1 Navigating with Landmarks: Computing Goal Locations from Place Codes

A. David Redish and *David S. Touretzky*
Carnegie Mellon University

To appear in *Symbolic Visual Learning*,
K. Ikeuchi and M. Veloso, eds.,
Oxford University Press.

Abstract

A computer model of rodent navigation, based on coupled mechanisms for place recognition, path integration, and maintenance of head direction, offers a way to operationally combine constraints from neurophysiology and behavioral observation. We describe how one such model reproduces a variety of experiments by Collett, Cartwright, and Smith [6] in which gerbils learn to find a hidden food reward, guided by an array of visual landmarks in an open arena. We also describe some neurophysiological predictions of the model; these may soon be verified experimentally. Portions of the model have been implemented on a mobile robot.

1. Introduction

Landmark-based navigation is a rich domain for exploring issues of visual and spatial cognition. At the behavioral level, there is a wealth of data on how animals use landmarks to locate food or return to their nests. At the neurophysiological level, hippocampal pyramidal cells called *place cells* have been discovered that fire when the animal is in a certain location in its environment [12, 20, 24]. These cells change their firing patterns in response to displacement or removal of prominent landmarks [21]. Cells in the subicular complex [34], the thalamus [19, 33], and the parietal cortex [4] whose firing rates are correlated with the animal's heading (*head direction cells*) have also been found to be controlled in part by visual cues.

In this chapter we describe a theory of navigation in rodents that is constrained by both behavioral and neurophysiological data. The theory is embodied in a computer model called CRAWL, allowing us to replicate various experiments in the rodent navigation literature and make predictions about place cell responses in novel situations. Some of this work has been previously reported in [37, 40, 41]. Portions of the model have also been implemented on a mobile robot [38, 40].

2. Experimental data from rodents

2.1. Representations

Place cells. Single-unit recordings from hippocampus have demonstrated that rats have an internal representation of space consisting of cells that fire when the animal is in a particular region of the environment (see [20] for a review). The activity in these cells seems to be dependent on the arrangement of visual cues: a rotation in visual cues causes a corresponding rotation of the firing fields, or *place fields* [15]. However, most place cells maintain their firing fields when some of the visual cues are removed [23, 25]. Place cells also continue to respond when the lights are turned off; thus they cannot be driven by visual cues alone. The response pattern may be rotated if the animal originally enters a symmetric environment such as a radial eight-arm maze in the dark, but the spatial relationships among place fields remain intact [28].

Head direction cells. Cells in other parts of the rat nervous system, including postsubiculum [34, 35], thalamus [19, 33], and parietal cortex [4], are unimodally tuned to head direction. The *preferred direction* for such a cell (i.e., the direction eliciting maximal response) is constant throughout an environment. Furthermore, the difference in preferred directions for any pair of cells is constant across environments. But a single cell's preferred direction measured with respect to true north may differ across environments. Head direction cells also maintain their activity in the dark, presumably by integrating vestibular cues, but the animal's directional sense will eventually drift if no external input is available [14].

There is evidence that rats use visual landmarks to correct for cumulative integration errors. If visual cues rotate while the rat is in a familiar environment, preferred directions of head direction cells rotate by a corresponding amount. On the other hand, rotation of an unfamiliar environment does not shift the cells' preferred directions [17].

Path integration. Mittelstaedt and Mittelstaedt [18] and Etienne [8] have shown that rodents are able to execute a direct path back to their starting location in a cue-controlled environment after having taken a complex path away from it. When the salient cues (such as the nest itself) are moved, the animals ignore those cues and return to the starting location. This suggests that the animals are maintaining their position relative to the start by means of path integration. There is no direct neurophysiological evidence for a path integration system, but there is some evidence that lesion of the caudate nucleus impairs such tasks [1, 26]. Other anatomical and physiological evidence has prompted Douglas [7] to posit that striatum and/or the parietal lobe are involved in updating an attention-related spatial vector by path integration.

Local view. A fourth representation of space comes from direct perception: it consists of a set of bearings and distances to currently visible landmarks. Although a great deal of work has been done on the neurobiology of the visual system, the mental representation of individual landmarks in a scene remains unknown. Nonetheless, the concept of "local view" is an important component of many hippocampal models [3, 16, 22].

2.2. Behavior

Collett, Cartwright, and Smith [6] trained gerbils to dig for a food reward at a fixed position relative to an array of circularly-symmetric cylindrical landmarks. In all of the experiments we will be concerned with, the array was translated but not rotated from trial to trial, and the animals were released from different starting points to ensure that the landmarks provided the only reliable cues to the reward location. The walls of the experimental chamber were painted black, and a single light bulb illuminated a central circular region, leaving the walls in darkness. This was done to limit the external cues the animal could glean from its environment, encouraging it to focus on the landmarks. The floor of the chamber was covered with wood chips.

Once the gerbils were trained to criterion, they could be tested by occasional probe trials in which the food was absent. The distribution of each animal's search effort during probe trials was plotted as a 2D histogram. In some probe trials the landmark array differed from that of the training trials. The distribution of search effort in these modified arrays is particularly informative, because it shows how the animals coped with discrepancies between memory and visual perception. Each animal was trained on only one landmark configuration and experienced probe trials only rarely, intermixed with the training trials.

Collett et al. trained one group of gerbils using a single landmark to indicate the food location (the *one-landmark* task). The gerbils learned to search for the food at the correct distance and bearing from the landmark: well-trained animals went directly to the food location upon release from the start box. (See Figure 1.) The fact that the animals could learn to search at the correct bearing as well as distance from a single symmetric landmark implies that they have some independent means of determining bearing information. Collett et al. supposed that the animals found some external cue that provided bearing information, despite their attempts to block this. But the subsequent discovery of an internal compass in rats (the head direction cells) suggests that external cues might not be necessary. Our model makes use of this compass.

Another group of gerbils were trained with two landmarks (the *two-landmark* task). Again, they successfully learned this task: well-trained animals went directly to the goal location. (See Figure 2.) With more than one landmark, the array can be modified; the animals' search patterns in response to the modified array can be very informative.

When one landmark was removed (the *two-minus-one* task), the gerbils searched

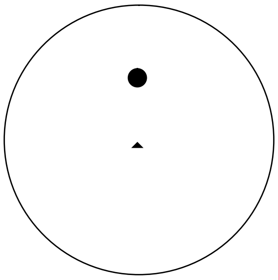

Copyrighted figure;
see [6] for original.

(a) (b)

Figure 1: *Distribution of gerbils' search effort during probe trials in Collett et al.'s one-landmark task. The large solid circle is the cylindrical landmark. The food reward, a sunflower seed, marked as a small triangle in the schematic (a), was located 50 cm south of the landmark during training trials. (b) Size of blobs shows time spent searching in each location during probe trials. Part b reprinted from [6] with permission of author and publisher.*

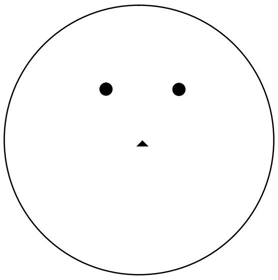

Copyrighted figure;
see [6] for original.

(a) (b)

Figure 2: *(a) Schematic of training with the two-landmark task. Food location is denoted by the small triangle. (b) Distribution of time spent by gerbils searching for food during probe trials in the two-landmark task. Part b reprinted from [6] with permission of author and publisher.*

Copyrighted figure;
see [6] for original.

Figure 3: *Distribution of time spent by gerbils searching for food after training with two landmarks and testing with only one. This is the two-minus-one task. Reprinted from [6] with permission of author and publisher.*

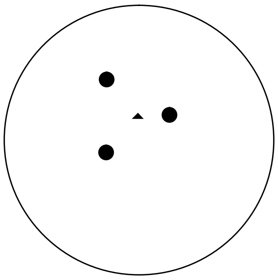

Copyrighted figure;
see [6] for original.

(a) (b)

Figure 4: *(a) Training in the three-landmark task. Food (small triangle) was located at the center of the landmark array. (b) Distribution of time spent by gerbils searching in the three landmark task. Part b reprinted from [6] with permission of author and publisher.*

Copyrighted figure;
see [6] for original.

(a)

Copyrighted figure;
see [6] for original.

(b)

Figure 5: *Distribution of time spent searching for food by gerbils trained with three landmarks and tested (a) with only two landmarks (the three-minus-one task), or (b) with only one landmark (the three-minus-two task). Reprinted from [6] with permission of author and publisher.*

alternately in two locations, each at the correct bearing and distance from one of the landmarks they observed during training (Figure 3). Collett et al. theorized that they were binding the cylinder first to one and then to the other of the two remembered landmarks. Note that the visual environment in the two-minus-one task is the same as in the one-landmark task, the difference in response results from the different training.

A third group of gerbils learned to find food at the center of a triangular array (the *three-landmark* task, Figure 4). Because there are now three landmarks, a number of manipulations can be made to the array. Similar to the two-landmark array, one or more landmarks can be removed. When only one landmark was removed, the gerbils concentrated their search in a single location (Figure 5a). The visual environment in the three-minus-one task is identical to that of the two-landmark task, but again the different training results in different responses. With two landmarks removed (Figure 5b) the visual environment reverts to that of the one-landmark task, but now the gerbils search three locations, equivalent to binding the remaining landmark to each remembered landmark in the three-

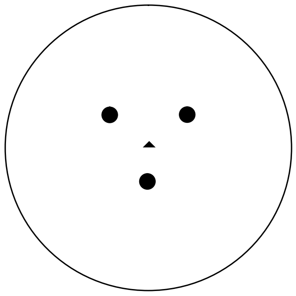

(a)

Copyrighted figure;
see [6] for original.

(b)

Figure 6: *(a) Training in the three-plus-1 task; same configuration as Figure 4a, save for orientation. (b) Distribution of search time during probe trials with a fourth landmark added. Part b reprinted from [6] with permission of author and publisher.*

Copyrighted figure;
see [6] for original.

Figure 7: *Distribution of time spent by gerbils searching when the one of the three landmarks is moved. This is the stretched-triangle task. Reprinted from [6] with permission of author and publisher.*

landmark task.

Collett et al. tested a third group of gerbils by adding a landmark to one side (the *three-plus-1* task). Adding a landmark produces two triangles with opposite orientations, one correct and one reversed (Figure 6). The gerbils concentrated their search at the center of the correctly oriented triangle and spent little time in the triangle with opposite orientation.

Some other manipulations of the triangular array produce inconsistent cues. For example, if one landmark is moved further away from the other two, it provides evidence for a different goal location than the other two landmarks. This is the *stretched-triangle* task (Figure 7). The gerbils concentrated their search at the location suggested by the two mutually consistent landmarks; they did not search at the learned bearing and distance from the outlier landmark.

To account for the gerbils' behavior on these tasks, Collett et al. posit that the gerbils predict the goal location by remembering vectors from each landmark to the food. Because the landmarks are identical, they apply every learned vector to every perceived landmark, as in Figure 8. The locations receiving the most votes are the ones searched. We shall refer to this as the *vector voting hypothesis*. It is

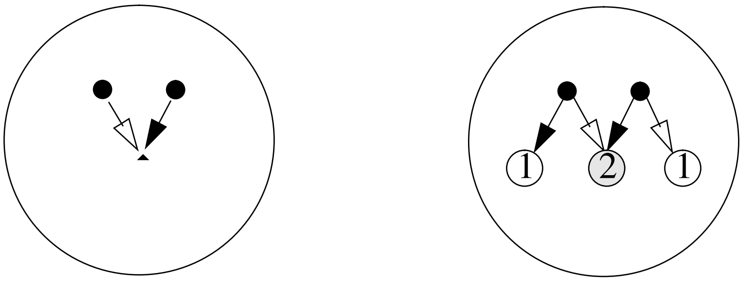

(a) (b)

Figure 8: *(a) Remembered vectors hypothesized by the vector voting hypothesis. Small triangle denotes food reward (sunflower seed) at the goal location. (b) Probe trial showing how vector voting determines the location to be searched. The location with two votes is searched while those with only one vote are ignored. From [37].*

compatible with all of the above results but requires $O(n^2)$ operations to compare perceived with remembered landmarks.

However, the vector voting hypothesis is not compatible with results from two additional manipulations performed by Collett et al.: the *split-array* task and the *rotated-triangle* task.

The two-landmark array can also be stretched. When trained on two landmarks and tested with the distance between them doubled (the *split array* task), the gerbils searched at two locations, each at the correct distance and bearing from one landmark (Figure 9a). The simplest form of the vector voting hypothesis incorrectly predicts search at four locations (Figure 9b). Collett et al. suggest that the gerbils may have used their perception of the array as a whole to disambiguate the two landmarks, but the mechanism for this is left unspecified.

When the three-landmark array was inverted, or equivalently, rotated by 60°, the gerbils first searched at the center of the rotated triangle and then at three exterior locations, as in Figure 11a. These three exterior locations can be derived by considering any two landmarks in the correct orientation. Figure 11b demonstrates how the vector voting hypothesis explains the three exterior search locations, but not the interior one.

The gerbils' tendency to search the center first can be explained by two possible strategies: At a distance, they may treat the entire array as a *beacon*, i.e. a single landmark co-localized with the goal. Alternatively, if a gerbil rotated its internal compass by 180°, the landmark array would appear in the "correct" orientation, and the animal would search the center of the array. However, when the animal failed to find food there, it would have to restore the previous alignment of its head direction system in order to generate search hypotheses at the three exterior points.

In the following sections we present an alternative model that reproduces

Copyrighted figure;
see [6] for original.

(a)

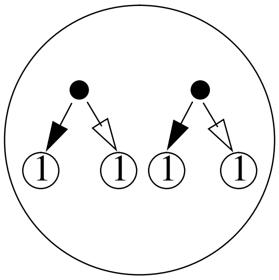

(b)

Figure 9: *(a) Distribution of search effort for gerbils trained on the two-landmark task (Figure 8a) and tested with distance between the landmarks doubled. (b) Vector voting applied to the split-array task. Part a reprinted from [6] with permission of author and publisher.*

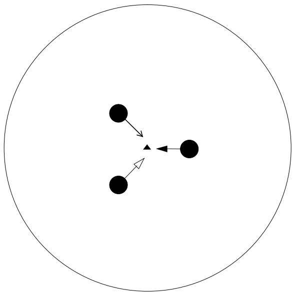

Figure 10: *The three-landmark task: remembered vectors hypothesized by the vector voting hypothesis. Small triangle denotes food reward (sunflower seed) at the goal location.*

Copyrighted figure;
see [6] for original.

(a)

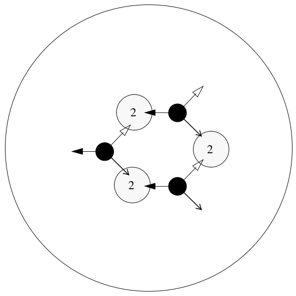

(b)

Figure 11: *(a) Distribution of search effort for gerbils trained on the three-landmark array but tested with the triangle inverted. (b) Vector voting applied to the rotated-triangle task. Part a reprinted from [6] with permission of author and publisher.*

the above tasks, including the two tasks that vector voting cannot. We will then show that our mechanism can be understood as a distributed connectionist implementation of vector voting, and discuss why it can explain the split-array and rotated-triangle tasks.

In all of these tasks, information from multiple landmarks must be combined to solve a problem: where to search for food. This involves an instance of the well-known *binding problem*: perceived landmarks must be bound to remembered landmarks to give a maximally consistent interpretation of the scene. Our model takes an indirect approach to the binding problem. Instead of asking "Which visual landmarks bind to which remembered landmarks?" our model asks "Where am I?" By answering this simpler question the model can derive answers to the more complex one, because once it knows its location, correspondences between visual percepts and remembered landmarks are highly constrained.

3. Overview of our theory of rodent navigation

The central points of our theory are (1) that behavior arises from an interaction between multiple representations of space and (2) that those representations mediate associations of internal and external stimuli. In our model, both the place cells and the head direction cells process combinations of internal and external information.

There is a well-established relationship between certain types of changes to the visual world and changes in place cell firing [20]. A number of computer models of visually-driven place cells have now been published [2, 3, 29, 30, 31, 39], but these fail to explain the persistence of place fields in the dark. In a circular arena, Quirk et al. observed that place fields in the dark are similar in size to fields in the light; their centers don't differ in distance to the wall, but the fields are often

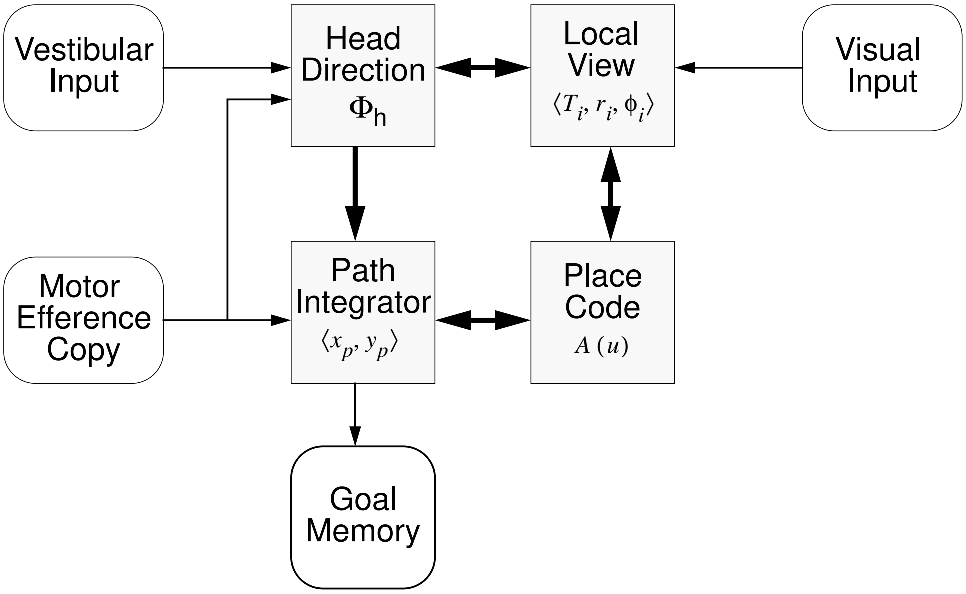

Figure 12: *Components of the* CRAWL *model.*

rotated [28]. If place cells were responding to non-visual cues, the size or shape of the fields might change, but the general location of the fields would not. The fact that place field locations change in a regular fashion in the dark suggests a source of internal cues. Our theory is that place cells receive an estimate of the animal's position from a path integration module, and the observed rotation of fields in the dark is a result of drift in the head direction system. Place cells and head direction cells have in fact been observed to drift in synchrony in the dark [15].

Head direction cells also associate internal with external cues. The preferred directions of cells in postsubiculum [34], thalamus [19, 33] and parietal cortex [4] rotate in response to the rotation of the visual environment. But cells also maintain their directional preferences in the dark, and therefore must be driven in part by internal (probably vestibular and motor efferent) signals. As mentioned earlier, rodents can integrate a complex exploratory path and return directly to their starting point. Mittelstaedt and Mittelstaedt [18] further demonstrated that gerbils could integrate angular acceleration as long as it was above the vestibular threshold of $0.24°\text{sec}^{-2}$. On the other hand, Etienne [8, 9] showed that when given an external orientation cue (a red light in a otherwise dark environment), hamsters preferred to rely on it rather than on internal cues.

4. The model

4.1. Structure

Figure 12 is a diagram of our model. Each box denotes a functional system; it should not be presumed to correspond to a single anatomical location. Visual input provides a tuple of information $\langle T_i, r_i, \theta_i \rangle$ specifying the landmark type and its range and bearing with respect to the animal. The animal's **head direction** Φ_h is updated by vestibular cues and efferent copies of motor commands as the animal moves. In the box labeled **local view**, allocentric bearings to landmarks are calculated by adding the head direction estimate to each egocentric bearing: $\phi_i = \theta_i + \Phi_h$.

The animal is also assumed to maintain an estimate of its position $\langle x_p, y_p \rangle$ in some internal coordinate system. This value is updated by the **path integrator** as the animal moves, based on vestibular input (via the head direction module) and motor efference copy.

The **place code** in our theory consists of place units[1] which maintain an association between the local view and path integrator coordinates. Place units are radial-basis units tuned to type, range, and allocentric bearing of two individual landmarks, plus retinal angle between one pair of landmarks and a point in path integrator coordinates. The activity of each unit is a product of univariate gaussian functions of each of these quantities.

We use redundant sources of cues in order to adequately cover a range of animal behaviors in a single model. For example, allocentric bearings to two landmarks are sufficient to uniquely localize a point in space only if the point is not co-linear with the landmarks. In environments with multiple identical landmarks, or only a single prominent landmark, combinations of range and bearing information help to efficiently disambiguate the visual input. Another type of information comes from the bearing difference between two landmarks, which is equivalent to their retinal angle. This defines a line of position, which can perhaps be measured more accurately than distance, and is especially useful when the internal compass has been disrupted and allocentric bearings cannot be determined.

Our place units compute a "fuzzy conjunction" of their inputs in which terms drop out when information is unavailable. Thus, in the dark, place units are driven solely by the path integrator. Conversely, each time the animal is reintroduced into a familiar lighted environment, path integrator coordinates are unavailable, so place units are driven solely by visual input until the animal has self-localized. Once the path integrator has been initialized, place units respond to a combination of internal and external cues.

Mathematically, we say that the activity of a place unit u is

$$A(u) = C(u) \cdot L(u)$$

where $C(u)$ is a two-dimensional gaussian tuning to path integrator coordinates centered at $\langle x_u, y_u \rangle$ and $L(u)$ is a product of univariate gaussian functions of perceptual variables. The width of these gaussians are determined by parameters

[1]We refer to the elements in our simulation as "units," reserving the term "place cells" for real hippocampal neurons. We do not assume a one-to-one correspondence between the model's place units and individual pyramidal cells.

which have been tuned to match neurophysiological data from [21]. We say that a cell with activity $A(u)$ greater than a dynamic threshold δ is *active*.

In the dark the $L(u)$ term drops out of the equation, leaving $A(u) = C(u)$, while on being re-introduced into a familiar lighted environment, the $C(u)$ term drops out, leaving $A(u) = L(u)$.

The local view term $L(u)$ can be further broken up into feature terms, corresponding to the univariate gaussian functions of spatial variables. If a feature term is tuned to a landmark of a type that is not present in the current environment, then that feature term will drop out of the equation for $L(u)$. We also allow feature terms to drop out when there are not enough landmarks of a type to fill the place unit's local view. For example, a place unit tuned to spatial variables of two cylinders in the two-landmarks task will, when faced with one cylinder, still show a significant amount of activity in its corresponding place field. One might ask why such a place unit will not show two place fields, corresponding to matching the one observed cylinder to each of the two remembered cylinders. It will only show a single field because it is constrained by the $C(u)$ term; in only one location in the environment will the path integrator coordinates match its remembered coordinates.

4.2. Training the model

In order for an environment to be familiar, the simulated animal (hereafter *sim-animal*) must have recruited enough place units that with high probability every location will activate some minimum required number. For the experiments reported here, it does not matter how these units were created. Therefore, we "train" the sim-animal using a method that is behaviorally unrealistic, but quickly produces enough place units to cover the space.

Training proceeds as follows: The sim-animal is deposited at a random spot in the arena, with a random heading. It is told its head direction Φ_h and path integrator coordinates[2] $\langle x_p, y_p \rangle$. Distances and angles to all visible landmarks are calculated and fed to the place units. If fewer than 20 units become active, a new place unit is recruited. Its field is centered at the present location by the following steps:

1. The response function to path integrator coordinates is tuned to the current coordinate values $\langle x_p, y_p \rangle$.

2. Two visible landmarks are chosen at random, and the response functions for allocentric bearings ϕ_i and distances r_i are tuned to their positions. We do not require that the two landmarks be distinct. Hence, for some units the two will be identical, making the unit effectively tuned to a single landmark.

[2] We assume the origin of the path integrator coordinate system to be at the food location. The actual location of the origin is irrelevant in the tasks described in this chapter, as long as it is consistent across trials.

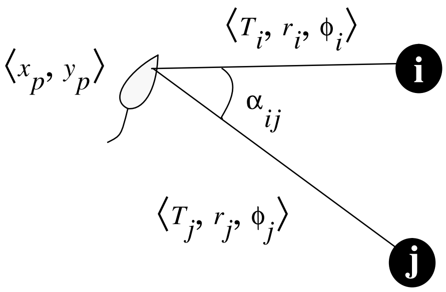

Figure 13: *Spatial variables used in tuning a place cell to two landmarks i and j at path integrator coordinates* $\langle x_p, y_p \rangle$.

3. Another two (possibly different) landmarks are chosen, and the response function for retinal angle α_{ij} is tuned to the difference in their bearings.

Figure 13 shows the relations of the spatial variables used in tuning a typical place cell.

The sim-animal is then moved to another location chosen at random and the process is repeated. Training continues until the sim-animal reaches criterion, defined as 20 locations visited in succession without recruiting a new place unit. For the Collett et al. tasks, approximately 10,000 units were required to adequately cover the environment.

4.3. Entering a Familiar Environment

On entering a familiar environment, the animal may need to reset either or both of its head direction and path integrator modules. If the animal has been disoriented (for example by being spun rapidly) before entering the environment, it may know that its head direction is incorrect. Similarly, if it has been passively moved (for example by being carried by an experimenter), it may know that its path integrator coordinates are incorrect. However, even if it believes its path integrator coordinates and head direction to be correct, it may still reset either or both to align itself with familiar cues.

Our model uses a measure of the *consistency* of the place code to determine whether the sim-animal is aligned with the familiar cues:

$$\mathcal{C} = \frac{\sum_u A(u)}{\sum_u C(u)}$$

When the current local view is compatible with the local view remembered for these path integrator coordinates, or conversely, when the current path integrator coordinates match those evoked by the current local view, $\mathcal{C}$ will be high. When the two are are not compatible, $\mathcal{C}$ will be low.

In our model, on entering a familiar environment, the sim-animal takes note of its current path integrator coordinates and head direction. It then attempts the following sequence:

1. If it does not know its path integrator coordinates, those must be reset.

2. If it does not know its head direction, that must be reset.

3. It then attempts a head direction reset, and, if the consistency $\mathcal{C}$ is improved, it uses the newly-derived value in place of the old one.

4. It checks whether a combined path integrator and head direction reset will improve the consistency. It does this by letting allocentric bearing terms drop out of the activation equation (followed necessarily by a head direction reset). If $\mathcal{C}$ is again improved, it uses the new values, otherwise it restores the remembered values.

In summary, the sim-animal performs a series of path integrator and head-direction resets, and when these improve the consistency of the place code, it uses the new values. In future work we hope to simplify this algorithm, but since it need only be performed on initial entry into an environment, the cognitive cost of even the present version does not seem unreasonable.

For the experiments discussed here, because the sim-animal has been placed in the environment at a random location, it does not know its path integrator coordinates. It therefore always resets the path integrator (step 1). However, the sim-animal is not disoriented, so it always knows its head direction at the start of a trial (step 2). It may still reset its head direction, depending on the environmental cues (steps 3-4). Once the sim-animal has determined its location and head direction, computing the trajectory to the goal can be done by vector subtraction.

4.4. Operations to Enforce Consistency Among Representations

We have defined several operations in the model that bring inconsistent internal and external cues into correspondence. They have been constrained to replicate a variety of behavioral and neurophysiological data [37, 38, 40, 41]. Here we will concentrate on just two operations: reset of the path integrator, and reset of the head direction module.

4.4.1. Reset of the Path Integrator

Reset of the path integrator is implemented by a parallel relaxation process that brings the locations represented by the place code and the path integrator coordinates into correspondence.

The simulator generates a perceptual tuple $\langle T_i, r_i, \theta_i \rangle$ for each landmark visible from the animal's present location. The local view component converts egocentric to allocentric bearings, producing $\langle T_i, r_i, \phi_i \rangle$. These tuples are input to the place units, whose activation values reflect the match between the tuples seen and the values to which each unit is tuned. When multiple landmarks of the required type are available, a unit's activation value is the maximum obtainable for any choice of landmarks.[3] We apply a dynamic thresholding operation (varying δ) to select a set of active place units of reasonable size.

Each place unit maintains an association between a specific location in space and the visual cues at that location when the unit was recruited. Activation of unit u can therefore be taken as evidence that the sim-animal's location is near the path integrator coordinates to which unit u is tuned. The set of active units therefore serves as a coarse-coded representation [10] of the sim-animal's current position. (In actual rats, Wilson and McNaughton have shown that the collective activity pattern of roughly 100 hippocampal place cells gives a highly accurate estimate of position in a rectangular arena [42].)

Because locations sharing visual features will be perceptually similar to the animal, the set of active place units might not represent a unique point in space. Figure 14a shows the place field centers of all units activated when the trained sim-animal was deposited at a random spot in the three-landmark environment. The active units have place fields centered at widely varying locations, meaning the place code is *incoherent*.

Coherency is achieved by parallel relaxation. First, weighted mean coordinate values are calculated based on the active place units. The contribution of each unit to the mean is scaled by its activation level. The path integrator is initialized to this mean value. Now path integrator coordinates can be used to further constrain the set of active place units. The width of each unit's path integrator tuning curve was set at some large value initially, and is tightened as the relaxation proceeds.

At each step of the relaxation process, place activation values are recalculated. (It is important that this be done for all units, not just those currently active, because some units that were below the dynamic threshold may rise above it given the new input from the path integrator.) Units whose place fields are too distant from the mean coordinate estimate are inactivated; this in turn causes the mean to shift a bit toward the centroid of the remaining active units. Repeating the process while successively tightening the tuning to path integrator coordinates causes the place code to quickly converge to a state in which all active units have highly overlapped place fields, meaning they are in agreement as to the animal's location. The place code is now *coherent*, as shown in Figure 14b.

4.4.2. Reset of the Head-Direction Module

The procedure for resetting head direction is as follows. Every place unit has associated with it the learned allocentric bearings of two landmarks. These are two of the feature terms of $L(u)$.

For each active place unit, the model computes the difference between the learned allocentric bearing $\phi_i(u)$ and the egocentric bearing θ_k of whichever landmark k it is presently using to compute that feature term. This gives roughly twice as many bearing differences as the number of active units. These are histogrammed, and the histogram segmented to find the peaks. The model then

[3] Note that this is an instance of the binding problem.

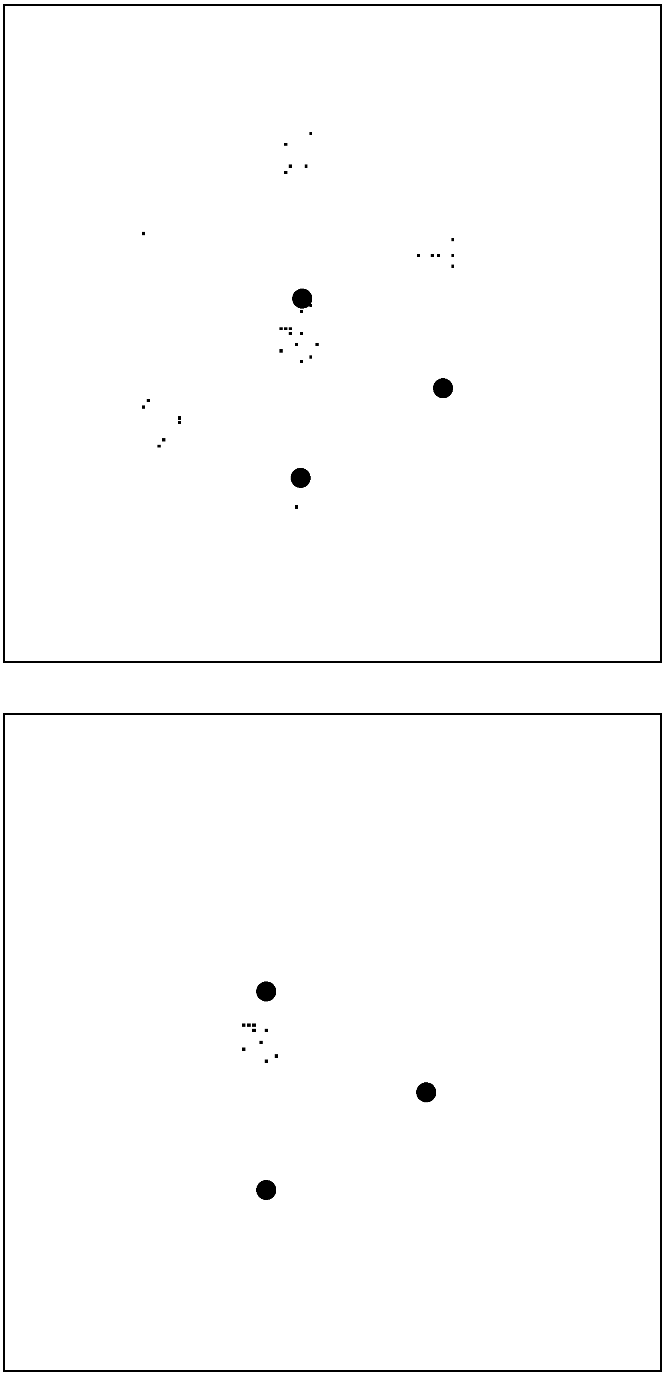

Figure 14: *Incoherency in the place code when the animal is first reintroduced into a familiar environment. The sim-animal was trained and tested on the three landmark task. Large filled circles are landmarks. Each small dot marks the center of the place field (defined by path integrator coordinates) of an active place unit. (top) Incoherent initial state: the fields are widely dispersed. (bottom) Coherent final state.*

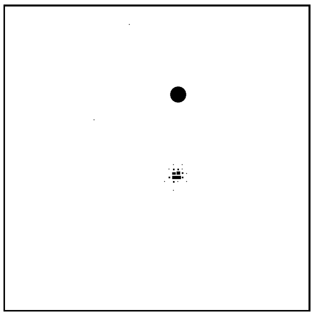

Figure 15: *Simulation of the one-landmark task. Blobs show distribution of goal prediction from various starting locations. Compare with Figure 1.*

probabilistically chooses one of these peaks for the new heading estimate Φ_h.

There are probably more natural ways to compute head direction from visual landmarks. The above procedure is easy to program and captures the essential claim of our theory: that the animal is sensitive to discrepancies between predicted and perceived landmark bearings.

5. Correspondence with data

Our model does not produce explicit behavior such as locomotion or search for a food reward. Thus, we cannot directly reproduce the sequence of motor actions a gerbil goes through to travel to a presumed goal location and dig about for buried food. Instead, to replicate Collett et al.'s behavioral results we run the model's self-localization procedure at 100 randomly-chosen locations in the environment. On each trial the model produces a position estimate, and from that it derives an estimate of the goal location. We plot 2D histograms of goal location estimates to compare our results with the gerbil behavioral data.

Collett et al. examined distribution of search effort over trials lasting 30 to 120 seconds. During this time, the animals must physically move between the different search locations. This is one reason why Collett et al.'s histograms should be noisier than ours and show broader peaks. In addition, it is likely that the animals were aware that the environment had changed [27, 36], so some of the time may have been spent exploring rather than searching for food. We do not model exploration. The main point, though, is that the analogy between their search time distributions and our goal prediction distributions is straightforward, and there is good agreement between the two sets of histograms.[4]

Figure 15 shows the distribution of goal predictions when the sim-animal was trained and tested on one landmark. As previously explained, the model was able to correctly estimate the bearing of the goal location from the landmark because

[4]Each of histograms from our simulations (Figures 15 through 20) is plotted to the same scale as the corresponding experimental figure (Figures 1 through 7). However, we have used a finer bin size to show additional detail.

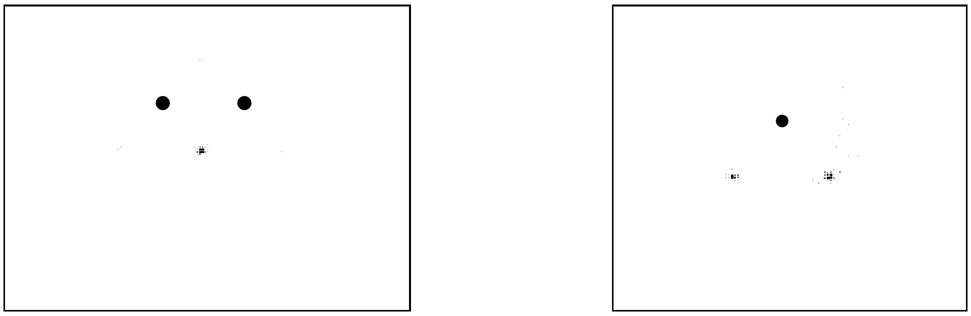

(a) (b)

Figure 16: *Simulation of (a) the two-landmark task and (b) the two-minus-one task. Compare with Figures 2 and 3.*

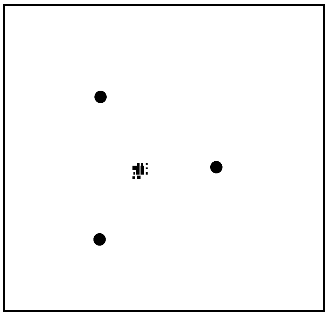

Figure 17: *Simulation results: distribution of goal predictions for the three-landmark task. Compare with Figure 4.*

its place units are tuned to allocentric bearings derived from the internal compass.

When trained on the two-landmark array, the model again predicts a single goal location (Figure 16a.) Because place units tune to allocentric bearing information, it does not search at the mirror image of this spot. When one of the two landmarks is removed, the model distributes its goal predictions between two locations (Figure 16b). Although the visual environment is identical to that of the one-landmark task, the place units are tuned differently and thus the goal predictions differ.

When trained on the three-landmark array, the model also produces a single goal estimate (Figure 17), and with one landmark removed, the goal predictions are still concentrated at a single location (Figure 18a). With two landmarks removed the model makes three goal predictions (Figure 18b), as did the gerbils.

When presented with an extraneous landmark, the model uses allocentric bearing information to distinguish the correct from the inverted triangle (Figure 19). When the triangle is stretched, it ignores the third landmark and searches only at the correct distance and bearing to the other two (Figure 20).

In the critical experiments that vector voting was unable to handle, the split-

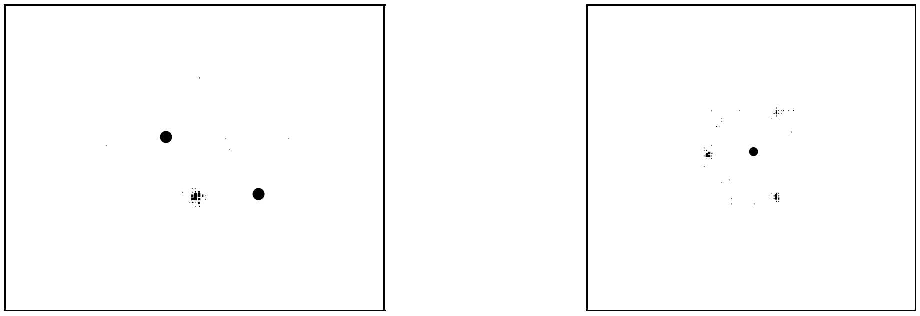

(a) (b)

Figure 18: *Simulation of (a) the three-minus-one task and (b) the three-minus-two task. Compare with Figure 5.*

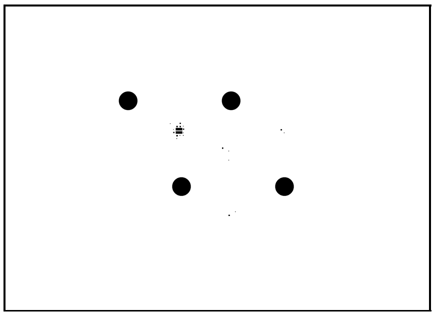

Figure 19: *Simulation of the three-plus-1 task. Compare with Figure 6.*

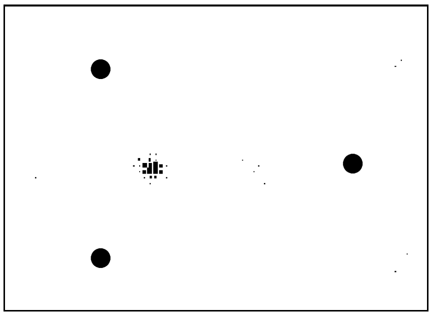

Figure 20: *Simulation of the stretched triangle task. Compare with Figure 7.*

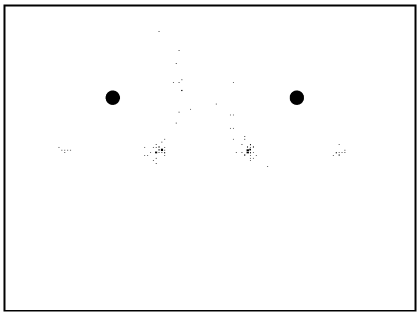

Figure 21: *Simulation of the split-array task. Compare with Figure 9.*

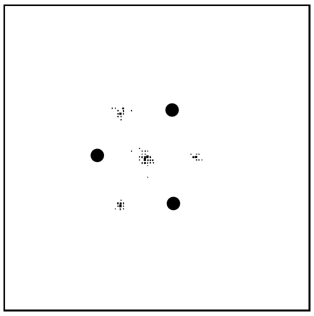

Figure 22: *Simulation of the rotated-triangle task. Compare with Figure 11.*

array task and rotated-triangle tasks, our model agrees with Collett et al.'s observations.

In the split-array task, the sim-animal concentrates search at the two locations interior to the array (Figure 21). Our model produces a concentration at the interior locations of the split-array for two reasons.

1. **Calculation of mean coordinate value.** The model estimates the sim-animal's position by computing a weighted mean coordinate value using all active place units. The result will normally be near the center of the array. The two exterior goal locations predicted by vector voting will be further from the mean than the two interior positions. Thus, place units coding for exterior goal locations are much more likely to be rejected as outliers by the parallel relaxation process, leading the model to prefer predictions near the centroid of the candidate set.

2. **Partial activation of place units.** In the split array condition, no place unit tuned to two distinct landmarks will see an exact distance and bearing match for both.[5] However, because place units are radial-basis functions

[5]Units tuned to single landmarks can still get perfect matches, but there will be an approximately equal number of these for each of the four candidate goal locations, so their effects cancel out.

with a smooth gaussian fall-off, those tuned to locations interior to the split array can still have an exact match for one landmark and a partial match for the other. Units tuned to exterior locations cannot get even a partial match, and will thus be at a disadvantage.

In the rotated-triangle task, the sim-animal divides its predictions between four locations: the three exterior predicted by vector voting and one interior at the center of the triangle. We believe the interior and exterior goal predictions occur because of a divergence in the possible outcomes of the algorithm described in Section 4.3..

1. **Interior goal predictions.** On some trials, resetting the head direction produces an improvement in consistency (steps 3-4). On trials in which the head direction is reset, the sim-animal will effectively see the three-landmark task (a correctly oriented triangle) and will predict the goal to be in the center of the array.

   It is unclear how such a reset could later be undone, yet Collett et al. report that the gerbils first searched in the center and then proceeded to search at exterior locations. Further work is need to determine whether our model is compatible with this observation.

2. **Exterior goal predictions.** On other trials, the improvement in consistency is too small to override the initial values (from step 1). On trials in which the head direction is not reset, the animal will use pairs of landmarks in the correct orientation to localize itself and will search the exterior locations as predicted by the vector voting hypothesis.

Our simulations are deterministic and might be sensitive to the starting position of the sim-animal. We do not know whether there are regions of the space in which one goal prediction is preferred over another or whether all goal predictions are mixed together randomly. It is also possible that adding noise to the system will reduce any systematic dependence on initial position. We are currently exploring these questions.

6. Our model and the vector voting hypothesis

According to the vector voting hypothesis, the animal remembers vectors from each landmark to the goal, and applies these n remembered vectors to each of the m perceived landmarks, generating $n \cdot m$ goal predictions. It then distributes its search effort over the locations with the most votes. The animal must somehow keep track of multiple candidate locations in order to tally votes for each of them.

Our self-localization operation can also be understood as a distributed form of vector voting. The goal location is not what is being determined; rather the animal's own location is. Each place unit can be viewed as voting for the animal being at the path integrator coordinates associated with that unit. Calculating the mean coordinate value for the set of active place units serves indirectly as a vote tallying mechanism. Deactivating units that are proposing locations too far from the current mean eliminates the decisively outvoted candidates.

The parallel relaxation algorithm always produces a single winner, while vector voting allows for ties. The latter therefore directly supports a division of search effort between multiple goal locations during the course of a single trial, which is what gerbils actually do. Our own model would require a different mechanism to produce alternation among tied candidates.

In the experiments reported here, the sim-animal was started from multiple locations and we simply plotted the distribution of goal predictions. A mechanism that could produce alternation among search candidates, and is supported by some preliminary simulations, is to inhibit the currently active place units when a search proves unsuccessful. If the sim-animal then performs a new relaxation based on currently perceived landmarks, a new set of place units will become active, indicating that its present position is not actually the goal as previously believed, but instead some other spot displaced from the goal. Calculating a trajectory from this spot to the goal (whose coordinates are always known) will take the animal to the next candidate search location.

7. Predictions

Our theory makes predictions about both behavior and neurophysiological activity. We give three examples here.

1. Viewing the self-localization procedure. Place unit activity in our model reflects where the animal believes it is. During normal navigation, the place code at any time is coherent, i.e. all active place cells have fields near each other. The self-localization procedure requires that many place cells be briefly activated at the beginning of a trial, before coherency has been enforced. We therefore predict that in a brief initial phase of each trial, the hippocampal activity patterns of real animals will form an incoherent place code, i.e., there will be simultaneously active place cells with fields far from each other. In ambiguous environments such as the split array task we might see a multi-modal distribution of activity during the transitory initial period.

2. Shifting goal locations. When the animal abandons one search location in favor of another, our theory requires that it first change its estimate of its own location. The place code must shift accordingly. Our theory therefore predicts a change in the hippocampal activity pattern will be observed *just prior* to the animal beginning a move to a new search location.

3. Place cell activity in the split array task. McNaughton, Knierim, and Wilson [15] recently proposed an alternative theory that also combines visual landmarks with path integration. Place cells in their formulation are only tuned to single landmarks. When trained on the two-landmark array, there would be two populations of cells, one tuned to the east landmark and one tuned to the west one. In the split array task, these populations would dissociate. This conflicts with the prediction of our own theory: that a single population of units will be active at both search locations.

In McNaughton et al.'s version, if place fields are small relative to the distance between the landmarks, the animal will exhibit distinct place codes as it searches the two locations during a probe trial. If, on the other hand, place fields are large enough that the fields of the two populations still overlap when the landmark array is stretched, McNaughton et al. predict that the active units should still be dissociable into two populations, with place field centers shifted by an amount equal to the shift in the landmarks [15]. That is, units tuned to the east landmark will have their fields shifted eastward relative to the center of the array, and conversely for units tuned to the west landmark. Demonstrating this shift may be difficult, however, because the animal would have to spend substantial time traveling around the arena in the split landmark condition in order for the experimenter to accurately measure firing rates to map the place fields. During this time the animal might modify its place code to take the new landmark configuration into account.

If our theory is correct, when hippocampal activity is recorded during search in the split-array task, as the animal alternates between goal predictions all place cells coding for the goal location will show two disjoint fields, or else a large field with two peaks, depending on the width of the fields relative to the inter-landmark distance. In either case, the population will not dissociate.

8. Discussion

Animals clearly use configurations of visual cues to navigate. In the tasks described here, rodents used range and bearing to nearby landmarks to locate food. However, their behavior cannot be accounted for by visual cues alone. The same environment can produce radically different behavior depending on the animal's past experience. For example, the two-minus-one and three-minus-two probe tasks use identical environments to the one-landmark task. In the two-minus-one task, gerbils searched in two locations, and in the three-minus-two task they searched in three locations, while in the one-landmark-task they searched in only one location. Thus we see the role of visual learning in guiding navigation behavior.

Our model relies on multiple representations of space: local views, an internal compass, coarse coded place units, and Cartesian coordinates maintained by path integration. Interactions among these representations maintain a consistent relationship between the visual environment and the animal's internal state. In addition to the data presented here, our model successfully replicates other neurophysiological data, such as from Muller et al. [20, 21] and Sharp et al. [32] on rats in small circular arenas, and behavioral experiments, such as data from Cheng, Gallistel, and Margules [5, 13] on rats in symmetric environments.

We are in the process of extending the model to incorporate yet more data (for example, differences in salience among cues, both for head direction [11, 15] and self-localization [6]) as well as considering how rodents really learn novel environments.

In novel environments, drift in the head direction system [15] cannot be detected by taking bearings to landmarks. Head direction errors can accumulate, leading to errors in path integration. And since the environment is unfamiliar, the place code cannot reset the path integrator. Rodents must have some way to correct for path integration errors simultaneous with tuning their place cells during exploration. Their solution to this problem would be of value to mobile robots.

References

[1] L. Abraham, M. Potegal, and S. Miller. Evidence for caudate nucleus involvement in an egocentric spatial task: Return from passive transport. *Physiological Psychology*, 11(1):11–17, 1983.

[2] N. Burgess, J. O'Keefe, and M. Recce. Using hippocampal 'place cells' for navigation, exploiting phase coding. In S. Hanson, J. Cowan, and L. Giles, editors, *Advances in Neural Information Processing Systems 5*. Morgan Kaufmann, San Mateo, CA, 1993.

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

[4] L. L. Chen. *Head-directional information processing in the the rat posterior cortical areas*. PhD thesis, University of Colorado, 1991.

[5] K. Cheng. A purely geometric module in the rat's spatial representation. *Cognition*, 23:149–178, 1986.

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

[7] R. J. Douglas. Spontaneous alternation behavior and the brain. In W. N. Dember and C. L. Richman, editors, *Spontaneous Alternation Behavior*, chapter 5, pages 73–108. Springer-Verlag, New York, 1989.

[8] A. S. Etienne. The control of short-distance homing in the golden hamster. In P. Ellen and C. Thinus-Blanc, editors, *Cognitive Processes and Spatial Orientation in Animals and Man*, pages 233–251. Martinus Nijhoff Publishers, Boston, 1987.

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

[10] G. E. Hinton, J. L. McClelland, and D. E. Rummelhart. Distributed representations. In D. E. Rummelhart and J. L. McClelland, editors, *Distribted Processing: Explorations in the Microstructure of Cognition*, volume 1, pages 77–109. MIT Press, 1986.

[11] J. J. Knierim, B. L. McNaughton, C. Duffield, and J. Bliss. On the binding of hippocampal place fields to the inertial orientation system. *Society for Neuroscience Abstracts*, 19:795, 1993.

[12] B. Leonard and B. L. McNaughton. Spatial representation in the rat: conceptual, behavioral, and neurophysiological perspectives. In R. P. Kesner and D. S. Olton, editors, *Neurobiology of Comparative Cognition*, chapter 13, pages 363–422. Lawrence Erlbaum Associates, Hillsdale, NJ, 1990.

[13] J. Margules and C. R. Gallistel. Heading in the rat: Determination by environmental shape. *Animal Learning and Behavior*, 16(4):404–410, 1988.

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

[15] B. L. McNaughton, J. J. Knierim, and M. A. Wilson. Vector encoding and the vestibular foundations of spatial cognition: Neurophysiological and computational mechanisms. In M. Gazzaniga, editor, *The Cognitive Neurosicences*. MIT Press, Boston, 1993.

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

[17] B. L. McNaughton, E. J. Markus, M. A. Wilson, and J. J. Knierim. Familiar landmarks can correct for cumulative error in the inertially based dead-reckoning system. *Society for Neuroscience Abstracts*, 19:795, 1993.

[18] M. L. Mittelstaedt and H. Mittelstaedt. Homing by path integration in a mammal. *Naturwissenschaften*, 67:566–567, 1980.

[19] S. J. Y. Mizumori and J. D. Williams. Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *Journal of Neuroscience*, 13(9):4015–4028, 1993.

[20] R. U. Muller, J. L. Kubie, E. M. Bostock, J. S. Taube, and G. J. Quirk. Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In J. Paillard, editor, *Brain and Space*, chapter 17, pages 296–333. Oxford University Press, New York, 1991.

[21] R. U. Muller, J. L. Kubie, and J. B. Ranck Jr. Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *Journal of Neuroscience*, 7:1935–1950, 1987.

[22] J. O'Keefe. Computations the hippocampus might perform. In L. Nadel, L. Cooper, P. Culicover, and R. M. Harnish, editors, *Neural Connections, Mental Computation*, chapter 8, pages 225–284. MIT Press, Cambridge, MA, 1989.

[23] J. O'Keefe and D. H. Conway. Hippocampal place units in the freely moving rat: Why they fire where they fire. *Experimental Brain Research*, 31:573–590, 1978.

[24] J. O'Keefe and L. Nadel. *The hippocampus as a cognitive map*. Clarendon Press, Oxford, 1978.

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

[26] M. Potegal. The caudatenucleus egocentric localization system. *Acta Neurobiological Experiments*, 32:479–494, 1972.

[27] B. Poucet, N. Chapuis, M. Durup, and C. Thinus-Blanc. A study of exploratory behavior as an index of spatial knowledge in hamsters. *Animal Learning and Behavior*, 14(1):93–100, 1986.

[28] G. J. Quirk, R. U. Muller, and J. L. Kubie. The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience*, 10(6):2008–2017, June 1990.

[29] N. A. Schmajuk and H. T. Blair. Place learning and the dynamics of spatial navigation: A neural network approach. *Adaptive Behavior*, 1:355–387, 1993.

[30] M. L. Shapiro and P. A. Hetherington. A simple network model simulates hippocampal place fields: Parametric analyses and physiological predictions. *Behavioral Neuroscience*, 107(1):34–50, 1993.

[31] P. E. Sharp. Computer simulation of hippocampal place cells. *Psychobiology*, 19(2):103–115, 1991.

[32] P. E. Sharp, J. L. Kubie, and R. U. Muller. Firing properties of hippocampal neurons in a visually symmetrical environment: Contributions of multiple sensory cues and mnemonic processes. *Journal of Neuroscience*, 10(9):3093–3105, Sep 1990.

[33] J. S. Taube. Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, 15(1), January 1995.

[34] J. S. Taube, R. I. Muller, and J. B. Ranck, Jr. Head direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *Journal of Neuroscience*, 10:436–447, 1990.

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

[36] C. Thinus-Blanc, L. Bouzouba, K. Chaix, N. Chapuis, M. Durup, and B. Poucet. A study of spatial parameters encoded during exploration in hamsters. *Journal of Experimental Psychology*, 13(4):418–427, 1987.

[37] D. S. Touretzky and A. D. Redish. Landmark arrays and the hippocampal cognitive map. In L. Niklasson and M. Boden, editors, *Current trends in connectionism - Proceedings of the 1995 Swedish Conference on Connectionism*, pages 1–13. Lawrence Erlbaum Associates, Hillsdale NJ, 1995. in press.

[38] D. S. Touretzky, H. S. Wan, and A. D. Redish. Neural representation of space in rats and robots. In Jacek M. Zurada, Robert J. Marks II, and C. J. Robinson, editors, *Computational Intelligence: Imitating Life*, pages 57–68. IEEE Press, 1994.

[39] A. Treves, O. Miglino, and D. Parisi. Rats, nets, maps, and the emergence of place cells. *Psychobiology*, 20(1):1–8, 1992.

[40] H. S. Wan, D. S. Touretzky, and A. D. Redish. Computing goal locations from place codes. In *Proceedings of the 16th annual conference of the Cognitive Science society*, pages 922–927. Lawrence Erlbaum Associates, Hillsdale NJ, 1994.

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

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

Acknowledgements

David Redish was supported by the National Science Foundation and the Neural Processes in Cognition program. Both authors are members of the Center for the Neural Basis of Cognition at Carnegie Mellon. Hank Wan also contributed to this work. We thank Bruce McNaughton and Neil Burgess, as well as two anonymous reviewers, for helpful discussions. We also thank T. S. Collett and Springer-Verlag publishers for their generous permission to reproduce the figures from [6].