampal formation in the rat: A review. *Hippocampus, 2*(1), 1–12.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience, 16*(6), 2112–2126.

Zola-Morgan, S., & Squire, L. R. (1993). Neuroanatomy of memory. *Annual Review of Neuroscience, 16*, 547–563.

Received October 7, 1996; accepted March 26, 1997.