

To appear in: L. Niklasson and M. Boden (eds.), Current Trends in Connectionism - Proceedings of the 1995 Swedish Conference on Connectionism. Hillsdale, NJ: Erlbaum. Copyright (c) 1995 Lawrence Erlbaum Associates.

Landmark Arrays and the Hippocampal Cognitive Map

David S. Touretzky and A. David Redish
Computer Science Department and
Center for the Neural Basis of Cognition
Carnegie Mellon University

Abstract

Animal navigation can be simultaneously studied externally, at the level of actions in the world, and internally, at the level of individual neuron firings. In this paper we describe a computational theory of rodent navigation that is constrained by data at both these levels. We report results of computer simulations of a split landmark array task that lead to interesting predictions about the behavior of hippocampal place cells.

1 Introduction

Landmark-based navigation is a rich domain for exploring issues of representation and processing in neural systems. At the behavioral level, there is a wealth of data on how animals use landmarks to locate food or return to their homes. At the neurophysiological level, the responses of hippocampal pyramidal cells, the well known “place” cells (O’Keefe and Dostrovsky, 1971; Leonard and McNaughton, 1990) and of head direction cells in thalamus (Mizumori and Williams, 1993) and the subicular complex (Taube et al., 1990) provide striking neural correlates to behavioral variables.

Systems-level theories fill the gap between these two modes of description. To construct such a theory for navigation, we must first determine a set of computations at some reasonably abstract level that can reproduce the observed behavior, and then show how these computations could be realized in neural tissue. We are of course a long way from this goal. However, in the present paper we describe a theory of landmark-based 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 animal navigation literature and make predictions about place cell responses in novel situations. Portions of the model have also been implemented on a mobile robot.

2 The Split Landmark Array Task

Collett, Cartwright, and Smith (1986) trained gerbils to search for a food reward at a fixed position relative to an array of identical cylindrical landmarks. The array was translated but not rotated from trial to trial, and the animals were released




Figure 1: Distribution of gerbils' search efforts; schematized version of experiments from (Collett et al., 1986), after (Leonard and McNaughton, 1990). Small dots are landmarks; large circles are areas where the majority of search time was spent. (a) training on one landmark; (b) training on two landmarks; (c) probe trial deleting one of the two landmarks; (d) probe trial doubling the distance between the two landmarks.




Figure 2: (a) Training on two landmarks; small triangle denotes food reward (sunflower seed) at the goal location. (b) Probe trial showing how vector voting determines the location to be searched.

from different starting points to ensure that the landmarks would provide the only reliable cues to the reward location. After training to criterion, probe trials were introduced in which the food was absent and the distribution of the animals' search efforts was measured. Figure 1 shows the results of four such experiments, in schematic form. The small black circles are landmarks, and the larger shaded circles denote peaks in the spatial distribution of search effort.

For a single cylinder (Figure 1a), the animals learned to search at the correct bearing and distance from that landmark. Bearing information was presumably measured with respect to the internal compass, because the arena was designed to minimize stimuli that could serve as directional cues. The room walls were painted black, and there was only a single overhead light, in the center of the ceiling. The light illuminated a circular region of the floor but left the walls in shadow.

Experiments with another group of animals using pairs of cylinders (Figure 1b) produced similar results: well-trained animals would go directly to the goal location. Now Collett et al. could test the animal's representation of the environment by introducing occasional probe trials with modifications to the landmark array.

When one landmark was missing on a probe trial (Figure 1c), the gerbils searched



Figure 3: Results of simulations where the model was trained on (a) one landmark or (b) two landmarks. Compare with Figures 1a and b, respectively.




Figure 4: Result when model was trained on two landmarks and then tested on one. Compare with Figure 1c.




Figure 5: Result when model was trained on two landmarks and tested with the landmarks twice as far apart (the "split landmark array" case.) Compare with Figure 1d.

alternately in two locations, each at the correct bearing and distance from one of the landmarks they had observed during training. It was as if they were binding the cylinder to one and then the other of the two remembered landmarks. Note that the physical environment in Figure 1c is the same as in Figure 1a; the difference in search behavior reflects the different training the two groups received.

When the distance between the two cylinders was doubled on a probe trial (the “split landmark array” case, Figure 1d), the animals also searched in two locations, each associated with one of the two landmarks. They did not search in the center of the expanded array.

Collett et al. proposed a mechanism that could account for these results, and several others involving more complex arrays of landmarks. We shall refer to it as the “vector voting hypothesis.” According to this hypothesis, when the gerbils reach the goal location and find food there, they note the vector between each landmark and their present position. Then, at the beginning of a new trial, when they first emerge from the “start box” and see the landmark array, they apply every learned vector to every currently perceived landmark, as in Figure 2. The locations receiving the most votes are the ones they choose to search.

This simple mechanism explains why gerbils search only one location in Figures 1a and 1b, and two locations in Figure 1c. However, the vector voting hypothesis alone does not account for the result in Figure 1d, where the gerbils searched in two locations instead of four. The four sites should receive one vote each when the landmark array is split. Collett et al. concluded that the gerbils must be using their perception of the entire array to distinguish the east from the west landmark, and binding remembered vectors to only the corresponding landmarks during the probe trials.

This explanation accounts for the data, but it introduces a binding mechanism whose characteristics are left unspecified. What internal representations permit the use of array configuration cues to disambiguate individual landmarks? How are such representations formed, and what are their limitations?

In a later section we will describe an apparent alternative to vector voting that reproduces all the effects in Figure 1 and a variety of similar experiments without resorting to explicit binding mechanisms to disambiguate landmarks. Subsequently we shall show that our mechanism is in fact a distributed implementation of the vector voting hypothesis.

3 Overview of Our Theory

Figure 6 is a diagram of our theory. Each of these blocks denotes a functional subsystem, and should not be presumed to correspond to a single anatomical site. Visual Input (upper right) provides range and egocentric bearing information (r_i, θ_i) for a set of currently visible landmarks. The animal’s Head Direction, i.e. its “internal compass,” is updated using vestibular cues; the current head direction estimate is the variable Φ_h . In the block labeled Local View, allocentric landmark

Figure 6: Components of our theory of rodent navigation.

bearings are calculated as $\phi_i = \theta_i + \Phi_h$.

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 cues and an efferent copy of motor commands. There is evidence for path integration abilities in a wide variety of animals; see (Gallistel, 1990) for general information and (Mittelstaedt and Mittelstaedt, 1980) for results on gerbils. Unfortunately, the neural substrate of path integration is presently unknown in any animal.

The role of the place units¹ in our theory is to maintain an association between perceived landmark positions and path integrator coordinates, so that either can be reconstructed from the other. This explains several questions raised by neural recording experiments using rats:

- *How do rats self-localize when reintroduced into a familiar environment at a random spot?* Our theory says they use visual landmarks to activate a place code, which in turn evokes a set of coordinates used to reinitialize the path integrator.
- *How are place fields able to persist in the dark?* The path integrator is updated with each motion the animal makes. Our theory says that the output

¹We refer to model components as “units,” reserving the term “cells” for biological neurons. We do not assume a one-to-one correspondence between the model’s place units and single hippocampal pyramidal cells; our units are probably closer to cell assemblies..

of the path integrator may be used to drive place cells. Errors will eventually accumulate, but the system may be kept reasonably calibrated if other sorts of cues are available, such as tactile information.

- *How is drift in the path integration and head direction systems corrected?*
The place units representing a location keep track of the allocentric bearings and distances of landmarks visible from that location. If landmarks appear at the correct distances but their bearings are off by a consistent amount, this indicates drift in the internal compass. If the path integrator's output differs somewhat from the coordinates derived from the place units (see below), this indicates drift in the path integrator.

For further details on these computations, see (Wan et al., 1994a; Touretzky et al., 1994; Wan et al., 1994b).

4 Training the Model

At the start of a simulation the model must be trained on the environment. Enough place units must be recruited so that, with high probability, every location will activate some minimum number of units. Training proceeds as follows. The simulated animal is deposited at a random spot in the arena. Distances and angles to all visible landmarks are calculated and fed to the place units. If fewer than ten units are activated, a new place unit is recruited with its field centered on the present location. Two landmarks are chosen at random from the currently visible set and the new place unit is tuned to their allocentric bearings and distances. Again two landmarks are chosen at random, possibly the same two as before, and the unit is tuned to the difference in their bearings. The unit's path integrator coordinates $\langle x_p, y_p \rangle$ are also tuned to the present location's coordinates. The animal is then deposited at a new random spot, and the process repeats. Training terminates when ten spots have been visited in succession without recruiting a new unit. The simulations reported here require approximately 3000 units.

5 Self-Localization by Relaxation

In our theory, upon reintroduction into a familiar environment the animal determines its position by observing visual landmarks. Since it also possesses a Cartesian representation of the environment in which the goal is at known coordinates, once it has self-localized it has access to its own coordinates and can calculate a trajectory to the goal by vector subtraction. We will now describe how self-localization is achieved.

The place units in our model are radial basis functions tuned to combinations of several quantities: distances and allocentric bearings for two landmarks chosen at random, (r_i, ϕ_i) and (r_j, ϕ_j) , plus the retinal angle between a possibly different random pair of landmarks, $\theta_m - \theta_n$. Retinal angle information is useful when the




Figure 7: Sample place field after training on the two-landmark environment.

internal compass has been disrupted, making allocentric bearings unavailable until head direction can be determined.

Place units are also tuned to coordinate values $\langle x_p, y_p \rangle$ output by the path integrator. When the animal is first released in the arena, path integrator coordinates are unavailable, so this term temporarily drops out of the equation for the place unit's activation. Conversely, when navigating in the dark the visual landmark terms drop out, and the place units are driven solely by the path integrator. Figure 7 shows the size and shape of a typical place field.

At the start of a probe trial, the simulated gerbil looks around and generates a set of perceptual vectors $\langle T_i, r_i, \phi_i \rangle$ describing the types and positions of the currently visible landmarks. These vectors provide input to the place units, whose activation levels reflect how well the vectors match the parameters each unit is tuned for. We apply a dynamic thresholding operation to select a set of "active" place units.

Each unit has associated with it coordinate values $\langle x_u, y_u \rangle$ representing the center of that unit's place field. The set of active units forms a coarse-coded representation (Hinton et al., 1986) of the animal's location in Cartesian space.

Because each place unit monitors only a few features of the environment, and also because multiple locations may appear perceptually similar to the animal, the set of initially active place units might not code unanimously for a unique region of space. Therefore we apply parallel relaxation to enforce consistency in the place code. First, a mean coordinate value is calculated for all active place units. The contribution of each unit to the position estimate is weighted by the unit's activation level. Then units whose coordinates are too distant from the mean are inhibited, and the mean is recomputed. This process quickly converges to a state in which all active units are in rough agreement as to the animal's coordinates,

forming a consistent place code.

6 Simulation Results

Figure 3 shows the predicted distribution of search effort when the model was trained on either one or two landmarks, and tested under the same conditions. Since our model doesn't actually spend time searching neighborhoods of points the way gerbils do, instead of plotting distribution of search time we plotted the distribution of goal location estimates for 100 random starting locations. The analogy between the two distributions is straightforward, and the agreement between the model and Collett et al.'s observations is excellent.

In ambiguous trials such as Figures 1c-d, the initial distribution of coordinate values for active place units will be multi-modal. Through the parallel relaxation process, a subset of place units is selected that provides consistent support for one of the peaks on each trial. Figure 4 shows the bimodal distribution observed over 100 trials using random starting locations with one of the two landmarks removed.

Figure 5 shows the distribution in the split landmark array case. Each of the two prominent peaks is located at the correct bearing and distance from one of the two landmarks. There are also a small number of outlying points such as would result from mis-identifying the east landmark as the west one (and assuming the east one to be missing), or vice versa. These outliers could be eliminated by adjusting the model's parameters, but we have retained them because we are using the same parameters for all the experiments we replicate. The current parameter set reproduces a diverse collection of other experimental results in addition to the gerbil behavioral results modeled here.

7 Relationship to Vector Voting

In the original formulation of the vector voting hypothesis, the animal applied each of n remembered goal-to-landmark vectors to each of m perceived landmarks, generating $n \cdot m$ votes for possible goal locations. The number of candidate goal locations and their positions relative to each other were a function of the landmark configuration on that probe trial. Thus, the animal must have had some way of keeping track of several distinct candidates and tallying votes for each of them. When there was a tie for first place, search effort was distributed among these locations. Otherwise only the winning location was searched.

Our parallel relaxation model can also be understood as a distributed form of vector voting. The goal location is not in dispute; rather, we are deciding among competing solutions for the animal's own location. Each place unit has some fraction of a vote, depending on its activation level.

Our calculation of a mean coordinate value for the active place code serves as a crude estimate of the best-supported candidate for the animal's location. And the step in our algorithm where units are deactivated if they propose coordinates too

distant from the mean results in the elimination of decisively outvoted candidates. Like many connectionist algorithms, this one takes a coarse coded approach that approximates the solutions produced by a fully discrete algorithm (vector voting.) The attraction of the connectionist approach is that it is more naturally implemented in a parallel architecture of simple units reflecting the way we think brains might do computation.

The principal difference between our model and the original vector voting hypothesis is in the handling of ambiguity. Because Collett et al.'s method has no difficulty representing tied candidates simultaneously, it directly supports the animal's dividing its search time between these locations. Our approach always produces a single winner; it requires a consistent place code as it can only represent a single estimate of the animal's Cartesian coordinates. Therefore we must find another way to reproduce the alternation among tied candidates.

In the simulations reported here, the animal was started from multiple locations and the distribution of goal estimates plotted. This reproduces the multi-modal quality of Collett et al.'s results, but does not in itself produce alternation between competing search locations. One method we believe could produce alternation is to suppress the currently active set of place units when a search is unsuccessful. If the animal then performs a new relaxation based on currently perceived landmarks, a new place code will form indicating that it isn't really at the goal after all, but is instead at some location displaced from the goal. Executing a trajectory from this recalculated position "back" to the goal should take the animal to the other search location. We are preparing a simulation to test this idea.

8 Place Fields in the Split Landmark Array Task

In Figure 5, our model produces a search distribution with two principal peaks, not four. Yet it does not contain a binding mechanism as suggested by Collett et al.'s proposal. Instead, interactions among place units are responsible for the system's sensitivity to the configuration of the array. These interactions are of three forms:

1. **Units tuned to single landmarks.** Since when units are first recruited their two controlling landmarks are chosen randomly, some units will choose the same landmark twice. These units will respond with maximum activation to a single landmark at the correct bearing and distance, no matter what the position of the other landmark.
2. **Calculation of mean coordinate value.** The model estimates the animal's position by computing a weighted mean coordinate value using all active place units. The result in the split landmark array case will normally be close to the center of the array, thus favoring the two interior goal locations over the two exterior ones.
3. **Partial activation of place units.** In the split array case, no place unit tuned to two distinct landmarks will have an exact bearing and distance

match for both. However, because these are radial basis units with a smooth gaussian falloff in response, some place units will have an exact match for one landmark and a partial match for the other. These units will be partially activated. The dynamic thresholding in the relaxation process, using a more lenient threshold, will maintain the number of active place units at the desired level.

Those place units seeing a perfect match for one landmark and no match for the other (as would be the case for units voting for search locations exterior to the stretched landmark array) will be less highly activated than units with a partial match for the second landmark, and hence will fall below even the reduced threshold.

The contributions of these three factors depend on the parameters of the model. We are presently exploring the parameter space.

9 Predictions

A consequence of our theory is that the hippocampal activity pattern should be nearly identical when the animal searches at either of the two locations in Figure 1c or 1d. The reason is that in each location the animal thinks it's at the goal, and the goal location's coordinates are fixed and represented by a single set of place cells.

Another proposal that combines visual landmarks with path integration was recently put forward by McNaughton, Knierim, and Wilson (1994). Place cells in their formulation are tuned to single landmarks. In Figure 1b, for example, the active place code would be composed of two populations: units tuned only to the east landmark and units tuned only to the west one. The McNaughton et al. theory predicts that in the split landmark array case these populations will dissociate.

If place fields are small enough relative to the distance between the split landmarks, the two populations will no longer be co-active. Instead the animal will exhibit distinct place codes as it searches at each of the two goal locations. This conflicts with the predictions of our own theory.

On the other hand, if the place fields of individual cells are large enough that they still largely overlap place fields from the other population when the array is split, then many cells will be active at both search locations, producing a similar place code. However, if one measures the entire extent of each place field and calculates its center, the cells should segregate into two populations. The centers of fields of cells in one population will be shifted relative to the centers of fields of cells in the other, by an amount equal to the stretching of the landmark array (McNaughton et al., 1994).

Another prediction of the McNaughton et al. single-landmark theory is that the loss of head direction information should result in an annular place field spanning all bearings at the learned distance from the controlling landmark (McNaughton et al., 1994). Our multiple-landmark theory would not produce such fields, because two

distance values or the difference in bearings between two landmarks ($\theta_m - \theta_n$) are enough to uniquely localize the animal.²

The single-landmark theory would potentially have difficulty with the split landmark array case of Figure 1d. But as with the vector voting hypothesis, additional mechanisms might be invoked to limit search to the two locations interior to the array. Simulations to test McNaughton et al.'s theory have not been done.

In both theories, place fields at locations more distant from the landmarks should show less distortion when the array is split, because the relative change in bearing and distance values will be reduced.

10 Expanded Environments

Muller and Kubie (1987) familiarized rats with a circular or rectangular arena whose walls were painted flat gray. The circular arena had a white cue card subtending approximately 100° of arc. The rectangular one had a white cue card covering one entire wall. There were no other prominent landmarks. After recording the place fields of various cells in the hippocampus, they placed the rats in an identical arena of twice the diameter. Place fields expanded by a somewhat smaller amount, but maintained their previous organization (Muller and Kubie, 1987).

This scaling of place fields suggests that the cells are not very sensitive to distance cues. Several place cell models make use of this assumption (Sharp, 1991; Burgess et al., 1994). The observed response pattern would seem to predict search at a single location in the split landmark array experiment, where the inter-landmark distance was also doubled. Yet that is not what Collett et al. reported.

There are difficulties in extrapolating from properties of individual place cells to properties of the animal's spatial representation system as a whole. Our model approaches this question by exploring possible modes of interaction among ensembles of active place units.

11 Conclusions

Modeling at the right level is important. We have not tried to faithfully model particular cells in CA3 or CA1. The wiring pattern of these areas is not understood in sufficient detail, and the input and output codes of the hippocampus are unknown. On the other hand, we have not ignored the single-cell recording data. Our model reproduces important properties of hippocampal place cells, including the persistence of place fields in the dark and the relationship between place field orientation and the head direction system.

There are several ways in which the theory could be extended. One current task is to

²If the two landmarks are indistinguishable, the animal's position cannot be uniquely determined, but will be constrained to two possible locations, on opposite sides of the landmark array.

incorporate additional classes of experiments, involving factors such as differences in salience among cues (Collett et al., 1986). A second task is to consider how rodents really learn new environments. Multi-electrode recording techniques now make it possible to observe real-time changes in place cell responses as a rat explores an environment for the first time (Wilson and McNaughton, 1993). The key problem for our theory appears to be how to correct for path integration errors simultaneous with tuning the place units.

Real hippocampal cells are considerably more complex than the idealized place cells we know how to model today. For example, in a two-landmark task similar to the one in Figure 1b, (Gothard et al., 1994) report a variety of response types for CA1 cells. Some were tied to the room, some to the landmarks or goal location (which translated from trial to trial), some fired when the animal was either inside, leaving, or reentering the start box no matter where it was located, and some appeared to be related to a mixture of features of the task. This suggests multiple representations of space exist simultaneously in the hippocampus. One approach to understanding the structure of these representations and the roles the various cell types play in navigation is to construct systems-level models of the sort described here.

Acknowledgements

The second author was supported by a fellowship from the National Science Foundation. Hank Wan also made important contributions to this work. We thank Bruce McNaughton and Kati Gothard for helpful discussions.

References

- Burgess, N., Recce, M., and O'Keefe, J. (1994). A model of hippocampal function. *Neural Networks*, 6/7:1065–1081.
- Collett, T. S., Cartwright, B. A., and Smith, B. A. (1986). Landmark learning and visuo-spatial memories in gerbils. *Journal of Comparative Physiology A*, 158:835–851.
- Gallistel, C. R. (1990). *The Organization of Learning*. MIT Press, Cambridge, MA.
- Gothard, K. M., Skaggs, W. E., Moore, K. M., and McNaughton, B. L. (1994). Behavioral correlates of hippocampal ca1 cells in a landmark navigation task. *Society for Neuroscience Abstracts*, 20:1207.
- Hinton, G. E., McClelland, J. L., and Rumelhart, D. E. (1986). Distributed representations. In Rumelhart, D. E. and McClelland, J. L., editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, volume 1, pages 77–109. MIT Press.

- Leonard, B. and McNaughton, B. L. (1990). Spatial representation in the rat: conceptual, behavioral, and neurophysiological perspectives. In Kesner, R. P. and Olton, D. S., editors, *Neurobiology of Comparative Cognition*, chapter 13, pages 363–422. Lawrence Erlbaum Associates, Hillsdale, NJ.
- McNaughton, B. L., Knierim, J. J., and Wilson, M. A. (1994). Vector encoding and the vestibular foundations of spatial cognition: Neurophysiological and computational mechanisms. In Gazzaniga, M., editor, *The Cognitive Neurosciences*. MIT Press, Boston.
- Mittelstaedt, M. L. and Mittelstaedt, H. (1980). Homing by path integration in a mammal. *Naturwissenschaften*, 67:566–567.
- Mizumori, S. J. Y. and Williams, J. D. (1993). Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *Journal of Neuroscience*, 13(9):4015–4028.
- Muller, R. U. and Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7:1951–1968.
- O’Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely moving rat. *Experimental Brain Research*, 34:171–175.
- Sharp, P. E. (1991). Computer simulation of hippocampal place cells. *Psychobiology*, 19(2):103–115.
- Taube, J. S., Muller, R. I., and Ranck, Jr., J. B. (1990). Head direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, 10:420–435.
- Touretzky, D. S., Wan, H. S., and Redish, A. D. (1994). Neural representation of space in rats and robots. In Zurada, J. M. and Marks, R. J., editors, *Computational Intelligence: Imitating Life*. IEEE Press.
- Wan, H. S., Touretzky, D. S., and Redish, A. D. (1994a). Computing goal locations from place codes. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*. Lawrence Erlbaum Associates.
- Wan, H. S., Touretzky, D. S., and Redish, A. D. (1994b). Towards a computational theory of rat navigation. In Mozer, M., Smolensky, P., Touretzky, D., Elman, J., and Weigend, A., editors, *Proceedings of the 1993 Connectionist Models Summer School*, pages 11–19. Lawrence Erlbaum Associates.
- Wilson, M. A. and McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055–1058.