To appear in the Proceedings of the 16th Annual Conference of the Cognitive Science Society, August 13-16, 1994, Atlanta, GA.

Computing Goal Locations from Place Codes

Hank S. Wan, **David S. Touretzky**, and **A. David Redish**
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213–3891
hsw+@cs.cmu.edu, dst+@cs.cmu.edu, dredish+@cs.cmu.edu

Abstract

A model based on coupled mechanisms for place recognition, path integration, and maintenance of head direction in rodents replicates a variety of neurophysiological and behavioral data. Here we consider a task described in [Collett et al. 1986] in which gerbils were trained to find food equidistant from three identical landmarks arranged in an equilateral triangle. In probe trials with various manipulations of the landmark array, the model produces behaviors similar to those of the animals. We discuss computer simulations and an implementation of portions of the model on a mobile robot.

Introduction

We have developed a model of rodent navigation based on tightly coupled mechanisms for place coding, path integration, and maintenance of head direction. Our theory reproduces a variety of behavioral and neurophysiological phenomena.[1] In this paper, we model behavioral data from [Collett et al. 1986] in which gerbils find a food reward among local landmarks in an otherwise cue impoverished environment. First we review some data on the mechanisms of place coding, head direction, and path integration.

Experimental data from rodents

Recordings from hippocampus have demonstrated that rats have an internal representation of place consisting of cells that fire when the animal is in a particular region of the environment (see [Muller et al. 1991] 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 place fields [McNaughton et al. 1993a]. However, in two visually similar but geometrically distinct arenas, one round and one rectangular, mostly disjoint sets of place cells are active [Muller & Kubie 1987], indicating sensitivity to the geometric structure of the environment. Once the rat has a notion of where it is, place cells continue to respond even when the lights are turned off [Quirk et al. 1990]. Most place cells also continue their activity when some of the visual cues are removed, although some fraction become inactive [O'Keefe & Conway 1978, O'Keefe & Speakman 1987].

Cells in various parts of the rat nervous system, including postsubiculum [Taube et al. 1990a, Taube et al. 1990b] and parietal cortex [Chen 1991], were found to be unimodally tuned to the animal's head direction. The *preferred direction* for such a cell (i.e., the direction eliciting maximal response) is constant throughout an environment. In addition, the difference in preferred direction for any pair of head direction (HD) cells is constant across all environments. But cells' preferred directions measured with respect to true north may differ across environments. HD cells 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 [McNaughton et al. 1991]. 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 HD cells rotate by a corresponding amount. On the other hand, rotation of an unfamiliar environment does not shift the cells' preferred directions [McNaughton et al. 1993b].

Mittelstaedt & Mittelstaedt [1980] and Etienne [1987] have shown that rodents performing a search task 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. This suggests that the animals are maintaining their position relative to the start by means of path integration. Direct neurophysiological evidence for a path integration system is lacking, but there is some evidence that lesion of the caudate nucleus impairs such tasks [Abraham et al. 1983]. Other anatomical and physiological evidence has prompted Douglas [1989] to posit that striatum and/or parietal lobe are involved in updating an attention-related spatial vector by path integration.

Open-field landmark-based navigation

Collett et al. [1986] report experiments in which gerbils were trained to find a food reward among cylindrical landmarks in an otherwise cue impoverished environment. We will concentrate on the response to landmark manipulations in a task where the reward is equidistant from an array of three cylinders arranged in an equilateral triangle. During training, the animal's starting location varied, and the landmark array was

[1] The list includes the following: (1) place cell activity normally reflects the rat's location; (1a) place fields are controlled by visual landmarks; (2) place cells maintain their activity in the dark; (3) reset of head direction is dependent on visual cues; (4) place cell firing is robust against deletion of landmarks; (5) places are not recognized after transposition of landmarks; (6) place cells can develop direction sensitivity in some tasks/environments; (7) loss of direction sensitivity can occur as a response to novelty; (8) in mazes with symmetry, rate of search error reflects the amount of maze rotation; (9) in unrotated mazes with symmetry, rodents show no confusion; and (10) repeated disorientations cause rats to stop resetting head direction by reference to visual cues. (See [Wan et al. 1994] and [Touretzky et al. 1994] for more details.)

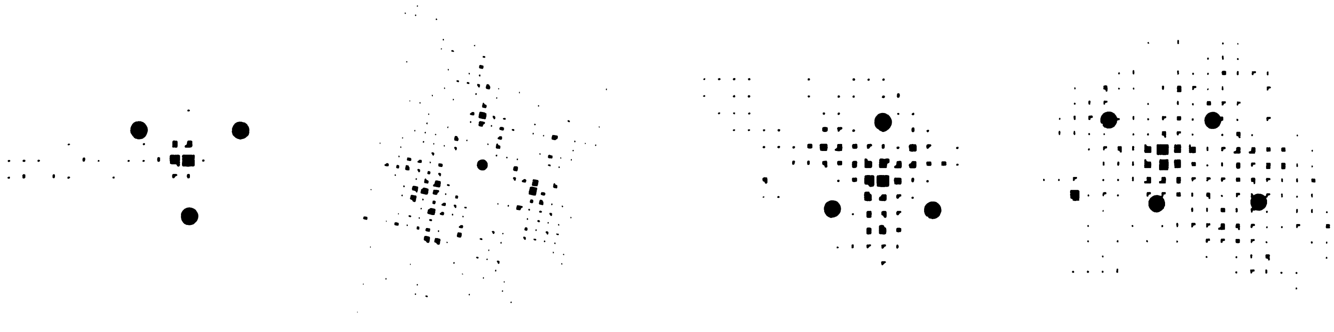

Figure 1: Gerbils were trained to obtain food at a location equidistant from three identical landmarks, drawn as solid circles, arranged in an equilateral triangle. Test trials shown include (a) original landmark array, (b) all but one landmark removed, (c) triangle rotated by 60 degrees, and (d) a fourth landmark added to form two triangles with opposite orientations. Size of blobs show distribution of search efforts. From [Collett et al. 86], reproduced with permission.

translated but not rotated between trials. In probe trials in which no food was present, the distribution of the animal's search time was recorded. Trained gerbils showed accurate performance in probe trials provided there were no changes to the landmark array. See Figure 1a.

When all but one landmark was removed from the array, gerbils searched three locations (Figure 1b). Collett et al. posit that gerbils store a representation containing information about distances and directions to landmarks from the goal location. The search pattern could then be explained by matching the one visible landmark against each of the three remembered landmarks. Each binding to a remembered landmark would predict the goal location to be at a different vector offset from the perceived landmark, thus accounting for the three observed search locations.

Collett et al. point out that since the single landmark is radially symmetric, the orientation of the landmark-to-goal vectors must be aligned with reference to some other directional cue. Rodents' internal sense of head direction, described earlier, could serve as this reference.

In another class of probe trials, the landmark array was rotated by 60°. Gerbils went first to the center of the triangle. Failing to find food there, they then searched briefly at three locations exterior to the triangle, as shown in Figure 1c. These locations can be derived by considering any two landmarks from the rotated array as parts of a triangle *with the correct orientation*.

To account for the gerbils' behavior, Collett et al. posit that the animals begin planning the path to goal by matching perceived to remembered landmarks. This is an instance of the *binding problem*. Given a candidate binding, the gerbil determines the goal location by applying remembered landmark-to-goal vectors to the landmarks it currently perceives. When the bindings are correct, all vectors point to the same goal location (left half of Figure 2.) In the rotated case, no choice of bindings can produce a totally consistent result, but some locations will get two votes while others get only one. Different bindings produce different results; the right half of Figure 2 shows the union of results from all possible bindings with the result from one set of bindings highlighted. The animal focuses its search on those places with the highest number of votes. We will refer to this as the *vector voting* hypothesis. Its predictions closely match the observed search patterns.

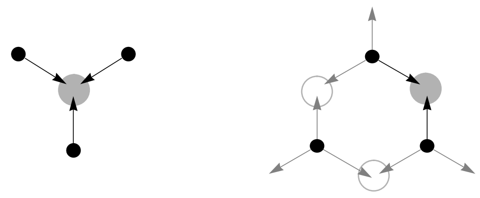

Figure 2: In the vector voting scheme, a vector from each landmark (solid circle) to the goal (shaded region) is learned during training trials. In later probe trials (right half of figure), these vectors are applied to each possible match between perceived and remembered landmarks. Places with the highest number of votes are searched. One search location and its corresponding vectors are shown as a shaded circle and solid arrows.

Although our own theory does not include explicit mechanisms for computing bindings and tallying votes, it produces behavior equivalent to vector voting on the tasks described.

In yet another experiment, a fourth landmark was added to form two triangles with opposite orientations. Gerbils spent most of their time searching at the center of the one with the correct orientation (Figure 1d). This shows that they prefer interpretations where the directional cues among the landmarks are consistent with the external directional cues. Our model also reproduces this result.

Structure of the model

The essence of our theory is in the interactions among mechanisms for maintenance of head direction, path integration, and place recognition. The structure of this theory is illustrated in Figure 3.

First, by combining head direction Φ_h with egocentric perception $\langle r_i, \theta_i \rangle$, the animal computes allocentric bearings to landmarks Φ_i. The identity, distance, and bearing of each landmark constitutes a local view of the environment unique

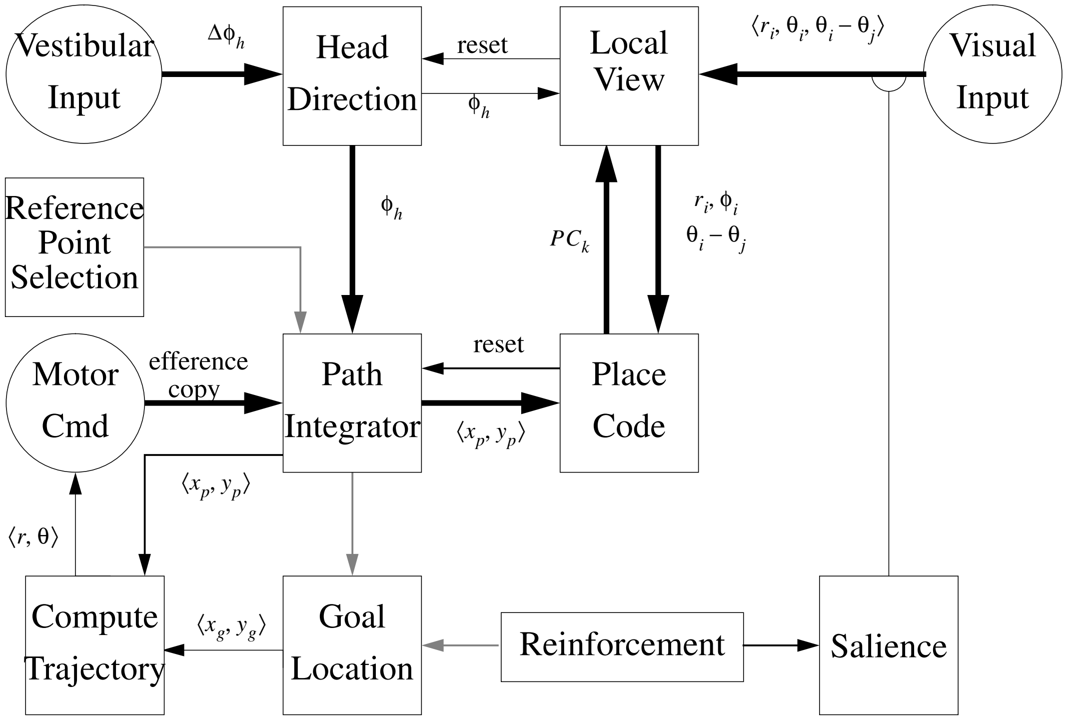

Figure 3: Suggested functional organization of orientation and recognition mechanisms in the rodent. Circles are input quantities, boxes are computational modules; they do not necessarily correspond to disjoint or unique brain areas. Thick lines denote main information pathways. Hashed lines are pathways not currently modeled. Φ_h is head direction, Φ_i and θ_i are allocentric and egocentric bearings to landmark i. r_i is distance to landmark i. (x_p, y_p) are coordinates with respect to reference point p, and (x_g, y_g) are coordinates of a goal location, relative to the same reference point.

to the current location.

Second, we postulate that animals keep track of their current location in terms of coordinates relative to *reference points*, by means of path integration. A similar proposal was recently made in [McNaughton et al. 1993a]. A reference point in our theory may be perceptually significant, such as a corner of the room or a place with distinct odor or texture, but it may also be distinguished only by the past occurrence of some significant event, such as the spot in the arena where the rodent first found food. There is some evidence that place cell firing in a visually symmetric environment is dependent on the point of entry [Sharp et al. 1990], which can be regarded as a reference point.

The place code in our theory is computed by radial-basis units tuned to distances, allocentric bearings, and angles between selected visual landmarks, and to a set of path integrator coordinates $\langle x_p, y_p \rangle$ measured with respect to a reference point. These units compute a fuzzy conjunction of their inputs, in which terms drop out when information is not available. Thus, in the dark, place units are driven solely by the path integrator. Conversely, when the animal is reintroduced into a familiar environment, path integrator coordinates are initially unavailable, and place units are driven by visual landmarks. Although they will show place fields similar to those of hippocampal place cells, individual units in our model do not necessarily correspond to single pyramidal cells; we prefer to think of them as cell assemblies.

Operations of the model

We model a rodent's spatial behavior in a variety of situations via a set of operations that include computing a place code, reset of head direction, reset of the path integrator, and computing a trajectory to a goal location. In addition, the model recruits new place units to represent its environment whenever too few are active at the current location.

The computation of place unit activity contributes to all these operations. The fuzzy conjunction that determines this activity is implemented as a product of Gaussian functions of the inputs, where the width of each Guassian can vary depending on the operation. A dynamic threshold allows units with the largest net input to remain active.

When the simulated rodent is first introduced into the experimental environment at the start of a trial, if it has been disoriented or the environment has been rotated, the model must realign its head direction sense. It does this by first computing a place code using just perceptual input (egocentric landmark positions.) Active place units then recall the learned allocentric bearings of landmarks in the current local view. Subtracting allocentric from egocentric bearings of all visible landmarks produces a set of estimates of the current head direction. The model stochastically chooses a heading according to this evidence and resets its head direction sense accordingly. Details of this operation and its use in simulating the tasks in [Cheng 1986] are given in [Touretzky et al. 1994].

When the environment is manipulated so that some landmarks are not at their expected positions, a larger number of

a b c d

Figure 4: Results of simulations. Blobs show distribution of estimates in the goal position across various starting locations. Compare with Figure 1.

place units may become active, but individual units are likely to be less strongly activated than in the unperturbed case. Each place field has an associated set of path integrator coordinates, and so, with a larger number of units active, there will be inconsistency in the estimate of the animal's position. The model averages the contributions of the active units to arrive at a rough position estimate for the path integrator. It then feeds this estimate back to the place units and sharpens their tuning to path integrator coordinates; units with coordinates very far from the estimated value are forced to drop out of the active set. The process repeats until the place system converges on a consistent position estimate. In the next section we show how this relaxation process is used for self-localization in the Collet et al. tasks.

Correspondence between simulations and behavioral data

A probe trial begins by starting the simulated rodent at a random spot in the environment, so that initially it has no knowledge of its path integrator coordinates, though its sense of direction remains intact. The model uses visual cues to activate the place code for the present location. The active place units then recall the path integrator coordinates associated with that place. Once the animal has reconstructed its position with respect to the reference point it can estimate a path to the goal by vector subtraction. The distribution of estimates gathered across probe trials is plotted in Figure 4.

With no change to the array (Figure 4a), the goal estimate is concentrated in the correct location, in the center of the triangle. With all but one landmark removed (Figure 4b), the goal estimate is equally distributed around three locations corresponding to the three possible matches between the perceived landmark and the three remembered landmarks. In the rotated triangle (Figure 4c), the goal estimate is concentrated in the three places receiving the most votes according to the vector voting hypothesis. And in the four landmark test (Figure 4d), the goal estimate is concentrated in the center of the triangle with the same orientation as seen during training.

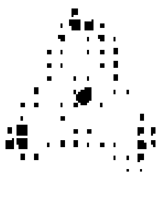

Figure 5: Results of simulating the task with one landmark remaining, invoking the self-localization operation near the landmark. Compare with Figure 1b and Figure 4b.

Our simulated search distributions are qualitatively similar to the behavioral data plotted in Figure 1. Although Figure 4b does show a concentration in the three locations corresponding to the three possible matches, the distributions are not as clean as in Figure 1b, due to limited accuracy of distance and angle measurements at a distance. If the rodent approaches the general vicinity of the landmark before invoking self-localization, the results are much cleaner, as shown in Figure 5.

Comparing Figure 4c with Figure 1c, our simulated rodent does not search the middle of the rotated triangle. This can be explained if the animal makes use of an alternate strategy: at a distance, it could treat the entire array as a "beacon," i.e., a single landmark co-located with the goal. Far away from the landmark array, aiming for this beacon would be easier than computing a trajectory by self-localization. In the case of the rotated triangle, when the animal treats the array as a beacon, it might not notice that the array has been rotated. It would go first to the center of the array; failing to locate reward in the center, it would then attempt to reorient itself based on its place representation, and would search the three exterior locations.

Alternatively, Collett et al. argue that the animal initially searched the middle of the rotated triangle because it preferred using directional cues based on the landmark array over those external to the array. This proposal is also compatible with our model. The animal could notice the rotation of the landmark array, and respond by realigning its sense of direction so that the landmarks appeared at their familiar bearings. Having done this, all three landmarks would match their remembered position, and so the place code would be retrieved that predicts the goal to be at the center of the triangle. However, when the animal failed to find food there, it would have to restore the previous setting of its head direction system in order to generate search hypotheses at the three points external to the landmark array. Our candidate mechanisms for accomplishing this reset are beyond the scope of this paper.

Predictions

The two competing explanations for search inside the rotated triangle can be distinguished by recording from head direction cells during probe trials. Our first proposal predicts that the tuning of these cells will not change when the animal gives up searching at the center of the array and moves to one of the three exterior search locations, while the second proposal requires such a change.

Our model also makes predictions about the behavior of place cells. Place unit activity reflects where the animal thinks it is in the coordinate system defined by some reference point. In any task where the animal abandons one search location in favor of an alternate, it does so by changing its estimate of its own position. Hence, the place code should be observed to shift accordingly. Recent advances in multi-cell recording techniques in freely-behaving rats [Wilson & McNaughton 1993] should permit a test of this prediction.

The place code relaxation process that initializes the path integrator implies that many place units should become briefly active at the beginning of a trial. Units will activate if some of their perceptual cues are approximately matched, even if the animal is not in the cell's normal place field as defined (in part) by path integrator coordinates. To our knowledge, no one has yet recorded from place cells of a disoriented rat while it is being reintroduced to a familiar environment.

Discussion

We have shown that when our model is applied to the manipulations of the equilateral triangle landmark arrangement described in [Collett et al. 1986], it produces behavior equivalent to vector voting. Yet it contains no mechanism for explicitly binding perceived to remembered landmarks, or for gathering votes from multiple independent vector computations.

Rather than explicitly computing a set of candidate goal locations, the place system assembles evidence from perceptual cues to refine the estimate of the animal's own location in terms of a coordinate frame on which the goal location has been established. Bindings are implictly achieved because once the position of the animal is determined, there is a unique correspondence between each remembered landmark and where it could be seen.

Implementation on a Mobile Robot

To show that our model is an adequate mechanism for rodent navigation, we have implemented portions of it on the Xavier mobile robot [Nourbakhsh et al. 1993] at CMU. This allows us to test our theory under real world constraints, such as sensor noise and uncertainty in the perceptual system.

Xavier is equipped with a ring of 24 ultrasonic range sensors, an infrared laser rangefinder, and a color camera mounted on a pan/tilt head. On-board computing power is provided by two i486 computers and a laptop, interconnected via Ethernet.

We ran the robot in a 22×23 foot irregularly-shaped maze whose walls were constructed from cardboard packing boxes. Ultrasonic range data was collected into an occupancy grid [Moravec 1988]; then standard edge detection and Hough transform algorithms [Ballard & Brown 1982] were applied to detect concave and convex corners. These corners served as landmarks for the model. The robot performed the self-localization operation described earlier, reset its path integrator, and computed a vector from its present position to a predefined goal location. This vector was then passed to a lower-level navigation controller for execution.

Parallel relaxation to compute a consistent place code appears to be robust against sensor noise and the false registration of features. Our preliminary results suggest that improved robot navigation algorithms may result from modeling animal behavior.

Acknowledgements

We thank Bruce McNaughton and Neil Burgess for helpful discussions, and T. S. Collett for permission to reproduce his figures. Hank Wan and David Redish were supported by NSF fellowships.

References

[Abraham et al. 1983] 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.

[Ballard & Brown 1982] D. H. Ballard and C. M. Brown. *Computer Vision*. Prentice-Hall, Englewood Cliffs, NJ, 1982.

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

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

[Collett et al. 1986] 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.

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

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

[McNaughton et al. 1991] 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.

[McNaughton et al. 1993a] 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, ed., *The Cognitive Neurosicences*. MIT Press, Boston, 1993.

[McNaughton et al. 1993b] 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.

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

[Moravec 1988] H. P. Moravec. Sensor fusion in certain grids for mobile robots. *AI Magazine*, 9(2):61–74, 1988.

[Muller & Kubie 1987] R. U. Muller and J. L. Kubie. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7:1951–1968, 1987.

[Muller et al. 1991] 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, ed., *Brain and Space*, ch. 17, pp. 296–333. Oxford University Press, New York, 1991.

[Nourbakhsh et al. 1993] I. Nourbakhsh, E. Gat, S. Morse, C. Becker, M. Balabanovic, R. Simmons, S. Goodridge, H. Potlapalli, D. Hinkle, K. Jung, and D. V. Vactor. The winning robots from the 1993 robot competition. *AI Magazine*, 14(4):51–62, 1993.

[O'Keefe & Conway 1978] 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.

[O'Keefe & Speakman 1987] 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.

[Quirk et al. 1990] 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.

[Sharp et al. 1990] 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.

[Taube et al. 1990a] 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.

[Taube et al. 1990b] 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.

[Touretzky et al. 1994] D. S. Touretzky, H. S. Wan, and A. D. Redish. Neural representation of space in rats and robots. In J. M. Zurada and R. J. Marks, eds., *Computational Intelligence: Imitating Life*. IEEE Press, 1994.

[Wan et al. 1994] 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, eds., *Proceedings of the 1993 Connectionist Models Summer School*, pp. 11–19. Lawrence Erlbaum Associates, 1994.

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