

Theory of Rodent Navigation Based on Interacting Representations of Space

David S. Touretzky¹ and A. David Redish²

¹Computer Science Department and ²Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, Pennsylvania

ABSTRACT: We present a computational theory of navigation in rodents based on interacting representations of place, head direction, and local view. An associated computer model is able to replicate a variety of behavioral and neurophysiological results from the rodent navigation literature. The theory and model generate predictions that are testable with current technologies. © 1996 Wiley-Liss, Inc.

KEY WORDS: place cells, head direction, local view, spatial learning, memory, path integration

INTRODUCTION

There is a wealth of data on rodent performance in spatial learning and memory tasks, but as yet there is no comprehensive theory of how rodents navigate. Toward that end, we present a computational model based on a collection of behavioral and neurophysiological observations that allows us to reproduce a variety of results within a single computer program. Ours is a systems-level model concerned with interactions among navigational subsystems, such as the place code, the head direction system, and the path integrator. Previous modeling efforts have generally focused on just one subsystem.

Our fundamental assumption, and a view held by other investigators as well (O'Keefe and Nadel, 1978; O'Keefe, 1989; Gallistel, 1990; McNaughton et al., 1991, 1996), is that at least two distinct spatial encodings are maintained in the rodent brain: a metric representation that supports a limited form of vector arithmetic, and a local view representation derived from bearing and/or distance relations to landmarks. The brain learns relationships between place descriptions in these two representations, and the hippocampal complex appears to operate on (and may be responsible for maintaining) the conjoint representation. We extend this core theory by showing how interactions among the two spatial representation systems, the hippocampal complex, and the head direction system can account for various facts about rodent navigation. Although only the hippocampal

place code is represented at the level of neuron-like elements, the model makes predictions about other brain areas that should be of interest to neuroscientists.

BACKGROUND

Place Cells

Single-unit recordings from CA3 and CA1 pyramidal cells show a high correlation between spike rate and the location of the animal in the environment. These cells fire at an elevated rate over a continuous, compact area, the *place field* of the cell (O'Keefe and Dostrovsky, 1971). See (Muller et al., 1991) and (McNaughton et al., 1994) for reviews.

Rats performing a spatial working memory task in which they must visit every arm of an eight-arm radial maze are sensitive to the positions of external landmarks, as can be demonstrated by rotating the landmarks part-way through the task (Suzuki et al., 1980). Hippocampal place fields rotate in correspondence with the rotation of prominent landmarks (Muller and Kubie, 1987; O'Keefe and Speakman, 1987; McNaughton et al., 1994; Knierim et al., 1995). In contrast, when landmarks are permuted rather than rotated, rats behave as if the environment were unfamiliar (Suzuki et al., 1980), and neural recordings suggest that the animal forms a different set of place codes¹ in the permuted environment that overlaps little with the previous set (Muller and Kubie, 1987; O'Keefe and Speakman, 1987).

These observations suggest that place cells can be modeled as visually driven pattern detectors tuned to bearings and/or distances to landmarks. A variety of such

Accepted for publication April 16, 1996.

Address correspondence and reprint requests to David S. Touretzky, Computer Science Department and Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213-3891.

¹A "place code" is a pattern of activity over the place cell population. It constitutes a distributed representation (Hinton et al., 1986) of the animal's position from which the true position can be estimated with high accuracy (Wilson and McNaughton, 1993).

models have been described in the literature, constructed by competitive learning (Sharp, 1991), recurrent back propagation learning (Shapiro and Hetherington, 1993; Hetherington and Shapiro, 1993), genetic algorithms (Treves et al., 1992), or radial basis functions (Zipser, 1985; Schmajuk and Blair, 1993; Burgess et al., 1994). These models all produce place cells with appropriately shaped fields, and since their responses are purely visually driven, the fields rotate with the visual cues. But this description of place cells is far from complete.

A second important feature of place cells is that their firing fields remain largely intact when prominent visual landmarks are removed (Muller and Kubie, 1987; O'Keefe and Speakman, 1987). An associative memory process was thought to be responsible for this. According to this view, a fully specified place code learned previously by the hippocampus could later be reconstructed from a partial pattern evoked by an incomplete set of visual cues (Rolls, 1989; McNaughton and Nadel, 1990). However, the memory models cited in these proposals use fixed point attractors (Marr, 1969; Kohonen, 1984). Such models store only discrete memories; they cannot produce the smoothly varying response pattern of place cells, some of whose firing rates fall off as Gaussian functions of distance from the place fields' centers. A smoothly varying response would appear to be fundamentally incompatible with the pattern completion function of an associative memory, because it is supposed to map many different partial patterns to the same whole. If the memory corrects a partial pattern when a landmark is missing, why would it not also correct the pattern (i.e., produce an identical firing rate in a place cell) when the landmark was at a slightly different distance or bearing?

Another problem with these memory models is that the stored patterns must be linearly independent to avoid mutual interference. If the animal is storing a separate memory for every viewing position it experiences, linear independence can only be maintained if the system is made extremely sensitive to differences in the original input patterns, and the patterns are then recoded to pull them apart, a function Marr (1969) ascribed to the dentate gyrus. But this degree of sensitivity would undercut any ability to do pattern completion in the event of missing landmarks.

To summarize, storing a separate memory for every possible view of the environment cannot be done without making the system overly sensitive to differences in input patterns, which blocks completion of partial patterns. Storing only a few memories is compatible with pattern completion, but in that case we would expect an identical response to similar input patterns, rather than the smoothly varying response exhibited by real place cells. Hence, matrix memory models using fixed point attractors are not viable models of the hippocampal place system.

A third critical observation about place cells is that they continue to fire when rats navigate in the dark (O'Keefe, 1976; McNaughton et al., 1989; Quirk et al., 1990; Markus et al., 1994). New place fields can even be recruited in the dark (Quirk et al., 1990), so the cells cannot be purely visual pattern detectors. This observation poses difficulties for the associative memory theory as well. The firing rates of the cells change as the rat moves through the environment. If associative recall were solely

responsible for place cell activity in the absence of visual input, the animal would have to memorize not only the place codes, but also a transition table showing for each combination of present location and motor action what its new location would be. Such a model was proposed by McNaughton and Nadel (1990), but no simulations were done for realistic environments. The combinatorics of this approach appear prohibitively expensive, both in the amount of memory such a table would require and the amount of exploration needed to fill in the entries.

An alternative hypothesis is that place cells in the dark are driven by path integration (O'Keefe, 1976; Wan et al., 1994a,b; McNaughton et al., 1996). Muller et al. (1991) and McNaughton et al. (1994) have proposed similar hypotheses in which animals remember the relative locations of landmarks even when they are out of sight. Place cells tuned to these virtual (invisible) landmarks would continue to fire in the dark. The proposal would require a landmark memory somewhere in the brain that interacts with the hippocampal formation, and further implies that the animal must maintain its position internally, by path integration. We see no need to posit the landmark memory: Place cells can be driven directly by the same process that allows the animal to perform path integration. We return to the discussion of the relation between virtual landmarks and path integration in Path Integration and Virtual Landmarks.

Path Integration

Mittelstaedt and Mittelstaedt (1980) showed that a female gerbil searching for a missing pup via an apparent random walk could execute a straight-line return to the nest once the pup was found. The experiment was performed in the dark to rule out visual homing. Displacement of the animal during its search caused the return path to be offset by a comparable amount, eliminating the possibility that auditory or olfactory cues guided the trajectory. Hence, they concluded that gerbils keep track of their position by path integration.

Etienne (1987) showed similarly that golden hamsters trained to find food at the center of a circular arena used path integration to return to the nest. When the nest was rotated 90° or 180° while the animal was at the center of the arena, the animal returned to where the nest had been originally, ignoring the rotation of the arena. The animals integrated their own active and passive rotation to determine the direction to the nest, but as was seen in the Mittelstaedt and Mittelstaedt experiments, the hamsters did not integrate passive translation. Alyan and Jander (1994) have shown that in the house mouse *Mus musculus* performing a pup retrieval task, path integration is initially the dominant influence on navigation behavior, but visual cues that conflict with path integration information due to rotation of the environment gradually acquire priority as the animal becomes more familiar with the arena.

Path integration has been demonstrated in a variety of other animals ranging from insects to humans (see Gallistel, 1990; Maurer and Seguinot, 1995, for reviews), although of course the mechanisms enabling these animals to path integrate may be quite different from those in rodents.

The neuroanatomical substrate for rodent path integration is presently under investigation. Recent proposals suggest that it arises from the interaction of several structures: the hippocampus (McNaughton et al., 1996), subiculum (Sharp et al., 1995), and parasubiculum (Redish and Touretzky, 1995; Redish and Touretzky, 1996a) may all be involved. Data demonstrating that striatal lesions impair a rat's ability to return from passive transport, but not to perform an odor-following task, suggest that the basal ganglia may also contribute (Potegal, 1972, 1982; Abraham et al., 1983).

Any representation subserving path integration must facilitate vector addition. In earlier work, we suggested a neural representation for direction and distance that supports several vector arithmetic operations (Touretzky et al., 1993). Wittmann and Schwegler (1995) subsequently showed how this encoding could be used to model path integration in ants. Another representation suggested for path integration in rodents is described by Samsonovich and McNaughton (1995), based on a moving hill of activation over a two-dimensional (2D) array of cells.

Head Direction

In order for path integration to be possible, the animal must have an internal representation of the direction of its motion. Cells that show a unimodal tuning to head direction independent of location have been reported in at least five areas of the rat brain: postsubiculum (Taube et al., 1990a,b), the lateral dorsal nucleus (LDN) of the thalamus (Mizumori and Williams, 1993), the anterior thalamic nuclei (ATN) (Blair and Sharp, 1995; Knierim et al., 1995; Taube, 1995), and to a lesser extent, striatum (Wiener, 1993; Mizumori and Cooper, 1995) and posterior parietal and cingulate cortex (Chen, 1991; Chen et al., 1994a,b). These cells (hereafter *head direction cells*) provide a sort of internal compass.

When landmarks in a familiar environment are rotated about the animal, tuning curves of head direction cells rotate accordingly (Taube et al., 1990b; McNaughton et al., 1994; Goodridge and Taube, 1995; Knierim et al., 1995; Taube, 1995). But in unfamiliar environments, vestibular information dominates and the rat does not respond to the rotation. Thus, like place cells, head direction cell responses can be controlled by visual landmarks when the configuration is familiar.

Head direction cells continue to fire in the dark. But if the animal wanders about in a darkened circular arena with no other positional cues available, they eventually drift.² Hippocampal place fields appear to drift in synchrony with the head direction system (McNaughton et al., 1994; Knierim et al., 1995).

By comparing the preferred direction of a head direction cell with its preferred direction at a previous time, we can measure the *precession* of the cell's preferred direction. The difference in preferred directions of any pair of head direction cells appears to

be constant (Taube et al., 1990b; Taube, 1995), so whenever one cell is seen to precess, the rest do likewise, by a similar amount. Thus, we can talk about the precession of the head direction system as a whole, abstracting away the individual cells.

Although there are some data suggesting different roles for the various head direction areas (Mizumori and Williams, 1993; Chen et al., 1994a,b; Blair and Sharp, 1995), we will not attempt to separate out those differences here. We will note, though, that the head direction system as a whole should make connections to both the path integration and local view systems, and that disruption of the head direction system should interfere with any behavior that depends on path integration.

Local View

The *local view* (McNaughton, 1989) of an environment is what can be seen from a particular viewing position. With a sufficiently rich set of cues and no pathological symmetries in the environment, local views describe unique places. In real animals, the local view module may be realized by areas of parietal cortex. But the neural representation of complex scenes is poorly understood at present, so place cell models encode visual input at a highly abstract level. Most treat landmarks as points, represented by their bearing and/or distance from the animal (Zipser, 1985; McNaughton et al., 1989; O'Keefe, 1989; Sharp, 1991).

In our model, visual landmarks are represented as triples $\langle \tau_i, r_i, \theta_i \rangle$, where τ_i indicates the type of the i th landmark, r_i its distance from the animal, and θ_i its bearing relative to the animal's midline (egocentric bearing). Types are not unique; there may be multiple objects in view that are of the same type. For example, in the Collet et al. tasks described in Navigation Using Landmark Arrays, several visually identical cylindrical landmarks are used. There is evidence that rats use allocentric bearings to solve the resulting "binding problem".³ In a rectangular arena, all corners may have the same type, but alternatively, following Gallistel (1990), we may distinguish two types of corners: one with the long wall on the left and the short wall on the right, and the other reversed; diagonally opposite corners would be of the same type.

This simple caricature of the rodent visual system is adequate for the present needs of our navigational model, but its limitations should be acknowledged. Perfect object recognition is assumed as a primitive. We have not considered how place cell activity might vary in response to objects that are of a slightly different shape, size, or color than the expected landmarks. And our theory assumes that rodents can recognize objects in unfamiliar orientations, despite a lack of data supporting this. However, the experiments we consider have all used either a limited range of viewing positions (e.g., a white cue card viewed only from the interior of a small arena) or circularly symmetric landmarks, so this limitation does not pose a problem for our simulations.

²Mizumori and Williams (1993) and Knierim et al. (1993) report that head direction cells in LDN and ATN drift after about 2 minutes. Taube (personal communication), recording from postsubiculum, has observed stable directional tuning over longer periods.

³This is the problem of matching perceived to remembered landmarks. Binding problems occur in many contexts in cognitive science. Connectionist models have often been proposed as solutions.

On the other hand, our model is more perceptually sophisticated than its predecessors in that it allows for surface-type as well as point-type landmarks. In this case $\langle r_i, \theta_i \rangle$ encodes the normal vector from the animal to the closest point on the surface. Using a surface-type landmark, a place cell tuned to the arena wall produces place fields contiguous with the wall and makes the model sensitive to the geometry of the environment, which is consistent with observations of CA1 cells by Muller et al. (1987) and Muller and Kubie (1987). It also assists us in modeling tasks in walled arenas where the cue card (represented by two point landmarks) has been displaced or removed (Sharp et al., 1990; see "Place Fields Controlled by Entry Point").

OVERVIEW OF THE THEORY

The major components of our theory are shown in Figure 1. Visual perception provides the type, range, and egocentric bearing $\langle \tau_i, r_i, \theta_i \rangle$ of each landmark. The local view system computes allocentric bearings ϕ_i by adding in the current head direction Φ_h , available from the head direction system. The retinal angle, or bearing difference, between pairs of landmarks $\alpha_{ij} = \phi_i - \phi_j$ is also computed as part of the local view.

Position $\langle x^*, y^* \rangle$ is maintained by the path integrator in Cartesian coordinates.⁴ The path integrator refers to Φ_h to update its position estimate. It is assumed to receive information about the animal's movements via an efferent copy of motor commands. It may also use vestibular sensations or visual cues such as optic flow to estimate the magnitude of movements.

The place code in our theory is realized as a set of place units⁵ whose activity levels are products of Gaussian response functions tuned to $\langle r_i, \phi_i \rangle$ values, retinal angles α_{ij} , and path integrator coordinates $\langle x^*, y^* \rangle$. Because these units are tuned to enough visuospatial parameters to localize a point in space, they exhibit visually controlled place fields.

The role of place units in our theory is to maintain a consistent association between landmark configurations (local views) and path integrator coordinates. Either representation of location

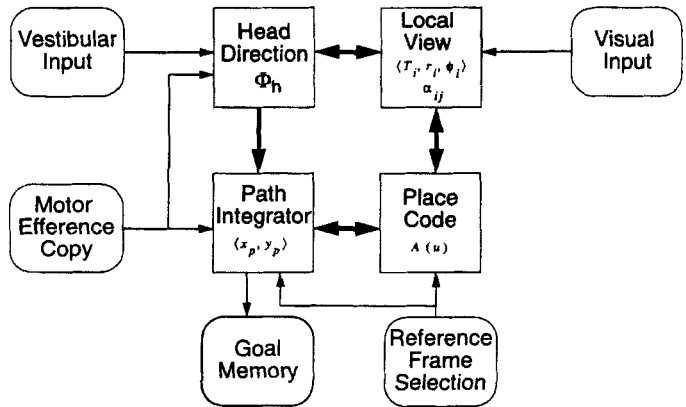


FIGURE 1. Components of our theory of rodent navigation.

may be used to reconstruct the other. In the remainder of this section we give an overview of how normal navigation is handled and discuss some broad issues in navigation before focusing on the details of the model.

Normal Navigation

During normal navigation through a familiar environment, rodents probably rely on path integration for short-term updating, because computing position by taking visual bearings to landmarks with each step would be too time consuming. That rodents have been observed to run at full speed into a barrier when it is placed across a familiar route (Carr and Watson, 1908) also suggests that they are not constantly attending to visual input. But at longer time scales drift in the path integrator will need to be corrected by perceptual feedback, perhaps when the animal has paused in its motion or is rearing and looking about.

When the animal does take visual bearings, the path integrator can contribute to the interpretation of visual cues by biasing place unit activity patterns based on the current rough position estimate. In other words, many place cells might match the bearings and distances of landmarks in the local view, but under the path integrator's influence, only those whose associated coordinate values are reasonably close to the current estimate would have a possibility of becoming active.

Consider a visually ambiguous environment such as an enclosed U-shaped maze (Fig. 2). Locations A and C in this maze should be represented by different place cell populations, but the visual cues available in the two locations are identical. If place cells are driven solely by visual inputs, both populations will become active. This could result in the perception of being in two places at once, or if the position estimate is the average of all votes, it might correspond to being in location B, which is not even in the maze. But since the two locations have distinct coordinates, our theory specifies that interaction between the metric (path integrator) and place code representations ensures that only hypotheses A and C will be considered, not their union or average.

Drift in the head direction system during navigation can be corrected visually in two ways. Landmarks at a substantial distance

⁴Path integration can also be done in polar coordinates, but Gallistel (1990) points out that this is computationally unstable, because each time the animal takes a step, the calculation of $\langle \Delta r, \Delta \phi \rangle$ depends on the present estimates of r and ϕ , whereas in a Cartesian system Δx and Δy are calculated independent of x and y . Errors cannot grow faster than linearly in the Cartesian system. A similar point is made in Mittelstaedt (1983).

We do not suggest that the path integrator in the rodent actually consists of a direct representation of Cartesian coordinates; instead we claim that the information represented by the path integration system is mathematically equivalent to a pair of Cartesian coordinates. Path integration in real biological systems is likely to be done using a distributed encoding that is not purely Cartesian or purely polar (see Touretzky et al., 1993, for an example). However, the independence requirements pointed out by Gallistel and Mittelstaedt still hold.

⁵We will always refer to theory components as "units," reserving the term "cells" for biological neurons.




FIGURE 2. U-shaped maze offers identical visual cues at locations A and C.

compared to the animal's normal range of travel will appear to have stable allocentric bearings. They can serve as "compass marks" that directly influence head direction cells. When compass marks are unavailable, remembered allocentric bearings of proximal landmarks may be used as directional cues. But since these values change with viewing position, the animal's location must first be accurately determined by settling on a consistent place code.

Other forms of sensory information may also be exploited by rodents as they move through the environment. When navigating a radial maze in the dark, drift in the head direction and path integration systems may be constrained by the narrowness and fixed orientations of the arms. The edges of the arms provide continuously available tactile cues to direction, whereas the ends of the arms provide unambiguous distance cues. In a circular arena, contact with the walls may be used to refine position estimates, but compass drift cannot be corrected this way, so error in the head direction system would continue to accumulate.

Reference Frames

The path integrator requires a stable reference point. The head direction system requires a stable reference orientation. We call the coordinate system defined by these quantities the *reference frame*. Neither the reference point nor reference orientation need be co-localized with an environmental cue, but they must be stable over time and space.

For a variety of reasons, we do not expect rodents to rely on a single global reference frame. As Kuipers et al. (1993) have noted, over large territories it would be difficult to maintain consistency of directional alignment and accuracy of metric information, as the same set of landmarks would not be visible from all locations. Their solution is a set of local coordinate systems linked together in some way. Worden (1992) has also presented a model of navigation linking multiple local views together.

We have suggested (Wan et al., 1994b; Touretzky et al., 1994) that the animal can maintain its position with respect to several simultaneously active reference frames. Place cells are tied to particular reference frames, and can only fire when one of these frames is active.⁶ If all of the reference frames with which a cell is associated are inactive, the cell will not fire. This aspect of our theory can explain task-specific fields as well as direction-tuned place cells. See "Changing Reference Frames" for details.

Trajectory Planning

In unobstructed environments such as open arenas, goal trajectories are simple vectors. Planning a path to the goal can be done by subtracting the present-position vector from the goal-position vector. Collett et al. (1986) report that gerbils can successfully execute a path to a goal even when the lights are extinguished shortly after they leave the start box, which suggests that they do maintain some sort of internal representation of their intended trajectory.

But in environments too large to support a global coordinate system, and in maze-like environments where paths contain multiple segments linked by heading changes, navigation becomes more complex. Some form of route learning is required, which is likely to involve additional brain areas. Our theory does not address route learning because of the relative lack of data on this phenomenon and the areas that facilitate it.

DETAILS OF THE MODEL

Level of Modeling

The place and head direction systems both have well-established anatomical correlates, and can in principle be modeled using connectionist or "neuron-like" units. The place system is indeed treated this way in our model: It is implemented as a population of up to 15,000 units, and their firing fields share key properties with real hippocampal place cells. Interactions among simultaneously active place units play an important role in determining the model's behavior.

The rodent head direction system is also sufficiently well understood that it could be modeled as a population of connectionist units, but we model head direction as a scalar quantity. Unlike place units, there is little to say at present about how interactions among units comprising a distributed representation of head direction would impact the animal's navigational behavior. There is much to say about how such interactions contribute to the maintenance of the head direction estimate itself (Skaggs et al., 1995a; Blair, 1996; Zhang, 1996), but these models do not address how the head direction estimate interacts with other systems (such as the place code) or how it drives behavior.

If a distributed head direction representation were used, the direction-dependent computations required by our theory, such as calculating allocentric bearings for landmarks, updating path integrator coordinates with self-motion, and updating the head direction estimate based on sightings of familiar landmarks, would require complex wiring schemes that cannot be grounded in anatomical data. Even though we do not model the head direction system as a population of neuron-like units, we can still make

⁶Hippocampal place cells show fields in multiple environments (Kubie and Ranck, 1983; Muller and Kubie, 1987; O'Keefe and Speakman, 1987; Thompson and Best, 1989). This implies that a single cell may be associated with multiple reference frames.

predictions about the results of neurophysiological experiments by predicting the value of the overall head direction estimate (see Predictions About the Head Direction System). We discuss the organization of the head direction system in considerably more detail in (Redish et al., 1996) and in a companion article to this one that is not concerned with modeling (Redish and Touretzky, 1996a).

Another crucial system, the path integrator, is also represented abstractly in our model. As with the head direction system, it would be possible to devise a distributed representation for path integrator coordinates and even invent machinery to implement the vector subtraction operation required for calculating a goal trajectory. In fact, we have done so previously (Touretzky et al., 1993). But the present model would not benefit from these details. When the anatomy and neurophysiology of the path integrator are better understood, comparable to the level of knowledge we now have about place cells, it will be profitable to move to a more biological representation to accommodate those data. In Redish and Touretzky (1996a) we develop a much more elaborate theory of path integration in which we suggest that this function is realized by a loop involving the hippocampus proper, the subiculum, the parasubiculum, and the superficial layers of entorhinal cortex. A computational model based on this proposal is planned for the future. For now, we model only the information represented in the path integrator: Coordinates in the current reference frame are represented in the model as vectors $\langle x, y \rangle$.

The model we describe here is a hybrid of connectionist and symbolic representations. Each subsystem is represented at a level of abstraction appropriate to the data being addressed, and the focus is on their interactions. For example, in some of the tasks we consider, the rodent uses its head direction sense plus visual cues to determine its position, whereas in other tasks visual cues and position information are used to determine head direction. What governs the direction in which information should flow? Furthermore, landmarks "out of place" sometimes cause a shift in the codes for both position and head direction, but in other situations displaced landmarks do not appear to influence these variables. What mechanism is sufficient to reproduce all these effects? Our model provides a working answer. Even though we model some subsystems at an abstract computational level, we can still make neurophysiological predictions, some of which are discussed below in Predictions.

Place Unit Activation Function

The activation level $A(u)$ of a place unit in our model is analogous to the firing rate of a neuron. It reflects the degree to which current visual inputs and path integrator coordinates match the values to which the unit is tuned. There are six features in the activation function: one set of coordinate values, two ranges and two allocentric bearings to landmarks, and one retinal angle (bearing difference) between a pair of landmarks. Each feature match is computed by a Gaussian function whose value ranges between 0 and 1.

$$A(u) = C(u) \cdot F_1(u) \cdot F_2(u) \cdot F_3(u) \cdot F_4(u) \cdot RA(u) \quad (1)$$

The activation level defined in Equation 1 is the product of these six feature match functions. The functions have variable gains (defined as $1/\sigma^2$ in the equations below) which become negligible when the feature values the function is trying to match are unavailable. For example, if visual information is removed, as when navigating in the dark, the gains of the five landmark-dependent functions are decreased to near 0 (by making the σ terms large), bringing the value of the exponential close to 1. Effectively, these features have dropped out of the activation equation, leaving the place unit to be controlled by just the match to path integrator coordinates. Conversely, when the animal first enters an environment in the light, it can see landmarks but will not know its path integrator coordinates, so the gain on that feature is minimized (large σ_c) until a coordinate estimate is obtained by self-localization, described in Self-Localization. Having terms drop out of the activation equation by reducing their gain allows us to distinguish a match *failure* (output approaches 0) from an inability to perform the match at all (0 gain; output close to 1).

Equation 1 should not be taken literally as a model of a pyramidal cell. Since we have not yet deciphered the rodent's representation of visual scenes, or any other aspect of the hippocampal input code, the computation performed by a single hippocampal place cell is difficult to specify. But it is probably safe to say that the six-term multiplication and the gain modulation mechanism just described are not realized by a synaptic integration process within one cell's dendritic tree. There are many layers of processing between visual receptors and place cells where parts of these tasks might be performed. Gain modulation for visual feature matching, for example, could take place in parietal cortex. Furthermore, conjunctions of feature matches need not be computed at the neural level by anything as simple as scalar multiplication. Many types of nonlinear, distributed operators might be suitable. The important point for the present model is that certain types of information must be combined in the hippocampal system in order to account for the full range of observed place cell responses.

Returning to the details of our activation equation, the first term, $C(u)$, is a Gaussian function tuned to the path integrator coordinates $\langle x_u, y_u \rangle$ associated with unit u . This function reaches its maximum value when the path integrator's position estimate $\langle x^*, y^* \rangle$ matches the coordinates to which the unit is tuned. The parameter σ_c controls the width of the Gaussian, as shown in Equation 2. The actual parameter values used in the simulations are given in the appendix.

$$C(u) = \exp\left(-\frac{\|\langle x_u, y_u \rangle - \langle x^*, y^* \rangle\|^2}{\sigma_c^2}\right) \quad (2)$$

Each F_i term in Equation 1 is a landmark feature from the following set: r^τ = distance to landmark of type τ , ϕ^τ = allocentric bearing to a landmark of type τ . If feature F_i uses a point landmark of type τ , then the feature value r_i (or ϕ_i) is compared to the distance (or bearing) from the animal to each visible landmark of type τ . If F_i uses a surface landmark,⁷ then the feature value is compared to the distance or bearing component of the

normal vector from the animal to the surface. Note that since ϕ terms are allocentric bearings, the place cell's activation level does not depend on the animal's heading. (The directionality of place cells in some environments is explained by another mechanism, reference frame selection, discussed in Changing Reference Frames). The parameters σ_r and σ_ϕ control the width of the Gaussians for r -type and ϕ -type F_i terms, respectively.

$$F_i(u) = \max_k \left\{ \begin{array}{l} \exp\left(-\frac{(r_i(u) - r_k)^2}{\sigma_r^2}\right) \text{ } r\text{-type feature} \\ \exp\left(-\frac{(\phi_i(u) - \phi_k)^2}{\sigma_\phi^2}\right) \text{ } \phi\text{-type feature} \end{array} \right\} \quad (3)$$

The $\max(k)$ selector of Equation 3 simulates the effect of an attentional mechanism that keeps landmarks distinct. It assures that when multiple landmarks of type τ are present, we choose the one that best matches the value to which the feature is tuned. If this selector were replaced by a simple summation, a feature that matched one object strongly could not be distinguished from a feature that matched many objects poorly: an obviously undesirable outcome. The neural implementation of this attentional process would presumably involve some type of inhibitory interactions between cells coding for competing objects, but the details do not concern us here.

The final term in the activation equation, $RA(u)$, is a Gaussian function tuned to the retinal angle $\alpha_{ij} = (\theta_i - \theta_j) = (\phi_i - \phi_j)$ between two landmarks. The parameter σ_α defines the width of the Gaussian tuning of $RA(u)$.

$$RA(u) = \max_{i,j} \left(\exp\left(-\frac{(\alpha(u) - [\theta_i - \theta_j])^2}{\sigma_\alpha^2}\right) \right) \quad (4)$$

The RA term is partly redundant with the distance and bearing features, but has its own distinct character. For example, although retinal angle information is implicit in the allocentric bearing features ϕ_i and ϕ_j , the latter are unavailable when head direction is unknown. Thus, retinal angle is especially useful for self-localization after rotational disorientation. Distance information is useful then too, but the distances r_i and r_j to two landmarks are not enough to uniquely localize the animal; there is a mirror image location, on the opposite side of the line joining the two landmarks, with identical distance values. This location will have a complementary retinal angle value, though, since the left landmark viewed from the original location will be the right landmark when viewed from the mirror image location. (This assumes the landmarks are distinguishable.) Another important point about the RA term is that retinal angle values change very rapidly, making the RA term very sensitive to position, when the animal is moving in the area between the landmarks, as when near the goal in the Collett et al. tasks. Our formula for r -type features does not show increased sensitivity for small r_i values, although

this could perhaps be modeled by making σ_r a function of $r_i(u)$. The RA term provides a more direct way of expressing the observation that, at least in humans, the angles between nearby landmarks are a highly salient feature of the local view.

The $\max(i,j)$ selector in Equation 4 evokes not only an attentional mechanism, but also a combinatorial search, since it selects the pair of landmarks that together give the best match to the angle α_{ij} . As expressed here this appears to involve $n(n-1)$ distinct computations, but less than $2n$ are required,⁸ and as a practical matter one would not choose landmarks of type τ to construct an RA feature when there were large numbers of indistinguishable landmarks of that type present. So a neural implementation of the RA term appears feasible.

The place field of a unit is defined as that portion of the environment for which the activation function $A(u)$ exceeds some threshold δ_A . Because all of the equations above decrease monotonically as the input values depart from the tuned values, place fields will be compact and continuous. Equations 1 through 4 predict that place fields can change their shapes when landmarks are slightly perturbed, but due to the way missing terms drop out of the product in the activation function, they should not change much when the lights go out or individual landmarks are removed.

Self-Localization

We now explain the process by which the model determines its coordinates based on visual cues. In most cases when an animal is deposited in an arena, path integrator coordinates are initially unavailable. The $C(u)$ term has low gain and thus little effect on Equation 1. Place units become active based only on how well their landmark tuning matches the current local view.

Because environments can be visually ambiguous (cf. the U-shaped maze in Fig. 2) and individual units sample only a few features, there is no guarantee that the initial place code will be *coherent*, i.e., that active place units will have highly overlapping place fields. If the code is incoherent, the active units disagree on the animal's location.⁹ In general, a visually derived place code cannot be guaranteed coherent unless every landmark is unique.

The model achieves a coherent place code by *parallel relaxation*.¹⁰ Each place unit has coordinates $\langle x_u, y_u \rangle$ associated with it. This is the center of the $C(u)$ term. The path integrator's location estimate tracks the weighted mean of the $\langle x_u, y_u \rangle$ values for

⁷We require the simulated animal to be within 10 cm of a surface (i.e., the arena wall) when tuning a place unit to a surface-type landmark.

⁸Order the landmarks from left to right as they appear in the visual field. Let i^*, j^* be the indices of the highest scoring pair found so far, and let i, j be the indices of the pair currently being examined, with $j \geq i$. As index j increases, index i must either stay the same or increase. Since with each comparison at least one of the two indices will increase, there must be fewer than $2n$ comparisons.

⁹*Coherency* can be defined as the inverse of the confidence interval of the location represented by the place units. The confidence interval can be measured by the standard "bootstrap" algorithm (Efron, 1982). Coherency is also inversely proportional to the variance in the distribution of the centers of the place fields of the active place cells.

all active units u whose activation level exceeds δ_A . Coordinate values are weighted by the activation levels of the units. As this estimate changes, so does the value of the $C(u)$ feature, so every unit continually recalculates its activation level. Units whose place field centers are too distant from the mean coordinate value may become inactive due to a poor match in the $C(u)$ term, whereas some inactive units may become active if their $C(u)$ value improves. As the relaxation proceeds, the initially low gain of the $C(u)$ term gradually increases toward its normal value. After just a few iterations, the place code converges to a coherent state in which all active units have highly overlapped place fields, forming a tight cluster as shown in Figure 3.

Muller et al. (1991) remark that because all place fields have centers in the arena, as the animal approaches the wall, a calculation of mean position based on the active place cells will show a uniform error away from the wall. Wilson and McNaughton (1993) claim that the place code provides an extremely accurate estimate of position (within 1 cm, using 141 simultaneously recorded place cells), but they used a different algorithm for determining the position represented and did not specifically examine error near the wall. In our simulations, we have not seen this uniform error. We believe this is because some of our place units are surface-type (tuned to the arena wall; see Place Unit Activation Function), and show crescent-shaped fields (see "Basic Place Fields").

Re-Orienting

We now consider versions of the self-localization problem where the animal's head direction estimate is invalid. There are two cases. If the animal lacks a heading estimate altogether, e.g., because the experimenter disoriented it by vigorous rotation before placing it in the arena, then it can ignore its internal compass temporarily. In the model, when both an initial heading and compass points are unavailable, the gain of the ϕ -type terms is reduced and place unit activations are computed using just distances and retinal angles. After the simulated animal's position has been determined, its compass is realigned as described below.

The more difficult case is where the animal has a usable heading estimate, but the value is wrong. Many experimenters disorient their animals by gently carrying them over a circuitous path for several minutes in an opaque box before placing them in the arena. Compass errors accumulate over this path because the compass is being updated solely from vestibular input. There is no motor feedback, no tactile information (e.g., from whiskers brushing the floor), and no visual cues such as optic flow that could be used to estimate the amount of each heading change. Nor are landmarks available that could be used to correct for drift. In this




FIGURE 3. Enforcing coherency in the place code when the animal is introduced into a familiar environment. Large solid circles are landmarks. Each small dot marks the center of a place field defined by the path integrator coordinates of an active unit. Top: Incoherent initial state: active units are widely dispersed. Bottom: Coherent final state.

¹⁰Anderson and Hinton (1981) explain parallel relaxation as "a constraint-satisfaction paradigm in which some input data must be given an interpretation that simultaneously satisfies a large set of local constraints. This interpretation corresponds to a pattern of activity over the units, and it is found by an interactive computation in which each unit affects many other units until the whole system settles down into a stable state."

case, when released into the arena the animal does not realize its heading estimate is in error. It will encounter difficulty reconciling its internal heading estimate with the bearings of whatever landmarks it recognizes. In our model, a measure of the *consistency* between local view and path integrator representations de-

termines whether the head direction estimate is in reasonable agreement with the bearings of familiar landmarks.

The consistency κ of the association between coordinates and local view maintained in the place code is given by

$$\kappa = \frac{\sum_u A(u)}{\sum_u C(u)}. \quad (5)$$

Because $A(u) < C(u)$ for all place units, κ ranges between 0 and 1. When the current local view is compatible with the local view predicted by the current path integrator coordinates, κ will be high. When the two are not compatible, κ will be low. In the latter case, the head direction estimate may be invalid.

The procedure for aligning head direction with the environment is as follows. Every place unit has associated with it the learned allocentric bearings of two landmarks: the two ϕ -type F_i terms. For each active unit u , the model computes the difference between the learned allocentric bearing $\phi_i(u)$ and the egocentric bearing θ_j of whichever landmark it is presently using to compute $F_i(u)$. This gives a set of roughly twice as many estimates of the animal's heading as the number of active place units. One peak in this distribution is chosen (stochastically, based on relative peak heights) to serve as the new heading estimate Φ_h .

Aligning Internal Representations With Familiar Cues

To complete the model, we must specify how the self-localization and re-orienting processes should interact to align the model's internal representations with visual cues. The details are constrained by experiments examining the effects of cue manipulations (Suzuki et al., 1980; Collett et al., 1986), disorientation (Cheng, 1986; Margules and Gallistel, 1988; Gallistel, 1990), and use of a consistent entry point (Sharp et al., 1990).

If the animal has initial estimates of its position and head direction when it enters a familiar environment, there are four possible outcomes:

1. Keep the initial estimates.
2. Reset head direction (re-orient).
3. Reset path integrator coordinates (self-localize).
4. Reset both head direction and path integrator coordinates.

These four choices can be prioritized by the Sharp et al. (1990) experiments, which we simulate in Figures 7 and 8.

¹¹Not all of the place fields in (Sharp et al., 1990) showed such a clean single place field when the second cue card was added. Five of the 18 cells changed their place fields over the various recording sessions, including showing paired fields during some sessions. Our model does not reproduce these effects, since it does not address changes in place cell tuning. Such changes could occur due to variations in the hippocampal input from one trial to the next as the animal's short-term memory accumulates new sensations and loses old ones. Or they could result from learning across trials. Although it would seem that place cells with paired fields would be incompatible with our model, we can explain them as cells not tuned to path integrator coordinates. The model requires that these cells be rare, which they do seem to be.

When rats are trained in a gray-walled cylinder with a white cue card subtending 90°, and then tested with an additional card opposite the original (a symmetrical environment), most (13/18) place fields do not double¹¹ (Sharp et al., 1990). Instead, the cells continue to fire at the original location. If the cards are visually indistinguishable, then the place fields must be controlled by something in addition to visual input. In our model, the extra factor is the path integrator.

The animals in this experiment were not disoriented before entering the arena, and they always entered at the same location: the northwest corner. Thus, when a well-trained animal enters the arena, it does not need to reset its path integrator; it already knows its path integrator coordinates. This stable entry point can serve as the origin of the reference frame for the arena.

Sharp et al. also tested the animals on the two-card arena by placing them at the opposite (southeast) corner of the arena. In this case the place fields rotated by 180 degrees. In our simulations, when the sim-animal¹² enters at the southeast corner, neither cue card is at the correct allocentric bearing associated with the initial coordinate estimate. The consistency in the place code association is therefore low, but it can be improved by resetting head direction as described in Re-Orienting. In this case, the only head direction compatible with the current local view (including the normal to the arena wall as one of the features) is 180° opposite of the initial orientation estimate. The sim-animal therefore resets its head direction, but not its path integrator coordinates, causing place fields to appear rotated by 180°.

Sharp et al. also tested their animals with a southeast entry point and a single cue card in the original position. In this case the place fields did not rotate. Our model reproduces this result as follows: Head direction reset will not improve the consistency of the place code. So the model resets its path integrator coordinates instead, by performing a self-localization operation. Thus, the sim-animal correctly determines that it is in the southeast corner of the arena.

When tested with cue cards rotated by $\pm 30^\circ$, Sharp et al. observed that place field locations were controlled by an interaction between choice of entry point and cue card positions. When the animal entered from the southeast with two cue cards at 30° and 210° (180° + 30°), the place field was rotated by 210°. Similar results were found for cue card rotations of -30° .

The relationship between head direction reset and path integrator reset can now be stated. In the rules below, κ^* is a constant whose value is determined empirically.

1. If initial position and head direction estimates prove correct by producing a good match to the local view, meaning the consistency κ is high, they are retained.
2. If realigning head direction improves the consistency by at least κ^* , the new head direction value should replace the initial estimate.
3. If self-localization based on visual input (using head direction to determine allocentric bearing in the local view representation)

¹²When necessary to differentiate our simulations from real animals, we refer to our simulation as a "sim-animal."

increases the consistency by more than κ^* , then the new path integrator coordinates (with the original head direction estimate) should replace the current position estimate.

4. If self-localization based on just visual input (head direction ignored, so no ϕ -type terms) followed by realignment of the head direction system improves consistency by at least κ^* more than the best value obtained in any of the other cases, then both the position and heading estimates should be changed.

Note that this prioritization predicts that given a choice between resetting either head direction or path integrator coordinates to achieve a consistent place code, as in the bottom right portion of Figure 7, the animal will reset its head direction, causing place fields to rotate.

CORRESPONDENCE WITH EXPERIMENTAL DATA

Reproducing Single-Cell Recording Experiments

Basic place fields

As we noted at the beginning of the article, cells in hippocampal areas CA3 and CA1 show firing fields correlated with the rodent's location. Any hippocampal navigation model must reproduce this fundamental result. Because the activation of place units in our model is controlled by Gaussian functions tuned to spatial features, these units will show compact, continuous place fields. We have tuned the parameters of our model to produce fields such as Figure 4 that resemble those seen by Muller et al. (1987) in rats. We model the environment of Muller et al. as a circular arena (radius 38 cm) with three landmarks: the arena wall (a surface landmark) and two distinguishable point landmarks: a "gray-to-white" and a "white-to-gray" transition, representing the left and right edges of a white cue card affixed to the wall. The cue-card subtends 90° of arc.

Crescent-shaped fields. Because some of our place units are tuned to surfaces, in an arena with perceptually prominent walls, some units will be tuned to distance and allocentric bearing to




FIGURE 4. A typical simulated place field in a small cylindrical arena, 38 cm radius, with a cue card along the east segment of the wall. Place cell activity ranges from 0.0 (white) to 1.0 (black).




FIGURE 5. A typical wall-tuned place field in the cylindrical arena with an east cue card. Scales as in Figure 4.

the arena wall. Such a unit will show a crescent-shaped field like those seen by Muller et al. (1987). Other models (such as Sharp, 1991) also include fields close to the wall, but they only produce convex fields. The crescent-shaped fields reported by Muller et al. are concave. Our model produces concave crescent-shaped fields, such as Figure 5.

Place fields in manipulated environments

Place fields tied to local landmarks. A number of investigators (Muller and Kubie, 1987; O'Keefe and Speakman, 1987; McNaughton et al., 1994; Knierim et al., 1995) have reported that place fields rotate with the rotation of prominent visual landmarks. Any model in which place cells are driven by visual inputs will reproduce this result. Our model does too (Fig. 6).

Taube et al. report that head direction cells [recorded from postsubiculum (Taube et al., 1990b) and anterior thalamic nuclei (Taube, 1995)] rotate their preferred orientation with the rotation of a prominent visual landmark. Knierim et al. (1995) and McNaughton et al. (1994) report that when head direction cells (also from anterior thalamic nuclei) are recorded simultaneously with place cells, the preferred directions of head direction cells and the locations of place fields rotate in correspondence. In our model, the interactions between the head direction and place code modules serve to keep these two representations in synchrony.

Moving vs. removing landmarks. Muller and Kubie (1987) and O'Keefe and Speakman (1987) found that place cells continue to show compact fields when visual landmarks are removed. Muller and Kubie (1987) removed the cue card from the cylindrical environment described earlier and found that although place fields sometimes rotated around the center of the arena, they did not change shape. O'Keefe and Speakman (1987) trained rats on a plus maze and then removed the landmarks. They found that place cells were not disrupted by this.

In our model, if the feature term F_i of a place unit is tuned to a landmark type τ_b , but no landmark of type τ_i is visible, then the gain of τ_i feature-matching functions is decreased, and that term drops out of the activation function (Equation 1). If all the visual feature terms are missing, the cell will be driven solely by the path integrator term, $C(u)$.




FIGURE 6. a–d. The place fields (and head direction) follow the cue card. Scales as in Figure 4.

However, O'Keefe and Conway (1978) and Pico et al. (1985) found that while removing some of the visual cues did not disrupt place fields, removing all of the cues did.

Our model is compatible with all four of these experiments. Because the animals in the Muller and Kubie experiment have the arena wall as a remaining cue, the shape of the place fields will not change, but because the arena wall is a circularly symmetric cue, the animals' head direction representation may drift. This would rotate the believed allocentric bearing to the wall, which would rotate the place fields around the center of the environment.

The difference between the O'Keefe and Speakman and the two latter experiments is that in the former experiment, the animals were first placed in the environment with the cues present. This would allow the system to reset the path integrator before the cues were removed. Even with some of the cues missing, the system would then use path integration to drive the place cells as the animal traversed the environment. On the other hand, in the latter experiments, the animals do not have the initial view of the cues to reset the path integrator and they may have to reexplore the environment. A similar effect was seen by Quirk et al. (1990) for animals navigating in the dark.

Place fields in the dark

One of the major observations that our model is among the first to simulate is the continued firing of place cells in the dark (O'Keefe, 1976; McNaughton et al., 1989; Quirk et al., 1990; Markus et al., 1994). According to our theory, because place units can be driven by the path integrator when visual cues are unavailable, place units will continue to fire in the dark, and maintain their usual fields. However, because the head direction sys-

tem can drift (Knierim et al., 1993), and the place code and head direction systems are always in synchrony (Knierim et al., 1995), place fields in the dark may drift, or may not always lie in the same orientation as in the light. This is exactly what is seen by Quirk et al. (1990).

Quirk et al. report that when the animal is placed into the arena in the light, and then the lights are extinguished, the place fields rarely change. If they do change, they continue to show the same shape and distance from the arena wall; it is their orientation that varies. On the other hand, if the animal is placed into the arena in the dark, the place fields often change radically. If the lights are then turned on, the fields sometimes snap back to their original location and sometimes not. Our theory says that when the animal is placed in the arena with the lights on, it can self-localize and set its path integrator correctly. When the lights are extinguished, the path integrator can drive the place code, and drift in the path integrator or head direction system may produce drift in the place field orientation. On the other hand, if the animal is placed in the arena in the dark, it cannot set its path integrator correctly, so the place fields may change dramatically.

Place fields controlled by entry point

As discussed in Aligning Internal Representations with Familiar Cues, Sharp et al. (1990) trained rats to find food scattered on the floor of a gray-walled cylinder with a white cue card subtending 90° after entering at a consistent entry point (the northwest corner), and then tested them with a variety of cue and entry point manipulations.

Figures 7 and 8 show a sample place field from our model for each of the conditions discussed: entering at northwest or southeast with one cue card at 0° (top of Fig. 7), two cue cards at 0° and 180° (bottom of Fig. 7), two cue cards at 330° and 150° (top of Fig. 8), and one cue card at 180° (bottom of Fig. 8). The model is compatible with all conditions tested in Sharp et al. (1990); the results are compared in Table 1.

When the animal entered from the NW and was given only one cue card at 180°, Sharp et al. report that the place field did not rotate. But in our model the place field does rotate, since place fields normally follow the cue card (see Fig. 6). The discrepancy might be explained by the fact that this particular manipulation was the last one in the sequence. McNaughton et al. (1994) and Knierim et al. (1995) have shown that if rats experience the cue card moving over a number of sessions, they eventually learn to ignore it and it loses control over the place fields. (Knierim et al. actually left the card in place, but rotationally disoriented the animals before placing them in the arena; we assume that moving the card would have a similar effect on Sharp et al.'s animals, which were not disoriented.) When we tested our model without a cue card, equivalent to a card being present but ignored, it produced a place field somewhat more diffuse than normal, but with no rotation (bottom-left of Fig. 8). The diffuseness could be remedied by a more sophisticated activation function that adjusted the gains on the remaining features to compensate for cue card-related F_i terms dropping out.

Muller et al. (1991) explain the results in (Sharp et al., 1990)




FIGURE 7. Simulation of (Sharp et al., 1990): Place fields controlled by both cue cards and entry point.

by treating the entry point as a virtual landmark. The animal then has three sources of information to judge its location: the cue card, the animal's position relative to its entry point, and static background cues. A majority voting scheme could discard whichever source was at odds with the other two. This has aspects in common with our treatment, but we use the consistency of the entry point during training only as a justification for initializing the path integrator to a known value. See Path Integration and Virtual Landmarks for a discussion of the relationship between virtual landmarks and path integration.

Changing reference frames

Markus et al. (1995) have reported task-specific cells, that is place cells that are active in one task but not in another, in the same environment. Similarly, Eichenbaum et al. (1987) have reported cells that show activity during certain aspects of a task, but not during other aspects, even while the rat traverses the same area of space. We have suggested that the reference frame is goal related, and therefore during different tasks (or different aspects of a task) the animal will activate different reference frames (Wan et al., 1994b). Changing the set of active reference frames will silence some place cells and enable others. A similar proposal was recently made by McNaughton et al. (1996).

Place cells show direction-dependent firing fields only in some specific types of tasks, particularly those with restricted movement paths. For example, place fields tend to be directional on a radial eight-arm maze (McNaughton et al., 1983), whereas directional place fields are rarely seen in an open arena (Muller et al., 1991,

1994). In the latter case, a task requiring restricted movement paths produces more directional place fields than one with unrestricted paths (Markus et al., 1995).

Because directional place fields are predominantly seen in tasks with restricted trajectories, Sharp (1991) suggested that the directionality results from a competitive learning process that sharpens the natural separation (in pattern space) between the views seen in one trajectory and those seen in its complement. But Markus et al. (1995) counters that directional place fields are also seen on the central dais of the radial maze, even though views from all directions are available there. Contrary to this result, Muller et al. (1994) report nondirectional fields on the central dais, and also some nondirectional fields on the arms.

We suggest that directionality is a consequence of changing reference frames. When the animal traverses a restricted path, it can define reference points at each end of the path, each of which serves as the origin for a reference frame. Repeated travel over the path may lead to indifference to one of the reference frames. If that frame becomes inactive and the path integrator no longer tracks the animal's position with respect to it, place cells tied only to that reference frame would be disabled. When the animal travels in the opposite direction, the frame will be active, and place cells tied to that frame will fire when the animal enters their fields.




FIGURE 8. Simulation of (Sharp et al., 1990): Place fields controlled by both cue cards and entry point. See text for discussion of lower left figure [NW, cue card at 180°].

TABLE 1.

*Comparison of Data From Sharp et al. (1990) With Our Model**

Entry point	Cue cards	Angle of rotation (Sharp et al., 1990)	HD precession (our model)
NW	0°	2.7°	0°
SE	0°	-6.0°	2°
NW	0° + 180°	-2.3°	0°
SE	0° + 180°	182.5°	178°
NW	150° + 330°	Not done	331°
SE	150° + 330°	158.3°	151°
NW	180°	-5.5°	188°
NW	None		0°
SE	180°	182.2°	179°

*Angle of rotation is the angle at which the place fields in the trial in question are maximally correlated with the corresponding place field in the training sessions; values taken from Sharp et al. (1990). HD precession is the precession of Φ_h in our model. Because the head direction and place code are coupled in our model, the precession of Φ_h indicates the rotation of the place code. The "NW, no cue card" is our simulation of the rat ignoring the cue card; see text.

Thus, the cells appear to be direction sensitive when in reality they are reference frame sensitive. Markus et al. (1995) have made a similar suggestion: that directionality is caused by the rat attending to only one end of the restricted path.

Our theory makes no prediction about the directionality of place cells on the central dais of the radial maze because we have not specified mechanisms for constructing new reference frames and switching between them. The only claim is that place cells are inherently nondirectional (which is why our F_i features are tuned to allocentric and not egocentric landmark bearings), and that directionality is an artifact of reference frame switching. The fact that some nondirectional cells are observed even on tasks where the majority are directional (Muller et al., 1994; Markus et al., 1995) poses an interesting problem which a more fully developed theory of reference frame selection will need to address.

One possibility is that multiple reference frames can be simultaneously active. If an animal performing an overtrained task begins attending to additional reference frames, perhaps due to increased arousal in response to an unexpected stimulus, then a sudden loss of directionality of place cells should be observed. However, the ability to track multiple reference frames simultaneously is unproven, and is not crucial to our theory.

Another interesting aspect of the reference frame selection hypothesis is that it allows us to explain both directionality of place cells and task-related place cells using the same mechanism. Eichenbaum et al. (1987) report cells tuned to place and behavioral correlates: after sniffing the stimulus at an odor release port, the animals would either receive a reward (odors S^+) or not (odors S^-). Independent of the particular odor used, they found that certain hippocampal cells only fired in a certain place if the animal had received an S^+ odor. Our theory suggests that upon receiving an S^+ odor, the animals attend to a reference frame that includes the reward location. Place cells associated with this reference frame would be enabled by the S^+ odor.

Bostock et al. (1991) recorded from place cells in an open cylindrical arena, first with a white cue card, and then with a black cue card. Sometimes the place fields were similar and sometimes they were unrelated (as if the two situations were encoded as different environments). However, once a place field changed dramatically between the two cue cards, then (1) all other place fields recorded from the same animal changed dramatically when the environment included the black cue card, and (2) when the white cue card was returned the place fields returned to their original configuration. Similar effects have been reported by Quirk et al. (1990) and Sharp et al. (1995). We suggest that at some point the animals began to encode the black cue card trials with a new reference frame.

Reproducing Simple Behavioral Experiments

We now turn to behavioral experiments. Because there is no overt behavior in our simulations, we look at unit responses and infer behavioral correlates. We will make clear in each case the analogy between our simulation results and the corresponding animal behavior.

Radial maze

Olton and Samuelson (1976) introduced the eight-arm radial maze as a spatial working memory task. Each arm of the maze was baited. Rats were allowed to visit three of the arms, then a delay was imposed, after which the rats could enter any arm. Well-trained rats skipped the three arms previously visited and chose each of the five remaining arms once.

At the end of each arm was a distinctive landmark. Suzuki et al. (1980) modified this task by manipulating the landmark locations. If, during the delay, the landmark configuration was rotated by 180°, the rats visited the five arms defined by the rotation (mean = 4.0 correct rotated choices; compare with mean =

4.3 correct choices in original configuration). If, however, the landmarks were permuted in a manner that did not correspond to any rotation, the rats were impaired in choosing unvisited arms (mean = 2.9 correct choices, chance level would be 3.13 correct choices).

This experiment demonstrates an interaction between head direction and place code. Because place units tune to allocentric bearings, changing the landmark configuration produces a conflict between the head direction predicted by the landmarks and the value in the head direction module. When the configuration was rotated by 180°, head direction could be reset to resolve the conflict, but when the configuration was permuted the conflicting information from different landmarks could not be reconciled.

To replicate this experiment, we assume that the animal has generated a place code for the environment with the original stimulus configuration, associating positions with allocentric landmark bearings based on accurate heading information. For simplicity, we place the reference point at the center of the arena. We model the arena as a cylinder (radius 65.5 cm) with seven distinct landmarks¹³ placed at seven of eight equally spaced points around the circle, at a distance of 100 cm from the center. See Figure 9.

We model the rotation result by demonstrating (1) that the simulation chooses to reset its head direction in the rotated case to 180° off normal and (2) that the place code association shows a high consistency after rotation by 180°. We model the permutation result by demonstrating that the consistency of the place code association is two orders of magnitude less in the permuted case. Table 2 shows the precession of the head direction chosen by the simulation upon entering the environment at the center of the arena and the consistency of the place code association based on that head direction for each of the three cases: (A) original stimulus configuration, (B) rotated configuration, and (C) permuted configuration.

Rectangular arena

Our theory makes an important distinction between disoriented and non-disoriented animals. This distinction has been demonstrated behaviorally in rats by Cheng (1986) and Margules and Gallistel (1988) (see also Gallistel, 1990). Cheng tried to train rats to find food at one corner of a rectangular arena. In order to make the corners as distinct as possible, he placed a panel at each one, covered with a different type of material. In addition, the panels had from zero to three pinholes through which light was visible, and two of the panels had unique odorants placed behind them. Cheng disoriented the rats before placing them in the arena at a random location, and found that although they were able to distinguish one pair of diagonally opposed corners from the other, they could not differentiate the two corners in each pair. Cheng reports that the animals chose the correct corner in approximately




FIGURE 9. Cue manipulations used in the Suzuki et al. (1980) radial maze experiment.

50% of the trials, and, in the other 50%, they chose the corner opposite it. This suggested that the rats were sensitive only to the geometric structure of the environment, and were ignoring other types of cues that could distinguish between the two opposite corners.

Margules and Gallistel replicated the experiment without disorientation and found that most animals had no difficulty selecting the correct corner over 75% of the time; some achieved better than 90% success rates.

We are unable to explain why rats ignore the odor, texture, and non-geometric visual cues, but we are able to model the effect of rotational disorientation on the ability to navigate with just ambiguous geometric information. In our model, the corners of the arena are modeled as identical point-type landmarks. If the sim-animal is disoriented (meaning a head direction estimate is not available), the ϕ -type F_i terms drop out of the activation equation (Equation 1). Because the sim-animal is placed into the environment at a random location, it also lacks an initial estimate of its path integrator coordinates, so the C term drops out as well. The only terms left with which the sim-animal can self-localize are r -type F_i terms and the RA term, both of which are rotationally symmetric in this environment. Therefore, the sim-animal has a 50% chance of localizing itself at the correct spot, and a 50% chance of choosing the mirror image location. When asked to calculate a trajectory to the goal, the sim-animal will choose the opposite of the goal corner 50% of the time. Figure 10a shows a histogram of the goal predictions produced when the model was disoriented.

TABLE 2. Simulation Results for Suzuki et al. (1980) Radial Maze Experiment

Configuration	Precession	Consistency
Original	0°	0.250
Rotated	182°	0.246
Permuted	0°	0.004

¹³In Suzuki et al.'s experiment (Suzuki et al., 1980), the eighth arm had no landmark.



FIGURE 10. Distribution of the simulation's goal predictions in each quadrant of a rectangular environment, (a) with disorientation and (b) without.

When the sim-animal is not disoriented, allocentric bearing information is available and the ϕ -type F_i terms do not drop out of the activation equation. The sim-animal therefore has enough information to differentiate the correct corner from its 180° opposite, and it always chooses the correct corner. Figure 10b shows a histogram of the goal predictions in the non-disoriented case.

Navigation using landmark arrays

Collett et al. (1986) report experiments in which gerbils were trained to find a food reward (a sunflower seed buried in gravel) among cylindrical landmarks in an otherwise impoverished environment. The landmarks were translated but not rotated from trial to trial, and they maintained the same orientation and distance from each other, thus forming a stable *landmark array*.

During training, both the gerbils' starting location and the location of the landmark array varied randomly; however, the food was always located at the same position relative to the array, and the start box always had the same orientation. Probe trials were run without a food reward, and the time spent searching was histogrammed. The left sides of Figures 11–20 show 2D histograms of search effort in which larger blobs correspond to more search time.

Our simulations of these tasks start with the simulated animal deposited at a random location in the environment; thus, initially it has no knowledge of its coordinates, but its sense of direction is correct. We assume the animal is already familiar with the task, i.e., it has place codes for all locations in the training environment; see Appendix A for details. The sim-animal uses visual cues to activate an initial place code and self-localize as described earlier. Once it has determined its coordinates, it estimates the path from its present location to the goal by vector subtraction.

We cannot measure search time directly in our simulations because there is no searching behavior. Therefore the distribution of goal estimates for 100 probe trials is what is plotted in the right halves of Figures 11–20. These simulations are discussed in more detail in Redish and Touretzky (1996b).

One-landmark experiment. Collett et al. (1986) trained one set of gerbils to find food placed near a single circularly symmetric landmark (Fig. 11a). The fact that the animals could learn to search at the correct bearing as well as distance implies that they have some independent means of determining bearing information. Collett et al. supposed that the animals were utilizing some external cue, despite their attempts to block this. The subsequent discovery of head direction cells suggests that this might not have




FIGURE 11. a: Histogram of search time of gerbils on the one-landmark task. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions by the simulation.

been the case. In our model, head direction is used to derive allocentric bearing information, and this plus distance is sufficient to localize the animal in the one-landmark case (Fig. 11b.)

Two-landmark experiments. Collett et al. (1986) also trained animals to find food at a location specified by a pair of landmarks, as shown in the search time distribution in Figure 12a. Our simulation produces a comparable result: a distribution of goal estimates with a single peak, centered at the correct location (Fig. 12b).

With two landmarks, the array can be manipulated to provide ambiguous or inconsistent cues. For example, one landmark can be removed. In this case, Collett et al. (1986) report that the gerbils searched alternately in two locations, each at the correct distance and bearing from one of the landmarks they had observed during training (Fig. 13). Collett et al. theorized that the animals were matching the landmark first to one and then to the other of the two they remembered.

In our model, when the animal first enters the environment the place code is derived solely from the local view. In the two-landmark task, the single landmark visible on the probe trial is insufficient to constrain the place code to a single location. Therefore, parallel relaxation produces a bi-modal distribution of goal predictions with peaks centered at the correct angle and distance from each of the two landmarks. See Figure 3 and Self-Localization for details.

The two-landmark array can also be stretched. When trained on two landmarks and tested with the distance between them doubled, the gerbils again searched at two locations, each at the correct distance and bearing from the corresponding landmark (Fig. 14a). They did not search at the two exterior locations that would be derived by matching the perceived west landmark against the learned east one and vice versa. Our simulation also concentrates its goal predictions at the interior locations (Fig. 14b). It prefers the interior points due to two factors:

1. The calculation of the mean coordinate value from all active units usually results in a goal prediction near the center of the array, so units coding for exterior positions will be treated as outliers and are likely to be deactivated.
2. Units tuned to locations interior to the array will have at least a partial match to both landmarks of the stretched array, giving




FIGURE 12. a: Histogram of search time of gerbils trained on the two-landmark task. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions by the simulation.

them higher activation levels than units tuned to exterior locations. This helps pull the goal prediction closer to the interior of the array.

The simulation does not predict the goal to be at the center of the split array (located between the two interior search locations) because of a third factor:

3. The four populations of units tuned to single landmarks (actually, the same landmark twice), which get a perfect match at all four possible search locations and become highly activated, contribute substantially to the total place unit activation. They therefore constrain the set of possible goal predictions to just those four locations. Factors 1 and 2 cause the model to choose the interior locations from this set.

Three-landmark experiments. Collett et al. (1986) trained a third group of gerbils to search for food at the center of a triangular array. As with the one- and two-landmark arrays, well-trained animals searched for food at the correct location (Fig. 15a). The goal predictions of our simulation (Fig. 15b) again correspond well to the behavioral data.

With three landmarks, additional manipulations are possible. Collett et al. deleted one or two landmarks (Figs. 16, 17), added a landmark to one side forming a second triangle with opposite orientation (Fig. 18), and moved one landmark away from the others, producing inconsistent cues (Fig. 19). The model is compatible with the behavioral data from each of these experiments.

When the triangle was inverted (or equivalently, rotated by




FIGURE 13. a: Histogram of search time of gerbils trained on the two-landmark task and tested with one landmark removed. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions by the simulation.




FIGURE 14. a: Histogram of search time of gerbils trained on the two-landmark task and tested with the inter-landmark distance doubled. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions from the simulation.

60°), gerbils first searched the center of the array and then proceeded to search three exterior points (Fig. 20a). Our model produces goal predictions at these four locations (see Fig. 20b). The interior and exterior goal predictions occur because of a divergence in the possible outcomes of the realignment process.

1. Interior goal predictions. On some trials, realigning the head direction produces an improvement in consistency. The sim-animal will then see a correctly oriented triangle and will predict the goal to be in the center of the array.
2. Exterior goal predictions. On other trials, the improvement in consistency is below the threshold necessary to sustain a head direction realignment. In this case the animal will retain its original head direction estimate and use pairs of landmarks in the correct orientation to localize itself, as in the missing-landmark probe trial. It will therefore search one of the three exterior locations.

It is unclear how the head direction realignment associated with searching an interior point could later be undone, yet Collett et al. (1986) report that the gerbils first searched in the center and then proceeded to search at exterior locations. Further work is needed to determine whether our model is compatible with this observation.

The vector voting hypothesis. To account for the gerbils' behavior in this collection of experiments, Collett et al. (1986) posit that the animals begin planning their path to the goal by applying every learned landmark-to-goal vector (Fig. 21) to all the landmarks they currently perceive. The animal tallies the votes, and




FIGURE 15. a: Histogram of search time of gerbils in the three-landmark task. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions from the simulation.




FIGURE 16. a: Histogram of search time of gerbils trained on the three-landmark task and tested with one landmark removed. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions from the simulation.

the best-supported location is taken to be the goal. When there is a tie, multiple locations are searched. We will refer to this as the *vector voting* hypothesis.

Both the parallel relaxation and vector voting mechanisms allow multiple hypotheses to be simultaneously active. Thresholding selects the best-supported hypothesis from the set. The difference between the two algorithms is that the vector voting scheme requires a separate representation for each candidate, so its votes can be tallied independently of the others. The parallel relaxation scheme is also a voting scheme, but votes come from individual place units trying to estimate a unique location for the animal by averaging coordinate values; ties are not permitted. The equivalence between relaxation and vector voting depends on two assumptions. First, place units must uniformly sample the landmarks available in the training environment, so that on probe trials, locations that should get more votes (due to their being consistent with a greater fraction of the landmarks present) will activate more units and exert more influence on the coordinate estimate. Second, there must be sufficient randomness in the simulation so that locations that receive equal support have a roughly equal chance of being selected by the relaxation process. Thus, when two candidates have roughly equal support, the distribution of winners across multiple trials will show two peaks.

Most of the tasks explored by Collett et al. can be explained by the vector voting hypothesis, but the split-array (Fig. 14) and the rotated-triangle (Fig. 20) tasks cannot. For the split-array task, the vector voting hypothesis predicts four search locations (Fig. 22). The animal's preference for the interior locations can be ex-




FIGURE 17. a: Histogram of search time of gerbils trained on the three-landmark task and tested with two landmarks removed. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions from the simulation.




FIGURE 18. a: Histogram of search time of gerbils trained on the three-landmark task and tested with a fourth landmark added. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions from the simulation.




FIGURE 19. a: Histogram of search time of gerbils trained on the three-landmark task and tested with one landmark moved twice as far from the goal. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions from the simulation.




FIGURE 20. a: Histogram of search time of gerbils trained on the three-landmark task and tested with a rotated array. From Collett et al. (1986), with permission of author and publisher. b: Histogram of goal predictions from the simulation.



FIGURE 21. In the vector voting scheme, a vector from each landmark (solid circle) to the goal (small triangle) is learned during training trials (a). In later probe trials (b), these vectors are applied to all the perceived landmarks. The location with the most votes is searched. The different arrow shapes denote different memory vectors.



FIGURE 22. Vector voting on the split-array task predicts four goal locations, but gerbils search only the two interior ones.

plained if it treated the landmark pair as a “beacon”: a single landmark co-located with the goal. When far from the landmark pair, aiming for this beacon would be easier than computing a trajectory by self-localization. However, if beaconing is involved, an additional mechanism would be required to explain why there is no central peak in the search-time histogram (Fig. 14a), as there is in the rotated-triangle task (Fig. 20a).

Collett et al. (1986) instead suggest that the animals use their perception of the split array as a whole to distinguish the east from the west landmark. This allows them to apply each of the two learned vectors to only the corresponding landmark instead of to both, generating just two candidates instead of four. The mechanisms required for this sort of holistic perception are left unspecified.

The vector voting hypothesis is also unable to explain the gerbils’ search-time spent in the center of the rotated-triangle. It predicts only the three exterior search locations (Fig. 23).

The animal’s predilection for the center of the rotated triangle early in its search can be explained if it treated the entire array as a beacon, as suggested for the split-array task. The animal might not notice at a distance that the triangle was rotated. Failing to find food at the center, it would attempt to reorient itself based on its place representation, and would thus be led to search the three exterior locations. However, this explanation is incompatible with the fact that the animals do not search at the center of other configurations, such as the split-array or stretched triangle tasks.




FIGURE 23. Vector voting on the rotated-triangle task predicts three exterior goal locations with two votes each. Gerbils also search in the center of the rotated array.

PREDICTIONS

Our model offers predictions which we hope to see confirmed experimentally. All of the effects described below have been observed in our simulations.

Predictions About the Place Code

1. The activation function for place cells predicts that the shapes of place fields should change when landmarks are perturbed by small amounts, but not when the landmarks are removed altogether.
2. Whenever the animal thinks it is at the goal in the Collett et al. tasks, its path integrator coordinates will be the same. Therefore, when rodents search in multiple locations, as in the split-array probe, the same place code should be observed, even though the local view (taking allocentric bearings into account) is different at each location.
3. In our theory, when the animal decides that its present search location is not really the goal, it repeats the self-localization procedure and gets back new coordinate values that do not match the goal location. This implies that the place code must change when the self-localization is repeated. Thus, our theory predicts that in the Collett et al. tasks where gerbils search in multiple locations, a change in the hippocampal place code will be observed just prior to the gerbil leaving one search location for another.
4. When re-entering a familiar environment, the number of active place cells should initially be larger than usual. This is because the path integrator will not yet be constraining the place cell activity pattern. We cannot predict the time course that the place code will require to settle into its normal mode, but we expect that it will be very fast. We are not aware of any study that has looked at place cell activity of a rat starting just before it is deposited into a familiar arena, so it is not surprising that the effect we predict has not been reported yet. Nevertheless, we believe that this effect will be observable by examining the number of initially active place cells over a number of trials.

5. When re-entering a familiar environment, the place code may initially be incoherent. Specifically, cells may fire outside of their normal place fields, and cells with non-overlapping place fields may be simultaneously active. On completion of the self-localization process, a normal, coherent place code should be seen. As with prediction 4, above, we cannot predict the time course of this initial incoherency; it may be very short-lived, and as far as we know, no one has looked for it. But we do expect it to be observable.

Alyan et al. (1995) report that rats appear to use passive path integration to keep track of their position while being transported from their home cage to a Morris water maze tank. The rats were carried on a wooden board along a straight path, and could see their surroundings. Another group of animals was transported in a black box, along a circuitous path, which presumably prevented passive path integration. This latter group scored significantly lower on the task according to three different performance measures. This suggests that in searching for the effects of self-localization, experimenters will need to disorient the animal and begin recording *before* it is removed from the transport box.

It is also likely that the degree of initial incoherency in the place code depends on the amount of symmetry in the environment and the number of distal cues available to resolve local ambiguities. A Hampton Court-style maze with many identical corridors, surrounded on all sides by a featureless black curtain, would provide for considerable ambiguity and might make the gradual formation of a coherent place code more observable.

Predictions About the Head Direction System

1. In the Collett et al. (1986) rotated triangle task, the tuning of the head direction system should be off by a multiple of 60° when the gerbil searches in the center of the triangle.
2. In the Cheng (1986) task, the tuning of the head direction system should be off by 180° when the rat searches in the mirror image to the goal location. Thus, by measuring the responses of head direction cells while the rat is still in the center of the arena, it should be possible to predict whether it will be successful in finding the goal on that trial.
3. In the Sharp et al. (1990) experiments, our model predicts which cue card and entry point configurations will cause the head direction system to precess, and by how much.

Interactions Between Place Code and Head Direction

1. Since place fields follow the head direction system, our model predicts when rotations of place fields will be observed in navigation tasks, for both rotationally disoriented and non-disoriented animals.
2. If Sharp et al.'s experiment using non-disoriented animals were repeated using several entry points for training (for example, placing the rat sometimes in the northwest, sometimes in the southwest, and sometimes in the northeast), our model predicts that place fields would not rotate when the rat entered from the south-east during a probe trial with two cards present. The reason is

that when trained with a varying entry point, the rat cannot estimate its path integrator coordinates upon entry. Therefore it would have to perform a self-localization to initialize its path integrator, and during this process place cells tuned to allocentric landmark bearings would allow it to distinguish between the two cards. Afterward, a head direction reset would not improve the consistency of the place code, so head direction would not change, and place fields would not rotate.

On the other hand, if the animals were disoriented prior to entry on the two-card probe trials, we predict that 50% of the time the place fields would appear rotated. A similar behavioral experiment by Cheng is described in Reproducing Simple Behavioral Experiments.

DISCUSSION

Our model demonstrates that a specific computational mechanism can account for a substantial body of navigation data. The mechanism is more complex than that of pure place cell models, as it includes the maintenance of head direction and involves both direct perception and an internal model of the world derived by path integration. But it is still relatively simple in comparison to the anatomy of the real hippocampal and head direction systems. It can be described by a few equations and parameter values.

One of the ways this work has furthered our understanding of rodent navigation is by calling attention to relationships between different experiments. For example, the split-array and rotated triangle tasks of Collett et al. (1986) are in conflict with respect to beaconing. The Sharp et al. (1990) one and two cue card tasks, together with the Cheng (1986) and Margules and Gallistel (1988) tasks, point out the complexity of interactions between the self-localization and head direction alignment operations. Modeling has also shown that vector voting can be accomplished in a distributed connectionist fashion, perhaps implemented in the hippocampal place code, that yields some additional emergent properties, such as only searching the two interior locations in the split-array task.

Are Place Cells Tied to Single Landmarks?

McNaughton et al. (1994) have proposed a theory of how visual landmarks can be used to locate a goal. Place cells in their formulation are only tuned to single landmarks, so that individual place cells can be identified with a specific landmark vector. Thus, when sitting at the goal in the two-landmark task (Collett et al., 1986) there would be two populations of cells active simultaneously: one tuned to the east landmark and one tuned to the west one. By subtracting a remembered goal-to-landmark vector from the currently perceived vector from the animal to the landmark, a vector from the present position to the goal can be derived.

In the absence of head direction information, these place cells are predicted to have annular place fields. McNaughton et al. note that the animal could self-localize by taking intersections of an-

nuli, but the mechanism for doing this is unspecified. In our model, place units fall back on distance and retinal angle information when allocentric landmark bearings cannot be determined. Some units may choose the same landmark twice, and thus might behave in a similar fashion to those of McNaughton et al. (1994) until the path integrator came online. However, once the path integrator is online, these cells will not show multiple or annular fields. Most units in our model are tied to two distinct landmarks, so there is no need to compute intersections of annuli when the model is rotationally disoriented.

Another difference between these two theories is revealed by the split-array probe trial. In the McNaughton et al. (1994) theory, the two populations of place cells coding for the goal location are predicted to dissociate when the landmarks are moved farther apart, so that only one population would be maximally active at each of the two search locations. We predict the same place cell activity pattern will be observed at both locations.

Place units tuned to single landmarks would generate four search locations in the split-array task. But as with the vector voting hypothesis, some additional mechanism might be invoked that would limit search to the two locations interior to the array.

Path Integration and Virtual Landmarks

Both Muller et al. (1991) and McNaughton et al. (1994) posit "invisible," "remembered," or "virtual" landmarks influencing place cell activity in tasks such as those with consistent entry points, those showing directionality of place fields, and navigation in the dark. In contrast, our model, and that of McNaughton et al. (1996), use path integrator information rather than virtual landmarks. These proposals are quite similar if one regards each virtual landmark as the origin of a reference frame for path integration, but there are some differences.

A place cell tuned to a single landmark should maintain its firing field when that landmark is removed from the environment. If this occurs by the animal's tracking a remembered landmark at that location, then as additional landmarks are removed the cognitive burden would increase. Eventually a limit should be reached. If the environment is rich in landmarks, then some place cells may have to shut down in the dark. Alternatively, an auto-associative interaction between place cells might activate a cell whose landmark was no longer driving it directly.

By including a separate path integration module referenced by all place units, our model is relieved of having to track an unbounded number of virtual landmarks, although it may track two or three reference frames simultaneously in some tasks.

Another difference between these proposals is that virtual landmarks, being landmarks, would most naturally be encoded in polar coordinates, whereas our model uses a Cartesian system for path integration. The latter is not only superior with regard to numerical stability, it also provides the means for vector subtraction necessary to calculate goal trajectories.

Finally, McNaughton et al. (1994) suggest that a more coarse-grained encoding would be used for large distance values relative to small ones. This is a reasonable assumption for representing perceived positions of distant landmarks, but it does not appear

compatible with path integration, where accurate track of position must be preserved so that when the animal returns to the vicinity of the (unseen) landmark it will have correct bearing information.

Future Work

Expanding environments

Two phenomena not addressed by our current activation function are the expansion of place fields in dilated environments and the introduction of barriers into the arena, both as reported by Muller and Kubie (1987). Place fields based purely on retinal angle features would expand linearly in a dilated environment. Place fields based on distances to pairs of landmarks would either be disrupted by dilation, or expanded and highly distorted, depending on the value of σ_r , the width of the Gaussian in the feature match function. What Muller and Kubie observed when the diameter of a circular arena was doubled from 38 to 76 cm, which quadrupled the surface area, was that the area of the place fields increased by an average factor of only 1.73. Place fields were defined as the region where the firing rate of the cell was within a certain percentage of the maximum firing rate observed for that cell in that arena. Although the place field area increased somewhat by this measure, the maximum firing rates of the cells decreased in absolute terms. This might be modeled by making the gain σ_r of the distance matching features dynamically adjustable within some range, suggesting a hippocampal gain control mechanism perhaps mediated by inhibitory interneurons. In the dilated environment, distance mismatches could be partially discounted by increasing σ_r whenever that improved certain qualities of the resulting place code.

Barriers in place fields

Muller and Kubie (1987) studied the effect that introducing a barrier into a cylindrical environment had on place cells. The barrier tended to disrupt a cell's response, but only if it was located inside the place field. A barrier outside the field but near the center of the cylinder (so that it blocked the animal's view of part of the cylinder) caused much less disruption than a barrier that bisected the field. Furthermore, transparent barriers had just as much effect on place cell firing as opaque ones. Although most cells decreased firing when the barrier was added, Muller and Kubie report one cell whose firing became more intense.

This suggests that place cell activity reflects non-visual aspects of the environment as well as visual ones. One hypothesis is that the rat maintains a representation of peripersonal space that is not primarily visual, and that place cells are tuned to features of this representation in addition to visual landmark features. This representation may be constructed in part from tactile information, such as from the vibrissae. This could explain how rats are able to maintain accurate distance information (but not heading) when navigating in a cylindrical arena in the dark. The representation of peripersonal space can be seen as an extension of the tuning to surface orientations available to place cells in the current model.

The role of hippocampal theta

The most prominent feature of hippocampal physiology not addressed by our model is the theta rhythm. Recent observations show that the firing phase of place cells with respect to the theta rhythm correlates with whether the place field is centered ahead of vs. behind the rat (O'Keefe and Recce, 1993; Skaggs et al., 1995b), and this property has been exploited in a model by Burgess et al. (1994). Several other proposals have also been made for a computational role for the theta rhythm in navigation (O'Keefe, 1989; O'Keefe, 1991). Our model presently makes no use of theta, but when its abstract computations are instantiated by some biological mechanism, it is possible that the usefulness of the phase relationship will become apparent.

Reference frames

In a two-landmark task similar to the split-array task, Gothard et al. (1994) report a variety of response types for CA1 cells. Some exhibited place fields tied to the room, some to the landmarks or goal location (which translated from trial to trial), some cells fired when the animal was either inside, leaving, or re-entering the start box no matter where it was located, and some appeared to be related to a mixture of these features. This suggests that multiple representations of space exist simultaneously in the hippocampus. We plan to extend our model to cover multiple reference frames and relationships between frames.

Learning

Perhaps the most important area yet to be addressed is the learning of new environments. Rats released in novel environments show place fields immediately (Hill, 1978), but the fields are strengthened and sharpened after an exploratory period of 10 minutes (Wilson and McNaughton, 1993). Place fields decrease in size and increase in directionality over the first few hours in an arena (Austin et al., 1993). A model based on realistic visual input and neurally plausible learning mechanisms that exhibited these sorts of properties would be a significant advance.

Acknowledgments

The first version of our model was built in collaboration with Hank Wan. We thank Bruce McNaughton, Kati Gothard, and Neil Burgess for helpful discussions and five anonymous referees for help in improving the manuscript. We are grateful to T.S. Collett and Springer-Verlag for permission to reproduce the figures from Collett et al. (1986), and Matt Wilson for use of his xviv software to create our graphic displays.

REFERENCES

- Abraham L, Potegal M, Miller S (1983) Evidence for caudate nucleus involvement in an egocentric spatial task: Return from passive transport. *Physiol Psychol* 11(1):11–17.
- Alyan SH, Jander R (1994) Short-range homing in the house mouse *Mus musculus*: stages in the learning of directions. *Animal Behavior* 48(2):285–298.
- Alyan SH, Touretzky DS, Taube JS (1995) The involvement of passive path integration in learning the Morris water maze. *Soc Neurosci Abstr* 21:1939.
- Anderson JA, Hinton GE (1981) Models of information processing in the brain. In: *Parallel models of associative memory* (Hinton, Anderson (eds), chapter 1, pp 9–48. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Austin KB, White LH, Shapiro ML (1993) Short- and long-term effects of experience on hippocampal place fields. *Soc Neurosci Abstr* 19:797.
- Blair HT (1996) A thalamocortical circuit for computing directional heading in the rat. In: *Advances in neural information processing systems* 8. Touretzky DS, Mozer MC, Hasselmo ME, (eds), pp 152–158. Cambridge, MA: MIT Press.
- Blair HT, Sharp PE (1995) Anticipatory head direction signals in anterior thalamus: evidence for a thalamocortical circuit that integrates angular head motion to compute head direction. *J Neurosci* 15(9): 6260–6270.
- Bostock E, Muller RU, Kubie JL (1991) Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* 1(2):193–206.
- Burgess N, Recce M, O'Keefe J (1994) A model of hippocampal function. *Neural Networks* 7(6/7):1065–1081.
- Carr H, Watson JB (1908) Orientation in the white rat. *J Comp Neurol Psychol* 18:27–44.
- Chen LL (1991) Head-directional information processing in the rat posterior cortical areas. PhD thesis, University of Colorado.
- Chen LL, Lin LH, Barnes CA, McNaughton BL (1994a) Head-direction cells in the rat posterior cortex: II. Contributions of visual and ideothetic information to the directional firing. *Exp Brain Res* 101: 24–34.
- Chen LL, Lin LH, Green EJ, Barnes CA, McNaughton BL (1994b) Head-direction cells in the rat posterior cortex: I. Anatomical distribution and behavioral modulation. *Exp Brain Res* 101:8–23.
- Cheng K (1986) A purely geometric module in the rat's spatial representation. *Cognition* 23:149–178.
- Collett T, Cartwright BA, Smith BA (1986) Landmark learning and visuo-spatial memories in gerbils. *J Comp Physiol [A]* 158:835–851.
- Efron B (1982) The jackknife, the bootstrap, and other resampling plans, volume 38 of regional conference series in applied mathematics. Philadelphia PA: Society for Industrial and Applied Mathematics.
- Eichenbaum H, Kuperstein M, Fagan A, Nagode J (1987) Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *J Neurosci* 7(3):716–732.
- Etienne AS (1987) The control of short-distance homing in the golden hamster. In: *Cognitive processes and spatial orientation in animals and man* (Ellen P, Thinus-Blanc C (eds), pp 233–251. Boston: Martinus Nijhoff Publishers.
- Gallistel CR (1990) *The organization of learning*. Cambridge, MA: MIT Press.
- Gaulin SJC (1992) Evolution of sex differences in spatial ability. *Yearbook Phys Anthropol* 35:125–151.
- Goodridge JP, Taube JS (1995) Preferential use of the landmark navigational system by head direction cells in rats. *Behav Neurosci* 109(1): 49–61.
- Gothard KM, Skaggs WE, Moore KM, McNaughton BL (1994) Behavioral correlates of hippocampal CA1 cells in a landmark navigation task. *Soc Neurosci Abstr* 20:1207.
- Hetherington PA, Shapiro ML (1993) A simple network model simulates hippocampal place fields: 2. Computing goal-directed trajectories and memory fields. *Behav Neurosci* 107(3):434–443.
- Hill AJ (1978) First occurrence of hippocampal spatial firing in a new environment. *Exp Neurol* 62:282–297.
- Hinton GE, McClelland JL, Rumelhart DE (1986) Distributed representations. In: *Parallel distributed processing: explorations in the microstructure of cognition* (Rumelhart DE, McClelland JL, eds), volume 1, pp 77–109. Cambridge, MA: MIT Press.
- Knierim JJ, McNaughton BL, Duffield C, Bliss J (1993) On the bind-

- ing of hippocampal place fields to the inertial orientation system. *Soc Neurosci Abstr* 19:795.
- Knierim JJ, Kudrimoti HS, McNaughton BL (1995) Place cells, head direction cells, and the learning of landmark stability. *J Neurosci* 15:1648–1659.
- Kohonen T (1984) *Self-organization and associative memory*. New York: Springer-Verlag.
- Kubie JL, Ranck JB Jr (1983) Sensory-behavioral correlates in individual hippocampus neurons in three situations: space and context. In: *Neurobiology of the hippocampus* (Seifert W, ed), chapter 22, pp 433–447. New York: Academic Press Inc.
- Kuipers BJ, Froom R, Lee W-Y, Pierce D (1993) The semantic hierarchy in robot learning. In: *Robot Learning* (Connell JH, Mahadevan S, eds), chapter 6, pp 141–170. Boston: Kluwer Academic Publishers.
- Margules J, Gallistel CR (1988) Heading in the rat: determination by environmental shape. *Animal Learning Behavior* 16(4):404–410.
- Markus EJ, Barnes CA, McNaughton BL, Gladden VL, Skaggs WE (1994) Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. *Hippocampus* 4:410–421.
- Markus EJ, Qin Y, Leonard B, Skaggs WE, McNaughton BL, Barnes CA (1995) Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J Neurosci* 15:7079–7094.
- Marr D (1969) Simple memory: a theory for archicortex. *Philos Trans R Soc Lond. [Biol.]* 262:23–81.
- Maurer R, Seguinot V (1995) What is modelling for? A critical review of the models of path integration. *J Theoret Biol* 175:457–475.
- McNaughton BL (1989) Neuronal mechanisms for spatial computation and information storage. In: *Neural connections, mental computation* (Nadel L, Cooper L, Culicover P, Harnish RM, eds), chapter 9, pp 285–350. Cambridge, MA: MIT Press.
- McNaughton BL, Nadel L (1990) Hebb-Marr networks and the neurobiological representation of action in space. In: *Neuroscience and connectionist theory* (Gluck MA, Rumelhart DE, eds), chapter 1, pp 1–63. Hillsdale, NJ: Lawrence Erlbaum Associates.
- McNaughton BL, Barnes CA, O'Keefe J (1983) The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* 52:41–49.
- McNaughton BL, Leonard B, Chen L (1989) Cortical-hippocampal interactions and cognitive mapping: a hypothesis based on reintegration of the parietal and inferotemporal pathways for visual processing. *Psychobiology* 17(3):230–235.
- McNaughton BL, Chen LL, Markus EJ (1991) "Dead reckoning," landmark learning, and the sense of direction: a neurophysiological and computational hypothesis. *J Cognitive Neurosci* 3(2):190–202.
- McNaughton BL, Knierim JJ, Wilson MA (1994) Vector encoding and the vestibular foundations of spatial cognition: neurophysiological and computational mechanisms. In: *The Cognitive Neurosciences* (Gazzaniga M, ed), chapter 37, pp 585–595. Boston: MIT Press.
- McNaughton BL, Barnes CA, Gerrard JL, Gothard K, Jung MW, Knierim JJ, Kudrimoti H, Qin Y, Skaggs WE, Suster M, Weaver KL (1996) Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J Exp Biol* 199(1):173–186.
- Mittelstaedt H (1983) The role of multimodal convergence in homing by path integration. In: *Multimodal convergences in sensory systems*, volume 28 of *fortschritte der zoologie* (Horn E, ed), pp 197–212. New York: Gustav Fischer Verlag.
- Mittelstaedt ML, Mittelstaedt H (1980) Homing by path integration in a mammal. *Naturwissenschaften* 67:566–567.
- Mizumori SJY, Cooper BG (1995) Spatial representations of dorsal caudate neurons of freely-behaving rats. *Soc Neurosci Abstr* 21:1929.
- Mizumori SJY, Williams JD (1993) Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *J Neurosci* 13(9):4015–4028.
- Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* 7:1951–1968.
- Muller RU, Kubie JL, Ranck JB Jr (1987) Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J Neurosci* 7:1935–1950.
- Muller RU, Kubie JL, Bostock EM, Taube JS, Quirk GJ (1991) Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In: *Brain and space* (Paillard J, ed), chapter 17, pp 296–333. New York: Oxford University Press.
- Muller RU, Bostock E, Taube JS, Kubie JL (1994) On the directional firing properties of hippocampal place cells. *J Neurosci* 14(12):7235–7251.
- O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Exp Neurol* 51:78–109.
- O'Keefe J (1989) Computations the hippocampus might perform. In: *Neural connections, mental computation* (Nadel L, Cooper L, Culicover P, Harnish RM, eds), chapter 8, pp 225–284. Cambridge, MA: MIT Press.
- O'Keefe J (1991) An allocentric spatial model for the hippocampal cognitive map. *Hippocampus* 1(3):230–235.
- O'Keefe J, Conway DH (1978) Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp Brain Res* 31:573–590.
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Exp Brain Res* 34:171–175.
- 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'Keefe J, Speakman A (1987) Single unit activity in the rat hippocampus during a spatial memory task. *Exp Brain Res* 68:1–27.
- Olton DS, Samuelson RJ (1976) Remembrance of places passed: spatial memory in rats. *J Exp Psychol [Anim Behav]* 2(2):97–116.
- Pico RM, Gerbrandt LK, Pondel M, Ivy G (1985) During stepwise cue deletion, rat place behaviors correlate with place unit responses. *Brain Res* 330:369–372.
- Potegal M (1972) The caudate nucleus egocentric localization system. *Acta Neurobiol Exp* 32:479–494.
- Potegal M (1982) Vestibular and neostriatal contributions to spatial orientation. In: *Spatial abilities: development and physiological foundations*. Potegal M (ed), chapter 15, pp 361–387. New York: Academic Press.
- Quirk GJ, Muller RU, Kubie JL (1990) The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J Neurosci* 10(6):2008–2017.
- Redish AD, Touretzky DS (1995) Revisiting the Papez circuit: the role of hippocampus and its afferent and efferent structures in rodent navigation. *Soc Neurosci Abstr* 21, p. 942.
- Redish AD, Touretzky DS (1996a) Cognitive maps beyond the hippocampus. Under review.
- Redish AD, Touretzky DS (1996b) Navigating with landmarks: Computing goal locations from place codes. In: *Symbolic visual learning* (Ikeuchi K, Veloso M, eds), Oxford University Press. In press.
- Redish AD, Elga AN, Touretzky DS (1996) A coupled attractor model of the rodent head direction system. under review.
- Rolls ET (1989) The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In: *The computing neuron* (Durbin R, Miall C, Mitchison G, eds) chapter 8, pp. 125–129. Reading, MA: Addison-Wesley.
- Samsonovich A, McNaughton BL (1995) Millisecond temporal structure of memory representations and hippocampal-dependent cognitive mechanisms. *Soc Neurosci Abstr* 21:942.
- Schmajuk NA, Blair HT (1993) Place learning and the dynamics of spatial navigation: a neural network approach. *Adaptive Behavior* 1:355–387.
- Shapiro ML, Hetherington PA (1993) A simple network model simulates hippocampal place fields: parametric analyses and physiological predictions. *Behav Neurosci* 107(1):34–50.

- Sharp PE (1991) Computer simulation of hippocampal place cells. *Psychobiology* 19(2):103–115.
- Sharp PE, Kubie JL, Muller RU (1990) Firing properties of hippocampal neurons in a visually symmetrical environment: contributions of multiple sensory cues and mnemonic processes. *J Neurosci* 10(9):3093–3105.
- Sharp PE, Blair HT, Erkin D, Tzanetos DB (1995) Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. *J Neurosci* 15(1):173–189.
- Skaggs WE, Knierim JJ, Kudrimoti HS, McNaughton BL (1995a) A model of the neural basis of the rat's sense of direction. In: *Advances in neural information processing systems 7* (Tesauro G, Touretzky DS, Leen TK eds), pp 173–180. Cambridge, MA: MIT Press.
- Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1995b) Relations between the hippocampal theta rhythm and spatial firing of CA1 complex spike cells. Manuscript.
- Suzuki S, Augerinos G, Black AH (1980) Stimulus control of spatial behavior on the eight-arm maze in rats. *Learning Motivation* 11:1–18.
- Taube JS (1995) Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *J Neurosci* 15(1):1953–1971.
- Taube JS, Muller RI, Ranck JB, Jr (1990a) Head direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J Neurosci* 10:420–435.
- Taube JS, Muller RI, Ranck JB, Jr (1990b) Head direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *J Neurosci* 10:436–447.
- Thompson LT, Best PJ (1989) Place cells and silent cells in the hippocampus of freely-behaving rats. *J Neurosci* 9(7):2382–2390.
- Touretzky DS, Redish AD, Wan HS (1993) Neural representation of space using sinusoidal arrays. *Neural Computation* 5(6):869–884.
- Touretzky DS, Wan HS, Redish AD (1994) Neural representation of space in rats and robots. In: *Computational intelligence: imitating life* (Zurada JM II, Robinson CJ, editors), pp 57–68. IEEE Press.
- Treves A, Miglino O, Parisi D (1992) Rats, nets, maps, and the emergence of place cells. *Psychobiology* 20(1):1–8.
- Wan HS, Touretzky DS, Redish AD (1994a) A rodent navigation model that combines place code, head direction, and path integration information. *Soc Neurosci Abstr* 20:1205.
- Wan HS, Touretzky DS, Redish AD (1994b) Towards a computational theory of rat navigation. In: *Proceedings of the 1993 connectionist models summer school* (Mozer M, Smolensky P, Touretzky D, Elman J, Weigend A, eds), pp 11–19. Hillsdale NJ: Lawrence Erlbaum Associates.
- Wiener SI (1993) Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *J Neurosci* 13(9):3802–3817.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.
- Wittmann T, Schwegler H (1995) Path integration—a network model. *Biol Cybernet* 73(6):569–575.
- Worden R (1992) Navigation by fragment fitting: a theory of hippocampal function. *Hippocampus* 2(2):165–187.
- Zhang K (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J Neurosci* 16(6):2112–2126.
- Zipser D (1985) A computational model of hippocampal place fields. *Behav Neurosci* 99(5):1006–1018.

APPENDIX A: CONSTRUCTING A PLACE CODE

One question the current theory does not address is: How is the place code learned? All the model requires is that a code exist with a sufficient number of place units associated with each location in the environment. Learning a place code has a number of complications for a real animal. Drift in the path integrator or in the head direction system must be corrected *while* the animal is learning the environment; any errors in those systems would result in incorrect associations between local views and path integrator coordinates. We have observed (unpublished observations) that gerbils exploring a large (10-foot diameter) open arena regularly return to their starting point. Perhaps they do so to realign their path integrator at a known location.

Our simulation uses an iterative procedure to construct a place code. At each step it picks a random location in the arena and initializes the path integrator to those coordinates. It then calculates the local view from that location and activates the current population of place units using both landmarks and path integrator input. If too few units become active, it recruits a new unit, sets its path integrator coordinates to the current location, and picks a

pair of currently visible landmarks at random to generate the distance, bearing, and retinal angle features. Then the process repeats.

The algorithm terminates when no new unit has been recruited for 20 successive steps, indicating that the set of place cells it has constructed is sufficient to cover the environment. Small environments, such as the cylindrical arena with 38 cm radius used by Muller et al., require approximately 2,000 place units to satisfy this criterion. The arena used by Collett et al., with a 175-cm radius, requires 10,000 to 15,000 place units. These numbers reflect the model's simplifying assumptions that the sizes of individual place fields do not scale with the environment, and that the whole arena is treated as a single environment. These assumptions probably do not apply in more naturalistic settings; rodents in the wild can have home ranges covering many hundreds of square meters (Gaulin, 1992).

This algorithm suffices to create a place code able to represent a familiar environment; when complete, each location of the environment is covered (with high probability) by a minimum number of place fields. A more realistic learning mechanism that also ensures that a sufficient number of place cells will be active at every location in the environment should not change any of the results discussed in Correspondence with Experimental Data. We are currently exploring such mechanisms.

APPENDIX B.

Model Parameters

Symbol	Value	Description
Place cell turning parameters		
σ_c	25 cm	Sharpness of $C(u)$ tuning to path integrator coordinates
σ_r (point)	15 cm	Sharpness of $F_i(u)$ tuning to range values for point landmarks
σ_r (surface)	10 cm	Sharpness of $F_i(u)$ tuning to range values for surface landmarks
σ_ϕ (point)	15°	Sharpness of $F_i(u)$ tuning to bearing values for point landmarks
σ_ϕ (surface)	50°	Sharpness of $F_i(u)$ tuning to bearing values for surface landmarks
σ_α	15°	Sharpness of $RA(u)$ tuning to retinal angles
Parallel relaxation parameters		
Initial σ_c	500 cm	Initial σ_c large enough to give $C(u)$ negligible gain at start of relaxation
ρ_c	0.9	Decay factor for tightening σ_c during relaxation
δ_A	0.75	Minimum activation level for including a place unit in coordinate calculation
Other parameters		
κ^*	0.4	Minimum acceptable improvement in consistency following a reset operation